

**Fluctuation-Induced Interactions and Nonlinear
Nanophotonics**

by

Alejandro Rodriguez-Wong

B.S., Massachusetts Institute of Technology (2006)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of


Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2010

ARCHIVES



© Massachusetts Institute of Technology 2010. All rights reserved.

Author
Department of Physics
May 13, 2010

Certified by
John D. Joannopoulos
Francis Wright Davis Professor of Physics
Thesis Supervisor

Certified by
Steven G. Johnson
Associate Professor of Applied Mathematics
Thesis Supervisor

Accepted by
Krishna Rajagopal
Associate Department Head for Education

Fluctuation-Induced Interactions and Nonlinear Nanophotonics

by

Alejandro Rodriguez-Wong

Submitted to the Department of Physics
on May 13, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

We present theoretical and numerical methods for studying Casimir forces and nonlinear frequency conversion in nanophotonic media consisting of arbitrary geometries and materials.

The first section of the thesis focuses on the study of various geometry-enabled resonant effects leading to strong nonlinear interactions. The starting point of this work is a coupled-mode theory framework for modeling a wide range of resonant nonlinear frequency-conversion processes in general geometries, ameliorating the need for repeated and expensive finite-difference time-domain simulations. We examine the predictions of the theory for two particular nonlinear processes: harmonic generation and difference-frequency generation. Our results demonstrate strong enhancement of nonlinear interactions at a “critical” input power leading to 100% frequency conversion, among many other interesting dynamical effects. Using a quantum-mechanical description of light, based on cavity quantum electrodynamics, similar enhancement effects are demonstrated at the single-photon level, leading to the possibility of achieving all-optical switching of a single signal photon by a single gating photon in a waveguide-cavity geometry consisting of pumped four-level atoms embedded in a cavity. Finally, we describe how one may tailor the geometry of certain materials to enhance their nonlinear susceptibilities by exploiting a consequence of the Purcell effect.

The second section of the thesis, the main contribution of this work, presents a new formulation for studying Casimir forces in arbitrary geometries and materials that directly exploits efficient and well-developed techniques in computational electromagnetism. To begin with, we present the step-by-step conceptual development of our computational method, based on a well-known stress tensor formalism for computing Casimir forces. A proof-of-concept finite-difference frequency-domain implementation of the stress-tensor method is described and checked against known results in simple geometries. Building on this work, we then describe the basic theoretical ingredients of a new technique for determining Casimir forces via antenna measurements in tabletop experiments. This technique is based on a (derived) correspondence between the complex-frequency deformation of the Casimir frequency-integrand for any given geometry and the real-frequency classical electromagnetic response of the same geometry, but with dissipation added everywhere. This correspondence forms the starting point of a numerical Casimir solver based on the finite-difference time-domain method, which we describe and then implement via an off-the-shelf time-domain solver, requiring no modifications. These numerical methods are then used to explore a wide range of geometries and materials, of various levels of complexity: First, a four-body piston-like geometry consisting of two cylinders next to adjacent walls, which exhibits a non-monotonic lateral Casimir force (explained via ray optics and the method of images);

Second, a zipper-like, glide-symmetric structure that leads to a net repulsive force arising from a competition between attractive interactions. Finally, we examine a number of geometries consisting of fluid-separated objects and find a number of interesting results. These include: stable levitation and suspension of compact objects, dispersion-induced orientation transitions, and strong non-zero temperature Casimir effects.

Thesis Supervisor: John D. Joannopoulos
Title: Francis Wright Davis Professor of Physics

Thesis Supervisor: Steven G. Johnson
Title: Associate Professor of Applied Mathematics

Dedication

I dedicate this thesis to the many friends and family members that I left in Cuba in search of a better life.

To my father Sergio Julian Rodriguez, for the sacrifice he made by allowing me to emigrate to the United States; for always being there for me despite the distance and obstacles that lay between us. To my mother Maria E. Wong, who has been a limitless source of strength and love throughout my life, and who will always be the little voice inside my head (for the uncountable scarificies you have made for me). To my step-father Lazaro Blanco, who has loved and guided me like a father since I was three years old (for imparting his love of physics to me and for loving me like a son). To the love of my life, Silvia M. Lopez, for being the most wonderful friend that anyone could ask for (because you are beautiful in every way). Thank you for your compassion and help in the most difficult of times.

To my grandpa Pedro Luis Soto (abo), for the many times that you walked me to school; for being the grandpa that every kid should have. To my late grandma Julia Pimentel (mima), for every “pomo de leche con chocolate” that she fed me in bed or the many days spent in the river catching “guajacones”; although you are no longer with me mima, you are also the little voice inside my head. To my aunt Beatriz Soto (zanahoria), uncle Caballo-yero and cousin Alejandro, for always making me feel that I am still a big part of the family, even though we have been separated for more than 13 years. To my late grandpa Adolfo Rodriguez (papo), who recently passed away, for giving me wonderful childhood memories—for allowing my cousin Dalena and I to frighten and throw water at him when he lay sleeping in the living-room rocking chair. To my grandma Dalia Gonzales (abuela), for allowing me to sleep between her and papo during the weekends, and for frequently frying “papas fritas” at my behest. To my tia Gisela Rodriguez and cousins Dalena and Daniela, for their affection and love despite the many years we’ve spent apart. To my tia Isabel Lopez (“my nephew”), for helping us settle in the US and for always being a great source of widsom (I love you my nephew).

Acknowledgements

About the Author

Alejandro Rodriguez Wong was born in Havana Cuba, and spent the early days of his youth in a little “reparto” known as Mantilla. He is an “aspiring physicist” and a member of a family of physics enthusiasts. The former was not always the case...

It all began with his grandma “mima”, who instilled in him a deep love and respect for flora and fauna, and initiated his desire to become a botanist by the age of four. Years later, he decided to become an “animal philanthropist”, and dreamt of the possibility of creating a farm where “all stray dogs could live happily ever after”. After seeing what he was convinced to be a “very sad dolphin” in the Cuban aquarium, his philanthropic dreams grew to incorporate an entire ocean. At the age of six, Alejandro decided to settle for a less noble cause, and instead dreamt of becoming the driver of a yellow cab. Alejandro’s family members often referred to him as “Armando Bianchi” (italian actor) or “Placido Domingo” (renowned spanish tenor) due to his propensity to dramatize situations and sing in a tenor voice (most commonly while showering in the bathroom), which led him to (briefly) consider a career in acting. After spending numerous hours playing Final Fantasy (the game) in middle school, Alejandro became fascinated with the idea of becoming a game designer. Years later, his game fixation motivated him to study basic computer programming, leading to the creation of two programs (*Calulator 2000* and *Student Database 2001*) which he went ahead and unsuccessfully tried to sell to various members of his high school, including his teachers. By his sophomore year in high school, Alejandro was a full-fledge physics enthusiast, like the rest of his family, and came to be known as “Captain physics” by many of his peers. Having graduated high school in 2003, Alejandro enrolled at MIT, where he met his current advisors Steven G. Johnson and John D. Joannopoulos.

Muchas Gracias

Being at MIT these last seven years, becoming part of a community of unique and talented people that love and shape science, has had a profound influence on me—I am overwhelmed with a strong desire to instill the same passion and awe in others, especially those that have been less fortunate than me. My journey in this capacity, as I see it, has just begun.

I would now like to thank all of those people that shaped my life throughout these seven

years:

I would like to thank my advisors John D. Joannopoulos (JJ) and Steven G. Johnson (maestro) for their immeasurable support and dedication; their contribution to my growth as a scientist cannot be overstated. Thank you John for introducing me to Steven the summer of 2004 and for always making sure that I felt welcomed and appreciated. Your ability to communicate science and your integrity will always be the standard by which I will measure myself. Thank you Steven for not only being my advisor, but also a good friend during hard times. Your passion for science, your broad understanding of many subject matters, your writing skills, your ability to convey the most difficult concepts in the simplest ways and most importantly, your willingness to help others, have guided me throughout these past six years and will also be the standard by which I will measure myself. Thank you both for treating me as part of your family: this has been a source of strength and pride for me. Thank you for teaching me how to perform great science and also how to be a great scientist (and for showing me the difference).

Many thanks to my academic advisor Marin Soljacic for his support and guidance, especially during the first few graduate years.

I would like to thank Adam Donovan, whom I met the first day of class at MIT (18.100B), and has been a very close friend since. Thank you buddy for your constant support throughout these past seven years. He is perhaps the smartest person I know and most importantly, among the nicest as well. In addition to sharing a deep love of computers and physics, Donny and I enjoy playing Age of Empires in the weekends, which has alleviated some of the stress associated with earning a PhD from MIT. His ability to diagonalize 9×9 matrices earned him the nickname “matlab” for a short period of time, although he is now known as “neutron”. Shortly after I met Donovan, I met another very close friend: “proton”, or as he is sometimes known, Lyel Resner. Lyel is blessed with the “gift of gab” and has also been a source of love and support. His desire to help others has always been a source of admiration from me and many others. Both of you have had a profound influence on me (thank you).

Many thanks to my close friends Manny Soto (mamberto perez) and Manuel Rivas (manolo) for their support and friendship (and for the many adventures) throughout these years.

Our group, the JDJ gang, also holds a special place in my heart: The first student I met, Mihai Ibanescu, who became an officemate and friend, and who played an instrumental role

during my early days in the group; Ardavan Farjadopour (“abishek”), who became a great friend and whom I share many wonderful memories with, from our biking trips throughout Boston, to the many practical jokes we played on each other, or our habitual political debates; Aristeidis Karalis (“kakalis”), who has been my officemate and friend throughout the past six years (you have been a constant source of support and friendship). Finally, many thanks to Peter Bermel (or as we call him “petey pablo”), Alexander McCauley (“myculo”), Hila Hashemi (“bobila”), and Maissam Barkeshli for your friendship. To the many other members of the group: thank you for making my time at MIT a wonderful experience.

A special thanks to Margaret O’Meara, Nancy Savioli and Crystal Young for their support and patience. Margaret and Nancy have known me since I was an undergraduate student and have become surrogate mothers (thank you for always brightening my day).

An important part of my thesis involves collaborations with many people from and outside of MIT, which I would also like to thank. These include: Mihai Ibanescu, Peter Bermel, Alexander McCauley, Hila Hashemi, Jorge Bravo, David Ramirez, Jaime Varela, Homer Reid, Jamal S. Rahi, Saad Zaheer, Jeremy Munday, David Woolf, Pui-Chuen Hui (Wallace), Ian Burgess, Murray McCutcheon, Yinan Zhang, Michael Levin, Ofer Shapira, Marin Soljacic, Robert L. Jaffe, Mehran Kardar, Davide Iannuzzi, Diego Dalvit, Marko Loncar, Federico Capasso.

Contents

1	Introduction	33
2	Classical Nonlinear Nanophotonics	45
2.1	Overview	45
2.1.1	Nonlinear Nanophotonics	47
2.1.2	Temporal Coupled-Mode Theory	49
2.2	Second and Third Harmonic Generation	53
2.2.1	Nonlinear Coupled-Mode Theory	53
2.2.2	Coupling Coefficients via Perturbation Theory	55
2.2.3	Complete Frequency Conversion at a Critical Power	58
2.2.4	Numerical Validation	65
2.2.5	Kerr Cross- and Self-Phase Modulation	67
2.2.6	Stability and Dynamics of XPM and SPM	75
2.2.7	Exciting High-Efficiency Solutions	78
2.2.8	Effect of Losses	83
2.2.9	Future Work	85
2.3	Difference Frequency Generation	85
2.3.1	Terahertz Generation	90
2.4	Three-Wave Mixing	94
3	Quantum Nonlinear Nanophotonics	97
3.1	Overview	97
3.2	Single-Photon Switching	98
3.3	Nonlinear Purcell Enhancement	107
3.4	Concluding Remarks	115

4	Casimir Forces: A Brief Introduction	117
4.1	Introduction	117
4.1.1	From van der Waals to Casimir-Polder Forces	117
4.1.2	Charge \leftrightarrow Electromagnetic Fluctuations	120
4.2	Experimental and Theoretical Progress	124
5	Casimir Forces in the Imaginary Frequency Domain	129
5.1	Overview	129
5.2	Mode Summation: A Simple Approach	131
5.3	Wick Rotation and Energy Density	134
5.4	Stress-Tensor Computational Approach	138
5.5	The Finite-Difference Method	140
5.6	Translation-Invariant Structures	142
5.6.1	Periodic Green's Function Calculation	144
5.7	Numerical Validation	146
5.7.1	Parallel Metal Plates	146
5.7.2	Cylinder and Plate	147
5.8	Beyond Finite-Difference Methods	148
5.8.1	Spectral and Monte-Carlo Approaches	150
5.9	Concluding Remarks	152
6	Casimir Forces in the Complex Time Domain	155
6.1	Overview	155
6.2	Revisiting the Stress Tensor Formulation	157
6.3	Correspondence: Contour \leftrightarrow Material Deformations	160
6.3.1	Complex-Frequency Green's Functions via Material Deformations	161
6.3.2	Numerical Experiment	164
6.3.3	A Casimir Analog Computer	165
6.4	Casimir Forces in the Time Domain	170
6.4.1	Time Domain Approach	172
6.4.2	Evaluation in the Time Domain	176
6.4.3	Properties of $g(-t)$	177
6.5	The Finite-Difference Time-Domain Method	180

6.5.1	Fields in Real Time	180
6.5.2	Parallel Plates	181
6.5.3	Force Convergence in Time	182
6.6	Implementing the FDTD Method	184
6.6.1	Temporal Discretization	184
6.6.2	The Magnetic Correlation Function	186
6.6.3	Material Dispersion	187
6.6.4	Overview of the Algorithm	188
6.6.5	Spatial Discretization	190
6.7	Concluding Remarks	191
7	Casimir Forces Between Vacuum-Separated Metals	193
7.1	Overview	193
7.2	Multi-Body Nonmonotonic Effects	194
7.2.1	Generalized Casimir Piston	194
7.2.2	Cylinders and Plates	198
7.2.3	Magnetic Sidewalls	206
7.3	Pistons via Ray Optics	207
7.3.1	Introduction to Ray Optics	210
7.3.2	Pairwise-Interaction Approximations	211
7.3.3	Path Contributions to Nonmonotonic Force	212
7.3.4	Details of the Ray-Optics Computation	216
7.3.5	Even Paths	216
7.3.6	Odd Paths	220
7.3.7	Casimir Piston	222
7.3.8	Numerical Evaluation	224
7.3.9	Concluding Remarks	225
7.4	Glide-Symmetric Geometry: Uni-Axial Stability	226
7.4.1	Numerical Results	230
7.4.2	Comparison to PFA	231
7.5	Concluding Remarks	234

8	Casimir Forces Between Fluid-Separated Dielectrics	237
8.1	Overview	237
8.2	Repulsive Forces via Fluids	239
8.3	Eccentric Cylinders: Uniaxial Stable Suspension	242
8.4	Stable Suspension and Confinement	249
8.4.1	Gravity-Induced Stability	249
8.4.2	Nontouching Plane-Separated Compact Objects	251
8.5	Strong Temperature Dependence	257
A	Disorder-Immune Confinement of Light in Photonic-Crystal Cavities	267
A.1	Influence of Disorder on Complete-Gap Modes	267
B	Computing Photonic-Quasicrystal Spectra via Bloch's Theorem	273
B.1	Quasicrystals via cut-and-project	275
B.2	Computations in Higher Dimensions	277
B.2.1	Bloch's Theorem and Numerics for Quasicrystals	279
B.2.2	The Spectrum of the Quasicrystal	280
B.3	One-Dimensional Results	281
B.3.1	Fibonacci Quasicrystal	282
B.3.2	Defect Modes	286
B.3.3	Continuously Varying the Cut Angle	289
B.4	Defining Irrational Slice	290
B.5	Concluding Remarks	292

List of Figures

- 2-1 (Top:) Schematic of narrow-band filter system, consisting of a cavity with a mode a of frequency ω_0 , coupled to two input/output ports $s_{1\pm}$ and $s_{2\pm}$, with decay rates $1/\tau_1$ and $1/\tau_2$, respectively. The cavity is filled with an absorptive medium and has an absorption rate $1/\tau'$. (Bottom:) transmission $T = |s_{2-}|^2/|s_{1+}|^2$ (red line) and reflection $R = |s_{1-}|^2/|s_{1+}|^2$ (blue line) coefficients of input light from the left (s_{1+}) as a function of ω , for $\tau' = \infty$ (solid line) and $\tau' \sim \tau_1$ (dashed line). The inset (green line) shows a generic Lorentzian line spectrum. 52
- 2-2 (Top:) Schematic of general scheme for third-harmonic generation, including dynamical variables for coupled-mode equations: a single input (output) channel (with incoming (outgoing) field amplitudes s_{\pm}) is coupled to a resonant cavity with two modes at frequencies ω_1 and $\omega_3 = 3\omega_1$ (with corresponding amplitudes a_1 and a_3). The two resonant modes are nonlinearly coupled by a Kerr ($\chi^{(3)}$) nonlinearity. Reflections at the fundamental frequency (s_{1-}) may also occur. (Bottom:) An example realization [445], in one dimension, using a semi-infinite quarter-wave stack of dielectric layers with a doubled-layer defect (resonant cavity) that is coupled to incident plane waves. Dielectric material is yellow, and the steady-state electric field E_z of third-harmonic mode is shown as blue/white/red for positive/zero/negative amplitude. 59
- 2-3 Plot of first and second harmonic efficiency, $|s_{1-}|^2/|s_{1+}|^2$ (black) and $|s_{2-}|^2/|s_{1+}|^2$ (red), vs. $\chi^{(2)}|s_{1+}|$. 100% power transfer from ω_1 to $\omega_2 = 2\omega_1$ is achieved at $\chi^{(2)}|s_{1+}| = 1.8 \times 10^{-3}$ 62

- 2-4 Plot of first and third harmonic steady-state efficiency, $|s_{1-}|^2/|s_{1+}|^2$ (black line) and $|s_{3-}|^2/|s_{1+}|^2$ (red line), for $\alpha = 0$ (no self/cross-phase modulation), as a function of input power $|s_{1+}|^2$ scaled by the Kerr coefficient $n_2 = 3\chi^{(3)}/4c\epsilon$. The reflected power at the incident frequency ω_1 is shown as a dashed black line. 100% power transfer from ω_1 to $\omega_3 = 3\omega_1$ is achieved at $n_2|s_{1+}|^2 = 2.8 \times 10^{-4}$. The parameters used in this plot are $Q_1 = 1000$, $Q_3 = 3000$, $\beta_1 = (4.55985 - 0.7244i) \times 10^{-5}$ in dimensionless units of $\chi^{(3)}/V\epsilon$. 64
- 2-5 Log-log plot of $|s_{3-}|^2/|s_{1+}|^2$ vs. $n_2|s_{1+}|^2$ for the coupled-mode theory (grey) and FDTD (black squares), where n_2 is being varied. Also shown are the corresponding $\Delta n/n$ (dashed blue) and $\Delta\omega_1/\omega_1$ (solid red) curves. 66
- 2-6 Shift in the resonant frequency ω_1^{NL} as a function of input power, due to self- and cross-phase modulation. (There is an identical shift in ω_3^{NL} .) If the cavity is designed so that the linear ($P_{\text{in}} \rightarrow 0$) frequencies are harmonics, the nonlinearity pushes the system out of resonance (lower blue line) as the power increases to the critical power for 100% efficiency. This is corrected by pre-shifting the cavity frequencies (upper green line) so that the nonlinear frequency shift pushes the modes into resonance at P_{crit} 72
- 2-7 Phase diagram of the nonlinear dynamics of the doubly-resonant nonlinear harmonic generation system from Fig. 6-8 as a function of the relative cavity lifetimes ($\tau_3/\tau_1 = 3Q_3/Q_1$) and the relative strength of SPM and XPM vs. harmonic generation (α/β_1) for input power equal to the critical power for 100% efficiency. For $\tau_3 < \tau_1$ there is always one *stable* 100%-efficiency solution, and for nonzero α the system may have additional stable solutions. For $\tau_3 > \tau_1$ the 100%-efficiency solution becomes unstable, but there are limit cycles and lower-efficiency stable solutions. Various typical points A–G in each region are labeled for reference in the subsequent figures. 74

2-8 An example of a limit-cycle solution, with a periodically oscillating harmonic-generation efficiency as a function of time, corresponding to point D in Fig. 2-7. Perturbations in the initial conditions produce only phase shifts in the asymptotic cycle. Here, the limit cycle has a period of around 3×10^4 optical cycles. *Inset:* Square of Fourier amplitudes (arbitrary units) for each harmonic component of the limit cycle in the Fourier-series expansion of the A_3 76

2-9 Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of α/β_1 for a fixed $\tau_3/\tau_1 = 0.7$, corresponding to the line ACF in Fig. 2-7 (see inset). The input power is the critical power P_{crit} , so there is always a 100%-efficiency stable solution, but as α/β_1 increases new stable and unstable solutions appear at lower efficiencies. 77

2-10 Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of τ_3/τ_1 for a fixed $\alpha/\beta_1 = 3$ (left) or $= 8$ (right), corresponding to the lines BCD or EFG, respectively, in Fig. 2-7 (see insets). The input power is the critical power P_{crit} , so there is always a 100%-efficiency steady-state solution, but it becomes unstable for $\tau_3 > \tau_1$ (a Hopf bifurcation leading to limit cycles as in Fig. 2-8). 78

2-11 (Left:) Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of input power $P_{\text{in}}/P_{\text{crit}}$ at fixed $\alpha/\beta_1 = 3$ and $\tau_3/\tau_1 = 0.7$, corresponding to point C in Fig. 2-7; the inset shows an enlarged view of the high-efficiency solutions. (Right:) Bifurcation diagram as a function of α/β_1 for fixed $P_{\text{in}}/P_{\text{crit}} = 1.35$ and fixed $\tau_3/\tau_1 = 0.7$; in this case, because it is not at the critical power, there are no 100%-efficiency solutions. 79

- 2-12 Asymptotic steady-state efficiency at point C (triply-stable) in the phase diagram (Fig. 2-7), with the initial conditions perturbed from the 100%-efficiency stable solution. The initial amplitudes A_{10} and A_{30} are perturbed by δA_{10} and δA_{30} , respectively, with $\delta A_{10}/A_1^{\text{crit}} = \delta A_{30}/A_3^{\text{crit}}$. The oscillation of the steady-state efficiency with the perturbation strength is an indication of the complexity of the phase space and the shapes of the basins of attraction of each fixed point. 80
- 2-13 One way of exciting the system into a controlled stable solution: the input power is the sum of an exponential turn-on (the blue curve, P_1) and a Gaussian pulse with amplitude P_0 and width δT . The amplitude P_0 is altered to control which stable solution the system ends up in. 81
- 2-14 (Left:) Steady-state efficiency at point C in Fig. 2-7 as a function of the transient input-pulse power P_0 from Fig. 2-13, showing how all three stable solutions can be excited by an appropriate input-pulse amplitude. (Right:) Same, but for an asymptotic input power $P_1 \approx 0.8P_{\text{crit}}$, for which the maximum efficiency is $\approx 90\%$ from Fig. 2-11(right), but is easier to excite. . . . 82
- 2-15 (Left:) Green line with arrows indicates instantaneous “efficiency” (harmonic output power / input power) as the input power is slowly decreased, starting at a power $\approx 1.7P_{\text{crit}}$. For comparison, Fig. 2-11(left) is superimposed as solid-red and dashed-blue lines. The solution “adiabatically” follows a steady state until the steady state becomes unstable, at which point it enters limit cycles, and then returns to a high-efficiency steady state, and finally goes drops to a low-efficiency steady-state if the power is further decreased. (Right:) Similar, but here the power is *increased* starting at the high-efficiency steady state solution for $P < P_{\text{crit}}$. In this case, it again enters limit cycles, but then it returns to a high-efficiency steady-state solution as the power is further increased, eventually reaching the 100%-efficiency stable solution. If the power is further increased, it drops discontinuously to the remaining lower-efficiency steady-state stable solution. 83

2-16 Schematic of general scheme for $\chi^{(2)}$ difference-frequency generation, including dynamical variables for coupled-mode equations: two input (output) channels (with incoming (outgoing) field amplitudes s_{\pm}) are coupled to a resonant cavity with three modes at frequencies ω_1 and ω_2 and $\omega_T = \omega_2 - \omega_1$ (with corresponding amplitudes a_1 , a_2 , and a_T). The two resonant modes are nonlinearly coupled by a Pockels ($\chi^{(2)}$) nonlinearity. All three modes leak energy at a rate determined by their corresponding quality factors Q 87

2-17 (a) Schematic illustration of triply-resonant monolithic (GaAs) geometry. $P_{1,\text{in}} = |s_{1+}|^2$ and $P_{2,\text{in}} = |s_{2+}|^2$ denote the input power of light at pump ω_1 and idler ω_2 frequencies, respectively, while $P_T = |s_{T-}|^2$ stands for the THz output power. The transmitted power at the pump and idler frequencies are denoted by $P_{1,\text{tr}}$ and $P_{2,\text{tr}}$, respectively. (b) Main panel: electric field profile E_z corresponding o the resonant mode appearing at 1 THz in the dielectric structure of (a). The value of the geometrical parameters are: $a = 102\mu\text{m}$, $d_1 = 40.8\mu\text{m}$, $d_2 = 25.1\mu\text{m}$, $d_3 = 18.8\mu\text{m}$, and $w = 0.8\mu\text{m}$. (Inset:) Enlarged view of the x -polarized electric field profile E_x , corresponding to a whispery-gallery mode with $m = 572$ circulating inside the dielectric ring, shown in the main figure. The internal/external radii of the ring resonator are $30.5\mu\text{m}$ and $40.1\mu\text{m}$, respectively. Shaded/white areas represent GaAs/air regions. . 91

2-18 Total efficiency (blue) of THz light $|s_{T-}|^2/(|s_{1+}|^2 + |s_{2+}|^2)$ emitted by the system, for $\omega_T = 1$ THz. The results also show the depletion efficiency $|s_{k-}^2|/(|s_{1+}|^2 + |s_{2+}|^2)$ of the pump ($k = 1$, green line) and idler ($k = 2$, red line) photons. The horizontal dashed line displays the maximum possible THz (quantum-limited) conversion efficiency, determined by the Manley-Rowe relations. The dotted line displays the effect of realistic absorption losses on the conversion efficiency. 93

3-1 Schematic illustration of the system investigated. A waveguide is coupled to a cavity with an EIT atom at its center. In the upper left-hand corner, an FDTD simulation that can be used to calculate the model parameters is shown. 100

- 3-2 Waveguide reflection for a lossless 3-level EIT atom for the four labeled values of the atomic coupling strength g_{13} , in GHz. The radiation rate $\Gamma_w = 21.5$ GHz and the ratio $g_{13}/\Omega_c = 2$ are fixed. Larger g_{13} produces larger peak separations (the blue curve shows Rabi peaks outside of the plot), favorable for switching. 104
- 3-3 Waveguide reflection (blue) and absorption (red) in the absence (solid) and presence (dashed) of an control photon, demonstrating nonlinear single-photon switching ($\Gamma_w = 21.5$ GHz, $g_{13} = 20.5$ GHz, $\Omega_c = 2$ GHz, $\Gamma_3 = 30$ GHz, $g_{24} = 8$ GHz, and $\Delta\tilde{\omega}_{24} = 30$ GHz). 105
- 3-4 Waveguide reflection with (dashed) and without (solid) a control photon, demonstrating lossless switching. (Left:) A higher quality factor has made the resonant peaks narrow enough to be shifted by more than their full width at half maximum by a single photon ($\Gamma_w = 3$ GHz, $g_{13} = 20.5$ GHz, $\Omega_c = 30$ GHz, $\Gamma_3 = 0$, $g_{24} = 30$ GHz and $\Delta\tilde{\omega}_{24} = 20$ GHz). (Right:) Multiple (49) EIT atoms have been used to push the Rabi-split peaks farther away in the presence of negligible loss ($\Gamma_w = 21.5$ GHz, $g_{13} = 143.5$ GHz, $\Omega_c = 210$ GHz, $\Gamma_3 = 0$, $g_{24} = 210$ GHz and $\Delta\tilde{\omega}_{24} = 20$ GHz). 106
- 3-5 A 2D triangular lattice of air holes in dielectric ($\varepsilon = 13$). On top of the dielectric structure in grey, the E_z field is plotted, with positive values in red, and negative values in blue. A small region of nonlinear material is placed exactly in the center of the structure. This material may be, for example, either two-level atoms, quantum wells, or some semiconductors such as InSb. 112
- 3-6 Relative enhancement of the TM local density of states for Fig. 3-5, as measured in the time-domain simulation rate of emission, Γ , normalized by the emission rate in vacuum, Γ_o 113
- 3-7 Contour plot of Kerr enhancement $\eta \equiv \text{Re} \chi_{\text{purcell}}^{(3)} / \text{Re} \chi_{\text{hom}}^{(3)}$ as a function of probe (ω_{ph}) and electronic transition (ω_{elec}) frequencies, for a single quantum well of GaAs-AlGaAs, (a) at $T = 200$ K, with $0.1\gamma_{\text{phase}} = 10\Gamma_{\text{nr}} = \Gamma_{\text{rad}}$, and (b) at $T = 225$ K, with $0.1\gamma_{\text{phase}} = \Gamma_{\text{nr}} = \Gamma_{\text{rad}}$ 113

4-1	(Left:) Two atoms separated by a distance d are interacting via fluctuation-induced dipole moments (polarizabilities α_1 and α_2). (Right:) Two macroscopic objects separated by a distance d are interacting via the electromagnetic field produced by the fluctuation-induced charges on the interior and surface of the bodies, described by macroscopic polarization fields \mathbf{P}_1 and \mathbf{P}_2 .	119
5-1	Schematic of one dimensional geometry, showing two 1d metal plates separated by a distance a , embedded in a computational cell of length $L = a + 4a$, with periodic boundary conditions, and resolution Δx .	132
5-2	(Left:) Plot of force summand, or spectral density, $\hbar\Delta\omega/2\Delta x$ vs. ω for 1d parallel metal plates from Fig. 5-1. (Right:) Plot of force partial sum $\sum^\omega \hbar\Delta\omega/2\Delta x$ vs. ω .	133
5-3	Plot of Casimir force integrand dU/da between two 1d parallel plates separated by a distance $a = 1$ versus imaginary frequency $w = \text{Im}\omega$, using the method of Sec. 5.3.	137
5-4	Schematic illustration of the finite-difference Yee grid: the vector components of the electric \mathbf{E} and magnetic \mathbf{H} fields are spatially staggered about a rectangular unit cell of a Cartesian computational grid so that each \mathbf{E} -field vector component is located midway between a pair of \mathbf{H} -field vector components, and vice versa. The electric field components form the edges of the cube, and the magnetic field components form the normals to the faces of the cube.	140
5-5	Schematic illustration of a possible contour around a body; the force on the body is given by an integral of the stress tensor around this contour.	141
5-6	Plot of Casimir force integrand $\oint\mathbf{T} \cdot d\mathbf{A}$ between two 1d parallel plates separated by a distance $a = 1$ versus imaginary frequency $w = \text{Im}\omega$, using the stress-tensor method of Sec. 5.4.	147

- 5-7 Casimir force between a 3d radius- R cylinder and a plate (inset), relative to the proximity-force approximation F_{PFA} , vs. normalized separation a/R . The solid lines are the Casimir force computed in Ref. [141] for TE (gray) and TM (blue) polarizations, along with results computed by our method with a simple finite-difference discretization (gray squares). Error bars were estimated for some data points by using computations at multiple spatial resolutions. Inset shows interaction stress tensor $\Delta\langle T_{xx} \rangle$ at a typical imaginary frequency $\omega = 2\pi c/a$, where red indicates attractive stress. 148
- 6-1 Schematic illustration of correspondence between two methods of calculating Casimir forces. (Left): Numerical method requiring evaluation of the force integrand over the imaginary-frequency axis (or *some* suitable complex-frequency contour). (Right): Antenna (S -matrix) measurements of the electromagnetic response at tabletop (e.g. cm) lengthscales. Here, the effect of a contour deformation is achieved by a material deformation which corresponds to the presence of a conductive fluid between the objects. 156
- 6-2 Complex-frequency ω plot of the Casimir force integrand ($\ln |\text{Re } dF_x/d\omega|$), where $dF_x/d\omega$ is in units of \hbar/d^2 , for the geometry of Fig. 6-8. As the real- ω axis is approached, the integrand becomes highly oscillatory, which is only partially revealed here due to the finite frequency resolution. Various integration contours of interest are labeled as black and dashed lines. (Inset:) Vacuum $\epsilon = 1$ contour deformations $\omega(\xi)$ and their corresponding (real-frequency) physical realizations $\epsilon_c(\xi)$ 160

6-3	(Top:) Partial force integral $\int_0^\xi dF_x$, normalized by F_x , as a function of ξ , for the various $\omega(\xi)$ -contours (equivalently, various $\varepsilon_c = \omega^2/\xi^2$) shown in Fig. 6-3. The solid green, red and black lines correspond to conductive media with $\sigma = 10, 10^2$ and 10^3 , respectively (σ has units of c/d). The dashed grey and solid blue lines correspond to $\phi = \pi/4$ and $\phi = \pi/2$ (Wick) rotations. (Bottom:) Illustration of the required frequency bandwidth for a possible realizations using saline solution at separation $d = 30\text{cm}$. The red lines plot the xx -component of the photon GF G_{xx} at a single location on the surface contour (see inset) as a function of ξ (GHz). The black line is the corresponding partial force integrand.	164
6-4	$ \text{Im} g(t) $ for various values of σ , illustrating the transition from t^{-1} to $t^{-1/2}$ power-law decay as σ increases. Because there are strong oscillations in $g(t)$ at the Nyquist frequency for intermediate σ , for clarity we plot the positive and negative terms in $g(t)$ as separate components.	178
6-5	Plot of the force error (force after a finite time integration vs. the force after a very long run time) for $g(t)$ determined from a numerical transform as in Ref. [449] and from the analytic transform of the high-frequency components. Inset: $\text{Im}[g(-t)]$ obtained without a cutoff, in which the high-frequency divergence is integrated analytically. Compare with Fig. 1 of Ref. [449]	180
6-6	$\Gamma_x^E(t) + \Gamma_x^H(t)$ for a set of one-dimensional parallel plates as the separation h is varied. The inset shows the physical setup.	181
6-7	Partial force as defined in Eq. (6.37) for one-dimensional parallel plates as a function of time t . (Inset): Relative error $\Delta(t)$ as a function of t on a semi log scale.	183
6-8	Schematic illustration of a possible contour around a body; the force on the body as computed by Eq. (6.49), involves an integral of the fields over the contour S	189

- 7-1 Casimir force between 2d (z -invariant fields) metal squares F/F_{PFA} , vs. distance from metal plate h (inset), normalized by the total force (TE+TM) obtained using the PFA, $F_{\text{PFA}} = \hbar c \zeta(3) s / 8\pi a^3$. The total force is plotted (black squares) along with the TE (red dots) and TM (blue circles) contributions. 195
- 7-2 (a-f): TM stress map of the geometry in Fig. 7-1 for various h . The interaction stress tensors $\langle T_{xx} \rangle$ (left) and $\langle T_{xy} \rangle$ (right) for: (a),(d): $h = 0.5a$; (b),(e): $h = a$; and (c),(f): $h = 2a$, where blue/white/red = repulsive/zero/attractive. 196
- 7-3 Casimir force between 2d (z -invariant fields) dielectric ($\epsilon = 4$, left) or gold (right) squares F/F_{PFA} , vs. distance from metal plate h (inset), normalized by the total force (TE+TM) obtained using the PFA (Here, the PFA force is computed for x -infinite slabs of corresponding dielectric ϵ). The total force is plotted (black squares) along with the TE (red dots) and TM (blue circles) contributions. 197
- 7-4 Casimir force per unit length between z -invariant 3d metal blocks F/F_{PFA} , vs. distance from metal plate h (inset), normalized by the total force (TE+TM) obtained using the PFA, $F_{\text{PFA}} = \hbar c s \pi^2 / 480 a^4$. The total force is plotted (black squares) along with the TE (red dots) and TM (blue circles) contributions. 198
- 7-5 Schematic of cylinder-plate geometry consisting of two z -invariant metallic cylinders of radius R , separated by a distance a from each other, and by a distance h from one (top) or two (bottom) metallic sidewall(s). 200
- 7-6 Casimir force per unit length between two cylinders (black) vs. normalized cylinder-sidewall separation h/R , at fixed $a/R = 2$, normalized by the total PFA force per unit length between two isolated cylinders $F_{\text{PFA}} = \frac{5}{2}(\hbar c \pi^3 / 1920) \sqrt{R/a^7}$ [425] for the two cases shown on the inset. Also shown are the individual TE (red) and TM (blue) forces. *Inset*: Schematic of geometry in the case of one (dashed lines) and two (solid lines) sidewalls. . . . 201

7-7	Casimir force per unit length between two cylinders of fixed radius R vs. normalized cylinder-sidewall separation h/R (for one plate), normalized by the total PFA force per unit length between two isolated cylinders $F_{\text{PFA}} = \frac{5}{2}(\hbar c\pi^3/1920)\sqrt{R/a^7}$. The force is plotted for different cylinder separations $a/R = 0.2, 0.6, 1.0, 2.0, 3.0,$ and 4.0	202
7-8	Casimir force per unit length between a plate and two cylinders of fixed radius R vs. normalized cylinder-cylinder separation a/R , normalized by the total PFA force per unit length between two isolated cylinders $F_{\text{PFA}} = \frac{5}{2}(\hbar c\pi^3/960)\sqrt{R/2a^7}$ [425]. The force is plotted for different plate separations $h/R = 0.28, 0.6, 1.0,$ and 2.0	203
7-9	Force for the double cylinders of [428] as a function of sidewall separation h/a , normalized by the PFA force $F_{\text{PFA}} = \hbar c\zeta(3)d/8\pi a^3$. Red/blue/black squares show the TE/TM/total force in the presence of metallic sidewalls, as computed by the FDTD method (squares). The solid lines indicate the results from the scattering calculations of [428], showing excellent agreement. Dashed lines indicate the same force components, but in the presence of perfect magnetic-conductor sidewalls (computed via FDTD). Note that the total force is nonmonotonic for electric sidewalls and monotonic for magnetic sidewalls.	207
7-10	Schematic of a two-dimensional geometry: two metal squares $s \times s$ separated by a distance a , and separated from two adjacent metal sidewalls by a distance h	208
7-11	Casimir even (red) and odd (blue) forces vs. sidewall separation h normalized by the PFA force $F_{\text{PFA}} = -\hbar c\zeta(3)s/8\pi a^3$ (dashed black), computed using the ray-optics (solid) and stress-tensor (dashed) methods. Note that the ray-optics results become exact as $h \rightarrow 0$	213
7-12	Casimir force vs. sidewall separation h normalized by the PFA force $F_{\text{PFA}} = -\hbar c\zeta(3)s/8\pi a^3$, computed using the ray-optics (solid) and stress-tensor (dashed) methods. The Neumann (green), Dirichlet (orange) and total (black) forces are all normalized by the total Neumann+Dirichlet PFA force.	214

7-13	Casimir force vs. separation between squares a with constant sidewall separation $h = 0.25$, normalized by the $F_{\text{PFA}} = -\hbar c \zeta(3) s / 8\pi a^3$ (top) and the corresponding $h = 0$ force (bottom). The force is computed for both even (red) and odd (blue) contributions, separately. <i>Inset</i> : schematic of geometry consisting of two isolated squares with adjacent sidewalls.	216
7-14	Schematic of general 2d squares+sidewalls lattice. Lines extending from solid circles unto solid circles represent even reflection paths. Lines extending from open solid circles unto open circles represent odd reflection paths. (Here, $h/a \approx 0.2$ and $s/a \approx 1$.) A possible even (red) and forbidden odd (blue) path is shown.	217
7-15	Schematic of three-reflection paths. Blue/red represent $(1, 2)/(2, 1)$ paths, and the lengths ℓ_i shown are used the calculation of the energy.	221
7-16	Convergence error $(\mathcal{E}_{r+1} - \mathcal{E}_r)/\mathcal{E}_r$ in the even (solid) and odd (dashed) path contributions vs. reflection order r for values of $h = 0$ (red), $h = 0.01$ (blue), and $h = 0.1$ (black) (here, \mathcal{E}_r means the energy computed up to order r). <i>Inset</i> : Absolute error $(\mathcal{E}_{\text{exact}} - \mathcal{E}_r)/\mathcal{E}_{\text{exact}}$ for both even (solid black) and odd (dashed red) contributions for $h = 0$	225
7-17	Three-dimensional schematic of the Casimir “zipper” geometry of interlocking metal brackets (shown in different colors for illustration only), along with a two-dimensional xy cross-section. The dashed lines extruding from the plates to the squares indicate their out-of-plane connectivity.	227
7-18	(Top:) Plot of the Neumann (blue, TE), Dirichlet (red, TM) and total (black, TE+TM) Casimir pressure (in units of $\hbar c/\Lambda^4$) between the objects of Fig. 7-17, as a function of a/s . The inset illustrates a two-dimensional cross-section. (Bottom:) Schematic indicating the various qualitatively different Casimir force regimes between the two structures.	229
7-19	Comparison of Casimir pressure (in units of $\hbar c/\Lambda^4$) as a function of a/s between the stress-tensor (exact) numerical results (black squares) and the proximity-force approximation (solid green). Also shown are the individual square-square (dashed blue) and square-plate (dashed orange) contributions to the PFA force. Inset: Schematic illustration of the chosen PFA “lines of interaction” between squares (dashed black lines).	232

8-1	Schematic of Lifshitz geometry, consisting of two semi-infinite plates (permittivities ε_1 and ε_2) separated by a dielectric with permittivity ε_3	239
8-2	Casimir force F_x in the x -direction, per unit z -length, on a SiO_2 cylinder suspended within a perfectly-metallic cylinder (inset), separated by fluid (ethanol), as a function of the x -displacement from equilibrium (eccentricity) d [in units of $a = 0.5(D - s)$] for both circular (solid-line) and square (dashed-line) cylinders. $d = 0$ is seen to be a stable equilibrium.	243
8-3	Casimir torque τ_z , per unit z -length, on the inner square rod of eccentric square cylinders as a function of the angle θ with respect to the x -axis (see right insets) for two material choices as shown in the top insets. Error bars are estimates of the effect of finite grid resolution.	244
8-4	Derivative of the Casimir torque in the z -direction with respect to θ , $d\tau_z/d\theta$, in units of $\hbar c/a^2$, evaluated at $\theta \approx 0^\circ$, as a function of lengthscale $a = 0.5(D - s)$. The solid line is a fit to the exact Casimir torque, computed by our numerical method (solid squares), and the dashed line is the torque as computed by a PFA approximation. Both curves display a change (dispersion-induced transition) in the stable orientation of the square.	245
8-5	Casimir force F_x in the x -direction, per unit area, between a planar SiO_2 slab suspended between two perfect-metal plates (solid red and black lines) or gold half-spaces (dashed red line), separated by a fluid (ethanol), as a function of the dimensionless x -displacement d from equilibrium. The force is plotted also at two lengthscales $a \equiv 0.5(D - s) = 0.0955 \mu\text{m}$ (red lines) and $a = 0.6 \mu\text{m}$ (black line). The solid blue line shows the same quantity for the case of a lithium-niobate slab and a gold half-space, for $a = 0.5 \mu\text{m}$. (<i>Insets:</i>) Top inset: $\varepsilon(i\xi)$ for SiO_2 (black), LiNbO_3 (solid and dashed blue), ethanol (orange) and gold (red) as a function of imaginary frequency ξ . Bottom inset: schematic of geometry.	247
8-6	Schematic of silica sphere levitated above a gold plate. At the equilibrium point, the force of gravity counters the Casimir force, while the Casimir force from the walls of the spherical indentation confine the sphere laterally (x - y plane).	249

8-7	(Left:) Total (Casimir + gravity) vertical (z) force on the silica sphere (depicted in the inset) as the height h of the sphere's surface above the indentation surface is varied. The point of vertical equilibrium occurs at $h_c \sim 450$ nm. (Right:) Casimir restoring force on the sphere as a function of lateral displacement Δx , when the vertical position is fixed at h_c , the height at which gravity balances the Casimir force.	250
8-8	Schematic schemes for stable suspension of fluid-separated objects, involving: (a) enclosed geometries; (b) gravity countering Casimir repulsion; and (c) material dispersion producing repulsive <i>and</i> attractive Casimir forces (here).	251
8-9	(Top:) Plot of the dielectric permittivity $\varepsilon(i\xi)$ of various materials evaluated at imaginary frequency ξ (units of $c/\mu\text{m}$). (Bottom:) Casimir force between a semi-infinite slab and a sphere (dots) and Casimir pressure between semi-finite slabs (lines), normalized by the corresponding perfect-metal PFA force $F_{\text{PFA}} = \hbar c\pi^3/720d^3$ (slab-sphere) and pressure $F_{\text{PFA}} = \hbar c\pi^2/240d^4$ (slab-slab). The force/pressure is plotted for various material configurations, described in the text.	253
8-10	Equilibrium separation d_c vs. sphere-radius (R) or slab-thickness (t) for various slab-sphere (dots) and slab-slab (lines) configurations. Dashed lines correspond to slab-slab geometries with materials swapped (semi-infinite silicon slab).	255
8-11	(Top:) Plot of the stable equilibrium center-surface (L_c) and surface-surface (h_c) separation between either a teflon (yellow) or silicon (blue) sphere and a semi-infinite gold slab (depicted schematically on the left inset), as a function of sphere radius R . The black lines also show the presence of an unstable equilibrium in the silicon-sphere case. (Bottom:) Plot of the force F_{SS} between two teflon/silicon spheres of radii R_S/R_T , showing the existence of a stable equilibrium.	256

8-12 Relative permittivity $\varepsilon(i\xi)$ of various materials as a function of imaginary frequency ξ (in units of $c/\mu\text{m}$) or “Matsubara” temperature $T = \hbar\xi/2\pi k_B$. Doped silicon corresponds (bottom to top) to doping density $\rho_d = \{1, 3, 5, 10, 10^2\} \times 10^{16}$, modeled via an empirical Drude model [127], as is gold [34]. Water, polystyrene, ethanol, teflon, and lithium niobate are all modeled via standard Lorentz-oscillator models [325]. 260

8-13 Equilibrium separation d_c (in units of μm) as a function of temperature T (in Kelvin), for a geometry consisting of fluid-separated semi-infinite slabs (no gravity). The various curves correspond to d_c for various material combinations. Solid/dashed lines correspond to stable/unstable equilibria, and shaded regions are T where ethanol is non-liquid at 1 atm [161]. Doped-silicon is plotted for various doping densities $\rho_d = \{1, 10, 100, 500, 10^3\} \times 10^{17}$. 261

8-14 Equilibrium position d_c (in units of μm) of a semi-infinite polystyrene (PS) slab immersed in ethanol (shaded $T = \text{non-liquid}$) and suspended against gravity by a repulsive Casimir force exerted by a doped-silicon (Si) slab. The solid/dashed lines correspond to stable/unstable d_c , and each color represents a different value of PS slab-thickness h (in units of μm). The inset shows the magnitude of the total energy $U_T(d)$ (in units of $k_B T$) as a function of d for $h = 150$ nm, at various T 262

8-15 Average separation $\langle d \rangle$ (circles) and equilibrium separation d_c (red line), in units of μm , vs. temperature T (in Kelvin), for a geometry consisting of a fluid-separated hollow PS sphere of inner/outer radius $r/R = 3.2/5 \mu\text{m}$ suspended in ethanol against gravity by a doped-silicon slab and subject to Brownian motion. Shaded region indicates where sphere is found with 95% probability. The thin black line is the average $\langle d \rangle_{140}$ if the Casimir energy at 140 K is used instead of the true temperature-dependent energy landscape. 263

A-1 Cavity mode E_z (blue/white/red = positive/zero/negative) for a 2d structure with (a) no disorder, (b) surface disorder, and (c) surface + position disorder, and for a (horizontal) cross-section of a 3d structure with (A) no disorder and (B) sidewall disorder. 268

A-2	Q vs. # layers for a 2d cavity (inset shows 2 layers), for disorder types (descending by Q): position (red), no disorder (black), ε (cyan), surface+position (blue), and surface (magenta). Error bars show std. dev. over 15 runs.	270
A-3	2d band gap and cavity mode frequencies vs. strength of ε disorder (maximum % variation of rod ε). Error bars show std. dev. for 10/100 runs for gap/cavity frequencies. Inset shows cavity structure (4 layers).	271
A-4	Q vs. # layers for cavity in the 3d structure (left inset), for disorder types (descending by Q): no disorder (black), height (blue), sidewall roughness (red). Right inset: mode $\varepsilon \mathbf{E} ^2$ in vertical cross-section (3 layers). Error bars show std. dev. over 15 runs.	271
B-1	Unit cell of the Fibonacci superspace dielectric. The physical dielectric is obtained by taking a slice at an angle $\tan \phi = \tau$. Black/white are the dielectric constants of the structure factor material and air, chosen to be $\varepsilon = 4.84$ and $\varepsilon = 2.56$, respectively.	276
B-2	(Left:) Frequency spectrum ω of the Fibonacci quasicrystal vs. “wave-vector” k_x . The blue lines indicate spurious states which arise due to finite resolution effects (see text). (Right:) Corresponding density of states $\rho(\omega)$ computed using a transfer-matrix technique with a supercell of 10^4 layers.	281
B-3	Integrated density of states (DOS) vs. band index (normalized by resolution), for various resolutions. The dashed red, diamond blue and solid black lines denote resolutions of 20, 50 and 200, respectively.	282
B-4	Plot of the magnetic field amplitude $ H_z $ for a band-edge state taken along a slice of the two-dimensional superspace (in the ϕ direction). <i>Inset:</i> Two-dimensional superspace field profile (red/white/blue indicates positive/zero/negative amplitude).	283
B-5	Electric field energy distribution of the band edge states of gaps 1 and 2 in Fig. B-2. Although they have a complex small-scale structure, the large-scale variation is easily understood in terms of the structure of the superspace.	284

B-6	Enlarged view of the Fibonacci spectrum (Fig. B-2) showing a gap with a spurious band crossing it. Insets show the magnetic field $ H_z $ for the spurious band at various k_x —the localization of this mode around the X -parallel edges of the dielectric indicate that this is a discretization artifact.	285
B-7	Dielectric for the Fibonacci chain with $\varepsilon = 2.56$ (light blue), and a defect—an additional $\varepsilon_d = 8.0$ layer, shown in gray.	286
B-8	Varying the defect epsilon for resolutions 50 (blue) and 100 (red). The thickness of the defect is fixed to 0.02 lattice constants. The number of spurious modes increases with the resolution, the true defect state being the lowest of these modes.	287
B-9	Semi-log plots of the magnetic field magnitude H_z for the lowest (top) and highest (bottom) defect state for the configuration shown in Fig. B-7. <i>Insets:</i> Two-dimensional superspace visualizations of the defect states. Note the additional node in the lower figure (corresponding to an unphysical oscillation).	288
B-10	Projected band structure vs. cut angle ϕ , showing different one-dimensional quasicrystal realizations. The vertical red line indicates the spectrum when the slope is the golden ratio τ (the spectra of ϕ and $\pi - \phi$ are equivalent)	289
B-11	Schematic showing the superspace slice X (left) and the projected slice modulo 1 into the unit cell \bar{X} (right), along with the intersection \bar{T} of \bar{X} with the $s = 0$ hyperplane.	291

List of Tables

6.1 Vacuum $\varepsilon = 1$ contour deformations $\omega(\xi)$ and their corresponding (real-frequency) physical realizations $\varepsilon_c(\xi)$	163
---	-----

Chapter 1

Introduction

Controlling and harnessing light has been a goal of many scientists and engineers for the past century. Beginning with the work of James Maxwell and others [8, 265, 347, 380], our understanding of light has led to a number of remarkable technological advances. By manipulating light using materials, scientists revolutionized the telecommunications industry and paved the way for many important technologies involving fiber-optic cables [336], antennas [20, 205], lasers [493], and dozens of different spectroscopic methods [114].

The physics of classical electromagnetic fields can be accurately described by a set of simple coupled partial differential equations, known as Maxwell's equations. Originally, research in classical electromagnetism was confined to the study of light in simple geometries, largely due to a lack of computational resources¹ and techniques capable of solving these equations in complicated settings. On the one hand, a basic understanding of the physics of light in simple geometries, such as in one-dimensional periodic structures or index-guided waveguides, was sufficient to usher a flood of discoveries and applications. On the other hand, if our experience in physics has taught us anything is that simple wave phenomena (e.g. interference, dispersion, diffraction) in complicated geometries can lead to new and unexpected phenomena. Toward that end, especially spurred by the growing availability of general-purpose computers in the 1960s, physicists and engineers set out to formulate ways by which to solve Maxwell's equations in arbitrary geometries. These efforts led to a wealth of computational techniques for solving various classes of complicated wave equations (even beyond the realm of electromagnetism) [25, 45, 91, 237, 252, 515, 531, 532, 548, 581, 596] that

¹The study of index-guiding, antennas and lenses predates the invention of computers.

prompted the discovery of many fundamentally new electromagnetic effects. Over the past few decades, the study of light in complicated geometries with features at the scale of the electromagnetic wavelength has revealed fundamentally new ways to manipulate and guide light ² [155, 222, 238, 241, 290, 512, 584], such as periodic structures with photonic band-gaps that allow confinement of light in small volumes for very long times [95, 158, 241, 395, 397, 420, 511, 537, 544, 545] as well as new forms of electromagnetic dispersion [188, 223, 224, 241, 335, 584]. These discoveries have in turn led to the development of new techniques and devices: guiding of light in air with low losses [58, 155, 156, 222, 247, 270, 290, 317], new types of optical insulators [241], lasers [316], linear filters [6, 7], and low-power logic gates [377], to name a few.

The realization of greater control and manipulation of light by way of geometry alone has also paved the way for new developments in the field of nonlinear optics. When the intensity of light interacting with a polarizable material exceeds a certain (material-dependent) threshold, its dielectric response can no longer be described by a linear permittivity, and the corresponding Maxwell’s equations become nonlinear [51, 469], coupling photons via material-mediated interactions ³. Although these nonlinearities are often relatively weak, and therefore negligible, they can give rise to a number of fascinating and important electromagnetic effects with many applications [14, 145, 308, 345, 366, 469, 477]. The key to strengthening nonlinear interactions is to require that the nonlinear “mixing” occur inside a resonant environment: a cavity or geometry in which light is confined for long times. The smaller the cavity (compared to the wavelength of light) and the longer the modal lifetime, the stronger the field intensity and the more time the light has to interact with it-

²The appendix section of the thesis provides examples of the kinds of novel classical electromagnetic phenomena that can arise in complicated geometries. For example, appendix A studies the influence of fabrication-induced disorder on a class of photonic microcavities formed by the creation of a defect in an otherwise perfectly-periodic photonic crystal environment (periodic structures whose period is on the order of the wavelength of light) [289, 419]. It is shown that the electromagnetic modes of these kinds of microcavities are immune to relatively large amounts of disorder, all the way up to the Anderson localization regime [444].

³A simple (heuristic) microscopic picture suffices to understand the origin of nonlinearities: macroscopic objects are composed of many atoms, each consisting of charged constituents. When incident light interacts with these charges, it gives rise to a redistribution of charge (momentary dipoles) that in turn radiate new light. The sum of the incident and radiated light can then be used to define an effective macroscopic electromagnetic field in the medium, which enters Maxwell’s equations via the polarization field. For low-intensity light (composed of few photons), this response is linear in the incident field amplitude; anthropomorphically speaking, the atom–photon ratio is so large that each photon has at least one atom all to itself. For high-intensity light (composed of many photons), the response becomes nonlinear in the incident amplitude; the atom–photon ratio is low that multiple photons begin to compete for the same atom’s attention. Namely, the interaction of a photon with an atom will induce a new charge redistribution on the atom, so that another photon simultaneously interacting with the atom will experience a different atom–photon interaction.

self. These two objectives were often at odds with each other in conventional cavities based on index-guiding (such as ring resonators). However, nanophotonic media (with complicated features at the scale of the electromagnetic wavelength), such as new generations of micro-cavities, can simultaneously achieve small modal volumes and long lifetimes, independent of one another [211, 248, 459, 509, 546, 549]. This ability not only allows a decrease in the power requirements of many nonlinear devices, but also leads to novel phenomena absent in conventional settings, such as optical bistability [81, 175, 387, 501, 507, 526], limit cycles [198], and many other dynamical effects. As the power requirements for nonlinear interactions continue to decrease, it becomes possible to even reach the single-photon level, where quantum-mechanical effects of light become important, leading to the possibility of observing strong photon-photon interactions. For example, recent developments in nanofabrication led to the observation of a phenomenon known as “photon blockade” in which a photon interacting with an atom “blocks” the interaction of another photon with the same atom [12, 149, 482]. The study of nonlinear interactions at low power levels has been at the frontier of many research areas, such as quantum information [379, 549]. Advances in micro-fabrication techniques and developments in our understanding of nonlinear processes in complicated geometries has led many researchers to dream about the not-so-distant possibility of realizing ultrafast integrated optical devices [57, 216, 346], similar in scope and functionality to the electronic transistors, switches and logic gates that preceded the birth of modern computers.

Although our ability to study light in complicated geometries has grown considerably, there remain many aspects of nonlinear optics yet unexplored, especially in the context of nanophotonic media (at the λ -scale). One problem of interest that has yet to be fully explored is a process known as nonlinear frequency conversion: input light at some frequency is converted to output light at a different frequency by a nonlinear interaction. These kinds of processes are at the heart of many up-and-coming technologies in the realm of optical devices, such as high-frequency sources [33], ultracompact-coherent optical sources [145, 308, 366], imaging [477], spectroscopy [345], and let us not forget the light-emitting diodes that permeate many of our television screens. Early attempts at reducing the power requirements of frequency conversion focused on large-etalon cavities that do not fully take advantage of the spatial and temporal confinement present in micro-photonic cavities, such as new generations of photonic-crystal nanobeam cavities [70, 115, 317]. More recently, there has

been renewed interest in studying these nonlinear processes in more complicated structures that offer greater confinement (in the context of frequency conversion, for both input and output light).

The first part of this thesis (chapters 2 and 3) studies how and to what extent one can enhance nonlinear interactions, both for classical and quantum light, by exploiting geometry-induced resonances.

Chapter 2 develops a universal framework for studying nonlinear frequency conversion in resonant environments consisting of arbitrary geometries and materials. This approach is based on a general class of methods known as coupled-mode theories [202, 520], applicable for geometries consisting of weakly-coupled components, e.g. waveguides and cavities. Using this formalism, we investigate some of the possible types of nonlinear behaviors that one can expect in strongly-resonant geometries: under coupled-mode theory, the precise features of the geometries are not important, but are rather captured by a small set of parameters related to the decay rates and modal volumes of the electromagnetic modes of the geometries. An important figure of merit for studying frequency conversion are the coupling coefficients describing the nonlinear interactions. We derive these coupling coefficients in general geometries (not restricted to large-etalon or specific structures, unlike previous work). We find that frequency conversion can be dramatically enhanced by the presence of simultaneous resonant cavities for both input and output light, in agreement with previous predictions in specific geometries. We focus on studying two important nonlinear processes, known as second- and third-harmonic generation, in which light at some input frequency is converted to output light at two and three times that input frequency. Frequency conversion is studied in a general class of geometries consisting of a waveguide directly coupled to a doubly-resonant cavity (resonant at the input and output frequencies). Our results demonstrate the existence of a critical power at which 100% of the input light is converted into output light. We examine the conditions under which this is achieved.

The possibility of achieving 100% conversion in doubly-resonant cavities was previously studied by a number of works (again for particular geometries or large-etalon cavities) in the case of second-harmonic generation. The third-harmonic generation case had only been studied in geometries consisting of waveguides or singly-resonant cavities. This latter case is sufficiently more complicated due to the presence of cross- and self-phase modulation effects: the frequencies of the input and output waves change as a function of the intensity

of the input and output light [51]. We demonstrate that these effects can play a dramatic role on the stability of the conversion process and give rise to dynamical behaviors unique to the doubly-resonant context, including limit cycles, multi-stability, and hysteresis. Obtaining and exciting efficient solutions turns out to require new (geometrical and operational) considerations, such as shifting of the cavity resonant frequencies and priming of the input light. We conclude our discussion of geometry-enabled nonlinear effects by studying two other types of frequency conversion schemes, one of which can be used to generate light at terahertz frequencies. Terahertz light is of consequence to many applications and extremely difficult to generate due to the lack of natural sources operating at room temperature.

Chapter 3 explores a different set of nonlinear problems. In particular, we explain how one may combine geometry-enabled resonances with novel materials to further enhance nonlinear interactions (decreasing their power requirements), so as to reach the single-photon operation regime. Light consisting of a single photon does not behave according to the standard rules of classical electromagnetism, requiring a different theoretical framework. In addition to exhibiting standard classical wave behavior, single photons must also conform to the rules of quantum mechanics, i.e. Heisenberg’s uncertainty principle, which can have dramatic consequences on their behavior [478, 551, 571].

Using a quantized version of the coupled-mode theory framework of chapter 2, a close cousin of the so-called Breit-Wigner scattering approach in quantum mechanics, we study some of the consequences of coupling single photons to single atoms embedded inside cavities. In particular, we consider a class of geometries similar to that of chapter 2, in which a waveguide is side-coupled (rather than directly coupled) to a cavity. An atom whose atomic transition energy is in resonance with the fundamental cavity mode is then embedded inside the cavity. This confluence of geometry and material resonances is a key component for achieving strong coupling between the atom and any resonant photon residing in the cavity [263, 323, 346, 488–490, 510, 534, 571]. In theory, this simple setup presents an excellent opportunity for entangling the fate of photons. In particular, one can imagine sending two photons into the cavity via the waveguide. In the absence of the atom the two photons “miss” each other, and are both simultaneously transmitted. In the presence of the atom, however, the photons can “compete” for the atom’s affection, with one photon “blocking” another photon from interacting with the atom, causing the two transmitted photons to be delayed relative to one another at the output due to the effective “repulsion” mediated by

the atom [12, 43, 149, 263, 482, 571]. Under other circumstances, the photons can instead be made to bond to one another via an atom-mediated “attraction” [117, 487]. A number of these effects have been observed experimentally under low temperature conditions.

Unfortunately, a big problem with many atom–photon interaction schemes is the fact that excited atoms tend to spontaneously radiate energy due to quantum and thermodynamic fluctuations. In particular, charge and electromagnetic fluctuations induced by quantum and thermal effects can exchange energy with otherwise static (stable) charge distributions (excited electrons in atoms) [286, 415, 479, 570, 579]. The consequence for strong coupling are devastating: the entanglement of a photon to incoherent forms of energy (radiative and non-radiative, e.g. phonons) via the atom results in a loss of phase information (possibly even the very photon). By “cooling” down single atoms, it is possible to reduce these undesired consequences of fluctuations, but this obviously presents a problem for applications. One way to go around this problem is to take advantage of an effect known as electromagnetically-induced transparency: an external classical oscillating field allows photons to interact with three-level atoms in a way that makes the photons “immune” to undesirable decay mechanisms [10, 133, 321]. By taking advantage of this effect (explored in chapter 3), we demonstrate the possibility of switching the behavior of a single signal photon by a single gating photon of a different frequency, via a type of cross-phase-modulation effect mediated by a four-level atom inside the cavity (a quantum analogue of the cross-phase modulation effect studied for classical light in chapter 2). The steady-state behavior of the system is solved exactly and analyzed using experimentally accessible parameters, which we use to confirm that the strong coupling regime is required for lossless two-photon quantum entanglement ⁴.

Following our discussion of single-photon switching, we go on to explore yet another consequence of quantum and thermodynamic fluctuations: we demonstrate the possibility of strengthening the nonlinear susceptibility (interaction) of certain classes of nonlinear materials by exploiting the Purcell effect [415, 569]. The Purcell effect, discovered by Edward Purcell in 1946, is based on the simple yet profound statement that the spontaneous emission rate of an atom depends on the photonic environment in which the atom resides—technically speaking, on the density of photonic modes available to the atom. Purcell discovered that

⁴The reason why single photons interact so strongly in the presence of an atom, as opposed to a medium composed of many atoms, is simple: a single photon can single-handedly saturate the response of an atom, whereas the same is not true of many atoms [51].

by placing the atom in an environment (geometry) with modes resonant at the atomic transition energies, one could enhance the rate of spontaneous emission. Specifically, by placing an atom in a microcavity with a resonance at the transition energy of the atom, it is possible to enhance the spontaneous emission rate of the atom by a factor proportional to the lifetime of the cavity resonance and inversely proportional the volume of the corresponding mode [381,382,585]. Note that the qualitative relationship between the lifetime and volume of a geometric resonance and the strength of nonlinear interactions (whether it be mediated via atoms or macroscopic materials) is universal [55, 57, 241, 252, 549, 574]: it is the basis of the nonlinear enhancement of chapter 2 and is directly related to the strong coupling regime of atom–photon interactions discussed above.

Shortly after Purcell’s discovery, another major discovery concerning the effects of fluctuation-induced interactions was made, this time by Dutch physicists Hendrick Casimir and Dirk Polder [78]. Casimir and Polder proposed the existence of a tiny, yet observable force between macroscopic neutral perfectly-metallic parallel plates that arises purely out of the quantum (zero-temperature) fluctuations of the electromagnetic field [77, 302], the same fluctuations that lead to some of the undesirable effects described in chapter 3.

The second part of this thesis (chapters 4–8) revolves around the study of Casimir forces in arbitrary geometries and materials.

The theory of fluctuation-induced interactions between atoms originated in the early 1930s with the works of Ref. [134, 318]. It was then recognized that thermodynamic and quantum fluctuations between atoms separated by large distances contributed significantly to what are now known as van der Waals forces [68, 77, 78, 116, 302, 318, 352]. In particular, even in its ground state, the electric charge on an atom will tend to fluctuate due to thermal agitations. The stochastically fluctuating charges will in turn cause electromagnetic radiation, and this radiation will consequently cause the otherwise neutral atoms to interact with one another. At zero-temperature, the fluctuations remain, but this time due to quantum-mechanical effects (the Heisenberg uncertainty principle). The first prediction of short-range van der Waals or London-dispersion forces between atoms was made in the late 1920s [134]. Casimir and Polder generalized these predictions to include the case of atoms separated by large distances (compared to the wavelength of their corresponding transitions), which required a full electro-dynamical treatment of the “retardation effects” associated with the finite speed of light [78]. Casimir however, generalized this prediction even further by

also considering the possible interactions between macroscopic object (consisting of many atoms). The predicted force between two parallel metallic plates is considerably different from that of two atoms, owing to the strong dependence of the equilibrium fluctuating electromagnetic fields on the shapes and material composition of macroscopic objects. For example, the force between two metal plates was shown to be attractive and decaying as a quartic power-law of the plate–plate separation [77]. This strong geometry dependence of the force has been a subject of considerable attention among physicists interested in fluctuation-induced interactions.

It took a number of decades, and many developments in the field of nanofabrication, for the tiny Casimir force to go from being an amusing and certainly fundamentally important consequence of quantum fluctuations, to being an object of concern for engineers working at the frontiers of nano-devices. In particular, the magnitude of the Casimir force between two $1\mu\text{m}\times 1\mu\text{m}$ parallel metallic plates separated by 100nm is comparable to one atmospheric pressure (100 kPa), surprisingly large considering the size and proximity of some of the new classes of micro-electromechanical systems [83]. Because the force is usually attractive, the Casimir force has been considered recently as a likely source of “stiction”, leading to the failure of many nano-devices operating in this regime. Deviations from the standard attractive, decaying force first predicted in 1948 are therefore sought after by both physicists and engineers. In fact, many researchers have become interested in exploring the shape and material dependence of the Casimir force with the goal of achieving repulsive forces or other exotic effects that may help to combat stiction, as well as reveal new insights into the Casimir effect.

There have been many experimental validations of the Casimir effect, albeit in parallel-plate geometries or simple-approximations thereof [169,268,352,358]. Theoretical tools were so under-developed just a few years ago that had experimentalists had the experimental tools needed to perform experiments on more complicated geometries back then, they would have had to go alone and blindly. Although the properties of the force were very well understood in simple parallel-plate geometries, based on the methods of Ref. [129, 302, 304], not until recently were there methods that allowed calculations of forces in arbitrary geometries ⁵. In fact, the most sophisticated theoretical calculation prior to 2007 was a calculation of

⁵Many unusual effects were actually discovered prior to advent of sophisticated numerical methods capable of handling arbitrary geometries. However, they all involved either exotic materials [54, 162, 189, 217, 219, 229, 259, 485, 491, 557, 594, 595] or fluids [74, 129].

the force between a plate and a cylinder [141]. On the one hand, this rather long hiatus between the work of Casimir and Ref. [141] mirrors some of the historical developments in the field of classical electromagnetism, which took painstakingly small steps in order to develop robust and efficient computational tools for studying complicated structures [ref]. On the other hand, the calculation of Casimir forces shares many similarities with standard problems in computational classical electromagnetism, as discussed in chapter 5, which begs the question: how can one bring all of the wisdom, work and tools of classical electromagnetism to bear on the problem of computing Casimir forces?

In 2007, inspired by the great theoretical and experimental leaps achieved in the field of classical electromagnetism, we set out to develop numerical methods for computing Casimir forces in arbitrary geometries and materials that could directly exploit standard techniques in classical electromagnetism. The main focus of this thesis is the development of two different numerical tools based on the finite-difference frequency- and time-domain methods that require very little if any modification of standard tools developed for computational electromagnetism and which allow calculations of Casimir forces in a very wide range of geometries (see description below). Furthermore, we demonstrate the proficiency of these tools by applying them to complicated geometries exhibiting unusual Casimir effects. This presentation is organized as follows:

Chapter 4 presents an overview of Casimir forces, emphasizing some of the theoretical and experimental developments of the last few years ⁶.

Chapter 5 explores several ways in which well-established, efficient techniques from standard computational electromagnetism can be brought to bear on the problem of computing Casimir forces, in order to predict forces for arbitrary geometries and materials with arbitrary accuracy with no uncontrolled approximations [443]. Starting from the simplest, most direct approaches, we show that practical considerations naturally lead towards a particular method involving the integral of the Minkowski stress tensor by repeated evaluation of the imaginary-frequency Green's function—a method previously developed for purely analytical calculations [129, 303, 304]. Because evaluation of the Green's function is such a standard problem in classical computational electromagnetism, this approach makes it possible to

⁶We apologize for any citation gaps and admit that the discussion of previous works is in no way a review of the literature. For a more holistic view of developments in the Casimir field starting from the early works of Casimir and Polder, the reader is encouraged to read Ref. [46, 68, 74, 169, 268, 352, 355, 358, 390] and references therein.

exploit many developments in fast solvers, based on finite-element [91, 237, 548, 596], spectral [50, 91], or boundary-element methods [90, 91, 191, 237, 430, 548], although our work centers only around the simplest of these methods: the finite-difference frequency-domain [93]. Our approach is tested for geometries with known solutions.

Chapter 6 describes two different methods for computing Casimir forces, both of which abandon the frequency domain altogether in favor of the time domain. As argued, there are many advantages to the time-domain perspective, most importantly, the fact that it makes it possible to obtain the entire frequency spectrum by iterating Maxwell's equations (few time steps) in response to a single dipole source. This time-domain approach is based on a correspondence between the calculation of Casimir forces for vacuum-separated objects and a similar electromagnetic-force calculation in which the objects are instead separated by a conducting fluid, also described in the chapter. The requirement that the geometry be mapped in this fashion turns out to be a practical consideration for any calculation based on the time domain (real-frequencies). In fact, it is the theoretical *equivalent* of a crucial and well-known technique for accurate numerical evaluation of Casimir forces, in which the force integrand is deformed via contour integration, and commonly evaluated over the imaginary-frequency axis [304, 443]. This mathematical equivalence between complex contour mappings and physical dielectric deformations reveals new opportunities for the experimental and theoretical study of Casimir interactions. On the theoretical side, this correspondence makes it easier to understand Casimir systems from the perspective of conventional classical electromagnetism, based on real-frequency responses, in contrast to the standard point of view based on Wick rotations (imaginary frequencies). Furthermore, it lead to a finite-difference time-domain numerical method for calculation of Casimir forces in arbitrary geometries and materials [344, 449], discussed at the end of the chapter.

Chapter 7 demonstrates the flexibility and capabilities of the finite-difference methods of chapter 5 and 6, by presenting results on a number of complicated geometries consisting of vacuum-separated metallic objects. Our methods are first used to obtain the first predictions of Casimir forces in a multi-body geometry. In particular, we study a structure consisting of two cylinders adjacent to two metallic walls, exhibiting an interesting non-monotonic dependence arising purely from the non-additivity of the Casimir force. The effect is explored for cylinders of different cross-sectional shape (squares, cylinders), and also for metallic and magnetic plates. Finally, we demonstrate how one may achieve one-dimensional Casimir

stability in a glide-symmetric geometry consisting of interlocked brackets in which repulsion is obtained via a competition of attractive surfaces that can be explained using simple pairwise approximations. While this type of repulsion is technically considered “cheating” due to its origins on standard attractive interactions, it does nevertheless suggest an interesting possibility for micro-electromechanical systems.

The final chapter of the thesis, chapter 8, examines a number of ways in which one can exploit fluids to obtain novel and interesting Casimir effects: these include stable suspension, aggregation of compact objects, and dramatic Casimir temperature effects.

The appendix of the thesis describes work on two selected topics in the field of classical electromagnetism that reveal some of the novel electromagnetic effects that can arise in complicated environments. Appendix A studies the influence of fabrication-induced disorder on microcavity modes formed by the presence of defects in complete-gap structures, such as two- and three-dimensional photonic crystals (periodic arrays of materials with period on the scale of the electromagnetic wavelength). The existence of band gaps in photonic crystals, contiguous range of electromagnetic frequencies for which propagation in the crystal is forbidden, is a direct consequence of the wave nature of light: the coherent (scatterless) interference of light propagating in periodic structures commensurate with its wavelength⁷. The presence of a defect on a photonic crystal allows light to be trapped inside the defect (under certain conditions), due to its inability to escape (propagate) through the surrounding medium. Because crystal disorder turns out to have a very small effect on the band gaps of the periodic structure, modes formed by defects in these structures are immune to scattering induced by the disorder [241], unlike the modes of conventional microcavities based on index-guiding principles [518]. We demonstrate that disorder merely reduces the lifetimes and shapes of complete-gap localized modes by performing large-scale finite-difference time-domain simulations on two particular photonic-crystal microcavities [444]. Appendix B presents work on a new numerical approach for studying a class of photonic structures known as photonic quasicrystals: aperiodic structures with long-range order that offer a number of unique properties of interest, compared to ordinary periodic structures [447]. Previous methods for determining photonic quasicrystal spectra relied on the use of large supercells to compute the eigenfrequencies and/or local density of states. We present a method by which the energy spectrum and the eigenstates of a photonic quasicrystal can be

⁷A similar effect is observed in crystal solids for propagation of electrons.

obtained by solving Maxwell's equations in higher dimensions, to which a generalization of Bloch's theorem applies. In addition, we demonstrate how one can compute band structures with defect states in the same higher-dimensional superspace. As a proof of concept, these general ideas are demonstrated for the simple case of one-dimensional quasicrystals, which can also be solved by simple transfer-matrix techniques.

Chapter 2

Classical Nonlinear Nanophotonics

“I’m astounded by people who want to ‘know’ the universe when it’s hard enough to find your way around Chinatown.” *Woody Allen*

2.1 Overview

In this chapter, we explore some of the consequences of geometric resonances on classical nonlinear electromagnetic interactions. In particular, we present a general description of nonlinear frequency conversion in resonant microcavities (structures with geometric or material features at the scale of the electromagnetic wavelength) which exploits the generality of a theoretical approach to resonant interactions known as temporal coupled-mode theory (CMT). This approach provides both quantitative and qualitative predictions for cavity-based frequency conversion, regardless of geometric structure (described by the symmetries, frequencies, and decay rates of the cavity).

The chapter is broken down into the following sections: Section 2.1.1 presents a general introduction to the field of nonlinear nanophotonics that emphasizes recent work involving photonic-crystal microcavity resonators. Section 2.1.2 provides a brief introduction to linear CMT—the formalism is illustrated via a simple example system known as a linear electromagnetic filter, in which light incident from a waveguide that is coupled to a microcavity is only transmitted for light frequencies lying in a specified frequency band.

Section 2.2 extends the linear CMT of Sec. 2.1.2 to include the presence of nonlinear processes occurring within the microcavity. In particular, we develop nonlinear coupled-mode equations for a process known as harmonic generation, in which input light at some

frequency ω is converted to light at some multiple of this frequency $k\omega$ due to the nonlinearities. However, because the construction of nonlinear coupled-mode theories hinges on very general symmetry arguments, independent of the details of the underlying wave equation, the formalism is easily applicable to other forms of nonlinear frequency conversion, as demonstrated in later sections. Section 2.2.2 describes how to compute all of the required CMT parameters, e.g. coupling coefficients, which requires performing a simple perturbation-theory calculation starting from the linear Maxwell’s equations. Once all of the basic components of the CMT are established, it becomes feasible to analyze harmonic generation in general geometries. Sections 2.2.3–2.2.8 are devoted to analyzing second- ($\chi^{(2)}$) and third- ($\chi^{(3)}$) harmonic generation in doubly resonant cavities. $\chi^{(2)}$ and $\chi^{(3)}$ harmonic generation involve the conversion of input light of frequency ω to output light of frequency 2ω and 3ω , respectively. Whereas it is well known that 100% harmonic conversion is possible, at least in $\chi^{(2)}$ media, we explore and identify the conditions under which this can be achieved in the doubly-resonant setting. Most importantly, we demonstrate the existence of a critical input power at which harmonic generation is maximized, in contrast to previous work that focused largely on the low-power (undepleted) limit, for which the efficiency increases monotonically with input power. As a check on the accuracy of our CMT, we examine third-harmonic generation (THG) via “exact” numerical experiments using an FDTD simulation in a simple one-dimensional geometry. The results of these numerical experiments allow us to quantify the accuracy of the CMT and identify its limits of applicability. In addition, they also reveal a number of important features crucial to THG, related to frequency-shifting effects due to cross- and self-phase modulation, the focus of Sec. 2.2.5–2.2.7.

Sections 2.2.5–2.2.7 are devoted to an investigation of the dynamics of THG. This is achieved by performing a simple stability analysis of the desired, high-efficiency solutions, which demonstrates that resonant harmonic generation can yield many rich dynamical behaviors absent in conventional (non-resonant) schemes, such as the existence of limit cycles and multistability. It is shown that achieving efficient THG requires additional considerations, such as shifting of the original cavity frequencies and proper priming of the steady-state solutions. The final sections on harmonic generation are devoted to a study of the effects of external losses (e.g. absorption or radiation loss) and a brief note on some of the outstanding questions for future work.

We conclude the chapter by discussing two additional types of nonlinear frequency conversion processes. Section 2.3 focuses on a process known as difference-frequency generation, in which incident light composed of two frequencies ω_1 and ω_2 is converted to light at a third frequency $\omega_T = \omega_2 - \omega_1$. A promising application of resonant DFG, explored in Sec. 2.3.1, is the possibility of efficiently generating an output signal $\omega_T \sim 1$ THz. In particular, we present the first predictions of all-optical THz generation in triply-resonant cavities, a current topic of interest due to the many important applications based on THz light, and the unfortunate scarcity of THz sources. Finally, Sec. 2.4 briefly surveys nonlinear three-wave mixing in triply-resonant cavities. Three-wave mixing is a process by which incident light composed of two frequencies ω_0 and $\omega_0 + \Delta$ is converted to light at a third frequency $\omega_0 - \Delta$.

2.1.1 Nonlinear Nanophotonics

When the intensity of light interacting with a polarizable material exceeds a certain (material-dependent) threshold, its dielectric response can no longer be described by a linear permittivity, and the corresponding Maxwell's equations become nonlinear [51, 469]. Although these nonlinearities are often relatively weak, and therefore negligible, they can give rise to a number of fascinating and important electromagnetic effects, with many promising applications [145, 308, 345, 366, 469, 477]. Not surprisingly, there has been intense theoretical and experimental work on the problem of strengthening nonlinear interactions. One particular approach involves either engineering or discovering novel materials with stronger nonlinear susceptibilities. A different approach, the focus of this chapter, involves tailoring the geometric features of nanophotonic media to obtain novel effects [241].

Since the early days of nonlinear optics, optical resonators have been seen as an attractive way to enhance nonlinear optical phenomena, such as frequency conversion processes [51] (in optical parametric oscillators) or optical bistability properties [175]. Traditionally, these nonlinear optical resonators have consisted of a nonlinear material located between two partially transmitting mirrors. Although interesting in their own right, the application of these nonlinear Fabry-Perot interferometers for designing all-optical logical devices is rather limited, as they cannot fulfill the requirements in size, switching time and operating power of practical integrated optical systems [57].

On a separate front, the rapid development of fabrication techniques at micro and nanometric scale has enabled the successful demonstration of optical micro- and nano-cavities

where light is strongly confined in a very small volume. This corresponds to the case where the ratio between the quality factor Q and the modal volume V is extremely large, such that $Q/V \gg 1/\lambda^3$). Here, it is important to emphasize that the modal volume is a measure of the spatial extent of the mode in the material of interest, but its precise definition will depend on the physical phenomenon of interest [57].

During the last decade, different types of optical cavities characterized by ultra-small modal volumes and extremely high quality factors have been successfully demonstrated (such as photonic-crystal (PhC) cavities or microtoroid and microsphere resonators, just to cite some examples [537]). Among them, due to their unique confinement mechanism [241], PhCs cavities have proved to be a versatile route to develop novel optical integrated devices. PhC optical cavities are usually created by introducing a small defect in an otherwise perfectly periodically modulated refractive index profile in either one, two or three dimensions. The linear properties of PhCs including such defects have been extensively studied both theoretically and experimentally [6, 148, 158, 246, 381, 395, 464, 508, 511, 549, 587]. Recently, it has been experimentally demonstrated that it is possible to design PhC microcavities with quality factors of $Q \sim 10^6$ and modal volumes of the order of a cubic wavelength [382].

As a consequence of these recent advances in nanophotonic fabrication, many of the nonlinear phenomena previously analyzed in conventional large-etalon structures are being revisited within the context of PhC cavities [5, 24, 36, 55, 81, 101, 102, 123, 150, 203, 337, 361, 387, 436, 437, 445, 500, 501, 504, 506, 507, 523, 526, 533, 575, 576]. Although it is true that the physical mechanisms producing nonlinear phenomena (such as optical bistability or nonlinear frequency conversion) in PhC cavities are similar to those observed in their conventional counterparts, it has been demonstrated that PhC microresonators enhance the performance of traditional nonlinear devices by several orders of magnitude. In addition, due to their design versatility, PhC cavities can be used as the basis of completely new configurations performing all-optical logic functions, such as all-optical transistor action [575]. Moreover, as detailed in chapter 3, PhC resonators offer new fundamental ways of tailoring optical nonlinearities by using the so-called Purcell effect [36, 569]. Finally, in this context, note that not only PhC resonators can lead to a strong enhancement of nonlinear phenomena, but also nonlinear effects can be enhanced by using slow-light properties of PhCs, via the corresponding band edges [5, 411] or by means of coupled-cavity waveguides [405, 504, 567, 580].

It is important to mention that, together with the aforementioned rapid development of

the experimental techniques, during the last decade there has also been an important growth in large-scale computing technologies. Thus, the combination of purely numerical methods, such as the nonlinear finite-difference-time domain method (FDTD) [523] (which simulates Maxwells equations with no approximation apart from discretization), in conjunction to analytical approaches such as coupled-mode theory [203] or perturbation theory [501], allow a complete characterization of the electromagnetic response of the nonlinear PhC cavities under study. In particular, all the theoretical calculations presented in this chapter have been obtained by using one of the tools just mentioned. However, since this work is mainly focused in explaining the physical mechanisms responsible for the observed results, we refer to the reader to more specialized references for details on the numerical calculations [5, 55, 150, 203, 523].

2.1.2 Temporal Coupled-Mode Theory

As mentioned in the previous section, coupled-mode theories are often constructed to describe weakly-coupled systems in terms of a set of idealized components (e.g. isolated waveguides and cavities) that are coupled in some fashion [203, 214, 241, 245, 495, 520, 574, 576, 578], similar to simple time-dependent perturbation theory (Breit-Wigner scattering) in quantum mechanics [286]. The idea is that, to a first approximation, the interaction of localized cavity modes with their environment is weak, making it possible to represent the solutions of the entire coupled system as a perturbative expansion in terms of the localized (component) modes. As discussed in Ref. [241], the construction of the linear theory relies only on a few, but general assumptions: weak coupling, linearity, time-invariant material/geometric properties, conservation of energy, and time-reversal invariance. In what follows, I will briefly describe some of the basic components of temporal coupled-mode theory. For a more in-depth presentation of the theory, the reader is encouraged to read chapter 10 in Ref. [241].

Let us first construct the coupled-mode equations describing the interaction of light in a multi-mode cavity coupled to a single set of input/output ports, from which light can couple in (s_+) and out (s_-) of the cavity. A schematic illustration of the system is shown in Fig. 6-8. Specifically, we follow the formalism described in Ref. [202], adapted to handle non-linearly coupled modes with frequencies ω_k , which is parameterized as follows.

We let a_k denote the time-dependent complex amplitude of the k th mode, conveniently

normalized so that $|a_k|^2$ is the electromagnetic energy stored in this mode. We let s_{\pm} denote the time-dependent amplitude of the incoming (+) or outgoing (-) wave, normalized so that $|s_{\pm}|^2$ is the power. (More precisely, $s_{\pm}(t)$ is normalized so that its Fourier transform $|\tilde{s}_{\pm}(\omega)|^2$ is the power at ω .)¹ By itself, a linear cavity mode decaying with a lifetime τ_k would be described by $da_k/dt = (i\omega_k - 1/\tau_k)a_k$, which corresponds to a mode a_k with solution:

$$a_k(t) = a_k(0)e^{-i\omega_k t - t/\tau_k}. \quad (2.1)$$

Technically, such a decaying mode is not a true eigenmode, but is rather a “leaky mode” [499], corresponding to a “quasi-bound state” in the Breit-Wigner scattering theory [286]. The decay rate $1/\tau_k$ can be decomposed into $1/\tau_k = 1/\tau_{e,k} + 1/\tau_{s,k}$ where $1/\tau_{e,k}$ is the “external” loss rate (absorption etc.) and $1/\tau_{s,k}$ is the decay rate into s_- . When the weak coupling ($\omega_k \tau_k \gg 1$) to s_{\pm} is included, energy conservation and similar fundamental constraints lead to equations of the form [203]:

$$\frac{da_k}{dt} = \left(i\omega_k - \frac{1}{\tau_k}\right) a_k + \sqrt{\frac{2}{\tau_{s,k}}} s_+ \quad (2.2)$$

$$s_- = -s_+ + \sqrt{\frac{2}{\tau_{s,k}}} a_k \quad (2.3)$$

This model can be generalized to incorporate multiple input/output ports, direct coupling between the ports, and so on [520], although the main situation(s) of interest involve either a single or two ports [483]. The only unknown parameters in this model are then the frequencies ω_k and the decay rates $1/\tau_k$, which can be determined by any numerical method to solve for the cavity modes (e.g. FDTD, below). (Note that instead of τ_k , one commonly uses the quality factor $Q_k = \omega_k \tau_k/2$.)

The fact that a_k is not a true eigenmode of the entire system is evident upon inspection

¹In 1d, the units of $|a_k|^2$ and $|s_{\pm}|^2$ are those of energy and power per unit area, respectively. More generally, in d dimensions, the units of $|a_k|^2$ and $|s_{\pm}|^2$ are those of energy and power per length^{3-d}.

of its Fourier spectrum \tilde{a}_k , which has a Lorentzian lineshape ²:

$$\tilde{a}_k(\omega) \sim \frac{a_k(0)/\tau_k^2}{(\omega - \omega_k)^2 + 1/\tau_k^2}, \quad (2.4)$$

In particular, a_k is a superposition of harmonic modes $\sim e^{-i\omega t}$ with frequencies in the vicinity of ω_k , and with bandwidth or spectral support inversely proportional to its lifetime, i.e., $\Delta\omega = 1/\tau_k$. The size of the domain over which frequencies $\omega \neq \omega_k$ in \tilde{a}_k contribute to a_k is a measure of the accuracy of the coupled-mode approximation: For small τ_k (corresponding to strong coupling to the waveguide), the number of modes contributing to a_k is large. As τ_k grows (corresponding to a more weakly-coupled cavity), the contribution of modes $\omega \neq \omega_k$ to a_k decreases, and it becomes more sensible to view a_k as a mode of the system. [Note that in the limit as $\tau_k \rightarrow \infty$, corresponding to an infinitely lived cavity mode with no coupling to the waveguide, a_k becomes a true eigenmode with Fourier transform given by $\tilde{a}_k \sim \delta(\omega - \omega_k)$.] In practice, temporal-coupled-mode theory (or this quasi-bound state approximation) becomes nearly exact for $Q_k > 30$.

Equation (2.2) allows us to predict the transmission/reflection spectrum of any weakly-coupled waveguide-cavity system. We now consider a simple example system for illustration purposes, known as a narrow-band filter. In particular, consider the system shown in Fig. 2-1: a lossy cavity of frequency ω_0 coupled to two waveguides, denoted by input ports s_{1+} and s_{2+} , with decay rates $1/\tau_1$ and $1/\tau_2$, respectively, and absorption rate $1/\tau'$. Let us assume that there is only incident light through the left waveguide ($s_{2+} = 0$). Since frequency is conserved in a linear system, if the input oscillates at a fixed frequency ω , then the field everywhere must oscillate as $e^{-i\omega t}$, and $da/dt = -i\omega a$. Plugging this ansatz, and $s_{2+} = 0$,

²An alternative, and more mathematical way to understand the leaky-mode approximation involves an analysis of the complex-frequency Green's function $G(\omega)$ of a weakly-coupled system of resonances, but this is outside of the scope of this thesis. The basic idea goes as follows: The eigenmodes of the individual components of the coupled mode theory, i.e. the cavities, are manifested as poles in G located along the real- ω axis. The act of coupling the cavity to an output port is equivalent to that of introducing a small absorption loss into the system, leading to a complex pole in G at $\omega + i\delta$, where δ is a small imaginary term that is proportional to the "coupling" to the port. The stronger the coupling, the farther the pole from the real-frequency axis, and the less it behaves like a true eigenmode.

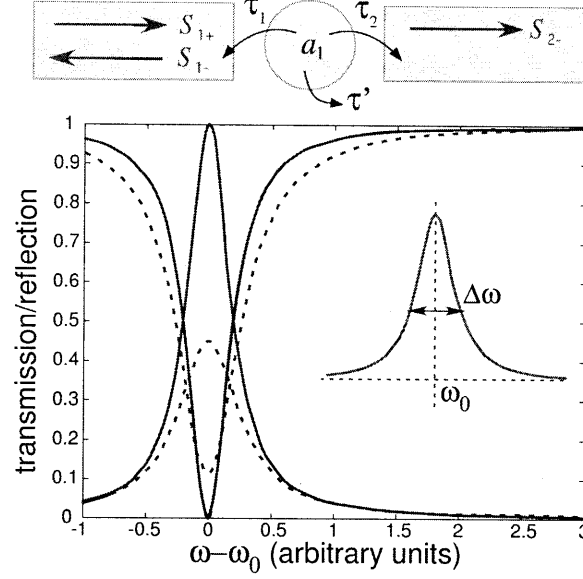


Figure 2-1: (Top:) Schematic of narrow-band filter system, consisting of a cavity with a mode a of frequency ω_0 , coupled to two input/output ports $s_{1\pm}$ and $s_{2\pm}$, with decay rates $1/\tau_1$ and $1/\tau_2$, respectively. The cavity is filled with an absorptive medium and has an absorption rate $1/\tau'$. (Bottom:) transmission $T = |s_{2-}|^2/|s_{1+}|^2$ (red line) and reflection $R = |s_{1-}|^2/|s_{1+}|^2$ (blue line) coefficients of input light from the left (s_{1+}) as a function of ω , for $\tau' = \infty$ (solid line) and $\tau' \sim \tau_1$ (dashed line). The inset (green line) shows a generic Lorentzian line spectrum.

into the coupled-mode equations, we obtain:

$$-i\omega a = -i\omega_0 a - \left(\frac{1}{\tau'} + \frac{1}{\tau_1} + \frac{1}{\tau_2} \right) a + \sqrt{\frac{2}{\tau_1}} s_{1+} \quad (2.5)$$

$$s_{1-} = s_{1+} + \sqrt{\frac{2}{\tau_1}} a \quad (2.6)$$

$$s_{2-} = \sqrt{\frac{2}{\tau_2}} a, \quad (2.7)$$

which yields the following transmission $T(\omega)$ and reflection $R(\omega)$ spectra:

$$T(\omega) = \frac{|s_{2-}|^2}{|s_{1+}|^2} = \frac{4/\tau_1\tau_2}{(\omega - \omega_0)^2 + (1/\tau_1 + 1/\tau_2 + 1/\tau')^2} \quad (2.8)$$

$$R(\omega) = \frac{|s_{1-}|^2}{|s_{1+}|^2} = \frac{(\omega - \omega_0)^2 + (1/\tau_1 - 1/\tau_2 - 1/\tau')^2}{(\omega - \omega_0)^2 + (1/\tau_1 + 1/\tau_2 + 1/\tau')^2}. \quad (2.9)$$

The (Lorentzian) transmission and reflection spectra for $\tau_1 = \tau_2$ are plotted in Fig. 2-1 for various values of τ' . In the absence of loss $\tau' = \infty$, it is easy to verify that $R + T = 1$ at all frequencies (energy is conserved) and that light is perfectly transmitted at $\omega = \omega_0$ if

and only if $\tau_1 = \tau_2$; equivalently, only when the rates at which the cavity decays into the two ports are equal, can there exist 100% transmission. In addition, if this “rate-matching” condition is satisfied, the reflection of the system is identically zero: what is occurring here is that the light reflected directly from the cavity is destructively canceled by the decay rate of the cavity mode, to yield zero reflection ³ The effect of the losses is to simply reduce the transmission and reflection of the system, as shown by the figure. Intuitively, the presence of loss in this system creates another port (decay mechanism) through which the light can escape. In general, frequency conversion or filtering in systems composed of multiple ports is undesirable because the strong “rate-matching” (interference) conditions that lead to 100% transmission/reflection is weakened by the presence of additional ports (energy-escape channels), leading to a degradation of the possible maximum efficiency of the device.

2.2 Second and Third Harmonic Generation

In this section, we extend the linear CMT developed in the previous section, to allow nonlinear interactions between cavity modes. The focus of this section is a particular type of nonlinear process known as harmonic generation, which we shall examine in two different cases, corresponding to $\chi^{(2)}$ and $\chi^{(3)}$ nonlinear interactions [51]. We derive nonlinear coupled-mode equations for both of these processes, investigate the possible steady-state solutions and corresponding stability (dynamics), and point out the most important considerations that go into achieving high-efficiency harmonic generation in both cases.

2.2.1 Nonlinear Coupled-Mode Theory

Nonlinearities modify the linear picture presented in Sec. 2.1.2 by introducing two new amplitude-dependent effects: a shift in the frequency (and decay rate) of the cavity, and a coupling of one cavity mode to another. We neglect nonlinear effects on the input/output ports, under the assumption that intense fields are only present in the cavity (due to spatial and temporal confinement). We will also make two standard assumptions of nonlinear systems. First, that the nonlinearities are weak, in the sense that we can neglect terms of order $(\chi^{(\ell)})^2$ or higher; this is true in practice because nonlinear index shifts are always

³As discussed below, a similar “rate-matching” criteria leads to 100% nonlinear frequency conversion efficiency at a critical power.

under 1% lost material breakdown occur. Second, we make the rotating wave approximation: since the coupling is weak, we only include terms for a_k that have frequency near ω_k . In particular, we suppose that $\omega_k \approx k\omega_1$, so that ω_k is the k th harmonic. Then, for a $\chi^{(2)}$ nonlinearity with two modes ω_1 and its second harmonic ω_2 , the coupled-mode equations must take the form:

$$\frac{da_1}{dt} = \left(i\omega_1 - \frac{1}{\tau_1} \right) a_1 - i\omega_1\beta_1 a_1^* a_2 + \sqrt{\frac{2}{\tau_{s,1}}} s_+ \quad (2.10)$$

$$\frac{da_2}{dt} = \left(i\omega_2 - \frac{1}{\tau_2} \right) a_2 - i\omega_2\beta_2 a_1^2 + \sqrt{\frac{2}{\tau_{s,2}}} s_+ \quad (2.11)$$

Similarly, for a $\chi^{(3)}$ nonlinearity with two modes ω_1 and its third harmonic ω_3 , the coupled-mode equations must take the form:

$$\frac{da_1}{dt} = \left(i\omega_1(1 - \alpha_{11}|a_1|^2 - \alpha_{13}|a_3|^2) - \frac{1}{\tau_1} \right) a_1 - i\omega_1\beta_1(a_1^*)^2 a_3 + \sqrt{\frac{2}{\tau_{s,1}}} s_+ \quad (2.12)$$

$$\frac{da_3}{dt} = \left(i\omega_3(1 - \alpha_{33}|a_3|^2 - \alpha_{31}|a_1|^2) - \frac{1}{\tau_3} \right) a_3 - i\omega_3\beta_3 a_1^3 + \sqrt{\frac{2}{\tau_{s,3}}} s_+ \quad (2.13)$$

In Eqs. (2.12–2.13), one sees two kinds of terms. The first are *frequency-shifting* terms, with coefficients α_{ij} , dependent on one of the field amplitudes. For $\chi^{(3)}$, this effect is known as self-phase and cross-phase modulation, which is absent for $\chi^{(2)}$ (under the first-order rotating-wave approximation). The second kind of term *transfers energy* between the modes, with coupling coefficients β_i , corresponding to four-wave mixing for $\chi^{(3)}$. Furthermore, we can constrain the coupling terms β_i by energy conservation: $\frac{d}{dt}(|a_1|^2 + |a_2|^2) = 0$. For $\chi^{(2)}$, the constraint that follows is: $\omega_1\beta_1 = \omega_2\beta_2^*$; for $\chi^{(3)}$, the constraint is $\omega_1\beta_1 = \omega_3\beta_3^*$.

The general process for construction of these coupled-mode equations is as follows. The underlying nonlinearity must depend on the physical, real part of the fields, corresponding to $(a_k + a_k^*)/2$. It then follows that the $\chi^{(\ell)}$ term will have ℓ powers of this real part, giving various product terms like $a_1^* a_2$ (for $\chi^{(2)}$) and $a_1^* a_1 a_1$ (for $\chi^{(3)}$). Most of these terms, however, can be eliminated by the rotating-wave approximation. In particular, we assume that each a_k term is proportional to $e^{ki\omega}$ multiplied by a slowly varying envelope, and we discard any product term whose total frequency differs from $k\omega$ for the da_k/dt equation. Thus, a term like $a_1^* a_3 a_3$ would be proportional to $e^{5i\omega}$, and would only appear in a da_5/dt equation. (We focus on the simpler case of doubly resonant cavities here.)

At this point, the equations are already useful in order to reason about what types of qualitative behaviors are possible in general. In fact, they are not even specific to electromagnetism and would also apply to other situations such as acoustic resonators. However, in order to make quantitative predictions, one needs to know the nonlinear coefficients α_{ij} and β_i (as well as the linear frequencies and decay rates). The evaluation of these coefficients requires a more detailed analysis of Maxwell's equations as described below.

2.2.2 Coupling Coefficients via Perturbation Theory

When a dielectric structure is perturbed by a small $\delta\varepsilon$, a well-known result of perturbation theory states that the corresponding change $\delta\omega$ in an eigenfrequency ω is, to first order [239]:

$$\frac{\delta\omega}{\omega} = -\frac{1}{2} \frac{\int d^3\mathbf{x} \delta\varepsilon |\mathbf{E}|^2}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}|^2} = -\frac{1}{2} \frac{\int d^3\mathbf{x} \mathbf{E}^* \cdot \delta\mathbf{P}}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}|^2} \quad (2.14)$$

where \mathbf{E} is the unperturbed electric field and $\delta\mathbf{P} = \delta\varepsilon\mathbf{E}$ is the change in polarization density due to $\delta\varepsilon$. In fact, Eq. (2.14) is general enough to be used with any $\delta\mathbf{P}$, including the polarization that arises from a nonlinear susceptibility. In particular, we can use it to obtain the coupling coefficients of the CMT.

To do so, we first compute the nonlinear first-order frequency perturbation due to the total field \mathbf{E} present from all of the modes. Once the frequency perturbations $\delta\omega_k$ are known, we can re-introduce these into the coupled-mode theory by simply setting $\omega_k \rightarrow \omega_k + \delta\omega_k$ in Eq. (2.2). By comparison with Eqs. (2.10–2.13), we can then identify the α and β coefficients.

We consider first a $\chi^{(2)}$ nonlinearity, with \mathbf{P} given by $P_i = \sum_{ijk} \varepsilon_0 \chi_{ijk}^{(2)} E_j E_k$, in a cavity with two modes \mathbf{E}_1 and \mathbf{E}_2 . As before, we require that the modes oscillate with frequency ω_1 and $\omega_2 \approx 2\omega_1$, respectively. Taking $\mathbf{E} = \text{Re}[\mathbf{E}_1 e^{i\omega_1 t} + \mathbf{E}_2 e^{i\omega_2 t}]$ and using the rotating-wave approximation, we can separate the contribution of $\delta\mathbf{P}$ to each $\delta\omega_k$, to obtain the following

frequency perturbations:

$$\frac{\delta\omega_1}{\omega_1} = -\frac{1}{4} \frac{\int d^3\mathbf{x} \sum_{ijk} \varepsilon_0 \chi_{ijk}^{(2)} \left[E_{1i}^* (E_{2j} E_{1k}^* + E_{1j}^* E_{2k}) \right]}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2} \quad (2.15)$$

$$\frac{\delta\omega_2}{\omega_2} = -\frac{1}{4} \frac{\int d^3\mathbf{x} \sum_{ijk} \varepsilon_0 \chi_{ijk}^{(2)} E_{2i}^* E_{1j} E_{1k}}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}_2|^2} \quad (2.16)$$

Similarly, for a centro-symmetric $\chi^{(3)}$ medium, \mathbf{P} is given by $\mathbf{P} = \varepsilon_0 \chi^{(3)} |\mathbf{E}|^2 \mathbf{E}$, with $\mathbf{E} = \text{Re}[\mathbf{E}_1 e^{i\omega_1 t} + \mathbf{E}_3 e^{i\omega_3 t}]$. We obtain the following frequency perturbations:

$$\begin{aligned} \frac{\delta\omega_1}{\omega_1} = & -\frac{1}{8} \left[\frac{\int d^3\mathbf{x} \varepsilon_0 \chi^{(3)} \left(|\mathbf{E}_1 \cdot \mathbf{E}_1|^2 + 2 |\mathbf{E}_1 \cdot \mathbf{E}_1^*|^2 + 2(\mathbf{E}_1 \cdot \mathbf{E}_1^*)(\mathbf{E}_3 \cdot \mathbf{E}_3^*) \right)}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2} \right. \\ & \left. + \frac{2 |\mathbf{E}_1 \cdot \mathbf{E}_3|^2 + 2 |\mathbf{E}_1 \cdot \mathbf{E}_3^*|^2 + 3(\mathbf{E}_1^* \cdot \mathbf{E}_1^*)(\mathbf{E}_1 \cdot \mathbf{E}_3)}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2} \right] \quad (2.17) \end{aligned}$$

$$\begin{aligned} \frac{\delta\omega_3}{\omega_3} = & -\frac{1}{8} \left[\frac{\int d^3\mathbf{x} \varepsilon_0 \chi^{(3)} \left(|\mathbf{E}_3 \cdot \mathbf{E}_3|^2 + 2 |\mathbf{E}_3 \cdot \mathbf{E}_3^*|^2 + 2(\mathbf{E}_1 \cdot \mathbf{E}_1^*)(\mathbf{E}_3 \cdot \mathbf{E}_3^*) \right)}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}_3|^2} \right. \\ & \left. + \frac{2 |\mathbf{E}_1 \cdot \mathbf{E}_3|^2 + 2 |\mathbf{E}_1 \cdot \mathbf{E}_3^*|^2 + (\mathbf{E}_1 \cdot \mathbf{E}_1)(\mathbf{E}_1 \cdot \mathbf{E}_3^*)}{\int d^3\mathbf{x} \varepsilon |\mathbf{E}_3|^2} \right] \quad (2.18) \end{aligned}$$

There is a subtlety in the application of perturbation theory to decaying modes, such as those of a cavity coupled to output ports. In this case, the modes are not truly eigenmodes, but are rather “leaky modes” [499], and are not normalizable. Perturbative methods in this context are discussed in more detail by Ref. [499, 520], but for a tightly confined cavity mode it is sufficient to simply ignore the small radiating field far away from the cavity. The field in the cavity is very nearly that of a true eigenmode of an isolated cavity.

As stated above, we can arrive at the coupling coefficients by setting $\omega_k \rightarrow \omega_k + \delta\omega_k$ in Eq. (2.2). However, the frequency perturbations $\delta\omega_k$ are time-independent quantities, and we need to connect them to the time-dependent a_k amplitudes. Therefore, to re-introduce the time dependence, one can use the slowly varying envelope approximation: a slowly varying, time-dependent amplitude $a_k(t)$ is introduced into the unperturbed fields $\mathbf{E}_k \rightarrow \mathbf{E}_k a_k(t)$. The eigenmode must be normalized so that $|a_k|^2$ is the energy, as assumed for the coupled-mode theory. Thus, we divide each \mathbf{E}_k by $\sqrt{\frac{1}{2} \int \varepsilon |\mathbf{E}_k|^2}$.

First, we consider the $\chi^{(2)}$ medium. Carrying out the above substitutions in Eq. (2.2) and grouping terms proportional a_k yields Eqs. (2.10–2.11) with α_{ij} and β_i given by:

$$\alpha_{ij} = 0 \quad (2.19)$$

$$\beta_1 = \frac{1}{4} \frac{\int d^3\mathbf{x} \sum_{ijk} \varepsilon_0 \chi_{ijk}^{(2)} \left[E_{1i}^* (E_{2j} E_{1k}^* + E_{1j}^* E_{2k}) \right]}{\left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2 \right] \left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_2|^2 \right]^{1/2}} \quad (2.20)$$

$$\beta_2 = \frac{1}{4} \frac{\int d^3\mathbf{x} \sum_{ijk} \varepsilon_0 \chi_{ijk}^{(2)} E_{2i}^* E_{1j} E_{1k}}{\left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2 \right] \left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_2|^2 \right]^{1/2}} \quad (2.21)$$

A similar calculation yields the $\chi^{(3)}$ coupled-mode equations with coefficients given by:

$$\alpha_{ii} = \frac{1}{8} \frac{\int d^3\mathbf{x} \varepsilon_0 \chi^{(3)} \left(|\mathbf{E}_i \cdot \mathbf{E}_i|^2 + |\mathbf{E}_i \cdot \mathbf{E}_i^*|^2 \right)}{\left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_i|^2 \right]^2} \quad (2.22)$$

$$\alpha_{31} = \frac{1}{4} \frac{\int d^3\mathbf{x} \varepsilon \chi^{(3)} \left(|\mathbf{E}_1|^2 |\mathbf{E}_3|^2 + |\mathbf{E}_1 \cdot \mathbf{E}_3|^2 + |\mathbf{E}_1 \cdot \mathbf{E}_3^*|^2 \right)}{\left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2 \right] \left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_3|^2 \right]} \quad (2.23)$$

$$\alpha_{13} = \alpha_{31}$$

$$\beta_1 = \frac{3}{8} \frac{\int d^3\mathbf{x} \varepsilon_0 \chi^{(3)} (\mathbf{E}_1^* \cdot \mathbf{E}_1^*) (\mathbf{E}_1^* \cdot \mathbf{E}_3)}{\left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2 \right]^{3/2} \left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_3|^2 \right]^{1/2}} \quad (2.24)$$

$$\beta_3 = \frac{1}{8} \frac{\int d^3\mathbf{x} \varepsilon_0 \chi^{(3)} (\mathbf{E}_1 \cdot \mathbf{E}_1) (\mathbf{E}_1 \cdot \mathbf{E}_3^*)}{\left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_1|^2 \right]^{3/2} \left[\int d^3\mathbf{x} \varepsilon |\mathbf{E}_3|^2 \right]^{1/2}} \quad (2.25)$$

Note that Eqs. (2.19–2.25) verify the conditions $\omega_1 \beta_1 = \omega_2 \beta_2^*$ and $\omega_1 \beta_1 = \omega_3 \beta_3^*$, previously derived from conservation of energy—for $\chi^{(2)}$, this requires that one apply the symmetries of the $\chi_{ijk}^{(2)}$ tensor, which is invariant under permutations of ijk for a frequency-independent $\chi^{(2)}$ [51]. Furthermore, we can relate the coefficients α and β to an effective modal volume V , similar to Ref. [501]. In particular, the strongest possible nonlinear coupling will occur if the eigenfields are a constant in the nonlinear material and zero elsewhere. In this case, any integral over the fields will simply yield the geometric volume V of the nonlinear material. Thus, for the $\chi^{(2)}$ effect we would obtain $\beta_i \sim \chi^{(2)}/\sqrt{V\varepsilon^3}$; similarly, for the $\chi^{(3)}$ effect we would obtain $\alpha_{ij}, \beta_i \sim \chi^{(3)}/V\varepsilon^2$. This proportionality to $1/\sqrt{V}$ and $1/V$ carries over to more realistic field profiles (and in fact could be used to *define* a modal volume for these effects).

2.2.3 Complete Frequency Conversion at a Critical Power

In what follows, we investigate the steady-state behavior of second- and third-harmonic generation in the particular system of Fig. 6-8, consisting of a single input/output port coupled to a doubly resonant cavity. Using a universal framework based on the nonlinear coupled-mode equations of Sec. 2.2.1, along with the coupling coefficients of Sec. 2.2.2, we generalize previous experimental and theoretical work on this subject, which had focused primarily on $\chi^{(2)}$ nonlinearities in large Fabry-Perot etalons or on specific geometries such as micro-ring resonators, to general geometries consisting of fully vectorial electromagnetic fields. We develop several results from this generalization: whereas it is well known that 100% harmonic conversion is possible, at least in $\chi^{(2)}$ media, in this section we further explore and identify the conditions under which this can be achieved. First, we demonstrate the existence of a critical input power at which harmonic generation is maximized. A number of works have derived exact equations of motion, albeit for specific structures, and determined the existence of this critical power [163, 225, 474]. However, these works have focused only on specific structures and the relationship between critical power and the cavity parameters seems to have been unappreciated. Second, while it is well known that harmonic conversion can be achieved at arbitrarily low powers given sufficiently long cavity lifetimes, this implies a narrow bandwidth—we show that, by combining moderate lifetimes (0.1% bandwidth) with tight spatial confinement, 100% second- and third-harmonic conversion can theoretically be achieved with sub-milliwatt power levels, without using gain or amplification of any kind. Such low-power conversion could find applications such as high-frequency sources [33], ultracompact-coherent optical sources [145, 308, 366], imaging [477], and spectroscopy [345].

Nonlinear frequency conversion has been commonly realized in the context of waveguides [4, 102, 128, 329], or even for free propagation in the nonlinear materials, in which light at one frequency co-propagates with the generated light at the harmonic frequency [5, 19, 383, 405]. A phase-matching condition between the two frequencies must be satisfied in this case in order to obtain efficient conversion [33]. Moreover, as the input power is increased, the frequency conversion eventually saturates due to competition between up and down conversion. Frequency conversion in a doubly-resonant cavity, as we shall derive, has three fundamental differences from this familiar case of propagating modes. First, light in a cavity




Figure 2-2: (Top:) Schematic of general scheme for third-harmonic generation, including dynamical variables for coupled-mode equations: a single input (output) channel (with incoming (outgoing) field amplitudes s_{\pm}) is coupled to a resonant cavity with two modes at frequencies ω_1 and $\omega_3 = 3\omega_1$ (with corresponding amplitudes a_1 and a_3). The two resonant modes are nonlinearly coupled by a Kerr ($\chi^{(3)}$) nonlinearity. Reflections at the fundamental frequency (s_{1-}) may also occur. (Bottom:) An example realization [445], in one dimension, using a semi-infinite quarter-wave stack of dielectric layers with a doubled-layer defect (resonant cavity) that is coupled to incident plane waves. Dielectric material is yellow, and the steady-state electric field E_z of third-harmonic mode is shown as blue/white/red for positive/zero/negative amplitude.

can be much more intense for the same input power, because of the spatial (modal volume V) and temporal (lifetime Q) confinement. We show that this enhances second-harmonic ($\chi^{(2)}$) conversion by a factor of Q^3/V and enhances third-harmonic ($\chi^{(3)}$) conversion by a factor of Q^2/V . Second, the phase-matching condition is replaced by two conditions: First, the cavity must support two modes of the requisite frequencies—the frequencies can be designed by tuning any of a number of cavity parameters; Second: a selection rule between the modes participating in the interaction, determined by the symmetry of the cavity. As opposed to the non-resonant case, frequency conversion in doubly-resonant cavities no longer saturates—instead, it peaks (at 100%, with proper design) for a certain critical input power satisfying a resonant condition, and goes to *zero* if the power is *either* too small or too large.

Previous experimental and theoretical work on second-harmonic generation in cavities has largely focused on cavities with a single resonant mode at the pump frequency [15, 16, 38, 65, 121, 123, 172, 192, 275, 312, 345, 366, 371, 406, 476, 498]. (A singly-resonant cavity at the harmonic frequency instead of the pump frequency was considered by [118].) Such structures require much higher powers than our proposed doubly-resonant cavity, however, approaching one Watt [393] and/or requiring amplification within the cavity. Second-harmonic generation in a doubly resonant cavity, with a resonance at both the pump and harmonic frequency, have previously been analyzed [33, 125, 309, 324, 393, 402, 437, 563, 602] both in the limit where nonlinear down-conversion can be neglected, and in the depleted-pump limit,

but in a non-universal framework specific to particular geometries [163, 225, 474]. Since we wish to obtain 100% conversion at minimal power without amplification, however, we have no choice but to include down-conversion processes. Previous work on third-harmonic generation in cavities, similarly, considered only singly resonant cavities [76, 264]; moreover, past work focused on the case of $\chi^{(2)}$ materials where 3ω is generated by cascading two nonlinear processes (harmonic generation and frequency summing) [271, 345]. Here, we examine third-harmonic generation using $\chi^{(3)}$ materials so that only a single resonant process need be designed. Furthermore, the previous theoretical work, with a few exceptions [118], focused on one-dimensional Fabry-Perot cavity geometries, in which the problem of obtaining cavity modes with the correct frequency ratio was posed as a problem of phase-matching, and addressed by methods such as using off-normal beams. In contrast, our formalism is appropriate for any cavity geometry, and we note that in an arbitrary geometry there are countless parameters that can be tuned in order to design a particular frequency ratio.

In particular, we consider the general situation depicted schematically in Fig. 6-8: a two-mode nonlinear cavity coupled to an input/output channel. For example, a one-dimensional realization of this is shown in Fig. 6-8: a Fabry-Perot cavity between two quarter-wave stacks [239], where the stack has fewer layers on one side so that light can enter/escape. For a nonlinear effect, we consider specifically a $\chi^{(\ell)}$ nonlinearity, corresponding essentially to a shift in the refractive index proportional to the nonlinear susceptibility $\chi^{(\ell)}$ multiplied by electric field \mathbf{E} to the $(\ell - 1)^{th}$ power. Most commonly, one would have either a $\chi^{(2)}$ (Pockels) or $\chi^{(3)}$ (Kerr) effect. Such a nonlinearity results in harmonic generation [51]: light with frequency ω is coupled to light with frequency $\ell\omega$. Therefore, if we design the cavity so that it supports two modes, one at ω and one at $\ell\omega$, then input power at ω can be converted, at least partially, to output power at $\ell\omega$.

In order to achieve 100% conversion from ω to $\ell\omega$ inside a cavity, it is not always sufficient to simply increase the input power in order to increase the rate of nonlinear transfer. Putting aside the question of at what point material breakdown will occur, increasing the input power for a $\chi^{(3)}$ medium (or any odd ℓ) also causes a shift in the resonant frequency (self-phase modulation) that, unchecked, will prevent 100% conversion by making the frequency ratio $\neq \ell$. To address this mismatch, we describe how one can use two materials with opposite-sign $\chi^{(\ell)}$ to cancel the frequency-shifting effect; it may also be possible to pre-shift the cavity resonant frequency to correct for the nonlinear shift. On the other hand, a $\chi^{(2)}$

medium has no self-phase modulation, and so in this case it is sufficient to increase the input power until 100% frequency conversion is reached. Regarding material breakdown, we show that it is sufficient to use modes with a large quality factor (lifetime) Q so that a slow conversion due to a weak nonlinear effect has enough time to occur.

Nonlinear frequency conversion has previously been studied with both direct simulation [39, 199, 324] and by semi-analytic perturbative methods [33, 579]. In particular, the most common perturbative approach is known as “coupled-wave” or “coupled-mode” theory (CMT): essentially, one writes a set of ordinary differential equations for the amplitudes of the linear modes, in which these amplitudes are weakly coupled by the nonlinearity. The most common variation of this CMT approach is for waveguides, in which the only degrees of freedom are two (or more) coupled waveguide modes, as described in Ref. [520]. In our case, the problem is modified by the fact that our modes are continuously “pumped” by an external input, and the quantity of interest is the power radiated to an external output. There are two approaches to developing a CMT for such a problem. First, in the “temporal” CMT of Ref. [202, 520], the most general possible CMT is derived from fundamental principles such as conservation of energy and reciprocity, parameterized by a few unknown frequencies and coupling factors that reflect the specific geometry. Second, one can apply perturbative expansions directly to Maxwell’s equations to derive explicit expressions for the coupling factors. Although both approaches have been successfully employed to describe various nonlinear phenomena [33, 312, 345, 366], frequency conversion in doubly-resonant cavities does not seem to have been fully addressed. In the following, we apply both approaches to derive both the most general CMT and also the specific coupling factors for $\ell = 2, 3$.

We now consider the conditions under which one may achieve *complete* frequency conversion: 100% of the incident power converted to output at the second or third harmonic frequency. As we shall see, this is easiest to achieve in the $\chi^{(2)}$ case, and requires additional design criteria in the $\chi^{(3)}$ case.

The key fact in a $\chi^{(2)}$ medium is that there are no frequency-shifting terms ($\alpha = 0$), so the resonance condition $\omega_2 = 2\omega_1$ is not spoiled as one increases the power. The only requirement that we must impose is that external losses such as absorption are negligible ($\tau_{e,k} \gg \tau_{s,k}$). (We consider the effect of losses in Sec. 2.2.8.) In this lossless case, 100% conversion corresponds to setting $s_{1-} = 0$ in the steady-state. Using this fact, Eqs. (2.10–

2.11) for an input source $s_+(t) = s_{1+} \exp(i\omega_1 t)$ yields the following condition on the input power for 100% conversion:

$$|s_{1+}^{\text{crit}}|^2 = \frac{4}{\omega_1^2 |\beta_1|^2 \tau_{s,2} \tau_{s,1}^2} = \frac{\omega_1}{|\beta_1|^2 Q_2 Q_1^2} \quad (2.26)$$

(A similar dependence of efficiency on $Q_1^2 Q_2$ was previously observed [33, 309], although a critical power was not identified.) Thus, we can always choose an input power to obtain 100% conversion. If $Q_1 \sim Q_2$, then this critical power $P_{\text{crit}} = |s_{1+}^{\text{crit}}|^2$ scales as V/Q^3 where V is the modal volume (recall that $\beta \sim 1/\sqrt{V}$).

This is limited, however, by our first-order approximation: if the input power becomes so large that second-order effects (or material breakdown) become significant, then this prediction of 100% conversion is no longer valid. However, if one chooses Q_1 and/or Q_2 to be sufficiently large, then P_{crit} can be made arbitrarily small in principle. Not only does the critical power decrease with Q^3 , but the field intensity in the cavity ($|a_i|^2$) decreases as $V/Q_1 Q_2$, and thus one can avoid material breakdown as well as lowering the power.




Figure 2-3: Plot of first and second harmonic efficiency, $|s_{1-}|^2/|s_{1+}|^2$ (black) and $|s_{2-}|^2/|s_{1+}|^2$ (red), vs. $\chi^{(2)}|s_{1+}|$. 100% power transfer from ω_1 to $\omega_2 = 2\omega_1$ is achieved at $\chi^{(2)}|s_{1+}| = 1.8 \times 10^{-3}$.

To illustrate second-harmonic conversion for a $\chi^{(2)}$ medium, we plot the solution to the coupled-mode equations as a function of input power in Fig. 3-3. The 100% conversion at the predicted critical power is clearly visible. For this calculation, we chose modal parameters similar to the ones from the FDTD computation of Sec. 6.3.2: $\omega_1 = 0.3$, $\omega_2 = 0.6$, $Q_1 = 10^4$, $Q_2 = 2 \times 10^4$, with dimensionless $\beta_1 = (4.55985 - 0.7244i) \times 10^{-5}$.

A $\chi^{(3)}$ medium, on the other hand, does suffer from nonlinear frequency shifts. For example, as demonstrated in Sec. 6.3.2, Fig. 6-8, which is by no means the optimal geometry, exhibits a maximal efficiency of $|s_{3-}|^2/|s_{1+}|^2 \approx 4 \times 10^{-3}$, almost three orders of magnitude away from complete frequency conversion. We consider the effects of frequency shifting in Sec. 2.2.5, but for simplicity, we focus first on the $\alpha_{ij} = 0$ case. In reality, achieving 100% conversion efficiency even for $\alpha_{ij} \neq 0$ could be done in two ways: First, one could employ *two* $\chi^{(3)}$ materials with opposite-sign $\chi^{(3)}$ values (e.g., as in Ref. [497]). For example, if the $\chi^{(3)}$ is an *odd* function around the cavity center, then the integrals for α_{ij} will vanish while the β integrals will not. (In practice, $\alpha \ll \beta$ should suffice.) Second, one could pre-compensate for the nonlinear frequency shifts: design the cavity so that the shifted frequencies, at the P_{crit} below, satisfy the resonant condition $\omega_3 + \Delta\omega_3 = 3(\omega_1 + \Delta\omega_1)$. Equivalently, design the device for $\alpha_{ij} = 0$ and then adjust the linear cavity frequencies *a posteriori* to compensate for the frequency shift at the critical power. (This last scheme is further explored in Sec. 2.2.5.)

If α_{ij} is thereby forced to be zero, and we can also neglect external losses (absorption, etc.) as above, then 100% third-harmonic conversion ($s_{1-} = 0$) is obtained when:

$$|s_{1+}^{\text{crit}}|^2 = \left[\frac{4}{\omega_1^2 |\beta_1|^2 \tau_{s,1}^3 \tau_{s,3}} \right]^{1/2} = \left[\frac{\omega_1 \omega_3}{4 |\beta_1|^2 Q_1^3 Q_3} \right]^{1/2} \quad (2.27)$$

If $Q_1 \sim Q_3$, then this critical power scales as V/Q^2 where V is the modal volume (recall that $\beta \sim 1/V$). This is precisely the scaling that was predicted for the power to obtain nonlinear bistability in a single-mode cavity [576]. Similarly, one finds that the energy density in the cavity ($|a_i|^2$) decreases proportional to $V/\sqrt{Q_1 Q_3}$.

We demonstrate the third-harmonic conversion for $\alpha_{ij} = 0$ by plotting the solution to the coupled-mode equations as a function of input power in Fig. 2-4. Again, 100% conversion is only obtained at a single critical power. Here, we used the same parameters as in the FDTD calculation of Sec. 6.3.2, but with $\alpha = 0$. In this case, comparing with Fig. 2-5, we observe that complete frequency conversion occurs at a power corresponding to $\Delta n/n \approx 10^{-2}$. This is close to the maximum power before coupled-mode theory becomes invalid (either because of second-order effects or material breakdown), but we could easily decrease the critical power by increasing Q .

For both the $\chi^{(2)}$ and the $\chi^{(3)}$ effects, in Figs. 3-3 and 2-4, we see that the harmonic




Figure 2-4: Plot of first and third harmonic steady-state efficiency, $|s_{1-}|^2/|s_{1+}|^2$ (black line) and $|s_{3-}|^2/|s_{1+}|^2$ (red line), for $\alpha = 0$ (no self/cross-phase modulation), as a function of input power $|s_{1+}|^2$ scaled by the Kerr coefficient $n_2 = 3\chi^{(3)}/4c\epsilon$. The reflected power at the incident frequency ω_1 is shown as a dashed black line. 100% power transfer from ω_1 to $\omega_3 = 3\omega_1$ is achieved at $n_2|s_{1+}|^2 = 2.8 \times 10^{-4}$. The parameters used in this plot are $Q_1 = 1000$, $Q_3 = 3000$, $\beta_1 = (4.55985 - 0.7244i) \times 10^{-5}$ in dimensionless units of $\chi^{(3)}/V\epsilon$.

conversion efficiency goes to zero if the input power (or χ) is either too small or too large. It is not surprising that frequency conversion decreases for low powers, but the decrease in efficiency for high powers is less intuitive. It corresponds to a well-known phenomenon in coupled-mode systems: in order to get 100% transmission from an input port to an output port, the coupling rates to the two ports must be matched in order to cancel the back-reflected wave [146, 202]. In the present case, the coupling rate to the input port is $\sim 1/Q_1$, and the coupling rate to the output “port” (the harmonic frequency) is determined by the strength of the nonlinear coupling. If the nonlinear coupling is either too small or too large, then the rates are not matched and the light is reflected instead of converted. (On the other hand, we find that for large input powers, while the conversion *efficiency* as a fraction of input power goes to zero, the *absolute* converted power ($|s_{2-}|^2$ or $|s_{3-}|^2$) goes to a constant.)

We have presented a rigorous coupled-mode theory for second- and third-harmonic generation in doubly resonant nonlinear cavities, accurate to first order in the nonlinear susceptibility. Our theory predicts several interesting consequences. First, it is possible to design the cavity to yield 100% frequency conversion in a passive (gain-free) device, even when nonlinear down-conversion processes are included, limited only by fabrication imperfections and losses. Second, this 100% conversion requires a certain critical input power—powers

either too large or too small lead to lower efficiency. Third, we describe how to compensate for the self-phase modulation in a $\chi^{(3)}$ cavity. The motivation for this work was the hope that a doubly resonant cavity would lead to 100% conversion at very low input powers, and so we conclude this section by estimating the critical power for reasonable material and geometry assumptions.

A typical nonlinear material is gallium arsenide (GaAs), with $\chi^{(2)} \approx 145$ pm/V and $n_2 = 1.5 \times 10^{-13}$ cm²/W at 1.5 μ m. (Al doping is usually employed to decrease nonlinear losses near resonance [543].) Although this has both $\chi^{(2)}$ and $\chi^{(3)}$ effects, we can selectively enhance one or the other by choosing the cavity to have resonances at either the second or third harmonic. Many well confined optical cavity geometries are available at these wavelengths and have been used for nonlinear devices, such as ring resonators [568] or photonic-crystal slabs [387]. We will assume conservative parameters for the cavity: a lifetime $Q_1 = 1000$, $Q_2 = 2000$, $Q_3 = 3000$, and a modal volume of 10 cubic half-wavelengths ($V \approx 10(\lambda/2n)^3$) with roughly constant field amplitude in the nonlinear material (worse than a realistic case of strongly peaked fields). In this case, the critical input power, from Eqs. (2.26–2.27), becomes approximately 70 μ W for second-harmonic generation and 2 mW for third-harmonic generation (with a moderate peak index shift $\Delta n/n \approx 10^{-3}$, justifying our first-order approximation).

In the following section, we check the correctness of our theoretical predictions (for the $\chi^{(3)}$ case) via numerical validation with exact FDTD simulations. In particular, we shall directly observe the consequences of ignoring the cross- and self-phase modulation effects present in THG.

2.2.4 Numerical Validation

To check the predictions of the $\chi^{(3)}$ coupled-mode equations [Eqs. (2.12–2.13)], we performed an FDTD simulation of the one-dimensional waveguide-cavity system shown in Fig. 6-8, whose analytical properties are uniquely suited to third-harmonic generation. This geometry consists of a semi-infinite photonic-crystal structure made of alternating layers of dielectric ($\epsilon_1 = 13$ and $\epsilon_2 = 1$) with period a and thicknesses given by the quarter-wave condition ($d_1 = \sqrt{\epsilon_2}/(\sqrt{\epsilon_1} + \sqrt{\epsilon_2})$ and $d_2 = a - d_1$, respectively). Such a quarter-wave stack possesses a periodic sequence of photonic band gaps centered on frequencies $\omega_1 = (\sqrt{\epsilon_1} + \sqrt{\epsilon_2})/4\sqrt{\epsilon_1\epsilon_2}$ (in units of $2\pi c/a$) for the lowest gap, and higher-order gaps centered on odd multiples of

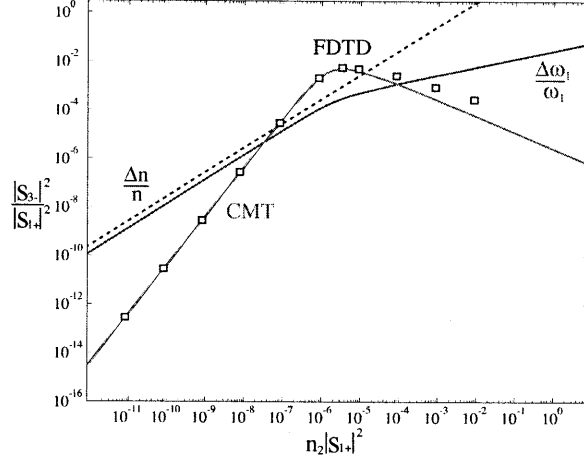


Figure 2-5: Log-log plot of $|s_{3-}|^2/|s_{1+}|^2$ vs. $n_2|s_{1+}|^2$ for the coupled-mode theory (grey) and FDTD (black squares), where n_2 is being varied. Also shown are the corresponding $\Delta n/n$ (dashed blue) and $\Delta\omega_1/\omega_1$ (solid red) curves.

ω_1 . Moreover, a defect formed by doubling the thickness of a ε_1 layer creates cavity modes at exactly the middle of every one of these gaps. Therefore, it automatically satisfies the frequency-matching condition for third-harmonic generation. In fact, it is too good: there will also be “ninth harmonic” generation from ω_3 to ω_9 . This unwanted process is removed, however, by the discretization error of the FDTD simulation, which introduces numerical dispersion that shifts the higher-frequency modes. To ensure the $\omega_3 = 3\omega_1$ condition in the face of this dispersion, we slightly perturbed the structure (increasing the dielectric constant slightly at the nodes of the third-harmonic eigenfield) to tune the frequencies. The simulated crystal was effectively semi-infinite, with many more layers on the right than on the left of the cavity. On the left of the cavity, after two period of the crystal the material is simply air ($\varepsilon = 1$), terminated by a perfectly matched layer (PML) absorbing boundary region.

We excite the cavity with an incident plane wave of frequency ω_1 , and compute the resulting reflection spectrum. The reflected power at ω_3 , the third-harmonic generation, was then compared with the prediction of the coupled-mode theory. The frequencies, decay rates, and α and β coefficients in the coupled-mode theory were computed from a linear FDTD simulation in which the eigenmodes were excited by narrow-band pulses. The freely available FDTD code of [150] was employed.

The results are shown in Fig. 2-5, in which the output power at ω_1 and $\omega_3 = 3\omega_1$ is denoted by $|s_{1+}|^2$ and $|s_{3-}|^2$, respectively, while the input power at ω_1 is denoted by $|s_{1+}|^2$.

In particular, we plot convenient dimensionless quantities: the third-harmonic conversion efficiency $|s_{3-}|^2 / |s_{1+}|^2$ as a function of the dimensionless product $n_2 |s_{1+}|^2$ in terms of the standard Kerr coefficient $n_2 = 3\chi^{(3)}/4c\epsilon$. There is clear agreement between the FDTD and CMT for small values of $n_2 |s_{1+}|^2$ (in which limit the conversion goes quadratically with $n_2 |s_{1+}|^2$). However, as the input power increases, they eventually begin to disagree, marking the point where second-order corrections are required. This disagreement is not a practical concern, however, because the onset of second-order effects coincides with the limits of typical materials, which usually break down for $\Delta n/n \equiv \chi^{(3)} \max|E|^2/2\epsilon > 1\%$. This is why we also plot the maximum index shift $\Delta n/n$ in the same figure.

Also shown in Fig. 2-5 is a plot of $\Delta\omega_1/\omega_1 = \text{Re}[\delta\omega_1/\omega_1]$. As expected, when $\Delta\omega_1$ is of the order of $1/Q_1 \sim 10^{-3}$, the frequency shift begins to destroy the frequency matching condition, substantially degrading the third-harmonic conversion. (It might seem that $\Delta n/n$ and $\Delta\omega_1/\omega_1$ should be comparable, but this is not the case because $\Delta n/n$ is the maximum index shift while $\Delta\omega_1/\omega_1$ is due to an average index shift.)

More specifically, the details of our simulation are as follows. To simulate a continuous wave (CW) source spectrum in FDTD, we employ a narrow-frequency Gaussian pulse incident from the air region. This pulse is carefully normalized so that the peak *intensity* is unity, to match the CMT. The field in the air region is Fourier transformed and subtracted from the incident field to yield the reflected flux. Using only two periods of quarter-wave stack on the left of the cavity we obtained two cavity modes with real frequencies $\omega_1 = 0.31818 (2\pi c/a)$, $\omega_2 = 0.95454 (2\pi c/a)$ and quality factors $Q_1 = 1286$ and $Q_3 = 3726$, respectively. Given these field patterns, we computed the α_{ij} and β_i coefficients. We obtained the following coupling coefficients, in units of $\chi^{(3)}/a$: $\alpha_{11} = 4.7531 \times 10^{-4}$, $\alpha_{22} = 5.3306 \times 10^{-4}$, $\alpha_{12} = \alpha_{21} = 2.7847 \times 10^{-4}$, $\beta_1 = (4.55985 - 0.7244) \times 10^{-5}$.

2.2.5 Kerr Cross- and Self-Phase Modulation

In this section, I present work described in Ref. [198], in which we extend the theoretical predictions of THG in cavities with Kerr nonlinearities described in Sec. 2.2.1–6.3.2, to the realistic case in which there exist self- and cross-phase modulation effects (which can drive the cavities out of resonance and thereby affect the dynamical stability of the high-efficiency steady-state solutions). In particular, we show that such doubly-resonant nonlinear optical systems can display a rich variety of dynamical behaviors, including multistability (dif-

ferent steady states excited by varying initial conditions, a richer version of the bistable phenomenon observed in single-mode cavities [501]), gap solitons [527], long-period limit cycles (similar to the “self-pulsing” observed for second-harmonic generation [125, 190]), and transitions in the stability and multiplicity of solutions as the parameters are varied. Furthermore, we demonstrate the possibility of achieving high-efficiency conversion with arbitrarily small power requirements for nonlinear devices, especially when compared against singly-resonant cavities or non-resonant structures [15, 16, 38, 65, 98, 121, 123, 172, 192, 275, 312, 345, 366, 371, 406, 470, 476, 498]. An appreciation and understanding of these behaviors is important to design efficient harmonic converters (the main focus of this chapter), but it also opens the possibility of new types of devices enabled by other aspects of the nonlinear dynamics. For example, strong Kerr nonlinearities are desired in the context of quantum information theory for use in low-loss photon entanglement and other single-photon applications [37, 173, 264, 377, 379].

As noted above, in a Kerr ($\chi^{(3)}$) medium, there is a change in the refractive index proportional to the square of the electric field; for an oscillating field at a frequency ω , this results in a shift in the index at the same frequency (self-phase modulation, SPM), generation of power at the third-harmonic frequency 3ω , and also other effects when multiple frequencies are present [cross-phase modulation (XPM) and four-wave mixing (FWM)] [51]. When the field is confined in a cavity, restricting to a small modal volume V for a long time given by the quality factor Q (a lifetime in units of the optical period) [241], such nonlinear effects are enhanced by both the increased field strength for the same input power and by the frequency sensitivity inherent in resonant effects (since the fractional bandwidth is $1/Q$). This enhancement is exploited, for example, in nonlinear harmonic and sum-frequency generation, most commonly for $\chi^{(2)}$ effects where the change in index is proportional to the electric field (which requires a non-centrosymmetric material) [51]. One can further enhance harmonic generation by using a cavity with *two* resonant modes, as demonstrated in Sec. 2.2.1, one at the source frequency and one at the harmonic frequency [33, 125, 126, 309, 324, 393, 402, 563, 602]. In this case, one must also take into account a nonlinear downconversion process that competes with harmonic generation [125, 393, 563], but it turns out to be theoretically possible to obtain 100% harmonic conversion for either $\chi^{(2)}$ ($\omega \rightarrow 2\omega$) or $\chi^{(3)}$ ($\omega \rightarrow 3\omega$) nonlinearities at a specific “critical” input power P_{crit} (both in a one-dimensional model of propagating waves for $\chi^{(2)}$ nonlinearities [474] and also the

general coupled-mode model for either $\chi^{(2)}$ or $\chi^{(3)}$ nonlinearities [445]).

In the previous sections, we studied the harmonic-generation and downconversion processes within the model system of Fig. 6-8: a single input channel (e.g. a waveguide) coupled to a nonlinear cavity with two resonant frequencies, where both reflected and harmonic fields are emitted back into the input channel. In that case, we predicted 100% harmonic generation at a critical power P_{crit} proportional to V/Q^3 for $\chi^{(2)}$ and V/Q^2 for $\chi^{(3)}$ [445]. However, we only looked at the steady-state solution of the system and not its dynamics or stability. Moreover, in the $\chi^{(3)}$ case there can also be an SPM and XPM effect that shifts the cavity frequencies out of resonance and spoils the harmonic-generation effect. Here, we consider both of these effects, describe how to compensate for SPM and XPM, and demonstrate the different regimes of stability in such $\chi^{(3)}$ doubly resonant systems. We show that the parameters and the initial conditions must be chosen within certain regimes to obtain a stable steady state with high conversion efficiency. In other regimes, we demonstrate radically different behaviors: not only low-efficiency steady states, but also limit-cycle solutions where the efficiency oscillates slowly with a repetition period of many thousands of optical cycles. With infrared light, these limit cycles form a kind of optical oscillator (clock) with a period in the hundreds of GHz or THz (and possibly lower, depending on the cavity parameters). Previously, limit-cycle (self-pulsing) behaviors have been observed in a number of other nonlinear optical systems, such as: doubly-resonant $\chi^{(2)}$ cavities coupled by second-harmonic generation [125]; bistable multimode Kerr cavities with time-delayed nonlinearities [2]; nonresonant distributed feedback in Bragg gratings [398]; and a number of nonlinear lasing devices [492]. However, the system considered in this work seems unusually simple, especially among $\chi^{(3)}$ systems, in that it only requires two modes and an instantaneous Kerr nonlinearity, with a constant-frequency input source, to attain self-pulsing, and partly as a consequence of this simplicity the precise self-pulsing solution is quite insensitive to the initial conditions. In other nonlinear optical systems where self-pulsing was observed, other authors have also observed chaotic solutions in certain regimes. Here, we did not observe chaos for any of the parameter regimes considered, where the input was a constant-frequency source, but it is possible that chaotic solutions may be excited by an appropriate pulsed input as in the $\chi^{(2)}$ case [125, 190].

Another interesting phenomenon that can occur in nonlinear systems is multistability, where there are multiple possible steady-state solutions that one can switch among by

varying the initial conditions. In Kerr ($\chi^{(3)}$) media, an important example of this phenomenon is bistable transmission through nonlinear cavities: for transmission through a *single*-mode cavity, output can switch discontinuously between a high-transmission and a low-transmission state in a hysteresis effect that results from SPM [501]. For example, if one turns on the power gradually from zero the system stays in the low-transmission state, but if the power is increased further and then decreased to the original level, the system can be switched to the high-transmission state. This effect, which has been observed experimentally [387], can be used for all-optical logic, switching, rectification, and many other functions [501]. In a cavity with multiple closely-spaced resonances, where the nonlinearity is strong enough to shift one cavity mode’s frequency to another’s, the same SPM phenomenon can lead to more than two stable solutions [153]. Here, we demonstrate a much richer variety of multistable phenomena in the doubly-resonant case for widely-separated cavity frequencies coupled by harmonic generation in addition to SPM—not only can there be more than two stable states, but the transitions between them can exhibit complicated oscillatory behaviors as the initial conditions are varied, and there are also Hopf bifurcations into self-pulsing solutions.

The discussion of the THG dynamics is structured as follows: In Sec. 2.2.6, we demonstrate the possibility of achieving 100% THG conversion by compensating for frequency shifting due to SPM and XPM, i.e. by pre-shifting the cavity frequencies. In Sec. 2.2.5, we analyze the stability of the 100%-efficiency solution, and demonstrate the different regimes of stable operation that are achieved in practice starting from that theoretical initial condition. We also present bifurcation diagrams that show how the stable and unstable solutions evolve as the parameters vary. Finally, in Sec. 2.2.7, we consider how to excite these high-efficiency solutions in practice, by examining the effect of varying initial conditions and uncertainties in the cavity parameters. In particular, we demonstrate the multistable phenomena exhibited as the initial conditions are varied. We close with some concluding remarks, discussing the many potential directions for future work that are revealed by the phenomena described here.

There are three different α_{ij} parameters in Eqs. (2.22–2.25) (two SPM coefficients α_{11} and α_{33} and one XPM coefficient $\alpha_{13} = \alpha_{31}$). All three values are different, in general, but are determined by similar integrals of the field patterns, produce similar frequency-shifting phenomena, and all scale as $1/V$. Therefore, in order to limit the parameter space analyzed

here, we consider the simplified case where all three frequency-shifting terms have the same strength $\alpha_{ij} = \alpha$.

One can also include various losses, e.g. linear losses correspond to a complex ω_1 and/or ω_3 , and nonlinear two-photon absorption corresponds to a complex α . As discussed in Sec. 2.2.8 and in the concluding remarks, however, we have found that such considerations do not qualitatively change the results (only reducing the efficiency somewhat, as long as the losses are not too big compared to the radiative lifetimes τ), and so in this section, we restrict ourselves to the idealized lossless case.

Figure 2-4 shows the steady-state conversion efficiency ($|s_{3-}|^2/|s_{1+}|^2$) versus input power of light that is incident on the cavity at ω_1^{cav} , for a given set of geometric parameters (i.e. assuming negligible self- and cross-phase modulation so that $\alpha = 0$), and not considering the stability of the steady state. As shown by the solid red curve, as one increases the input power, the efficiency increases, peaking at 100% conversion for a critical input power $P_{\text{crit}} = |s_{1+}^{\text{crit}}|^2$, where

$$|s_{1+}^{\text{crit}}|^2 = \left(\frac{4}{|\omega_1 \beta_1|^2 \tau_1^3 \tau_3} \right)^{1/4}. \quad (2.28)$$

The efficiency decreases if the power is either too low (in the linear regime) or too high (dominated by down-conversion). P_{crit} scales as V/Q^2 , so one can in principle obtain very low-power efficient harmonic conversion by increasing Q and/or decreasing V [445]. Again, including absorption or other losses decreases the peak efficiency, but does not otherwise qualitatively change this solution [445].

There are two effects that we did not previously analyze in detail, however, which can degrade this promising solution: nonlinear frequency shifts and instability. Here, we first consider frequency shifts, which arise whenever $\alpha \neq 0$, and consider stability in the next section. The problem with the α terms is that efficient harmonic conversion depends on the cavities being tuned to harmonic frequencies $\omega_3 = 3\omega_1$; a nonlinear shift in the cavity frequencies due to self- and cross-phase modulation will spoil this resonance. In principle, there is a straightforward solution to this problem, as depicted in Fig. 2-6. Originally (for $\alpha = 0$), the cavity was designed to have the frequency ω_1 in the linear regime, but with $\alpha \neq 0$ the effective cavity frequency ω_1^{NL} (including self- and cross-phase modulation terms) is shifted away from the design frequency as shown by the blue line. Instead, we can simply design the linear cavity to have a frequency ω_1^{cav} slightly different from the

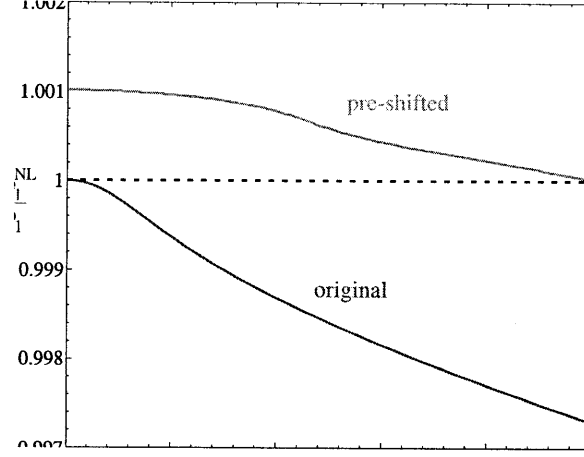


Figure 2-6: Shift in the resonant frequency ω_1^{NL} as a function of input power, due to self- and cross-phase modulation. (There is an identical shift in ω_3^{NL} .) If the cavity is designed so that the linear ($P_{\text{in}} \rightarrow 0$) frequencies are harmonics, the nonlinearity pushes the system out of resonance (lower blue line) as the power increases to the critical power for 100% efficiency. This is corrected by pre-shifting the cavity frequencies (upper green line) so that the nonlinear frequency shift pushes the modes into resonance at P_{crit} .

operating frequency ω_1 , so that self- and cross-phase modulation shifts ω_1^{NL} exactly to ω_1 at the critical input power, as depicted by the green line in Fig. 2-6. Exactly the same strategy is used for ω_3^{NL} , by pre-shifting ω_3^{cav} .

More precisely, to compute the required amount of pre-shifting, we examine the coupled-mode equations Eqs. (2.12–2.13). First, we solve for the critical power P_{crit} assuming $\alpha = 0$, as in Ref. [445], and obtain the corresponding critical cavity fields a_k^{crit} :

$$|a_1^{\text{crit}}|^2 = \left(\frac{1}{\omega_1^2 |\beta_1|^2 \tau_3 \tau_{1,s}} \right)^{1/2} \quad (2.29)$$

$$|a_3^{\text{crit}}|^2 = \left(\frac{\omega_3 \beta_3 \tau_3}{(\omega_1 \beta_1 \tau_{1,s})^3} \right)^{1/2}. \quad (2.30)$$

Then, we substitute these critical fields into the coupled-mode equations for $\alpha \neq 0$, and solve for the new cavity frequencies ω_k^{cav} so as to cancel the α terms and make the a_k^{crit} solutions still valid. This yields the following transformation of the cavity frequencies:

$$\omega_1^{\text{cav}} = \frac{\omega_1}{1 - \alpha_{11}|a_1^{\text{crit}}|^2 - \alpha_{13}|a_3^{\text{crit}}|^2} \quad (2.31)$$

$$\omega_3^{\text{cav}} = \frac{\omega_3}{1 - \alpha_{13}|a_1^{\text{crit}}|^2 - \alpha_{33}|a_3^{\text{crit}}|^2}. \quad (2.32)$$

By inspection, when substituted into Eqs. (2.12–2.13) at the critical power, these yield the same steady-state solution as for $\alpha = 0$. (There are two other appearances of ω_1 and ω_3 in the coupled-mode equations, in the β_k terms, but we need not change these frequencies because that is a higher-order effect, and the derivation of the coupled-mode equations considered only first-order terms in $\chi^{(3)}$.)

The nonlinear dynamics turn out to depend only on four dimensionless parameters: $\tau_3/\tau_1 = Q_3/3Q_1$, α_{11}/β_1 , α_{33}/β_1 , and $\alpha_{13}/\beta_1 = \alpha_{31}/\beta_1$. The overall scale of Q , α , etcetera, merely determines the absolute scale for the power requirements: it is clear from the equations that multiplying all α and β coefficients by an overall constant K can be compensated by dividing all a and s amplitudes by \sqrt{K} [which happens automatically for s at the critical power by Eq. (2.28)]; the case of scaling $\tau_{1,3}$ by an overall constant is more subtle and is considered below. As mentioned above, for simplicity we take $\alpha_{11} = \alpha_{33} = \alpha_{13} = \alpha_{31} = \alpha$. Therefore, in the subsequent sections we will analyze the asymptotic efficiency as a function of τ_3/τ_1 and α/β_1 .

So far, we have found a steady-state solution to the coupled-mode equations, including self- and cross-phase modulation, that achieves 100% third-harmonic conversion. In the following sections, we consider under what conditions this solution is stable, what other stable solutions exist, and for what initial conditions the high-efficiency solution is excited.

To understand the dynamics and stability of the nonlinear coupled-mode equations, we apply the standard technique of identifying fixed points of the equations and analyzing the stability of the linearized equations around each fixed point [521].

By a “fixed point,” we mean a steady-state solution corresponding to an input frequency ω_1 ($s_{1+} \sim e^{i\omega_1 t}$) and hence $a_1(t) = A_1 e^{i\omega_1 t}$ and $a_3(t) = A_3 e^{i3\omega_1 t}$ for some unknown constants A_1 and A_3 . [An input frequency ω_1 can also generate higher harmonics, such as $9\omega_1$ or $5\omega_1$, but these are negligible: both because they are higher-order effects ($\sim [\chi^{(3)}]^2$, and all such terms were dropped in deriving the coupled-mode equations), and because we assume there is no resonant mode present at those frequencies.] By substituting this steady-state form into Eqs. (2.12–2.13), one obtains two coupled polynomial equations whose roots are the fixed points. We already know one of the fixed points from the previous section, the 100% efficiency solution, but to fully characterize the system one would like to know all of the fixed points (both stable and unstable). We solved these polynomial equations using Mathematica, which is able to compute all of the roots, but some transformations were

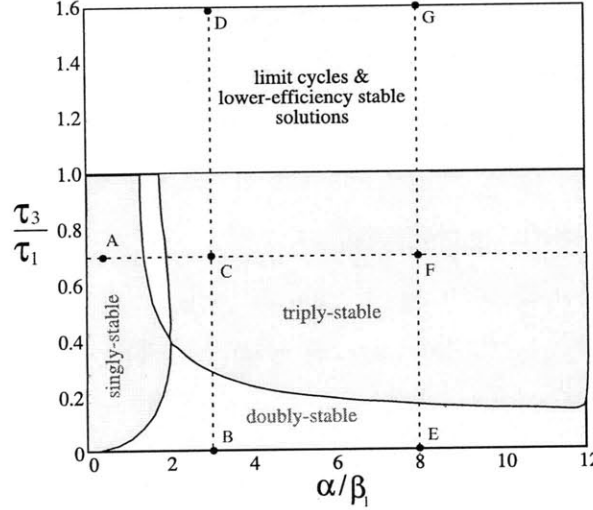


Figure 2-7: Phase diagram of the nonlinear dynamics of the doubly-resonant nonlinear harmonic generation system from Fig. 6-8 as a function of the relative cavity lifetimes ($\tau_3/\tau_1 = 3Q_3/Q_1$) and the relative strength of SPM and XPM vs. harmonic generation (α/β_1) for input power equal to the critical power for 100% efficiency. For $\tau_3 < \tau_1$ there is always one *stable* 100%-efficiency solution, and for nonzero α the system may have additional stable solutions. For $\tau_3 > \tau_1$ the 100%-efficiency solution becomes unstable, but there are limit cycles and lower-efficiency stable solutions. Various typical points A–G in each region are labeled for reference in the subsequent figures.

required to put the equations into a solvable form, as explained in more detail in the appendix.

As mentioned above, the dynamics are independent of the overall scale of $\tau_{1,3}$, and depend only on the ratio τ_3/τ_1 . This can be seen from the equations for $A_{1,3}$, in which the $\omega_{1,3}$ oscillation has been removed. In these equations, if we multiply τ_1 and τ_3 by an overall constant factor K , after some algebra it can be shown that the $A_{1,3}$ equations are invariant if we rescale $A_1 \rightarrow A_1/\sqrt{K}$, $A_3 \rightarrow A_3/\sqrt{K}$, rescale time $t \rightarrow Kt$, and rescale the input $s_{1+} \rightarrow s_{1+}/K$ [which happens automatically for the critical power by Eq. (2.28)]. Note also that the conversion efficiency $|s_{3-}/s_{1+}|^2 = (2/\tau_3)|A_3/s_{1+}|^2$ is also invariant under this rescaling by K . That is, the powers and the timescales of the dynamics change if you change the lifetimes, unsurprisingly, but the steady states, stability, etcetera (as investigated in the next section) are unaltered.

2.2.6 Stability and Dynamics of XPM and SPM

Given the steady-state solutions (the roots), their stability is determined by linearizing the original equations around these points to a first-order linear equation of the form $dx/dt = Ax$; a stable solution is one for which the eigenvalues of A have negative real parts (leading to solutions that decay exponentially towards the fixed point) [521].

The results of this fixed-point and stability analysis are shown in Fig. 2-7 as a “phase diagram” of the system as a function of the relative lifetimes $\tau_3/\tau_1 = 3Q_3/Q_1$ and the relative strength of self-phase-modulation vs. four-wave mixing α/β_1 . Our original 100%-efficiency solution is always present, but is only stable for $\tau_3 < \tau_1$ and becomes unstable for $\tau_3 > \tau_1$. The transition point, $\tau_3 = \tau_1$, corresponds to equal energy $|a_1|^2 = |a_3|^2$ in the fundamental and harmonic mode at the critical input power. The unstable region corresponds to $|a_3|^2 > |a_1|^2$ (and the down-conversion term is stronger than the up-conversion term)—intuitively, this solution is unstable because, if any perturbation causes the energy in the harmonic mode to decrease, there is not enough pumping from up-conversion to bring it back to the 100%-efficiency solution. Conversely, in the stable $|a_3|^2 < |a_1|^2$ ($\tau_3 < \tau_1$) regime, the higher-energy fundamental mode is being directly pumped by the input and can recover from perturbations. Furthermore, as α/β_1 increases, additional lower-efficiency stable solutions are introduced, resulting in regimes with two (doubly stable) and three (triply stable) stable fixed points. These different regimes are explored in more detail via bifurcation diagrams below, and the excitation of the different stable solutions is considered in the next section.

For $\tau_3 > \tau_1$, the 100%-efficiency solution is unstable, but there are lower-efficiency steady-state solutions and also another interesting phenomenon: limit cycles. A limit cycle is a stable oscillating-efficiency solution, one example of which (corresponding to point D in Fig. 2-7) is plotted as a function time in Fig. 2-8. (In general, the existence of limit cycles is difficult to establish analytically [521], but the phenomenon is clear in the numerical solutions as a periodic oscillation insensitive to the initial conditions). In fact, as we shall see below, these limit cycles result from a “Hopf bifurcation,” which is a transition from a stable fixed point to an unstable fixed point and a limit cycle [516]. In this example at point D, the efficiency oscillates between roughly 66% and nearly 100%, with a period of several thousand optical cycles. As a consequence of the time scaling described in the




Figure 2-8: An example of a limit-cycle solution, with a periodically oscillating harmonic-generation efficiency as a function of time, corresponding to point D in Fig. 2-7. Perturbations in the initial conditions produce only phase shifts in the asymptotic cycle. Here, the limit cycle has a period of around 3×10^4 optical cycles. *Inset*: Square of Fourier amplitudes (arbitrary units) for each harmonic component of the limit cycle in the Fourier-series expansion of the A_3 .

last paragraph of the previous section, the period of such limit cycles is proportional to the τ 's. If the frequency ω_1 were $1.55 \mu\text{m}$, for a Q_1 of 500 optical cycles, this limit cycle would have a frequency of around 70 GHz, forming an interesting type of optical “clock” or oscillator. Furthermore, the oscillation is not sinusoidal and contains several higher harmonics as shown in the inset of Fig. 2-8; the dominant frequency component in this case is the fourth harmonic (~ 280 GHz), but different points in the phase diagram yield limit cycles with different balances of Fourier components.

To better understand the phase diagram of Fig. 2-7, it is useful to plot the efficiencies of both the stable and unstable solutions as a function of various parameters. Several of these bifurcation diagrams (in which new fixed points typically appear in stable-unstable pairs) are shown in Figs. 2-9–2-11. To begin with, Figs. 2-9 and 2-10 correspond to lines connecting the labeled points ACF, BCD, and ECG, respectively, in Fig. 2-7, showing how the stability changes as a function of α/β_1 and τ_3/τ_1 . Figure 2-9 shows how first one then two new stable fixed points appear as α/β_1 is increased, one approaching zero efficiency and the other closer to 50%. Along with these two stable solutions appear two unstable solutions (dashed lines). (A similar looking plot, albeit inverted, can be found in Ref. [153] for SPM-coupled closely-spaced resonances.) In particular, the fact that one of the unstable

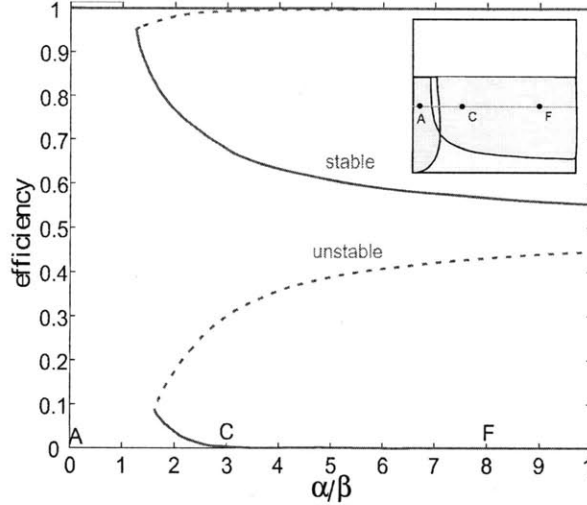


Figure 2-9: Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of α/β_1 for a fixed $\tau_3/\tau_1 = 0.7$, corresponding to the line ACF in Fig. 2-7 (see inset). The input power is the critical power P_{crit} , so there is always a 100%-efficiency stable solution, but as α/β_1 increases new stable and unstable solutions appear at lower efficiencies.

solutions approaches the 100%-efficiency stable solution causes the latter to have a smaller and smaller basin of attraction as α/β_1 increases, making it harder to excite as described in the next section. The next two plots, in Fig. 2-10, both show the solutions with respect to changes in τ_3/τ_1 at two different values of α/β_1 . They demonstrate that at $\tau_1 = \tau_3$, a Hopf bifurcation occurs where the 100%-efficiency solution becomes unstable for $\tau_3 \geq \tau_1$ and limit cycles appear, intuitively seeming to “bounce between” the two nearby unstable fixed points. (The actual phase space is higher dimensional, however, so the limit cycles are not constrained to lie strictly between the efficiencies of the two unstable solutions.) It is worth to note that the remaining nonzero-efficiency stable solution (which appears at a nonzero τ_3/τ_1) becomes less efficient as τ_3/τ_1 increases.

The above analysis and results were for the steady-state-solutions when operating at the critical input power to obtain a 100%-efficiency solution. However, one can, of course, operate with a different input power—although no other input power will yield a 100%-efficient steady-state solution, different input powers may still be useful because (as noted above and in the next section) the 100%-efficiency solution may be unstable or practically unattainable. Figure 2-11(left) is the bifurcation diagram with respect to the input power $P_{\text{in}}/P_{\text{crit}}$ at fixed α/β_1 and fixed τ_3/τ_1 , corresponding to point C in Fig. 2-7. This power bifurcation diagram displays a number of interesting features, with the steady-state solutions

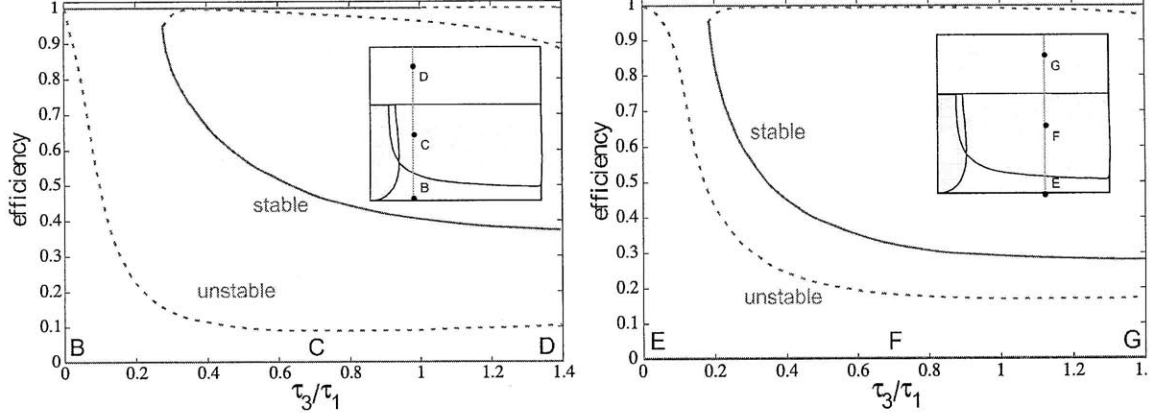


Figure 2-10: Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of τ_3/τ_1 for a fixed $\alpha/\beta_1 = 3$ (left) or $= 8$ (right), corresponding to the lines BCD or EFG, respectively, in Fig. 2-7 (see insets). The input power is the critical power P_{crit} , so there is always a 100%-efficiency steady-state solution, but it becomes unstable for $\tau_3 > \tau_1$ (a Hopf bifurcation leading to limit cycles as in Fig. 2-8).

transitioning several times from stable to unstable and vice versa. As we will see in the next section, the stability transitions in the uppermost branch are actually supercritical (reversible) Hopf bifurcations to/from limit cycles. Near the critical power, there is only a small region of stability of the near-100%-efficiency solution, as shown in the inset of Fig. 2-11(left). In contrast, the lower-efficiency stable solutions have much larger stable regions of the curve while still maintaining efficiencies greater than 70% at low powers comparable to $P_{\text{crit}} \sim V/Q^2$, which suggests that they may be attractive regimes for practical operation when α/β_1 is not small. This is further explored in the next section, and also by Fig. 2-11(right) which shows the bifurcation diagram along the line ACF in Fig. 2-7 [similar to Fig. 2-9], but at 135% of the critical input power. For this higher power, the system becomes at most doubly stable as α/β_1 is increased, and the higher-efficiency stable solution becomes surprisingly close to 100% as $\alpha/\beta_1 \rightarrow 0$.

2.2.7 Exciting High-Efficiency Solutions

One remaining concern in any multistable system is how to excite the desired solution—depending on the initial conditions, the system may fall into different stable solutions, and simply turning on the source at the critical input power may result in an undesired low-efficiency solution. If α/β is small enough, of course, then from Fig. 2-7 the high-efficiency solution is the only stable solution and the system must inevitably end up in this state




Figure 2-11: (Left:) Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of input power P_{in}/P_{crit} at fixed $\alpha/\beta_1 = 3$ and $\tau_3/\tau_1 = 0.7$, corresponding to point C in Fig. 2-7; the inset shows an enlarged view of the high-efficiency solutions. (Right:) Bifurcation diagram as a function of α/β_1 for fixed $P_{in}/P_{crit} = 1.35$ and fixed $\tau_3/\tau_1 = 0.7$; in this case, because it is not at the critical power, there are no 100%-efficiency solutions.

no matter how the critical power is turned on. Many interesting physical systems will correspond to larger values of α/β , however [445], and in this case the excitation problem is complicated by the existence of other stable solutions. Moreover, the basins of attraction of each stable solution may be very complicated in the phase space, as illustrated by Fig. 2-12, where varying the initial cavity amplitudes $A_{1,3}$ from the 100%-efficiency solution causes the steady state to oscillate in a complicated way between the three stable solutions (at point C in Fig. 2-7). We have investigated several solutions to this excitation problem, and found an “adiabatic” excitation technique that reliably produces the high-efficiency solution without unreasonable sensitivity to the precise excitation conditions.

First, we considered a simple technique similar to the one described in Ref. [501] for exciting different solutions of a bistable filter: as shown in Fig. 2-13, we “turn on” the input power by superimposing a gradual exponential turn-on (asymptoting to $P_1 = P_{crit}$) with a Gaussian pulse of amplitude P_0 and width δT . The function of the initial pulse is to “kick” the system into the desired stable solution. We computed the eventual steady-state efficiency (after all transient effects have disappeared) as a function of the pulse amplitude P_0 at point C in Fig. 2-7, where there are three stable solutions. The results are shown in Fig. 2-14(left), and indeed we see that all three stable solutions from point C in Fig. 2-9: one at near-zero efficiency, one at around 47% efficiency, and one at 100% efficiency. Unfortunately,




Figure 2-12: Asymptotic steady-state efficiency at point C (triply-stable) in the phase diagram (Fig. 2-7), with the initial conditions perturbed from the 100%-efficiency stable solution. The initial amplitudes A_{10} and A_{30} are perturbed by δA_{10} and δA_{30} , respectively, with $\delta A_{10}/A_1^{\text{crit}} = \delta A_{30}/A_3^{\text{crit}}$. The oscillation of the steady-state efficiency with the perturbation strength is an indication of the complexity of the phase space and the shapes of the basins of attraction of each fixed point.

the 100% efficiency solution is obviously rather difficult to excite, since it occurs for only a very narrow range of P_0 values. One approach to dealing with this challenge is to relax the requirement of 100% efficiency (which will never be obtained in practice anyway due to losses), and operate at a power $P_1 < P_{\text{crit}}$. In particular, Fig. 2-11(left) shows that there is a much larger stable region for $P_1 \approx 0.8P_{\text{crit}}$ with efficiency around 90%, leading one to suspect that this solution may be easier to excite than the 100%-efficiency solution at $P_1 = P_{\text{crit}}$. This is indeed the case, as is shown in Fig. 2-14(right), plotting efficiency vs. P_0 at point C with $P_1 \approx 0.8P_{\text{crit}}$. In this case, there are only two stable solutions, consistent with Fig. 2-11(left), and there are much wider ranges of P_0 that attain the high-efficiency ($\approx 90\%$) solution.

There are also many other ways to excite the high-efficiency solution (or whatever steady-state solution is desired). For example, because the cavity is initially detuned from the input frequency, as described in Sec. 2.2.5, much of the initial pulse power is actually reflected during the transient period, and a more efficient solution would vary the pulse frequency in time to match the cavity frequency as it detunes. One can also, of course, vary the initial pulse width or shape, and by optimizing the pulse shape one may obtain a more robust solution.

In particular, one can devise a different (constant-frequency) input pulse shape that




Figure 2-13: One way of exciting the system into a controlled stable solution: the input power is the sum of an exponential turn-on (the blue curve, P_1) and a Gaussian pulse with amplitude P_0 and width δT . The amplitude P_0 is altered to control which stable solution the system ends up in.

robustly excites the high-efficiency solution, insensitive to small changes in the initial conditions, by examining the power-bifurcation diagram in Fig. 2-11(left) in more detail. First, we observe that input powers $\gtrsim 1.45P_{\text{crit}}$ have only one stable solution, meaning that this stable solution is excited regardless of the initial conditions or the manner in which the input power is turned on. Then, if we slowly decrease the input power, the solution must “adiabatically” follow this stable solution in the bifurcation diagram until a power $\approx 0.95P_{\text{crit}}$ is reached, at which point that stable solution disappears. In fact, by inspection of Fig. 2-11(left), at that point there are *no* stable solutions, and solution jumps into a limit cycle. If the power is further decreased, a high-efficiency stable solution reappears and the system must drop into this steady state (being the only stable solution at that point). This process of gradually decreasing the power is depicted in Fig. 2-15(left), where the instantaneous “efficiency” is plotted as a function of input power, as the input power is slowly decreased. (The efficiency can exceed unity, because we are plotting instantaneous output vs. input power, and in the limit-cycle self-pulsing solution the output power is concentrated into pulses whose peak can naturally exceed the average input or output power.) Already, this is an attractive way to excite a high-efficiency ($> 90\%$) solution, because it is insensitive to the precise manner in which we change the power as long as it is changed slowly enough—this rate is determined by the lifetime of the cavity, and since this lifetime is likely to be sub-nanosecond in practice, it is not difficult to change the power “slowly” on that




Figure 2-14: (Left:) Steady-state efficiency at point C in Fig. 2-7 as a function of the transient input-pulse power P_0 from Fig. 2-13, showing how all three stable solutions can be excited by an appropriate input-pulse amplitude. (Right:) Same, but for an asymptotic input power $P_1 \approx 0.8P_{\text{crit}}$, for which the maximum efficiency is $\approx 90\%$ from Fig. 2-11(right), but is easier to excite.

timescale. However, we can do even better, once we attain this high-efficiency steady state, by then *increasing* the power adiabatically. As we increase the power, starting from the high-efficiency steady-state solution below the critical power, the system first enters limit-cycle solutions when the power becomes large enough that the stable solution disappears in Fig. 2-11(left). As we increase the power further, however, we observe that these limit cycles *always* converge adiabatically into the 100%-efficiency solution when $P \rightarrow P_{\text{crit}}$. This process is shown in Fig. 2-15(right). What is happening is actually a supercritical Hopf bifurcation at the two points where the upper branch changes between stable and unstable: this is a reversible transition between a stable solution and a limit cycle (initially small oscillations, growing larger and larger away from the transition). This is precisely what we observe in Fig. 2-15, in which the limit cycle amplitudes become smaller and smaller as the stable solutions on either side of the upper branch are approached, leading to the observed reversible transitions between the two. The important fact is that, in this way, by first decreasing and then increasing the power to P_{crit} , one always obtains the 100%-efficiency solution regardless of the precise details of how the power is varied (as long as it is “slow” on the timescale of the cavity lifetime).

The discussion thus far has revolved around idealized lossless systems—our main intention has been to examine the fundamental behaviors (steady-state solution and dynamics) of these systems rather than any specific experimental realization. However, as explained in




Figure 2-15: (Left:) Green line with arrows indicates instantaneous “efficiency” (harmonic output power / input power) as the input power is slowly decreased, starting at a power $\approx 1.7P_{\text{crit}}$. For comparison, Fig. 2-11(left) is superimposed as solid-red and dashed-blue lines. The solution “adiabatically” follows a steady state until the steady state becomes unstable, at which point it enters limit cycles, and then returns to a high-efficiency steady state, and finally goes drops to a low-efficiency steady-state if the power is further decreased. (Right:) Similar, but here the power is *increased* starting at the high-efficiency steady state solution for $P < P_{\text{crit}}$. In this case, it again enters limit cycles, but then it returns to a high-efficiency steady-state solution as the power is further increased, eventually reaching the 100%-efficiency stable solution. If the power is further increased, it drops discontinuously to the remaining lower-efficiency steady-state stable solution.

the subsequent section, we have performed preliminary calculations that include both linear losses (such as radiation or material absorption) and nonlinear two-photon absorption, and we find that these losses do not qualitatively change the observed phenomena. In fact, one still obtains multistability, limit cycles, and bifurcations, merely at a reduced peak efficiency depending on the strength of the losses.

2.2.8 Effect of Losses

In practice, a real device will have some additional losses, such as linear or nonlinear absorption and radiative scattering. Such losses will lower the peak conversion efficiency below 100%. As we show in this section, their quantitative effect depends on the ratio of the loss rate to the total loss rate $1/Q$. We also solve for the critical input power to achieve maximal conversion efficiency in the presence of losses.

For a $\chi^{(2)}$ medium with a linear loss rate $1/\tau_{k,e}$, we solve Eqs. (2.10–2.11) for $|s_{2-}|^2$ and enforce the condition for maximal conversion efficiency: $\frac{d}{dt}(|s_{2-}|^2/|s_{1+}|^2) = 0$. We thus

obtain the following optimal input power and conversion efficiency:

$$|s_{1+}|^2 = \frac{4\tau_{s,1}}{\omega_1^2 |\beta_1|^2 \tau_1^3 \tau_2} \quad (2.33)$$

$$\frac{|s_{2-}|^2}{|s_{1+}|^2} = \frac{\tau_1 \tau_2}{\tau_{s,1} \tau_{s,2}} \quad (2.34)$$

It immediately follows that for zero external losses, i.e. $\tau_k = \tau_{s,k}$, Eq. (2.34) gives 100% conversion and Eq. (2.33) reduces to Eq. (2.26). For small external losses $\tau_{s,k} \ll \tau_{e,k}$, the optimal efficiency is reduced by the ratio of the loss rates, to first order:

$$\frac{|s_{2-}|^2}{|s_{1+}|^2} \approx 1 - \left(\frac{\tau_{s,2}}{\tau_{e,2}} + \frac{\tau_{s,1}}{\tau_{e,1}} \right). \quad (2.35)$$

(A similar transmission reduction occurs in coupled-mode theory when any sort of loss is introduced into a resonant coupling process [203].)

The same analysis for $\chi^{(3)}$ yields the following critical input power and optimal efficiency:

$$|s_{1+}|^2 = \left[\frac{4\tau_{s,1}^2}{\omega_1^2 |\beta_1|^2 \tau_1^5 \tau_3} \right]^{1/2} \quad (2.36)$$

$$\frac{|s_{3-}|^2}{|s_{1+}|^2} = \frac{\tau_1 \tau_3}{\tau_{s,1} \tau_{s,3}} \quad (2.37)$$

where by comparison with Eq. (2.34), a first-order expansion for low-loss yields an expression of the same form as Eq. (2.35): the efficiency is reduced by the ratio of the loss rates, with τ_2 replaced by τ_3 .

A $\chi^{(3)}$ medium may also have a nonlinear “two-photon” absorption, corresponding to a complex-valued $\chi^{(3)}$, which gives an absorption coefficient proportional to the field intensity. This enters the coupled-mode equations as a small imaginary part added to α , even if we have set the real part of α to zero. (The corresponding effect on β is just a phase shift.) That yields a nonlinear (NL) $\tau_{e,k}$ of the following form, to lowest order in the loss:

$$\frac{1}{\tau_{s,1}^{\text{NL}}} \approx \omega_1 \text{Im} \left[\alpha_{11} \frac{\tau_{s,1}}{2} |s_{1+}|^2 + \alpha_{13} \frac{\tau_{s,3}^2 \tau_{s,1}^3}{8} \omega_3^2 |\beta_3|^2 |s_{1+}|^6 \right] \quad (2.38)$$

$$\frac{1}{\tau_{e,3}^{\text{NL}}} \approx \omega_3 \text{Im} \left[\alpha_{31} \frac{\tau_{s,1}}{2} |s_{1+}|^2 + \alpha_{33} \frac{\tau_{s,3}^2 \tau_{s,1}^3}{8} \omega_3^2 |\beta_3|^2 |s_{1+}|^6 \right]. \quad (2.39)$$

These loss rates can then be substituted in the expression for the losses above, in which case

one obtains the following optimal efficiency of third-harmonic generation, to lowest-order, not including linear losses:

$$\frac{|s_{3-}|^2}{|s_{1+}|^2} \approx 1 - \frac{\tau_{s,3}}{|\beta_1|} \sqrt{\frac{\tau_{s,3}}{\tau_{s,1}}} \operatorname{Im} \left[\frac{\alpha_{11} + 3\alpha_{13}}{\tau_{s,3}} + \frac{\alpha_{13} + 3\alpha_{33}}{\tau_{s,1}} \right] \quad (2.40)$$

(The linear and nonlinear losses can be combined by simply multiplying Eq. (2.35) and Eq. (2.40).) Thus, the nonlinear loss is proportional to the ratio $\operatorname{Im} \alpha / |\beta|$, which is proportional to $\operatorname{Im} \chi^{(3)} / |\chi^{(3)}|$.

2.2.9 Future Work

Future work on harmonic generation will involve designing specific doubly-resonant cavity geometries and more precise power predictions. Using our expressions for α and β , optimized cavities for harmonic generation can be designed using standard methods to compute the linear eigenmodes. In practice, experimentally achieving cavity modes with “exactly” harmonic frequencies, matched to within the fractional bandwidth $1/Q$, is a challenge and may require some external tuning mechanism. For example, one could use the nonlinearity itself for tuning, via external illumination of the cavity with an intense “tuning” beam at some other frequency. Understanding the experimental limitations for achieving frequency matching is particularly important for the $\chi^{(3)}$ case, where pre-correcting the frequency to compensate the nonlinear frequency shift (self-phase modulation) may require some care to ensure a stable solution. In particular, to obtain widely-spaced resonant modes ω and 3ω in a nanophotonic (wavelength-scale) context (as opposed to macroscopic Fabry-Perot cavities with mirrors), the most promising route seems to be a ring resonator of some sort [469], rather than a photonic crystal [241] (since photonic band gaps at widely separated frequencies are difficult to obtain in two or three dimensions). Although such a cavity will naturally support more than the two ω_1 and ω_3 modes, only two of the modes will be properly tuned to achieve the resonance condition for strong nonlinear coupling.

2.3 Difference Frequency Generation

The previous sections were mainly concerned in developing a theoretical framework based on coupled-mode theory for second- and third-harmonic generation in doubly resonant cavities, valid in the high-efficiency limit and not restricted to any specific geometry. As noted, this

type of self-consistent treatment is possible in nonlinear mixing where all relevant modes are resonant (e.g. fundamental and second- harmonic) because the resonant enhancement of the density of optical modes at each frequency causes the nonlinear interaction of these modes to be strongly favored over other processes [203,445]. Using CMT, we proved the existence of a critical power (that depends on the cavity parameters, e.g. frequencies, quality factors (Q), and mode-overlaps) at which 100% conversion is possible, and beyond which efficiency decreases. In addition, we showed that in the case of $\chi^{(3)}$ harmonic generation, the inclusion of self-phase modulation and cross- phase modulation allow for a rich diversity of dynamics, including limit cycles and multi-stable solutions.

Nonlinear $\chi^{(2)}$ difference-frequency generation (DFG), the subject of the next two sections, is yet another type of nonlinear frequency-conversion process in which two input sources, a pump beam at frequency ω_1 and a probe beam at frequency ω_2 interact via a $\chi^{(2)}$ nonlinearity to yield a third frequency $\omega_T = \omega_2 - \omega_1$ [51]. DFG processes have been considered recently for the generation of coherent radiation in the THz frequency range [9,17,42,194,227,367,547]. There has been some prior theoretical work on nonlinear frequency conversion, and especially DFG, in resonant systems [2,9,16,33,42,126,194,324,341,393,398,402,445,563]; however, as in the harmonic generation case, these processes are generally treated in the undepleted pump limit, in which reverse processes are neglected, and which lead to incorrect results in the limit of high conversion efficiency. In this section, we extend the framework presented in the previous section [198,445,474] to the case of difference-frequency generation (DFG) caused by a second-order $\chi^{(2)}$ nonlinear polarization in a resonant photonic structure. In contrast to previous work on DFG, we consider the case where both the input pump (highest frequency, ω_1) and idler waves ω_2 are resonant, as well as the generated signal at the difference frequency ω_T . We show that CW conversion from pump to signal with quantum-limited efficiency (complete depletion of the pump) is possible in such a system for any power at the pump frequency, provided there is a critical input power $P_{\text{crit}} = |s_{2+}^{\text{crit}}|^2$ at the idler frequency ω_2 . Furthermore, as in the harmonic-generation case, the steady-state solutions to the coupled wave equations and their corresponding conversion efficiencies are shown to depend only on universal dimensionless parameters, indicating that their structure is qualitatively the same in all geometries.

The general framework of the temporal CMT model for DFG is shown in Fig. 2-16. Our cavity has three resonant modes at frequencies, ω_1 (pump), ω_2 (idler), and ω_T (signal) which




Figure 2-16: Schematic of general scheme for $\chi^{(2)}$ difference-frequency generation, including dynamical variables for coupled-mode equations: two input (output) channels (with incoming (outgoing) field amplitudes s_{\pm}) are coupled to a resonant cavity with three modes at frequencies ω_1 and ω_2 and $\omega_T = \omega_2 - \omega_1$ (with corresponding amplitudes a_1 , a_2 , and a_T). The two resonant modes are nonlinearly coupled by a Pockels ($\chi^{(2)}$) nonlinearity. All three modes leak energy at a rate determined by their corresponding quality factors Q .

satisfy $\omega_T = \omega_1 - \omega_2$. As in harmonic generation, energy is coupled in/out of each cavity mode through a designated input/output port (e.g. waveguide). The propagating modes in the input/output ports are similarly represented. The relevant incoming and outgoing wave amplitudes are again represented by the variables $s_{k+}(t)$ and $s_{k-}(t)$, respectively, where $|s_{k\pm}(t)|^2$ is the power of the propagating mode. By extending the previous formalism to non-degenerate three-wave mixing, where we assume there is no input at the signal frequency ($s_{T+} = 0$), the coupled-mode equations can be written as:

$$\dot{a}_1 = (i\omega_1 - 1/\tau_1) a_1 - i\omega_1 \beta_1 a_2 a_T + \sqrt{\frac{2}{\tau_{1,e}}} s_{1+} \quad (2.41)$$

$$\dot{a}_2 = (i\omega_2 - 1/\tau_2) a_2 - i\omega_2 \beta_2 a_1 a_T^* + \sqrt{\frac{2}{\tau_{1,e}}} s_{2+} \quad (2.42)$$

$$\dot{a}_T = (i\omega_T - 1/\tau_T) a_T - i\omega_2 \beta_T a_1 a_2^*, \quad (2.43)$$

where $|a_k|^2$ again yields the energy in the cavity, the τ_k are the corresponding total cavity lifetimes (where $\tau_{k,e}$ are the coupling lifetimes of the modes), and where $|s_{k+}|^2$ correspond to the input power for the k th mode.

The resulting coupling coefficients are:

$$\beta_{1,2} = \frac{1}{4} \frac{\int d^3 \mathbf{x} \sum_{ijk} \chi_{ijk}^{(2)} E_{1,i}^* (E_{T,j} E_{2,j} + E_{2,j} E_{T,k})}{\sqrt{\int d^3 \mathbf{x} \epsilon |\mathbf{E}_1|^2} \sqrt{\int d^3 \mathbf{x} \epsilon |\mathbf{E}_2|^2} \sqrt{\int d^3 \mathbf{x} \epsilon |\mathbf{E}_T|^2}} \quad (2.44)$$

$$\beta_T = \beta_{1,2}^* \quad (2.45)$$

(We do not derive these coupling coefficients here, since the derivation is simple and has been laid out in the previous sections). Equation (2.45) shows that $\beta \equiv \beta_1 = \beta_2 = \beta_T^*$ depends

strongly on the overlap between the various modes: in the limit where $\chi^{(2)}$ is constant, and ignoring polarization effects, then the relevant quantity of interest is $\sim E_1^* E_T E_2$. From this expression, it becomes obvious why it is generally difficult to achieve large values of β . Namely, because E_1 and E_2 will tend to oscillate in space much more rapidly than E_T , their contribution to the integral will yield a vanishingly small coefficient.

The DFG efficiency can be assessed again by analyzing the steady state solutions ($da_k(t)/dt = ds_{k\pm}(t)/dt = 0$) of Eqs. (2.41–2.43). In contrast with degenerate nonlinear frequency conversion (e.g. harmonic generation), the non-degenerate case does not allow for 100% conversion of power from ω_1 to ω_T . This arises because the destruction of a photon at ω_1 and the creation of a photon at ω_T is always accompanied by the creation of a second photon at ω_2 . This fundamental constraint is described by the Manley-Rowe relations [51]. Considering also losses ($\tau_k \neq \tau_{k,e}$), the quantum limit of conversion efficiency is given by:

$$\frac{|s_{T-}|^2}{|s_{1+}|^2} \leq \frac{\omega_T \tau_{1,e} \tau_{T,e}}{\omega_1 \tau_1 \tau_T} \quad (2.46)$$

This condition can be derived directly from Eqs. (2.41–2.43) in the steady-state, and is omitted here. However, the fact that it is possible to derive constraints imposed by quantum-mechanical selection rules from a purely classical equation of motion is quite extraordinary, and represents yet another triumph of the coupled-mode theory formalism.

From Eq. (2.46), it is clear that the conversion efficiency is maximized by maximizing the ratios, $\tau_{1,e}/\tau_1$ and $\tau_{T,e}/\tau_T$. This implies that over-coupling the cavity modes to the input/output waveguide is essential to achieving high conversion efficiency. In other words, it is more important that the Q-factors be limited by leakage to the desired output port than that they have a higher value but lack a well-defined output channel. This is the same condition for low-loss operation that we obtained in the previous sections on harmonic generation. However, notice that the loss ratio in the idler, ω_2 mode, does not appear in Eq. (2.46), i.e. does not affect the overall efficiency. This is due to the fact that the input wave at ω_2 is not down-converted, but rather amplified; therefore, losses can be compensated for by pumping in more power at this frequency. One can find a critical relationship between the input powers, $|s_{1+}|^2$ and $|s_{2+}|^2$ that allows for maximum conversion efficiency (complete depletion of the pump) by imposing either of the following equivalent constraints: $s_{1-} = 0$

or an equality in Eq. (2.46), yielding:

$$|s_{2+}|^2 = \frac{\omega_2 \left| 1 - \frac{4}{\omega_1} |\beta|^2 Q_1 Q_2 Q_T \frac{\tau_{1,e}}{\tau_1} |s_{1+}|^2 \right|^2}{16 |\beta|^2 Q_1 Q_2 Q_T \frac{\tau_{2,e}}{\tau_2}} \quad (2.47)$$

Notice that the power coupled into the cavity at the idler frequency, $|s_{2+}|^2$, depletes the signal (ω_T) when nonlinearly converted (via sum-frequency generation), and produces power at the pump ω_1 . Therefore, in the case in which the pump is completely depleted ($s_{1-} = 0$), no net power from the idler frequency is converted and $|s_{2-}|^2 = |s_{2+}|^2 + \tau_{1,e} \tau_{2,e} \omega_2 |s_{1+}|^2 / \tau_1 \tau_2 \omega_1$. The conversion of the total input power, $|s_{1+}|^2 + |s_{2+}|^2$, is thus maximized in the limit $|s_{2+}|^2 \rightarrow 0$, $|s_{1+}|^2 = \omega_1 / (4 |\beta|^2 Q_1 Q_2 Q_T \tau_{1,e} / \tau_1)$. Note, however, that Eqs. (2.41–2.43) require a non-zero value (this can be arbitrarily small) of $|s_{2+}|^2$ for non-zero conversion efficiency. Note also that the power at which total conversion is optimal decreases with increasing Q -factors and vice-versa. This means that lower Q -factors are desirable for high power applications. As an example, consider the case where $\lambda_1 = 1\mu\text{m}$, $\lambda_2 \approx 1.1\mu\text{m}$, $\lambda_T \approx 10\mu\text{m}$, in a cavity composed of a typical III-V semiconductor such as GaAs. Taking the magnitude of the second-order effective nonlinear susceptibility to be $\text{deff} \approx 274\text{pm/V}$ [494], and assuming a constant field in each cavity with an optimal overlap, the critical power for optimal total power conversion (where $|s_{2+}|^2 \rightarrow 0$) would be reached at $|s_{1+}|^2 = 4.4\text{mW}$ for $Q_1 = Q_2 = Q_T = 104$. The actual implementation of a triply-resonant system characterized by high Q/V ratios for the three frequencies involved in the nonlinear frequency mixing process and high modal overlap is a challenging problem and is out of the scope of this thesis.

Equation (2.47) also allows for quantum-limited conversion for an arbitrarily small value of $|s_{1+}|^2$ as $|s_{2+}|^2 \rightarrow \omega_2 / (16 |\beta|^2 Q_1 Q_2 Q_T \tau_{2,e} / \tau_2)$. This is a unique property of the non-degeneracy of DFG, and could prove useful for low power applications [346, 469]. A similar result was found in the quantum limit for doubly resonant DFG, specifically that a single photon could be converted with high fidelity from ω_1 to ω_T , given an appropriate pump power at ω_2 . Note that in both limits the optimal input power is proportional to $1/Q_1 Q_2 Q_T \propto 1/Q^3$. A similar result was found in the degenerate case of second-harmonic generation, where the critical power was shown to scale as $1/Q_1^2 Q_2$. A related enhancement factor was also reported earlier in the undepleted-pump limit [309]. In the specific case of THz generation through DFG of telecom-band modes ($\omega_1 \approx \omega_2$), assuming the two

telecom-band modes have the same mode volume ($V_1 \approx V_2$), then under the conditions of best possible overlap, the overlap would scale as $\beta \sim 1/V_T^{1/2}$. This means that in triply resonant THz generation, the mode volumes of the telecom-band cavities do not affect the strength of the nonlinear coupling as long as they match each other (i.e. $V_1 \approx V_2$).

Finally, note that we have not considered the stability of the critical DFG solutions, which we require in order to verify whether the solution can in fact be excited, as discussed in the previous sections. However, we defer the reader to Ref. [70], upon which the discussion in this section is based.

2.3.1 Terahertz Generation

The generation of terahertz (THz) or submillimeter radiation has been a subject of considerable interest in recent decades [28, 530]. An important criterion for applications based on THz radiation is that the THz sources operate at sufficiently high power levels, near room temperature and using compact turn-key sources [28, 30, 547, 560, 583]. These constraints have presented a theoretical and experimental challenge that has been difficult to overcome with conventional approaches [272, 295, 367, 462]. The traditional approach to generate THz radiation involves some type of electrical pumping, for instance FIR/quantum cascade lasers [418, 472, 535, 550, 561], while more recent methods have explored a wider range of options, including mechanical shocks [433] and nonlinear optical processes [9, 204, 341]. Nanophotonics and in particular nonlinear DFG offer a promising route to the realization of high power, efficient THz source. Recently, different approaches to resonant enhancement of difference-frequency nonlinear coupling processes, including those involving a final frequency in the THz regime, have been proposed. However, these previous approaches were unable to yield conversion efficiencies close to the quantum-limit in a compact device operating at room temperature [9, 17, 42, 227, 367, 547]. In this section, we describe a scheme, described in the work of Ref. [56], that enables enhancement of THz power generation via second-order nonlinear frequency down-conversion by up to three orders of magnitude compared to conventional nonresonant approaches. By using a combination of accurate numerical simulations and the coupled-mode equations in the previous section, we show how the unique properties of photonic microresonators to confine light in subwavelength volumes for many optical periods enable the implementation of highly-efficient compact on-chip continuous-wave THz sources operating at room temperature and pumped by sub-W

pulses, which could contribute to the practical realization of efficient THz sources.

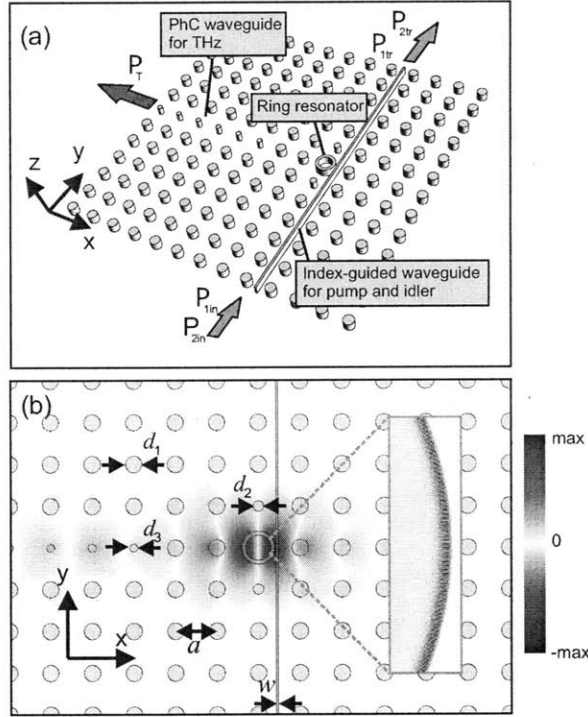


Figure 2-17: (a) Schematic illustration of triply-resonant monolithic (GaAs) geometry. $P_{1,in} = |s_{1+}|^2$ and $P_{2,in} = |s_{2+}|^2$ denote the input power of light at pump ω_1 and idler ω_2 frequencies, respectively, while $P_T = |s_{T-}|^2$ stands for the THz output power. The transmitted power at the pump and idler frequencies are denoted by $P_{1,tr}$ and $P_{2,tr}$, respectively. (b) Main panel: electric field profile E_z corresponding to the resonant mode appearing at 1 THz in the dielectric structure of (a). The value of the geometrical parameters are: $a = 102\mu\text{m}$, $d_1 = 40.8\mu\text{m}$, $d_2 = 25.1\mu\text{m}$, $d_3 = 18.8\mu\text{m}$, and $w = 0.8\mu\text{m}$. (Inset:) Enlarged view of the x -polarized electric field profile E_x , corresponding to a whispery-gallery mode with $m = 572$ circulating inside the dielectric ring, shown in the main figure. The internal/external radii of the ring resonator are $30.5\mu\text{m}$ and $40.1\mu\text{m}$, respectively. Shaded/white areas represent GaAs/air regions.

To date, most resonant THz DFG schemes have been designed and analyzed in the singly- and doubly-resonant case, with cavities designed at the $\omega_{1,2}$ frequencies [9, 17, 42, 227, 367, 547]. The motivation for extending this work to include triply-resonant cavities comes from the additional enhancement factor $\sim Q_T/V_T$ arising from confinement of the THz mode. The difficulty with a triply resonant scheme comes from the fact that introducing a cavity at ω_T requires a multiscale design for the geometry which can yield a number of significant experimental difficulties: in particular, DFG THz generation often involves

input GHz (μm) light being converted to THz (millimeter) light. In addition to posing experimental challenges, in many cases, these multiscale cavities will often yield low conversion efficiency due to the relatively small coupling coefficient between the modes, i.e. there is a tradeoff between the Q/V enhancement obtained for each individual mode due to confinement, and the coupling coefficient β .

To explore the extent to which one can exploit resonant DFG to solve the current lack of room-temperature efficient THz sources, we investigate THz generation in a specific (realistic), shown in the schematic of Fig. 2-17. In this configuration, the THz field is confined by means of a THz-scale photonic crystal cavity, whereas the pump and idler fields are resonantly confined using two whispering gallery modes (WGM) supported by a dielectric ring resonator. More specifically, the power carried by two near-infrared (IR) beams of wavelengths λ_1 and λ_2 correspond to the pump and idler beams, respectively, with corresponding input power $P_{1,\text{in}}$ and $P_{2,\text{in}}$, respectively. Both modes are coupled, by means of an index-guided waveguide, to two high-order WGM, characterized by angular momenta m_1 and m_2 , respectively. The ring resonator also acts as a dipole-like defect at λ_T , when embedded in an otherwise perfectly periodic THz-wavelength scale PhC formed by a square lattice of dielectric rods (see the corresponding electric field profile in Fig. 2-17). Thus, the $\chi^{(2)}$ nonlinear frequency down-conversion interaction that takes place between the two NIR WGMs circulating inside the ring resonator yields a current distribution that radiates inside the PhC cavity at the frequency difference $\omega_T = \omega_2\omega_1$. In order to extract the THz output power ($P_T = |s_{T-}|^2$), we introduce into the system a PhC waveguide created by reducing the radius of a row of rods [see Fig. 2-17(a-b)]. In addition, due to the symmetry of the cavity, we are forced to break the degeneracy of the resulting x - and y -oriented dipole defect modes: this is achieved by reducing the radius of two of the nearest-neighbor rods of the ring resonator with respect to the radius of the other PhC rods. The entire configuration permits a large Q_T/V factor, along with high- Q resonant confinement for the pump and idler frequencies. Figure 2-17 shows the structure that results from optimizing the geometrical parameters of the system for efficient THz generation of light with $\omega_T \sim 1\text{THz}$, along with the corresponding electric field profiles, as obtained from repeated finite-difference time-domain (FDTD) simulations. In these calculations we have assumed pump and idler beams of wavelength $\lambda_1 = 1550\text{nm}$ and $\lambda_2 = 1542\text{nm}$, respectively; the structure is implemented in Gallium Arsenide (GaAs), for which the relevant component of the nonlinear susceptibility

tensor is $d_{14} = 274\text{pm/V}$ [272]. At these wavelengths, and for this geometry, GaAs yields a nonlinear coupling coefficient β that is maximized if the structure is designed to support two WGM with $m_1 = 572$ and $m_2 = 575$ at λ_1 and λ_2 respectively [see inset of Fig. 2-17], along with a dipole defect mode in the THz-scale PhC. We emphasize that in conventional phase-matching schemes, the overall efficiency of a difference-frequency generation process relies entirely on finding a suitable nonlinear material whose dispersion relation allows one to simultaneously fulfill both frequency-matching and phase-matching (selection rules) conditions.

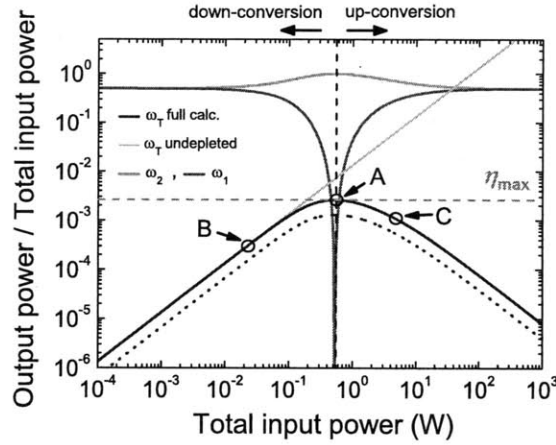


Figure 2-18: Total efficiency (blue) of THz light $|s_{T-}|^2/(|s_{1+}|^2 + |s_{2+}|^2)$ emitted by the system, for $\omega_T = 1$ THz. The results also show the depletion efficiency $|s_{k-}|^2/(|s_{1+}|^2 + |s_{2+}|^2)$ of the pump ($k = 1$, green line) and idler ($k = 2$, red line) photons. The horizontal dashed line displays the maximum possible THz (quantum-limited) conversion efficiency, determined by the Manley-Rowe relations. The dotted line displays the effect of realistic absorption losses on the conversion efficiency.

We now apply the coupled-mode theory formalism of our previous sections to the system of Fig. 2-17. Figure 2-17 summarizes the main results obtained in the CW regime. These calculations assume equivalent input powers $|s_{1+}|^2 = |s_{2+}|^2$, and quality factors $Q_1 = Q_2 = 3.5 \times 10^5$, $Q_T = 103$. These values of Q are reasonable, and desirable, considering the absorption coefficient of GaAs at 1 THz (the linear absorption rate of GaAs corresponds to a Q factor 1.5×10^3) [204] and moreover, have been recently obtained experimentally in similar geometrical configurations for both the ring resonator and the PhC cavity [7, 382]. As shown in Fig. 2-18, for total input power $|s_{1+}|^2 + |s_{2+}|^2 \geq 70\text{mW}$, the conversion efficiency (defined here as the ratio $|s_{T-}|^2/(|s_{1+}|^2 + |s_{2+}|^2)$, as opposed to the quantum-limited efficiency) begins departing from the conversion efficiency predicted by

the undepleted approximation, eventually reaching the maximum value predicted by the ManleyRowe relations [51]. It is clear from inspection of Fig. 2-18 that the critical power in this system is reached at $P_{\text{in}} = 0.54\text{W}$, at which point the pump power that is coupled to the ring resonator is completely down-converted to power at the THz and idler frequencies. This represents a dramatic reduction (about three orders of magnitude) in the total required input power for THz generation with respect to the most efficient schemes in nonlinear crystals reported up to date [547, 583]. Furthermore, we emphasize that, in addition to powerful lasers, current efficient schemes for THz generation require intricate phase-matching considerations, whereas in the system introduced here the maximum theoretically possible efficiency can be achieved in an integrated structure having a total area $\sim 1\text{mm}^2$ using less than 1 W input powers that are readily accessible with compact turn-key sources. As noted above, since $|s_{1+}^{\text{crit}}|^2 \sim 1/Q_1Q_2Q_T$, the critical power can be adjusted by varying the total bandwidth product $Q_1Q_2Q_T$.

2.4 Three-Wave Mixing

Three-wave mixing (TWM) is yet another type of nonlinear process, often studied in $\chi^{(2)}$ media, in which input light composed of frequencies ω_0 and $\omega_0 + \Delta$ gets converted to light at frequency $\omega_0 - \Delta$. (The special limit $\Delta \rightarrow -\omega_0$ corresponds to degenerate TWM, or second-harmonic generation.) Previous work on TWM focused primarily in singly-resonant or large-etalon structures with stringent phase-matching requirements [41, 236, 394, 496, 582, 597], as were most of the other nonlinear processes studied above. In this section, we write down the basic coupled-mode equations of this three-wave mixing in a triply-resonant cavity system, and furthermore solve for the corresponding steady-state behavior.

Consider a triply resonant $\chi^{(3)}$ cavity with resonances at $\omega_1, \omega_2 = \omega_1 - \Delta$ and $\omega_3 = \omega_1 + \Delta$. Consider also an up-conversion process in which one is pumping energy into the cavity at ω_0 (the pump signal) and ω_2 (the idler signal), to generate light at the higher frequency ω_3 ⁴. Following the prescription of Sec. 2.2.1, one obtains the following coupled-

⁴Note that in order to consider the alternative process, corresponding to a down-conversion process in which ω_3 and ω_2 represent the idler and generated signals, one need not consider a separate set of CMEs. The down-conversion process can be analyzed simply by setting $\Delta \rightarrow -\Delta$ in Eqs. (2.48–2.50) and Eq. (2.57).

mode equations:

$$\frac{da_1}{dt} = \left(i\omega_1(1 + \alpha_{11}|a_1|^2 + \alpha_{21}|a_2|^2 + \alpha_{31}|a_3|^2) - \frac{1}{\tau_1} \right) a_1 + 2i\omega_1\beta^* a_1^* a_2 a_3 + \sqrt{\frac{2}{\tau_{s,1}}} s_{1+} \quad (2.48)$$

$$\frac{da_2}{dt} = \left(i\omega_2(1 + \alpha_{12}|a_1|^2 + \alpha_{22}|a_2|^2 + \alpha_{32}|a_3|^2) - \frac{1}{\tau_2} \right) a_2 + i\omega_2\beta a_1^2 a_3^* + \sqrt{\frac{2}{\tau_{s,2}}} s_{2+} \quad (2.49)$$

$$\frac{da_3}{dt} = \left(i\omega_3(1 + \alpha_{13}|a_1|^2 + \alpha_{23}|a_2|^2 + \alpha_{33}|a_3|^2) - \frac{1}{\tau_3} \right) a_3 + i\omega_3\beta a_1^2 a_2^*. \quad (2.50)$$

Furthermore, by following the simple perturbation-theory analysis described in Sec. 2.2.2, one also obtains the corresponding coupling coefficients β_i and α_{ij} :

$$\beta_1 = \frac{1}{8N} \int d^3\mathbf{x}\varepsilon\chi^{(3)}\mathbb{E}^* \quad (2.51)$$

$$\beta_2 = \beta_3 = \frac{1}{16N} \int d^3\mathbf{x}\varepsilon\chi^{(3)}\mathbb{E} \quad (2.52)$$

$$N = \left[\int d^3\mathbf{x}\varepsilon|\mathbf{E}_1|^2 \right] \left[\int d^3\mathbf{x}\varepsilon|\mathbf{E}_2|^2 \right]^{1/2} \left[\int d^3\mathbf{x}\varepsilon|\mathbf{E}_3|^2 \right]^{1/2} \quad (2.53)$$

$$\mathbb{E} = (\mathbf{E}_1 \cdot \mathbf{E}_1)(\mathbf{E}_2 \cdot \mathbf{E}_3) + 2(\mathbf{E}_1 \cdot \mathbf{E}_2)(\mathbf{E}_1 \cdot \mathbf{E}_3) \quad (2.54)$$

$$\alpha_{ii} = \frac{1}{8} \frac{\int d^3\mathbf{x}\varepsilon\chi^{(3)} [|\mathbf{E}_i \cdot \mathbf{E}_i^*|^2 + |\mathbf{E}_i \cdot \mathbf{E}_i|^2]}{\left[\int d^3\mathbf{x}\varepsilon|\mathbf{E}_i|^2 \right]^2} \quad (2.55)$$

$$\alpha_{ij} = \alpha_{ji} = \frac{1}{8} \frac{\int d^3\mathbf{x}\varepsilon\chi^{(3)} [|\mathbf{E}_i \cdot \mathbf{E}_j|^2 + |\mathbf{E}_i \cdot \mathbf{E}_j^*|^2 + |\mathbf{E}_i|^2|\mathbf{E}_j|^2]}{\left[\int d^3\mathbf{x}\varepsilon|\mathbf{E}_i|^2 \right] \left[\int d^3\mathbf{x}\varepsilon|\mathbf{E}_j|^2 \right]}, \quad i \neq j, \quad (2.56)$$

which again are seen to satisfy the corresponding energy-conservation constraint, here given by $\omega_1\beta_1^* = \omega_2\beta_2 + \omega_3\beta_3$.

Assuming again, for simplicity, that $\alpha_{ij} \ll \beta$, the maximum quantum-limited conversion efficiency $\eta_3 = |s_{3-}|^2/(|s_{1+}|^2 + |s_{2+}|^2)$ can be obtained by setting $s_{1,-} = 0$, and as in the DFG example of the preceding section, taking the limit $s_{2,+} \rightarrow 0$. (Since TWM is a non-degenerate nonlinear process, the Manley-Rowe relations demand an efficiency cut-off determined by the possible rate of photon up- and down-conversion.) The resulting

maximum quantum-limited efficiency is given by:

$$\eta_3 = \frac{\omega_3}{2\omega_1} = \frac{1}{2} \left(1 + \frac{\Delta}{\omega_1} \right), \quad (2.57)$$

corresponding to a critical power $P_{\text{crit}} = |s_{1,+}^{\text{crit}}|^2$ determined by the solution of the following quadratic equation:

$$\left(1 - \frac{\tau_1^2 \mu_2 \mu_3^*}{4} |s_{1,+}^{\text{crit}}|^4 \right)^2 = \tau_1 \tau_2 \mu_1 \mu_3 |s_{1,+}^{\text{crit}}|^2 |s_{2,+}|^2; \quad (2.58)$$

Solving this equation (for the limit as $s_{2,+} \rightarrow 0$) numerically yields two solutions: a stable and an unstable solution. For this (desired) stable solution, maximal efficiency is obtained at a critical power:

$$P_{\text{crit}} = \frac{2}{\tau_1 |\beta| \sqrt{\tau_2 \tau_3 \omega_2 \omega_3}}. \quad (2.59)$$

Chapter 3

Quantum Nonlinear Nanophotonics

“It is impossible to travel faster than the speed of light, and certainly not desirable, as one’s hat keeps blowing off.” *Woody Allen*

3.1 Overview

In the previous chapter, we explored the possibility of enhancing nonlinear interactions by tailoring the geometry of various devices (modifying the photon density of states). We argued that microresonators, such as PhC microcavities, offer unique fundamental ways of enhancing a variety of nonlinear optical processes, stemming primarily from their unique ability to confine light over long lifetimes Q and small modal volumes V . This places many previously known nonlinear optical processes in a new operating regime. For example, this enhancement improves the performance of nonlinear optical devices to such an extent that their corresponding operation powers and switching times are suitable for their implementation in realistic ultrafast integrated optical devices [57]. In this chapter, based on work performed in collaboration with Peter Bermel described in Ref. [36,37], we explain how one may combine these geometric-based approaches with novel materials to further enhance nonlinear interactions, in some cases allowing us to even reach the single-photon (quantum) operation regime.

The chapter is composed of two major sections. Section 3.2 demonstrates switching of a single signal photon by a single gating photon of a different frequency, via a type of cross-phase-modulation (XPM). This XPM effect is mediated by materials exhibiting electromagnetically induced transparency (EIT), which are embedded in a photonic crystal

cavity coupled to a waveguide. An analytical model based on waveguide-cavity quantum electrodynamics, a quantized version of the coupled-mode theory of chapter 2, is constructed for our system: a PhC waveguide and a PhC microcavity containing a four-level EIT atom. The steady-state behavior of the system is solved exactly and analyzed using experimentally accessible parameters. Our results confirm that the strong coupling regime is required for lossless two-photon quantum entanglement.

Section 3.3 presents an approach to enhancing nonlinear interactions that focuses on strengthening (tailoring) the nonlinear optical susceptibility of typical nonlinear materials by designing geometries at the scale of the material resonances. In particular, it is shown that one can exploit the Purcell effect to enhance the Kerr susceptibility of certain materials. While this is a general physical principle that applies to a wide variety of nonlinearities, we specifically investigate the Kerr effect arising from the interaction of light with certain material systems. This idea is illustrated by explicit FDTD calculations in a system composed of a discrete system (e.g. atom or a quantum dot) embedded in a PhC cavity, wherein it is shown that by manipulating the Purcell effect for frequencies close to the atomic resonance one can substantially influence the resulting Kerr nonlinearity for light of all (even highly detuned) frequencies. Enhancements of the Kerr coefficient of one or two orders of magnitude could be achieved in realistic homogeneous systems (semiconductors). As pointed out in Sec. 3.4, this last topic lies at the intersection of classical and quantum electromagnetism, and finds its origins in the fluctuation-induced interactions of the electromagnetic field, the subject of the remainder of this thesis.

3.2 Single-Photon Switching

Several emerging technologies, such as integrated all-optical signal processing and all-optical quantum information processing, require strong and rapid interactions between two distinct optical signals [31, 176, 314, 378]. Achieving this goal is a fundamental challenge because it requires a unique combination of large nonlinearities and low losses, i.e. strong and rapid nonlinear photon–photon interactions. The weak nonlinearities found in conventional media mean that large powers are required for switching. On the classical end, experimental and theoretical studies demonstrate that bistable switching and many other strong nonlinear effects can be achieved by placing a Kerr nonlinear material inside a photonic crystal cav-

ity [80, 386, 502]. Unfortunately, although promising for many classical applications, most of these approaches are not scalable to the quantum (single-photon) regime, due primarily to high losses and low-power limitations. In this section, we explore the possibility of exploiting electromagnetically-induced transparency (EIT), previously explored in the context of classical nonlinear nanophotonics [ref], to enable efficient low-power (single-photon) switching.

Recent work showed that nonlinearities up to 12 orders of magnitude larger than those observed in common materials [200] with low losses can be achieved using EIT materials [14, 44, 200]. One can then envision inducing strong interactions between two very weak signals of different frequencies by placing a 4-level EIT atom in a high- Q cavity, so that a very small signal at a specific atomic transition frequency could shift another resonant frequency of the system by a measurable amount [505]. This approach differs from several optical switching schemes for small numbers of photons that have previously been discussed in the literature. One of the pioneering papers in this area used a single three-level atom with a V -level structure in an optical cavity to induce a cross-phase modulation of 16° between two photons [534]. EIT offers even further opportunities in terms of larger nonlinearities and greater tunability, which has directed much subsequent work in this direction. EIT materials have been predicted to cause a photon blockade effect, where the state of a cavity can be switched by the self-phase modulation of a single photon [226, 432, 556] or several photons [174, 187]. This effect has recently been observed experimentally [43]. Ref. [29] predicts that ensembles of EIT atoms can be modulated to create quantum entangled states for a small number of photons. An alternative method is discussed in Ref. [242], whereby a laser beam can control the relative populations of a two-state system embedded in a PhC, which switches its transmission properties at low power levels.

Ref. [505] semi-classically demonstrates the strong interaction of very low intensity fields which can be mediated by EIT materials. This work extends that idea to the quantum regime, by writing down the waveguide-cavity QED Hamiltonian for a system consisting of one or a few 4-level EIT atoms strongly coupled to a photonic crystal (PhC) cavity mode, which in turn is coupled to a PhC waveguide, and solving it exactly. Furthermore, an approach to calculating the relevant parameters from first principles is demonstrated. It should be experimentally feasible, with EIT having already been demonstrated in a Pr doped Y_2SiO_5 crystal [193, 536]. Note that compared to EIT systems such as Na BEC's




Figure 3-1: Schematic illustration of the system investigated. A waveguide is coupled to a cavity with an EIT atom at its center. In the upper left-hand corner, an FDTD simulation that can be used to calculate the model parameters is shown.

displaying narrow bandwidths (e.g., 2 MHz [200]), switching can occur over much larger bandwidths even for single photon power levels (e.g., 2 GHz, using the parameters from Ref. [586]) because the PhC cavity compensates for weaker nonlinearities, as demonstrated here. Furthermore, this approach utilizes PhC's, which offer confinement of light to high quality factor microcavities with low modal volumes, which facilitates strong coupling between light and matter. The emergence of new phenomena associated with the quantization of the probe and gate fields (e.g., Rabi-splitting) is discussed. Finally, it is shown that switching behavior can be achieved with single probe and gate photons, and the physical parameters needed to achieve such operations are calculated.

Consider the following design, illustrated in Fig. 3-1. There is a cavity which supports two resonant modes, one with a resonant frequency ω_{RES} and the other with a control frequency ω_{CON} , enclosing a single four-level EIT atom with coupling strengths g_{ij} and atomic transition frequencies ω_{ij} , where i and j refer to the initial and final atomic states, respectively. The EIT dark state is created by adding a classical coupling field to the cavity with frequency ω_{23} and Rabi frequency $2\Omega_c$ – all other quantities are treated quantum mechanically. In general, any number of coupling schemes between the cavity and one or more waveguides could be utilized. However, in this work, the ω_{RES} cavity mode is side-coupled to an adjacent single-mode waveguide with a radiative linewidth $\Gamma_w = \omega_{RES}/2Q_w = V_w^2/v_g$, where Q_w is the quality factor of the ω_{RES} cavity mode, V_w is the coupling strength and v_g is the group velocity in the waveguide – its dispersion relation $\omega(k)$ is assumed to be approximately linear near the ω_{RES} resonance. For relatively strong cavity-waveguide couplings, radiative couplings out of the system are much smaller and may be neglected. Also, the ω_{CON} resonance is designed to have a much smaller decay rate $\Gamma_{CON} = \omega_{CON}/2Q_{CON}$.

This can be achieved by starting with two dipole modes, one with an even symmetry coupled strongly to the waveguide, and one with an odd symmetry exactly decoupled from the waveguide. A slight shift in the cavity position can then create a slight coupling that nonetheless creates a substantial disparity in quality factors, i.e., $Q_{CON} \gg Q_w$ (see, e.g., Refs. [240, 251]). Alternatively, one could use two cavities to create even and odd modes with substantially different quality factors [519]. In the absence of an atom, this design produces a Lorentzian lineshape for the *reflection* (because of the side-coupling), centered around ω_{RES} [201]. A photonic crystal implementation of this is shown in the upper left-hand corner of Fig. 3-1—a triangular lattice of air holes in silicon with radius $0.48a$ that has a complete 2D photonic bandgap. A similar geometry has been used for quantum dots in photonic crystal microcavities, as in Ref. [586]. That experimental system exhibits a critical photon number $m_0 = \Gamma_3^2/2g^2 = 0.55$ and critical atom number $N_0 = 2\Gamma_w\Gamma_3/g^2 = 4.2$. Ideally both of these numbers would be less than one for quantum information processing [323]; it should be possible to achieve this goal with improvements in Q or modal volume V_{mode} , or by placing several atomic or quantum dot systems in the same microcavity. Note that it could also be possible to achieve similar behavior with other physical systems, such as high-finesse Fabry-Perot optical microcavities [212], or ultrahigh-Q toroidal microresonators [510].

Combining Ref. [556]’s Hamiltonian for an EIT atom in a cavity and Ref. [486]’s Hamiltonian for a waveguide interacting with a cavity yields:

$$\begin{aligned}
H/\hbar = & \sum_k \omega_k a_k^\dagger a_k + \omega_{RES} a^\dagger a + \omega_{CON} b^\dagger b + \sum_k V_w (a_k^\dagger + a_k)(a^\dagger + a) + \omega_{21} \sigma_{22} \\
& + (\omega_{13} - i\Gamma_3) \sigma_{33} + (\omega_{14} - i\Gamma_4) \sigma_{44} + \Omega_c (\sigma_{32} + \sigma_{23}) \cos(\omega_{23}t) \\
& + g_{13} (a^\dagger \sigma_{13} + a \sigma_{31}) + g_{24} (b^\dagger \sigma_{24} + b \sigma_{42})
\end{aligned} \tag{3.1}$$

where a_k are the annihilation operators for waveguide states of wavevector k and frequency ω_k , a and b are the annihilation operators for cavity photon states of frequencies ω_{RES} and ω_{CON} , respectively (which are considered here to be singly occupied), σ_{ij} are the projection operators that take the atomic state from j to i , Γ_3 is the nonradiative decay rate of the third level, Γ_4 is the nonradiative decay rate of the fourth level, $\Delta\tilde{\omega}_{13} = \omega_{13} - \omega_{RES} - i\Gamma_3$ is the complex detuning of the 1→3 transition from ω_{RES} , and $\Delta\tilde{\omega}_{24} = \omega_{24} - \omega_{CON} - i\Gamma_4$ is the complex detuning of the 2→4 transition from ω_{CON} . In this work, the cavity resonance

is designed to match the 1→3 transition, i.e., $\omega_{RES} = \omega_{13}$, so that $\Delta\tilde{\omega}_{13} = i\Gamma_3$. Also, although $\Delta\tilde{\omega}_{24}$ is predominantly real, in general there is an imaginary part corresponding to absorption losses in the fourth level. However when the detuning greatly exceeds the decay rate of the upper level, this contribution may be neglected. Losses from the second atomic level are also neglected, since typically it is a metastable state close to the first atomic level in energy. Finally, although in general the two cavity modes will have at least slightly different frequencies, we set $\omega_{CON} = \omega_{RES}$ for simplicity.

The Hamiltonian in equation Eq. (3.2) can then be rewritten in real space and separated into a diagonal part:

$$H_o/\hbar = \omega_{RES} \int dx \left[a_R^\dagger(x)a_R(x) + a_L^\dagger(x)a_L(x) \right] \quad (3.2)$$

$$+ \omega_{RES} (a^\dagger a + b^\dagger b + \sigma_{33} + \sigma_{44}) + \omega_{21}(\sigma_{22} + \sigma_{44}),$$

where a_L and a_R refer to left and right moving waveguide photons, respectively, as well as an interaction part:

$$H_I/\hbar = \int dx \left[a_R^\dagger(x)(-iv_g\partial_x - \omega_{RES})a_R(x) + a_L^\dagger(x)(iv_g\partial_x - \omega_{RES})a_L(x) \right] \quad (3.3)$$

$$+ V_w\delta(x)(a_R^\dagger(x)a + a_R(x)a^\dagger + a_L^\dagger(x)a + a_L(x)a^\dagger) + \Omega_c(\sigma_{23} + \sigma_{32})$$

$$+ g_{13}(a^\dagger\sigma_{13} + a\sigma_{31}) - i\Gamma_3\sigma_{33} + \Delta\tilde{\omega}_{24}\sigma_{44} + g_{24}(b\sigma_{42} + b^\dagger\sigma_{24})$$

via the interaction picture (using the rotating-wave approximation [478]), where the total system Hamiltonian is given by $H = H_o + H_I$. The eigenstate for the system can be written as:

$$|\psi_k\rangle = \left\{ \int dx \left[\phi_{k,R}^+(x)a_R^\dagger(x) + \phi_{k,L}^+(x)a_L^\dagger(x) \right] + e_k a^\dagger \right. \quad (3.4)$$

$$\left. + f_k\sigma_{31} + h_k\sigma_{21} + p_k\sigma_{41}b \right\} |0, 0, 1\rangle_{phc} \otimes |1\rangle_{atom}$$

where:

$$\phi_{k,R}^+(x) = e^{ikx} [\theta(-x) + t\theta(x)] \quad (3.5)$$

$$\phi_{k,L}^+(x) = r e^{-ikx}\theta(-x),$$

e_k is the probability amplitude of the cavity photon at ω_{RES} , and f_k , h_k , and p_k are the occupations of the 3rd, 2nd, and 4th atomic levels, respectively. t and r are the waveguide transmission and reflection amplitudes, respectively. All of these parameters are determined when the eigenequation is solved below. $|0, 0, 1\rangle_{phc} \otimes |1\rangle_{atom}$ is an eigenstate consisting of a direct product of a photonic state (*phc*) and an atomic state (*atom*). The photonic state consists of zero photons in the waveguide, zero photons in the cavity at ω_{RES} , and one photon in the cavity at ω_{CON} , respectively. The atomic state consists of a single atom in its ground state. (In this model, the number of ω_{CON} photons can be fixed to be either zero or one, since the commutator $[H, b]$ vanishes when applied to $|\psi_k\rangle$.) Note that $|\psi_k\rangle$ is written in terms of an annihilation operator b in order to simplify the notation, which would otherwise require b^\dagger operators in all but one term. In the case of n ω_{CON} photons, the coupling constant $g_{24} \rightarrow g'_{24} = g_{24}\sqrt{n}$.

Applying the Hamiltonian, equation Eq. (3.4), to the time-independent eigenvalue equation $H_I |\psi_k\rangle = \hbar\epsilon_k |\psi_k\rangle$, where $\epsilon_k = \omega - \omega_{RES}$, and solving for the reflection coefficient yields $|r(\epsilon_k)|^2 = |\Gamma_w/(\xi - i\Gamma_w)|^2$, where:

$$\xi = \epsilon_k - \frac{g_{13}^2}{\epsilon_k + i\Gamma_3 - \frac{\Omega_c^2}{\epsilon_k - g_{24}^2/(\epsilon_k - \Delta\tilde{\omega}_{24})}} \quad (3.6)$$

The parameters g_{13} , V_w (or Γ_w), v_g , and Ω_c of equation Eq. (3.4) can be determined from a numerical solution to Maxwell's equations (e.g., via [522]) as follows. First, the cavity mode is excited by a source, and the modal volume of the cavity is found from the field patterns by $V_{mode} = \left(\int_{mode} d^3x \epsilon |\mathbf{E}|^2\right) / \epsilon |\mathbf{E}_{max}|^2$. One can then apply the formula $g_{13} = \sqrt{\pi e^2 f_{13} / m \epsilon V_{mode}}$ [11], where e is the elementary electric charge, ϵ is the dielectric constant of the medium in which the atomic system is embedded, m is the free electron mass, and f_{13} is the oscillator strength for the $|1\rangle \rightarrow |3\rangle$ transition (1/2 in Na [200]). The linewidth Γ_w can be calculated by examining the decay rate of the field in the cavity mode. The waveguide group velocity is given by $v_g = \left.\frac{d\omega(k)}{dk}\right|_{\omega=\omega_{RES}}$. Finally, the Rabi frequency Ω_c can be estimated from quantum mechanics by first determining the vacuum Rabi splitting for the 2 \rightarrow 3 atomic transition g_{23} , and then multiplying by \sqrt{n} , where n is the number of ω_{23} photons.

First, consider the case of a 2-level atomic system (i.e., $\Omega_c = 0$, $g_{24} = 0$), with a waveguide coupling Γ_w and a non-radiative decay rate Γ_3 . For a fixed atom-photon coupling




Figure 3-2: Waveguide reflection for a lossless 3-level EIT atom for the four labeled values of the atomic coupling strength g_{13} , in GHz. The radiation rate $\Gamma_w = 21.5$ GHz and the ratio $g_{13}/\Omega_c = 2$ are fixed. Larger g_{13} produces larger peak separations (the blue curve shows Rabi peaks outside of the plot), favorable for switching.

g_{13} and zero non-radiative absorption, the single resonant mode at $\epsilon_k = 0$ experiences a Rabi splitting into two orthogonal linear superpositions of the cavity and atom modes at $\epsilon_k = \pm g_{13}$. As long as one remains in the strong coupling regime $g_{13} > \Gamma_3/2$, the absorption for all frequencies increases nearly linearly with Γ_3 .

However, in the opposite regime of weak coupling ($g_{13} < \Gamma_3/2$), the normal modes of the system are mostly photonic (lossless) or mostly atomic (very lossy). This phenomenon eliminates the Rabi splitting and gives rise to a reflection nearly indistinguishable from a system without an atom for sufficiently large Γ_3 .

Now, consider a 3-level atomic system without losses. Compared to the 2-level system, a third mode, corresponding to the dark state of the EIT atom, will emerge at $\epsilon_k = 0$ between the previously observed Rabi-split peaks. The dark eigenstate is given by $|\psi\rangle_{dark} = [a^\dagger - (g_{13}/\Omega_c)\sigma_{21}] |0, 0, 0\rangle_{phc} \otimes |1\rangle_{atom}$. The width of the central peak is expected to scale as $(\Omega_c/g_{13})^2$ for small Ω_c/g_{13} [29]. If one substitutes the expression given in Ref. [11] for g_{13} , one obtains the classical results found in Refs. [200, 505]. Meanwhile, the width of the side peaks is set by Γ_w and remains roughly constant as one tunes the parameters of the system.

In Fig. 3-2, $g_{13}/\Omega_c = 2$ while g_{13} is varied. It is shown that as g_{13} is decreased, the central resonance width stays constant, while the distance between the central and Rabi-split peaks becomes smaller. For use in applications, it therefore seems optimal to have a large Rabi

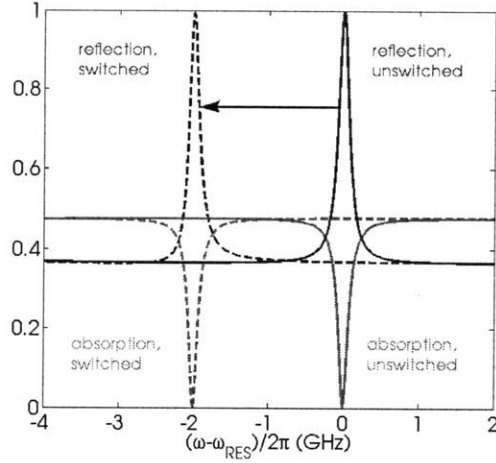


Figure 3-3: Waveguide reflection (blue) and absorption (red) in the absence (solid) and presence (dashed) of a control photon, demonstrating nonlinear single-photon switching ($\Gamma_w = 21.5$ GHz, $g_{13} = 20.5$ GHz, $\Omega_c = 2$ GHz, $\Gamma_3 = 30$ GHz, $g_{24} = 8$ GHz, and $\Delta\tilde{\omega}_{24} = 30$ GHz).

splitting, corresponding to the very strong coupling limit, which can also be viewed as corresponding to critical photon and atom numbers much less than one. The experimental values for a system with a single quantum dot emitting a single photon observed in Ref. [586] correspond to a regime where $g_{13} \approx \Gamma_w$ – specifically, they find that for operation at $\lambda = 1.182 \mu\text{m}$, $g_{13} = 20.5$ GHz and $\Gamma_w = 21.5$ GHz; note that PhC microcavities are optimal for simultaneously decreasing Γ_w and increasing g_{13} .

Now, consider a 4-level system with a control photon present. Two possible effects can be induced by the control photon. When the control frequency ω_{CON} is close to the electronic transition frequency ω_{24} , an Autler-Townes doublet is observed; upon detuned, an AC-Stark shift will be induced in this system instead [505, 556]. The latter effect has been suggested as a switching mechanism in Refs. [196, 475, 505]. This can be shown by using equation Eq. (3.6) to calculate the poles of the EIT term in the reflection, i.e., set $\epsilon_k - g_{24}^2/(\epsilon_k - \Delta\tilde{\omega}_{24}) = 0$, which yields $\epsilon_k = \pm g_{24}$ for no detuning, and $\epsilon_k \approx -g_{24}^2/\Delta\tilde{\omega}_{24}$ for a large detuning, matching the semi-classical result found in Ref. [505].

Single-photon switching is obtained when the reflection peak is shifted by an amount greater than its width, via the presence or absence of one control photon. In order to achieve this goal, one can take two different approaches. First, in the regime where $g_{13} \approx \Gamma_w$, as in Ref. [586], one can introduce an absorption via $\Gamma_3 \neq 0$, and thus absorb the majority of light not coupled to the dark state. In Fig. 3-3, the reflection and absorption are plotted for an

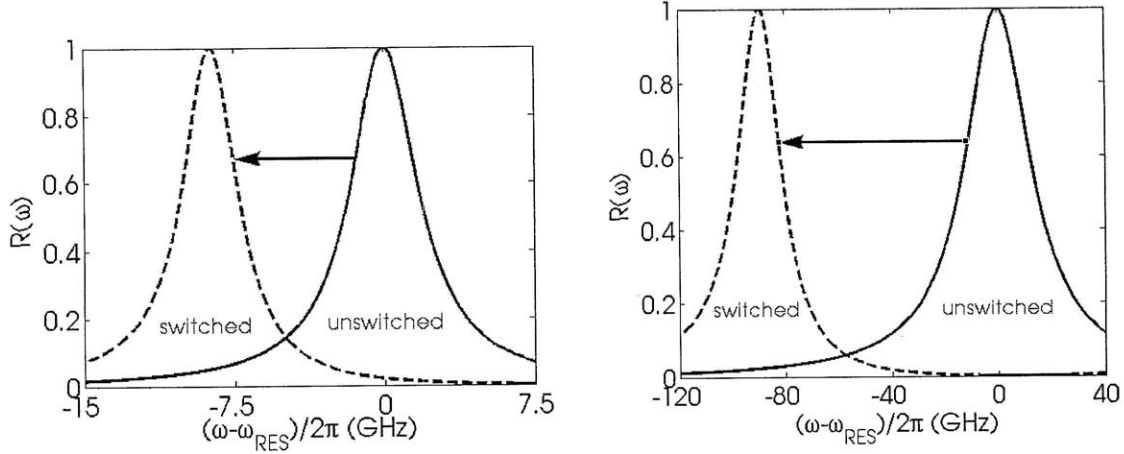


Figure 3-4: Waveguide reflection with (dashed) and without (solid) a control photon, demonstrating lossless switching. (Left:) A higher quality factor has made the resonant peaks narrow enough to be shifted by more than their full width at half maximum by a single photon ($\Gamma_w = 3$ GHz, $g_{13} = 20.5$ GHz, $\Omega_c = 30$ GHz, $\Gamma_3 = 0$, $g_{24} = 30$ GHz and $\Delta\tilde{\omega}_{24} = 20$ GHz). (Right:) Multiple (49) EIT atoms have been used to push the Rabi-split peaks farther away in the presence of negligible loss ($\Gamma_w = 21.5$ GHz, $g_{13} = 143.5$ GHz, $\Omega_c = 210$ GHz, $\Gamma_3 = 0$, $g_{24} = 210$ GHz and $\Delta\tilde{\omega}_{24} = 20$ GHz).

optimized value of $\Gamma_3 = 30$ GHz, both before and after switching. As shown, reflections at the Rabi-split frequencies are decreased substantially (to about 40%), while full reflection is still observed at the central, EIT-narrowed peak. Furthermore, in the presence of a single detuned control photon, it is possible to switch the peak reflection frequency by an amount greater than the EIT-narrowed central peak width. A second, lossless approach, appropriate if producing a large nonradiative decay Γ_3 or small Ω_c is difficult in a single-atom device, is to enhance the ratio g_{13}/Γ_w . This goal can be achieved by either decreasing Γ_w or V_{mode} , or by increasing the number of atoms from one to N . The first example of switching by decreasing the waveguide coupling is shown in Fig. 3-4, where the waveguide coupling width Γ_w is decreased by about a factor of 7 to $\Gamma_w = 3$ GHz. Now the peaks are narrow enough that a single photon of frequency ω_{CON} can shift the peak by more than the full width at half-maximum. The second example of switching, by increasing the number of atoms is illustrated in Fig. 3-4(Bottom). In general, it is clear that increasing the number of atoms collectively oscillating will improve the coupling strength; in the special case where each atom has equal coupling to the field, the N -atom treatment in Ref. [571] shows that the coupling constant $g_{13} \rightarrow g'_{13} = g_{13}\sqrt{N}$. Furthermore, one can generalize the arguments of Ref. [571] to a four-level system of N atoms to show that the other coupling constants g_{24}

and Ω_c will scale in an identical fashion (i.e., $g_{24} \rightarrow g'_{24} = g_{24}\sqrt{N}$, $\Omega_c \rightarrow \Omega'_c = \Omega_c\sqrt{N}$). This collective Rabi oscillation separates the Rabi-split peaks much further from the central peak. Figure 3-4(Bottom) shows switching exploiting this phenomenon based on parameters from Ref. [586] and using $N = 49$. The advantage of this lossless switching scheme is that one obtains a substantially greater tuning range and contrast (the difference between the peaks and the troughs) than with the lossy ($\Gamma_3 \neq 0$) scheme.

As we have shown, the reflection peak of a waveguide-cavity system can be switched in and out of resonance by a single gating photon, assuming realistic experimental parameters. Thus, one photon can be used to gate another photon of a different frequency, via a Kerr cross-phase modulation. This approach is distinct from the photon blockade system where self-phase modulation is responsible for the switching behavior. Under proper circumstances, this can give rise to two-photon entangled states. The integration of microcavities and waveguides in the same photonic crystal means that the entanglement could be preserved, in principle, throughout the system, which could be of use for quantum information processing [323]. Moreover, since these devices have an extremely small size, one can imagine creating integrated devices using thousands of these basic building blocks to process large numbers of qubits. This could have applications in quantum computing or the construction of quantum networks. If two atoms are used per node, it should even be possible to achieve quantum error correction [541].

In the following section, we switch gears and begin to probe the topic of fluctuation-induced interactions, by exploiting the Purcell effect in nanophotonic structures to enhance the nonlinear susceptibility of certain materials.

3.3 Nonlinear Purcell Enhancement

Optical nonlinearities have fascinated physicists for many decades because of the variety of intriguing phenomena that they display, such as frequency mixing, supercontinuum generation, and optical solitons [52, 124]. Moreover, they enable numerous important applications such as higher-harmonic generation and optical signal processing [3, 52, 378]. On a different note, the Purcell effect has given rise to an entire field based on studying how complex dielectric environments can strongly enhance or suppress spontaneous emission from a dipole source [35, 144, 267, 416, 466, 570] [ref]. Below, we demonstrate that the Purcell effect can

also be used to tailor the effective nonlinear optical susceptibility. While this is a general physical principle that applies to a wide variety of nonlinearities, we specifically investigate the Kerr nonlinearity, in which the refractive index is shifted by an amount proportional to intensity. This effect occurs in most materials, modeled here as originating from the presence of a collection of two-level systems. We show theoretically that using the Purcell effect for frequencies close to an atomic resonance can substantially influence the resultant Kerr nonlinearity for light of all (even highly detuned) frequencies.

Specifically, optical nonlinearities are caused by atomic or molecular resonances: the closer the frequency ω becomes to the resonant frequency ω_{ba} , the stronger the atom-photon coupling becomes and the larger the nonlinear effects. On the other hand, operating too close to resonance generally leads to large absorption loss, so instead one operates at a detuned frequency $\omega = \omega_{ba} + \Delta$. Intuitively, modifying the spontaneous-emission rate (SE) of the resonance changes a property of the resonance which causes nonlinearity, so it should modify the nonlinear process in some way. The key point is that the atomic resonance and the optical mode lie at different frequencies: this makes it possible for a photonic crystal to suppress SE at ω_{ba} via a photonic band gap, which increases $\chi^{(3)}$, as shown below, while simultaneously operating at a probe frequency ω outside the gap. In fact, one can additionally simultaneously engineer optical resonant cavities at ω , which further enhance nonlinear effects by concentrating the elds. Moreover, as we show below, this enhancement effect displays some unexpected properties. For example, while increasing SE strengthens the resonance by enhancing the interaction with the optical field, it actually makes the optical nonlinearity weaker. Furthermore, phase damping (e.g., through elastic scattering of phonons), which is detrimental to most optical processes, plays an essential role in this scheme, because in its absence, these effects disappear for large detunings (i.e., the regime in which low loss switching can take place).

A simple, generic model displaying Kerr nonlinearity is a two-level system. Its susceptibility has been calculated to all orders in both perturbative and steady state limits [52]. However, this derivation is based on a phenomenological model of decay observed in a homogeneous medium, and does not necessarily apply to systems in which the density of states is strongly modified, such as a cavity or a photonic crystal bandgap. Following an approach similar to Ref. [244], the validity of this expression can be established from a more fundamental point of view. Start by considering a collection of N two-level systems per unit

volume in a photonic crystal cavity, whose levels are labeled a and b . The corresponding Hamiltonian is given by the sum of the self-energy and interaction terms (H_0 and $V(t)$, respectively). Using the electric dipole approximation, one obtains:

$$H = H_0 + V(t) = \hbar [\omega_a \sigma_{aa} + \omega_b \sigma_{bb} + \Omega(t) \sigma_{ab} + \Omega^*(t) \sigma_{ba}], \quad (3.7)$$

where $\sigma_{ij} = c_i^\dagger c_j$ is the operator that transforms the fermionic state j to the fermionic state i , $\Omega(t) = -\mu \cdot \mathbf{E}(t)/\hbar$ is the Rabi amplitude of the applied field as a function of time, and the scalar dipole moment μ is defined in terms of its projection along the applied field $\mathbf{E}(t)$. In general, if this system is weakly coupled to the environmental degrees of freedom, then the timescale for the observable dynamics of the system is less than the timescale of the “memory” of the environment. In this case, information sent into the environment is irretrievably lost – this is known as the Markovian approximation [414]. The dynamics of this system can then be modeled by the Lindbladian \mathcal{L} , which is a superoperator defined by $\dot{\rho} \equiv \mathcal{L}[\rho]$. In general, one obtains the following master equation from the Lindbladian:

$$\dot{\rho} = -(i/\hbar) [H, \rho] + \sum_{\mu} \left[L_{\mu} \rho L_{\mu}^{\dagger} - \frac{1}{2} L_{\mu}^{\dagger} L_{\mu} \rho - \frac{1}{2} \rho L_{\mu}^{\dagger} L_{\mu} \right] \quad (3.8)$$

Using the only two quantum jump operators that are allowed in this system on physical grounds – $L_1 \equiv \sigma_{ab}/\sqrt{T_1}$ and $L_2 \equiv \sigma_{bb}\sqrt{\gamma_{\text{phase}}}$ [414] – one can obtain the following dynamical equations:

$$\frac{d\rho_{ba}}{dt} = -(i\omega_{ba} + T_2^{-1})\rho_{ba} + i\Omega(t)(\rho_{bb} - \rho_{aa}) \quad (3.9)$$

$$\frac{d(\rho_{bb} - \rho_{aa})}{dt} = -\frac{(\rho_{bb} - \rho_{aa}) + 1}{T_1} - 2i [\Omega(t)\rho_{ab} - \Omega^*(t)\rho_{ba}], \quad (3.10)$$

where $\omega_{ba} \equiv \omega_b - \omega_a$, T_1 is the rate of population loss of the upper level, and $T_2 = (1/2)T_1^{-1} + \gamma_{\text{phase}}$ is the rate of polarization loss for the off-diagonal matrix elements. The prediction of exponential decay via spontaneous emission is known as the Wigner-Weisskopf approximation [97]. Although it has been shown that the atomic population can display unusual oscillatory behavior in the immediate vicinity of the photonic band edge [243, 282], theoretical [244] and experimental considerations [26, 315] show that this approximation is fine for resonant frequencies well inside the photonic bandgap. This is assumed to be the case throughout the rest of the section. Next, one can make the rotating wave approxima-

tion for Eqs. (3.9–3.10), and then solve for the steady state. If the polarization is defined by $P = N\mu(\rho_{ba} + \rho_{ab}) = \chi E$, where χ is the total susceptibility to all orders, one obtains the following well-known expression for the susceptibility [52, 244]:

$$\chi = -\frac{N\mu^2(\omega - \omega_{ba} - i/T_2)T_2^2/\hbar}{1 + (\omega - \omega_{ba})^2T_2^2 + (4/\hbar^2)\mu^2|E|^2T_1T_2}. \quad (3.11)$$

In general, equation (??) may be expanded in powers of the electric field squared. Of particular interest is the Kerr susceptibility, also in Ref. [52]:

$$\chi^{(3)} = \frac{4}{3}N\mu^4\frac{T_1T_2^2(\Delta T_2 - i)}{\hbar^3(1 + \Delta^2T_2^2)^2}, \quad (3.12)$$

where $\Delta \equiv \omega - \omega_{ba}$ is the detuning of the incoming wave from the electronic resonance frequency. For large detunings $\Delta T_2 \gg 1$, one obtains the approximation that:

$$\text{Re } \chi^{(3)} \approx \frac{4}{3}N\mu^4\left(\frac{1}{\hbar\Delta}\right)^3\frac{T_1}{T_2}. \quad (3.13)$$

Of course, there are many types of materials to which a simple model of noninteracting two-level systems does not apply. However, it has been shown that some semiconductors such as InSb (a III-V direct bandgap material) can be treated as a collection of independent two-level systems with energies given by the conduction and valence bands, and yield reasonable agreement with experiment [350]. If the parameter Δ is defined in terms of the bandgap energy such that $\Delta \equiv \omega_G - \omega$, then one can look at the regime $\Delta T_2 \gg 1$ studied above, and obtain the following equation:

$$\text{Re } \chi^{(3)} \approx -\frac{1}{30\pi\hbar^3}\left(\frac{eP}{\hbar\omega}\right)^4\left(\frac{2m_r}{\hbar\Delta}\right)^{3/2}\frac{T_1}{T_2}, \quad (3.14)$$

where P is a matrix element discussed in Ref. [350], ω_G is the direct bandgap energy of the system, and m_r is the reduced effective mass of the exciton. This equation displays the same scaling with lifetimes as Eq. (3.13), so the considerations that follow should also apply for such semiconductors.

Now, consider the effects of changing the spontaneous emission properties for systems modeled by Eqs. (3.13–3.14). When spontaneous emission is suppressed, as in the photonic bandgap of a photonic crystal, T_1 will become large while T_2 remains finite, thus enhancing

$\chi^{(3)}$ by up to one or more orders of magnitude (for materials with the correct properties). For large detunings (where $\Delta T_{2,\text{vac}} \gg 1$), we expect that $\chi^{(3)}$ will scale as T_1/T_2 . The enhancement of the real part of $\chi^{(3)}$ is defined to be $\eta \equiv \text{Re} \chi_{\text{purcell}}^{(3)} / \text{Re} \chi_{\text{hom}}^{(3)}$, where $\chi_{\text{purcell}}^{(3)}$ is the nonlinear susceptibility in the presence of the Purcell effect, while $\chi_{\text{hom}}^{(3)}$ is the nonlinear susceptibility in a homogeneous medium. Since $T_1^{-1} = \Gamma_{\text{rad}} + \Gamma_{\text{nr}}$, the maximum enhancement is predicted to be roughly:

$$\eta \approx \frac{T_{1,\text{purcell}} T_{2,\text{vac}}}{T_{1,\text{hom}} T_{2,\text{purcell}}} = \frac{\frac{1}{2} \Gamma_{\text{nr}} + \gamma_{\text{phase}}}{\frac{1}{2} (\Gamma_{\text{nr}} + \Gamma_{\text{rad}}) + \gamma_{\text{phase}}} \frac{\Gamma_{\text{rad}} + \Gamma_{\text{nr}}}{\Gamma_{\text{nr}}} \quad (3.15)$$

where Γ_{rad} is the radiative decay rate in vacuum. Since the Purcell effect increases the amplitude of $\chi^{(1)}$, one might also expect it to increase the amplitude of $\chi^{(3)}$; however, according to Eq. (3.15), the opposite is true. This can be understood by noting that Purcell enhancement decreases the allowed virtual lifetime, and thus, the likelihood of nonlinear processes to occur [467]. Moreover, since the Purcell factor [416] is calculated by only considering the photon modes [466], one would not necessarily expect phase damping effects to also play a role. However, the results of Eq. (3.15) show the contrary to be true, and can be explained as follows: when phase damping is large, the polarization will decay quickly, thus giving rise to a small average polarization. However, as phase damping is lessened, polarization decay slows down and allows the average polarization to rise. In the limit that phase damping is controlled exclusively by the spontaneous emission rate, the two competing effects will cancel, and the nonlinearity will revert to its normal value.

Furthermore, the presence of large phase damping effects makes T_2 effectively constant, which means that suppression of spontaneous emission (caused by the absence of photonic states at appropriate energies [267]) can enhance Kerr nonlinearities by one or more orders of magnitude, while enhancement of spontaneous emission via the Purcell effect [144, 466] can suppress these nonlinearities. For the case where Purcell enhancement takes place, T_1 decreases while T_2 may not change as rapidly, due to the constant contribution of phase damping effects. This applies in the regime where $T_1 \gg \gamma_{\text{phase}}^{-1}$. Otherwise, for sufficiently small T_1 , T_2 will scale in the same way and $\chi^{(3)}$ will remain approximately constant for large detunings, where $\Delta T_2 \gg 1$. This opens up the possibility of suppressing nonlinearities in photonic crystals (to a certain degree). For processes such as four-wave mixing or cross phase modulation, $\chi^{(3)}$ will generally involve a detuning term and will differ from Eq. (3.12).

It is also interesting to note that this enhancement scheme will generally not increase nonlinear losses, which are a very important consideration in all-optical signal processing. If the nonlinear switching figure of merit ξ is defined by $\xi = \text{Re} \chi^{(3)} / (\lambda \text{Im} \chi^{(3)})$ [296], then $\xi_{\text{purcell}} / \xi_{\text{vacuum}} = T_{2,\text{purcell}} / T_{2,\text{vacuum}} \geq 1$, for all cases of suppressed spontaneous emission.

The general principal described thus far should apply for any medium where the local density of states is substantially modified. In what follows, we show how this effect would manifest itself in one such exemplary system: a photonic crystal. This example serves as an illustration as to how strong nonlinearity suppression / enhancement effects could be achieved in realistic physical systems. It consists of a 2D triangular lattice of air holes in dielectric ($\epsilon = 13$), with a two-level system placed in the middle, as in Fig. 3-5.

Note that the vast majority of photonic crystal literature is generally focused on modification of dispersion relations at the frequency of the light that is sent in as a probe. By contrast, in the current work, it is only essential to modify the dispersion relation for the frequencies close to the atomic resonances; the dispersion at the frequency of the light sent in as a probe can remain quite ordinary.




Figure 3-5: A 2D triangular lattice of air holes in dielectric ($\epsilon = 13$). On top of the dielectric structure in grey, the E_z field is plotted, with positive values in red, and negative values in blue. A small region of nonlinear material is placed exactly in the center of the structure. This material may be, for example, either two-level atoms, quantum wells, or some semiconductors such as InSb.

First, consider the magnitude of the enhancement or suppression of spontaneous emission in this system. Clearly, since there are several periods of high contrast dielectric, two effects are to be expected. First, there will be a substantial but incomplete suppression of emission inside the bandgap. Second, there will be an enhancement of spontaneous emission outside the bandgap (since the density of states is shifted to the frequencies surrounding the bandgap). For an atom polarized in the direction out of the 2-D plane, only the TM polarization need be considered.

We numerically obtain the enhancement of spontaneous emission by performing two time-domain simulations in Meep [151], a finite difference time-domain code which solves Maxwell's equations exactly with no approximations, apart from discretization (which can be systematically reduced) [522]. First, we calculate the spontaneous emission of a dipole placed in the middle of the photonic crystal structure illustrated in Fig. 3-5, then divide by the spontaneous emission rate observed in vacuum. The resulting values of T_1 and T_2 are calculated numerically, and Eq. (3.12), in conjunction with the definition of the enhancement factor η , is used to plot Fig. 3-7. The results are plotted in Fig. 3-6.

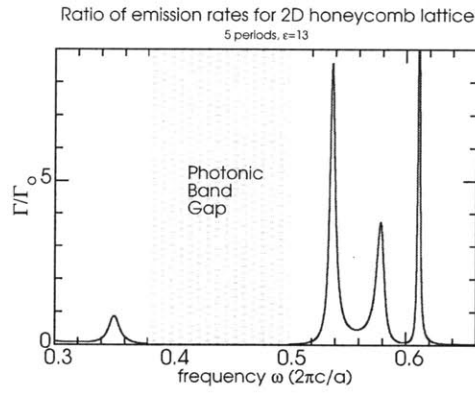


Figure 3-6: Relative enhancement of the TM local density of states for Fig. 3-5, as measured in the time-domain simulation rate of emission, Γ , normalized by the emission rate in vacuum, Γ_0 .




Figure 3-7: Contour plot of Kerr enhancement $\eta \equiv \text{Re} \chi_{\text{purcell}}^{(3)} / \text{Re} \chi_{\text{hom}}^{(3)}$ as a function of probe (ω_{ph}) and electronic transition (ω_{elec}) frequencies, for a single quantum well of GaAs-AlGaAs, (a) at $T = 200$ K, with $0.1\gamma_{\text{phase}} = 10\Gamma_{\text{nr}} = \Gamma_{\text{rad}}$, and (b) at $T = 225$ K, with $0.1\gamma_{\text{phase}} = \Gamma_{\text{nr}} = \Gamma_{\text{rad}}$.

A GaAs-AlGaAs single quantum well can lie in the interesting regime discussed above,

where the radiative loss rate Γ_{rad} dominates the non-radiative loss rate Γ_{nr} as well as the overall loss rate of the quantum well, for certain temperatures [276]. Equation (3.14) implies that one can see substantial enhancement of the Kerr coefficient in that regime.

At a temperature of about 200 K, $\Gamma_{\text{nr}} \approx 0.1\Gamma_{\text{rad}}$ [276]. Although experimental measurements for γ_{phase} are unavailable to the authors, the presence of a substantial phonon bath at that temperature leads one to expect a fairly large value, which may be conservatively estimated by $10\Gamma_{\text{rad}}$. These results are displayed in Fig. 3-7(a). Note that enhancement is primarily observed inside the photonic bandgap (e.g., Fig. 3-6). We observe an enhancement in the real part of the Kerr coefficient up to a factor of 12, close to the predicted maximum enhancement factor of 10.48 in the regime of large detunings ($\Delta T_2 \gg 1$).

Also, at a temperature of about 225 K, $\Gamma_{\text{nr}} \approx \Gamma_{\text{rad}}$ [276], and again we take $\gamma_{\text{phase}} = 10\Gamma_{\text{rad}}$. These results are displayed in Fig. 3-7(b). In this case, we observe an enhancement up to a factor of 2.5, close to the predicted maximum enhancement factor of 1.91 in the regime of large detunings ($\Delta T_2 \gg 1$).

Finally, we note that close to room temperature (285 K), the system in Ref. [276] displays $\Gamma_{\text{nr}} \approx 10\Gamma_{\text{rad}}$, which is predicted to yield a maximum enhancement factor of 1.06. Since this number is fairly negligible, it illustrates that this approach has little impact when non-radiative losses dominate the decay of the electronic system.

On the other hand, some recent work has demonstrated that a single quantum dot can demonstrate predominantly radiative decay in vacuum even at room temperature, e.g., single CdSe/ZnS core-shell nanocrystals with a peak emission wavelength of 560 nm, with $\Gamma_{\text{rad}} \approx 39\Gamma_{\text{nr}}$ [66]. Even bulk samples of similar nanocrystals have been shown to yield a significant radiative decay component, corresponding to $\Gamma_{\text{rad}} \approx \Gamma_{\text{nr}}$ [525]. Thus, we predict that with strong suppression of radiative decay, nonlinear enhancement of a factor of two, or more, could be observed at room temperature.

We now discuss the implications of this effect on previous work describing nonlinearities in photonic crystals, such as Ref. [503], and the references therein. Most past experiments should not have observed this effect, because they were designed with photonic bandgaps at optical frequencies significantly smaller than the frequencies of the electronic resonances generating the nonlinearities, in order to operate in a low-loss regime. Furthermore, in most materials, non-radiative decays will dominate radiative decays at room temperature. Finally, all the previous *analyses* are still valid as long as one considers the input parameters

to be effective nonlinear susceptibilities, which come from natural nonlinear susceptibilities modified in the way described in this section.

In conclusion, we have shown that the Purcell effect can be used to tailor optical nonlinearities. This principle manifests itself in an exemplary two-level system embedded in a photonic crystal; for realistic physical parameters, enhancement of Kerr nonlinearities by more than an order of magnitude is predicted. The described phenomena is caused by modifications of the local density of states near the resonant frequency. Thus, this treatment can easily be applied to analyze the Kerr nonlinearities of two-level systems in almost any geometrical structure in which the Purcell effect is substantial (e.g., photonic crystal fibers [310], optical cavities). It also presents a reliable model for a variety of materials, such as quantum dots, atoms, and certain semiconductors. Future investigations will involve extending the formalism described here to other material systems.

3.4 Concluding Remarks

The preceding section marks a turning point in our discussion of nonlinear nanophotonics and is a segway for the subject of the remainder of this thesis: fluctuation-induced interactions. Specifically, it was shown that by modifying the density of states around an atomic resonance (e.g. increasing Q/V), one could dramatically enhance the SE rate of the atom and thereby strengthen the atom–photon interaction [96, 119, 286, 288, 348, 416, 468]. Interestingly, the SE rate of an atom is a direct consequence of quantum electromagnetic fluctuations, as ordinary (first-quantized) quantum mechanics forbids such spontaneous transitions from ever occurring [479, 585]. In the language of quantum mechanics, the existence of SE comes from the “mixing” or interaction of the vacuum (ground-state) electromagnetic fluctuations with the eigenstates of the atomic system (computed in the absence of fluctuations) [348, 585].

Because the interaction of the vacuum electromagnetic field with polarizable atoms is the underlying basis of the Purcell effect, one may view the aforementioned effect of Sec. 3.3 as a consequence of a type of fluctuation-induced interaction. As we shall see in the upcoming chapters, fluctuation-induced interactions not only play an important role in determining the properties atoms and their effect on atom–photon interactions, but they can also have dramatic consequences on the behavior of macroscopic objects.

Chapter 4

Casimir Forces: A Brief

Introduction

“I am thankful for laughter, except when milk comes out of my nose.” *Woody Allen*

4.1 Introduction

Quantum fluctuations can give rise to a wide variety of important phenomena, such as the Purcell effect studied in the previous chapter, or even intermolecular forces between neutral atoms [68,77,352]. One of the most dramatic manifestations of quantum mechanics observed in the last half-century is the Casimir force: a tiny force on an uncharged, source-free body due to changes in the zero-point energy associated with quantum vacuum fluctuations (virtual photons) [22,68,77,78,304,352,370,409]. Aside from being a fascinating application of quantum electrodynamics, Casimir forces also play a significant role in the physics of neutral, macroscopic bodies at micrometer separations, such as in new generations of micro-electromechanical systems (MEMS) [83,86,481]. Casimir forces are also studied in order to place constraints on non-Newtonian theories of gravity [358,369,390] and are becoming a relevant component of hadronic models [355,357]. In this chapter, we give a brief description of the Casimir effect and emphasize some of the most important developments in the field.

4.1.1 From van der Waals to Casimir-Polder Forces

The theory of fluctuation-induced interactions between atoms originated in the early 1930s with the works of Ref. [134,318]. It was then recognized that thermodynamic and quantum

fluctuations between atoms separated by large distances (compared to typical atomic sizes) contributed significantly to what are now known as van der Waals forces [68, 77, 78, 116, 302, 318, 352]. As was first shown by Ref. [318], they originate due to induced atomic dipole moments $\mathbf{p} = \alpha\mathbf{E}$ arising from atomic charge fluctuations. The interaction of these induced dipole moments leads to van der Waals [77] and London-dispersion forces [134, 318]. In particular, for atoms separated by a distance d , the potential energy of interaction was shown to scale $\sim 1/d^6$ (using simple perturbation theory). For atoms separated by large distances (compared to the wavelength λ of relevant atomic transitions), retardation effects associated with the finite speed of light become non-negligible and can dramatically change the behavior of the long-range interaction. These retardation effects were first calculated by Casimir and Polder in 1948 [78]. In that seminal work, they demonstrated that long-range forces decayed more rapidly than their short-range (van der Waals) counterparts. In particular, the Casimir-Polder potential between two atoms separated by a distance $d \gg \lambda$ (shown schematically in Fig. 4-1) can be written in terms of their static polarizabilities α_1 and α_2 , and yields an attractive interaction between the atoms, given by [78]:

$$U = -\frac{23\hbar c}{4\pi d^7}\alpha_1\alpha_2 \quad (4.1)$$

Direct measurements of this retarded van der Waals force were made in the 1950s [116, 230], and more recently as well [105, 491, 538]. The passage from the short-range to the long-range regime was an important achievement in the understanding of fluctuation-induced forces [230], and paved the way for the study of objects of macroscopic extent, which is the subject of the remainder of this thesis. For a review of retarded atom-atom Casimir-Polder interactions, the reader is referred to Ref. [68, 129, 230].

Macroscopic objects are composed of many atomic constituents. These objects are therefore naturally described in terms of averaged quantities that vary continuously in space (thus the term continuous media). In particular, charge fluctuations on macroscopic objects are described by effective macroscopic charge and current densities $\rho(\mathbf{r}, t)$ and $\mathbf{J}(\mathbf{r}, t)$, or more commonly via the polarization field $\mathbf{P}(\mathbf{r}, t) \sim \partial\mathbf{J}/\partial t$ ¹. Figure 4-1(Right) illustrates this idea schematically: two macroscopic objects are interacting via the electromagnetic

¹Determining the induced macroscopic polarization from the microscopic properties of materials is a beautiful subject that is outside the scope of this thesis, although our discussion of the nonlinear Purcell effect does involve such a procedure. For a more detailed discussion, the reader is encouraged to read Ref. [51].




Figure 4-1: (Left:) Two atoms separated by a distance d are interacting via fluctuation-induced dipole moments (polarizabilities α_1 and α_2). (Right:) Two macroscopic objects separated by a distance d are interacting via the electromagnetic field produced by the fluctuation-induced charges on the interior and surface of the bodies, described by macroscopic polarization fields \mathbf{P}_1 and \mathbf{P}_2 .

fluctuations produced by fluctuating charges on the interior (and surface) of the bodies, described by effective polarization fields \mathbf{P}_1 and \mathbf{P}_2 . It is this macroscopic regime of the Casimir-Polder interaction that is known as the Casimir effect ².

Although Casimir-Polder forces between macroscopic objects have their origin in the simple intermolecular forces between neutral atoms, their behavior can be dramatically different ³. The Casimir effect is more complicated in at least two major ways: First, the mutual interaction of the many atoms comprising the participating objects leads to dramatically-different material properties, more complicated than individual atomic polarizabilities ⁴. Second, Casimir forces depend dramatically on the shape of the participating objects due to the boundary conditions that the electromagnetic field must satisfy at the object boundaries [231]. For example, the electric field between two perfectly-metallic plates must satisfy the standard Dirichlet boundary conditions imposed by Maxwell's equations (if the plates span the x - y plane, then $\mathbf{E} \cdot \hat{\mathbf{z}} = 0$ at the plates' surface). Therefore the kinds of charge or electromagnetic fluctuation that can be excited will vary depending on the separation between the plates. Consequently, the electromagnetic energy density $\sim \langle \mathbf{E}^2 \rangle$ and therefore the force between the plates will depend strongly and non-additively on the separation. To a large extent, the focus of much research on the Casimir effect is in understanding

²In the dilute limit (of rarefied media), the dielectric permittivity of a collection of atoms of density n can be defined in terms of the atomic polarizability as $\epsilon = 1 + 4\pi n\alpha$ [78]. In this limit, one can arrive at the force between individual atoms, given by Eq. (4.1), from the macroscopic formula describing the interaction energy between two dielectric half-spaces, known as the Lifshitz formula, via an appropriate limiting procedure and the above relationship between ϵ and α .

³Casimir-Polder atom-atom interactions serve only as a reasonable heuristic for understanding the kinds of effects that one can expect from macroscopic Casimir forces, as we shall see in chapter 7.

⁴The validity and accuracy of macroscopic dielectric models in Casimir physics has been a topic of recent debate in the field [352].

its dependence on geometry (boundary conditions) and materials. In the following section, we briefly describe two equivalent ways through which it is possible to relate the geometry and material properties of macroscopic objects to the equilibrium charge/electromagnetic fluctuations.

4.1.2 Charge \leftrightarrow Electromagnetic Fluctuations

In thermodynamic equilibrium, the fluctuating charges and photons are both at equilibrium with each other, so that one can view Casimir forces as resulting from either fluctuating charges or from fluctuating electromagnetic fields.

In the absence of thermodynamic fluctuations, macroscopic neutral objects have vanishing charge and current densities, and consequently do not interact. The presence of fluctuations however, produces charge and current fluctuations that cancel on average, i.e., $\langle \rho \rangle = \langle \mathbf{J} \rangle = 0$. In turn, these give rise to fluctuating electromagnetic fields that also cancel on average, i.e., $\langle \mathbf{E} \rangle = 0$, whose interaction with the same fluctuating charges result in long-range forces between the objects [22, 352, 400]. (The average force $\sim \langle \mathbf{J} \cdot \mathbf{E} \rangle$ does not vanish because it is a quadratic function of the charge fluctuations, and the correlation of the fluctuations is non-zero.) This simple picture, based on radiation-reaction [233], offers a simple explanation for the existence of van der Waals forces, and also forms the basis of the Casimir effect. If our only goal is to describe the thermodynamic (steady-state) behavior of the entire system, then all that remains is to consider the resulting thermodynamic distribution of either the charge or the electromagnetic field (because they are at equilibrium, then one immediately leads to the other).

One way to describe the equilibrium electromagnetic distribution, independent of the sources, is to describe the electromagnetic field as a collection of harmonic oscillators. In particular, the total electromagnetic energy, a quadratic quantity of the electric and magnetic fields ($\sim \langle \mathbf{E}^2 \rangle + \langle \mathbf{B}^2 \rangle$) integrated over all space (see chapter 5), can be described as a collection of photons whose distributions as a function of frequency and temperature is determined by the Bose-Einstein distribution $\langle n \rangle_{\omega, T} = \hbar\omega / [\exp(\hbar\omega/kT) + 1]$ [257, 287, 435]. The average total electromagnetic energy, perhaps the most popularly recognized expression of the Casimir effect, is given by [352]:

$$\langle U \rangle = \sum_{\omega} \hbar\omega \left(\langle n \rangle_{\omega, T} + \frac{1}{2} \right). \quad (4.2)$$

Again, because the distribution and energy of the photons depends on the allowed electromagnetic frequencies ω , the total energy will also depend on the shape and material properties of the objects in the geometry. At zero temperature, $\langle n \rangle = 0$ and the temperature fluctuations are of a purely quantum character. These corresponds to the zero-point oscillations of the electromagnetic field, or the so-called “virtual photons”, that arise due to the Heisenberg uncertainty principle [352]. For perfectly-metallic parallel plates, Eq. (4.2) yields the well-known plate–plate force per unit area $F \equiv \partial\langle U \rangle / \partial d = -\hbar c \pi^2 / 240 d^4$.

A large focus of fluctuational electrodynamics concerns the relationship between the equilibrium correlation functions of the electromagnetic field, $\langle E_i E_j \rangle$ and $\langle H_i H_j \rangle$, and the properties of the geometries/materials that produce them. From these quantities, one can obtain much if not all of the equilibrium physics of the electromagnetic field (at all temperatures). For example, the electromagnetic energy at any point in space depends on the self-correlations $\langle E_i E_i \rangle$ and $\langle H_i H_i \rangle$, which can be used to derive the expression in Eq. (4.2).

One way to arrive at the correlation functions of the electromagnetic field is to exploit a powerful theorem, known as the fluctuation-dissipation theorem (FDT), based on linear-response theory [131, 288, 465]. The essence of the fluctuation-dissipation theorem lies in the direct correspondence between fluctuations and dissipation [131, 257, 435, 465]. Fluctuation-dissipation theorems in electromagnetism relate the fluctuations of dissipative, polarizable media, represented by the stochastically varying macroscopic electric and magnetic polarization fields \mathbf{P} and \mathbf{M} , to the fluctuations of the electromagnetic field. These relations come in two flavors, known as fluctuation-dissipation theorems of the first and second kinds, denoted by $\text{FDT}^{(1)}$ and $\text{FDT}^{(2)}$, respectively. The relationship between these two types of FDTs is described in Ref. [130]. In what follows, we merely write down the basic results and point out their applications to Casimir physics. For a detailed derivation of these results, the reader is encouraged to read Ref. [130] and references therein.

The origin of the electromagnetic fluctuations are the electric and magnetic charge density fluctuations produced by dissipative bodies. Given these external quantities, \mathbf{P} and \mathbf{M} , one can immediately describe the resulting electromagnetic field by way of Maxwell’s equations. In the frequency domain, they are written as:

$$\nabla \times \mathbf{E} = i\omega(\mu\mathbf{H} + 4\pi\mathbf{M}), \quad \nabla \times \mathbf{H} = -i\omega(\varepsilon\mathbf{E} + 4\pi\mathbf{P}), \quad (4.3)$$

where ε and μ are the dielectric permittivities and permeabilities of the objects in the geometry, the imaginary parts of which correspond to absorption.

The first FDT allows us to infer the equilibrium correlation functions of the electromagnetic field in terms of the electromagnetic Green's function (GF), or the response functions of the Maxwell operator. They are given, in the frequency domain, by [304, 465]:

$$\langle E_i(\mathbf{r})E_j(\mathbf{r}') \rangle_\omega = -\hbar\omega^2 \coth(\hbar\omega/2kT) \text{Im} G_{ij}(\mathbf{r}, \mathbf{r}', \omega) \quad (4.4)$$

$$\langle H_i(\mathbf{r})H_j(\mathbf{r}') \rangle_\omega = \hbar \coth(\hbar\omega/2kT) \text{Im}(\nabla \times)_{il}(\nabla' \times)_{jm} G_{lm}(\mathbf{r}, \mathbf{r}', \omega) \quad (4.5)$$

$$\langle E_i(\mathbf{r})H_j(\mathbf{r}') \rangle_\omega = \hbar\omega \coth(\hbar\omega/2kT) [-\text{Re}(\nabla' \times)_{ji} G_{il}(\mathbf{r}, \mathbf{r}', \omega)], \quad (4.6)$$

where $\langle \mathbf{E} \rangle = \langle \mathbf{H} \rangle = 0$, and G_{ij} is the standard photon Green's function, determined by the solution of the equation:

$$[\nabla \times \nabla \times - \varepsilon\omega^2] \mathbf{G}_j(\mathbf{r}, \mathbf{r}', \omega) = \delta(\mathbf{r} - \mathbf{r}') \hat{\mathbf{e}}_j \quad (4.7)$$

Note that the properties of the materials (the sources of radiation, ε and μ) do not appear explicitly in Eqs. (4.4–4.6). The presence of dissipation in the system is “hidden” inside the photon GF, whose imaginary part is zero in the absence of loss⁵. (Note also that the FDT⁽¹⁾ applies only at equilibrium. Thus, the time average of $\langle E_i H_j \rangle$ will be identically zero everywhere in space, i.e. there will be no net flux or heat radiated between the bodies.)

The second FDT involves the correlation functions of the underlying source of the fluctuations (charges and currents). In this case, one applies the FDT to the sources of radiations, which yields the following relations:

$$\langle P_i(\mathbf{r})P_j(\mathbf{r}') \rangle_\omega = \hbar \coth(\hbar\omega/2kT) \delta(\mathbf{r} - \mathbf{r}') \text{Im} \varepsilon_{ij}(\mathbf{r}, \omega), \quad (4.8)$$

$$\langle M_i(\mathbf{r})M_j(\mathbf{r}') \rangle_\omega = \hbar \coth(\hbar\omega/2kT) \delta(\mathbf{r} - \mathbf{r}') \text{Im} \mu_{ij}(\mathbf{r}, \omega), \quad (4.9)$$

where ε and μ are the linear response functions of the materials (here, they are considered to be anisotropic for general purposes). By employing Eqs. (4.8–4.9) in conjunction with Maxwell's equations, one can determine the field correlations, this time expressed as

⁵In the previous section, we mentioned that there exist a finite force between perfectly-metallic plates, which would seem to contradict the underlying basis of the fluctuation-dissipation theorem. For a brief comment on this apparent contradiction, the reader is referred to the discussion of the Lifshitz formula in chapter 8.

integrals over the sources of radiation:

$$\begin{aligned} \langle E_i(\mathbf{r})E_j(\mathbf{r}') \rangle_\omega &= \hbar\omega^4 \coth(\hbar\omega/2kT) \\ &\int_V d^3\mathbf{r}'' \operatorname{Im} \varepsilon_{ij}(\mathbf{r}'', \omega) G_{ik}(\mathbf{r}, \mathbf{r}'', \omega) G_{jk}^*(\mathbf{r}', \mathbf{r}'', \omega) \end{aligned} \quad (4.10)$$

$$\begin{aligned} \langle H_i(\mathbf{r})H_j(\mathbf{r}') \rangle_\omega &= -\hbar\omega^2 \coth(\hbar\omega/2kT) \\ &\int_V d^3\mathbf{r}'' \operatorname{Im} \varepsilon_{ij}(\mathbf{r}'', \omega) (\nabla \times)_{il} G_{lk}(\mathbf{r}, \mathbf{r}'', \omega) (\nabla' \times)_{km} G_{jm}^*(\mathbf{r}', \mathbf{r}'', \omega) \end{aligned} \quad (4.11)$$

$$\begin{aligned} \langle E_i(\mathbf{r})H_j(\mathbf{r}') \rangle_\omega &= -i\hbar\omega^3 \coth(\hbar\omega/2kT) \\ &\int_V d^3\mathbf{r}'' \operatorname{Im} \varepsilon_{ij}(\mathbf{r}'', \omega) G_{ik}(\mathbf{r}, \mathbf{r}'', \omega) (\nabla' \times)_{km} G_{jm}^*(\mathbf{r}', \mathbf{r}'', \omega), \end{aligned} \quad (4.12)$$

FDT⁽¹⁾ [Eqs. (4.4–4.10)] and FDT⁽²⁾ [Eqs. (4.10–4.12)] can be shown to be equivalent by using a number of the symmetries of the photon GF, in conjunction with the relation: $\operatorname{Im} \varepsilon \delta(\mathbf{r} - \mathbf{r}') = -1/2 \omega^2 i [G^{-1}(\mathbf{r}, \mathbf{r}', \omega)^* - G^{-1}(\mathbf{r}, \mathbf{r}', \omega)]$. In FDT⁽²⁾, however, the correlations depend nonlinearly on the photon GFs, and arise from the cumulative contribution of all of the dissipative sources in the geometry (volume V). The explicit dependence of the field correlation functions on the dissipative current sources is an advantage of this formulation, as it directly reveals the origins of the fluctuations. In fact, because one can distinguish the contributions to the correlations with respect to their origins, FDT⁽²⁾ can also be employed in certain non-equilibrium situations [75, 130, 376]. For example, from these relations, one can also readily compute the thermal radiation emanating from bodies [75, 322, 465].

The two fluctuation-dissipation theorems above reduce the problem of computing quantum fluctuations to a classical calculation: they relate the quantum-mechanical excitations of the electromagnetic field to the classical response of a dissipative medium. Whether one chooses to employ the first or second FDT is a matter of convenience, since they both offer equivalent representations of the fluctuations, although again FDT⁽²⁾ does apply to more general circumstances. For the remainder of this thesis, we shall only focus on the use of FDT⁽¹⁾ for Casimir calculations at equilibrium, and leave the problem of computing forces out of equilibrium for future work.

In the following section, we give a brief review of some of the problems that motivated the work in this thesis. In particular, we describe some of the previous work on the Casimir effect, along with a number of the outstanding questions in the field.

4.2 Experimental and Theoretical Progress

There have been many experimental verifications of the Casimir force reported in recent decades, by delicate experiments at micron and sub-micron lengthscales [48, 67, 74, 88, 110, 269, 283, 284, 352, 358, 365, 370, 390, 409], but always restricted to simple geometries (parallel plates [59], spheres and plates [59, 82, 83, 87, 109, 111, 195, 218, 230, 283, 307, 326, 365, 460, 538], or crossed cylinders [132, 230, 326]). Many of the early Casimir-force experiments were motivated by the early theoretical predictions of H. Casimir, which applied only to finite planar geometries [77]. Not surprisingly, the majority of works on the Casimir force have focused in this particular case, as evidenced by the large number of experiments involving spheres and plates. (The sphere–plate structure is preferable to the original plate–plate geometry due to some of the challenges associated with keeping two surfaces parallel to one another [46].) Unfortunately, the force in these particular geometries is almost always attractive [260] and monotonically decreasing with separation, leading to “stiction” and other experimental difficulties [82, 83].

Thus, one might ask whether it is possible to obtain more complicated behavior (such as non-monotonic or repulsive forces) in more complex structures, and more generally whether complex geometries might give rise to unexpected force phenomena. This question has begun to be addressed in experiments [84] and by promising new theoretical methods [105, 122, 138, 143, 177, 180, 261, 281, 404, 425, 428, 434, 438, 443]. Theoretical efforts to predict Casimir forces for geometries very unlike the standard case of parallel plates have begun to yield fruit, having demonstrated a number of interesting results for strong-curvature structures [46, 141, 143, 178, 359, 362, 442, 452]. For example, recent works have shown that it is possible to find unusual effects arising from many-body interactions or from systems exhibiting strongly coupled material and geometric dispersion [13, 298, 360, 442, 446, 452, 590]. Some of these effects are described in chapter 7 and 8. (Only a few months ago, the first repulsive geometry between vacuum-separated metallic objects, consisting of a spheroidal particle suspended above a plate with an air hole, was discovered [298].)

The behavior of Casimir forces can also be changed by employing unusual material systems [54, 74, 162, 217, 219, 229, 259, 485, 491, 557]. As discussed in chapter 8, it is even possible to suspend objects, create non-touching configurations of compact objects, and observe strong temperature corrections to the Casimir force in geometries consisting of

fluid-separated objects [450, 452, 453]. Recent works (performed within the last decade) have also demonstrated a large number of interesting Casimir-enabled effects of consequence to micro-electromechanical systems and amenable to experiments [46, 67, 122, 142, 169, 220, 268, 327, 362, 374, 375, 390]. Overall, there are many aspects of the Casimir problem that are only recently beginning to be investigated by experiments. These include various geometry [269, 279, 461] and material effects [364], skin-depth [221, 307] and surface roughness [40, 396] effects, forces between anisotropic dielectrics [401, 457], critical forces [162], repulsive forces [294] and also Casimir torques [373, 401]. Most of these works have, nevertheless, been mainly focused in two-dimensional [46, 140, 208, 441] or simple three-dimensional constant-cross-section geometries [107, 143, 442] for which numerical calculations are tractable. Theoretical challenges still remain, more so in geometries involving multiple bodies and/or multiple lengthscales [443]. For more complicated geometries, however, calculations become extremely cumbersome and often require drastic approximations, a limitation that has hampered experimental and theoretical work beyond the standard geometries.

Until recently, the common approach to predicting the Casimir force had been to consider approximations for small perturbations around known solutions, such as parallel plates or dilute gases. For parallel plates in d dimensions, separated by a distance a , there is a well-known attractive force that scales as $1/a^{d+1}$, first predicted by Casimir [77] and later extended to formulas for any planar-multilayer dielectric distribution $\varepsilon(x, \omega)$ via the generalized Lifshitz formula [529]. A direct, intuitive extension of this result is the proximity force approximation (PFA) [48], which treats the force between two surfaces as a two-body interaction given by the sum of “parallel plate” contributions. Valid in the limit of small curvature, PFA provides an easy way to conceptualize the Casimir force in complex geometries as a simple two-body force law, but unfortunately it may also be deceptive: outside its range of applicability, the Casimir force is not additive [352] and may be qualitatively different from PFA’s predictions [71, 135, 138, 140, 168, 179, 327, 442]. Other perturbative approaches include renormalized Casimir-Polder [480, 524] or semi-classical interactions [473], multiple scattering expansions [21, 280], classical ray optics approximations [234], higher-order PFA corrections [46, 184], and other perturbative techniques [139]. The ray-optics approach is especially interesting because, although it is only strictly valid in the small-curvature limit, it captures multiple-body interactions (see Chapter 6) and can therefore sometimes predict

the qualitative behavior in cases where other approximations fail [62, 234]. Nevertheless, none of these methods provide any guarantees of accuracy in arbitrary geometries, where they involve uncontrolled approximations. Therefore, for complex new geometries, where one might hope to encounter behaviors very different from those in the parallel-plate limit, a different approach is required.

Toward this end, researchers have sought “exact” numerical methods applicable to arbitrary geometries—that is, methods that converge to the exact result with arbitrary accuracy given sufficient computational resources. One such method was proposed by Ref. [72], based on a path-integral representation for the effective action, which was first applied to obtain predictions of forces between cylinders and a plates [141], and between corrugated surfaces [135, 140]. It is based on a surface parametrization of the fields coupled via vacuum Green’s functions, requiring $O(N^2)$ storage and $O(N^3)$ time for N degrees of freedom, making scaling to three dimensions problematic. Another exact method is the “world-line approach” [177], based on Monte-Carlo path-integral calculations. The scaling of the world-line method involves a statistical analysis, determined by the relative feature sizes in the geometry, and is discussed below. Until recently, the methods of Ref. [72] and Ref. [177] had only been demonstrated for perfect-metallic z -invariant structures—in this case, the vector unknowns can be decomposed into TE ($\mathbf{E} \cdot \hat{\mathbf{z}} = 0$) and TM ($\mathbf{H} \cdot \hat{\mathbf{z}} = 0$) scalar fields with Neumann and Dirichlet boundary conditions, respectively—although generalizations have been proposed [136, 177].

In the last few years, the number of theoretical methods for computing Casimir forces in arbitrary geometries has grown considerably. Thus far, there are two different, yet complementary approaches.

Spectral methods often employ a complete geometry-dependent spectral basis to represent the objects of interest, allowing rapid (exponential) convergence as a function of the number of basis functions [50, 186]. One promising type of spectral method for Casimir-force calculations are scattering methods, developed by a number of works [21, 69, 137–139, 143, 170, 184, 186, 207, 232, 261, 280, 299, 360, 426], which has been shown to be very efficient for spherical, cylindrical or nearly-spherical geometries. Spectral methods can offer a great deal of insight and can be extremely efficiency in certain asymptotic limits (e.g. the multipole method of [143] converges very rapidly for nearly-spherical objects at large separations). Using a multi-pole scattering approach, Ref. [427] proved a recent theorem constraining the

types of stability that may exist between vacuum-separated objects. The disadvantage of spectral methods is their poor performance in general geometries consisting of complicated objects (e.g. objects with corners) or involving many objects at close separations. Because the choice of spectral basis is often dependent on the particular geometry of interest, these methods may also require new codes for every new geometry. The alternative is to employ the same spectral basis for every geometry, but this reduces to a brute-force method that does not generally achieve spectral (i.e. exponential) accuracy. Nevertheless, sometimes these generic (fixed-basis) spectral methods are convenient (e.g. the electromagnetic solver MPB is a spectral method, although it is not exponentially accurate except for smooth periodic epsilon [249]). In the case of integral-equation methods (below), however, most of the fast-solver algorithms have been designed for localized basis functions (BEM methods), so a spectral basis may not be the most convenient choice for applying known algorithms.

A different class of numerical method, the focus of this thesis, are methods which employ localized basis functions that are independent of the particular geometry or materials of interest. There are three main such kinds of methods, whose efficiency depend on the type of problem being studied. First, one can employ boundary-element methods [45, 91] or integral-equation methods, which discretize the surfaces of the objects into planar triangle and introduce localized basis functions on these surface elements in order to represent the electromagnetic unknowns (the scattering coefficients of the scattering methods). Recent implementations of the Casimir-force problem using boundary-element methods were demonstrated recently by various authors, for very complicated three-dimensional geometries [434, 566]. Second, one can employ finite element methods, which discretize the interior and surface of the objects in the geometry [237, 548]. Finally, the simplest class of methods are finite-difference methods [93, 150, 278, 515, 523], the subject of this thesis, which discretize the entire geometry by defining the electromagnetic unknowns on a grid (often chosen to be uniform for simplicity). An example of this discretization procedure, known as a Yee grid, is shown in chapter 5. The advantage of these numerical methods is that their convergence properties can be made independent of the shape and materials of the objects in the geometry. They also converge rapidly enough for accurate computations involving large three-dimensional objects, albeit less so than spectral methods in geometries with a very high degree of symmetry for which spectral methods can be made extremely efficient. Furthermore, modifying the geometry does not require any additional modifications

or “code”. Finally, because these methods have been studied for over four decades, in the context of classical electromagnetic calculations, they are well understood and require little, if any, modification.

In the following chapters, we formulate the calculation of Casimir forces by exploiting well-known techniques in classical electromagnetism, based on the finite-difference approach.

Chapter 5

Casimir Forces in the Imaginary Frequency Domain

“To love is to suffer. To avoid suffering, one must not love. But then, one suffers from not loving. Therefore, to love is to suffer; not to love is to suffer; to suffer is to suffer. To be happy is to love. To be happy, then, is to suffer, but suffering makes one unhappy. Therefore, to be happy, one must love or love to suffer or suffer from too much happiness.”

Woody Allen

5.1 Overview

In this chapter, we explore several ways in which well-established, efficient techniques from standard computational electromagnetism can be brought to bear on the problem of computing Casimir forces, in order to predict forces for arbitrary geometries and materials with arbitrary accuracy with no uncontrolled approximations [443]. Starting from the simplest, most direct approaches, we show that practical considerations naturally lead towards a particular method involving the integral of the Minkowski stress tensor by repeated evaluation of the imaginary-frequency Green’s function—a method previously developed for purely analytical calculations [129, 303, 304]. For our initial volume discretization with N degrees of freedom and an efficient iterative solver, this requires $O(N)$ storage and at best $O(N^{2-1/d})$ time in d dimensions. Furthermore, because evaluation of the Green’s function is such a standard problem in classical computational electromagnetism, it will be possible to exploit many developments in fast solvers, based on finite-element [91, 237, 548, 596],

spectral [50, 91], or boundary-element methods [90, 91, 191, 237, 430, 548]. As we argue below, an implementation using boundary-element methods (BEM) should attain nearly $O(N \log N)$ time. To illustrate the method, however, our initial implementation is based on the much simpler finite-difference frequency-domain method (FDFD) [93] with either a conjugate-gradient [18] or direct [25] solver, as described below. However, evaluation of the Green's function is so standard that many other sophisticated techniques are immediately applicable, and we discuss what techniques are likely to be optimal for large, complicated geometries. Our approach is tested for geometries with known solutions, and subsequently applied to more complicated geometries leading to surprising nonmonotonic effects in chapter 7. This method is also capable of handling dispersive dielectric materials, and was the first such method to depart from idealized metals. The key advantage of exploiting standard computational approaches is not merely that existing code, error analyzes, and other experience can be applied to the Casimir problem, but also that these methods have been proved to scale to large three-dimensional problems, which previously seemed out of reach of exact methods for the Casimir force.

In the following sections, we describe the step-by-step conceptual development of our finite-difference frequency-domain computational method. Our purpose here is to start back at the beginning, with the earliest theoretical descriptions of the Casimir force, and analyze these formulations from the point of view of their suitability for purely numerical calculations. Although the final technique, in terms of the stress tensor integrated over space and imaginary frequency, can be viewed as a numerical implementation of a textbook result due to Dzyaloshinskiĭ *et al.* [129, 304, 408], it is illustrative to derive it as the culmination of a sequence of simpler approaches, in order to show how it circumvents a number of numerical obstacles that would hinder a more direct method. To begin with, we illustrate the methods using the well-known case of parallel plates where they can be compared to analytical expressions, but a more rigorous test is subsequently provided by the situation of a cylinder and plate, recently solved numerically [141]. We conclude by analyzing the scaling of the method compared to previous approaches and discussing the application of more sophisticated finite-element and boundary-element techniques.

5.2 Mode Summation: A Simple Approach

Perhaps the simplest approach to computing the Casimir force is to think of it as the derivative of the zero-point energy U expressed as a sum of ground-state photon energies $\hbar\omega/2$. For each photon mode with frequency ω , the zero-point energy is $\hbar\omega/2$, and thus the total Casimir energy, at least for non-dissipative systems where ω is real [529], is formally given by the sum over all modes [77, 352]:

$$U = \sum_{\omega} \frac{1}{2} \hbar\omega \quad (5.1)$$

This sum is formally infinite, because the classical harmonic modes ω have unbounded frequencies. There is some controversy over the physical interpretation of this divergence [358], but in practice it is removed by regularizing the sum in some fashion, for example multiplying by $e^{-s\omega}$ for $s > 0$, and taking $s \rightarrow 0$ only *after* the sum is differentiated to obtain the force $F = -dU/da$ between two bodies with separation a [342]. This approach, which was an early method to analytically compute the force between perfect-metal plates [77], might seem to provide the most direct computational method as well. After all, the computation of electromagnetic eigenmodes is routine even in very complicated geometries, and efficient methods are known [91, 237, 249, 548]. Unfortunately, it turns out not to be practical for this problem (except in one-dimensional geometries [539]), as explained below, but the *reason* why it is impractical points the way to more efficient methods.

To illustrate the difficulty in directly evaluating Eq. (5.1), let us consider the simplest one-dimensional geometry: two parallel perfect-metal plates, separated by a distance a , in which case one can predict analytically an attractive force $F = \pi\hbar c/24a^2$ [358]. Ignoring this analytical result, let us apply a numerical method that, conceptually, we could apply to an arbitrary geometry:

1. Discretize space with resolution Δx using a finite-difference approximation, with space truncated to a finite computational cell (e.g. with periodic boundaries).
2. Solve numerically for the eigenmode frequencies ω and sum to obtain $U(a)$.
3. Shift one body (one plate) by one pixel Δx and thus compute $U(a + \Delta x)$.
4. Obtain the force $F \approx -[U(a + \Delta x) - U(a)]/\Delta x$.

Note that this method automatically provides its own regularization: the number of modes ω in a discretized computational cell is finite (the frequencies are bounded by the Nyquist limit), and hence U is finite for $\Delta x > 0$. The periodic boundaries will lead to artificial “wrap-around” forces, but since Casimir forces decay rapidly with distance, the contribution of these forces can be made negligible for a sufficiently large computational cell.




Figure 5-1: Schematic of one dimensional geometry, showing two 1d metal plates separated by a distance a , embedded in a computational cell of length $L = a + 4a$, with periodic boundary conditions, and resolution Δx .

This method, for the one-dimensional parallel-plate geometry, is illustrated in Fig. 5-1. Here, we have two plates with separation a and an overall computational cell size $L = 5a$ (which will contribute an erroneous wrap-around force 1/16 of the physical force). Maxwell’s equations, in one dimension, can be written as the scalar eigenproblem $\nabla^2 E_z = \omega^2 E_z$ (in $c = 1$ units), which is discretized to

$$\frac{E_{n+1} - 2E_n + E_{n-1}}{\Delta x^2} = \omega_n^2 E_n \quad (5.2)$$

in a center-difference approximation for $E_z(n\Delta x) = E_n$. For two metal plates with separation d and Dirichlet boundary condition $E_z = 0$, the discrete eigenvalues ω_n can be found analytically:

$$\omega_n = \frac{2}{\Delta x} \sin\left(\frac{n\pi\Delta x}{2d}\right), \quad (5.3)$$

for $n = 0, \dots, d/\Delta x$. The energy U is then given by summing ω_n in Eq. (5.1) for $d = a$ and $d = L - a$, and the force F by the discrete derivative of U as above.

Applying this procedure numerically for $\Delta x = 0.05a$, one obtains the correct force to within 5% (and to any desired accuracy by increasing L and decreasing Δx), so at first glance it may seem that the method is successful. However, its impracticality is revealed if we examine the contribution of each frequency ω_n to the force.

The top panel of Fig. 5-2 shows the contribution $\hbar\Delta\omega/2\Delta x$ [$\Delta\omega = \omega(a + \Delta x) - \omega(a)$] to the force summation, and the bottom panel shows the corresponding partial sum (for




Figure 5-2: (Left:) Plot of force summand, or spectral density, $\hbar\Delta\omega/2\Delta x$ vs. ω for 1d parallel metal plates from Fig. 5-1. (Right:) Plot of force partial sum $\sum^{\omega} \hbar\Delta\omega/2\Delta x$ vs. ω .

frequencies up to ω). We see that *every* frequency (of the regularized/finite-resolution problem) makes a non-negligible contribution to the force, and the summation is of a wildly-oscillating quantity that leaves a tiny remainder at the end. The reason for these oscillations is quite simple: as a is increased, the frequencies on the $4a$ side of the plates increase slowly, while the smaller number of frequencies on the a side of the plate decrease more rapidly, and these lead to the positive and negative contributions in the top panel of Fig. 5-2, respectively.

These two features, which are an intrinsic feature not limited to this particular discretized geometry [157], combine to make this method impractical in higher-dimensional structures. Because *every* frequency contributes to the force, in a numerical method one must compute every eigenvalue of the Maxwell eigenproblem. In one dimension, that is not so bad, but in general if there are N degrees of freedom (N grid points), then computing every eigenvalue of an $N \times N$ matrix requires $O(N^2)$ storage and $O(N^3)$ time. This is impractical in three dimensions where N may be in the millions. Furthermore, the wild oscillations of the summand imply that the eigenvalues must be computed quite accurately, and may exacerbate numerical difficulties in larger problems.

However, these undesirable features are avoidable, because we have not yet exploited a key property of Maxwell's equations: causality. If we ignore the causality constraint, then the oscillatory spectrum would be an observable effect: one would simply employ a material that is a perfect metal in some frequency range and transparent otherwise, in order to obtain the force spectrum integrated only in that range. Such a material, however, would violate

the Kramers-Kronig constraints that follow from causality considerations [231]. Thus, we are motivated to exploit causality in some fashion to avoid the oscillatory spectrum ¹

5.3 Wick Rotation and Energy Density

In order to exploit causality, we will rewrite Eq. (5.1) in terms of the electromagnetic Green's function via an integral over the density of states. Causality implies that the Green's function has no poles in the upper-half plane, so one can perform a contour integration, or *Wick rotation*, to transform the sum over real frequencies into an integral along the imaginary-frequency axis. The result of this standard trick turns out to be a well-known expression: an integral of the mean electromagnetic energy density, evaluated by the fluctuation-dissipation theorem using the temperature (Matsubara) Green's function. Again, we focus on the method's suitability as a purely *numerical* approach, for arbitrary geometries, and we will find that it still leaves something to be desired.

First, we can express the zero-point energy of Eq. (5.1) as an integral over the local density of states $D(\mathbf{x}, \omega)$:

$$U = \frac{\hbar}{2} \int_0^\infty d\omega \int \omega D(\mathbf{x}, \omega) d^3\mathbf{x}. \quad (5.4)$$

Since we are solving for the eigenstates of Maxwell's equations $(\nabla \times \nabla \times -\omega^2 \epsilon)\mathbf{E} = 0$, the local density of states $D(\mathbf{x}, \omega)$ can be expressed in terms of the Green's tensor G_{jk} [286]:

$$D(\mathbf{x}, \omega) = \frac{1}{\pi} \frac{d(\omega^2 \epsilon)}{d\omega} \sum_{k=1}^3 \text{Im} \langle \mathbf{x}; \hat{\mathbf{e}}_k | \frac{1}{\nabla \times \nabla \times -\omega^2 \epsilon + i0^+} | \mathbf{x}; \hat{\mathbf{e}}_k \rangle = \frac{1}{\pi} \frac{d(\omega^2 \epsilon)}{d\omega} \text{Im tr } G(\omega; \mathbf{x} - \mathbf{x})$$

where G_{jk} solves: $(\nabla \times \nabla \times -\omega^2 \epsilon)\mathbf{G}_k(\omega; \mathbf{x} - \mathbf{x}') = \delta^3(\mathbf{x} - \mathbf{x}')\hat{\mathbf{e}}_k$, with $\hat{\mathbf{e}}_k$ denoting the unit vector in the k th direction. For non-dissipative systems in which Eq. (5.1) is valid, ϵ is real and we can therefore pull the Im outside of the integral. (The generalization to dissipative materials is discussed below.)

Furthermore, we know from causality requirements that the Green's function has no poles in the upper half plane in ω -space [231, 352]. This means that one can perform a

¹Unfortunately, abandoning the real-frequency formulation of the Casimir problem is not entirely fruitful. On the one hand, we are avoiding a number of numerical difficulties [339, 340]. On the other hand, this requires researchers to abandon some of the basic intuition associated with the real-frequency domain, such as the presence of geometric resonances and other important effects from classical electromagnetism [229].

contour integration to relate $\int_0^\infty dw$ to the integral $\int_0^\infty dw$ along the imaginary-frequency axis $\omega = iw$, also known as a Wick rotation. We therefore obtain:

$$U = \frac{\hbar}{2\pi} \int_0^\infty dw \int w \frac{d[w^2 \varepsilon(iw)]}{dw} \text{tr} G(iw; \mathbf{x} - \mathbf{x}') d^3 \mathbf{x}, \quad (5.5)$$

where the *new* problem to be solved is that of finding the solutions to the imaginary-time Green's function ($c = 1$ units):

$$[\nabla \times \nabla \times + w^2 \varepsilon(iw, \mathbf{x})] \mathbf{G}_k(iw; \mathbf{x} - \mathbf{x}') = \delta^3(\mathbf{x} - \mathbf{x}') \hat{\mathbf{e}}_k. \quad (5.6)$$

As usual, this is formally infinite, because the Green's function is singular at $\mathbf{x} = \mathbf{x}'$, but one typically regularizes the problem by subtracting the vacuum Green's function, which removes the singularity without changing the net force.

Equation (5.5) is not a new result, nor is it limited to non-dissipative materials (unlike our derivation) [529]. In fact, it is equivalent to the mean energy in the fluctuating electromagnetic field, derived from the fluctuation-dissipation theorem via the temperature Green's functions [304]. Our purpose in deriving it this way is to emphasize the connection to the simplistic approach of Eq. (5.1). In particular, the mean energy in the electromagnetic fields (for the case of non-magnetic materials $\mu = 1$) is given by [231]:

$$\begin{aligned} U &= \int_0^\infty dw \int \left[\frac{d(w\varepsilon)}{dw} \langle \mathbf{E}^2 \rangle_w + \langle \mathbf{H}^2 \rangle_w \right] d^3 \mathbf{x} \\ &= \int_0^\infty dw \int \frac{d(w^2 \varepsilon)}{dw} \langle \mathbf{E}^2 \rangle_w d^3 \mathbf{x}, \end{aligned} \quad (5.7)$$

using the fact that $\int \varepsilon \langle \mathbf{E}^2 \rangle_w = \int \langle \mathbf{H}^2 \rangle_w$. Here, the key point is that the mean values of the fluctuating fields are given, via the fluctuation-dissipation theorem, in terms of the imaginary-frequency Green's function:

$$\langle E_j(\mathbf{x}) E_k(\mathbf{x}') \rangle_w = \hbar w^2 G_{jk}(iw; \mathbf{x} - \mathbf{x}'). \quad (5.8)$$

where the dyadic Green's function G_{ij} solves Eq. (5.6) and obeys the usual boundary conditions on the electric field from classical electromagnetism [231]. Substituting Eq. (5.8) into Eq. (5.7), one recovers Eq. (5.5).

From a computational perspective, the imaginary-frequency integral of Eq. (5.5) turns

out to be greatly superior to the real-frequency summation of Eq. (5.1), for two reasons. First, while every real frequency ω contributed to the force $-dU/da$, the same is not true for the derivative of the imaginary-frequency integrand. In particular, as discussed below, the force integrand in imaginary frequencies turns out to be a smooth, non-oscillatory, strongly peaked function of w , meaning that one can integrate it via a smooth-quadrature method that evaluates the integrand at only a small number of w values. Second, the imaginary-frequency Green's function turns out to be quite easy to obtain by relatively standard methods, including for dissipative systems (where obtaining the eigenmodes is harder because it involves a non-Hermitian eigenproblem). These two favorable features are closely related.

Consider Eq. (5.6) for the imaginary-frequency Green's function. Unlike its real-frequency counterpart, the linear operator on the left-hand side of this equation is *real-symmetric* and *positive-definite* (for $w > 0$). This is true even for dissipative materials where $\varepsilon(\omega)$ is complex, since causality requirements imply that $\varepsilon(iw)$ is purely real and positive (in the absence of gain) [231, 352]. For one thing, this implies that the most powerful numerical methods are applicable to solving the linear system of Eq. (5.6)—many of these methods (e.g. the conjugate-gradient method) are restricted to Hermitian positive-definite operators [18]. Also, the resulting Green's function is particularly well-behaved: it is exponentially decaying and non-oscillatory. This transforms the highly oscillatory real- ω force integrand into a mostly non-oscillatory integrand, and also makes the force integrand exponentially decaying for large w (for large w , the interactions between bodies become exponentially small). (In addition, as we will discuss in Sec. 5.8, the exponentially decaying Green's function is especially favorable for boundary-element numerical methods.)

Again, considering the simplest possible finite-difference scheme, this leads us to the following numerical algorithm to compute the force:

1. For a given imaginary frequency w :
 - (a) For every grid point \mathbf{x} , solve Eq. (5.6) for each polarization k to obtain $G_{kk}(iw; \mathbf{x} - \mathbf{x})$.
 - (b) Sum over \mathbf{x} to compute the spatial integral in Eq. (5.5).
 - (c) Repeat the above for the body shifted by one pixel Δx and subtract to obtain the force integrand at w .




Figure 5-3: Plot of Casimir force integrand dU/da between two 1d parallel plates separated by a distance $a = 1$ versus imaginary frequency $w = \text{Im } \omega$, using the method of Sec. 5.3.

2. Employ a smooth quadrature scheme to integrate the above function over w to obtain the force.

Again, the spatial discretization provides its own regularization (G_{kk} is finite), and thus no additional regularization is required (the contribution of the vacuum Green's function to the net force is zero). Again, one can truncate the computational cell in a number of ways, for example with periodic boundaries, and the artifacts thereby introduced will decrease rapidly with cell size. Again, there are also many other ways that one could potentially solve for G_{kk} besides a finite-difference approximation, but we will delay discussion of those techniques until we have formulated the final method in the next section.

The above procedure can again be applied to the one-dimensional problem of the force between two plates, as in Fig. 5-2, to illustrate its basic features. The result is shown in Fig. 5-3, plotting the force integrand as a function of imaginary frequency w , and the difference from Fig. 5-2 is striking. The integrated force is the same (correct) result as before. Unlike Fig. 5-2, the force integrand has no sign oscillations, is exponentially decaying for large w , and is strongly peaked around a characteristic $w = 2\pi c/a$, corresponding to a "wavelength" of a (the separation). These features imply that the force can be accurately integrated by an adaptive Gauss-Kronrod quadrature scheme [407] using at most a few dozen w points.

Although this method is much more efficient than the one described in the previous section, and is potentially practical at least in two dimensions, it still has some undesirable features. Suppose that we have N grid points in our discretized operator (N may be very

large in 3d). Even if we have an ideal iterative solver for the sparse linear system of Eq. (5.6), such as an ideal multigrid solver [532, 596], each evaluation of the Green's function takes at best $O(N)$ time with $O(N)$ storage. However, we must evaluate the Green's function $3N$ times in order to perform the spatial integration, resulting in $O(N^2)$ complexity. [A recent implementation of this approach was demonstrated in Ref. [403, 404] using a direct solver [25] to evaluate Eq. (5.6).] As is discussed in the next section, one can do much better than this by using the *stress tensor* instead of the energy density. In fact, as is discussed in Sec. 5.8, it should ultimately be possible to obtain the force with nearly $O(N \log N)$ work using a boundary-element method to compute the stress tensor, in which N is only the number of degrees of freedom required to represent the *interfaces* between materials.

5.4 Stress-Tensor Computational Approach

After having analyzed the feasibility of several techniques to solve the Casimir problem through the lens of numerical electromagnetism, we are ready to appreciate and explore the most feasible of the methods thus far presented: an approach based on the Maxwell stress tensor.

As derived by Dzyaloshinskiĭ *et al.* [129, 304, 408], the net Casimir force on a body can be expressed as an integral over any closed surface around the body of the mean electromagnetic stress tensor $\langle \mathbf{T}_{ij} \rangle$, integrated over all frequencies. Again, using the same arguments as above, it is computationally convenient to perform a Wick rotation, expressing the net force as an integral over imaginary frequencies $\omega = iw$. (The original derivation used imaginary frequencies to start with, via the temperature Green's function, but the result is equivalent to a Wick rotation of the real-frequency expression. A related analytical treatment, but using purely real ω and therefore unsuitable for numerical computation because of the oscillations discussed above, has also been examined [339, 340].) The resulting net force is:

$$\mathbf{F} = \int_0^\infty \frac{dw}{\pi} \oint_{\text{surface}} \langle \mathbf{T}(\mathbf{r}, iw) \rangle \cdot d\mathbf{A}. \quad (5.9)$$

In two or one dimensions, one or two of the spatial integrals are omitted, respectively, but the result still has the units of force; the change in dimensions of $d\mathbf{A}$ is balanced by a change in the dimensions of the delta function in Eq. (5.6). On the other hand, for a 3d structure

that is z -invariant [a constant 2d cross-section $\varepsilon(x, y)$] or yz -invariant [a univariate $\varepsilon(x)$], the integrals over the invariant directions are replaced by integrals over the corresponding wavevectors, resulting in a net force per unit length or per unit area, respectively. (These wavevector integrals are discussed in more detail in Sec. 5.6.) The stress tensor is defined as usual by:

$$\begin{aligned} \langle \mathbf{T}_{ij}(\mathbf{r}, iw) \rangle &= \mu(\mathbf{r}, iw) \left[\langle H_i(\mathbf{r}) H_j(\mathbf{r}) \rangle - \frac{1}{2} \delta_{ij} \sum_k \langle H_k(\mathbf{r}) H_k(\mathbf{r}) \rangle \right] \\ &+ \varepsilon(\mathbf{r}, iw) \left[\langle E_i(\mathbf{r}) E_j(\mathbf{r}) \rangle - \frac{1}{2} \delta_{ij} \sum_k \langle E_k(\mathbf{r}) E_k(\mathbf{r}) \rangle \right], \end{aligned}$$

where μ and ε are the relative permeability and permittivity, respectively, although in most cases we set $\mu = 1$ for simplicity (since most materials have negligible magnetic response at short wavelengths, and in any case the stress tensor is normally evaluated over a surface lying in vacuum). As before, the connection to quantum mechanics arises from the correlation functions of the fluctuating fields, given via the fluctuation-dissipation theorem in terms of the imaginary- ω dyadic Green's function. The correlation function for the electric field $\langle E_i E_j \rangle$ is given in Eq. (5.8). In this case, however, we also need the magnetic-field correlation functions, which can be obtained by differentiating the electric-field Green's function:

$$\langle H_i(\mathbf{r}) H_j(\mathbf{r}') \rangle = -\hbar (\nabla \times)_{i\ell} (\nabla' \times)_{jm} G_{\ell m}(iw; \mathbf{r} - \mathbf{r}'), \quad (5.10)$$

(Alternatively, the $\langle H_i H_j \rangle$ correlation function can be computed from the magnetic Green's function, which is the magnetic field in response to a given magnetic-dipole current.) The above expressions are given at zero temperature T ; the nonzero-temperature force is found by changing $\int dw$ in Eq. (5.9) into a discrete summation [280, 304]. Although the Green's function (and thus \mathbf{T}) is formally infinite at $\mathbf{r} = \mathbf{r}'$, this divergence is conventionally removed by subtracting the vacuum Green's function; in a numerical method with discretized space, as below, there is no divergence and no additional regularization is required. (The vacuum Green's function gives zero net contribution to the $d\mathbf{A}$ integral, and therefore need not be removed as long as the integrand is finite.)

Historically, this stress-tensor expression was used to derive the standard Lifshitz formula for parallel plates, where G_{ij} is known analytically. Its adaptability and suitability as a purely computational method does not seem to have been recognized, however. As in

the previous section, the method involves computing the Green's function for many imaginary frequencies w and spatial points \mathbf{x} , integrated over w and \mathbf{x} . However, a quick glance at Eq. (5.9) will suggest at least two obvious computational advantages compared to the method discussed in Sec. 5.3. First, in framing the problem in terms of the stress tensor, we have reduced the spatial integral over the whole volume (Eq. (5.7)) to a surface integral around the body of interest. This implies that, for a d -dimensional geometry, the computational effort due to spatial integration is reduced from $O(N^2)$ to $O(N^{2-1/d})$. Second, the force is now given *directly* in terms of the dyadic Green's function (via the stress tensor), rather than its derivative, which avoids another layer of computation. Moreover, although our derivation is only valid when the stress tensor is evaluated at points within lossless dielectrics (regardless of whether the bodies themselves are dissipative), one can also extend it for evaluation in absorbing media [408]. However, the case discussed above (bodies separated by vacuum) is the most common.

So far, we have presented the step-by-step development of an efficient approach to computing Casimir forces. In what follows, we illustrate our new approach using a proof-of-concept finite-difference implementation, and present results that demonstrate its flexibility and utility.

5.5 The Finite-Difference Method




Figure 5-4: Schematic illustration of the finite-difference Yee grid: the vector components of the electric \mathbf{E} and magnetic \mathbf{H} fields are spatially staggered about a rectangular unit cell of a Cartesian computational grid so that each \mathbf{E} -field vector component is located midway between a pair of \mathbf{H} -field vector components, and vice versa. The electric field components form the edges of the cube, and the magnetic field components form the normals to the faces of the cube.

At this point, all that remains is the numerical computation of the Green's function

via Eq. (5.6) for an imaginary frequency $\omega = iw$. This is simply the inversion of a linear operator $[\nabla \times \nabla \times + w^2 \epsilon(\mathbf{r}, iw)]$ that has the convenient properties of being real-symmetric and positive-definite, as stated above. Almost any technique developed for computational electromagnetism is applicable here, modified to operate at an imaginary frequency. To illustrate our approach, we used a very simple, yet extremely general, method: finite-difference frequency-domain (FDFD) discretization of Eq. (5.6) in a staggered Yee grid [93], shown schematically in Fig. 5-4, which we then invert by a conjugate-gradient method [18]. Although the Yee grid in principle allows second-order-accurate finite-difference approximations, unfortunately the whole scheme becomes only first-order-accurate once a discontinuous dielectric function ϵ is included. (There are ways to treat interfaces more accurately [150], but we did not implement them here.) Moreover, a very high resolution is often required to resolve the stress tensor close to a dielectric boundary due to the Green's function divergences as a boundary is approached [191]. Despite its shortcomings, however, we found FDFD to be sufficient to obtain accurate results (to within a few percent in a reasonable time) for two-dimensional, and three-dimensional z -invariant, geometries. The key advantage of FDFD is its flexibility: with very little effort, we were able to implement support for arbitrary geometric shapes and arbitrary materials (both perfect metals and dispersive/dissipative dielectrics).

Again choosing the simplest possible approach, we apply periodic boundary conditions at the edges of the computational cell, which are accurate as long as the boundaries are sufficiently far compared to the separation between the interacting bodies. That is, the periodicity leads to artificial “wrap-around” forces that decay rapidly with cell size L (at least as $1/L^3$ in 2d); we chose cell sizes large enough to make these contributions negligible ($< 1\%$).




Figure 5-5: Schematic illustration of a possible contour around a body; the force on the body is given by an integral of the stress tensor around this contour.

The computational process (using a simple finite-difference scheme) goes as follows:

1. Pick a contour/surface around the body of interest, as in Fig. 5-5 (which will typically *not* coincide with the boundary of the body).
2. For a given frequency w :
 - (a) For every grid point \mathbf{x} on the discretized contour/surface, solve Eq. (5.6) for each polarization k to obtain $G_{jk}(iw; \mathbf{x} - \mathbf{x})$.
 - (b) Integrate the resulting stress tensor T_{jk} over the surface, as in Eq. (5.9).
3. Integrate the above function over w to obtain the force; since the integrand is a smooth function of w , an efficient adaptive quadrature scheme can be employed [407].

Although this scheme does not require any additional regularization (the integrand is finite for a finite resolution, and the integral of the vacuum stress tensor over the contour is zero), we have found that numerical convergence can be accelerated by subtracting the stress-tensor integral over the isolated bodies. For example, in the schematic of Fig. 5-5, we would first compute the stress-tensor integral for the two bodies as shown, then subtract the integral of the stress tensor over the *same* surface with one body removed, and then subtract again for the stress tensor with the other body removed. Of course, these subtracted quantities are zero in the limit of infinite spatial resolution—there is no net force on an isolated body. However, at a finite resolution the discretization error at the interface between two materials can lead to a finite force that vanishes as resolution is increased. By subtracting this error term from the force, we find that the numerical error is typically reduced by an order of magnitude or so. We emphasize, however, that this is merely an optimization—even without subtraction, the force converges to the correct result, and merely requires a somewhat higher resolution.

5.6 Translation-Invariant Structures

It is common to solve for the Casimir force between bodies that are translation-invariant in one or more directions; for example, between a cylinder and a plate [141] that are invariant in the z direction. More generally, one might consider structures that are periodic in some direction with a non-zero period Λ , where $\Lambda \rightarrow 0$ corresponds to translation invariance. Intuitively, in such cases one need only perform computations in the unit cell of the periodicity, and the spatial integration in the invariant direction(s) is replaced by integration over

a wavevector \mathbf{k} from Bloch's theorem [71, 239]. Although special cases of this familiar idea are well known in Casimir computations [71, 72], here we provide a review of this approach for an arbitrary periodicity in the context of stress-tensor computational methods. We also mention a useful optimization for the common special case of perfect-metallic z -invariant structures.

Let us consider a single direction of periodicity: suppose that the structure is periodic in z with period Λ . In this case, it is natural to choose a surface for our stress-tensor integral that is also periodic in z . For example, imagine that Fig. 5-5 depicts a two-dimensional (xy) cross-section of a z -invariant structure, and the dashed line depicts a cross-section of the corresponding z -invariant stress-tensor surface. Because the total force is infinite, the quantity of interest is the force per unit z . It is convenient to consider the net force \mathbf{F} from a finite length $L = N\Lambda$ with periodic boundaries, and obtain the force per unit length as $\lim_{N \rightarrow \infty} \mathbf{F}/L$. Naively, \mathbf{F}/L can be written directly via Eq. (5.9), where we break the integral over z into a summation over the unit cells:

$$\frac{\mathbf{F}}{L} = \frac{1}{L} \int_0^\infty \frac{dw}{\pi} \sum_{n=0}^{N-1} \iint \mathbf{T}(iw; \mathbf{r} - n\Lambda\hat{\mathbf{z}}) \cdot d\mathbf{A}, \quad (5.11)$$

$$= \frac{1}{\Lambda} \int_0^\infty \frac{dw}{\pi} \iint \mathbf{T}(iw; \mathbf{r}) \cdot d\mathbf{A}, \quad (5.12)$$

where the surface integral is over the portion of the surface lying in the unit cell only, and in the second line we have used the fact that the stress tensor is periodic. This expression is inconvenient, however, because the direct evaluation of $\mathbf{T}(iw, \mathbf{r})$ requires the response to a *single* point source in the large- L structure, and a single point source does *not* produce a periodic field (requiring a full three-dimensional calculation even for a z -invariant structure). Rather, one would like to consider the field in response to *periodic* point sources, which produce a periodic field that can be treated by a small computational cell with periodic boundary conditions. This is accomplished by Fourier-transforming the expressions in Eq. (5.11) and taking the $N \rightarrow \infty$ limit, as described in detail in the following subsection. The resulting force per unit length is:

$$\frac{\mathbf{F}}{L} = \frac{1}{\Lambda} \int_0^\infty \frac{dw}{\pi} \int_{-\pi/\Lambda}^{\pi/\Lambda} \frac{dk_z}{2\pi} \iint \mathbf{T}(iw, k_z; \mathbf{r}) \cdot d\mathbf{A}, \quad (5.13)$$

where the surface integral is still over the portion of the surface lying in the unit cell.

Here, $\mathbf{T}(iw, k_z; \mathbf{r})$ denotes the stress tensor computed from the Green's functions for Bloch-periodic boundaries—that is, from the fields in response to a periodic set of point-dipole sources with phase $e^{ik_z \Lambda n}$ in the n th unit cell. This stress tensor can be computed using a computational cell that is only one unit cell in the z direction, e.g. by a two-dimensional computational cell for a z -invariant structure. [Equivalently, $\mathbf{T}(iw, k_z; \mathbf{r})$ could instead be computed from the Green's function for ordinary periodic boundaries, but with ∇ replaced by $\nabla + ik_z \hat{\mathbf{z}}$ [239].] Just as for w , the stress tensor is a smooth function of k_z and therefore the k_z integral can be computed by an efficient quadrature scheme (e.g. Gaussian quadrature).

If the structure is periodic (or invariant) in more than one direction, one simply repeats the above procedure: for each periodic direction, we only consider the portion of the stress-tensor integral in the unit cell, with Bloch-periodic boundary conditions, and integrate over the corresponding Bloch wavevector component. Also, by symmetry one only need integrate over the irreducible Brillouin zone of the structure [239], e.g. in one dimension (where time-reversal symmetry normally equates k_z and $-k_z$) the $\int_{-\pi/\Lambda}^{\pi/\Lambda} dk_z$ integral can be replaced by $2 \int_0^{\pi/\Lambda} dk_z$.

For the common case of a z -invariant perfect-metal structure (i.e. one has a homogeneous ε surrounded by perfect metal), there are several important simplifications. First, the (k_z, w) Green's function is *exactly* the same as the $(0, \sqrt{k_z^2 + w^2})$ solution [141]. Therefore, we need only compute the $k_z = 0$ solutions at each w , and weight the dw integrand by a factor of πw (the circumference of a semi-circle of radius w) divided by the 2π that would appear in the dk_z integral. (For two directions of translational symmetry, one would weigh the integral by the area of a hemisphere of radius w , and in general, in d dimensions, by the area of a radius- w hypersphere.) Furthermore, at $k_z = 0$, the solutions can be divided into two scalar polarizations, TE ($\mathbf{E} \cdot \hat{\mathbf{z}} = 0$) and TM ($\mathbf{H} \cdot \hat{\mathbf{z}} = 0$) [231].

5.6.1 Periodic Green's Function Calculation

In what follows, we derive the force per unit length in translation-invariant structures in terms of an integral over the solutions of Bloch-periodic problems. This is not a new idea, but an explicit general derivation starting from the spatial integral of the stress tensor, including the case of finite periodicity, seems difficult to find in the literature. As given

above by Eq. (5.11), the force per unit length can be written as:

$$\frac{\mathbf{F}}{L} = \frac{1}{L} \int_0^\infty \frac{dw}{\pi} \sum_{n=0}^{N-1} \iint \mathbf{T}(iw; \mathbf{r} - n\Lambda\hat{\mathbf{z}}) \cdot d\mathbf{A}, \quad (5.14)$$

for a periodic structure in the z -direction with period Λ , and size $L = N\Lambda$ with periodic boundary conditions (ultimately, we will take the limit as $N \rightarrow \infty$). As in Eq. (5.10), the stress tensor is expressed via the Green's function, and therefore Eq. (5.14) can be decomposed into individual terms of the form:

$$\frac{1}{L} \int_0^\infty \frac{dw}{\pi} \sum_{n=0}^{N-1} \iint \mathbf{G}(iw; \mathbf{r}_n, \mathbf{r}_n) dA, \quad (5.15)$$

where $\mathbf{r}_n = \mathbf{r} - n\Lambda\hat{\mathbf{z}}$ and the Green's function is given as in Eq. (5.6) by the solution of:

$$\mathbf{G}_k(iw; \mathbf{x}, \mathbf{x}') = \hat{O}^{-1} \delta^3(\mathbf{x} - \mathbf{x}') \hat{\mathbf{e}}_k, \quad (5.16)$$

in which \hat{O} denotes the linear operator $\hat{O} = (\nabla \times \nabla \times + w^2 \varepsilon)$. In the following, we will focus on the periodic z direction and leave the x and y coordinates implicit for simplicity. That is, we will write e.g. $\delta(z - n\Lambda - z')$ instead of $\delta(\mathbf{r}_n - \mathbf{r}')$.

At this point, we can re-express the delta function over the periodic direction z in terms of the Fourier identity:

$$\delta(z - z' - n\Lambda) = \frac{1}{N} \sum_{\ell=0}^{N-1} \sum_{m=0}^{N-1} \delta(z - z' - \ell\Lambda - n\Lambda) e^{\frac{2\pi i}{N} \ell m}, \quad (5.17)$$

Substituting this into Eq. (5.15), we will move the \sum_m outside the \sum_n and consider the action of \hat{O}^{-1} on the remaining summation:

$$J_{n,m} = \sum_{\ell=0}^{N-1} \delta(z - z' - \ell\Lambda - n\Lambda) e^{\frac{2\pi i}{N} \ell m} = J_{0,m} e^{-\frac{2\pi i}{N} mn}, \quad (5.18)$$

where we have used the periodic boundary conditions in L to realize that $J_{n,m}$ is a cyclic shift of $J_{0,m}$ with a phase factor. Now, we must operate \hat{O}^{-1} on $J_{0,m} \hat{\mathbf{e}}_k$ and evaluate at $z' = z - n\Lambda$. However, this corresponds to finding the field from a Bloch-periodic current

source, and such a field is also Bloch-periodic. Therefore:

$$\left(\hat{O}^{-1}J_{0,m}\hat{\mathbf{e}}_k\right)\Big|_{z'=z-n\Lambda} = \left(\hat{O}^{-1}J_{0,m}\hat{\mathbf{e}}_k\right)\Big|_{z'=z} e^{\frac{2\pi i}{N}mn}, \quad (5.19)$$

At this point, we have completely eliminated the n -dependence from the evaluation of the Green's function $\hat{O}^{-1}J_{0,m}\hat{\mathbf{e}}_k$, and the phase factors from Eq. (5.18) and Eq. (5.19) cancel. The remaining summation \sum_n simply yields N , which cancels the $1/N$ factor from Eq. (5.17). Equation (5.15) therefore becomes:

$$\frac{1}{L} \int_0^\infty \frac{dw}{\pi} \sum_{m=0}^{N-1} \iint \mathbf{G}(iw, m; \mathbf{r}_n, \mathbf{r}_n) dA, \quad (5.20)$$

where $\mathbf{G}_k(iw, m) = \hat{O}^{-1}J_{0,m}\hat{\mathbf{e}}_k$, the field from a Bloch-periodic sum of delta-function sources. Finally, we can now take the limit $N \rightarrow \infty$ by turning \sum_m into an integral:

$$\lim_{N \rightarrow \infty} \frac{1}{L} \sum_{m=0}^{N-1} = \frac{1}{2\pi\Lambda} \int_{-\pi/\Lambda}^{\pi/\Lambda} dk_z, \quad (5.21)$$

where k_z is the Bloch wavevector ($k_z = 2\pi m/N$). We therefore obtain Eq. (5.13).

5.7 Numerical Validation

In the following sections we demonstrate our method's validity by checking it against known results for perfect metals, and in particular for the case of a cylinder adjacent to a plate.

5.7.1 Parallel Metal Plates

First, for comparison to the one-dimensional parallel-plate integrands plotted in Fig. 5-2 and Fig. 5-3, we plot the corresponding integrand for the stress tensor in Fig. 5-6. Again, this integrates to the correct force ($\hbar c\pi/24a^2$ for each polarization), but we note that the integrand is not identical to the integrand from the imaginary-frequency energy derivative.

Since the one-dimensional parallel-plate force is commonly derived from the Lifshitz formula, which in turn is derived from the stress tensor, this cannot be regarded as a rigorous validation of our method (except in the most basic sense of checks for bugs in our code).




Figure 5-6: Plot of Casimir force integrand $\iint \mathbf{T} \cdot d\mathbf{A}$ between two 1d parallel plates separated by a distance $a = 1$ versus imaginary frequency $w = \text{Im } \omega$, using the stress-tensor method of Sec. 5.4.

5.7.2 Cylinder and Plate

A more complicated geometry, consisting of a perfect metallic cylinder adjacent to a perfect metallic plate in three dimensions, was solved numerically by Ref. [141], to which our results are compared in Fig. 5-7. Ref. [141] used a specialized Fourier-Bessel basis specific to this cylindrical geometry, which should have exponential (spectral) convergence. Our use of a simple uniform grid was necessarily much less efficient, especially with the first-order accuracy, but was able to match the Ref. [141] results within $\sim 3\%$ using reasonable computational resources. A simple grid has the advantage of being very general, as illustrated below, but other general bases with much greater efficiency are possible using finite-element or boundary-element methods; the latter, in particular, could use a spectral Fourier basis similar to Ref. [141] and exploit a fast-multipole method or similar $O(N \log N)$ solver technique. This is discussed in greater detail by Sec. 5.8.

Also shown, in the inset of Fig. 5-7, is a plot of the interaction stress-tensor component $\Delta\langle T_{xx} \rangle$ at a typical imaginary frequency $w = 2\pi c/a$. By “interaction” stress tensor $\Delta\langle T_{ij} \rangle$, we mean the total $\langle T_{ij} \rangle$ of the full geometry minus the sum of the $\langle T_{ij} \rangle$ ’s computed for each body in isolation. Here, the stress tensors of the isolated cylinder and plate have been subtracted, giving us a way to visualize the force due to the interaction. As described further below, such stress plots reveal the spatial regions in which two bodies most strongly affect one another, and therefore reveal where a change of the geometry would have the most impact. (In contrast, Ref. [178] plots an interaction-energy density that does not directly




Figure 5-7: Casimir force between a 3d radius- R cylinder and a plate (inset), relative to the proximity-force approximation F_{PFA} , vs. normalized separation a/R . The solid lines are the Casimir force computed in Ref. [141] for TE (gray) and TM (blue) polarizations, along with results computed by our method with a simple finite-difference discretization (gray squares). Error bars were estimated for some data points by using computations at multiple spatial resolutions. Inset shows interaction stress tensor $\Delta\langle T_{xx} \rangle$ at a typical imaginary frequency $\omega = 2\pi c/a$, where red indicates attractive stress.

reveal the force, since the force requires the energy to be differentiated with respect to a . For example, Ref. [178]’s subtracted energy density apparently goes to zero as a metallic surface is approached, whereas the stress tensor cannot since the stress integration surface is arbitrary.)

5.8 Beyond Finite-Difference Methods

Above, we implemented the stress-tensor integration using a finite-difference frequency-domain approach to compute the Green’s function. While sufficient for a proof-of-concept implementation, one would like to use more sophisticated methods in order to explore complex geometries more quickly, especially in three dimensions. The primary drawbacks of the finite-difference scheme are threefold. First, the material discontinuities imply that the error converges only linearly with resolution, although there are techniques to improve this to quadratic convergence [150]. Second, non-uniform resolution would be desirable to handle small features, such as the narrow channels in the piston structure for small h . Third, the stress-tensor integrand is a discretized, non-smooth function of space, meaning that we must evaluate it at a number of grid points proportional to the resolution (or resolution squared, in three dimensions); in general, the integrand must therefore be evaluated $O(N^{(d-1)/d})$ times

for a $d-1$ dimensional surface in d dimensions, leading to at best $O(N^{2-1/d})$ complexity. To address these drawbacks, we consider two standard approaches to solving partial differential equations in a more efficient manner in complex geometries: finite-element and boundary-element methods.

Finite-element methods can employ a non-uniform volume discretization, via an unstructured mesh, that can both put more resolution where it is needed and conform to the material interfaces to obtain higher-order accuracy [91, 237, 548, 596]. However, from the perspective of Casimir-force calculations, finite-element methods seem to have two potential drawbacks. First, because space is still discretized, the stress tensor is again not a smooth function of space and its accurate integration requires that the Green's function be evaluated at many mesh points. Second, there may be a problem with regularization: although the Green's function does not diverge in discretized space, with a non-uniform resolution this regularization varies at different points. Unless there is a way to locally regularize the problem (subtracting a vacuum Green's function computed at the local spatial resolution), this may lead to unphysical, non-convergent forces.

Boundary-element methods (BEMs) involve a discretization in terms of unknowns only at the interfaces between different materials—these surface unknowns are coupled to one another via the (known) Green's functions of the homogeneous regions [90, 91, 191, 237, 430, 548]. They are thus ideal for open problems, in which the bodies are surrounded by infinite volumes of empty space, since those infinite regions are treated analytically. Using the fast-multipole method (FMM) or other fast integral-equation methods [91, 191, 237, 548], BEMs can solve for the surface unknowns, and hence the Green's function, in $O(N \log N)$ time for N discretized unknowns, multiplied by a number of iterations ($\ll N$) that depends on the condition number of the matrix and the preconditioning [18]. Furthermore, BEMs have two unique advantages when applied to the problem of Casimir forces. First, the regularization can be performed analytically: since BEMs express the Green's function as the sum of the vacuum Green's function plus a set of contributions from surface currents, the vacuum Green's function can be trivially subtracted analytically. Second, because space is not discretized, the stress tensor in a BEM will be a *smooth* (infinitely differentiable) function of space—this means that the spatial integral can be performed with exponentially-convergent smooth quadrature methods. Therefore, the number of times that the Green's function must be computed is determined only by the convergence of the smooth multidimensional

quadrature in $2 + 1$ dimensions (space + frequency), *independent* of N .

For these reasons, we believe that BEMs are in fact the most efficient methods to compute Casimir forces for complicated structures in three dimensions. Moreover, all that needs to be done is to take an existing BEM Green's function solver and change it to solve for the imaginary- ω Green's function. Because the imaginary- ω Green's functions are exponentially decaying (and approach the familiar Poisson kernel as $\omega \rightarrow 0$), such a fast solver should actually be simpler than the corresponding real- ω solver, nor are the singularities in the Green's function any worse. And, as mentioned previously, the resulting matrix equation is real-symmetric and positive-definite for imaginary ω , unlike the real- ω case, even for dissipative materials. A BEM implementation of the stress tensor method was developed for perfectly-metallic geometries [565, 566]. A generalization for dielectric objects is currently being developed by Reid. *et. al.* A BEM implementation based on the formalism of Ref. [137, 143, 426] was recently implemented and used to obtain accurate predictions of forces between complicated three-dimensional non-spherical dielectric objects with sharp corners [434]. The method is considerably more efficient than the simple finite-difference method describe here, allowing calculations of large three-dimensional objects using a single computer and only tens to hundreds of CPU hours.

5.8.1 Spectral and Monte-Carlo Approaches

In this section, we briefly compare the stress-tensor approach with other known exact numerical methods applicable to arbitrary geometries, focusing mostly on the computational aspects. In particular, we examine the methods of Emig *et al.* [136] and Gies *et al.* [177].

Emig's method, applicable to both separable and non-separable geometries and not limited to perfect metallic structures (although currently only demonstrated in perfect metallic separable geometries), involves a surface parametrization of the Green's function. Specifically, the Casimir energy is given in terms of an integral over imaginary frequencies of the change in the photon density of states (DOS) $\delta\rho(iw)$:

$$U = \int \frac{\hbar w}{2} \delta\rho(iw) dw, \quad (5.22)$$

similar to the expression in Eq. (5.4). However, the crucial aspect of this method lies in the

evaluation of the DOS, given by [136]:

$$\delta\rho(iw) = \frac{1}{\pi} \frac{\partial}{\partial w} \text{tr} \ln (M_\infty^{-1} M). \quad (5.23)$$

Here M is an $N \times N$ dense matrix, where N is the number of surface degrees of freedom, whose entries are in terms of the imaginary- ω vacuum Green's function $\mathbf{G}(iw; \mathbf{x}, \mathbf{x}')$ evaluated on the surface of each body. M_∞ is the same matrix for the case where the bodies are infinitely far apart. Thus, the trace in Eq. (5.23) is analogous to an integration over the surfaces of all the bodies. Inverting a dense matrix, multiplying two dense matrices, and taking the log of a dense matrix all require $O(N^3)$ time (for practical algorithms) and $O(N^2)$ storage [185].

We should comment however, that Emig's method is closely related to a boundary-element method (BEM) as discussed in the previous section. BEMs also involve parametrization in terms of surface degrees of freedom, which are also coupled in terms of vacuum Green's functions, leading to a dense matrix which must be inverted to compute the inhomogeneous Green's function. By recognizing this relationship, one should be able to exploit fast-multipole and similar $O(N \log N)$ techniques to accelerate Emig's method. In particular, computations such as $M_\infty^{-1} M$ involve the solution of N linear equations, each of which can employ iterative methods with $O(N)$ storage, and similar iterative methods may also be available for computing matrix logarithms [185]. At best, this leads to $O(N^2 \log N)$ time. However, this is still much less efficient than the stress-tensor BEM approach discussed in the previous section, because the latter requires only $\ll N$ linear equations to be solved (the number of linear equations to be solved is determined by a smooth spatial quadrature, independent of N).

A second exact computational method available is that of Ref. [177], based on a "world-line" approach. In this method, the Casimir energy is represented via a scalar field in a smooth background potential, and the effective action is obtained via a Feynman path integral over all proper time worldlines using a Monte-Carlo approximation. The method has only been formulated for the TM polarization with perfect-metal bodies, although preliminary generalizations have been suggested [180]. Specifically, the method expresses the Casimir energy between two bodies as the integral of a functional Θ_Σ over closed paths $\mathbf{x}(\tau)$

of length (proper time) T and center of mass \mathbf{y} :

$$U = -\frac{1}{8\pi^2} \int_0^\infty \frac{dT}{T^3} \int d^3\mathbf{y} \langle \Theta_\Sigma[\mathbf{x}(\tau)] \rangle_{\mathbf{x}} \quad (5.24)$$

where $\Theta_\Sigma[\mathbf{x}(\tau)]$ is a worldline functional (similar to a step function) defined to be $\Theta_\Sigma[\mathbf{x}(\tau)] = 1$ if the path $\mathbf{x}(\tau)$ intersects a surface and 0 otherwise [180]. This integration is performed by generating an ensemble of n_L random N -point paths $\mathbf{x}(\tau)$ and evaluating Θ_Σ for each one. The computational complexity $O(n_L \cdot N \cdot \#\mathbf{y} \cdot \#T)$ (where $\#\mathbf{y}$ and $\#T$ are the number of \mathbf{y} and T integrand evaluations, respectively), therefore, depends on quantities such as the precise statistical rate of convergence of this integral (error $\sim 1/\sqrt{n_L}$), which in turn depends on the geometry and on the manner in which the path ensembles are generated and integrated. Ref. [177] does not present a general analysis of these quantities, nor will we do so here. However, in the specific case of the force between a radius- R sphere and a plate separated by a distance a , Ref. [177] shows that $N \gg a^2/R^2$ for large a/R , corresponding to at least a quadratic scaling of the number of degrees of freedom and time with the relevant feature size a/R , even without the $O(n_L \cdot \#\mathbf{y} \cdot \#T)$ factors. In comparison, a BEM for the same geometry (exploiting the cylindrical symmetry to reduce it to a 2d problem similar to Ref. [177]) should require degrees of freedom N that never scale worse than linearly with the relevant lengthscale, and time that scales with $N \log N$ rather than N^2 . A general comparison seems difficult, however, and a detailed statistical study is outside the scope of this thesis.

5.9 Concluding Remarks

The general considerations involved in designing a purely computational method are often quite different from those involved in designing an analytical method. For this reason, we believe it is most fruitful to start back at the earliest possible formulations and proceed using the new computational perspective, rather than attempting to add more and more corrections to analytical methods for specific geometries. Moreover, since decades of research have gone into the development of numerical methods for classical electromagnetism, culminating in methods applicable to complicated inhomogeneous three-dimensional geometries, it is desirable to seek approaches for the Casimir force that exploit these developments. We believe that the stress-tensor approach, developed for analytical calculations several decades ago,

provides the ideal formulation for a computational approach exploiting standard numerical techniques.

The large and growing number of interesting applications of the Casimir effect, as well as the ongoing experimental work on non-standard geometries [74], provide an environment in which the generality and strengths of the stress-tensor method could be exploited. In chapter 7, we employ the stress-tensor approach to study Casimir forces and torques in more realistic, unusual structures. The extension of this work to study the *torque* on a body, instead of the net force, is in fact quite trivial. Classically, given the stress tensor \mathbf{T} , one can compute the torque by integrating $\mathbf{r} \times (\mathbf{T} \cdot d\mathbf{A})$ [285]. This has been exploited by several authors to compute classical electromagnetic torques [99, 311, 410]. Similarly, the Casimir torque can be obtained by using the mean $\langle \mathbf{T} \rangle$ from the fluctuation-dissipation theorem, just as for the net force [303, 452]. Therefore, one can compute torques by almost the same method as above, via repeated evaluations of the classical Green's function, as is done in chapter 8. Several authors have computed Casimir torques in parallel-plate and perfect-metallic wedge geometries [63, 373, 431, 484, 540]; other interesting structures include the Casimir “pendulum” [471] as well as corrugated surfaces [440], although these two structures have only been evaluated by methods with uncontrolled approximations.

Chapter 6

Casimir Forces in the Complex Time Domain

“Magnetism, as you recall from physics class, is a powerful force that causes certain items to be attracted to refrigerators.” *Dave Barry*

6.1 Overview

In the previous chapter, we described a frequency-domain method for computing Casimir forces in arbitrary geometries and materials, based on well-established classical electromagnetism techniques. Starting from the simplest, most direct approaches, we showed that practical considerations naturally lead towards a particular method involving the integral of the Minkowski stress tensor that requires repeated evaluation of the imaginary-frequency Green’s function. In this chapter, we describe two different methods for computing Casimir forces, both of which abandon the frequency domain altogether in favor of the time domain. As argued, there are many advantages to the time-domain perspective, most importantly, the fact that it makes it possible to obtain the entire frequency spectrum by iterating Maxwell’s equations (few time steps) in response to a single dipole source.

The first section reviews a well-known formulation of the Casimir-force problem, in which the Casimir force is expressed in terms of an integral of the classical Green’s function (GF) via the electromagnetic stress tensor (ST) and the fluctuation-dissipation theorem [304]. Although numerical constraints normally require integration over either the real or imaginary frequency axes, we abandon both of these standard choices, and instead consider mappings




Figure 6-1: Schematic illustration of correspondence between two methods of calculating Casimir forces. (Left): Numerical method requiring evaluation of the force integrand over the imaginary-frequency axis (or *some* suitable complex-frequency contour). (Right): Antenna (S -matrix) measurements of the electromagnetic response at tabletop (e.g. cm) lengthscales. Here, the effect of a contour deformation is achieved by a material deformation which corresponds to the presence of a conductive fluid between the objects.

over the general complex-frequency (ω) plane.

In Sec. 6.3, we describe a correspondence between the calculation of Casimir forces for vacuum-separated objects and a similar electromagnetic-force calculation in which the objects are instead separated by a conducting fluid, as illustrated in Fig. 6-8. The requirement that the geometry be mapped in this fashion is a practical consideration for any calculation based on the time domain (real-frequencies). In fact, it is the theoretical *equivalent* of a crucial and well-known technique for accurate numerical evaluation of Casimir forces, in which the force integrand is deformed via contour integration, and commonly evaluated over the imaginary-frequency axis [304, 443]. Our formulation circumvents difficulties with all previous expressions of Casimir forces in terms of frequency-integrals of classical Green's functions, which required either formally infinite bandwidth (when evaluated over real frequencies) or exponentially growing fields (when evaluated over imaginary frequencies)—instead, we exploit the moderate-bandwidth real-frequency response of a physical, dissipative system. This mathematical equivalence between complex contour mappings and physical dielectric deformations reveals new opportunities for the experimental and theoretical study of Casimir interactions. On the theoretical side, this correspondence makes it easier to understand Casimir systems from the perspective of conventional classical electromagnetism, based on real-frequency responses, in contrast to the standard point of view based on Wick rotations (imaginary frequencies). Furthermore, it has already led to a finite-difference time-domain numerical method for calculation of Casimir forces in arbitrary geometries and materials [344, 449], discussed in Sec. 6.4.1. A third application of this correspondence, described in Sec. 6.3.3 is a new technique for calculation of Casimir forces based on experimental S -matrix (microwave antenna) measurements in centimeter-scale

models. An experiment (centimeter-scale model) of this sort is not a Casimir “simulator,” in that one is not measuring forces, but rather a quantity that is mathematical related to the Casimir force—in this sense, it is an *analog computer*.

6.2 Revisiting the Stress Tensor Formulation

The Casimir force can be expressed as an integral over any closed surface S (enclosing a body) of the mean electromagnetic ST $\langle T_{ij}(\mathbf{r}, \omega) \rangle$ over all frequencies [304]. The mean ST is determined simply from the classical GF (the fields in response to current sources at a fixed frequency), a consequence of the fluctuation-dissipation theorem. In particular, the force in the i th direction is given by:

$$F_i = \int_0^\infty d\omega \oint_S \sum_j \langle T_{ij}(\mathbf{r}, \omega) \rangle dS_j, \quad (6.1)$$

The stress tensor is expressed in terms of correlation functions of the field operators $\langle E_i(\mathbf{r}, \omega) E_j(\mathbf{r}', \omega) \rangle$ and $\langle H_i(\mathbf{r}, \omega) H_j(\mathbf{r}', \omega) \rangle$:

$$\begin{aligned} \langle T_{ij}(\mathbf{r}, \omega) \rangle = \mu(\mathbf{r}, \omega) & \left[\langle H_i(\mathbf{r}) H_j(\mathbf{r}) \rangle_\omega - \frac{1}{2} \delta_{ij} \sum_k \langle H_k(\mathbf{r}) H_k(\mathbf{r}) \rangle_\omega \right] \\ & + \varepsilon(\mathbf{r}, \omega) \left[\langle E_i(\mathbf{r}) E_j(\mathbf{r}) \rangle_\omega - \frac{1}{2} \delta_{ij} \sum_k \langle E_k(\mathbf{r}) E_k(\mathbf{r}) \rangle_\omega \right], \end{aligned} \quad (6.2)$$

where both the electric and magnetic field correlation functions can be written as derivatives of a vector potential operator $\mathbf{A}^E(\mathbf{r}, \omega)$:

$$E_i(\mathbf{r}, \omega) = -i\omega A_i^E(\mathbf{r}, \omega) \quad (6.3)$$

$$\mu H_i(\mathbf{r}, \omega) = (\nabla \times)_{ij} A_j^E(\mathbf{r}, \omega) \quad (6.4)$$

We explicitly place a superscript on the vector potential in order to refer to our choice of gauge [Eqs. (6.3–6.4)], in which \mathbf{E} is obtained as a time-derivative of \mathbf{A} . The fluctuation-dissipation theorem relates the correlation function of \mathbf{A}^E to the photon Green’s function $G_{ij}^E(\omega; \mathbf{r}, \mathbf{r}')$:

$$\langle A_i^E(\mathbf{r}, \omega) A_j^E(\mathbf{r}', \omega) \rangle = -\frac{\hbar}{\pi} \text{Im} G_{ij}^E(\omega, \mathbf{r}, \mathbf{r}'), \quad (6.5)$$

where G_{ij}^E is the vector potential A_i^E in response to an electric dipole current \mathbf{J} along the $\hat{\mathbf{e}}_j$ direction:

$$\left[\nabla \times \frac{1}{\mu(\mathbf{r}, \omega)} \nabla \times - \omega^2 \varepsilon(\mathbf{r}, \omega) \right] \mathbf{G}_j^E(\omega; \mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}') \hat{\mathbf{e}}_j, \quad (6.6)$$

Given G_{ij}^E , one can use Eqs. (6.3–6.4) in conjunction with Eq. (6.5) to express the field correlation functions at points \mathbf{r} and \mathbf{r}' in terms of the photon Green's function:

$$\langle E_i(\mathbf{r}, \omega) E_j(\mathbf{r}', \omega) \rangle = \frac{\hbar}{\pi} \omega^2 \text{Im} G_{ij}^E(\omega; \mathbf{r}, \mathbf{r}') \quad (6.7)$$

$$\langle H_i(\mathbf{r}, \omega) H_j(\mathbf{r}', \omega) \rangle = -\frac{\hbar}{\pi} (\nabla \times)_{il} (\nabla' \times)_{jm} \text{Im} G_{lm}^E(\mathbf{r}, \mathbf{r}', \omega), \quad (6.8)$$

In order to find the force via Eq. (6.1), we must first compute $G_{ij}^E(\mathbf{r}, \mathbf{r}' = \mathbf{r}, \omega)$ at every \mathbf{r} on the surface of integration S , and for every ω [304]. It turns out, however, that this frequency integral is badly behaved from the perspective of numerical calculations (or experiments, below). In particular, the integral is formally infinite, requiring regularization, and highly oscillatory. Thus, as it is, Eq. (6.6) is of limited computational utility because it gives rise to an oscillatory integrand with non-negligible contributions at all frequencies, making numerical integration difficult [443]. However, the integral over ω can be re-expressed as the imaginary part of a contour integral of an analytic function by commuting the ω integration with the Im operator in Eqs. (6.7–6.8). Physical causality implies that there can be no poles in the integrand in the upper complex plane. The integral, considered as a complex contour integral, is then invariant if the contour of integration is deformed above the real frequency axis and into the first quadrant of the complex frequency plane, via some mapping $\omega \rightarrow \omega(\xi)$, conveniently parametrized by a real parameter ξ . This allows us to add a positive imaginary component to the frequency, which causes the force integrand to decay rapidly with increasing ξ [448]. In particular, upon deformation, Eqs. (6.7–6.8) are mapped to:

$$\langle E_i(\mathbf{r}, \omega) E_j(\mathbf{r}', \omega) \rangle = \frac{\hbar}{\pi} \omega^2 G_{ij}^E(\omega; \mathbf{r}, \mathbf{r}') \quad (6.9)$$

$$\langle H_i(\mathbf{r}, \omega) H_j(\mathbf{r}', \omega) \rangle = -\frac{\hbar}{\pi} (\nabla \times)_{il} (\nabla' \times)_{jm} G_{lm}^E(\mathbf{r}, \mathbf{r}', \omega), \quad (6.10)$$

where the deformed GF is given by:

$$\left[\nabla \times \frac{1}{\mu(\mathbf{r}, \omega)} \nabla \times - \omega^2(\xi) \varepsilon(\mathbf{r}, \omega) \right] \mathbf{G}_j^E(\xi; \mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}') \hat{\mathbf{e}}_j, \quad (6.11)$$

Similarly, the force in the i -th coordinate [Eq. (6.1)] becomes:

$$F_i = \text{Im} \int_0^\infty d\xi \frac{d\omega(\xi)}{d\xi} \int_{\text{surface}} \sum_j \langle T_{ij}(\mathbf{r}, \omega(\xi)) \rangle dS_j. \quad (6.12)$$

(Note that a finite spatial grid (as used in the present approach) requires no further regularization of the integrand, and the finite value of all quantities means there is no difficulty in commuting the Im operator in Eqs. (6.7–6.8) with the integration.) The standard Wick rotation corresponds to the particular choice $\omega(\xi) = i\xi$ and yields a smooth and rapidly decaying integrand [443].

Equation (6.11) can be solved in a number of ways, for example analytically in one dimension [304], or more generally by a finite-difference discretization [443]: this involves discretizing space and solving the resulting matrix eigenvalue equation using standard numerical linear algebra techniques [25, 531]. Of course, the diagonal ($\mathbf{r}' = \mathbf{r}$) part of the GF is formally infinite, but this singularity is not relevant because its surface integral is zero, and it is typically removed by some regularization (e.g. by the finite discretization or by a finite antenna size in the proposed experiments below). A crucial step for evaluation of Eq. (6.12), as mentioned above, is the passage to imaginary frequencies $\omega(\xi) = i\xi$. For real frequencies $\omega(\xi) = \xi$, the GF is oscillatory, leading to a highly oscillatory ST integrand that does not decay—even when a regularization (ultraviolet cutoff) is imposed, integrating a highly oscillatory function over a broad bandwidth is problematic. When evaluated over imaginary frequencies, on the other hand, the GF is exponentially decaying, due to the operator in Eq. (6.11) being positive-definite ($\nabla \times \nabla \times + \varepsilon\xi^2$) [443], leading to a decaying non-oscillatory integrand.

However, the Wick rotation is not the only contour in the complex plane that leads to a well-behaved decaying integrand. This is illustrated by Fig. 6-2, which shows the force integrand in the complex plane for the piston-like geometry of Fig. 6-8, recently shown to exhibit non-monotonic variations in the force as a function of plate-separation h [442, 443]. In particular, we calculate the (x -direction) force integrand $dF_x/d\omega = \int_{\text{surf}} \sum_j T_{xj}(\mathbf{r}, \omega) dS_j$

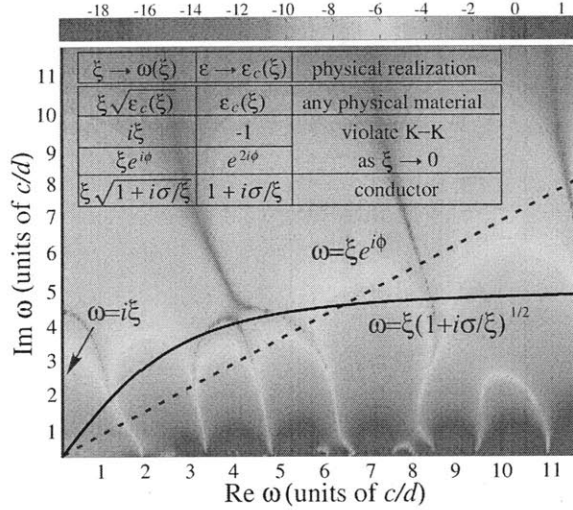


Figure 6-2: Complex-frequency ω plot of the Casimir force integrand ($\ln |\operatorname{Re} dF_x/d\omega|$), where $dF_x/d\omega$ is in units of \hbar/d^2 , for the geometry of Fig. 6-8. As the real- ω axis is approached, the integrand becomes highly oscillatory, which is only partially revealed here due to the finite frequency resolution. Various integration contours of interest are labeled as black and dashed lines. (Inset:) Vacuum $\varepsilon = 1$ contour deformations $\omega(\xi)$ and their corresponding (real-frequency) physical realizations $\varepsilon_c(\xi)$.

on one square, for $h = 0.5d$ and $s = d$, where d is the separation between the blocks. Here, we plot $\ln |\operatorname{Re} dF_x/d\omega|$, which illustrates the basic features of the integrand $d\omega/d\xi dF_x/d\omega$ (while $\operatorname{Im} dF_x/d\omega$ also contributes to the total force, it is qualitatively similar, and therefore omitted). As described above, this integrand is oscillating along the $\operatorname{Re} \omega$ axis and decaying along the $\operatorname{Im} \omega$ axis. Moreover, it also decays along any contour where $\operatorname{Im} \omega$ is increasing (such as the three contours shown, to be considered in more detail below).

6.3 Correspondence: Contour \leftrightarrow Material Deformations

In this section, we derive an equivalence between the complex-frequency GF (the GF evaluated over a complex frequency contour) of a geometry, and the real-frequency (standard) GF of an identical geometry with a transformed electromagnetic medium ε_c . As argued below, this correspondence presents a way to realize this contour-deformation technique in a real experiment. From this point of view, it becomes clear that only certain contours in the complex-frequency plane correspond to physically realizable ε_c , and that neither the real or imaginary frequency axes are suitable for experiments (the latter due to the fact that it is equivalent to a gain medium). Instead, we identify a suitable contour whose response is

equivalent to a physically realizable, conventional conducting medium. Section 6.3.2 demonstrates that the response of such a medium yields the correct Casimir force by performing numerical experiments on the geometry shown in Fig. 6-8, and then go on to describe the properties of the corresponding integrand. Finally, we consider the implications of the above mapping for practical experiments, including possible materials, lengthscales, and sources of experimental errors. The key point is that, once the Casimir force is expressed in terms of the response of a realizable medium over a reasonably narrow bandwidth, the scale-invariance of Maxwell's equations permits this response to be measured at any desired lengthscale, e.g. in a tabletop microwave experiment.

6.3.1 Complex-Frequency Green's Functions via Material Deformations

The fact that ω and ε appear together in Eq. (6.6), as $\varepsilon\omega^2$, immediately suggests that, instead of changing ω via a complex mapping $\omega(\xi)$ (again parametrized by real ξ), we can instead operate at real ξ by transforming ε . Specifically, evaluating Eq. (6.6) over a complex contour $\omega(\xi)$ is formally equivalent to evaluating Eq. (6.6) over real frequencies $\omega(\xi) = \xi$, but for a complex dielectric permittivity given by:

$$\varepsilon_c(\mathbf{r}, \xi) = \frac{\omega^2(\xi)}{\xi^2} \varepsilon(\mathbf{r}, \omega(\xi)) \quad (6.13)$$

Thus, given a medium $\varepsilon(\mathbf{r}, \omega)$ and a complex-frequency contour $\omega(\xi)$ over which we wish to compute the ST, we can transform to an equivalent problem involving a real-frequency ST evaluated over a complex dielectric $\varepsilon_c(\mathbf{r}, \xi)$, determined by Eq. (6.13). (Conversely, the GF of any frequency-dependent material $\varepsilon_c(\xi)$ at a given point in space can be related to the GF for vacuum ($\varepsilon = 1$) at that point by going from the real frequency ξ to a complex frequency $\omega = \xi\sqrt{\varepsilon_c(\xi)}$.) Equation (6.13) yields an intuitive explanation for why transforming to the complex-frequency plane is numerically advantageous: complex frequency deformations correspond to lossy materials that act to damp out the frequency oscillations. Furthermore, this equivalence is also of practical value to Casimir-force calculations (beyond being of pedagogical value): First, it forms an essential ingredient of a numerical finite-difference time-domain (FDTD) method, formulated in Ref. [449] and demonstrated for various geometries in Ref. [344]. The other application, discussed below, involves the possibility of computing the force via real experiments (in contrast to numerical experiments), in the

spirit of analog computations.

Because the Casimir ST is expressed in terms of the electromagnetic GF, and the GF of a system at microwave lengthscales is merely a rescaling of the GF of the system at micron lengthscales ¹, one can conceivably measure the GF (electric field in response to a current source) in an experiment via the S -matrix elements of antennas at centimeter scales, and in doing so determine the Casimir force via integration of the ST. The passage to complex frequencies is essential to an experiment, as in numerics, because measuring the GF in the original geometry $\varepsilon(\mathbf{r}, \xi)$ (real frequency ξ), as the electromagnetic response to an oscillating current source $\sim e^{i\xi t}$, will yield narrow peaks in its spectrum, requiring integration of a highly oscillatory force integrand over an infinite bandwidth, and thus imposing significant experimental challenges. Thus, one would like a way to *implement* the effect of complex frequency deformations in an experiment, which is the subject of the remaining text.

There are at least three ways to obtain the complex-frequency GF $G_{ij}(\mathbf{r}, \mathbf{r}', \omega)$ of a geometry $\varepsilon(\mathbf{r}, \omega)$ in an experiment [the electric field $\mathbf{E}(\mathbf{r})$ at a point \mathbf{r} in response to a dipole current source $\mathbf{J} = \delta(\mathbf{r} - \mathbf{r}')e^{-i\omega t}$, evaluated at *complex frequency* $\omega = \omega(\xi)$]: First, one may directly measure the electromagnetic field in response to a current source with time dependence $\sim e^{i\omega(\xi)t}$ (the most straightforward interpretation of complex-frequency GFs [304]). However, an $\omega(\xi)$ in the upper-half complex plane corresponds to both sources and fields with exponential growth in time, which is difficult for experiments. A second possibility involves measuring the standard real-frequency GF (the response to an oscillating source $e^{-i\xi t}$) over a large frequency bandwidth, and then performing a *numerical analytic continuation* of the form $\xi \rightarrow \omega(\xi)$. However, it is already known that inferring the Casimir force from real- ω GF measurements is extremely difficult, requiring the GF over a very wide bandwidth to a high accuracy [443]. A third alternative, and the subject of this section, is to instead exploit the correspondence above, in order to measure the real-frequency GF (the response to an oscillating source $\sim e^{-i\xi t}$) of a physical medium with a complex permittivity given by Eq. (6.13). As noted above, due to the presence of loss [corresponding to the imaginary part of $\omega(\xi)$], the integrand will be both smooth and decaying as a function of

¹The invariance of the GF to changes in lengths-scale is a consequence of the scale-invariance of Maxwell's equations for geometries composed of perfect-metal objects. When either the perfect-metal approximation breaks down or real materials are involved, changing the lengthscales of the problem also requires an appropriate dielectric re-scaling, i.e. materials with the appropriate dielectric response at the relevant lengthscales

Table 6.1: Vacuum $\varepsilon = 1$ contour deformations $\omega(\xi)$ and their corresponding (real-frequency) physical realizations $\varepsilon_c(\xi)$.

$\xi \rightarrow \omega(\xi)$	$\varepsilon \rightarrow \varepsilon_c(\xi)$	physical realization
$\xi\sqrt{\varepsilon_c(\xi)}$	$\varepsilon_c(\xi)$	any physical material
$i\xi$	-1	violate K-K as $\xi \rightarrow 0$
$\xi e^{i\phi}$	$e^{2i\phi}$	
$\xi\sqrt{1+i\sigma/\xi}$	$1+i\sigma/\xi$	conductor

real-frequency ξ .

In order for this approach to be experimentally viable, ε_c must satisfy two properties: it should correspond to an ω contour where the ST integrand is rapidly decaying, and it should be physically realizable. For simplicity, we begin by considering complex contours for a geometry in which the objects are separated by vacuum, so that the ST need only be evaluated in vacuum ($\varepsilon = 1$). This implies that, at those points where the ST needs to be evaluated, the real-frequency medium corresponding to a contour deformation $\omega(\xi)$ is determined by a dielectric function of the form $\varepsilon_c(\xi) = \omega(\xi)^2/\xi^2$.

We begin by noting that, in order for $\varepsilon_c(\xi)$ to correspond to a physical medium, it must satisfy the complex-conjugate property $\varepsilon_c(-\xi) = \varepsilon_c(\xi)^*$ as well as the Kramers–Kronig (K–K) relations [231]. It is most important to satisfy these conditions for small ξ , since the ST integrand is dominated by long-wavelength contributions. One should also prohibit gain media, which would lead to the exponentially growing fields we are trying to avoid by not using sources with complex ω . For example, a Wick rotation corresponds to a medium with dispersion $\varepsilon_c(\xi) = -1$, and this is only possible at $\xi = 0$ in a gain medium, since in a dissipative medium, ε_c is real and positive along the whole imaginary- ω axis (this is implied by K–K). The generalization to arbitrary rotations in the complex plane $\varepsilon = e^{2i\phi}$ yields similar difficulties, since these act both as gain media and also violate $\varepsilon_c(-\xi) = \varepsilon_c(\xi)^*$ near $\xi = 0$. Thus, no realizable material can emulate these standard contours even in a narrow bandwidth around $\xi = 0$, as summarized in Table 6.1 (also shown on the inset of Fig. 6-2).

Although traditional Wick rotations correspond to unphysical materials (as argued in the preceding paragraph), there are obviously many physical lossy materials to choose from, each of which corresponds to a contour in the complex plane, and one merely needs to find such a “physical” contour on which the ST is rapidly decaying so that experiments can be performed over reasonable bandwidths. A simple and effective lossy material for this purpose is a conductor with conductivity σ . Because the integral will turn out to be

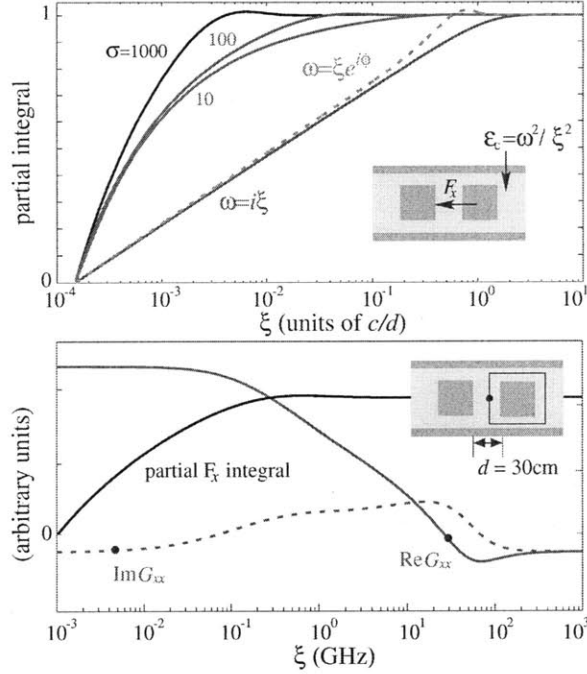


Figure 6-3: (Top:) Partial force integral $\int_0^\xi dF_x$, normalized by F_x , as a function of ξ , for the various $\omega(\xi)$ -contours (equivalently, various $\epsilon_c = \omega^2/\xi^2$) shown in Fig. 6-3. The solid green, red and black lines correspond to conductive media with $\sigma = 10, 10^2$ and 10^3 , respectively (σ has units of c/d). The dashed grey and solid blue lines correspond to $\phi = \pi/4$ and $\phi = \pi/2$ (Wick) rotations. (Bottom:) Illustration of the required frequency bandwidth for a possible realizations using saline solution at separation $d = 30\text{cm}$. The red lines plot the xx -component of the photon GF G_{xx} at a single location on the surface contour (see inset) as a function of ξ (GHz). The black line is the corresponding partial force integrand.

dominated by the contributions near zero frequency, it is sufficient to consider σ to be a constant (the DC conductivity), although of course the full experimental permittivity $\epsilon_c(\xi)$ could also be used. Specifically, we consider the general class of conductors defined by dispersion relations of the form $\epsilon_c(\xi) = 1 + i\sigma/\xi$, corresponding to vacuum with a complex contour $\omega_\sigma(\xi) = \xi\sqrt{1 + i\sigma/\xi}$. As shown in Fig. 6-2, the integrand of this contour is in fact well behaved, rapidly decaying and exhibits few oscillations. Moreover, because conductive fluids are ubiquitous, this choice is especially promising for experiments (see the section on analog computations).

6.3.2 Numerical Experiment

We now consider the Casimir force for the same structure as in Fig. 6-8, still calculated by the same finite-difference method as in Fig. 6-3, but we now focus on the properties along

different contour choices for both physical and unphysical media. In particular, Fig. 6-2(top) plots the partial integral $\int_0^\xi (d\omega/d\xi)(dF_x/d\omega)d\xi$, normalized by the total force $\int_0^\infty dF_x$, as a function of ξ . [As it must, the total integral over ξ , the force F_x , is invariant regardless of the contour $\omega(\xi)$ and agrees with previous results [443]; Specifically, $F_x = 0.0335 (\hbar c/d^3)$.] We now comment on two important features of the ω_σ contour that are relevant to experiments.

First, the Jacobian factor for ω_σ is given by $d\omega_\sigma/d\xi = 0.5(2+i\sigma/\xi)/\sqrt{1+i\sigma/\xi}$ and turns out to be very important at low ξ . The ST integrand itself goes to a constant as $\xi \rightarrow 0$ (due to the constant contribution of zero-frequency modes), but the Jacobian factor diverges in an integrable square-root singularity $\sim \sqrt{\sigma/\xi}$. Since this singularity is known analytically, however, separate from the measured or calculated GF, integrating it accurately poses no challenge. Second, the larger the value of σ , the more rapidly the ST integrand decays with ξ , and as a consequence the force integral for larger σ is dominated by smaller ξ contributions. In comparison, previous calculations of Casimir forces along the imaginary-frequency axis revealed that the relevant ξ bandwidth was determined by some characteristic lengthscale of the geometry such as body separations [443]. Here, we have introduced a new parameter σ that can squeeze the relevant ξ bandwidth into a narrower region. This “spectral squeezing” effect is potentially useful for experiments, as it partially decouples the experimental lengthscale of the geometry from the required frequency bandwidth.

6.3.3 A Casimir Analog Computer

The use of tabletop models and analog computers in physics, though previously unexplored in the context of quantum vacuum fluctuations, continues to play an important role in contemporary research areas like quantum evolution [320] and quantum information [313]. For example, many analog computers have been developed to simulate fluid flow problems [384,389], electromagnetic and acoustic wavefields [94], cell electrolysis [338], and many of the problems governed by Laplace’s equations [258]. More recently, advances in the field of microwave electromagnetics (which offers unprecedented control over sources and detection of microwaves), have spurred the development of classical electromagnetic analog computations for studying complex quantum problems, including the calculation of the energy levels of Bloch electrons in magnetic fields [277] and the dynamics of certain classes of chaotic quantum systems [320], to name a few. Computing fluctuation-induced effects like the Casimir force is substantially different from previous analog-computation problems

in quantum simulation, in that it involves not the evolution of a single quantum state but rather the combined effect of fluctuations over a broad (formally infinite) bandwidth.

As a consequence of the above results, we can now outline a possible experiment at centimeter lengthscales that determines the Casimir force at micron lengthscales, a Casimir analog computer (CAC). Suppose that one wishes to compute the Casimir force between perfect-metal objects separated by vacuum, such as the geometry in Fig. 6-8. One would then construct a scale model of this geometry at a tabletop scale (e.g., centimeters) out of metallic objects (which can be treated as perfect metals at microwave and longer wavelengths). To determine the ST integrand along a complex- ω_σ contour, one would measure the GF at real frequencies ξ for the model immersed in a conducting fluid. The GF is related to the S -matrix of pairs of antennas, and the diagonal of the GF to the S -matrix diagonal of a single antenna [noting that the finite size of the antenna automatically regularizes the integrand, as noted after Eq. (6.6)]. It is important for the model structure to be large enough that the introduction of a small dipole-like antenna does not significantly alter the electromagnetic response. In general, the ST must be integrated in space over a closed surface around the object, and correspondingly the antenna's S -matrix spectrum must be measured at a number of antenna positions (2D quadrature points) around this surface. (Unless one is interested in computing the force on a single atom, which requires a single antenna measurement.) The different components of the GF tensor correspond to different antenna orientations. The magnetic GF can be determined from the photon GF by Eq. (6.7), or possibly by employing "magnetic dipole" antennas formed by small current loops.

We now consider a particular CAC (at the cm scale) that employs realistic geometric and material parameters. Many available fluids exhibit almost exactly the desired material properties from above. One such example is saline water, which has $\varepsilon(\xi) = \varepsilon_s + i\sigma/\xi$, where $\varepsilon_s \approx 80$ and $\sigma \approx 5$ S/m for relatively small values of salt concentration [266]. A calculation using these parameters, based on the geometry of Fig. 6-8, assuming object sizes and separations at the centimeter to meter scale (we choose $d = 0.3$ m for the structure in Fig. 6-8, corresponding to a frequency of 1 GHz), reveals that it is only required to integrate the stress tensor up to small GHz frequencies ξ , which is well within the reach of conventional antennas and electronics. This is illustrated in Fig. 6-3(bottom), which plots the G_{xx} component of the GF (red lines) as well as the partial force integrand (black line), showing

the high $\xi < 1$ GHz cancellations that occur once the ST is integrated along a surface (inset). We note that most salts exhibit additional dispersion for $\xi > 10$ GHz [266], but we do not need to reach those frequency scales. (Nevertheless, should there be substantial dispersion in the conducting fluid, one could easily take it into account as a different complex- ω contour.)

Some attention to detail is required in applying this correspondence correctly. For instance, using network analyzers, what is measured in such an experiment is not the photon GF \mathbf{G} , but rather the electric S -matrix \mathbf{S}^E (the currents in a set of receiver antennas due to currents in the source antennas), related to the electric GF (the \mathbf{E} -field response to an electric current \mathbf{J}) by a factor depending on the antenna geometry alone (relating \mathbf{J} to \mathbf{E}). The electric GF will differ from the photon GF by a factor of the real frequency $i\xi$. To summarize, the photon GF will be given in terms of the measured \mathbf{S}^E by $G_{ij}(\omega) = (\alpha/i\xi)S_{ij}^E(\xi)$, where α is the antenna-dependent geometric factor (which can be measured with high accuracy). To obtain the ST from G_{ij} , one multiplies by factors of $\omega(\xi)^2$ as in Eq. (6.7). Fig. 6-3(bottom) illustrates the expected behavior of $S_{xx} \sim G_{xx}/\xi$ in a realistic system employing a saline solution with $d = 30$ cm.

We believe that such an experiment is feasible and capable of yielding accurate Casimir forces. In particular, the accuracy of the resulting force would only be limited by the magnitude of any measurement error, as the remaining post-processing steps can be performed with very high accuracy using well-known techniques [443]. With respect to measurement errors, we believe these to be well within the bounds necessary to obtain forces with accuracies better than or similar to those of our previous (and current) numerical experiments [344, 442, 443, 446, 449, 452]. Some of these may include: electronic noise, finite-size effects associated with any measurement device (e.g. antenna width), and surface roughness. It should be possible to make surface roughness negligible at cm scales, and in any case roughness is present in real Casimir experiments and is known to have a small effect if the roughness amplitude is small compared to the surface separations [40, 327] (although roughness could be intentionally introduced in the microwave system to study its effect). Finite-size effects are expected to be small because they are closely analogous to the effect of discretization in our numerical simulations (which give a finite size to current sources, among other effects), which do not contribute significantly to the force if the finite-size is small enough. Specifically, if a is a typical lengthscale in the problem (here, centimeters), for resolutions = 50 pixels/ a , corresponding to a smallest representable lengthscale of $0.02a$,

we obtain at least 1% accuracy in the force. For a tabletop experiment with a around 10 cm, this corresponds to antenna sizes < 2 mm, which is easily achievable. One must also deal with electronic sources of noise in the microwave regime. These include interference effects coming from sources in the vicinity of the experiment as well as thermal noise, both of which become significant at frequencies above 100 MHz [413]. There are, however, many effective standard ways to reduce interference errors, such as shielding and grounding of the entire apparatus [368]. Finally, though insignificant in the microwave regime, one may encounter noise originating in the receiver circuit, e.g. shot and pink noise. These absolute noise floors scale as the measurement bandwidth and can therefore be reduced if these measurements are performed over smaller frequency ranges (equivalent to averaging) [413]. Fortunately, both of these sources of noise are independent of the power at which one operates the antenna (measurement) device—therefore, since the GF measurement upon which this experiment is based is independent of the source amplitude (within the limits of the linear response of the materials), it is possible to increase the source power in order to significantly reduce the signal-to-noise ratio. For example, typical network analyzers can yield noise in the range (95, 120) dB at these frequencies for powers in the mW range [413], which is negligible even when accumulated over hundreds of measurements in a random walk. Moreover, numerical methods have the same sort of noise in the form of accumulated roundoff error, which is also on the order of 10^{-9} or larger (proportional to machine precision multiplied by the condition number of the matrix [531]), and we have not observed any appreciable effect of such roundoff noise on the accuracy (which is mostly limited by discretization/finite-size effects).

The calculation of Casimir forces and other fluctuation-induced interactions presents a challenge for both theory and experiment [169,268,352,358]. Especially for three-dimensional geometries, tabletop experiments offer an alternative route to rapidly exploring many different geometric configurations that are only now starting to become accessible to conventional numerical calculation (although complex geometries involving many objects remain challenging). Although many details of such an experiment remain to be developed, we believe that the basic ingredients are both clear and feasible, at least when restricted to perfect-metal bodies. The most difficult case to realize seems to be the force between imperfect-metal or dielectric bodies with a permittivity $\varepsilon(\omega)$, as the corresponding tabletop system requires materials with a specified dispersion relation $\varepsilon_c^{\text{body}}(\xi)/\varepsilon_c^{\text{fluid}}(\xi) = \varepsilon(\omega(\xi))$ relative

to the conducting fluid. This may be an opportunity for specially designed meta-materials with the desired frequency response. Analog computations at large lengthscales should also enable scientists to study Casimir forces in geometries composed of real dielectrics whose low-frequency response is of interest, such as recent predictions of repulsive forces between dielectric and magnetic [53, 356, 577] or magnetoelectric [189, 595] bodies which currently only exist at long wavelengths [164].

The correspondence (developed above) between complex-frequency deformations of the ST and causal material transformations is an indispensable ingredient for the development of a CAC. However, this equivalence is useful in at least two other ways: First, it was used to design new numerical methods for Casimir-force calculations based on the finite-difference time-domain method [344, 449], explored in the subsequent section; Second, it offers an alternative perspective on the Casimir effect. Specifically, one of the difficulties involved in understanding Casimir forces comes from the wide-bandwidth aspect of the real-frequency domain, which makes it impossible to separate contributions coming from individual electromagnetic modes, and therefore discourages the use of ideas from standard classical electromagnetism based on geometric or material resonances. For example, electromagnetic metamaterials are commonly designed to have effective materials properties that differ dramatically from their constituent materials, but only over a narrow bandwidth because they rely on strong resonant effects in subwavelength structures [599]. Physicists familiar with Casimir theory typically consider the Casimir-force contributions decomposed along the imaginary frequency axis, where the contributions are exponentially decaying and are mostly non-oscillatory because they are far removed from the resonant modes (all of which lie below the real-frequency axis) [68, 352, 443]. However, this viewpoint is foreign to most researchers from the classical electromagnetism community, which focuses on real frequencies and is often concerned with ideas based on geometric or material resonances in narrow bandwidths. The imaginary-frequency viewpoint indicates that resonant effects on Casimir phenomena are weak and that broad real- ω bandwidths are important to Casimir interactions, but this insight is far removed from the classical photonics way of thinking. On the other hand, absorption loss is well understood and familiar in classical electromagnetism, so our theoretical framework may provide a gentler pathway into Casimir physics from classical electromagnetics research: by realizing that Casimir forces are determined by a narrow-bandwidth response to a system with artificial dissipation added everywhere,

it becomes immediately clear even from the classical-photonics viewpoint why strong resonance effects are damped out. For example, although metamaterials can be designed to have effective materials properties that differ dramatically from their constituent materials over a narrow bandwidth, the fact that these effects mostly disappear at imaginary frequencies [456] can be difficult to convey to traditional metamaterials researchers. At the same time, it is a familiar fact in the metamaterial community that useful metamaterial properties are rapidly overwhelmed as strong dissipation is added to the system, and by expressing Casimir interactions in these terms more familiar classical analytical tools become applicable.

6.4 Casimir Forces in the Time Domain

In this section, we present a simple and general method to compute Casimir forces in arbitrary geometries and for arbitrary materials that is based on a finite-difference time-domain (FDTD) scheme in which Maxwell's equations are evolved in time [344, 449, 523]. A time-domain approach offers a number of advantages over previous methods. First, and foremost, it enables researchers to exploit powerful free and commercial FDTD software with no modification. The generality of many available FDTD solvers provides yet another means to explore the material and geometry dependence of the force, including calculations involving anisotropic dielectrics [457] and/or three-dimensional problems. Second, this formulation also offers a fundamentally different viewpoint on Casimir phenomena, and thus new opportunities for the theoretical and numerical understanding of the force in complex geometries.

Our time-domain method is again based on the stress tensor formulation of the previous sections. As noted above, evaluation at imaginary frequencies is unsuitable in the time domain. We overcome this difficulty by exploiting the equivalence of Sec. 6.3 between the system for which we wish to compute the Casimir force and a transformed problem in which all material properties are modified to include dissipation [448]. To illustrate this approach, we consider a simple choice of contour, corresponding to a conductive medium, that leads to a simple and efficient time-domain implementation. Finally, using a free, widely-available FDTD code [150], we compute the force between two vacuum-separated perfectly-metallic plates, a geometry that is amenable to analytical calculations and which we use to analyze

various important features of our method. An illustration of the power and flexibility of this method is provided in Ref. [344], in which we demonstrate computations of the force in a number of non-trivial (dispersive, three-dimensional) geometries as well as further refinements to the method.

In what follows, we describe the conceptual development of the FDTD method. The basic steps involved in computing the force on a body are:

- (1) Map the problem exactly onto a new problem with dissipation given by a frequency-independent conductivity σ .
- (2) Measure the electric \mathbf{E} and magnetic \mathbf{H} fields in response to current pulses placed separately at each point along a surface enclosing the body of interest.
- (3) Integrate these fields in space over the enclosing surface and then integrate this result, multiplied by a known function $g(-t)$, over time t , via Eq. (6.29).

The result of this process is the exact Casimir force (in the limit of sufficient computational resolution), expressed via Eq. (6.29) and requiring only the time-evolution of Eqs. (6.15–6.16).

In what follows, we describe the mathematical development of our time-domain computational method, starting from Eq. (6.12), in which the Casimir force is expressed as a contour integral of the frequency-domain stress tensor. We consider the frequency domain for derivation purposes only, since the final technique outlined above resides entirely in the time domain. In this framework, computing the Casimir force involves the repeated evaluation of the photon Green’s function G_{ij} over a surface S surrounding the object of interest. Our goal is then to compute G_{ij} via the FDTD method. The straightforward way to achieve this involves computing the Fourier transform of the electric field in response to a short pulse. However, in most methods a crucial step for evaluating the resulting frequency integral is the passage to imaginary frequencies, corresponding to imaginary time. We show that, in the FDTD, this gives rise to exponentially growing solutions and is therefore unsuitable. Instead, we describe an alternative formulation of the problem that exploits a recently proposed equivalence in which contour deformations in the complex frequency-domain $\omega(\xi)$ correspond to introducing an effective dispersive, dissipative medium at a real “frequency” ξ . From this perspective, it becomes simple to modify the FDTD Maxwell’s equations for the purpose of obtaining well-behaved stress tensor frequency integrands. We illustrate our

approach by considering a contour corresponding to a medium with frequency-independent conductivity σ . This contour has the advantage of being easily implemented in the FDTD, and in fact is already incorporated in most FDTD solvers. Finally, we show that it is possible to abandon the frequency domain entirely in favor of evaluating the force integral directly in the time domain, which offers several conceptual and numerical advantages.

6.4.1 Time Domain Approach

We begin our discussion of time-domain Casimir calculations starting with the frequency-domain approach discussed in the previous section, in which the force is computed by repeated evaluation of Eq. (6.11). In evaluating Eq. (6.11), we can choose from a general class of contours, provided that they satisfy $\omega(0) = 0$ and remain above the real ξ axis. The standard contour $\omega(\xi) = i\xi$ is a Wick rotation, which is known to yield a force integrand that is smooth and exponentially decaying in ξ [304]. In general, the most suitable contour will depend on the numerical method being employed. A Wick rotation guarantees a strictly positive-definite and real-symmetric Green's function, making Eq. (6.6) solvable by the most efficient numerical techniques (e.g. the conjugate-gradient method) [531]. One can also solve Eq. (6.6) for arbitrary $\omega(\xi)$ [448], but this will generally involve the use of direct solvers or more complicated iterative techniques [25]. However, the class of contours amenable to an efficient time-domain solution is more restricted. For instance, as discussed in Sec. 6.3.1, a Wick rotation turns out to be unstable in the time domain because it implies the presence of gain [448].

It is possible to solve Eq. (6.11) in the time domain by evolving Maxwell's equations in response to a delta-function current impulse $\mathbf{J}(\mathbf{r}, t) = \delta(\mathbf{r} - \mathbf{r}')\delta(t - t')\hat{\mathbf{e}}_j$ in the direction of $\hat{\mathbf{e}}_j$. G_{ij}^E can then be directly computed from the Fourier transform of the resulting \mathbf{E} field. However, obtaining a smooth and decaying force integrand requires expressing the mapping $\omega \rightarrow \omega(\xi)$ in the time-domain equations of motion. A simple way to exploit to see the effect of this mapping is to notice that Eq. (6.11) can be viewed as the Green's function at real "frequency" ξ and complex dielectric [448] given by Eq. (6.13), where for simplicity we have taken μ and ε to be frequency-independent. We assume this to be the case for the remainder of the chapter. At this point, it is important to re-emphasize that the original physical system ε at a frequency ω is the one in which Casimir forces and fluctuations appear; the dissipative system ε_c at a frequency ξ is merely an artificial

technique introduced to compute the Green's function.

Integrating along a frequency contour $\omega(\xi)$ is therefore equivalent to making the medium dispersive in the form of Eq. (6.13). Consequently, the time domain equations of motion under this mapping correspond to evolution of the fields in an effective dispersive medium given by $\varepsilon_c(\mathbf{r}, \xi)$.

To be suitable for FDTD, this medium should have three properties: it must respect causality, it cannot support gain (which leads to exponential blowup in time-domain), and it should be easy to implement. A Wick rotation is very easy to implement in the time-domain: it corresponds to setting $\varepsilon_c = -\varepsilon$. However, a negative epsilon represents gain (the refractive index is $\pm\sqrt{\varepsilon}$, where one of the signs corresponds to an exponentially growing solution). We are therefore forced to consider a more general, frequency-dependent ε_c .

Implementing arbitrary dispersion in FDTD generally requires the introduction of auxiliary fields or higher order time-derivative terms into Maxwell's equations, and can in general become computationally expensive [523]. The precise implementation will depend strongly on the choice of contour $\omega(\xi)$. However, almost any dispersion will suit our needs, as long as it is causal and dissipative (excluding gain). A simple choice is an $\varepsilon_c(\mathbf{r}, \xi)$ corresponding to a medium with frequency-independent conductivity σ :

$$\varepsilon_c(\mathbf{r}, \xi) = \varepsilon(\mathbf{r}) \left(1 + \frac{i\sigma}{\xi} \right) \quad (6.14)$$

This has three main advantages: first, it is implemented in many FDTD solvers currently in use; second, it is numerically stable; and third, it can be efficiently implemented without an auxiliary differential equation [523]. In this case, the equations of motion in the time domain are given by:

$$\frac{\partial \mu \mathbf{H}}{\partial t} = -\nabla \times \mathbf{E} \quad (6.15)$$

$$\frac{\partial \varepsilon \mathbf{E}}{\partial t} = \nabla \times \mathbf{H} - \sigma \varepsilon \mathbf{E} - \mathbf{J} \quad (6.16)$$

Writing the conductivity term as $\sigma \varepsilon$ is slightly nonstandard, but is convenient here for numerical reasons. In conjunction with Eqs. (6.3–6.4), and a Fourier transform in ξ , this

yields a photon Green's function given by:

$$\left[\nabla \times \frac{1}{\mu(\mathbf{r})} \nabla \times - \xi^2 \varepsilon(\mathbf{r}) \left(1 + \frac{i\sigma}{\xi} \right) \right] \mathbf{G}_j(\xi; \mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}') \hat{\mathbf{e}}_j, \quad (6.17)$$

This corresponds to picking a frequency contour of the form:

$$\omega(\xi) \equiv \xi \sqrt{1 + \frac{i\sigma}{\xi}}, \quad (6.18)$$

Note that, in the time domain, the frequency of the fields is ξ , and not ω , i.e. their time dependence is $e^{-i\xi t}$. The only role of the conductivity σ here is to introduce an imaginary component to Eq. (6.17) in correspondence with a complex-frequency mapping. It also explicitly appears in the final expression for the force, Eq. (6.12), as a multiplicative (Jacobian) factor.

The standard FDTD method involves a discretized form of Eqs. (6.15–6.16), from which one obtains \mathbf{E} and \mathbf{B} , not G_{ij}^E . However, in the frequency domain, the photon Green's function, being the solution to Eq. (6.6), solves exactly the same equations as those satisfied by the electric field \mathbf{E} , except for a simple multiplicative factor in Eq. (6.3). Specifically, G_{ij}^E is given in terms of \mathbf{E} by:

$$G_{ij}^E(\xi; \mathbf{r}, \mathbf{r}') = -\frac{E_{i,j}(\mathbf{r}, \xi)}{i\xi \mathcal{J}(\xi)}, \quad (6.19)$$

where $E_{i,j}(\mathbf{r}, \xi)$ denotes the field in the i th direction due to a dipole current source $\mathbf{J}(\mathbf{r}, t) = \mathcal{J}(t)\delta(\mathbf{r} - \mathbf{r}')\hat{\mathbf{e}}_j$ placed at \mathbf{r}' with time-dependence $\mathcal{J}(t)$, e.g. $\mathcal{J}(t) = \delta(t)$.

In principle, we can now compute the electric- and magnetic-field correlation functions by using Eqs. (6.9–6.10), with $\omega(\xi)$ given by Eq. (6.18), and by setting $\mathbf{r} = \mathbf{r}'$ in Eq. (6.10). Since we assume a discrete spatial grid, no singularities arise for $\mathbf{r} = \mathbf{r}'$, and in fact any \mathbf{r} -independent contribution is canceled upon integration over S . This is straightforward for Eq. (6.9), since the \mathbf{E} -field correlation function only involves a simple multiplication by $\omega^2(\xi)$. However, the \mathbf{H} -field correlation function, Eq. (6.10), involves derivatives in space. Although it is possible to compute these derivatives numerically as finite differences, it is conceptually much simpler to pick a different vector potential, analogous to Eqs. (6.3–6.4), in which \mathbf{H} is the time-derivative of a vector potential \mathbf{A}^H . As discussed in the Appendix, this choice of vector potential implies a frequency-independent magnetic conductivity, and

a magnetic, instead of electric, current. The resulting time-domain equations of motion are:

$$\frac{\partial \mu \mathbf{H}}{\partial t} = -\nabla \times \mathbf{E} + \sigma \mu \mathbf{H} - \mathbf{J} \quad (6.20)$$

$$\frac{\partial \varepsilon \mathbf{E}}{\partial t} = \nabla \times \mathbf{H} \quad (6.21)$$

In this gauge, the new photon Green's function $G_{ij}^H = \langle A_i^H(\mathbf{r}, \xi) A_j^H(\mathbf{r}', \xi) \rangle$ and the field \mathbf{H} in response to the current source \mathbf{J} are related by:

$$G_{ij}^H(\xi; \mathbf{r}, \mathbf{r}') = -\frac{H_{i,j}(\mathbf{r}, \xi)}{i\xi \mathcal{J}(\xi)}, \quad (6.22)$$

where the magnetic-field correlation function:

$$\langle H_i(\mathbf{r}, \xi) H_j(\mathbf{r}', \xi) \rangle = \frac{\hbar}{\pi} \omega^2(\xi) G_{ij}^H(\xi; \mathbf{r}, \mathbf{r}'), \quad (6.23)$$

is now defined as a frequency multiple of G_{ij}^H rather than by a spatial derivative of G_{ij}^E .

This approach to computing the magnetic correlation function has the advantage of treating the electric and magnetic fields on the same footing, and also allows us to examine only the field response at the location of the current source. The removal of spatial derivatives also greatly simplifies the incorporation of discretization into our equations (see Appendix for further discussion). The use of magnetic currents and conductivities, while unphysical, are easily implemented numerically. Alternatively, one could simply interchange ε and μ , \mathbf{E} and \mathbf{H} , and run the simulation entirely as in Eqs. (6.15–6.16).

The full force integral is then expressed in the symmetric form:

$$F_i = \text{Im} \frac{\hbar}{\pi} \int_{-\infty}^{\infty} d\xi g(\xi) (\Gamma_i^E(\xi) + \Gamma_i^H(\xi)), \quad (6.24)$$

where

$$\Gamma_i^E(\xi) \equiv \iint_S \sum_j \varepsilon(\mathbf{r}) \left(E_{i,j}(\mathbf{r}) - \frac{1}{2} \delta_{ij} \sum_k E_{k,k}(\mathbf{r}) \right) dS_j \quad (6.25)$$

$$\Gamma_i^H(\xi) \equiv \iint_S \sum_j \frac{1}{\mu(\mathbf{r})} \left(H_{i,j}(\mathbf{r}) - \frac{1}{2} \delta_{ij} \sum_k H_{k,k}(\mathbf{r}) \right) dS_j \quad (6.26)$$

represent the surface-integrated field responses in the frequency domain, with $E_{i,j}(\mathbf{r}) \equiv$

$E_{i,j}(\mathbf{r}; \xi)$. For notational simplicity, we have also defined:

$$g(\xi) \equiv \frac{\omega^2}{i\xi \mathcal{J}(\xi)} \frac{d\omega}{d\xi} \Theta(\xi) \quad (6.27)$$

Here, the path of integration has been extended to the entire real ξ -axis with the use of the unit-step function $\Theta(\xi)$ for later convenience.

The product of the fields with $g(\xi)$ naturally decomposes the problem into two parts: computation of the surface integral of the field correlations Γ , and of the function $g(\xi)$. The Γ_i contain all the structural information, and are straightforward to compute as the output of any available FDTD solver with no modification to the code. This output is then combined with $g(\xi)$, which is easily computed analytically, and integrated in Eq. (6.24) to obtain the Casimir force. As discussed in Sec. 6.6.1, the effect of spatial and temporal discretization enters explicitly only as a slight modification to $g(\xi)$ in Eq. (6.24), leaving the basic conclusions unchanged.

6.4.2 Evaluation in the Time Domain

It is straightforward to evaluate Eq. (6.24) in the frequency domain via a dipole current $\mathcal{J}(t) = \delta(t)$, which yields a constant-amplitude current $\mathcal{J}(\xi) = 1$. Using the frequency-independent conductivity contour Eq. (6.18), corresponding to Eqs. (6.15–6.16), we find the following explicit form for $g(\xi)$:

$$g(\xi) = -i\xi \left(1 + \frac{i\sigma}{\xi} \right) \frac{1 + i\sigma/2\xi}{\sqrt{1 + i\sigma/\xi}} \Theta(\xi) \quad (6.28)$$

One important feature of Eq. (6.28) is that $g(\xi) \rightarrow \sqrt{i\sigma^3/\xi}$ becomes singular in the limit as $\xi \rightarrow 0$. Assuming that $\Gamma^E(\xi)$ and $\Gamma^H(\xi)$ are continuous at $\xi = 0$ (in general they will not be zero), this singularity is integrable. However, it is cumbersome to integrate in the frequency domain, as it requires careful consideration of the time window for calculation of the field Fourier transforms to ensure accurate integration over the singularity.

As a simple alternative, we use the convolution theorem to re-express the frequency (ξ) integral of the product of $g(\xi)$ and $\Gamma^E(\xi)$ arising in Eq. (6.24) as an integral over time t of their Fourier transforms $g(-t)$ and $\Gamma^E(t)$. Technically, the Fourier transform of $g(\xi)$ does not exist because $g(\xi) \sim \xi$ for large ξ . However, the integral is regularized below using

the time discretization, just as the Green's function above was regularized by the spatial discretization. (As a convenient abuse of notation, ξ arguments will always denote functions in the frequency domain, and t arguments their Fourier transforms in the time domain.)

Taking advantage of the causality conditions ($\Gamma^E(t), \Gamma^H(t) = 0$ for $t < 0$) yields the following expression for the force expressed purely in the time domain:

$$F_i = \text{Im} \frac{\hbar}{\pi} \int_0^\infty dt g(-t) (\Gamma_i^E(t) + \Gamma_i^H(t)) \quad (6.29)$$

The advantage of evaluating the force integral in the time domain is that, due to the finite conductivity and lack of sources for $t > 0$, $\Gamma(t)$ will rapidly decay in time. As will be shown in the next section, $g(-t)$ also decays with time. Hence, although dissipation was originally introduced to permit a natural high-frequency cutoff to our computations, it also allows for a natural time cutoff T . We pick T such that, for times $t > T$, knowledge of the fields will not change the force result in Eq. (6.29) beyond a predetermined error threshold. This approach is very general as it requires no precise knowledge of how the fields decay with time.

6.4.3 Properties of $g(-t)$

Given $g(\xi)$, the desired function $g(-t)$ is a Fourier transform. However, the discretization of time in FDTD implies that the frequency domain becomes periodic and that $g(t) = g(n\Delta t)$ are actually Fourier series coefficients, given by:

$$g(n\Delta t) = \int_0^{2\pi/\Delta t} d\xi g_d(\xi) e^{-i\xi n\Delta t}, \quad (6.30)$$

where $g_d(\xi)$ is the discretized form of Eq. (6.27) and is given in the Appendix by Eq. (6.43). These Fourier series coefficients are computed by a sequence of numeric integrals that can be evaluated in a variety of ways. It is important to evaluate them accurately in order to resolve the effect of the $\xi = 0$ singularity. For example, one could use a Clenshaw-Curtis scheme developed specifically for Fourier integrals [407], or simply a trapezoidal rule with a large number of points that can be evaluated relatively quickly by an FFT (e.g. for this particular $g(\xi)$, 10^7 points is sufficient).

Since it is possible to employ strictly-real current sources in FDTD, giving rise to real Γ , and since we are only interested in analyzing the influence of $g(t)$ on Eq. (6.29), it suffices to




Figure 6-4: $|\text{Im } g(t)|$ for various values of σ , illustrating the transition from t^{-1} to $t^{-1/2}$ power-law decay as σ increases. Because there are strong oscillations in $g(t)$ at the Nyquist frequency for intermediate σ , for clarity we plot the positive and negative terms in $g(t)$ as separate components.

look at $\text{Im } g(-t)$. Furthermore, $g(t)$ will exhibit rapid oscillations at the Nyquist frequency due to the delta-function current, and therefore it is more convenient to look at its absolute value. Figure 6-4, below, plots the envelope of $|\text{Im } g(-t)|$ as a function of t , where again, $g(t)$ is the Fourier transform of Eq. (6.27). As anticipated in the previous section, $g(t)$ decays in time. Interestingly, it exhibits a transition from $\sim t^{-1}$ decay at $\sigma = 0$ to $\sim t^{-1/2}$ decay for large σ . The slower decay at long times for larger σ arises from a transition in the behavior of Eq. (6.28) from the singularity at $\xi = 0$.

Computing $g(t)$ straightforwardly, as was done in Fig. 6-4, is possible, provided one impose a high-frequency Nyquist cutoff. However, as argued in Ref. [344], there are a number of important simplifications that make it easier to compute both $g(t)$ and $F(t)$. (The remainder of this subsection is work described in Ref. [344].)

In particular, as noted above, $g(\xi)$ diverges in the high-frequency limit. For large ξ , $g(\xi)$ has the form:

$$g(\xi) \rightarrow g_1(\xi) \equiv \frac{\xi}{i} \Theta(\xi) + \sigma \Theta(\xi) \quad \text{as } \xi \rightarrow \infty \quad (6.31)$$

Viewing $g_1(\xi)$ as a function, we can only compute its Fourier transform $g_1(t)$ by introducing a cutoff in the frequency integral at the Nyquist frequency, since the time signal is only defined up to a finite sampling rate and the integral of a divergent function may appear to be undefined in the limit of no cutoff. (This was in fact the case in Fig. 6-4.)

Applying this procedure to compute $g(-t)$ yields a time series that has strong oscillations at the Nyquist frequency. The amplitude of these oscillations can be quite high, increasing

the time needed to obtain convergence and also making any physical interpretation of the time series more difficult.

These oscillations are entirely due to the high-frequency behavior of $g(\xi)$, where $g(\xi) \sim g_1(\xi)$. However, $g(t)$ and $g(\xi)$ only appear when they are being integrated against smooth, rapidly decaying field functions $\Gamma(\mathbf{x}, t)$ or $\Gamma(\mathbf{x}, \xi)$. In this case, g can be viewed as a tempered distribution (such as the δ -function) [463]. Although $g(\xi)$ diverges for large ξ , this divergence is only a power law, so it is a tempered distribution and its Fourier transform is well-defined without any truncation. In particular, the Fourier transform of $g_1(\xi)$ is given by:

$$g_1(-t) = \frac{i}{2\pi} \left(\frac{1}{t^2} + \frac{\sigma}{t} \right) \quad (6.32)$$

Adding and subtracting the term $g_1(\xi)$ from $g(\xi)$, the remaining term decays to zero for large ξ and can be Fourier transformed numerically without the use of a high-frequency cutoff, allowing $g(-t)$ to be computed as the sum of $g_1(t)$ plus the Fourier transform of a well-behaved function. This results in a much smoother $g(-t)$ which will give the same final force as the $g(-t)$ used in Ref. [449], but will also have a much more well-behaved time dependence. In Fig. 6-5 we plot the convergence of the force as a function of time for the same system using the $g(-t)$ obtained by use of a high-frequency cutoff and for one in which $g_1(\xi)$ is transformed analytically and the remainder is transformed without a cutoff. The inset plots $\text{Im} g(-t)$ obtained without using a cutoff (since the real part is not used in this work) for $\sigma = 10$.

In addition to the treatment of the high-frequency divergence, we find it convenient to also Fourier transform the low-frequency singularity of $g(\xi)$ analytically. As discussed above, the low-frequency limit of $g(\xi)$ is given by:

$$g(\xi) \rightarrow g_2(\xi) \equiv \frac{\sqrt{i} \sigma^{3/2}}{2 \xi^{1/2}} \Theta(\xi) \text{ as } \xi \rightarrow 0 \quad (6.33)$$

The Fourier transform of $g_2(\xi)$, viewed as a distribution, is:

$$g_2(-t) = \frac{i \sigma^{3/2}}{4\sqrt{\pi} t^{1/2}} \quad (6.34)$$

After removing both the high- and low-frequency divergences of $g(\xi)$, we perform a numerical Fourier transform on the function $\delta g(\xi) \equiv g(\xi) - g_1(\xi) - g_2(\xi)$, which is well-




Figure 6-5: Plot of the force error (force after a finite time integration vs. the force after a very long run time) for $g(t)$ determined from a numerical transform as in Ref. [449] and from the analytic transform of the high-frequency components. Inset: $\text{Im}[g(-t)]$ obtained without a cutoff, in which the high-frequency divergence is integrated analytically. Compare with Fig. 1 of Ref. [449]

behaved in both the high- and low-frequency limits. The imaginary part of $g(-t)$ is thus given by:

$$\text{Im}[g(-t)] = \text{Im}(\delta g(-t)) + \frac{1}{2\pi} \left(\frac{1}{t^2} + \frac{\sigma}{t} \right) + \frac{1}{4\sqrt{\pi}} \frac{\sigma^{3/2}}{t^{1/2}} \quad (6.35)$$

6.5 The Finite-Difference Time-Domain Method

In this section we discuss the practical implementation of the time-domain algorithm (using a freely-available time domain solver [150] that required no modification). We analyze its properties applied to the simplest parallel-plate geometry [Fig. 6-6], which illustrate the essential features in the simplest possible context. In particular, we analyze important computational properties such as the convergence rate and the impact of different conductivity choices.

6.5.1 Fields in Real Time

The dissipation due to positive σ implies that the fields, and hence $\Gamma^E(t)$, will decay exponentially with time. Below, we use a simple one-dimensional example to understand the consequences of this dissipation for both the one-dimensional parallel plates and the two-dimensional piston configuration. The simplicity of the parallel-plate configuration allows us to examine much of the behavior of the time-domain response analytically. (The understanding gained from the one-dimensional geometry can be applied to higher dimensions.)




Figure 6-6: $\Gamma_x^E(t) + \Gamma_x^H(t)$ for a set of one-dimensional parallel plates as the separation h is varied. The inset shows the physical setup.

Furthermore, we confirm that the error in the Casimir force due to truncating the simulation at finite time decreases exponentially (rather than as t^{-1} , as it would for no dissipation).

6.5.2 Parallel Plates

To gain a general understanding of the behavior of the system in the time domain, we first examine a simple configuration of perfectly metallic parallel plates in one dimension. The plates are separated by a distance h (in units of an arbitrary distance a) in the x dimension, as shown by the inset of Fig. 6-6. The figure plots the field response $\Gamma_x^E(t) + \Gamma_x^H(t)$, in arbitrary units, to a current source $\mathcal{J}(t) = \delta(t)$ for increasing values of h , with the conductivity set at $\sigma = 10(2\pi c/a)$.

Figure 6-6 shows the general trend of the field response as a function of separation. For short times, all fields follow the same power-law envelope, and later rapidly transition to exponential decay. Also plotted for reference is a $t^{-3/2}$ curve, demonstrating that the envelope is in fact a power law.

We can understand the power law envelope by considering the vacuum Green's function G^E in the case $h \rightarrow \infty$ (analogous conclusions hold for G^H). In the case $h \rightarrow \infty$, one can easily solve for the vacuum Green's function $G^E(\xi, \mathbf{r} - \mathbf{r}')$ in one dimension for real frequency ξ :

$$G^E(\xi, \mathbf{r} - \mathbf{r}') = \frac{e^{i\xi|\mathbf{r}-\mathbf{r}'|}}{i\xi} \quad (6.36)$$

We then analytically continue this expression to the complex frequency domain via Eq. (6.18) and compute the Fourier transform $\int d\xi e^{i\xi t} G^E(\omega(\xi))$. Setting $\mathbf{r} = \mathbf{r}'$ in the final expression, one finds that, to leading order, $G^E(t) \sim t^{-3/2}$. This explains the behavior of the envelope

in Fig. 6-6 and the short-time behavior of the Green's functions: it is the field response of vacuum.

Intuitively, the envelope decays only as a power in t because it receives contributions from a continuum of modes, all of which are individually decaying exponentially (this is similar to the case of the decay of correlations in a thermodynamic system near a critical point [183]). For a finite cavity, the mode spectrum is discrete — the poles in the Green's function of the non-dissipative physical system are pushed below the real frequency axis in this dissipative, unphysical system, but they remain discretely spaced.

At short times, the field response of a finite cavity will mirror that of an infinite cavity because the fields have not yet propagated to the cavity walls and back. As t increases, the cavity response will transition to a discrete sum of exponentially decaying modes. From Eq. (6.18), higher-frequency modes have a greater imaginary-frequency component, so at sufficiently long times the response will decay exponentially, the decay being determined by the lowest-frequency cavity mode. The higher the frequency of that mode, the faster the dissipation.

This prediction is confirmed in Fig. 6-6: as h decreases, the source “sees” the walls sooner. From the standpoint of computational efficiency, this method then works best when objects are in close proximity to one another (although not so close that spatial resolution becomes an issue), a situation of experimental interest.

6.5.3 Force Convergence in Time

We now examine the force on the parallel plates. From the above discussions of the field decay and the decay of $g(t)$, we expect the time integral in Eq. (6.29) to eventually converge exponentially as a function of time. In the interest of quantifying this convergence, we define the time dependent “partial force” $F_i(t)$ as:

$$F_i(t) \equiv \text{Im} \frac{\hbar}{\pi} \int_0^t dt' g(-t') (\Gamma_i^E(t') + \Gamma_i^H(t')) \quad (6.37)$$

Letting $F_i(\infty)$ denote the $t \rightarrow \infty$ limit of $F_i(t)$, which is the actual Casimir force, we define the relative error $\Delta_i(t)$ in the i -th component of the force as:

$$\Delta_i(t) \equiv \left| \frac{F_i(t) - F_i(\infty)}{F_i(\infty)} \right| \quad (6.38)$$




Figure 6-7: Partial force as defined in Eq. (6.37) for one-dimensional parallel plates as a function of time t . (Inset): Relative error $\Delta(t)$ as a function of t on a semi log scale.

We plot $F_x(t)$ in Fig. 6-7 for the one-dimensional parallel-plate structure with different values of σ . The inset plots $\Delta(t)$ for the same configuration. As expected, the asymptotic value of $F_x(t)$ is independent of σ , and $\Delta(t)$ converges exponentially to zero.

For σ near zero, the force is highly oscillatory. In one dimension this gives the most rapid convergence with time, but it is problematic in higher dimensions. This is because, in higher-dimensional systems, S consists of many points, each contributing a response term as in Fig. 6-7. If σ is small, every one of these terms will be highly oscillatory, and the correct force Eq. (6.37) will only be obtained through delicate cancellations at all points on S . Small σ is thus very sensitive to numerical error.

Increasing σ smooths out the response functions, as higher frequency modes are damped out. However, somewhat counter-intuitively, it also has the effect of slowing down the exponential convergence. One can understand the asymptotic behavior of the force by considering the equations of motion Eq. (6.17) as a function of σ and ξ . When the response function exhibits few if any oscillations we are in the regime where $\sigma \gg \xi$. In this limit, the approximate equations of motion are:

$$\left[\nabla \times \frac{1}{\mu(\mathbf{r})} \nabla \times - i\sigma \xi \varepsilon(\mathbf{r}) \right] \mathbf{G}_j(\xi; \mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}') \hat{\mathbf{e}}_j \quad (6.39)$$

In the limit of Eq. (6.39), the eigenfrequency ξ of a given spatial mode scales proportional to $-i/\sigma$. The lowest-frequency mode therefore has a time-dependence $\sim e^{-Ct/\sigma}$, for some constant $C > 0$. Since the decay of the force at long times is determined by this mode, we expect the decay time to scale inversely with σ in the limit of very high σ . This is suggested

in Fig. 6-7 and confirmed by further numerical experiments.

Additionally, from Eq. (6.39) we see that in the case of a homogeneous one-dimensional cavity, the solutions have a quadratic dispersion $\xi \sim ik^2$, for spatial dependence e^{ikx} , and so the lowest cavity frequency scales as the inverse square of the cavity size. This means that the rate of exponential convergence of Fig. 6-6 should vary as $\sim h^{-2}$ in the limit of very large σ . This scaling is approximately apparent from Fig. 6-6, and further experiments for much larger σ confirm the scaling. We thus see that in this limit, the effect of increasing σ by some factor is analogous to increasing the wall spacing of the cavity by the square root of that factor.

The present analysis shows that there are two undesirable extremes. When σ is small, rapid oscillations in $F_i(t)$ will lead to large numerical errors in more than one dimension. When σ is large, the resulting frequency shift will cause the cavity mode to decay more slowly, resulting in a longer run time. The optimal σ lies somewhere in between these two extremes and will generally depend on the system being studied. For the systems considered in this thesis, with a typical scale $\approx a$, $\sigma \sim 1$ ($2\pi c/a$) appears to be a good value for efficient and stable time-domain computation.

6.6 Implementing the FDTD Method

In this section, we discuss some of the considerations that went into implementing the above time-domain approach in the FDTD method.

6.6.1 Temporal Discretization

FDTD algorithms approximate both time and space by a discrete uniform mesh. Bearing aside the standard analysis of stability and convergence [523], this discretization will slightly modify the analysis in the preceding sections. In particular, the use of a finite temporal grid (resolution Δt) implies that all continuous time derivatives are now replaced by a finite-difference relation, which in the case of Maxwell's equations is commonly taken to be a forward difference:

$$\frac{\partial f}{\partial t} \approx \frac{f_i(\mathbf{r}, t + \Delta t) - f_i(\mathbf{r}, t)}{\Delta t} \equiv \partial_t^{(d)} f \quad (6.40)$$

where $f(t)$ is an arbitrary function of time. The effect of temporal discretization is therefore to replace the linear operator $\partial/\partial t$ with $\partial_t^{(d)}$. The representation of this operator is simple

to compute in the frequency domain. Letting $\partial_t^{(d)}$ act on a Fourier component of $f(t)$ yields:

$$\partial_t^{(d)} e^{-i\xi t} = -i\xi_d e^{-i\xi t}, \quad (6.41)$$

where

$$\xi_d(\xi) \equiv \frac{2}{\Delta t} \sin\left(\frac{\xi \Delta t}{2}\right) e^{-i\frac{\xi \Delta t}{2}} \quad (6.42)$$

The effect of discretization on the system is thus to replace $i\xi$ by $i\xi_d$ in the derivatives, which corresponds to numerical dispersion arising from the ultraviolet (Nyquist) frequency cutoff $\pi/\Delta t$. This not only affects the solutions to Eqs. (6.15–6.16), but also the complex-frequency mapping of Eq. (6.18): the mapping $\xi \rightarrow \omega(\xi) = \xi\sqrt{1+i\sigma/\xi}$ becomes $\xi \rightarrow \omega_d(\xi) = \xi_d\sqrt{1+i\sigma/\xi_d}$. Note that ξ is still the frequency parameter governing the time dependence of the Fourier components of $f(t)$ and $\xi_d \rightarrow \xi$ in the limit of infinite resolution ($\Delta t \rightarrow 0$). In this limit, our original mapping is restored and FDTD yields the correct temporal evolution of the fields. In practice, accurate convergence is achieved for reasonable resolutions, because Casimir forces are dominated by frequency contributions much smaller than the Nyquist frequency. The dominant frequency contributions to the Casimir force are bounded from above by a frequency $\sim 1/\ell$, where ℓ is determined by a typical lengthscale, e.g. body sizes and separations.

Because FDTD is convergent [$\xi_d = \xi + O(\Delta t^2)$], most of the analysis can be performed (as in this thesis) in the $\Delta t \rightarrow 0$ limit. However, care must be taken in computing $g(t)$ because the Fourier transform of $g(\xi)$, Eq. (6.27), does not exist as $\Delta t \rightarrow 0$. We must compute it in the finite Δt regime. In particular, the finite resolution requires, via Eq. (6.42), that we replace $g(\omega)$ in Eq. (6.27) by:

$$g_d(\xi) \equiv \frac{\omega_d^2}{i\xi_d \mathcal{J}(\xi)} \frac{d\omega}{d\xi} \quad (6.43)$$

Note that the Jacobian factor $d\omega/d\xi$ involves ω and ξ , not $\omega_d(\xi) = \xi_d\sqrt{1+i\sigma/\xi_d}$ and ξ_d , although of course the latter converge to the former as $\Delta t \rightarrow 0$. The basic principle is that one must be careful to use the discrete analogues to continuous solutions in cases where there is a divergence or regularization needed. This is the case for $g(\xi)$, but not for the Jacobian.

Similarly, if one wished to subtract the vacuum Green's function from the Green's function, one needs to subtract the vacuum Green's function as computed in the discretized

vacuum. Such a subtraction is unnecessary if the stress tensor is integrated over a closed surface (vacuum contributions are constants that integrate to zero), but is useful in cases like the parallel plates considered here. By subtracting the (discretized) vacuum Green's function, one can evaluate the stress tensor only for a single point between the plates, rather than for a “closed surface” with another point on the other side of the plates [304].

As was noted before Eq. (6.30), the Nyquist frequency $\pi/\Delta t$ regularizes the frequency integrations, similar to other ultraviolet regularization schemes employed in Casimir force calculations [303, 342]. Because the total frequency integrand in Eq. (6.1) goes to zero for large ξ (due to cancellations occurring in the spatial integration and also due to the dissipation introduced in our approach), the precise nature of this regularization is irrelevant as long as Δt is sufficiently small (i.e., at high enough resolution).

6.6.2 The Magnetic Correlation Function

One way to compute the magnetic correlation function is by taking spatial derivatives of the electric Green's function by Eq. (6.8), but this is inefficient because a numerical derivative involves evaluating the electric Green's function at multiple points. Instead, we compute the magnetic Green's function directly, finding the magnetic field in response to a magnetic current. This formulation, however, necessitates a change in the choice of vector potentials [Eqs. (6.3–6.4)] as well as a switch from electric to magnetic conductivity, for reasons explained in this section.

Eqs. (6.3–6.4) express the magnetic field \mathbf{B} as the curl of the vector potential \mathbf{A}^E , enforcing the constraint that \mathbf{B} is divergence-free (no magnetic charge). However, this is no longer true when there is a magnetic current, as can be seen by taking the divergence of both sides of Faraday's law with a magnetic current \mathbf{J} , $\partial\mathbf{B}/\partial t = -\nabla \times \mathbf{E} - \mathbf{J}$, since $\nabla \cdot \mathbf{J} \neq 0$ for a point-dipole current \mathbf{J} . Instead, because there need not be any free electric charge in the absence of an electric current source, one can switch to a new vector potential \mathbf{A}^H such that

$$\varepsilon E_i(\mathbf{r}, \omega) = (\nabla \times)_{ij} A_j^H(\mathbf{r}, \omega) \quad (6.44)$$

$$H_i(\mathbf{r}, \omega) = -i\omega A_i^H(\mathbf{r}, \omega). \quad (6.45)$$

The desired correlation function is then given, analogous to Eq. (6.7), by

$$\langle H_i(\mathbf{r}, \omega) H_j(\mathbf{r}', \omega) \rangle = \frac{\hbar}{\pi} \omega^2 \text{Im} G_{ij}^H(\omega; \mathbf{r}, \mathbf{r}'), \quad (6.46)$$

where the photon magnetic Green's function G^H solves [similar to Eq. (6.6)]

$$\left[\nabla \times \frac{1}{\varepsilon(\mathbf{r}, \omega)} \nabla \times - \omega^2 \mu(\mathbf{r}, \omega) \right] \mathbf{G}_j^H(\omega; \mathbf{r}, \mathbf{r}') = \delta(\mathbf{r} - \mathbf{r}') \hat{\mathbf{e}}_j. \quad (6.47)$$

Now, all that remains is to map Eq. (6.47) onto an equivalent real-frequency (ξ) system that can be evaluated in the time domain, similar to Sec. 6.4.1, for $\omega(\xi)$ given by Eq. (6.18). There are at least two ways to accomplish this. One possibility, which we have adopted here, is to define an effective magnetic permeability $\mu_c = \mu \omega^2(\xi) / \xi^2$, corresponding to a *magnetic* conductivity, similar to Eq. (6.13). Combined with Eq. (6.18), this directly yields a magnetic conductivity as in Eq. (6.20).

A second possibility is to divide both sides of Eq. (6.47) by $\omega^2 / \xi^2 = 1 + i\sigma / \xi$, and absorb the $1 + i\sigma / \xi$ factor into ε via Eq. (6.13). That is, one can compute the magnetic correlation function via the magnetic field in response to a magnetic current with an *electric* conductivity. However, the magnetic current in this case has a frequency response that is divided by $1 + i\sigma / \xi$, which is simply a rescaling of $\mathcal{J}(\xi)$ in Eq. (6.22). There is no particular computational advantage to this alternative, but for an experimental realization [448], an electric conductivity is considerably more attractive. [Note that rescaling $\mathcal{J}(\xi)$ by $1 + i\sigma / \omega$ will yield a new $g(\xi)$ in Eq. (6.27), corresponding to a new $g(t)$ that exhibits slower decay.]

6.6.3 Material Dispersion

In this section, we extend the time-domain formalism presented above to cases where the dielectric permittivity of the medium of interest is dispersive. To begin with, note that in this case the dissipative, complex dielectric ε_c of Eq. (6.48) is given by:

$$\varepsilon_c(\mathbf{r}, \xi) = \frac{\omega^2(\xi)}{\xi^2} \varepsilon(\mathbf{r}, \omega(\xi)), \quad (6.48)$$

where $\varepsilon(\mathbf{r}, \omega(\xi))$ denotes the permittivity of the geometry of interest evaluated over the complex contour $\omega(\xi)$.

This complex dielectric manifests itself as a convolution in the time-domain equations

of motion, i.e. in general, $\mathbf{D}(t) = \int dt' \varepsilon_c(t - t') \mathbf{E}(t')$. The standard way to implement this in FDTD is to employ an auxiliary equation of motion for the polarization [588]. For the particular contour chosen here [Eq. (6.18)], the conductivity term already includes the prefactor ω^2/ξ^2 and therefore one need only add the dispersion due to $\varepsilon(\mathbf{r}, \omega(\xi))$.

The only other modification to the method comes from the dependence of $\Gamma^E(\xi)$ in Eq. (6.25) on ε . We remind the reader that our definition of Γ was motivated by our desire to interpret Eq. (6.24) as the Fourier transform of the convolution of two quantities, and thus to express the Casimir force directly in terms of the electric and magnetic fields $\mathbf{E}(t)$ and $\mathbf{H}(t)$, respectively. A straightforward generalization of Eq. (6.25) to dispersive media entails setting $\varepsilon(\mathbf{r}) \rightarrow \varepsilon(\mathbf{r}, \omega)$. However, in this case, the Fourier transform of Eq. (6.25) would be given by a convolution of $\mathbf{E}(\xi)$ and $\varepsilon(\mathbf{r} \in S, \omega(\xi))$ in the time domain, making it impossible to obtain $\Gamma^E(t)$ *directly* in terms of $\mathbf{E}(t)$. This is not a problem however, because the stress tensor *must* be evaluated over a surface S that lies entirely within a uniform medium (otherwise, S would cross a boundary and interpreting the result as a force on particular objects inside S would be problematic). The dielectric appearing in Eq. (6.25) is then at most a function of $\omega(\xi)$, i.e. $\varepsilon(\mathbf{r} \in S, \omega) = \varepsilon(\omega)$, which implies that we can simply absorb this factor into $g(\xi)$, modifying the numerical integral of Eq. (6.30). Furthermore, the most common case considered in Casimir-force calculations is one in which the stress tensor is evaluated in vacuum, i.e. $\varepsilon(\mathbf{r} \in S, \omega) = 1$, and thus dispersion does not modify $g(\xi)$ at all.

6.6.4 Overview of the Algorithm

In what follows, we summarize the time-domain Casimir method presented above, which leads to the following numerical algorithm to compute the Casimir force on an object:

- (1) Pick a contour $\omega(\xi)$ over which to integrate the Casimir force. (For the class of contours considered here, Eq. (6.18), this corresponds to choosing a value for σ .)
- (2) Pick a surface of integration S around the body of interest, depicted by the black contour in Fig. 6-8.
- (3) Compute $g(n\Delta t)$, given by Eq. (6.30), by numerically integrating $g_d(\xi)$, given by Eq. (6.43). Note that the particular form of $g_d(\xi)$ will depend on $\omega(\xi)$, the temporal discretization Δt , the dielectric response at the surface of integration $\varepsilon(\mathbf{r} \in S, \omega) =$




Figure 6-8: Schematic illustration of a possible contour around a body; the force on the body as computed by Eq. (6.49), involves an integral of the fields over the contour S .

$\varepsilon(\omega)$, and the particular choice of current source—here, we only consider temporal delta-function current sources, i.e. $\mathcal{J}(\xi) = 1$. (Note that this step of the algorithm does not depend on geometry or the solution of the field equations. Thus, one can compute $g(n\Delta t)$ with very high accuracy and also store it for future use.)

- (4) For every grid point $\mathbf{r} \in S$ on the discretized surface S , and for each polarization \hat{e}_j :
 - (a) Time-evolve Eqs. (6.15–6.16) and Eqs. (6.20–6.21) in response to a dipole current source $\mathbf{J}(\mathbf{r}, n\Delta t) = \mathcal{J}(n\Delta t)\hat{e}_j$, where $\mathcal{J}(n\Delta t) = \delta_{n,0}/\Delta t$ is the discretized version of a delta function.
 - (b) Obtain the electric $E_{i,j} = E_i(\mathbf{r}, n\Delta t)$ and magnetic $H_{i,j} = H_i(\mathbf{r}, n\Delta t)$ fields at position \mathbf{r} , and at each time step $n\Delta t$.
 - (c) Terminate the simulation once the fields have decayed beyond a desired error threshold (a time $T = N\Delta t$ determined by σ , as discussed in Sec. 6.5.3).
- (5) Integrate $E_{i,j}$ and $H_{i,j}$ over the contour S to obtain $\Gamma_i^E(n\Delta t)$ and $\Gamma_i^H(n\Delta t)$, as in Eqs. (6.25–6.26). (Note that the spatial integral over the contour S is in fact a summation.)
- (6) The force in the i th direction is given by a discretized form of Eq. (6.29):

$$F_i = \text{Im} \frac{\hbar}{\pi} \sum_{n=0}^N g(n\Delta t) (\Gamma_i^E(n\Delta t) + \Gamma_i^H(n\Delta t)) \quad (6.49)$$

6.6.5 Spatial Discretization

The FDTD algorithm involves a discretization of both space and time. The former leads to some subtleties in the implementation of the above algorithm, because the standard spatial discretization of FDTD is a *Yee lattice* in which different field components are discretized at different point in space [523]. For example, the grid for E_x is offset by half a pixel from the grids for E_y and H_z . This “staggered” grid is important to obtain second-order accuracy in discretizing the spatial derivatives, and also leads to other important properties of the algorithm [523]. However, it complicates the computation of the Casimir force, because evaluating the stress tensor requires all components of the electric and magnetic fields to be evaluated at the same point.

The standard solution, which we use [150] and is also commonly implemented in FDTD software, is to simply linearly interpolate the different field components to the same point by averaging fields at adjacent grid points. As long as one interpolates to a point that is centered with respect to the different grids, such as the center of a pixel/voxel, this interpolation is second-order accurate. Accurate centered interpolation is especially important for Casimir forces because the net force involves delicate cancellations.

When a discontinuous boundary between two materials is discretized, the standard FDTD method becomes only first-order accurate, although there are schemes to restore second-order accuracy for both dielectric [120, 150] and perfect-conductor [523, 589] interfaces. There is an additional difficulty for Casimir calculations, however, because they involve a singular Green’s function. The Green’s function itself is always singular since we are evaluating the source and field at the same point, but ordinarily this singularity is subtracted out (for analytical regularizations) or in any case integrates to zero contribution to the net force [443]. Near an interface, however, there is an additional singularity due to reflection, which can be thought of as the singularity in the Green’s function of an image source on the other side of an interface [231]. Analytically, this additional term is not problematic, since the stress-tensor integral must yield the same result independent of the integration surface S , even if S is very close to a material interface. Numerically, on the other hand, we have observed that the additional singularity in the Green’s function exacerbates the spatial discretization error—the solution still converges at the same rate, $O(\Delta x)$ or $O(\Delta x^2)$ depending on which interface-discretization scheme is used, but the con-

stant factor becomes larger as S approaches an interface. We dealt with this problem in two ways: First, we typically choose S to be midway between interfaces in order to maximize the distance from the interfaces. Second, we implemented the subtraction scheme described in Ref. [443], subtracting the stress integral for a second computation involving an isolated object for which the net force is theoretically zero, and we again observe that this greatly improves the constant coefficient in the convergence rate.

Finally, we should mention one pitfall to avoid in implementing the magnetic Green's function. Two approaches to computing the magnetic Green's function were described in Sec. 6.6.2, but there is also a third approach. Since Maxwell's equations are equivalent under interchange of \mathbf{E} and \mathbf{H} and ε and μ , one could compute the magnetic Green's function by simply recomputing the electric Green's function with ε and μ swapped (exchanging any perfect electric conductors for perfect magnetic conductors). However, because ε and μ are discretized slightly differently in a Yee-scheme FDTD implementation (due to the different grids of \mathbf{E} and \mathbf{H}), this approach can lead to subtle errors because it could change the volume of the objects being discretized. Even though the change in volume is of $O(\Delta x)$, the contribution from this change of volume need not converge to zero as $\Delta x \rightarrow 0$, because of the formally infinite energy density in the vacuum energy (due to the singularity of the Green's function). Therefore, we recommend computing the magnetic Green's function via the response to a magnetic current as described earlier.

6.7 Concluding Remarks

An algorithm to compute Casimir forces in FDTD has several practical advantages. FDTD algorithms that solve Maxwell's equations with frequency-independent conductivity, and even more complicated dispersions, are plentiful and well-studied. They are stable, convergent, and easily parallelized. Although the current formulation of our method requires the evaluation of $G_{ij}(\mathbf{r})$ along a surface S , requiring a separate calculation of the fields for each dipole source in S , all of these sources can be simulated in parallel, with no communication between different simulations until the very end of the computation. In addition, many FDTD solvers will allow the computational cell for each source to be parallelized, providing a powerful method capable of performing large computations.

The calculations of this chapter employed non-dispersive materials in the original (ω)

system. However, the theoretical analysis applies equally well to materials of arbitrary dispersion. Any materials that can be implemented in an FDTD solver (e.g. a sum of Lorentzian dielectric resonances [523]) can also be included, and existing algorithms have demonstrated the ability to model real materials [523, 588]. Existing FDTD implementations also handle anisotropy in ϵ and μ , multiple types of boundary conditions, and other complications [523].

In principle, the computational scaling of this FDTD method is comparable to finite-difference frequency-domain (FDFD) methods [443]. In both cases, each solver step (either a time step for FDTD or an iterative-solver step for FDFD) requires $O(N)$ work for N grid points. The number of time steps required by an FDTD method is proportional to the diameter of the computational cell, or $N^{1/d}$ in d dimensions. With an ideal multigrid solver, FDFD can in principle be solved by $O(1)$ solver steps, but a simpler solver like conjugate gradient requires a number of steps proportional to the diameter as well [443]. In both cases, the number of points to be solved on the surface S is $O(N^{1-1/d})$. Hence, the overall complexity of the simplest implementations (not multigrid) is $O(N^2)$. We believe that future boundary-element methods [434, 443] will achieve better efficiency, but such methods require considerable effort to implement and their implementation is specific to the homogeneous-medium Green's function, which depends on the boundary conditions, dimensionality and types of materials considered [430].

In Ref. [344] we illustrate the method in various non-trivial two- and three-dimensional geometries, including dispersive dielectrics. In addition, we introduce an optimization of our method (based on a rapidly converging series expansion of the fields) that greatly speeds up the spatial integral of the stress tensor. We also compute forces in three-dimensional geometries with cylindrical symmetry, which allows us to take advantage of the cylindrical coordinates support in existing FDTD software [150] and employ a two-dimensional computational cell.

Chapter 7

Casimir Forces Between Vacuum-Separated Metals

“At the moment it’s just a notion, but with a bit of backing I think I could turn it into a concept, and then an idea.” *Woody Allen*

7.1 Overview

The previous chapters described a number of computational methods, based on the finite-difference method, that directly exploit standard techniques from numerical electromagnetism, and which allow calculations of Casimir forces in arbitrary geometries and materials. In this chapter, we demonstrate the flexibility and capabilities of these numerical methods by presenting results on a number of complicated geometries consisting of vacuum-separated metallic objects. Section 7.2 describes a geometry composed of four z -invariant metallic objects, two cylinders adjacent to two metallic walls, exhibiting an interesting non-monotonic dependence arising purely from the non-additivity of the Casimir force—in fact, as argued below, the effect is impossible to explain, even qualitatively, using simple two-body (pairwise-interaction) approximations. The effect is explored for cylinders of different cross-sectional shape (squares, cylinders), and also for metallic and magnetic plates. Section 7.3 explains this effect using a perturbative Casimir-force computational method known as ray-optics that captures many-body effects. Finally, Sec. 7.4 presents calculations demonstrating one-dimensional Casimir stability in a glide-symmetric geometry consisting of interlocked brackets in which repulsion is obtained via a competition of attractive sur-

faces that can be explained using simple pairwise approximations. This type of repulsion may be of relevance to micro-electromechanical system.

7.2 Multi-Body Nonmonotonic Effects

In what follows, we study a pistonlike geometry introduced in Ref. [442] involving two infinite dielectric or metallic cylinders adjacent to two metallic plates (a distance h away), in two and three dimensions, and demonstrate nonadditive and nonmonotonic changes in the force due to these lateral plates. Rahi *et. al* showed this effect to be a consequence of the screening (enhancement) of the interactions by the fluctuating charges (currents) on the two cylinders and their images on the nearby plates, using a simple argument based on the method-of-images [428]. This work involved circular cylinders, and is described in Sec. 7.2.2. Finally, using a ray-optics approximation that is exact in the $h = 0$ limit, we examine the $h > 0$ case (not exact due to the fact that ray optics neglects diffraction effects), and show that ray optics is able to qualitatively reproduce the non-monotonicity identified in the exact calculation. In this sense, the ray optics is shown to capture the most essential features of the multi-body interaction, unlike simpler pairwise-interaction approximation, such as the proximity-force approximation (PFA). Using ray-optics, it becomes possible to quantify the impact of diffraction on the Casimir force in this geometry.

7.2.1 Generalized Casimir Piston

We consider a complicated geometry in which there are interactions between multiple bodies: a two-dimensional “piston”-like structure, shown in Fig. 7-1, consisting of two metal $s \times s$ squares separated by a distance a from one another (here, $s = a$) and separated by a distance h from infinite metal plates on either side. In particular, we compute the Casimir force between the two squares, in two dimensions (that is, for z -invariant fields, unlike the standard Casimir force in which z oscillations are included), as a function of the separation h . We perform these calculations using a proof-of-concept implementation of the FDFD method of chapter 5. The result for perfect conductors is shown in Fig. 7-1, plotted for the TE and TM polarizations and also showing the total force. In the limit of $h \rightarrow 0$, this structure approaches the “Casimir piston,” which has been solved analytically for the 2d TM polarization [79, 208], and our TM force, linearly extrapolated to $h = 0$, agrees with




Figure 7-1: Casimir force between 2d (z -invariant fields) metal squares F/F_{PFA} , vs. distance from metal plate h (inset), normalized by the total force (TE+TM) obtained using the PFA, $F_{\text{PFA}} = \hbar c \zeta(3) s / 8\pi a^3$. The total force is plotted (black squares) along with the TE (red dots) and TM (blue circles) contributions.

this result to within 4% (although we have computational difficulties for small h due to the high resolution required to resolve a small feature in FDFD). For $h > 0$, however, the result is surprising in at least two ways. First, the total force is *non-monotonic* in h , due to a competition between the TE and TM contributions to the forces. Second, the h dependence of the force is a *lateral* effect of the parallel plates on the squares, which would be zero by symmetry in PFA or any other two-body-interaction approximation.

Although lateral forces can still arise qualitatively in various approximations, such as in ray optics or in PFA restricted to “line-of-sight” interactions, it may not be immediately clear how these could predict non-monotonicity. We also note that, in the large- h limit, the force remains different from PFA due to finite- s “edge” effects [178], which are captured by our method. It turns out that one can qualitatively predict the non-monotonic behavior, due to the competition between TE and TM forces, using the ray-optics approximation, although this approximation is not quantitatively accurate except for $h = 0$; this is described in Sec. 7.3, based on the work in Ref. [590].

To further explore the source of the h -dependence, we plot the TM interaction-stress maps $\Delta\langle T_{xx} \rangle$ and $\Delta\langle T_{xy} \rangle$ in Fig. 7-2, for the perfect-metal squares at a typical frequency $\omega = 2\pi c/a$, and for varying distances from the metal plates ($h = 0.5, 1.0, 2.0$). As shown, the magnitudes of both the xx (a–c) and xy (d–f) components of the stress tensor change dramatically as the metal plates are brought closer to the squares. For example, one change in the force integral comes from T_{xy} , which for isolated squares has an asymmetric pattern

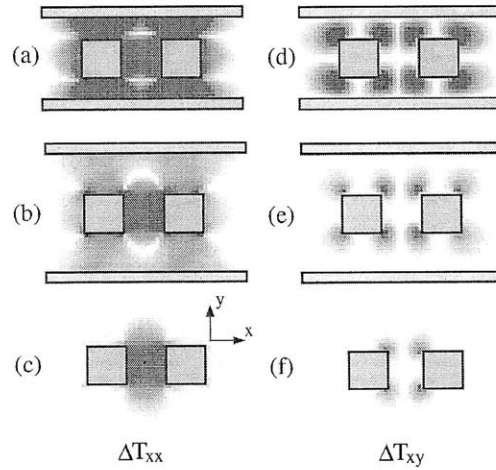


Figure 7-2: (a–f): TM stress map of the geometry in Fig. 7-1 for various h . The interaction stress tensors $\langle T_{xx} \rangle$ (left) and $\langle T_{xy} \rangle$ (right) for: (a),(d): $h = 0.5a$; (b),(e): $h = a$; and (c),(f): $h = 2a$, where blue/white/red = repulsive/zero/attractive.

at the four corners that will contribute to the attractive force, whereas the presence of the plates induces a more symmetric pattern of stresses at the four corners that will have nearly zero integral. This results in a decreasing TM force with decreasing h as observed in Fig. 7-1. Because stress maps indicate where bodies interact and with what signs, it may be useful in future work to explore whether they can be used to design unusual behaviors such as non-additive, non-monotonic, or even repulsive forces.

The FDFD method is also capable, without modification, of handling arbitrary dielectric materials. The calculation of general dispersive media can be performed with minor or no additional computational effort, since the computations at different ω are independent. Unlike most previously published techniques, which do not easily generalize to non-perfect metals, the stress tensor approach does not distinguish between the two regimes, and computational methods for inhomogeneous dielectric materials are widely available. Furthermore, we reiterate that along the imaginary- ω axis, ε is purely real and positive even for dissipative materials (which have complex ε on the real- ω axis), greatly simplifying computations.

The method's ability to handle dielectric structures is demonstrated below, where the Casimir force between the two squares is shown for two different cases:

First, we compute the force between two squares made of dielectric material with $\varepsilon = 4$ (an artificial mathematical choice for illustration purposes), whereas the parallel plates are still perfect metal. The result is shown on the plot of Fig. 7-3(Left). As might be expected, the dielectric squares have a weaker interaction than the perfect-metal squares, but are still

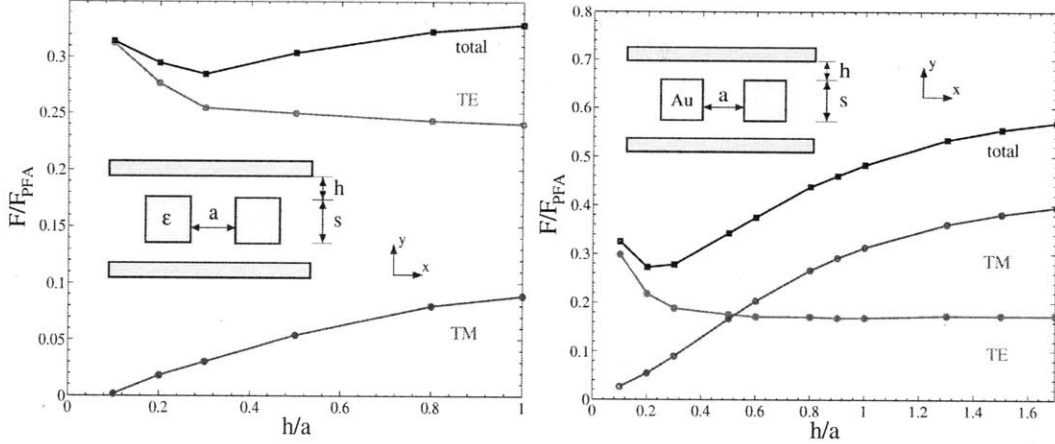


Figure 7-3: Casimir force between 2d (z -invariant fields) dielectric ($\epsilon = 4$, left) or gold (right) squares F/F_{PFA} , vs. distance from metal plate h (inset), normalized by the total force (TE+TM) obtained using the PFA (Here, the PFA force is computed for x -infinite slabs of corresponding dielectric ϵ). The total force is plotted (black squares) along with the TE (red dots) and TM (blue circles) contributions.

non-monotonic.

Second, as a more interesting example, the squares are made of gold with a Drude dispersion taken from experiment, again with adjacent perfect metallic plates. In particular, the following Drude model is used for the material dispersion of gold [61]:

$$\epsilon(\omega) = 1 - \frac{\omega_p^2}{\omega(\omega + i\Gamma_p)} \quad (7.1)$$

with $\omega_p = 1.37 \times 10^{16}$ Hz and $\Gamma_p = 5.32 \times 10^{13}$ Hz, corresponding to $\omega_p = 7.2731$ and $\Gamma_p = 0.028243$ in our units of $2\pi c/a$, for $a = 1 \mu\text{m}$. For $\omega = iw$, this is real and positive, as expected. The resulting force is shown in Fig. 7-3(Right). Not surprisingly, the gold squares have a weaker interaction than perfect-metal squares, since at large $w = \text{Im} \omega$ the dielectric constant ϵ goes to 1.

The previous 2d calculations are important in at least two ways: first, they allow us to check the stress tensor method against previous piston calculations in the $h \rightarrow 0$ limit, while exploring an interesting new geometry; second, based on their results, one might predict a similar behavior for the force per unit length in the analogous 3d z -invariant geometry, shown in the inset of Fig. 7-4. Indeed, this is the case: the non-monotonic force in three dimensions is shown in Fig. 7-4. As discussed in chapter 5, the integrand for a z -invariant perfect-metallic structure differs from the two-dimensional integrand only by a factor of πw ,

and therefore the 3d force is obtained from the 2d calculations with very little additional computation.




Figure 7-4: Casimir force per unit length between z -invariant 3d metal blocks F/F_{PFA} , vs. distance from metal plate h (inset), normalized by the total force (TE+TM) obtained using the PFA, $F_{\text{PFA}} = \hbar c s \pi^2 / 480 a^4$. The total force is plotted (black squares) along with the TE (red dots) and TM (blue circles) contributions.

Again, in the $h = 0$ limit there is a known analytical result for this geometry [208], based on the ray-optics method. Linearly extrapolating our plot to $h = 0$, we reproduce this result to within $\approx 2\text{--}3\%$.

7.2.2 Cylinders and Plates

The previous results on non-monotonic plates focused on square cylinders. Here, we replace the square cylinders with perfect-metal cylinders [141] and also explore what happens when one removes one of the sidewalls. Our results, based on work performed in collaboration with Jamal S. Rahi, Mehran Kardar, and Robert Jaffe (MIT) and Thorsten Emig (Univeriste Paris Sud), show that the non-monotonic effect is not limited to squares (i.e., it does not arise from sharp corners or parallel flat surfaces), nor does it require two sidewalls. Instead, the nonmonotonicity stems from a competition between forces from transverse electric (TE) and transverse magnetic (TM) field polarizations: In the latter case, the interaction between fluctuating charges on the cylinders is screened by opposing image charges, in the former case it is enhanced by analogous fluctuating image currents. Furthermore, we show that a related nonmonotonic variation arises for the force between the cylinders and a sidewall as a function of separation between the cylinders, a geometry potentially amenable to

experiment.

Casimir forces are not two-body interactions: quantum fluctuations in one object induce fluctuations throughout the system which in turn act back on the first object. However, both the PFA and CPI view Casimir forces as the result of attractive two-body (“pairwise”) interactions between surfaces that decrease rapidly as the surfaces are separated, and are reasonable approximations only in certain limits (e.g., low curvature for PFA), and can fail qualitatively as well as quantitatively otherwise, such as in geometries where the multi-body character of the Casimir force is important. Pairwise estimates fail to account for two important aspects of the Casimir forces in the geometry we consider [590]. First, as already mentioned, a monotonic pairwise attractive force clearly cannot give rise to the nonmonotonic effect of the sidewalls. Second, the application of either method here would include two contributions to the force on one of the cylinders: attraction to the other cylinder and attraction to the sidewall(s). If the latter contribution is restricted to the portion of the sidewall(s) where the other cylinder does not interpose (“line of sight” pairwise interactions), the cylinder will experience greater attraction to the portion of the sidewalls away from the other cylinder, thereby *reducing* the net attractive force between the two cylinders [590]. In contrast, exact calculations predict a nonmonotonic force that is *larger* in the limit of close sidewalls than for no sidewalls. These important failures illustrate the need for caution when applying uncontrolled approximations to new geometries even at the qualitative level. (On the other hand, it turns out that a ray-optics approximation, which can incorporate non-additive many-body effects, at least qualitatively predicts these features for the two-square/sidewall case [590].)

Figure B-1 is a schematic of the three-dimensional geometries that we consider: two infinite, parallel, perfect-metal cylinders of radius R separated by a distance a (center-to-center separation $2R + a$) and oriented along the z axis, with one (top panel) or two (bottom panel) infinite metal sidewall(s) parallel to the cylinders and separated from both by equal distance h ($h + R$ to the cylinder axes). Because the geometry consists of perfectly conducting objects (metals) and has continuous z -translational symmetry, the electric (\mathbf{E}) and magnetic (\mathbf{H}) fields can be decomposed into TE and TM polarizations, described by scalar fields ψ satisfying Neumann (TE, $\psi = H_z$) and Dirichlet (TM, $\psi = E_z$) boundary conditions at the metallic surfaces [100].

To analyze these geometries, we employ two complementary and exact computational




Figure 7-5: Schematic of cylinder-plate geometry consisting of two z -invariant metallic cylinders of radius R , separated by a distance a from each other, and by a distance h from one (top) or two (bottom) metallic sidewall(s).

methods, based on path integrals (PI) and the mean stress tensor (ST). The methods are *exact* in that they involve no uncontrolled approximations and can yield arbitrary accuracy given sufficient computational resources. They are *complementary* in the sense that they have different strengths and weaknesses. The PI method is most informative at large separations where it leads to analytical asymptotic expressions. The ST method, while relatively inefficient for large separations or for the specific geometries where PI is exponentially accurate, is formulated in a generic fashion that allows it to handle arbitrary complex shapes and materials without modification. The PI method is described in detail in Ref. [72], while the ST method was described in chapter 5, and so we only summarize them briefly here. The present geometry provides an arena where both methods can be applied and compared.

In the PI approach, the Casimir force is calculated via the constrained partition function. The Dirichlet (Neumann) constraints on the TM (TE) fields are imposed by auxiliary fields on the metallic surfaces [140] which can be interpreted as fluctuating charges (currents). The interaction between these charges is related to the free-space Green's functions—the addition of metallic sidewall(s) merely requires using image charges (currents) to enforce the appropriate boundary conditions. The calculations are further simplified by using Euclidean path integrals and the corresponding imaginary-frequency $\omega = iw$ Green's function. In the specific case of infinite cylinders, these surface fields can be represented in terms of a spectral basis: their Fourier series, leading to Bessel functions in the Green's function [46, 107, 141]. An important property of such a spectral basis is that its errors go to zero exponentially




Figure 7-6: Casimir force per unit length between two cylinders (black) vs. normalized cylinder-sidewall separation h/R , at fixed $a/R = 2$, normalized by the total PFA force per unit length between two isolated cylinders $F_{\text{PFA}} = \frac{5}{2}(\hbar c\pi^3/1920)\sqrt{R/a^7}$ [425] for the two cases shown on the inset. Also shown are the individual TE (red) and TM (blue) forces. *Inset:* Schematic of geometry in the case of one (dashed lines) and two (solid lines) sidewalls.

with the number of degrees of freedom.

We also use a method based on integration of the mean ST, evaluated in terms of the imaginary-frequency Green's function via the fluctuation-dissipation theorem [443]. The Green's function can be evaluated by a variety of techniques [443], but here we use a simple finite-difference frequency-domain method (FDFD) [93] that has the advantage of being very general and simple to implement at the expense of efficiency—it is much less efficient for this specific geometry than the PI method above. The results from both methods are shown in Fig. 7-6, with the PI method indicated by solid/dashed lines for two/one sidewalls, and the ST method indicated by data points. Both results agree to within the numerical accuracy, as expected, although we have fewer data points (and larger error bars) from the ST method because it is less efficient for this geometry. In the remaining text, therefore, we focus on the interpretation of the results rather than on the computational techniques.

To begin with, we compute the force between the two cylinders (with $a/R = 2$) as a function of the sidewall separation h/a , for fixed a . The results, for both one sidewall (solid lines) and two sidewalls (dashed lines) are shown in Fig. 7-6 for the total force (black), and also the forces for the individual polarizations TE (red) and TM (blue). The lines and dots indicate the two independent approaches discussed above. The forces are all normalized relative to the PFA force between two cylinders [425], which is independent of




Figure 7-7: Casimir force per unit length between two cylinders of fixed radius R vs. normalized cylinder-sidewall separation h/R (for one plate), normalized by the total PFA force per unit length between two isolated cylinders $F_{\text{PFA}} = \frac{5}{2}(\hbar c \pi^3 / 1920) \sqrt{R/a^7}$. The force is plotted for different cylinder separations $a/R = 0.2, 0.6, 1.0, 2.0, 3.0,$ and 4.0 .

h and therefore is merely a constant scaling that does not affect the shape of the curves in Fig. 7-6. Two interesting observations can be made from this figure. First, the total force (for both one and two sidewalls) is a nonmonotonic function of h/R : at first decreasing and then increasing towards the asymptotic limit of the force between two isolated cylinders.

The extremum for the one-sidewall case occurs at $h/R \approx 0.27$, and for the two-sidewall case is at $h/R \approx 0.46$ (similar to the $h/R \approx 0.5$ that was observed for squares between two sidewalls [442]). Second, the total force for the two-sidewall case in the $h = 0$ limit is larger than for $h \rightarrow \infty$. As might be expected, the h -dependence for one sidewall is weaker than for two sidewalls, and the effects of the two sidewalls are not additive: not only is the difference from $h \rightarrow \infty$ force not doubled for two sidewalls compared to one, but the two curves actually intersect at one point. When the forces are decomposed into the TE and TM contributions, it is clear that the nonmonotonicity results from a competition between them: the TE force is quickly decreasing with h while the TM force is slowly increasing.

Since nonmonotonic sidewall effects appear to occur for a variety of shapes (both for square [442] and circular cross sections), it is natural to seek a simple generic argument to explain this phenomenon. As we see in Fig. 7-6, and also in our earlier work [442] (Sec. 7.2), the nonmonotonicity arises from a competition between the TE and TM force contributions. Therefore, an underlying question is why the TE force increases as the sidewall comes closer, while the TM force decreases.




Figure 7-8: Casimir force per unit length between a plate and two cylinders of fixed radius R vs. normalized cylinder-cylinder separation a/R , normalized by the total PFA force per unit length between two isolated cylinders $F_{\text{PFA}} = \frac{5}{2}(\hbar c\pi^3/960)\sqrt{R/2a^7}$ [425]. The force is plotted for different plate separations $h/R = 0.28, 0.6, 1.0,$ and 2.0 .

An intuitive perspective of the effects of the metallic sidewall(s) on the TE/TM forces can be obtained from the “method of images,” whereby the boundary conditions at the plate are enforced by appropriate image sources. The Green’s function, the field in response to a point current source, arises in many formulations of the Casimir force. For example, in the stress-tensor formulation, the change in the Green’s function of the whole geometry compared to the Green’s function of vacuum appears directly as a “pressure” exerted on the object (cylinder) [425]. The change in the Green’s function arises because any current source induces currents on the surfaces of nearby objects, which in turn affect the total field—the question becomes, how does the presence of the plate change the effect of these induced surface currents? By the well-known method of images, the plate boundary conditions are satisfied by adding “image” currents below the plate. For Dirichlet boundary conditions (TM), these image currents must be of opposite sign, and therefore they act to “screen” the surface currents on the cylinders, *reducing* the force as h is decreased.¹ (This is analogous to a metal plate screening charges near the plate and reducing the classical force between them.) For Neumann boundary conditions (TE), the image currents must be of the *same* sign, and therefore *increase* the effect of the surface currents, *increasing* the force as h is decreased. (An analogous additive effect occurs for the classical force between current loops

¹A key fact is that the Green’s functions in Casimir forces are naturally evaluated at imaginary frequencies [443], which means that they are decaying and not oscillating. If the Green’s function were oscillating, one could not as easily infer whether opposite-sign image currents add or subtract to the field.

near a conducting plane.) The effect of the image currents in the TE case is not simply to double the force, however: on one hand, the image currents are farther away (the distance from one cylinder to another is a , whereas the distance from a cylinder to the image of the other is $> a$ for all h); on the other hand, the method of images is not additive in any case (the image currents induce currents on the original objects, which induce images that must be solved self-consistently), and indeed Fig. 7-6 shows the TE force more than tripling for small h .

The method of images thus explains the basic competition between TE and TM that underlies the nonmonotonic effects observed here and previously, but further considerations are required to insure that the sum of these two competing forces is nonmonotonic rather than being dominated by one or the other. As an extreme example, if the two TE and TM force decay rates with h were equal and opposite, they would cancel and there might be no h -dependence at all. However, we have recently shown [425] that the TE forces decay faster than the TM forces. Applied to the force between a cylinder and the image of the other cylinder, this suggests that the TE force can dominate for small h while the TM force should dominate for larger h , which is what we observe in Fig. 7-6.

While the method of images explains the competition between TE and TM that underlies the nonmonotonic effects observed here and previously, further considerations are required to ensure that the sum of these two competing forces is nonmonotonic. For example, if the two TE and TM force decay rates with h were equal and opposite, they would cancel and there might be no h -dependence at all. That this is not the case can be checked by examining the limit $a \gg h \gg R$. In this limit the forces are dominated by the lowest spectral (Fourier) mode, s -wave for TM and p -wave for TE [141]. The former is stronger and leads to an asymptotic form (for one plate):

$$\frac{F}{L} = -\frac{4\hbar c}{\pi} \frac{h^4}{a^7 \log^2(R/h)}, \quad (7.2)$$

While the logarithmic dependence on R could have been anticipated [141], the h^4 scaling is a remarkable consequence of the multi-body effect. Each cylinder and its mirror image can be considered as a dipole of size $\sim h$. The interaction of the two dipoles should hence scale as the interaction between two cylinders of radius $\sim h$ with Neumann boundary conditions, i.e. $\sim a^{-3} (h/a)^2$. We know that for $a \gg h$ the answer for the latter problem scales as $\sim h^4/a^7$,

explaining the above result, up to the log. To analytically establish the nonmonotonic character, we also need to show that the TE force is dominant in the opposite limit of $h \ll R$. So far, we only have numerical arguments in favor of this [425].

In Fig. 7-7, we show the total force vs. h/R for a variety of different values of the cylinder separation a in the presence of a single sidewall. The value $a/R = 2$ corresponds to our previous results in Fig. 7-6. Figure 7-7 shows two interesting effects. First, if a is too large or too small, the degree of non-monotonicity (defined as the difference between the minimum force and the $h = 0$ force) decreases. (For small a , the force is accurately described by the PFA force between the tangent cylinder surfaces, while for large a the TM mode dominates as indicated in Eq. (7.2).) The separation $a/R = 2$ from Fig. 7-6 seems to correspond to the largest value of non-monotonicity.

Consider the single-sidewall case from the top panel of Fig. B-1. A nonmonotonic force F_x between the cylinders means that there is a h where $dF/dh = 0$. The force is the derivative of the energy, $F_x = -\partial\mathcal{E}/\partial a$, so there is a point where $\partial^2\mathcal{E}/\partial a\partial h = 0$. These two derivatives, of course, can be interchanged to obtain $\partial(\partial\mathcal{E}/\partial h)/\partial a = 0$. But this means that $dF_y/da = 0$ at the same point, where $F_y = -\partial\mathcal{E}/\partial h$ is the force between the cylinders and the sidewall. That is, the force between the cylinders and the sidewall must also be a nonmonotonic function of the cylinder separation a . Precisely this cylinders-sidewall force is plotted in Fig. 7-8 as a function of a/R for various values of h/R . (It is not surprising that the effect of a small cylinder on the force between two bodies is smaller than the effect of an infinite plate. This is also reflected in the fact that the cylinder-cylinder force is generally less than the cylinder-plate force [425].)

The advantage of the cylinder-plate force compared to the cylinder-cylinder force is that it seems closer to the sphere-plate geometries that have been realized experimentally. In order to measure the cylinder-cylinder force, one would need to suspend two long cylinders in vacuum nearly parallel to one another. To measure the cylinder-plate force, the cylinders need not be separated by vacuum—we expect that a similar phenomenon will arise if the cylinders are separated by a thin dielectric spacer layer. Unfortunately, the nonmonotonic effect in Fig. 7-8 is rather small (roughly 0.2%), but it may be possible to increase this by further optimization of the geometry.

In previous research, many authors have sought unusual Casimir force phenomena by considering parallel plates with exotic materials: for example, repulsive forces have been

predicted using magnetic conductors [259], combinations of different dielectrics [228], fluids between the plates [375], and even negative-index media with gain [297]. A different approach is to use ordinary materials with more complicated geometries: as illustrated in this and previous [442] work, surprising nonmonotonic (attractive) effects can arise by including only interactions between three objects. A recent work studied the presence of non-monotonic effects in a variant of the z -invariant geometry considered here. In particular, Ref. [454] studied the force between two spheres of equal radii adjacent to two metallic walls (i.e. the same structure studied here but with spheres instead of cylinders). As turns out, the non-monotonicity persists in the case of spheres even though in that case there is no TE/TM decomposition.

7.2.3 Magnetic Sidewalls

In this section we consider a variant of the piston-like structure, shown as the inset to Fig. 7-9, studied using the FDTD method of chapter 6 in Ref. [344]. This system consists of two cylindrical rods sandwiched between two sidewalls, and is of interest again due to its non-monotonic dependence on h/a . The case of perfect metallic sidewalls ($\varepsilon(x) = -\infty$)² was studied in the previous section; here we also treat the case of perfect magnetic conductor sidewalls ($\mu(x) = -\infty$) as a simple demonstration of the ability of our approach to handle magnetic materials. In addition, as explained below, the presence of magnetic boundary conditions completely removes any non-monotonic dependence.

While three-dimensional in nature, the system is translation-invariant in the z -direction and involves only perfect metallic or magnetic conductors. As discussed in Ref. [443] this situation can actually be treated as the two-dimensional problem in the frequency domain. The same applies to the time domain, except that one must employ a slightly different form for the temporal Kernel $g(-t)$ discussed in chapter 6 [344]. As shown, our results agree with the high-precision scattering calculations of Ref. [428] (which uses a specialized exponentially convergent basis for cylinder/plane geometries).

The force, as a function of the vertical sidewall separation h/a , and for both TE and TM field components, is shown in Fig. 7-9 and checked against the results from the previous section [428]. We also show the force (dashed lines) for the case of perfect magnetic conductor

²Formally, in imaginary frequency, a perfect metal satisfies $\varepsilon(i\xi) = +\infty$. However, for a finite difference method either sign will give identical answers, and we find $\varepsilon(i\xi) = -\infty$ more suitable for our computational purposes.




Figure 7-9: Force for the double cylinders of [428] as a function of sidewall separation h/a , normalized by the PFA force $F_{\text{PFA}} = \hbar c \zeta(3) d / 8\pi a^3$. Red/blue/black squares show the TE/TM/total force in the presence of metallic sidewalls, as computed by the FDTD method (squares). The solid lines indicate the results from the scattering calculations of [428], showing excellent agreement. Dashed lines indicate the same force components, but in the presence of perfect magnetic-conductor sidewalls (computed via FDTD). Note that the total force is nonmonotonic for electric sidewalls and monotonic for magnetic sidewalls.

sidewalls. In the case of metallic sidewalls, the force is nonmonotonic in h/a . As explained in Ref. [428], this is due to the competition between the TM force, which dominates for large h/a but is suppressed for small h/a , and the TE force, which has the opposite behavior, explained via the method of images for the conducting walls. Switching to perfect magnetic conductor sidewalls causes the TM force to be enhanced for small h/a and the TE force to be suppressed, because the image currents flip sign for magnetic conductors compared to electric conductors. As shown in Fig. 7-9, this results in a monotonic force for this case.

7.3 Pistons via Ray Optics

As evidenced by the previous results, the ability to predict Casimir forces in arbitrary geometries can give rise to surprising and new effects. In fact, theoretical efforts to predict Casimir forces for geometries very unlike the standard case of parallel plates have begun to yield fruit, with several promising “exact” numerical methods (i.e., which can attain arbitrary accuracy given sufficient computing power) having been demonstrated for a few strong-curvature structures [46, 141, 143, 178, 427, 434, 442], in addition to yielding low-order analytical asymptotic solutions [46, 141, 143]. On the other hand, exact numerical methods are slow and difficult for complex geometries, and it is desirable to form a simple *qualitative*

intuition based on simpler methods that are quantitatively accurate only in certain limits. In this section, we explore the ability of a simple approximate method, the ray-optics technique [234,471], to bridge the gap between analytical calculations for simple geometries and brute-force numerics for complex structures. Unlike pairwise-interaction approximations such as the proximity-force approximation (PFA) [48], ray optics can capture multi-body interactions and thus has the potential to predict phenomena that simpler techniques cannot (despite the fact that both methods are quantitatively accurate only in a low-curvature limit). In particular, we show that the ray-optics approach can qualitatively predict the non-monotonic effect described in the previous sections [442,443]. For a sidewall separation $h = 0$, this is known as a “Casimir piston,” and in that case has been solved exactly [79,209]. While the ray optics approach is exact in this structure only for the piston case, its ability to capture the essential qualitative features for $h > 0$ suggests a wider utility as a tool to rapidly evaluate different geometries in order to seek interesting force phenomena. Furthermore, by comparison to the exact brute-force numerical FDFD method [443], we can evaluate the precise effect of diffraction (which is neglected by ray optics) on the Casimir force in this geometry.

The ray optics approximation expresses the Casimir force in terms of a sum of contributions from all possible classical ray paths (loops) with the same starting and ending point [234]. While it is strictly valid only in the limit of low surface curvature, since it neglects diffraction effects, the rays include multiple-body interactions because there exist loops that bounce off multiple objects. In contrast, most other low-curvature approximations, such as PFA [48] or other perturbative expansions [473,480], are essentially pairwise-interaction laws, and can therefore miss interesting physics that occurs when multiple bodies are brought together (in addition to the inevitable quantitative errors due to large curvature). One example occurs in the structure depicted in Fig. B-1, where there is a force




Figure 7-10: Schematic of a two-dimensional geometry: two metal squares $s \times s$ separated by a distance a , and separated from two adjacent metal sidewalls by a distance h .

between two square ($s \times s$) ideal metallic blocks separated by a distance a that is affected by the presence of two infinite parallel metallic sidewalls, separated from the blocks by a distance h . For perfect metals in the $h = 0$ limit, this geometry was solved analytically in both two dimensions [79] for Dirichlet boundary conditions and in three dimensions [210,334] for electromagnetic fields. (By “two dimensions,” we mean three-dimensional electromagnetism restricted to z -invariant fields; equivalently, a combination of scalar waves with Dirichlet and Neumann boundary conditions, corresponding to the two polarizations.) For $h > 0$, this geometry was recently solved by an exact computational method (that is, with no uncontrolled approximations) based on numerical evaluation of the electromagnetic stress tensor [442,443]. In this case, an unusual effect was observed: as h is increased from 0, the force between the blocks varies *non-monotonically* with h . The attractive force between the squares actually decreases with h down to some minimum and then increases toward the asymptotic limit of two isolated squares $h \rightarrow \infty$. These effects were obtained for both perfect-metal squares and for more realistic dispersive dielectrics, but here we focus on the simpler case of perfect metals (neglecting the non-zero skin depths and short-wavelength cutoffs of real metals). In the numerical solution, this non-monotonic effect arose as a competition between the TM polarization (electric field in the z direction, with Dirichlet boundary conditions) and the TE polarization (magnetic field in the z direction, with Neumann boundary conditions), which have opposite dependence on h . As explained below, it is unclear how this non-monotonic effect could arise in PFA or similar approximations—even if sidewall effects are included by restricting the pairwise force due to line-of-sight interactions, it seems that the effect of the sidewall must always decrease monotonically with h . When we analyze this structure in the ray-optics approximation, however, we find that a similar competition between the loops with an even and odd number of reflections again gives rise to a non-monotonic h dependence.

Below, we first give a general outline of the ray-optics approach, explain why pairwise approximations such as PFA must fail qualitatively in this geometry, and then present our results for the structure of Fig. B-1 in two dimensions. This is followed by a detailed description of the ray-optics analysis for this structure, which involves a combination of analytical results for certain (even-reflection) paths and a numerical summation for other paths.

7.3.1 Introduction to Ray Optics

Following the framework of Ref. [234], we express the two-dimensional Casimir energy via the ray-optics approximation. The ray-optics approach recasts the Casimir energy as the trace of the (scalar) electromagnetic Green's function $G(\mathbf{x}, \mathbf{x}') = [\nabla^2 + \omega^2 \varepsilon(\mathbf{x}, \omega)]^{-1} \delta(\mathbf{x} - \mathbf{x}')$, which is in turn expressed as a sum over contributions from classical “optical” paths via saddle-point integration of the corresponding path-integral (this is also referred to as the “classical optical approximation”). The optical paths follow straight lines and are labeled by the number of specular reflections from the surfaces of perfectly conducting objects. In particular, the Casimir energy between flat surfaces for Dirichlet ($\eta = -1$) or Neumann ($\eta = 1$) boundary conditions is given approximately by [234]:

$$\mathcal{E}_r = -\frac{\hbar c}{4\pi} \sum_r \eta^r \int_{D_r} \frac{d^2 x}{\ell_r(\mathbf{x})^3} \quad (7.3)$$

Here, the length of a closed geometric path starting and ending at a point \mathbf{x} is denoted by $\ell_r(\mathbf{x})$. D_r is the set of points that contribute to a closed optical path reflecting r times from the conducting surfaces. The Casimir energy above is thus the integral over the whole domain of such points. This problem reduces to computing a term-by-term contribution from each possible closed path, as determined by the specific geometrical features of the system under consideration. Because the Neumann (TE) and Dirichlet (TM) boundary conditions are given by the sum and difference of the even and odd paths (paths with even/odd numbers r of reflections), respectively, it is convenient to compute the contribution of even and odd paths separately. The Casimir force is then obtained by the derivative of the energy with respect to the object separation a .

Equation (7.3) is exact for objects with zero curvature (flat surfaces). In the presence of curved surfaces (or sharp corners, which in general have measure-zero contribution to the set of ray paths), however, the energy will include additional diffractive effects that are not taken into account by Eq. (7.3). One can include low-order corrections for small curvature [471], but this is obviously not applicable to the case of sharp corners. There is one special exception, the $h = 0$ “piston”: in this case, the sum over optical paths reduces to the method of images, which is exact for the interior of rectangular structures. These limitations are to be expected, however, since the optical theorem is a stationary-phase approximation.

7.3.2 Pairwise-Interaction Approximations

There are various pairwise-interaction force laws that were derived from various limits, and have been subsequently employed as approximate methods to compute Casimir forces in arbitrary geometries. Although they are controlled approximations only in certain limits (e.g., low curvature), it is tempting to use these methods to build a qualitative intuition in more general structures, in order to roughly understand Casimir interactions as a simple attractive force law. The most well-known of these is the proximity-force approximation (PFA), which treats the force between two bodies as a pairwise sum of “parallel-plate” contributions [48]. PFA is exact for parallel plates, and may have low-order corrections for small curvature [46], but is an uncontrolled approximation for strong curvature where it can sometimes give qualitatively incorrect results [138, 140, 168, 179, 327, 442]. Another pairwise interaction is the Casimir-Polder $1/r^7$ potential, valid in the limit of dilute media, which has recently been proposed as a simple (uncontrolled) approximation for arbitrary geometries by renormalizing it for the parallel-plate case [78, 480]. In this section, we briefly argue why no such pairwise-interaction approximation can give rise to the non-monotonic dependence on h that we observe in the structure of Fig. B-1.

Of course, if one considers the pairwise interaction as a true two-body force, for each pair of bodies in isolation, then the sidewalls in the structure can have no effect whatsoever: the force from one sidewall on one square will be exactly vertical (and canceled by the force from the other sidewall). However, a “lateral force” from the sidewalls, or equivalently an h -dependent change in the attractive force between the two squares, can be obtained by restricting the pairwise interactions to “line of sight” forces. For example, when considering the force on one vertical edge of a square from one of the sidewalls, one would include contributions only from the portion of the sidewall that is “visible” from that edge (connected in a straight line from a point on the edge to the point on the sidewall without passing through either square)³. For a fixed h , the line-of-sight force on the left edge of a square will be different from the force on the right edge, since one edge will have a portion of the sidewall blocked by the other square, and hence there will be an h -dependence of the horizontal attractive force.

³Such a restriction is at best *ad hoc*, since the points within each body are not ignored even though they are “blocked” by the covering material. Alternatively, one could formulate the two-body interaction as a force between surface elements exposed to one another, but then there are ambiguities about how to treat the surface orientation.

In particular, since the outside edges of the squares “see” (and are attracted to) a greater portion of the sidewalls than the inside edges, the net force from the sidewalls will *always reduce* the attractive force between the squares. Already, this contradicts the exact numerical calculations, in which both the Neumann force and the total force are *greater* at $h = 0$ than for $h \rightarrow \infty$.

Moreover, the effect of the sidewalls in a pairwise approximation must always *decrease* with h , again contradicting our results and making non-monotonic effects impossible. As h increases, two things happen: first, the inner edges of the square “see” a larger portion of the sidewalls, with area proportional to h ; second, the distance from the sidewalls to the squares increases proportional to h . The latter contribution must always dominate, however, because any pairwise force must decrease at least as fast as $1/h^3$ in two dimensions in order to reproduce the parallel-plate result. Therefore, the sidewall contribution must decrease monotonically at least as fast as $1/h^2$ in any pairwise-interaction approximation.

Unlike pairwise-interaction approximations, we show below that the ray-optics approximation correctly reproduces both qualitative behaviors: the total force is larger for $h = 0$ than for $h \rightarrow \infty$, and the total force is non-monotonic in h .

7.3.3 Path Contributions to Nonmonotonic Force

Here, we present the results of our calculations for the general $h \geq 0$ structure shown in Fig. B-1, and compare with the numerical results from the stress-tensor method [442, 443]. It turns out that the ray-optics technique indeed captures the non-monotonic dependence of the force with h , although of course the quantitative predictions differ from the exact calculations. By definition of ray optics, these quantitative corrections can be attributed to diffraction from the corners. Because we wish to emphasize the results of the ray-optics approach, rather than the details of the calculation of the different loop-lengths ℓ_r in Eq. (7.3), we defer those calculational details until Sec. 7.3.4 and here discuss the results.

Figures 7-11 and 7-12 show two different plots of the force vs. distance from the metal sidewalls h , computed via both Eq. (7.3) (solids) and the numerical stress-tensor method (dashed). All results are normalized by the PFA force between isolated squares (see captions), which are independent of h . The bottom panel shows the contributions from Neumann boundaries (TE polarization) and Dirichlet boundaries (TM polarization), along with the total Neumann+Dirichlet force. Recall that, in the ray optics approximation, the




Figure 7-11: Casimir even (red) and odd (blue) forces vs. sidewall separation h normalized by the PFA force $F_{\text{PFA}} = -\hbar c \zeta(3) s / 8\pi a^3$ (dashed black), computed using the ray-optics (solid) and stress-tensor (dashed) methods. Note that the ray-optics results become exact as $h \rightarrow 0$.

Neumann/Dirichlet forces are given in terms of the even- and odd-path contributions by $(\text{even} \pm \text{odd})/2$, respectively, and thus the total Neumann+Dirichlet force is equal to the contribution of the even paths alone. Because the even/odd decomposition is more natural, in the ray optics approximation, than Neumann/Dirichlet, the top panel shows the even/odd contributions from the same calculations. (Although the stress-tensor calculation does not decompose naturally into even and odd “reflection” contributions, here we simply define the even/odd components as $(\text{Neumann} \pm \text{Dirichlet})/2$, respectively.)

As h goes to zero, the ray-optics results become exact. The numerical computation of the stress-tensor force becomes difficult for small h due to our implementation’s uniform grid, but nevertheless the linear extrapolation of the numerical calculations to $h = 0$ agree with ray optics to within a few percent. For $h > 0$, the total force for both the ray-optics and stress-tensor results displays a minimum in the range $h \sim 0.2\text{--}0.3$. In particular, the extrema lie at $h \approx 0.3$ and $h \approx 0.25$, respectively. Not only is this striking non-monotonic behavior captured by the ray-optics approximation, but the agreement in the location of the extremum is also excellent.

Thus far, Figs. 7-11 and 7-12 reveal two significant differences between the ray-optics and stress-tensor results. First, the forces when h is not small differ quantitatively, by about 30% as $h \rightarrow \infty$. Since the ray-optics approximation is essentially obtained by dropping terms due to diffraction (from curved surfaces and corners), we can attribute this quantitative difference to the diffractive contribution to the Casimir force from the finite




Figure 7-12: Casimir force vs. sidewall separation h normalized by the PFA force $F_{\text{PFA}} = -\hbar c \zeta(3) s / 8\pi a^3$, computed using the ray-optics (solid) and stress-tensor (dashed) methods. The Neumann (green), Dirichlet (orange) and total (black) forces are all normalized by the total Neumann+Dirichlet PFA force.

size. In the large- h limit, where the sidewalls become irrelevant and the ray-optics result approaches PFA, the differences compared to the exact solution are sometimes called *edge effects* [178, 442, 443]. Second, although the exact and ray-optics results match in the $h = 0$ limit as discussed above, the functional forms for small h are quite different, at least for Dirichlet boundaries. In the ray optics expressions, the odd contributions have a logarithmic singularity at $h = 0$, which lead to corresponding singularities in the Neumann and Dirichlet forces. However, in the exact stress-tensor calculation, only the Neumann force seems to display a sharp upturn in slope as $h = 0$ is approached (although it is impossible to tell whether it is truly singular); the stress-tensor Dirichlet force seems to be approaching a constant slope (which is why we were able to linearly extrapolate it to $h = 0$ with good accuracy).

Another qualitative difference appears if we look at the even and odd contributions in Fig. 7-11: whereas ray optics and the stress-tensor method give a similar non-monotonic shape for the even force, the odd forces are quite different. (The stress-tensor odd force is monotonic while the ray-optics odd force is not, while the latter goes to zero for large h and the former does not.) Again, we attribute this to a greater sensitivity to diffraction effects, this time for the odd forces compared to the even forces. As will be argued in Sec. 7.3.5, the domain of integration and the length of even ray-optics paths have a weaker dependence on the corners of the squares than the odd paths, and thus should be less sensitive to corner-based diffractive effects. Fortunately, the total force depends solely on the even path

contributions, which helps to explain why ray-optics ultimately does effectively capture the non-monotonic behavior and the location of the extremum.

Having explored the h -dependence of the force using both ray-optics and numerical stress-tensor methods, we now turn to Fig. 7-13 to study the behavior of the force as a function of the square separation a . Figure 7-13 shows the Casimir force vs. square separation a at constant $h/s = 0.25$, normalized by the PFA force between isolated squares (top panel) or by the exact force at $h = 0$ (bottom panel). Note that in both cases the normalization is a -dependent, unlike in Figs. 7-11 and 7-12, and any non-monotonicity in Fig. 7-13 is only an artifact of this normalization. The normalization by the PFA force allows us to gauge both the sidewall/edge effects (which disappear for $a \rightarrow 0$) and whether there is a difference in scaling from PFA's $1/a^3$ dependence. For the bottom panel, we normalize against the exact $h = 0$ force, which tells us whether the finite sidewall separation makes a difference for the large- a scaling. In both cases, we show a few points of the stress-tensor calculation (which became very expensive for small or large a), to get a sense of the accuracy of the ray-optics method at different a .

We should expect that as $a \rightarrow 0$, both the PFA and the ray-optics solution should approach the exact solution, because the sidewall contribution becomes negligible. This agreement, as compared to the extrapolated numerical stress-tensor results, can be observed in Fig. 7-13(top). In contrast, for the large- a limit the ray-optics force appears to decay as $1/a^2$ instead of $1/a^3$ for PFA, leading to the apparent linear growth in the top panel of Fig. 7-13. If we compare to the $h = 0$ dependence in Fig. 7-13(bottom), it appears to be asymptoting to a constant for large a , indicating that the power laws for $h = 0$ and $h > 0$ may be identical. However, even if we had more data it would be difficult to distinguish the presence of, for example, logarithmic factors in this dependence. For the $h = 0$ case, we have analytical results for even- and odd-path forces in Sec. 7.3.7: from the analytical expressions, the odd-path $h = 0$ force clearly goes as $1/a^2$, and the even-path $h = 0$ force also turns out to have the same $1/a^2$ dependence⁴. The exact stress-tensor computation appears to be quite different from both the PFA and the ray-optics force as a function of a , but we were not able to go to large enough computational cells to estimate the asymptotic power law. It is striking that h/s as small as 0.25 is already large enough to yield substantial

⁴The asymptotic behavior of the Epstein Zeta function in the even force (Sec. 7.3.7) seems little studied, but it is straightforward to show that it goes as $1/a^2$ for large a .

diffraction effects in the force.

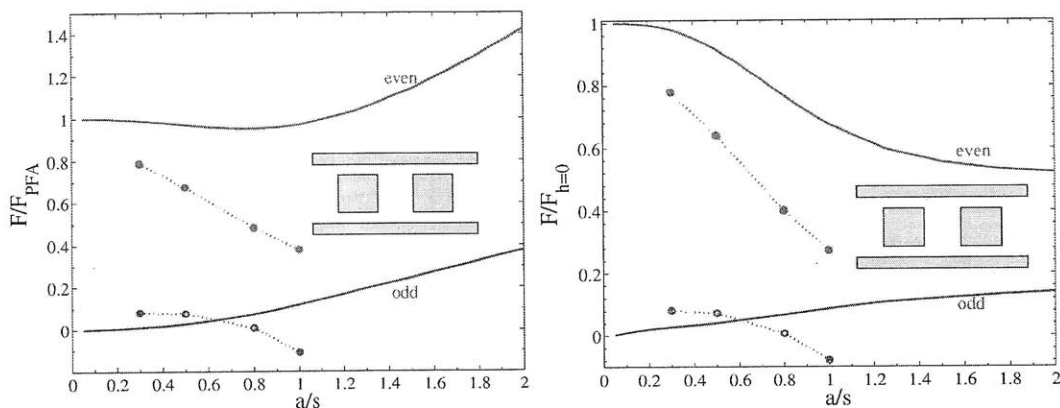


Figure 7-13: Casimir force vs. separation between squares a with constant sidewall separation $h = 0.25$, normalized by the $F_{\text{PFA}} = -\hbar c \zeta(3) s / 8\pi a^3$ (top) and the corresponding $h = 0$ force (bottom). The force is computed for both even (red) and odd (blue) contributions, separately. *Inset*: schematic of geometry consisting of two isolated squares with adjacent sidewalls.

7.3.4 Details of the Ray-Optics Computation

In this section, we describe the computation of the ray-optics approximation for the $h > 0$ squares+sidewalls structure, according to Eq. (7.3). This involves systematically identifying all of the possible closed ray loops, integrating them for a given number of reflections over the spatial domain, and then summing over the number of reflections. It turns out that the contributions of any even reflection order can be integrated analytically as shown below, although the final summation over reflection order is still numerical. On the other hand, the integrals from the odd reflection orders become increasingly difficult as the reflection order is increased, and so we resorted to numerical integration for odd orders $r > 5$.

7.3.5 Even Paths

Because we are dealing with perfect metals, and because the geometry has reflection-symmetry about the x and y axes, it is helpful to represent the optical paths using an infinite periodic lattice, shown in Fig. 7-14, similar to the construction in Ref. [210]. The reason for this construction is that, because of the equal-angle law for specular reflections in geometric optics, a reflected ray is equivalent to a linearly extended ray in a mirror-image structure. This allows us to visualize and count the set of possible closed paths in a

straightforward fashion. Specifically, a closed path which starts at a point \mathbf{x} and ends on itself is fully determined by the set of lines that start at \mathbf{x} and end in the corresponding set of image points on the extended lattice. The unit cell of this periodic construction is just two vertical black lines (of length s and separation a) that represent the parallel walls of the two squares. These are repeated with a horizontal period a and a vertical period $s + 2h$. Any path that passes through the gap between one of these lines and the horizontal sidewalls escapes from between the two squares and therefore is not counted among the closed loops.




Figure 7-14: Schematic of general 2d squares+sidewalls lattice. Lines extending from solid circles unto solid circles represent even reflection paths. Lines extending from open solid circles unto open circles represent odd reflection paths. (Here, $h/a \approx 0.2$ and $s/a \approx 1$.) A possible even (red) and forbidden odd (blue) path is shown.

In order to construct all of the closed loops that originate at a given point \mathbf{x} in the unit cell, we proceed as follows. First, we construct the mirror reflections of this point through the vertical lines (the boundaries of the squares) and the horizontal lines (the sidewalls), corresponding to reflections from these metallic walls. This gives us a set of points in the nearest-neighbor cells. Then, we construct the reflections of the nearest-neighbor points through their sidewalls, and so on, corresponding to reflections of higher and higher order. A closed loop is simply a line segment from the original \mathbf{x} to one of the reflected points, as long as it does not escape through one of the gaps between the squares and sidewalls as explained above. Figure 7-14 was generated from a unit cell with $h/a \approx 0.2$ and $s/a \approx 1$, and shows both even-reflection (solid red) and odd-reflection (dashed blue) paths, where in this case the odd path shown escapes and therefore would not be counted. In the figure,

the points are labeled according to the number of reflections that generate them from the original point: solid circles when the numbers of horizontal and vertical reflections are both even, open squares when the numbers are both odd, and open circles otherwise. It follows that lines connecting solid circles to solid circles are even paths, and lines connecting solid to open circles are odd paths; the rays connecting solid circles and squares always escape and therefore do not contribute..

To compute the Casimir energy from these paths, the key quantity in Eq. (7.3) is the length ℓ of the path. Let us label each unit cell by (n, m) according to its horizontal (n) and vertical (m) offset from the cell $(0, 0)$ where the original point \mathbf{x} resides. For an even path, \mathbf{x} must be connected to an even-indexed image $(2n, 2m)$, for which the length of the path is:

$$\ell_{n,m} = \sqrt{(2na)^2 + [2m(2h + s)]^2} \quad (7.4)$$

and the angle of the path, determined by m/n , is:

$$\tan \theta_{n,m} = \frac{m(2h + s)}{n a}. \quad (7.5)$$

The only things left to figure out are the domain of integration of \mathbf{x} and the allowed (n, m) for non-escaping paths. If $h = 0$, it is obvious that our expression reduces to the expression of Ref. [210], since both the whole spatial domain within the unit cell and *all* (n, m) are allowed. However, when $h > 0$, each (n, m) will be non-escaping only for \mathbf{x} in a subset of the unit cell. To determine these subsets, the domains of integration, we take advantage of the closed-loop nature of the paths to cast Eq. (7.3) in a different light. For a given (n, m) , instead of integrating over x and y , it is convenient to change variables to integrate over y and a coordinate $\xi = x/\cos \theta_{n,m}$ measuring displacement along the path (along the $\theta_{n,m}$ direction). It might seem that one should integrate ξ along the whole line from $(0, 0)$ to (n, m) , but this may involve counting the same point \mathbf{x} in the unit cell multiple times. Instead, to avoid over-counting loops that wrap around on themselves, one instead integrates from $(0, 0)$ to (\tilde{n}, \tilde{m}) where $\tilde{n} = n/\text{gcd}(n, m)$ and $\tilde{m} = m/\text{gcd}(n, m)$ are reduced to lowest terms. One then obtains the following equation for the energy:

$$\mathcal{E}_{\text{even}} = -\frac{\hbar c}{4\pi} \sum_{n,m} \int dy \int_{\xi(0,0)}^{\xi(\tilde{n},\tilde{m})} d\xi \cos \theta'_{\tilde{n},\tilde{m}} \frac{1}{\ell_{n,m}^3} \quad (7.6)$$

Although not so obvious from looking at Eq. (7.6), the integral in the y direction simply counts the number of (\tilde{n}, \tilde{m}) paths that exist in the unit cell. Because the length and angle of such paths are independent of y , we can integrate over $d\xi$ to obtain:

$$\mathcal{E}_{\text{even}} = -\frac{\hbar c}{4\pi} \sum_{n,m} \int dy \frac{2\tilde{n}a}{\ell_{n,m}^3} \quad (7.7)$$

where we used $\int d\xi = \ell_{\tilde{n},\tilde{m}}$ and $\cos \theta_{\tilde{n},\tilde{m}} = 2\tilde{n}a/\ell_{\tilde{n},\tilde{m}}$, between which the $\ell_{\tilde{n},\tilde{m}}$ cancels. All that is left to figure out is the integral in the y direction.

To carry the integral in the y direction, we must determine the limits of integration, or equivalently, the range over which we can displace the path so that it does not escape. We go back to Fig. 7-14 for reference. Again, as outlined above, we extend a line from a solid circle in the $(0,0)$ cell to another solid circle in the $(2n, 2m)$ -th cell. For the path to be allowed, it must intercept all of the vertical black line segments that lie between the two points, i.e. at each horizontal reflection from the squares. Because each interception (horizontal reflection) occurs at periodic intervals, these end up partitioning the unit cell in the y direction into \tilde{n} sets of length $(2h + s)/\tilde{n}$. Note that that we divide by \tilde{n} , rather than n , because as explained above the topologically distinct paths are uniquely specified by (n, m) reduced to lowest terms. From this simple argument, we obtain that the vertical displacement is $(2h + s)/\tilde{n} - 2h$, provided that $(2h + s)/\tilde{n} - 2h > 0$. To help visualize this result, it is best to think of the problem on a circle. That is, consider a circle of length $2h + s$ and partition it into \tilde{n} sets of length $(2h + s)/\tilde{n}$, as well as into two regions of length s and $2h$. If the path is to exist, each of the $(2h + s)/\tilde{n}$ points on the circle must not intercept the region marked as belonging to $2h$ (the air gaps between the squares). The result follows directly by considering the distance that one can displace the points before any of them intercepts the $2h$ region.

Thus, the final expression for the even path energies is left as a sum over n and m :

$$\mathcal{E}_{\text{even}} = -\frac{\hbar c}{4\pi} \sum_{n,m>0} \Theta \left(\frac{2h + s}{\tilde{n}} - 2h \right) \left[\frac{2h + s}{\tilde{n}} - 2h \right] \frac{4\tilde{n}a}{[(2na)^2 + (2m(2h + s))]^{3/2}}$$

An extra factor of 2 was included in the numerator since the contributions of the $(-n, -m)$ paths are identical to those of the (n, m) paths by symmetry. In the limit $h = 0$, we recover the even energy expression of Ref. [209], given in Sec. 7.3.7.

Although we are almost done with the even reflection paths, we are missing a very important contribution to the energy: the PFA terms, i.e. the $(n, 0)$ and $(0, m)$ paths. The PFA energy between two parallel finite metal regions is a well-known result, and we include here only $(n, 0)$ for completeness:

$$\begin{aligned}\mathcal{E}_{\text{PFA}} &= -\frac{\hbar csa}{16\pi} \sum_{n>0} \frac{1}{(na)^3} \\ &= -\frac{\hbar cs\zeta(3)}{16\pi a^2}\end{aligned}\tag{7.8}$$

The final expression Eq. (7.3.5) must still be summed numerically over (n, m) up to some upper cutoff, but this is a simple computation and its convergence with the cutoff is discussed in Sec. 7.3.8.

7.3.6 Odd Paths

Unlike even reflections, odd reflection paths are quite tedious to compute because they have less symmetry. For one thing, the length of a path depends not only on (n, m) (the offset of the unit cells being connected), but also on the position \mathbf{x} of the starting point. More importantly, the identification of the domain of \mathbf{x} for non-escaping paths depends in a much more complicated way on (n, m) , making it difficult to write down a single expression that works for all (n, m) . Therefore, we analytically solved for the odd-path contributions only up to five-reflection paths, where each order requires a separate analysis, and treated higher-order paths by a purely numerical approach. Below, the analytical solution for the third-order paths is given, both to illustrate the types of computations that are involved and also to demonstrate the logarithmic singularity in the force as $h \rightarrow 0$.

The results of Ref. [210] give an upper limit for the number of odd paths $r \geq 3$ that exist in this geometry for $h = 0$ ($\#$ paths = $2^{(r+3)/2}$). The same upper limit holds for $h > 0$, but in this case the number of paths is actually reduced because some of the $h = 0$ paths now escape. An analytical solution for any particular order must begin by drawing all paths for $h = 0$ and then perturbing them for $h > 0$ to eliminate any impossible paths. At least for low-order paths, simple geometrical arguments can then determine the domain of integration.

The coordinate dependence of odd paths arises from the fact that any path that reflects

an odd number of times from the planar surfaces will be non-periodic: if we extend the path beyond its endpoint (= starting point), it will not repeat. This greatly complicates the analysis. For example, Fig. 7-14 shows one such escaping odd path. It turns out that as h grows, the domain of integration for odd paths shrinks and becomes harder to visualize. Moreover, if we vary h , we notice that paths for some (n, m) , regardless of their origin \mathbf{x} , always escape. The problem of determining the domain of integration and the allowed (n, m) seems difficult and does not seem to have a general closed-form solution.

From Fig. 7-14 we verify that there are eight possible three reflection paths, $(\pm 2, \pm 1)$ and $(\pm 1, \pm 2)$ according to the notation of Sec. 7.3.5, shown in Fig. 7-15. The only two nonequivalent cases are $(2, 1)$ and $(1, 2)$, which have different domains of \mathbf{x} integration.




Figure 7-15: Schematic of three-reflection paths. Blue/red represent $(1, 2)/(2, 1)$ paths, and the lengths ℓ_i shown are used in the calculation of the energy.

Inspection of Fig. 7-15 and Fig. 7-14 yields the length of these paths $\ell_{2,1} = 2\sqrt{a^2 + (2h + s - y)^2}$. Moreover, we see that $\ell_{2,1}$ is invariant along the x -axis. For this particular path, the maximum displacement in the vertical direction y_{\max} occurs when $\ell_1 \sin \theta = \ell_2 \sin \theta = h$, while the minimum x_{\min} and maximum x_{\max} horizontal displacements occur when $\ell_2 \sin \theta = h$ and $\ell_1 \sin \theta = h$, respectively. Therefore, the range of integration is:

$$0 \leq y \leq s \quad (7.9)$$

$$\left(\frac{ah}{2h + s - y} \right) \leq x \leq \frac{a(h + s - y)}{2h + s - y}, \quad (7.10)$$

giving us the following expression for the energy:

$$\mathcal{E}_{2,1} = -\frac{\hbar c}{32\pi a^2} \left\{ \frac{sa}{\sqrt{a^2 + (2h + s)^2}} - 2h \log \left[\left(\frac{2h + s}{2h} \right) \frac{1 + \sqrt{1 + \left(\frac{2h}{a} \right)^2}}{1 + \sqrt{1 + \left(\frac{2h+s}{a} \right)^2}} \right] \right\} \quad (7.11)$$

In order to compute $\mathcal{E}_{1,2}$, we apply similar geometrical considerations. Once again, from Fig. 7-14 we obtain the length to be $\ell_{1,2} = 2\sqrt{(a-x)^2 + (2h+s)^2}$. Similarly, we see that $\ell_{1,2}$ is invariant along y -axis, and approaches a minimum when $x = a$. Therefore, the range of integration is:

$$h \leq y \leq h + s \quad (7.12)$$

$$0 \leq x \leq a, \quad (7.13)$$

giving us the following expression for the energy:

$$\mathcal{E}_{1,2} = -\frac{\hbar c a s}{32\pi (2h+s)^2} \frac{1}{\sqrt{a^2 + (2h+s)^2}}. \quad (7.14)$$

Adding Eq. (7.11) to Eq. (7.14), and multiplying by four to account for the different \pm sign possibilities in (n, m) yields the total three-reflection contribution to the energy. This contribution has two types of terms: a polynomial term from Eq. (7.14) and the first term of Eq. (7.11), and a logarithmic term from the second term of Eq. (7.11). The polynomial term remains at $h = 0$. The more intriguing component of this result is the logarithmic term, which falls as $O(h \log h)$ for small h , vanishing completely at $h = 0$. This (and similar terms for higher-order reflections) is the source of the logarithmic singularity in slope of the ray-optics odd-path Casimir force at $h = 0$ observed in Fig. 7-11.

The analytical solution of the five-reflection contribution is rather complicated and is not reproduced here. However, it has a few interesting features that we summarize here. First, just as for the three-reflection paths, the five-reflection contribution has an $O(h \log h)$ term that contributes to the singularity we observe in the ray-optics odd-path Casimir force at $h = 0$. Second, one also obtains $O(\log h)$ terms, which should seem to unphysically diverge as $h \rightarrow 0$. It turns out, however, that for a path that gives a $(\log h)$ contribution at small h , there exists another path with a $-(\log h)$ contribution, resulting in exact cancellation of any divergences.

7.3.7 Casimir Piston

The $h = 0$ limit is the well-known ‘‘Casimir piston’’ geometry. In this limit, all of the optical paths contribute to the Casimir energy, making it possible to compute Eq. (7.3)

analytically. Ref. [210] performs this calculation in three-dimensions, and a similar unpublished result was obtained in two-dimensions [209]. The two-dimensional geometry was also solved analytically for Dirichlet boundaries by another method [79]. Here, we reproduce this calculation as a check on both the even- [Eq. (7.3.5)] and odd-path [Eqs. (7.14–7.11)] contributions to the energy, and also because the Neumann result is useful and previously unpublished.

We begin by computing the $h = 0$ even-path contributions. The expression for the even energy Eq. (7.3.5) is much simpler than for $h > 0$, because the Θ function disappears and all that is left is a polynomial function in terms of n, m . The even-path energy, not including the PFA contribution (terms where $n = 0$ or $m = 0$, but not both), is given by:

$$\begin{aligned}\mathcal{E}_{\text{even}} &= -\frac{\hbar c}{8\pi} \sum_{n,m>0} as \left((na)^2 + (ms)^2 \right)^{-3/2} \\ &= -\frac{\hbar c}{8\pi} as Z_2(a, s; 3)\end{aligned}\tag{7.15}$$

we have identified the summation above as the second order Epstein Zeta function Z_2 :

$$Z_2(a, b; p) = \sum_{n,m>0} \left((na)^2 + (mb)^2 \right)^{-p/2}\tag{7.16}$$

The same simplification occurs in the case of odd paths, the lengths of which can be found, again, by inspection of the lattice in Fig. 7-14. Again, as described in Sec. 7.3.5, given a point on the unit cell \mathbf{x} , one can determine all possible odd-reflection paths by drawing straight lines from the solid circles (\mathbf{x}) to the open circles in the lattice.

There are two types of open circles, each of which denote two different types of paths: those that reflect across the x axis and have y -invariant length, as well as those that reflect across the y axis and have x -invariant length. Whichever coordinate is the invariant one gives a constant integral over the unit cell, so the double integral is reduced to a single integral. For example, consider those that reflect across the x axis and have y -invariant length: for these paths, we need to integrate over x in the unit cell and perform a double sum over (n, m) . However, the double sum can be reduced to a single sum by eliminating the sum over n in favor of integrating x along the whole real line instead of just the unit cell. Thus, we are left with a single integral and a single summation. Similarly paths that reflect across the y axis reduce to a single integral over y and a single sum over n . As a

result of all this manipulation, the odd integral becomes:

$$\mathcal{E}_{odd} = -\frac{\hbar cs}{32\pi} \sum_{m>0} \int_{-\infty}^{\infty} \frac{dx}{[(ms)^2 + x^2]^{3/2}} - \frac{\hbar ca}{32\pi} \sum_{n>0} \int_{-\infty}^{\infty} \frac{dy}{[(na)^2 + y^2]^{3/2}}. \quad (7.17)$$

Again, we restrict the sum to $n, m > 0$, because horizontal and vertical paths are divergent terms that contribute only to the self-energy of the metallic walls [471]. Carrying out the above integrals yields the following expression for the odd-path energy:

$$\mathcal{E}_{odd} = -\frac{\hbar c\pi}{48} \left(\frac{1}{s} + \frac{1}{a} \right). \quad (7.18)$$

The odd-path contribution to the force is therefore $\propto 1/a^2$ for large a .

7.3.8 Numerical Evaluation

To evaluate Eq. (7.3) numerically, we used an adaptive two-dimensional quadrature (cubature) algorithm [171] to perform the $\mathbf{x} = (x, y)$ integration for each (n, m) . (For each \mathbf{x} and (n, m) it is easy to numerically check whether the path is allowed, and set the integrand to zero otherwise. Unfortunately, this makes the integrand discontinuous and greatly reduces the efficiency of high-order quadrature schemes; an adaptive trapezoidal rule might have worked just as well. For very small h , this requires some care because the energy depends on a tiny remainder between two diverging terms, as we saw in Sec. ??, but for most h there was no difficulty.

We repeat this calculation for increasing reflection order r until the total energy \mathcal{E}_r converges to the desired accuracy. From general considerations, one expects the error $|\mathcal{E} - \mathcal{E}_r|$ for the energy from a finite r to decrease as $O(1/r)$. In particular, the path lengths ℓ increase proportional to r (the radius in the extended lattice), and the number of paths with a given length also increase proportional to r (the circumference in the extended lattice), so the $\sum 1/\ell^3$ for a given r goes as $O(1/r^2)$. The error in the energy is the sum over all paths of order $> r$, and this therefore goes as $O(1/r)$. This scaling is verified in Fig. 7-16, which plots the relative error $(\mathcal{E}_{r+1} - \mathcal{E}_r)/\mathcal{E}_r$ between the order- r and the order- $(r + 1)$ energy computations for the particular case of $a = s = 1$. In general, if the energy converged as $O(1/r^n)$ for some power n , one would expect this difference to converge as $O(1/r^{n+1})$, and so we expect Fig. 7-16 to asymptotically go as $O(1/r^2)$. This is precisely what is observed,




Figure 7-16: Convergence error $(\mathcal{E}_{r+1} - \mathcal{E}_r)/\mathcal{E}_r$ in the even (solid) and odd (dashed) path contributions vs. reflection order r for values of $h = 0$ (red), $h = 0.01$ (blue), and $h = 0.1$ (black) (here, \mathcal{E}_r means the energy computed up to order r). *Inset:* Absolute error $(\mathcal{E}_{\text{exact}} - \mathcal{E}_r)/\mathcal{E}_{\text{exact}}$ for both even (solid black) and odd (dashed red) contributions for $h = 0$.

for both even and odd paths, and for both $h = 0$ and $h > 0$: all of the curves asymptotically approach straight lines (on a log-log scale) with slope -2 .

An interesting though unfortunate result is that the odd-path energy requires larger r in order to obtain the same accuracy as the even-path energy. Though this may not be obvious from looking at the convergence error, it is clear from the inset of Fig. 7-16, where we plot the absolute error $(\mathcal{E}_{\text{exact}} - \mathcal{E}_r)/\mathcal{E}_{\text{exact}}$ instead (at $h = 0$). The constant offset observed in the absolute errors imply that, given a desired accuracy constrain on the even and odd energy calculations, one would have to compute roughly twice the number of odd paths in order to obtain equivalent accuracy.

7.3.9 Concluding Remarks

By comparing the ray-optics approximation with the exact brute-force calculation of chapter 6, we have been able to study both the successes and limitations of the ray-optics approximation. On the positive side, the ray-optics approximation is capable of capturing surprising behaviors that arise in closed geometries involving multiple bodies, qualitatively matching phenomena identified in exact brute-force calculations. In particular, the ray-optics approximation captures the non-monotonic sidewall effects observed for metallic squares between parallel sidewalls, generalized from the Casimir piston geometry. This effect is clearly a manifestation of the multi-body character of the interaction, since it does not arise in sim-

ple two-body force laws such as PFA. Ray optics appears to be unique among the current simple approximations for Casimir force in that it can capture such multi-body effects, even though it cannot be quantitatively accurate in geometries with strong curvature. On the negative side, diffractive effects set in rather quickly when h is increased from zero, marking the agreement between ray-optics and the exact results only qualitative.

This makes the ray optics approximation a promising approach to quickly search for unusual Casimir phenomena in complicated geometries. However, since it is an uncontrolled approximation in the presence of strong curvature, nor does it include material dispersion, any prediction by ray optics in such circumstances must naturally be checked against more expensive exact calculations. There will undoubtedly be complex structures in the future where ray optics fails qualitatively as well as quantitatively. For instance, ray optics has more difficulty with open geometries—e.g., for two squares with only one sidewall, only PFA paths are present. On the other hand, the reach of the ray optics technique seems in some sense to be larger than that of simpler approximations such as PFA. Even qualitative *failures* of such methods, however, can reveal interesting things about the physics of Casimir interactions, such as highlighting situations where multi-body interactions or diffraction effects are central.

7.4 Glide-Symmetric Geometry: Uni-Axial Stability

In this section, we describe a metallic, glide-symmetric, “Casimir zipper” structure (depicted in Fig. 7-17) in which both repulsive and attractive Casimir forces arise, including a point of stable equilibrium with respect to perpendicular displacements (a saddle-point equilibrium in the single direction in which repulsion is achieved). Here, the forces are “repulsive” in the sense that they act to separate the two structures, but in some sense they are a combination of attractive interactions as discussed below. We compute the force using the FDFD method, which is an “exact” computational method (i.e. with no uncontrolled approximations, so that it yields arbitrary accuracy given sufficient computational resources), and compare these results to the predictions of an *ad hoc* attractive interaction based on the proximity-force approximation (PFA).

One of the major goals of current research on Casimir forces revolves around the interesting question of whether the Casimir force can manifest itself in ways very different




Figure 7-17: Three-dimensional schematic of the Casimir “zipper” geometry of interlocking metal brackets (shown in different colors for illustration only), along with a two-dimensional xy cross-section. The dashed lines extruding from the plates to the squares indicate their out-of-plane connectivity.

from the usual monotonically decaying attractive force, and especially under what circumstances the force can become repulsive. It has been proven that the Casimir force is always attractive in a mirror-symmetric geometry of dielectric materials with $\varepsilon \geq 1$ on the imaginary-frequency axis [260], but there remains the possibility of repulsive forces in asymmetric structures and/or with different materials. For example, repulsive forces arise in exotic asymmetric material systems, such as a combination of theoretical magnetic and electric materials [54, 259, 485] (with some suggestions of metamaterials as route to realization [206, 297, 458]), fluid-separated dielectric plates [375], and possibly also in metamaterials with gain [297] or for excited atoms [491] (although this result of Ref. [297] is problematic because the Lifshitz formula may not be applicable to excited [491] or amplifying [423, 458] media). It has been suggested that vacuum-separated chiral metamaterials may exhibit repulsive interactions and stable repulsive/attractive transitions [595], although no specific metamaterial geometries (chiral or otherwise) exhibiting repulsion have yet been proposed—in any case, the predicted chiral repulsive forces arise only for small separations where the metamaterial approximation cannot be trusted, and recent exact calculations indicate that they appear to be attractive [343]. Moreover, recent theoretical work has shown that vacuum-separated objects can never form stable configurations [427].

Another route to unusual Casimir phenomena is to use conventional materials in complex geometries, which have been shown to enable asymmetrical lateral [139] or “ratchet”

effects [142, 362], in addition to the nonmonotonic dependence on external parameters of the previous sections [442]. Both of the previous lateral effects [139, 142] and the repulsive interaction described here intuitively arise from competing attractive contributions to the net force. The basic idea of using geometry to enable unusual Casimir phenomena provides an interesting supplement to the materials approach and could be useful in the design of experiments and technology involving the Casimir effect.

Here, we employ the FDFD method to study the geometry depicted schematically in Fig. 7-17: we have two periodic sequences of metal “brackets” attached to parallel metal plates, which are brought into close proximity in an interlocking “zipper” fashion. In Fig. 7-17, we have colored the two plates/brackets red and blue to distinguish them, but they are made of the same metallic material. This structure is not mirror symmetric (and in fact is glide-symmetric, although the glide symmetry is not crucial), so it is not required to have an attractive Casimir force by Ref. [260]. Furthermore, the structure is connected and the objects can be separated via a rigid motion parallel to the force (a consideration that excludes interlocking “hooks” and other geometries that trivially give repulsive forces). This structure is best understood by considering its two-dimensional cross-section, shown in Fig. 7-17(right) for the middle of the brackets: in this cross-section, each bracket appears as an $s \times s$ square whose connection to the adjacent plate occurs out-of-plane. (Here, the brackets are repeated in each plate with period $\Lambda = 2s + 2h$ and are separated from the plates by a distance d . The plates are separated by a distance $2d + s + a$, so that $a = 0$ is the point where the brackets are exactly aligned.) The motivation for this geometry is an intuitive picture of the Casimir force as an attractive interaction between surfaces. When the plates are far apart and the brackets are not interlocking, the force should be the ordinary attractive one. As the plates move closer together, the force is initially dominated by the attractions between adjacent bracket squares, and as these squares move past one another ($a < 0$ in Fig. 7-17), one might hope that their attraction leads to a net repulsive force pushing the plates apart. Finally, as the plates move even closer together, the force should be dominated by the interactions between the brackets and the opposite plate, causing the force to switch back to an attractive one. This intuition must be confirmed by an exact numerical calculation, however, because actual Casimir forces are not two-body attractions, are not generally additive, and can sometimes exhibit qualitatively different behaviors than a two-body model might predict [179, 259, 442, 590]. Such a computation of the total force

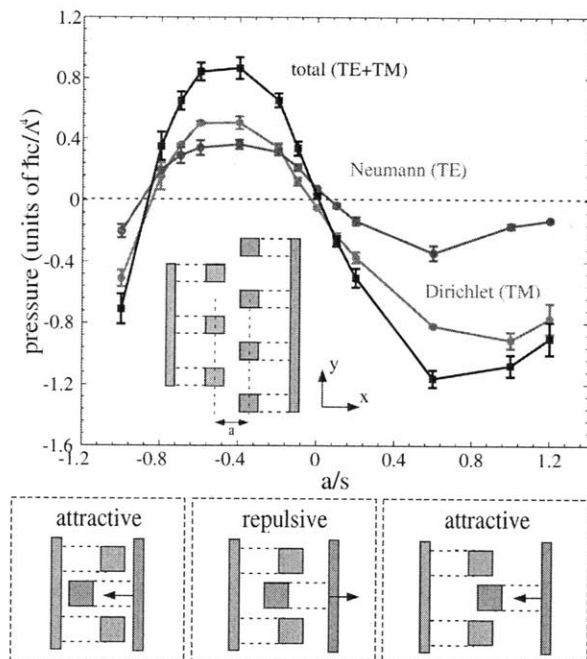


Figure 7-18: (Top:) Plot of the Neumann (blue, TE), Dirichlet (red, TM) and total (black, TE+TM) Casimir pressure (in units of $\hbar c/\Lambda^4$) between the objects of Fig. 7-17, as a function of a/s . The inset illustrates a two-dimensional cross-section. (Bottom:) Schematic indicating the various qualitatively different Casimir force regimes between the two structures.

per unit area is shown in Fig. 7-18, and demonstrates precisely the expected sign changes in the force for the three separation regimes. These results are discussed in greater detail below.

Previous theoretical studies of Casimir forces in geometries with strong curvature have considered a variety of objects and shapes. Forces between isolated spheres [143] and isolated cylinders [428], or between a single sphere [46], or cylinder [46, 141] and a metal plate, all exhibit attractive forces that decrease monotonically with separation. When a pair of squares [442] or cylinders [428] interacts in the presence of two adjacent metal sidewalls, the force is still attractive and monotonic in the square/square or cylinder/cylinder separation, but is a nonmonotonic function of the sidewall separation. When two corrugated surfaces are brought together in a way that breaks mirror symmetry (i.e., the corrugations are not aligned between the two surfaces), a lateral force can arise [140, 440], and an asymmetric lateral force from asymmetric corrugations can lead to a “ratchet” effect in which random forces preferentially displace the plates in one direction [142]. Such a lateral force has also been observed experimentally [88]. In the geometry of Fig. 7-17, in contrast, there is no

lateral force (due to a mirror-symmetry plane perpendicular to the plates), and hence we consider only the normal force between the plates. Because of the strong curvature of the surfaces relative to their separations, simple parallel-plate approximations are not valid (although we consider their qualitative accuracy below), and the force must be computed numerically.

7.4.1 Numerical Results

In order to simplify the calculations, we assume the length of the brackets in the z direction L to be sufficiently long to make their contributions to the force negligible (we estimate the minimum length below). We can therefore describe the geometry as both z -invariant and y -periodic (with period Λ). This implies that it is only necessary to compute the Green's function using an xy unit cell, with the periodic/invariant directions handled by integrating over the corresponding wavevectors [443]. Furthermore, we approximate the bracket/plate materials by perfect metals, valid in the limit of small lengthscales (which are dominated by long-wavelength contributions where the skin depth is negligible). In this case, the contributions to the force can be separated into two polarizations: transverse electric (TE) with the electric field in the xy plane (a scalar magnetic field with Neumann boundary conditions); and transverse magnetic (TM) with the magnetic field in the xy plane (a scalar electric field with Dirichlet boundary conditions) [443], and these two contributions are shown separately in Fig. 7-18.

The resulting force per unit area between the plates, for the chosen parameters $d/s = 2$ and $h/s = 0.6$, is plotted as a function of a/s in Fig. 7-18 (Top); error bars show estimates of the numerical accuracy due to the finite spatial resolution. A number of unusual features are readily apparent in this plot. First, the sign of the force changes not only once, but twice. The corresponding zeros of the force lie at $a/s \approx -0.8$ and $a/s \approx -10^{-2}$. The first zero, $a/s \approx -0.8$, is a point of unstable equilibrium, to the left of which the force is attractive and to the right of which the force is repulsive. The second zero at $a/s \approx -10^{-2}$ corresponds to a point of stable equilibrium, with respect to perpendicular displacements, for which the force is attractive to the right and repulsive to the left. (This point is still unstable with respect to lateral displacements, parallel to the plates and perpendicular to the brackets, however: any such lateral displacement will lead to a lateral force that pulls the red and blue brackets together.) In between these equilibria, the repulsive force has a local maximum at $a/s \approx$

−0.5. Finally, at $a/s \approx 0.6$ the magnitude of the attractive force reaches a local maximum (a local minimum in the negative force on the plot), and then decreases asymptotically to zero as $a/s \rightarrow \infty$. Thus, as the two objects move apart from one another, the force between them varies in a strongly nonmonotonic fashion (distinct from the nonmonotonic dependence on an external parameter shown in our previous work [428, 442, 590]). These three different sign regimes are shown schematically in Fig. 7-18 (Bottom), as predicted by the intuitive picture described above.

7.4.2 Comparison to PFA

Since the qualitative features of the Casimir force in this geometry correspond to the prediction of an intuitive model of pairwise surface attractions, it is reasonable to ask how such a model compares quantitatively with the numerical results. The most common such model is the proximity-force approximation (PFA), which treats the force as a summation of simple “parallel-plate” contributions [46]. (Another pairwise power-law heuristic is the “Casimir-Polder interaction” approximation, strictly valid only in the limit of dilute media [524].) Applied to a geometry with strong curvature and/or sharp corners such as this one, PFA is an uncontrolled approximation and its application is necessarily somewhat *ad hoc* (due to an arbitrary choice of which points on the surfaces to treat as “parallel plates”), but it remains a popular way to quantify the crude intuition of Casimir forces as pairwise attractions.

Applying the PFA approximation to the two objects in Fig. 7-17, we treat the net force as a sum of three contributions: the force between the two parallel plates, the force between each square and the opposite plate, and the force between adjacent red and blue squares. Namely,

$$P_{\text{PFA}} = \frac{1}{\Lambda L} (F_{\text{pl-pl}} + 2F_{\text{sq-pl}} + 2F_{\text{sq-sq}}), \quad (7.19)$$

where the first term is the pressure between two parallel plates (pl-pl), and the two remaining terms correspond to the square-plate (sq-pl) and square-square (sq-sq) interactions. The factors of Λ and L are introduced because these expressions are computed per unit length in the z direction, and per period in the y direction.

The first two PFA contributions are relatively simple to calculate because they are between parallel metal surfaces, and thus (in the PFA approximation) are the ordinary




Figure 7-19: Comparison of Casimir pressure (in units of $\hbar c/\Lambda^4$) as a function of a/s between the stress-tensor (exact) numerical results (black squares) and the proximity-force approximation (solid green). Also shown are the individual square–square (dashed blue) and square–plate (dashed orange) contributions to the PFA force. Inset: Schematic illustration of the chosen PFA “lines of interaction” between squares (dashed black lines).

Casimir force weighted by the respective areas:

$$P_{\text{pl-pl}} = -\frac{\hbar c \pi^2 h}{120\Lambda} \frac{1}{(2d + a + s)^4} \quad (7.20)$$

$$P_{\text{sq-pl}} = -\frac{\hbar c \pi^2 s}{240\Lambda} \frac{1}{(d + a)^4} \quad (7.21)$$

Computing the square–square force is less straightforward, since there is some ambiguity as to what the PFA approximation even means for two non-parallel surfaces (separate from the question of its accuracy). In PFA, one adds up “parallel plate” contributions to the force between two objects by including a force between each point on one surface and a corresponding point on the other surface, where corresponding points are connected by parallel “lines of interaction.” In this geometry, we take the lines of interaction to lie parallel to the center-to-center displacement between two squares, as depicted by the inset in Fig. 7-19, but of course this choice is somewhat arbitrary. (A similar choice was made by Ref. [106,107] to define the PFA force between two eccentric cylinders.) The PFA force

between one pair of squares is then:

$$\begin{aligned}
P_{\text{sq-sq}} = & - \frac{\hbar c \pi^2 a}{240 \Lambda D^5} \left\{ \left[\frac{2|a|}{3} (H^3 - 1) + \frac{sH^3}{h} (Hh - |a|) \right] \Theta(Hh - a) \right. \\
& \left. + \left[\frac{2Hh}{3} (A^3 - 1) + \frac{sA^3}{|a| - s} (|a| - Hh) \right] \Theta(|a| - Hh) \right\} \quad (7.22)
\end{aligned}$$

where $D \equiv \sqrt{a^2 + (h + s)^2}$, $H \equiv 1 + s/h$ and $A \equiv 1 - s/|a|$. The resulting net force is shown in Fig. 7-19, along with the contributions due to the isolated square-square and square-plate PFA forces (a separate line for the plate-plate contributions is not shown because this contribution is always very small).

For comparison, Fig. 7-19 also shows the exact total force from Fig. 7-18, and it is clear that, while PFA captures the qualitative behavior of the oscillating force sign, in quantitative terms it greatly overestimates the magnitude of the repulsive force. Of course, since it is an uncontrolled approximation in this regime there is no reason to expect quantitative accuracy, but the magnitude of the error illustrates how different the true Casimir force is from this simple estimate. The PFA estimate for the square-plate force, however, does help us to understand one feature of the exact result. If there were no plates, only squares, then the force would be zero by symmetry exactly at $a = 0$, and indeed the exact result including the plates has zero force at $a \approx 0$; clearly, the contribution to the force from the plates is negligible for $a \approx 0$, and this is echoed by the PFA $P_{\text{sq-pl}}$ force. Also, using a PFA approximation, one can attempt to estimate the order of magnitude of the force contribution from the ends of the bracket, which was neglected in the exact calculation. This contribution to the total force must decrease as $\sim 1/L$ for a fixed a , and is estimated to be less than 1% of the peak repulsive force for $L \gtrsim 60\Lambda$.

Because the basic explanation for the sign changes in the force for this structure is fundamentally geometrical, we expect that the qualitative behavior will be robust in the face of imperfect metals, surface roughness, and similar deviations from the ideal model here. The main challenge for an experimental realization (for example, to obtain a mechanical oscillator around the equilibrium point) would appear to be maintaining a close parallel separation of the brackets (although it may help that in at least one direction this parallelism is a stable equilibrium). Furthermore, although in this section we demonstrated one realization of a geometry-based repulsive Casimir force, this opens the possibility that future work will reveal similar phenomena in many other geometries, perhaps ones more

amenable to experiment.

Since the repulsion between the two objects here appears to be a result of attractive interactions between surfaces, our result also invites the following unanswered question: is it possible to design a geometry with a repulsive Casimir interaction, with ordinary vacuum-separated metallic materials, that *cannot* be understood as a combination of surface attractions? More precisely, can one obtain a repulsive Casimir interaction between a metallic/dielectric object lying in $x < 0$ and another lying in $x > 0$, i.e. two objects with a separating plane in vacuum, so that the distances between the surfaces are strictly nondecreasing as they are moved apart? The geometry here does not have a separating plane in the repulsive regime, while Ref. [260] only eliminated the mirror-symmetric case. It turns out, as mentioned in previous chapters, that the question of whether there exists a genuine repulsive Casimir force between vacuum-separated metals has been answered in the affirmative by Levin *et. al.* via a geometry consisting of an elongated spheroidal metallic object suspended on top of a metallic plate with an air hole [298].

7.5 Concluding Remarks

In this chapter, we presented a number of interesting results involving Casimir nonmonotonic effects and repulsion in complex geometries composed of vacuum-separated metallic objects. Our results are a consequence of the strong geometry dependence of the Casimir effect and shed light on some of the Casimir phenomena that can arise in complicated structures. As we have shown, the numerical methods developed in chapter 5 and 6 are instrumental for obtaining accurate predictions in these complicated geometries, especially when multi-body effects are present, where approximations may fail even qualitatively. There remain many complicated structures to be explored via these numerical methods: just recently, Ref. [298] predicted a novel repulsive Casimir force between vacuum-separated objects in which uniaxial repulsion and stability were found for plane-separated objects (in contrast to the repulsion in the previous section, which arose from a competition between attractive forces). The FDTD method of chapter 6 was employed in a number of the calculations in Ref. [298].

The last chapter of this thesis, chapter 8, examines the possibility of achieving stable suspension of objects via the Casimir force. The work of Ref. [427] demonstrates that stable

equilibria (in all directions) between vacuum-separated objects is impossible. However, as demonstrated in the subsequent chapter, suspension is possible when fluids are involved...

Chapter 8

Casimir Forces Between Fluid-Separated Dielectrics

“I don’t want to achieve immortality through my work. I want to achieve it through not dying.” *Woody Allen*

8.1 Overview

The Casimir force between vacuum-separated metallic or dielectric objects is most typically attractive and decaying as a function of object separation, an undesirable feature for experiments that may contribute to “stiction” in micro-electromechanical systems [83]. A repulsive interaction would be desirable to combat stiction as well as for frictionless suspension, among other applications. One route to obtaining repulsive forces is to employ complex geometries, such as the interleaved metallic structure of the previous chapter [446] or a recently-proposed geometry consisting of a small anisotropic (elongated) object suspended on top of a plate with a hole [298]. In this chapter, we shall focus on a different scheme, based on fluids.

The possibility of obtaining repulsion between fluid-separated asymmetric plates has been known for decades [129, 302]. More recently, due to increased interest on the application of Casimir forces in MEMS, it was suggested that repulsive Casimir forces between fluid-separated objects could lead to stable mechanical equilibria, and hence frictionless static bearings or other useful passive-suspension devices [74, 373, 374, 400]. However, until recently, calculations and experiments involving fluids only focused on achieving repul-

sion and were restricted to geometries involving parallel plates or simple approximations thereof [375] (similar to work on air-separated metals, as reviewed in several recent papers, e.g. Ref. [48]), leaving a wide range of fluid-based Casimir effects unexplored.

In this chapter, we examine a number of ways in which one can exploit fluids to obtain novel and interesting Casimir effects: these include stable suspension, aggregation of compact objects, and dramatic Casimir temperature effects. In Sec. 8.2, we briefly review the basic mechanisms enabling Casimir repulsion in fluids, based on the work of Ref. [77, 129, 302]. We consider three main kinds of stable-suspension schemes: First, Sec. 8.3 demonstrates stable suspension of z -invariant eccentric objects—the inner cylinder is confined in-plane by the repulsive Casimir force exerted by the outer cylinder. We investigate confinement for both circular and square cross sections, and show that the latter exhibits both a stable orientation and stable position. These calculations are performed using a modified version of the FDFD method of chapter 4 that allows calculations of Casimir torques for finite objects. In addition, the stable orientation of the square cylinder is shown to undergo a 45° transition as the separation lengthscale is varied, which is explained as a consequence of material dispersion; Second, Sec. 8.4.1 presents calculations of gravity-induced suspension—here, a sphere is suspended above a hemispherical indentation by a repulsive Casimir force, and the same repulsion is used to confine the sphere vertically against gravity [344]. Finally, Sec. 8.4 presents a scheme for obtaining suspension of dielectric nontouching objects—here, stability is induced by the dispersion properties of real dielectric (monolithic) materials. The consequence of this effect is the possibility of obtaining stable configurations (clusters) of compact objects, which is illustrated via a molecular two-sphere dycluster geometry consisting of two spheres levitated above a slab. Our calculations also reveal a strong interplay between material and geometric dispersion, and this is exemplified by the qualitatively different stability behavior observed in planar versus spherical geometries [450].

We believe that fluids are a promising setting to investigate new Casimir physics (of potential application in microfluidics). Not only do fluids enable suspension and passive control of compact objects, e.g. spheres, but they also provide new opportunities for exploring Casimir effects previously out of reach for experiments. For example, measuring the temperature dependence of the Casimir force has been an outstanding experimental problem—for typical vacuum-separated geometries, the T dependence of the Casimir force becomes observable only at micron separations (and above), where Casimir forces are rel-

atively weak and therefore difficult to observe experimentally. We address this problem in Sec. 8.5, which describes a method of achieving large temperature T sensitivity in the Casimir force that involves measuring the stable separation between dielectric objects immersed in a fluid [451]. In particular, we study the Casimir force between slabs and spheres using realistic material models, and find large > 2 nm/K variations in their stable separations (most typically in the hundreds of nanometers) near room temperature. In addition, we analyze the effects of Brownian motion on suspended objects, and show that their average separation is also sensitive to changes in T . Finally, we show that this approach also leads to rich qualitative phenomena, such as the existence of irreversible transitions, from suspension to stiction, as T is varied.

8.2 Repulsive Forces via Fluids

Most of the discussion in this section on repulsive Casimir forces is carefully laid out in [129, 304]; here, we only focus on the most basic and important results relevant to the work on stable suspension, the subject of this entire chapter.

The possibility of obtaining repulsive Casimir forces in fluids can be readily understood from the analogous Casimir-Polder force between polarizable atoms [77]. In particular, as reviewed in [129, 304] and in chapter 4, the potential U of two polarizable atoms with polarizability α_1 and α_2 is proportional to the product of their polarizabilities $U \propto \alpha_1\alpha_2$. A repulsive CP force arises if this product is negative or equivalently, if the atoms have *opposite* polarizability. Physically, this requires that there be an intermediate medium between the atoms that “screen” the charge on one atom and enhance the charge on the other. In the rarefied-medium (dilute) approximation, this conclusion can be rigorously extended to the case of macroscopic half-spaces [129], which we describe below using the Lifshitz formula.




Figure 8-1: Schematic of Lifshitz geometry, consisting of two semi-infinite plates (permittivities ϵ_1 and ϵ_2) separated by a dielectric with permittivity ϵ_3 .

A generalization of the standard perfectly-metallic vacuum-separated plates geometry,

first studied by Ref. [302], is the system shown schematically in Fig. 8-1: two semi-infinite dielectric objects of permittivities ϵ_1 and ϵ_2 , separated by a distance d , and embedded in a medium of dielectric ϵ_3 . A semi-analytical expression of the force in this case was derived by Ref. [302] and Ref. [129] in the 1950s by way of the mean stress-tensor formulation, and also more recently by way of the standard dispersive electromagnetic energy [529], applicable even in the presence of loss for subtle reasons ¹.

The Lifshitz theory provides a semi-analytical expression for the imaginary-frequency Casimir energy per unit area U of the system in Fig. 8-1 by expanding the Green's function of the entire system in terms of an exponentially-convergent (Fourier) basis. The result is an expression for U involving a double integral over imaginary frequencies ξ and perpendicular wave vector \mathbf{k}_\perp , given by [304, 426, 529]:

$$U = \frac{\hbar c}{2\pi} \int_0^\infty d\xi \int_0^\infty \frac{dk_\perp}{2\pi} k_\perp \log \prod_{q=s,p} \left(1 - r_1^q r_2^q e^{-2n_3 \xi d \sqrt{1+k_\perp^2/\epsilon_3 \xi^2}} \right), \quad (8.1)$$

where $k_\perp = |\mathbf{k}_\perp|$, $n_i = \sqrt{\epsilon_i \mu_i}$, and where the r_i^q are standard TE/TM (s/p) Fresnel coefficients [231]:

$$r_i^s(i\xi, x) = \frac{\mu_i(i\xi) - \mu_3(i\xi) \sqrt{1 + (n_i^2/n_3^2 - 1)x^2}}{\mu_i(i\xi) + \mu_3(i\xi) \sqrt{1 + (n_i^2/n_3^2 - 1)x^2}} \quad (8.2)$$

$$r_i^p(i\xi, x) = \frac{\epsilon_i(i\xi) - \epsilon_3(i\xi) \sqrt{1 + (n_i^2/n_3^2 - 1)x^2}}{\epsilon_i(i\xi) + \epsilon_3(i\xi) \sqrt{1 + (n_i^2/n_3^2 - 1)x^2}}, \quad (8.3)$$

evaluated at imaginary frequencies ($i\xi$) and for $x = \sqrt{1 + k_\perp^2/\epsilon_3 \xi^2}$. Equation (8.1) is often re-written via a simple change of variables $p = \sqrt{1 + \mathbf{k}_\perp^2/\epsilon_3 \xi^2}$ as:

$$U = \frac{\hbar c}{(2\pi)^2} \int_0^\infty (n_3 \xi)^2 d\xi \int_1^\infty p dp \log \left[\left(1 - r_1^s r_2^s e^{-2n_3 \xi p d} \right) \left(1 - r_1^p r_2^p e^{-2n_3 \xi p d} \right) \right], \quad (8.4)$$

¹The justifications for employing the standard dispersive forms of the energy and Lifshitz stress tensor expressions in the case of fluctuations are subtle [60, 354, 408]. Criticisms of these expressions have been raised based on the superficial similarity of the Lifshitz stress tensor to the Minkowski stress tensor, which is problematic in dissipative media [422, 424]. However, the justification for employing the Lifshitz theory and the use of the standard dispersive-energy formalism have been clearly laid out in Ref. [64, 408], and are based on equilibrium considerations of fluctuations at all frequencies [304].

which yields the following force $F = -\partial U/\partial d$ between the plates:

$$F = \frac{\hbar c}{2\pi^2} \int_0^\infty (n_3 \xi)^3 d\xi \int_1^\infty p^2 dp \sum_{q=s,p} \frac{r_1^q r_2^q \exp(-2n_3 \xi p d)}{1 - r_1^q r_2^q \exp(-2n_3 \xi p d)} \quad (8.5)$$

Equation (8.5) is a complicated expression of the dielectric properties of the materials and d , and therefore does not lend itself to a simple closed-form solution. Nevertheless, the expression does permit us to extract a great deal of insight². In particular, the sign of the integral of Eq. (8.5) depends on the product of the reflection (Fresnel) coefficients [Eqs. (8.2–8.3)] of the materials. The integrand is negative (leading to repulsion) at a frequency ξ if the reflection coefficients of the two half-spaces evaluated at ξ have *opposite* sign, analogous to the requirement that the two atoms (above) have opposite polarizabilities. By inspection of Eqs. (8.2–8.3) (assuming non-magnetic materials $\mu = 1$), the sign of the force integrand $f(\xi)$ will therefore depend on the ordering of the dielectrics as follows:

$$\text{sgn}(f(\xi)) = \begin{cases} -1, & \varepsilon_1(i\xi) < \varepsilon_3(i\xi) < \varepsilon_2(i\xi) \\ 1, & \text{otherwise,} \end{cases} \quad (8.6)$$

where $+/-$ denotes an attractive/repulsive force. Equation (8.6) can only be satisfied if the intervening medium ε_3 is a fluid—the alternative is that the medium be a solid, but then it becomes problematic to even define the force on objects 1 and 2. Because the Casimir force depends on all frequencies, from $\xi = 0$ up to some frequency $\xi \sim 2\pi c/a$, related to some lengthscale a (corresponding to a typical geometrical feature, see chapter 6), in order for the total force to be repulsive, the condition of Eq. (8.6) must be satisfied over a sufficiently large frequency bandwidth, as we shall demonstrate in the subsequent sections. (In the dilute-limit (rarefied-medium) approximation, it is possible to simplify Eq. (8.5) to obtain the original Casimir-Polder force between atoms, leading to the analogous condition for repulsion noted above [129].)

Repulsive forces in fluids have been measured in various experiments involving spheres

²The familiar expression for the Casimir pressure between two vacuum-separated perfectly-metallic plates can be arrived at by taking the limit as $\varepsilon_1, \varepsilon_2 \rightarrow \infty, \varepsilon_3 = 1$ in Eq. (8.5), to obtain $F = \hbar c/240\pi^2 d^4$. Note that Eq. (8.5) yields a non-zero force between lossless materials (e.g. perfect metals). This is not an unphysical result, nor does it violate the fluctuation-dissipation theorem, since Eq. (8.5) is the result of assuming some infinitesimal amount of loss in the system. Mathematically, this is equivalent to setting $\varepsilon \rightarrow \varepsilon + i\delta$ inside the integral (for a lossless dielectric, i.e. $\text{Im} \varepsilon(\xi) = 0$), and then taking the limit as $\delta \rightarrow 0$ outside of the integral. It is this procedure that allowed Ref. [302] to equate the real and imaginary-frequency integrals (Wick rotation), by pushing all of the poles of the system below the upper-half complex- ξ plane.

and plates, both in the short-range (van der Waals) [152, 292, 293, 349, 351] and long-range [374] regimes. Fluid experiments on macroscopic objects such as those of Ref. [374] involve capacitive measurements of a sphere attached to a cantilever and suspended on top of a slab. While these types of experiments represent a significant step toward the development of tools and ways in which to manipulate objects via the Casimir effect (with potential applications in microfluidics and MEMS), they leave something to be desired: a tantalizing possibility, yet unexplored by experiments and recently suggested by a number of theoretical calculations [344, 429, 450, 452] is the possibility of achieving stable (passive) suspension of objects immersed in fluids, the subject of the remainder of this chapter.

8.3 Eccentric Cylinders: Uniaxial Stable Suspension

We will first focus on a scheme for obtaining uniaxial stable suspension that is based on the repulsive effect of the previous section (for fluid-separated objects) and that exploits the strong geometry dependence of the Casimir force. Using recently developed numerical techniques [443], we present theoretical calculations of Casimir forces/torques between fluid-separated objects with finite square and circular cross sections (Fig. 8-2 insets) that rigorously demonstrate stable positional/orientational Casimir equilibria. (Vacuum-separated perfect-metal circular cylinders were shown to exhibit an unstable positional equilibrium [107].) In the case of square cross sections, a surprising result is obtained for the rotational equilibria: the stable orientation changes by 45° depending on the length-scale, which we show to be a consequence of material dispersion. In particular, for certain fluid-separated materials, the frequency dependence of the permittivity ε (material dispersion) causes the Casimir force to switch from repulsive to attractive at some critical separation [129]; this leads to the orientation transition described here, and also produces other lengthscale-based qualitative transitions, as shown in the subsequent section. In what follows, We apply the FDFD method of chapter 5 to compute Casimir torques on finite objects (in contrast to previous uncontrolled approximations [471] or analytical results for planar cases [373, 400, 401, 431, 440, 484, 540]); we supplement the accurate results with a heuristic model based on the proximity force approximation (PFA) that turns out to capture qualitative behaviors of the orientation transition, and therefore provides some simple insight.




Figure 8-2: Casimir force F_x in the x -direction, per unit z -length, on a SiO_2 cylinder suspended within a perfectly-metallic cylinder (inset), separated by fluid (ethanol), as a function of the x -displacement from equilibrium (eccentricity) d [in units of $a = 0.5(D - s)$] for both circular (solid-line) and square (dashed-line) cylinders. $d = 0$ is seen to be a stable equilibrium.

As noted in the previous section, a repulsive Casimir force arises between parallel plates of $\epsilon_{1,2}$ separated by a fluid of ϵ_f , if $\epsilon_1(i\xi) < \epsilon_f(i\xi) < \epsilon_2(i\xi)$ for a sufficiently wide range of imaginary frequencies $\xi = \text{Im } \omega$ [129]. Three such materials are silica (SiO_2) and metal (Au) separated by ethanol [74], as discussed below. We studied the three-dimensional constant cross-section geometries shown in the inset of Fig. 8-2: square or circular SiO_2 cylinders of diameter s surrounded by a metal cylinder of diameter D , separated by ethanol. For computational ease, we use a perfect metal [$\epsilon(i\xi) \rightarrow \infty$] for the outer cylinder, but we use experimental ω -dependent ϵ for the SiO_2 and ethanol. In such cases, with object sizes and curvatures comparable to their separation, approximations of the Casimir force as a pairwise attraction between surfaces are not valid and their qualitative predictions (e.g. stability) may be incorrect [106, 179, 443].

Figure 8-2 shows the force per unit z -length on the inner cylinder as a function of the displacement d from equilibrium in units of a , where $a \equiv 0.5(D - s) = 0.0955 \mu\text{m}$ so that $d/a = \pm 1$ for touching surfaces, for parameters $s/D = 0.25$. It demonstrates a stable equilibrium for both square (solid black) and circular (dashed grey) cross sections. The computational method is based on integration of the mean electromagnetic stress tensor (valid even for fluids [408]) evaluated in terms of the imaginary-frequency Green's function via the fluctuation-dissipation theorem [443]. Additional details are provided below.

Having demonstrated positional stability (Fig. 8-2), we now explore an unusual effect

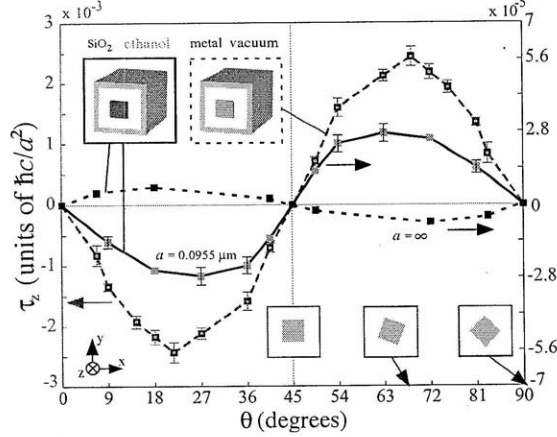


Figure 8-3: Casimir torque τ_z , per unit z -length, on the inner square rod of eccentric square cylinders as a function of the angle θ with respect to the x -axis (see right insets) for two material choices as shown in the top insets. Error bars are estimates of the effect of finite grid resolution.

arising from material dispersion. The square-cylinder geometry (Fig. 8-3 inset) must, by symmetry, exhibit two equilibrium orientations ($\theta = 0^\circ$ and $\theta = 45^\circ$). *A priori*, it is not clear which of the two configurations is stable. To determine this, Fig. 8-3 plots the Casimir torque per unit z -length on the inner square, as a function of the rotation angle θ between the inner and outer squares. In addition to computing the torque for the SiO₂-ethanol-perfect-metal case, we also analyze the torque for vacuum-separated perfect-metal cylinders, in which case the forces are purely attractive and there is no stable positional equilibrium [107].

The resulting torque per unit z -length is shown in Fig. 8-3, for both the repulsive (solid, filled-square-lines) and attractive (solid, open-square-line) cases, as θ is varied from -45° to 45° . The stable orientation for both cases is $\theta = 45^\circ$, a surprising result considering the difference in the sign of the force. A stable orientation of 45° in the perfect-metal case is not so surprising, since an attractive force that decreases with separation should intuitively (in the heuristic picture of pairwise attractions between surfaces) favor an orientation where the surfaces are as close as possible, pulling the corners of the inner square towards the outer surfaces. On the other hand, a repulsive force should intuitively push the surfaces as far apart as possible, which would suggest a 0° stable orientation. The reason for this apparent contradiction between intuition and the numerical results lies in the effects of material dispersion—as explained below, $\epsilon_{\text{Si}} - \epsilon_{\text{eth}}$ (and hence the force) switches sign at some lengthscales (some $1/\xi$). For example, if we neglect the material dispersion of SiO₂

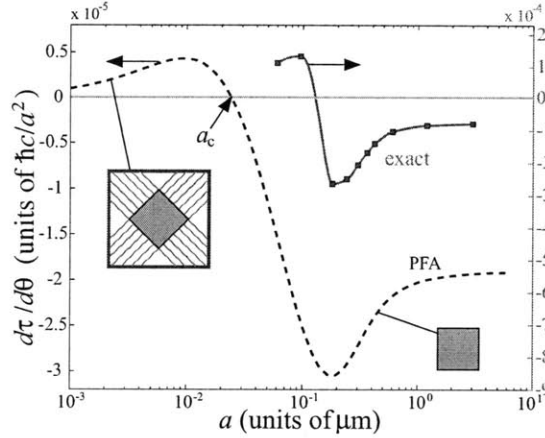


Figure 8-4: Derivative of the Casimir torque in the z -direction with respect to θ , $d\tau_z/d\theta$, in units of $\hbar c/a^2$, evaluated at $\theta \approx 0^\circ$, as a function of lengthscale $a = 0.5(D - s)$. The solid line is a fit to the exact Casimir torque, computed by our numerical method (solid squares), and the dashed line is the torque as computed by a PFA approximation. Both curves display a change (dispersion-induced transition) in the stable orientation of the square.

and ethanol, and use only the $\xi \rightarrow 0$ dielectric constants, the torque as a function of θ (shown as the dashed black line in Fig. 8-3) indeed exhibits a stable orientation at 0° , as expected from the pairwise intuition: in this case, the force is repulsive at *all* lengthscales.

Since $\theta = 0^\circ$ is unstable for the $a = 0.0955 \mu\text{m}$ fluid case and stable for $a \rightarrow \infty$ (equivalent to $\xi \rightarrow 0$), there *must* be a transition at some critical intermediate lengthscale a_c . One way to determine a_c is to calculate the derivative of the torque at $\theta = 0^\circ$ as a function of a (noting that the dimensionless torque $\tau a^2/\hbar c$ and its derivative go to a nonzero value as $a \rightarrow \infty$), and to look for a change in the sign. This derivative is plotted as the solid line in Fig. 8-4 and displays the expected transition from unstable ($d\tau_z/d\theta > 0$, left) to stable ($d\tau_z/d\theta < 0$, right) at $a_c \approx 0.1 \mu\text{m}$, a consequence of material dispersion. A better understanding of this transition can be gained by inspecting a simple heuristic (PFA). PFA is only an *ad hoc* model, in which the force on each point of the surface (and hence the torque) is treated as simply the parallel-plate (Lifshitz) force between fluid-separated half-spaces (with “lines of interaction” perpendicular to the inner square, as depicted in the left inset of Fig. 8-4). This simple model turns out to capture some qualitative features of the orientation transition, as shown by the dashed line in Fig. 8-4, although it is of course quantitatively incorrect [440].

PFA provides an explanation for why orientational stability need not coincide with positional stability. Stability, in general, arises from the competing interactions of the inner

and outer surfaces, e.g. in the 45° orientation for repulsive interactions, the nearest surface “pushes” the corner away, while the other surface “pushes” the corner back. Which of these competing forces dominates depends on their sign and power law, but in the torque $\boldsymbol{\tau} = \mathbf{r} \times \mathbf{F}$ case, there is an additional effect (from the $\mathbf{r} \times$ dependence): the force “from” the nearest surface elements is parallel to the radial direction and does *not* produce a torque, whereas it *does* contribute to the total force. Therefore, the net force and torque can respond differently to the same material dispersion, because the lengthscales of the dominant contributing separations differ. More generally, in the exact model where the force cannot be decomposed into additive contributions, the dominant imaginary-frequency contribution can differ between torque and force.

There is another interesting feature in Fig. 8-4: for $a > a_c$, the derivative of the torque is nonmonotonic, *decreasing* in magnitude towards $a \rightarrow \infty$. This may seem counter-intuitive because the ε contrast is maximum for $\xi \rightarrow 0$ (Fig. 8-5 inset), which would seem to predict greater forces. However, because the net torque arises from a competition between nearby and faraway surfaces, the magnitude depends not only on the force but also on the power law: if the force decreases more rapidly with distance, then faraway surfaces contribute less and the net torque is larger. This is precisely what is happening here: at intermediate a , material dispersion acts to increase the power law [129] compared to the $a \rightarrow \infty$ case where dispersion plays no role. Preliminary work shows similar effects in the positional stability. This also explains the relative torque magnitudes for $a \rightarrow \infty$ and $a = 0.0955 \mu\text{m}$ in Fig. 8-3.

Since the qualitative behavior of the orientation transition is captured by PFA, it makes sense to examine a simple one-dimensional model in detail, depicted on the inset of Fig. 8-5: a SiO_2 slab of thickness s , separated from two metallic half-spaces, a (surface-to-surface) distance $a \equiv 0.5(D - s)$ from the SiO_2 surface, with ethanol (fluid) in between. This geometry is analyzed using a generalization of the Lifshitz formula [529]. Figure 8-5 shows the force per unit area on the inner slab (again using $s/D = 0.25$), as a function of the displacement d , for two lengthscales: $a = 0.0955 \mu\text{m}$ (red), and $a = 0.6 \mu\text{m}$ (black). There is a stable equilibrium at $d = 0$, just as for the two-dimensional case in Fig. 8-2. In the case of perfect-metal half-spaces, the force changes sign at a critical distance d_c that depends on the lengthscale a . However, if we use a realistic metal (Au), shown by the dashed line for $a = 0.0955 \mu\text{m}$, this transition disappears and the force is always repulsive. These features are explained by the inset of Fig. 8-5, which shows $\varepsilon(i\xi)$ for SiO_2 , ethanol,




Figure 8-5: Casimir force F_x in the x -direction, per unit area, between a planar SiO_2 slab suspended between two perfect-metal plates (solid red and black lines) or gold half-spaces (dashed red line), separated by a fluid (ethanol), as a function of the dimensionless x -displacement d from equilibrium. The force is plotted also at two lengthscales $a \equiv 0.5(D - s) = 0.0955 \mu\text{m}$ (red lines) and $a = 0.6 \mu\text{m}$ (black line). The solid blue line shows the same quantity for the case of a lithium-niobate slab and a gold half-space, for $a = 0.5 \mu\text{m}$. (Insets:) Top inset: $\epsilon(i\xi)$ for SiO_2 (black), LiNbO_3 (solid and dashed blue), ethanol (orange) and gold (red) as a function of imaginary frequency ξ . Bottom inset: schematic of geometry.

and Au. The key point is that the contributions to the Casimir force come primarily from imaginary “wavelengths” $2\pi/\xi$ larger than some lengthscale set by the separation, while very short wavelengths (large ξ) on this scale are exponentially cut off in the force integral [303]. Thus, for large lengthscales a , the force is dominated by small- ξ contributions where $\epsilon_{\text{silica}} < \epsilon_{\text{ethanol}} < \epsilon_{\text{metal}}$, giving a repulsive force and a stable equilibrium. On the other hand, for small lengthscales and separations, large- ξ contributions become more important, for which $\epsilon_{\text{silica}} > \epsilon_{\text{ethanol}}$ (for $\xi > 2.3 c/\mu\text{m}$, marked by the vertical dashed line in the inset of Fig. 8-5), leading to attractive forces. For the case of perfect-metal cylinders, these attractive contributions are large enough to flip the sign of the total force for small separations, whereas for Au the attractive high- ξ contributions are suppressed by the diminishing $\epsilon(i\xi)$ of the Au and the force remains repulsive. While the transition is absent in the SiO_2 -ethanol-Au configuration, we have calculated the force between ethanol-separated lithium-niobate LiNbO_3 and Au plates and found a sign transition at separations $\sim 0.05 \mu\text{m}$. This is illustrated in Fig. 8-5 (solid blue line); because LiNbO_3 's anisotropy greatly complicates the modeling, we used a PFA approximation based on summing the Lifshitz forces between semi-infinite Au and LiNbO_3 slabs (neglecting the finite LiNbO_3 thickness). The two principal values of the ϵ tensor of LiNbO_3 [34] are plotted in the

inset. There are also transitions for ethanol-separated barium-titanate and calcite plates at separations $\sim 0.01 \mu\text{m}$ [373]. Since these two cases are much more difficult to compute (the outer cylinder is not perfect metal and LiNbO_3 is anisotropic), we focused on the SiO_2 -ethanol-metal case here, which has similar qualitative behaviors.

The experimental ε s of SiO_2 and ethanol were fit to a standard multiple-oscillator model [325] accurate from infrared to ultraviolet wavelengths: $\varepsilon(i\xi) = 1 + \sum_{n=1}^N C_n \left[1 + (\xi/\omega_n)^2 \right]^{-1}$. The parameters we used are [34, 351]: ethanol ($N = 2$) $\omega_n = \{6.6, 114\} \times 10^{14}$ Hz and $C_n = \{23.84, 0.852\}$; SiO_2 ($N = 3$) $\omega_n = \{0.867, 1.508, 203.4\} \times 10^{14}$ Hz and $C_n = \{0.829, 0.095, 1.098\}$. We model the dielectric constant of Au by a Drude model $\varepsilon(i\xi) = 1 + \omega_p^2/\xi(\xi + \gamma)$, where $\omega_p = 1.367 \times 10^{14}$ Hz and $\gamma = 5.320 \times 10^{13}$ Hz [34]. The Casimir force \mathbf{F} is computed as an integral $\mathbf{F} \sim \int_0^\infty d\xi \oint \langle \mathbf{T} \rangle d\mathbf{A}$ [443], where the mean stress tensor $\langle \mathbf{T} \rangle$ is computed at each position and frequency from the Green's function as described in Ref. [443]. The Casimir torque $\boldsymbol{\tau}$, was computed in the same framework, with a minor modification, $\boldsymbol{\tau} \sim \int_0^\infty d\xi \oint \mathbf{r} \times (\langle \mathbf{T} \rangle d\mathbf{A})$, also proposed in our earlier work [443].

We are hopeful that these phenomena will be amenable to experiment. There are some experimental advantages over planar geometries: the concentric configuration is a stable equilibrium, static charges on the outer metallic cylinder are screened in the interior, and the equilibrium is easily distinguished from electrostatic effects (which cannot produce stability [106]). (If we inverted the geometry, to have the SiO_2 on the outside and the metal on the inside, then the larger static ε of the fluid compared to the SiO_2 would lead to a classical stable equilibrium for a charged inner cylinder [92, 106, 353].) We expect that similar stability will be obtained for real metals, as in our Au calculation in Fig. 8-5, and there are several materials that exhibit repulsive-attractive transitions (unlike SiO_2 -ethanol-Au) that should display the orientation transition. We have also uncovered an intriguing question to explore in future work: in the PFA heuristic, it appears that the critical lengthscale at which the orientation transition occurs is *smaller* than the corresponding positional transition lengthscale (changing from stable to unstable suspension), whereas the exact calculations give the opposite result.

Stable suspension and dispersion-induced transitions are certainly not confined to the z -invariant structures studied in this section. In fact, subsequent to this work, Ref. [429] explored a geometry involving eccentric spheres embedded in a fluid that also exhibits a stable configuration. As opposed to the saddle-point equilibrium described here (stable




Figure 8-6: Schematic of silica sphere levitated above a gold plate. At the equilibrium point, the force of gravity counters the Casimir force, while the Casimir force from the walls of the spherical indentation confine the sphere laterally (x - y plane).

suspension restricted to the x - y plane), the eccentric-sphere structure of Ref. [429] exhibit an honest-to-god stable equilibrium (suspension in all directions).

8.4 Stable Suspension and Confinement

In this section, we consider two types of fluid-based suspension schemes. First, we consider a scheme explored in Ref. [344] using the FDTD method of chapter 6, and presented schematically in Fig. 8-6: a dielectric sphere is suspended above a gold plate and confined laterally (in the x - y plane) by a spherical indentation and vertically (in the z direction) by a competition between a repulsive Casimir force and gravity. We then describe a method of obtaining stable suspension of compact objects that is entirely based on material considerations (independent of geometry or external forces) and that suggests an approach for obtaining non-touching configurations, or diclusters, of compact objects.

8.4.1 Gravity-Induced Stability

This section demonstrates stable confinement via a combination of geometric and material dispersion coupled with gravity: a dielectric sphere is levitated above a gold plate with hemispherical indentation, as in Fig. 8-6. The basic geometry considered here, illustrated in Fig. 8-6 and studied using the FDTD method of chapter 6, was first proposed by Ref. [62] in the unstable (vacuum) case. With this example, we present a setup similar to that used previously to measure repulsive Casimir forces [374], with the hope that this system may be experimentally feasible.

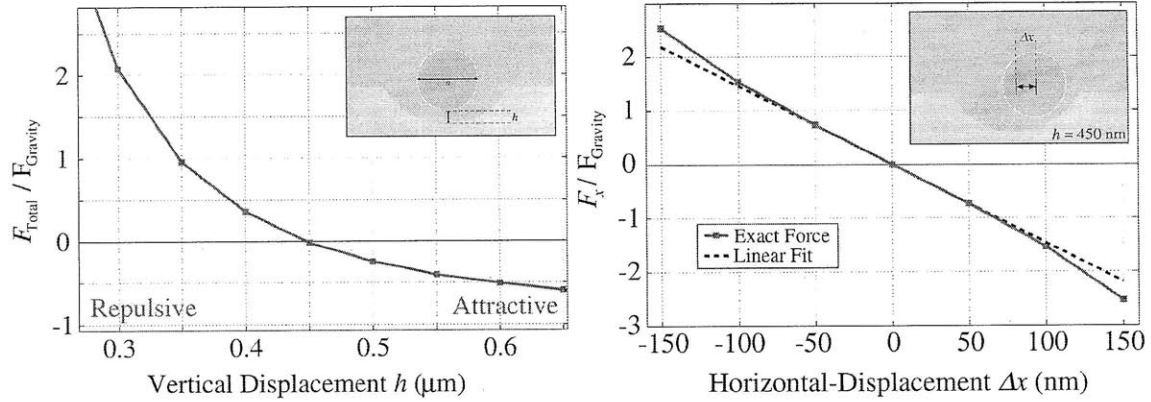


Figure 8-7: (Left:) Total (Casimir + gravity) vertical (z) force on the silica sphere (depicted in the inset) as the height h of the sphere's surface above the indentation surface is varied. The point of vertical equilibrium occurs at $h_c \sim 450$ nm. (Right:) Casimir restoring force on the sphere as a function of lateral displacement Δx , when the vertical position is fixed at h_c , the height at which gravity balances the Casimir force.

A silica sphere is suspended above a perfect metal plane with a spherical indentation in it. The sphere is immersed in bromobenzene. Because the system satisfies $\epsilon_{\text{sphere}} < \epsilon_{\text{fluid}} < \epsilon_{\text{plane}}$, the sphere feels a repulsive Casimir force (in the $+z$ direction) [374]. This is balanced by the downward force of gravity, which confines the sphere vertically. In addition, the Casimir repulsion from the sides of the spherical indentation confine the sphere in the lateral (x - y plane) direction. The radius of the sphere is $R = 1\mu\text{m}$, and the circular indentation in the metal is formed from a circle of radius $2\mu\text{m}$, with a center $1\mu\text{m}$ above the plane. For computational simplicity, in this model we neglect dispersion and use the zero-frequency values for the dielectrics, as the basic effect does not depend upon the dispersion (the precise values for the equilibrium separations will be changed with dispersive materials). These are $\epsilon = 2.02$ for silica and $\epsilon = 4.30$. For the gravitational force we use densities of 1.49 g/cm^3 for bromobenzene and 1.96 g/cm^3 for silica.

An efficient strategy to determine the stable point is to first calculate the force on the glass sphere when its axis is aligned with the symmetry axis of the indentation. This configuration is cylindrically-symmetric and can be efficiently computed as in the previous section. Results for a specific configuration, with a sphere radius of 500 nm and an indentation radius of $1\mu\text{m}$, are shown in Fig. 8-7. The force of gravity is balanced against the Casimir force at a height of $h = 450\text{ nm}$. To determine the strength of lateral confinement, we perform a fully three-dimensional computation in which the center of the sphere is displaced laterally from equilibrium by a distance Δx (the vertical position is held fixed at the equilibrium




Figure 8-8: Schematic schemes for stable suspension of fluid-separated objects, involving: (a) enclosed geometries; (b) gravity countering Casimir repulsion; and (c) material dispersion producing repulsive *and* attractive Casimir forces (here).

value $h = 450\text{nm}$). The results are shown in Fig. 8-7(Right). It is seen that over a fairly wide range ($|\Delta x| < 100\text{ nm}$) the linear term is a good approximation to the force, whereas for larger displacements the Casimir force begins to increase more rapidly. Of course, at these larger separations the vertical force is no longer zero, due to the curvature of the indentation, and so must be re-computed as well.

8.4.2 Nontouching Plane-Separated Compact Objects

In this section, we demonstrate stable Casimir suspension of realistic dielectric/metallic objects immersed in a fluid. Unlike suspension in the previous sections (explored in Ref. [344, 452]), this suspension does not involve one object enclosing another or external forces, but instead occurs between objects on opposite sides of an imaginary separating plane. This effect is a consequence of material dispersion, and is here validated in various experimentally accessible geometries consisting of ethanol-separated dielectric spheres and semi-infinite slabs. Furthermore, we show the possibility of achieving Casimir “molecular” clusters in which objects can form stable non-touching configurations in space—this is illustrated in a “diatomic” or “dicluster” geometry involving two dielectric spheres of different radii bound into a non-touching pair and levitated above a gold slab. Finally, our calculations reveal interesting effects related to the interplay of geometric and material dispersion, in which stability responds to finite size in a way that is qualitatively different in planar vs spherical geometries.

Prior to this work, Casimir stability had only been studied in the following four different contexts: first, in geometries involving mutually enclosed fluid-separated objects, in which the inner object is repelled by the outer object [452]; second, the slab–sphere geometry of the previous section, in which fluid-induced repulsion counteracts the force of gravity [344]; third, interleaved structures like the zipper geometry of the previous chap-

ter, studied by Ref. [446], in which the surfaces of two complicated objects interleave so that their mutual attraction acts to separate the objects; and fourth, metamaterial proposals [595] that currently have no clear physical realization. The first two approaches (involving fluids) are illustrated by the schematics in Fig. 8-8(a–b). While all of these examples clearly demonstrate the possibility of Casimir stability, they leave something to be desired: they are limited to enclosed/complex geometries or require that stability lie only along a single direction (e.g. direction of gravity). A less constrained and previously unexplored form of stability is one involving compact objects on either side of an imaginary separating plane, as illustrated in Fig. 8-8(c) for two spheres: in this case (involving fluids), we will show that the objects form stable configurations that are independent of external forces, and seem more accessible to experiment, opening up new possibilities for the creation of multi-body clusters based on the Casimir force.

From our discussion in the previous sections, we know that Casimir force between two dielectric objects embedded in a fluid can become repulsive if their dielectric permittivities satisfy Eq. (8.6) over a sufficiently wide range of imaginary frequencies ξ [129]. The possibility of stable separations arises if the force transitions from repulsive at small separations (conceptually dominated by large- ξ contributions) to attractive at large separations (conceptually dominated by small- ξ contributions). A criterion for obtaining stability is therefore that Eq. (8.6) be violated at small ξ , and satisfied for $\xi > \xi_c$, with the transition occurring at some $\xi_c \sim 2\pi c/\lambda_c$ roughly related to the lengthscale λ_c at which the repulsive/attractive transition occurs. This criterion is only heuristic, but helps guide our intuition. The real system is more complicated, as we shall see, because the sign of the force also depends on many other factors such as the relative strength of the contributions of different frequencies (related to the strength of the ε contrast) as well as on finite-size effects.

After considering a number of possibilities, we identified several material combinations that satisfy Eq. (8.6) for large ξ (small separations). Figure 8-9(top) plots the dielectric permittivity $\varepsilon(i\xi)$ of various materials (Si, teflon, SiO₂, and ethanol) satisfying Eq. (8.6) over some large range of ξ . In order to establish the existence of stable separations, we first compute the force between semi-infinite slabs separated by ethanol, using the Lifshitz formula [304]. Figure 8-9(bottom) plots the Casimir force between semi-infinite slabs for different material arrangements, normalized by the corresponding force between perfectly-

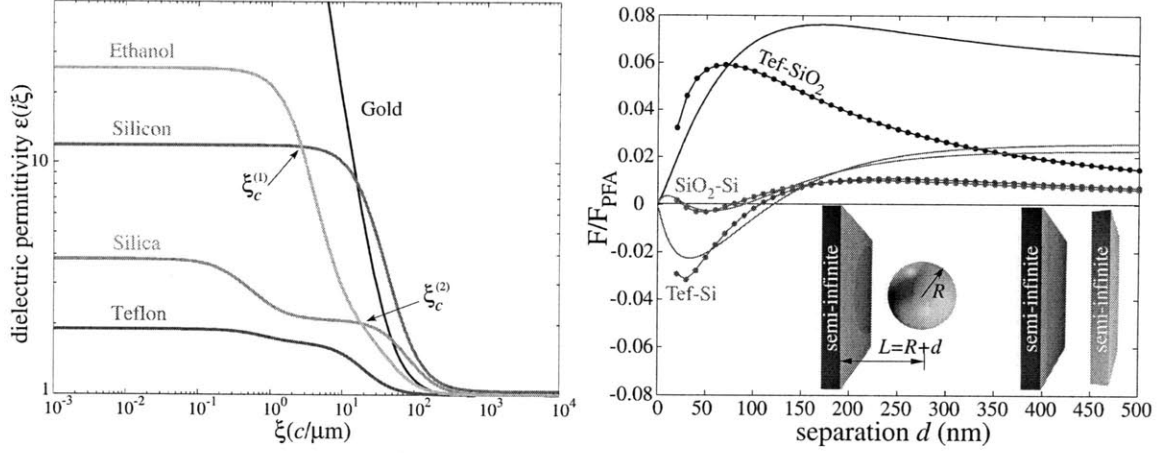


Figure 8-9: (Top): Plot of the dielectric permittivity $\varepsilon(i\xi)$ of various materials evaluated at imaginary frequency ξ (units of $c/\mu\text{m}$). (Bottom): Casimir force between a semi-infinite slab and a sphere (dots) and Casimir pressure between semi-finite slabs (lines), normalized by the corresponding perfect-metal PFA force $F_{\text{PFA}} = \hbar c \pi^3 / 720 d^3$ (slab-sphere) and pressure $F_{\text{PFA}} = \hbar c \pi^2 / 240 d^4$ (slab-slab). The force/pressure is plotted for various material configurations, described in the text.

metallic slabs. Our results show that both teflon-Si (yellow) and SiO_2 -Si exhibit stable equilibria at $d_c \approx 120\text{nm}$ and $d_c^{(s)} \approx 90\text{nm}$, respectively. SiO_2 -Si exhibits a finite region of stability coming from the existence of an *unstable* equilibrium at a smaller $d_c^{(u)} \approx 29\text{nm}$, a consequence of its two dielectric crossings $\xi_c^{(1)} \approx 2.6 \cdot 2\pi c/\mu\text{m}$ and $\xi_c^{(2)} \approx 26.4 \cdot 2\pi c/\mu\text{m}$ labeled in Fig. 8-9(top). As mentioned above, this prescription for obtaining stability is only heuristic: for example, the force between teflon- SiO_2 slabs is always attractive, even though their permittivities satisfy Eq. (8.6) at large ξ . In that case, while there is a crossing of the form of Eq. (8.6) at $\xi_c^{(2)}$, the repulsive contributions coming from $\xi > \xi_c^{(2)}$ are overwhelmed by the attractive contributions coming from $\xi < \xi_c^{(2)}$, since SiO_2 and teflon become transparent at relatively small ξ . Thus, the repulsive region of the frequency spectrum merely *reduces* the attractive force between the objects at small separations.

The existence of stable separations for semi-infinite slabs is promising, especially because this occurs in the 100nm range where Casimir forces are easily observable, leading us to investigate whether similar phenomena occur for finite-size objects: finite-thickness slabs, slab-sphere, and sphere-sphere configurations.

For slab-sphere and sphere-sphere geometries, rapid exact calculations are performed using the spherical-harmonic scattering formulation of [73, 426]. Given the material properties, the scattering-matrix formulation (derived from a path-integral evaluation of the

Casimir energy) yields the exact Casimir force, with the only approximation being the numerical truncation of the sum over spherical-harmonic contributions—these contributions decay exponentially fast, and we only required harmonics up to order $\ell = 20$ to obtain better than 1% accuracy in the cases studied here. For convenience, we evaluated the force at zero temperature T ; because the room-temperature Matsubara wavelength $\hbar/kT \sim 50 \mu\text{m}$ is much larger than the $< 200 \text{ nm}$ separations considered here, the finite-temperature corrections to the equilibrium separations are small [304], e.g. $\ll 1\%$ for the $R_S = 100 \text{ nm}$ teflon–Si sphere–sphere case in Fig. 8-11. Intuitively, one might expect the finite size or thickness of an object to suppress the contributions from small ξ (large “wavelengths”), and therefore to change (or eliminate) the separation at which the repulsive/attractive transition occurs. Here, where attraction comes from small- ξ contributions, one might expect the finite size to decrease the attractive contributions and therefore increase the equilibrium separation d_c .

Figure 8-9(bottom) plots the force between a sphere of radius $R = 200\text{nm}$ and a semi-infinite slab for different material arrangements, normalized by the corresponding perfect-metal proximity-force approximation (PFA) force. (When referring to a geometry consisting of a semi-infinite α and finite β object, we shall denote the combination by α – β . For example, teflon–SiO₂ refers to a semi-infinite teflon slab and finite SiO₂ sphere configuration.) The results in this slab–sphere geometry look qualitatively similar to those in the slab–slab structure. In particular, both teflon–Si and SiO₂–Si exhibit stable equilibria at $d_c \approx 105\text{nm}$ and $d_c \approx 78\text{nm}$, respectively, roughly 15nm smaller than the d_c in the slab-slab case.

The fact that d_c is smaller in the slab–sphere case than for semi-infinite slabs was initially unexpected since it contradicts the intuition described above. However when we plot d_c for the various materials (solid dotted lines) as a function of $R \in (0, 350)\text{nm}$ in Fig. 8-10, we indeed observe the expected behavior: as R decreases, d_c increases, asymptoting to a constant at $R = 0$ that corresponds to the Casimir–Polder force between a spherical nanoparticle and a slab [426]. Similar increases in d_c as thickness t is decreased are observed in *some* of the finite-slab geometries (solid lines) in Fig. 8-10, at least in all of the configurations where the thickness of the silicon is varied. In the $t \rightarrow \infty$ limit for the slab–slab case, the semi-infinite result is recovered. For the slab–sphere case with $R \rightarrow \infty$, the asymptotic d_c occurs for a smaller separation than for semi-infinite slabs: in this limit, where PFA is valid, the curvature of the spheres yields an average separation that is larger than d_c .




Figure 8-10: Equilibrium separation d_c vs. sphere-radius (R) or slab-thickness (t) for various slab-sphere (dots) and slab-slab (lines) configurations. Dashed lines correspond to slab-slab geometries with materials swapped (semi-infinite silicon slab).

An interesting question is to what extent one may tailor object geometries in order to change the qualitative features of the force, e.g. its sign. With this question in mind, the dashed lines in Fig. 8-10 shows d_c for slab-slab geometries in which the two solid materials have been exchanged so that the finite-thickness slab is no longer silicon. Apparently, changing which slab is finite qualitatively reverses the dependence on d_c in some cases: for Si-Au and Si-SiO₂ slabs, the equilibrium separation d_c *decreases* with decreasing thickness t , corresponding to an *increased* attractive force despite the fact that the attractive contributions arise from small ξ (which are intuitively cut off by the finite thickness). This reversal, however, does not happen in all cases: it does not occur for the Si-teflon slab-slab geometry or for *any* of the slab-sphere geometries (and we do not plot d_c in these reversed cases because the results are very similar to the original results). Evidently, the finite *lateral* size of the spheres has a dramatic qualitative interaction with the material dispersion. In future work, we plan to investigate this interesting interplay between material and geometric dispersion, in addition to the finite-size behavior exhibited by the finite gold [279] and SiO₂ slabs.

Another feature worth noting in Fig. 8-10 also stems from the anomalous response of the Si-SiO₂ slab-slab geometry to changes in the SiO₂ thickness t : the stable $d_c^{(s)}$ and unstable $d_c^{(u)}$ equilibria bifurcate at a radius R_c .

In what follows, we illustrate an interesting corollary of this type of stability: the possibility of obtaining stable noncontact configurations of compact objects at a nonzero sep-

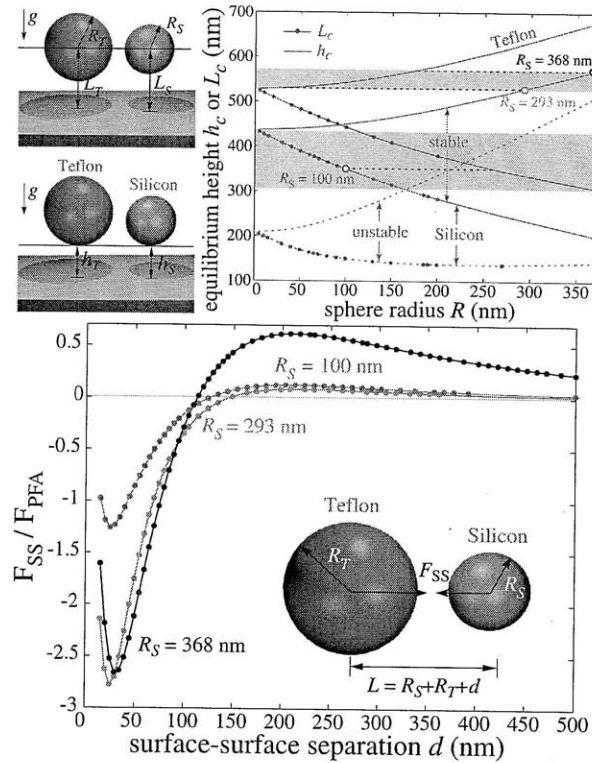


Figure 8-11: (Top:) Plot of the stable equilibrium center–surface (L_c) and surface–surface (h_c) separation between either a teflon (yellow) or silicon (blue) sphere and a semi-infinite gold slab (depicted schematically on the left inset), as a function of sphere radius R . The black lines also show the presence of an unstable equilibrium in the silicon-sphere case. (Bottom:) Plot of the force F_{SS} between two teflon/silicon spheres of radii R_S/R_T , showing the existence of a stable equilibrium.

ation. In particular, we calculate the Casimir force in a nanoparticle-dicluster system consisting of teflon and silicon spheres, of different radii R_T and R_S , respectively, immersed in ethanol. The force F_{SS} between the spheres is plotted in Fig. 8-11(bottom) for different sets of radii $R_S = \{99.69, 293.67, 368.39\}$ nm and $R_T = \{262.09, 32.67, 176.64\}$ nm, respectively. This choice of the radii was motivated by one possible experimental configuration, in which the pair of spheres are also levitated above a planar slab: in this geometry, discussed below, the sphere radii are chosen so that both spheres are suspended at the same height above the slab. With this choice of materials, the spheres are again attractive at large separations and repulsive for small separations, leading to a stable (orientation-independent) surface–surface separation in the 100–150nm range.

To make such a dicluster pair easier to observe in experiments, one could simultaneously suspend them at a known distance above a planar substrate, using the interplay between

the repulsive Casimir force and gravity to create stable levitation. For simplicity, we investigate this possibility within the additive approximation: the slab–sphere and sphere–sphere interactions are considered independently. Because Casimir forces are not additive, the presence of the slab will change the stable separation of the spheres (and vice versa for the stable height of the spheres). However, this approximation forms a useful starting point for the design of such an experiment and should even be accurate in the limit where the sphere diameter is much larger than the stable surface–surface separations (here, most of the diameters are at least twice the stable separations). Figure 8-11(top) shows a plot of the equilibrium surface–surface (h_c) and center-surface ($L_c = h_c + R$) separations of the teflon (yellow) and silicon (blue) spheres suspended above a semi-infinite gold slab, as a function of sphere radius R . As shown by the figure, decreasing R acts to increase L_c and decrease h_c . The decrease in h_c occurs much more rapidly than in Fig. 8-10 due to the fact that, in addition to geometric dispersion, the force of gravity scales as the mass ρV of a sphere (where the density $\rho \approx \{2330, 2200, 789\}$ kg/m³, for {Si, teflon, ethanol}, respectively). The center–center separation $L_c = h_c + R$ increases with increasing R because $\partial h_c / \partial R > -1$. In addition to a stable equilibrium h_c , Au–Si exhibits an unstable equilibrium at smaller d_c due to the transition to an attractive Casimir force for small separations; if the sphere were ever pushed below the unstable equilibrium point, it would continue downward and adhere to the slab. This would be a concern for experiments if fluctuations in the sphere height could push it below the unstable equilibrium, but as seen from Fig. 8-11 the distance between the stable and unstable equilibria is over 200nm for $R < 50$ nm nanoparticles.

The gray areas in Fig. 8-11(top) depict regions in which the L_c or h_c of the two spheres can be made equal by an appropriate choice of radii, as shown schematically in Fig. 8-11(top-left). This determines the stable configuration of the two-sphere dicluster when they are brought together above the surface; the three dashed horizontal lines in Fig. 8-11(top) correspond to the radii used for the force calculation in Fig. 8-11(bottom).

8.5 Strong Temperature Dependence

Our discussion of fluctuation-induced interactions throughout the entire thesis has been focused on the zero-temperature limit of thermodynamic electromagnetic fluctuations (the $T = 0$ limit of the Bose–Einstein distribution of the photon fluctuations) [304]. Here, we

focus in the nonzero temperature T case, where the Casimir force is predicted to change as a consequence of the changing photon thermal distribution (the presence of “real photons”). Unfortunately, this change is typically negligible near room temperature and submicron separations [213, 358] and is only a few percent for ~ 100 K changes in T at $1\text{--}2\ \mu\text{m}$ separations where Casimir forces are barely observable [47, 49, 213, 284, 358]. Therefore, despite theoretical interest in these T effects [32, 358, 555, 573], it has proved difficult for experiments to unambiguously observe T corrections to the Casimir force [47, 49, 284]. Other attempts to measure T Casimir corrections have focused on nonequilibrium situations that differ conceptually from forces due purely to equilibrium fluctuations [388]. A clear experimental verification of a T Casimir correction would be important in order to further validate the foundation of Lifshitz theory for Casimir effects [213, 358]. A related experimental difficulty centers around observing small (few percent at micron separations) corrections to the Casimir force due to different models of material dispersion [61, 358].

In this last section of the thesis, we shift our attention to the $T > 0$ Casimir effect. In particular, we propose a method for obtaining strongly temperature-dependent Casimir effects by exploiting geometries involving fluid-separated dielectric objects (with separations in the hundreds of nanometers). As described in the previous sections, the Casimir force between fluid-separated objects can be repulsive [129, 304], and can even lead to stable suspensions of objects due to force-sign transitions from material dispersion [450, 452] or gravity [344, 450]. We show that, by a proper choice of materials/geometries, this stable separation d can depend dramatically on T ($2\ \text{nm/K}$ is easily obtainable), and there can even be transitions where d jumps discontinuously at some T . Essentially, a stable separation arises from a delicate cancellation of attractive and repulsive contributions to the force from fluctuations at different frequencies, and this cancellation is easily altered or upset by the T corrections. This appears to be the first prediction of a strong T -dependent Casimir phenomenon at submicron separations where Casimir effects are most easily observed. We present the phenomenon in simple parallel-plate geometries, but we believe that the basic idea should extend to many other geometries and materials combinations that have yet to be explored. Finally, we also point out that the same systems that are strongly T -dependent can also be very sensitive to the precise details of the material dispersion at low frequencies, a property that we plan to exploit in the future; for example, for certain choice of materials, the difference between a Drude and plasma model for gold leads to a $15\ \text{nm}$ change in the

stable separation (and larger sensitivity may be possible in other geometries).

From the discussion in chapter 5 and 6, we know that the Casimir force between two bodies is a combination of fluctuations at all frequencies ω , and at $T = 0$ can be expressed as an integral $F(0) = \int_0^\infty f(\xi)d\xi$ over imaginary frequencies $\omega = i\xi$ [304]. The contributions $f(\xi)$ from each imaginary frequency are a complicated function of the geometry and materials, and can be computed by way of stress tensors (see chapter 5 and 6) [129] (valid in fluids for subtle reasons [354, 408]) or via the Casimir energy [354]. At a finite T , this integral is replaced by a sum over “Matsubara frequencies” $2\pi n k T / \hbar = n \xi_T$ for integers n :

$$F(T) = \frac{2\pi k T}{\hbar} \left[\frac{f(0^+)}{2} + \sum_{n=1}^{\infty} f\left(\frac{2\pi k T}{\hbar} n\right) \right]. \quad (8.7)$$

Physically, this arises as a consequence of the $\coth(\hbar\omega/2kT)$ Bose–Einstein distribution of fluctuations at real frequencies—when one performs a contour integration in the upper-half complex- ω plane, the residues of the \coth poles at $\hbar\omega/2kT = ni\pi$ lead to the summation [284]. Mathematically, Eq. (8.7) corresponds exactly (including the $1/2$ factor for the zero-frequency contribution) to a trapezoidal-rule approximation to the $F(0)$ integral, which allows one to use the well-known convergence properties of the trapezoidal rule [50] to understand the magnitude of the T correction. In particular, the difference between the trapezoidal rule and the exact integral scales as $O(T^2)$ for smooth $f(\xi)$ with nonzero derivative $f'(0)$ [358] (typical for Casimir forces between metals [213]). More specifically, $f(\xi)$ is exponentially decaying with a decay-length $2\pi c/a$ for some characteristic lengthscale a (e.g. the separation or size of the participating objects), while the discrete sum of Eq. (8.7) corresponds to a lengthscale given by the Matsubara wavelength $\lambda_T = 2\pi c/\xi_T = \hbar c/kT$, in which case one would expect the T correction to scale as $O(\lambda_T^2/a^2)$. Unfortunately, at $T = 300$ K, $\lambda_T \approx 7.6 \mu\text{m}$, which is why the T corrections are typically so small unless $a > 1 \mu\text{m}$ [358]. We cannot change the smoothness of $f(\omega)$ since it arises from the analyticity of the classical electromagnetic Green’s function in the upper-half complex- ω plane [304], so the only way to obtain a larger T correction is to introduce a longer lengthscale Λ into the problem that dominates over other lengthscales such as the separation a . One way of achieving this is to make the $f(\xi)$ integrand *oscillatory* with an oscillation period $\Delta\xi \sim 2\pi c/\Lambda$ that is much shorter than the decay length $\sim 2\pi c/a$. Intuitively, discretizing an oscillatory integral induces much larger discretization effects than for a non-oscillatory integral, and this intuition




Figure 8-12: Relative permittivity $\varepsilon(i\xi)$ of various materials as a function of imaginary frequency ξ (in units of $c/\mu\text{m}$) or “Matsubara” temperature $T = \hbar\xi/2\pi k_B$. Doped silicon corresponds (bottom to top) to doping density $\rho_d = \{1, 3, 5, 10, 10^2\} \times 10^{16}$, modeled via an empirical Drude model [127], as is gold [34]. Water, polystyrene, ethanol, teflon, and lithium niobate are all modeled via standard Lorentz-oscillator models [325].

can be formalized by a Fourier analysis of the convergence rate of the trapezoidal rule [50]. The question then becomes: how does one obtain an oscillatory Casimir integral?

One way to obtain oscillatory frequency contributions to the Casimir force is to employ a system where there are combinations of attractive and repulsive contributions, as described in Sec. 8.4.2. In particular, it was shown that the sign and magnitude of the total force at any given separation can be changed by a proper choice of material dispersion, leading to the possibility of obtaining Casimir equilibria between objects at multiple separations, explored in the preceding section. [Basically, dispersion can cause the relative ordering of the dielectric permittivities of the objects to change, leading to a change in the sign of $f(\xi)$.] In this section, for the purpose of achieving a strong T -dependence at short (submicron) separations, we search for materials or geometries with dielectric crossings occurring at sufficiently small $\xi = 2\pi c/\Lambda \sim \xi_T$, close to the room-temperature Matsubara-frequency scale ξ_T .

To begin with, we compute the Casimir force between semi-infinite slabs, computed via a generalization of the Lifshitz formula [529] that can handle multi-layer dielectric objects, with relative permittivities ε plotted in Fig. 8-12 as a function of imaginary frequency $i\xi$ (bottom axis) or “Matsubara temperature” $T = \hbar\xi/2\pi k$ (top axis). Figure 8-13 shows the




Figure 8-13: Equilibrium separation d_c (in units of μm) as a function of temperature T (in Kelvin), for a geometry consisting of fluid-separated semi-infinite slabs (no gravity). The various curves correspond to d_c for various material combinations. Solid/dashed lines correspond to stable/unstable equilibria, and shaded regions are T where ethanol is non-liquid at 1 atm [161]. Doped-silicon is plotted for various doping densities $\rho_d = \{1, 10, 100, 500, 10^3\} \times 10^{17}$.

equilibrium separation d_c (in units of μm) as a function of temperature $T \in (0, 400)$ K (in Kelvin) for some of the material combinations (solid/dashed lines correspond to stable/unstable equilibria), and demonstrates various degrees of T sensitivity. The previously studied [450] material combination of teflon/ethanol/silicon (data not shown) shows very little T -dependence: d_c varies $< 1\%$ over 400 K. More dramatic behavior is obtained for lithium niobate (LiNbO3) or doped silicon (doping density $\rho_d = \{1, 10, 100, 500, 10^3\} \times 10^{17}$), whose low- ξ ϵ crossings with ethanol lead to the desired oscillatory $f(\xi)$ in Eq. (8.7): the stable-equilibrium separation $d_c^{(s)}$ for both cases decreases by $> 2 \mu\text{m}$ over 400 K, and $\frac{d}{dT}d_c^{(s)} \approx -2 \text{ nm/K}$ near $T = 300 \text{ K}$ (where $d_c^{(s)} \sim 300\text{--}700 \text{ nm}$). The sign of $\frac{d}{dT}d_c^{(s)}$ comes from the increasing domination of the repulsive small- ξ (large-separation) contributions to $f(\xi)$ as T increases. Varying the silicon doping density dramatically changes the T dependence because it tunes the low- ξ silicon/ethanol ϵ crossing in Fig. 8-12. Doped-silicon exhibits another interesting behavior: d_c disappears at a critical temperature T_c (determined by ρ_d) due to a bifurcation between the stable/unstable equilibria. (T_c is determined not only by ρ_d but also by the LiNbO3 layer-thickness, here $\sim 50 \text{ nm}$, in the absence of which $T_c \approx 150 \text{ K}$.) Experimentally, such a bifurcation yields an *irreversible* transition from




Figure 8-14: Equilibrium position d_c (in units of μm) of a semi-infinite polystyrene (PS) slab immersed in ethanol (shaded $T = \text{non-liquid}$) and suspended against gravity by a repulsive Casimir force exerted by a doped-silicon (Si) slab. The solid/dashed lines correspond to stable/unstable d_c , and each color represents a different value of PS slab-thickness h (in units of μm). The inset shows the magnitude of the total energy $U_T(d)$ (in units of $k_B T$) as a function of d for $h = 150 \text{ nm}$, at various T .

suspension ($T < T_c$) to stiction ($T > T_c$). Figure 8-13 also shows a small sample of the many other material possibilities. The shaded regions in Fig. 8-13 correspond to T above the boiling point (320 K) or below the freezing point (159 K) of ethanol at 1 atm [161].

The inclusion of gravity/buoyancy introduces another force into the system and leads to the possibility of additional phenomena, such as additional stable equilibria due to gravity/Casimir competition [450]. For example, Fig. 8-14 shows the equilibrium separations d_c of a polystyrene (PS) slab of thickness h in ethanol above a semi-infinite doped-silicon slab ($\rho_d = 1.1 \times 10^{15}$), including gravity (mass density $\rho_{\text{PS}} - \rho_{\text{ethanol}} = 0.264 \text{ g/cm}^3$ [161]). As in Fig. 8-13, d_c varies dramatically with T : $\frac{d}{dT}d_c \approx 1.2 \text{ nm/K}$ near $T = 300 \text{ K}$. Gravity becomes increasingly important as h grows: compared to $h = 0$ (leftmost line), it creates an additional stable equilibrium (solid lines) at large d_c (hundreds of nm) where the downward gravity dominates. With gravity, there are three stable/unstable bifurcations instead of two, leading to three critical temperatures where qualitative transitions occur: T_g refers to the temperature of the topmost bifurcation, created by gravity, and the other two temperatures are labeled T_1 ($\approx 100 \text{ K}$) and T_2 ($\approx 180 \text{ K}$). If $T_g < T_1$ ($h < 40 \text{ nm}$), there exists an irreversible transition from suspension to stiction as T is decreased below

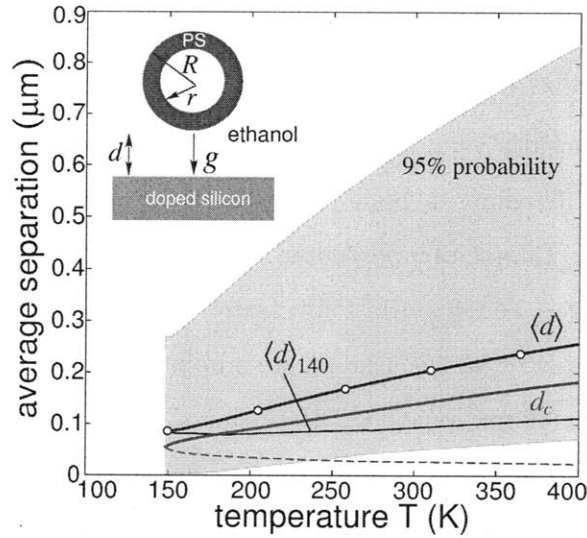


Figure 8-15: Average separation $\langle d \rangle$ (circles) and equilibrium separation d_c (red line), in units of μm , vs. temperature T (in Kelvin), for a geometry consisting of a fluid-separated hollow PS sphere of inner/outer radius $r/R = 3.2/5 \mu\text{m}$ suspended in ethanol against gravity by a doped-silicon slab and subject to Brownian motion. Shaded region indicates where sphere is found with 95% probability. The thin black line is the average $\langle d \rangle_{140}$ if the Casimir energy at 140 K is used instead of the true temperature-dependent energy landscape.

T_g . If $T_1 < T_g < T_2$, there are two irreversible transitions from suspension to suspension (smaller d_c) to stiction as T is lowered from $T > T_g$ to $T < T_1$ starting in the large- d_c equilibrium. Finally, when $T_g \rightarrow T_2$ ($h \approx 300 \text{ nm}$) the two stable equilibria merge and only the T_1 bifurcation remains. Perhaps most interestingly, when this merge occurs the slope $\frac{d}{dT}d_c$ can be made arbitrarily large but finite, corresponding to an arbitrarily large (but reversible) temperature dependence. For example, $\Delta d_c \approx 130 \text{ nm}$ for a small change $\Delta T \approx 5 \text{ K}$ around T_2 , for $h = 300 \text{ nm}$.

In a real experiment, the situation is further complicated by Brownian motion, which will cause the separation to fluctuate around stable equilibria and will also lead to random transitions between equilibria [439]. In the example of Fig. 8-14, the attractive interaction at small separations means that there is a nonzero probability that the slabs will fluctuate past the unstable-equilibrium energy barrier ΔU_T into stiction, but the rate of such a transition decreases proportional to $\exp(-\Delta U_T/k_B T)$ [439]—here, assuming a $50 \times 50 \mu\text{m}^2$ PS slab, $\Delta U_T/k_B T \approx 10^4$, so the stiction rate is negligible. The energy landscape $U_T(d)/k_B T$ is plotted for several cases in the inset to Fig. 8-14: the general prediction of experimental observations involves a viscosity-damped Langevin process [439] that is beyond the scope of this paper to model, but by choosing T one can make the potential barrier between the

two stable equilibria arbitrarily small and therefore should be able to reach an experimental regime in which “hopping” is observable.

Alternatively, we consider a simpler example system with only a single stable equilibrium and a single degree of freedom: a hollow PS sphere (experimentally available at similar scales [559]), filled with ethanol, of inner/outer radius $r/R = 3.2/5 \mu\text{m}$ suspended in ethanol above a doped-silicon ($\rho_d = 1.1 \times 10^{15}$) substrate, shown on the inset of Fig. 8-15. (To compute the Casimir energy in this system, we employ a simple PFA approximation that is sufficiently accurate for our purpose. Here, for $d \approx 500 \text{ nm}$, the exact energy is $\approx 85\%$ of the PFA energy.) For this example, in Fig. 8-15 we plot the mean surface–surface separation $\langle d \rangle \sim \int dz z \exp[U_T(z)/k_B T]$ (determined only by the energy landscape and the Boltzmann distribution [439]), corresponding to an experiment averaging d over a long time, along with a confidence interval (shaded region) indicating the range of d where the particle is found with 95% probability. The sphere experiences an attractive interaction at small separations, but again we find that the unstable-equilibrium energy barrier is sufficiently large ($\Delta U/k_B T \approx 50$) to prevent stiction for T near 300 K. As T varies, two factors affect $\langle d \rangle$: the T -dependence of the Casimir energy $U_T(z)$, and the explicit $k_B T$ in the Boltzmann factor. To distinguish these two effects, we also plot (thin black line) $\langle d \rangle_{140} \sim \int dz z \exp[U_{140}(z)/k_B T]$ where the $T = 140 \text{ K}$ (bifurcation point) Casimir energy is used at all temperatures. Comparing $\langle d \rangle$ with $\langle d \rangle_{140}$, it is evident that most of the positive-slope T dependence of $\langle d \rangle$ ($\approx 0.8 \text{ nm/K}$ around 300 K) is due to U_T , and therefore $\langle d \rangle$ offers a direct measure of the Casimir-energy T dependence.

Experimentally, measuring hundreds of nm changes in separation over tens or hundreds of Kelvins appears very feasible, perhaps even easier than traditional measurements of Casimir forces. (In a fluid, static-charge effects can be neutralized by dissolving electrolytes in the fluid [372], which also have the added benefit of significantly reducing/increasing the freezing/boiling point of the fluid [161].) Such temperature-dependent suspensions may even have practical applications in microfluidics. We believe that the examples shown in this letter only scratch the surface of the possible temperature/dispersion effects that can be obtained in Casimir-suspension systems. Not only are there many other possible materials and geometries to explore in the fluid context (along with more detailed calculation of the Brownian dynamics), and by no means are the effects shown here the maximum possible, but similar principles should apply in other systems where competing attractive/repulsive

contributions to Casimir forces can be found, such as the vacuum-separated “zipper” structure [446] or a recent repulsive metallic geometry [298].

Appendix A

Disorder-Immune Confinement of Light in Photonic-Crystal Cavities

A.1 Influence of Disorder on Complete-Gap Modes

In this appendix, we show that a photonic-crystal cavity can confine light with a lifetime that is *not* limited by fabrication disorder and is independent of modal volume. A photonic crystal is a periodic dielectric structure with a photonic band gap, a range of wavelengths in which light cannot propagate in the crystal; a “point defect” introduced in such a crystal traps a localized mode. [239] Dielectric cavities can also be made by ring resonators [469] as well as by point defects in photonic-crystal “slabs” that combine band gaps in one or two directions with index-guiding; [250] all such cavities, however, suffer from intrinsic radiation losses whose reduction ultimately requires a tradeoff in modal volume [250, 316, 464, 508] (except for Ref. [554]), and any disorder should also scatter light and limit the lifetime. With a complete photonic band gap (*i.e.*, for all directions and, in 3d, all polarizations), the dimensionless lifetime “ Q ” of the cavity increases exponentially with the number of crystal periods surrounding the cavity, independent of its modal volume (which can approach a half-wavelength diameter). [239] Moreover, it has been shown that the photonic bandgap is preserved even when a small (bounded) amount of disorder is introduced into the crystal, [147, 160, 256, 300] leading to the hope that optical confinement is maintained in a disordered cavity (with some evidence in 2d provided by Ref. [572]). Here, we show that this hope is justified: by finite-difference time-domain (FDTD) [523] simulations for crystals in both




Figure A-1: Cavity mode E_z (blue/white/red = positive/zero/negative) for a 2d structure with (a) no disorder, (b) surface disorder, and (c) surface + position disorder, and for a (horizontal) cross-section of a 3d structure with (A) no disorder and (B) sidewall disorder.

2d and 3d, we demonstrate that the only effect of the disorder is a slight change in the exponential *rate* at which Q increases with the number of crystal periods, and also a small shift in the cavity frequency ω that could be compensated by introducing a single external tuning parameter into the cavity [397, 417, 562] We obtain a Q of 10^8 optical periods and a half-wavelength single-mode diameter in a small (finite) 3d crystal with unphysically large fabrication disorder.

Thus, with a complete-gap crystal, the only factors limiting the lifetime of a cavity are the size of the crystal that can be fabricated and also the intrinsic absorption losses of the materials. For example, Si at $\lambda = 1.55\mu\text{m}$ has a bulk extinction coefficient k so small it is difficult to measure, but an extrapolation of the band-edge k tail [528] is $\sim 10^{-12}$, giving a maximum $Q > 10^{12}$; before this point, however, electronic surface states and similar effects determined by fabrication details may dominate absorption. The resulting high Q and small modal volume are important especially for nonlinear devices to minimize the power requirements. [501] Note that, although $1/Q$ is a fractional bandwidth, one actually wants Q much larger than a design bandwidth might suggest, since the Q here is only that of radiative loss: the ratio of twice the total Q ($\sim \omega/\text{bandwidth}$) to the loss Q is the

fractional loss. [202] For comparison, we point out that past dielectric cavities with very high Q values required a substantial sacrifice in modal volume compared to the theoretical optimum $\approx (\lambda/2n)^3$, such as microspheres ($Q \approx 10^{10}$, modal volume $> 10^4 \times$ optimum) [542] or photonic-crystal heterostructures ($Q \approx 10^6$, modal volume $\approx 10 \times$ optimum). [508]

We focus on two structures: in 2d, a square lattice of dielectric rods in air (period a , radius $0.2a$, $\varepsilon = 12 \sim \text{Si}$), shown in the Fig. A-2 inset, which has a gap (frequencies $0.286\text{--}0.421 \ 2\pi c/a$) for the TM polarization (\mathbf{E} out of the plane); [239] in 3d, an fcc lattice (lattice constant a) of air cylinders in dielectric ($\varepsilon = 12$) oriented along the 111 direction [248] (radius $0.293a$, height $0.93a$), forming an alternating stack of rods-in-air and holes-in-dielectric layers with a complete gap (frequencies $0.507\text{--}0.628 \ 2\pi c/a$), shown in the Fig. A-4 inset. The 3d structure has an $\varepsilon = 12$ substrate below and to the sides, and air above with a hole layer topmost, similar to experiment. [419]

In 2d, a cavity is formed by removing a single rod (as in the Fig. A-2 inset), giving a single confined mode [239] with $\omega = 0.3779 \ 2\pi c/a$ shown in Fig. A-1(a). In 3d, we also remove a rod from one of the “rods-in-air” layers, giving a mode with $\omega = 0.5689 \ 2\pi c/a$ that resembles the TM mode of the analogous 2d cavity; [412] a horizontal cross-section of this mode is shown in Fig. A-1(A) and a vertical cross section is given in the inset of Fig. A-4. Using FDTD with PML absorbing boundaries, [523] we excite the cavity modes by a Gaussian pulse and extract ω and Q from the response by a filter diagonalization method. [250,331,332] Q is defined [202,469] as the number of optical periods for the energy to decay by $\exp(-2\pi)$, and modal volume V is computed as $\int \varepsilon |\mathbf{E}|^2 / \max \varepsilon |\mathbf{E}|^2$. [415] For disorder, we compute the mean and standard deviation (std. dev.) for each result over 15 random configurations (except where otherwise specified).

In 2d, with a resolution of 16 pixels/ a , we examine four types of random disorder, all with unrealistically large magnitudes: surface (Fig. A-1b, described below), position (random rod-center offsets in $[-0.15a, 0.15a]$), surface+position (Fig. A-1c), and rod ε variations within $[-15\%, 15\%]$. For surface disorder, 20 cylinders of random radius in $[0, 0.0625a]$ are added/removed from the surface of each rod. Fig. A-2 shows the resulting Q vs. number n of layers around the cavity (e.g. $n = 2$ in the inset). Neither the exponential increase of Q (Fig. A-2) nor the localized field pattern (Fig. A-1) is substantially degraded by even such large disorder. (Indeed, sometimes the disorder happens to improve Q , since the perfect structure is not optimized for Q .) ω was $0.3779 \ 2\pi c/a$ for the perfect cavity and had std.




Figure A-2: Q vs. # layers for a 2d cavity (inset shows 2 layers), for disorder types (descending by Q): position (red), no disorder (black), ϵ (cyan), surface+position (blue), and surface (magenta). Error bars show std. dev. over 15 runs.

dev. 1%/2%/3%/1% for surface/position/surf.+pos./ ϵ disorder, respectively. The modal volume (area) in all cases was roughly $(0.64\lambda/2)^2$ (std.dev. < 6%).

This insensitivity of confinement to disorder derives from the fact that the bandgap is preserved (only shrunk slightly): disorder introduces states only near the gap edges. We illustrate this by computing the 2d gap as a function of ϵ disorder strength, plotted in Fig. A-3 (using a planewave expansion [249] in a 9×9 supercell and extracting the largest gap, averaged over 10 runs). Superimposed is the mean FDTD mode ω and its std. dev. (over 100 runs), corresponding to a well-defined localized mode until the gap shrinks substantially (requiring over 50% ϵ disorder). Once the bandgap closes entirely, which is not shown, any observed localized modes have no relation to the original cavity mode, and may correspond instead to Anderson localization.

In 3d, with a resolution of 25 pixels/ a (~ 17.7 pixels per in-plane lattice constant of the 2d cross sections), we examine two types of random disorder: height disorder via variations in $[-5\%, 5\%]$ in the height of the air cylinders (or $\pm 26\%/42\%$ variation in the thickness of a rod/hole layer, respectively), and sidewall disorder (adding/removing cylinders as in 2d) shown in Fig. A-1(B). The resulting Q vs. number n of surrounding layers is shown in Fig. A-4, and the field pattern for sidewall disorder is in Fig. A-1(B). (n is defined so that the number of nearest-neighbor rods *surrounding* the cavity is $n - 1$ in-plane and n vertically; the field plots show $n = 3$. The *total* number of vertical periods both above and below is $\sim 2n/3$, where a vertical period comprises three layers of holes.) Again, Q is only




Figure A-3: 2d band gap and cavity mode frequencies vs. strength of ϵ disorder (maximum % variation of rod ϵ). Error bars show std. dev. for 10/100 runs for gap/cavity frequencies. Inset shows cavity structure (4 layers).




Figure A-4: Q vs. # layers for cavity in the 3d structure (left inset), for disorder types (descending by Q): no disorder (black), height (blue), sidewall roughness (red). Right inset: mode $\epsilon|\mathbf{E}|^2$ in vertical cross-section (3 layers). Error bars show std. dev. over 15 runs.

slightly affected by disorder, modal volume is roughly fixed at $(0.68\lambda/2)^3$, and frequency $\omega = 0.5689 \, 2\pi c/a$ has std. dev. 0.3%/2% for height/sidewall disorder. (The sidewall disorder, here, also causes a 3.6% systematic ω shift, because adding/removing cylinders does not have a symmetric first-order effect in 3d [253])

We believe that this robust confinement is a general feature of complete-gap systems for any sufficiently bounded disorder, not limited to the examples here.

Appendix B

Computing Photonic-Quasicrystal Spectra via Bloch's Theorem

In this appendix, we describe a computational method to solve for the spectra and eigenstates of quasicrystalline electromagnetic structures by directly solving a periodic eigenproblem in a higher-dimensional lattice. Such photonic quasicrystals (PQCs) have a number of unique properties compared to ordinary periodic structures [1, 85, 89, 104, 112, 154, 159, 165, 167, 262, 274, 306, 363, 385, 399, 455, 553, 558, 564, 591, 593, 600, 601], especially in two or three dimensions where they can have greater rotational symmetry and therefore offer some hope of complete photonic band gaps with lower index contrast [197, 255, 592, 601] than the roughly 2:1 contrast currently required for periodic structures [328]. However, the study of two- and three-dimensional photonic quasicrystals has been hampered by the computational difficulty of modeling aperiodic structures, which has previously required large “supercell” calculations that capture only a portion of the infinite aperiodic lattice. Our method, in contrast, captures the entire infinite aperiodic structure in a single higher-dimensional unit cell, and we believe that this approach will ultimately be much more computationally tractable for two- and three-dimensional quasicrystals. The idea that many quasicrystals can be constructed by an irrational slice of a higher-dimensional lattice is well known [235, 513, 517], and in fact is the most common formulation of quasicrystals in two and three dimensions [291, 330, 552], but the possibility of direct numerical calculations within the higher-dimensional space seems to have been little explored outside of some tight-binding calculations in quantum systems [108, 319]. As a proof of concept, we

demonstrate a first implementation of the technique applied to one-dimensional quasicrystals, such as the well known Fibonacci structure. Not only can we reproduce the spectrum from transfer-matrix calculations, but we also show that the higher-dimensional picture provides an interesting way to visualize the eigenmodes and compute defect states in the infinite aperiodic structure.

There have been several previous numerical approaches to simulating quasicrystal structures in electromagnetism and quantum mechanics. In one dimension, a typical quasicrystal is an aperiodic sequence of two or more materials, determined either by a slice of a higher-dimensional lattice [513] or by some “string concatenation” rule [235]. In either case, efficient 2×2 transfer-matrix methods are available that allow one to quickly compute the transmission spectra and density of states for supercells consisting of many thousands of layers [181, 215]. Two- and three-dimensional quasicrystals are almost always defined as an irrational slice (i.e., incommensurate Miller indices) of a higher-dimensional lattice; for example, the famous Penrose tiling can be viewed as a two-dimensional slice of a five-dimensional cubic lattice or of a four-dimensional root lattice A_4 [517]. In such cases, supercell computations of a finite portion of the infinite aperiodic structure (or a rational approximant thereof [181, 513]) require slower numerical methods, most commonly finite-difference time-domain (FDTD) simulations [165, 166, 262] or planewave expansions [113, 254]. Unfortunately, these methods become very expensive for large supercells, nearly prohibitively so for three-dimensional quasicrystals—there have been experiments for 3D PQC [291, 330], but as yet few theoretical predictions [514, 598]. With FDTD methods, for example, the PQC local density of states is typically integrated in Monte-Carlo fashion via random sources or initial conditions [1, 363, 553], but many simulations are required to sample all possible modes in a large supercell. Also, the finite domain of a supercell becomes even more significant in higher dimensions where a tractable supercell is necessarily smaller, as there can be localized states [112, 165, 262, 363] whose presence is dependent on the particular region of the PQC considered. Our method of computing the spectrum directly in the higher-dimensional unit cell, on the other hand, requires no supercell to capture the infinite aperiodic structure—it uniformly samples (up to a finite resolution) every possible supercell of the infinite quasicrystal, rather than any particular subsection. The influence of finite-resolution on the convergence of the spectrum can be systematically understood: one is not “missing” any part of the quasicrystal, so much as resolving the entire quasicrystal with

lower resolution.

The structure of this paper is as follows: in Sec. B.1 we review the “cut-and-project” method for defining a PQC as a slice of a higher-dimensional lattice, followed in Sec. B.2 by a description of our computational method in the higher-dimensional lattice. There, we describe the extension of Maxwell’s equations to higher dimensions and also describe its solution in terms of a higher-dimensional Bloch plane-wave expansion. As a proof of concept, we present a sequence of one-dimensional examples in Sec. B.3. First, we compare results for a one-dimensional “Fibonacci sequence” with standard one-dimensional transfer-matrix techniques. Second, as mentioned above, we demonstrate how one can use the same technique to study defects in the quasicrystal, as demonstrated in the one-dimensional “Fibonacci” example. Finally, we demonstrate the ease with which one can construct and explore different quasicrystals by continuously varying the cut angle.

B.1 Quasicrystals via cut-and-project

Given a periodic lattice, any lower dimensional cross-section of that lattice may be either periodic or quasi-periodic, depending upon the angle of the cross-section. For example, the periodic 2D cross-sections of a 3D crystal are the lattice planes, defined in crystallography by integer Miller indices. If the Miller indices have irrational ratios, on the other hand, the cross-section is aperiodic but still has long-range order because of the underlying higher-dimensional periodicity. This is what is known as a “cut-and-project” method of defining a quasicrystalline structure: as a slice of a periodic structure in a higher-dimensional “superspace” [235, 513]. (For a thorough discussion of quasicrystals via cut-and-project, see Ref. [235].) Cut-and-project defines a specific class of quasicrystals; equivalently, and more abstractly, cut-and-project corresponds to structures whose Fourier transform has support spanned by a finite number of reciprocal basis vectors (the projection of the reciprocal lattice vectors from higher dimensions) [235, 552]. This class includes most commonly considered quasicrystals in two or three dimensions, including the Penrose tiling [517] and the 2D Fibonacci quasicrystal [305], as well as many one-dimensional quasicrystals including a 1D Fibonacci structure.

For example, consider the Fibonacci PQC in one dimension formed from two materials $\epsilon_A = 4.84$ and $\epsilon_B = 2.56$ in layers of thickness A and B , respectively, similar to a recent




Figure B-1: Unit cell of the Fibonacci superspace dielectric. The physical dielectric is obtained by taking a slice at an angle $\tan \phi = \tau$. Black/white are the dielectric constants of the structure factor material and air, chosen to be $\epsilon = 4.84$ and $\epsilon = 2.56$, respectively.

experimental structure [104]. The Fibonacci structure S is then defined by the limit $n \rightarrow \infty$ of the string-concatenation rule $S_n = S_{n-2}S_{n-1}$ with starting strings $S_0 = B$ and $S_1 = A$ [104], generating a sequence $BABAABABAABA \dots$. In the case where B/A is the golden ratio $\tau = (1 + \sqrt{5})/2$, exactly the same structure can be generated by a slice of a two-dimensional lattice as depicted in Fig. B-1 [235]. The slice is at an angle ϕ with an irrational slope $\tan \phi = 1/\tau$, and the unit cell of the 2D lattice is an $A \times A$ square at an angle ϕ in a square lattice with period $(A + B) \sin \phi = a$. Because the slope is irrational, the offset/intercept of the slice is unimportant: any slice at an angle ϕ intercepts the unit cell at infinitely many points, filling it densely.

For thickness ratios $B/A \neq \tau$, the Fibonacci structure cannot be constructed by cut-and-project, and in general string-concatenation rules can produce a different range of structures (such as the Thue-Morse PQC [103]) than cut-and-project. This is partly a question of definition—some authors reserve the term “quasicrystal” for cut-and-project structures [517]. In any case, cut-and-project includes a wide variety of aperiodic structures, including most of the structures that have been proposed in two or three dimensions (where they can be designed to have n -fold rotational symmetry for any n), and are the class of quasicrystals that we consider in this paper.

In general, let $d \leq 3$ be the number of physical dimensions of a quasicrystal structure

generated by a d -dimensional “slice” of an n -dimensional periodic structure ($n > d$). Denote this slice by X (the physical space) with coordinates $\mathbf{x} \in \mathbb{R}^d$, and denote the remaining $n - d$ coordinates by $\mathbf{y} \in \mathbb{R}^{n-d}$ in the “unphysical” space Y (so that the total n -dimensional superspace is $Z = X \oplus Y$). The primitive lattice vectors $\mathbf{R}_i \in Z$ define the orientation of the lattice with respect to the slice (rather than vice versa), with corresponding primitive reciprocal vectors \mathbf{G}_i defined by the usual $\mathbf{R}_i \cdot \mathbf{G}_j = 2\pi\delta_{ij}$ [235]. (The concept of an “irrational slice” is commonly used in the quasicrystal literature. However, a general definition of what is meant by an “irrational slice” seems difficult to find, and less evident in dimensions $d > 2$. A more precise definition of “irrational slice” in general dimensions and a proof that it is dense in the unit cell is given in the last section.)

The physical dielectric function $\varepsilon(\mathbf{x})$ is then constructed by starting with a periodic dielectric function $\varepsilon(\mathbf{x}, \mathbf{y})$ in the superspace and evaluating it at a fixed \mathbf{y} (forming the slice). Because an irrational slice is dense in the unit cell of the superspace [235], it doesn’t matter what value of \mathbf{y} one chooses, as discussed below. In principle, one could define the unit cell of ε in the superspace to be any arbitrary n -dimensional function, but in practice it is common to “decorate” the higher-dimension unit cell with extrusions of familiar d -dimensional objects [235, 517]. More precisely, “cut-and-project” commonly refers to constructions where a set of lattice points within a finite window of the cut plane are projected onto the cut plane, and this is equivalent to a simple cut where objects at the lattice points are extruded in the y direction by the window width [235]. In particular, the extrusion window is commonly an inverted projection (shadow) of the unit cell onto the y directions [235], although this is not the case for the Fibonacci construction of Fig. B-1.

(Note that the higher-dimensional lattice need not be hypercubic. For example, the Penrose tiling can be expressed as a two-dimensional slice of either a five-dimensional hypercubic lattice or of a non-orthogonal four-dimensional root lattice A_4 [517]. For computational purposes, the lower the dimensionality the better.)

B.2 Computations in Higher Dimensions

Although the cut-and-project technique is a standard way to *define* the quasicrystal structure, previous computational studies of photonic quasicrystals then proceeded to simulate the resulting structure only in the projected (d -dimensional) physical space. Instead, it is

possible to extend Maxwell’s equations into the periodic n -dimensional superspace, where Bloch’s theorem applies to simplify the computation. By looking at only the unit cell in n dimensions one can capture the infinite d -dimensional quasicrystal. Our development of this technique was inspired by earlier research on analogous electronic quasicrystals that applied a tight-binding method in two dimensions to compute the spectrum of a one-dimensional electronic quasicrystal [108,319].

Let us start with Maxwell’s equations in the physical space X for the quasicrystal $\varepsilon(\mathbf{x}, \mathbf{y})$ at some fixed \mathbf{y} (that is, \mathbf{y} is viewed as a parameter, not a coordinate). Maxwell’s equations can be written as an eigenproblem for the harmonic modes $\mathbf{H}(\mathbf{x}, \mathbf{y})e^{-i\omega t}$ [239], where again \mathbf{y} appears as a parameter.

$$\nabla_{\mathbf{x}} \times \frac{1}{\varepsilon(\mathbf{x}, \mathbf{y})} \nabla_{\mathbf{x}} \times \mathbf{H} = (\omega/c)^2 \mathbf{H}, \quad (\text{B.1})$$

where $\nabla_{\mathbf{x}} \times$ denotes the curl with respect to the \mathbf{x} coordinates. Assuming that the structure is quasicrystalline, i.e. that X is an irrational slice of the periodic superspace Z , then ω should not depend upon \mathbf{y} [108]. The reason is that \mathbf{y} only determines the offset of the “initial” slice of the unit cell (for $\mathbf{x} = 0$), but as we reviewed above the slice (considered in all copies of the unit cell) fills the unit cell densely. Therefore, any change of \mathbf{y} can be undone, to arbitrary accuracy, merely by offsetting \mathbf{x} to a different copy of the unit cell. An offset of \mathbf{x} doesn’t change the eigenvalues ω , although of course it offsets the eigenfunctions \mathbf{H} .

The fact that ω is independent of \mathbf{y} allows us to re-interpret Eq. (B.1), without actually changing anything: we can think of \mathbf{y} as a coordinate rather than a parameter, and the operator on the left-hand side as an operator in d -dimensional space. Note that \mathbf{H} is still a three-component vector field, and $\nabla_{\mathbf{x}} \times$ is still the ordinary curl operator along the \mathbf{x} directions, so this is not so much a higher-dimensional version of Maxwell’s equations as an extension of the unmodified ordinary Maxwell’s equations into a higher-dimensional parameter space. The \mathbf{y} coordinate appears in the operator only through ε . Because ω is independent of \mathbf{y} , i.e. it is just a number rather than a function of the coordinates, the equation Eq. (B.1) in higher dimensions is still an eigenproblem, and its spectrum of eigenvalues ω is the same as the spectrum of the d -dimensional quasicrystal, since the equations are identical. The physical solution is obtained by evaluating these higher-dimensional so-

lutions at a fixed \mathbf{y} , say $\mathbf{y} = 0$ (where a different \mathbf{y} merely corresponds to an offset in \mathbf{x} as described above).

For a real, positive ε , both the physical operator and the extended operator in Eq. (B.1) are Hermitian and positive semi-definite, leading to many important properties such as real frequencies ω [239].

B.2.1 Bloch’s Theorem and Numerics for Quasicrystals

Because the superspace eigenproblem is periodic, Bloch’s theorem applies: the eigenfunctions $\mathbf{H}(\mathbf{x}, \mathbf{y})$ can be written in the Bloch form $\mathbf{h}(\mathbf{z})e^{i\mathbf{k}\cdot\mathbf{z}}$, where \mathbf{h} is a *periodic* function defined by its values in the unit cell, and \mathbf{k} is the n -dimensional Bloch wavevector [239].

Here, \mathbf{k} determines the phase relationship between \mathbf{H} in different unit cells of the superspace, but it does not have a simple interpretation once the solution is projected into physical space. The reason is that \mathbf{h} , viewed as a function of \mathbf{x} , is again only quasiperiodic: translation in \mathbf{x} “wraps” the slice into a different portion of the unit cell, so both \mathbf{h} and $e^{i\mathbf{k}\cdot\mathbf{z}}$ change simultaneously and the latter phase cannot be easily distinguished. This prevents one from defining a useful phase or group velocity of the PQC modes.

The key point is that Bloch’s theorem reduces the eigenproblem to a finite domain (the n -dimensional unit cell), rather than the infinite domain required to describe the quasicrystal solutions in physical space. This means that standard numerical methods to find the eigenvalues of differential operators are immediately applicable. For example, since the solution \mathbf{h} is periodic, one can apply a planewave expansion method [249] for \mathbf{h} :

$$\mathbf{h}(\mathbf{z}) = \sum_{\mathbf{G}} \tilde{\mathbf{h}}_{\mathbf{G}} e^{i\mathbf{G}\cdot\mathbf{z}}, \quad (\text{B.2})$$

where the summation is over all n -dimensional reciprocal lattice vectors \mathbf{G} . Because the curl operations only refer to the \mathbf{x} coordinates, $\nabla_{\mathbf{x}} \times \mathbf{h}$ is replaced by a summation over $\mathbf{g}_{\mathbf{x}} \times \tilde{\mathbf{h}}_{\mathbf{G}}$, where $\mathbf{g}_{\mathbf{x}}$ denotes \mathbf{G} projected into X . The resulting eigenproblem for the Fourier coefficients $\tilde{\mathbf{h}}$ (once they are truncated to some wavevector cutoff) can be computed either by direct dense-matrix methods [185] or, more efficiently, by iterative methods exploiting fast Fourier transforms [249]. In the present paper, we do the former, which is easy to implement as a proof of concept, but for higher-dimensional computations an iterative method will become necessary.

We should also remind the reader that there is a constraint $\nabla_{\mathbf{x}} \cdot \mathbf{H} = 0$ on the eigenfunctions, in order to exclude unphysical solutions with static magnetic charges. In a planewave method, this leads to a trivial constraint $(\mathbf{k}_{\mathbf{x}} + \mathbf{g}_{\mathbf{x}}) \cdot \tilde{\mathbf{h}} = 0$, again with \mathbf{k} and \mathbf{G} projected into X .

B.2.2 The Spectrum of the Quasicrystal

With a familiar eigenproblem arising from Bloch's theorem, such as that of a periodic physical structure, the eigenvalues form a band structure: discrete bands $\omega_n(\mathbf{k})$ that are continuous functions of \mathbf{k} , with a finite number of bands in any given frequency range [23]. For a finite-resolution calculation, one obtains a finite number of these bands ω_n with some accuracy that increases with resolution, but even at low resolutions the basic structure of the low-frequency bands is readily apparent. The eigenvalues of the higher-dimensional quasicrystal operator of Eq. (B.1), on the other hand, are quite different.

The underlying mathematical reason for the discrete band structure of a physical periodic structure is that the Bloch eigenoperator for a periodic physical lattice, $(\nabla + i\mathbf{k}) \times \frac{1}{\varepsilon}(\nabla + i\mathbf{k}) \times$, is the inverse of a compact integral operator corresponding to the Green's function, and hence the spectral theorem applies [182]. Among other things, this implies that the eigenvalues at any given \mathbf{k} for a finite unit cell form a discrete increasing sequence, with a finite number of eigenvalues below any finite ω . The same nice property does not hold for the operator extended to n dimensions, because along the \mathbf{y} directions we have no derivatives, only a variation of the scalar function ε . Intuitively, this means that the fields can oscillate very fast along the \mathbf{y} directions without necessarily increasing ω , allowing one to have infinitely many eigenfunctions in a finite bandwidth. More mathematically, an identity operator is not compact and does not satisfy the spectral theorem [182], and since the operator of Eq. (B.1) is locally the identity along the \mathbf{y} directions the same conclusion applies. This means that, when the \mathbf{y} direction is included as a coordinate, it is possible to get an infinite number of bands in a finite bandwidth at a fixed \mathbf{k} .

In fact, as we shall see below, this is precisely what happens, and moreover it is what *must* happen in order to reproduce the well-known properties of quasicrystal spectra. It has been shown that quasicrystal spectra can exhibit a fractal structure [235], with infinitely many gaps (of decreasing size) in a finite bandwidth, and such a structure could not arise from an ordinary band diagram with a finite number of bands in a given bandwidth. Of




Figure B-2: (Left:) Frequency spectrum ω of the Fibonacci quasicrystal vs. “wave-vector” k_x . The blue lines indicate spurious states which arise due to finite resolution effects (see text). (Right:) Corresponding density of states $\rho(\omega)$ computed using a transfer-matrix technique with a supercell of 10^4 layers.

course, once the unit cell is discretized for numerical computation, the number of degrees of freedom and hence the number of eigenvalues is finite. However, as the resolution is increased, not only do the maximum frequency and the accuracy increase as for an ordinary computation, but also the number of bands in a given bandwidth increases. Thus, as the resolution is increased, more and more of the fractal structure of the spectrum is revealed.

B.3 One-Dimensional Results

As a proof of concept implementation of cut-and-project, we construct a Fibonacci quasicrystal in Sec. B.3.1 using the projection method described above, compute the band structure as a function of the projected wave-vector k_d and compare to a transfer-matrix calculation of the same quasicrystal structure. We also demonstrate the field visualization enabled by the projection method, both in the superspace (n dimensions) as well as in the physical space (d dimensions). In Sec. B.3.2, we demonstrate how this method can accommodate systems with defects. Finally, we explore several one-dimensional quasicrystal configurations in Sec. B.3.3 by varying the cut angle ϕ .




Figure B-3: Integrated density of states (DOS) vs. band index (normalized by resolution), for various resolutions. The dashed red, diamond blue and solid black lines denote resolutions of 20, 50 and 200, respectively.

B.3.1 Fibonacci Quasicrystal

We solved Eq. (B.1) numerically using a planewave expansion in the unit cell of the 2D superspace, as described above, for the 1D Fibonacci quasicrystal structure depicted in Fig. B-1. The resulting band diagram is shown in Fig. B-2(left), along with a side-by-side comparison of the local density of states in Fig. B-2(right) calculated using a transfer-matrix approach with a supercell of 10^4 layers [301]. The two calculations show excellent agreement in the location of the gaps, except for one or two easily-identified spurious bands inside some of the gaps, which are discussed in further detail below. The most important feature of Fig. B-2(left) is the large number of bands even in the finite bandwidth $\omega \in [0, 0.4]$, with the number of bands increasing proportional to the spatial resolution (planewave cutoff). This is precisely the feature predicted abstractly above, in Sec. B.2.2: at a low resolution, one sees only the largest gaps, and at higher resolutions further details of the fractal spectrum are revealed as more and more bands appear within a given bandwidth, very different from calculations for periodic physical media. These features are illustrated in Fig. B-3. The important physical quantity is not so much the band structure, since \mathbf{k} has no simple physical meaning as discussed previously, but rather the density of states formed by projecting the band structure onto the ω axis. In this density of states, the small number of spurious bands within the gaps, which arise from the discretization as discussed below, plays no significant role: the density of states is dominated by the huge number of flat bands (going to infinity




Figure B-4: Plot of the magnetic field amplitude $|H_z|$ for a band-edge state taken along a slice of the two-dimensional superspace (in the ϕ direction). *Inset:* Two-dimensional superspace field profile (red/white/blue indicates positive/zero/negative amplitude).

as the resolution is increased), and the addition of one or two spurious bands is negligible.

Computing the eigenmodes in the higher-dimensional superspace immediately suggests a visualization technique: instead of plotting the quasiperiodic fields as a function of the physical coordinates x by taking a slice, plot them in the two-dimensional superspace. This has the advantage of revealing the entire infinite aperiodic field pattern in a single finite plot [108]. Such plots are also used below, to aid in understanding the spurious modes localized at staircased interfaces. A typical extended mode profile is shown in Fig. B-4, plotted both as a function of the physical coordinate x for large supercell and also in the unit cell of the superspace (inset). In the inset superspace plot, one can clearly see the predicted field oscillations perpendicular to the slice plane, as well as a slower oscillation rate (inversely proportional to the frequency) parallel to the slice. In the plot versus x , one can see the longer-range quasi-periodic structure that arises from how the slice wraps around the unit cell in the superspace. The factor of three to four long-range variations in the field amplitude are suggestive of the critically localized states (power-law decay) that one expects to see in such quasicrystals [104, 273, 391].

By visualizing the bands in this way, we can demonstrate the origin of the quasicrystal band gap in an interesting way. In an ordinary photonic crystal, the gap arises because the lowest band concentrates its electric-field energy in the high-dielectric regions (due to the variational principle), while the next band (above the gap) is forced to have a nodal plane in these regions (due to orthogonality) [239]. A very similar phenomenon can be observed

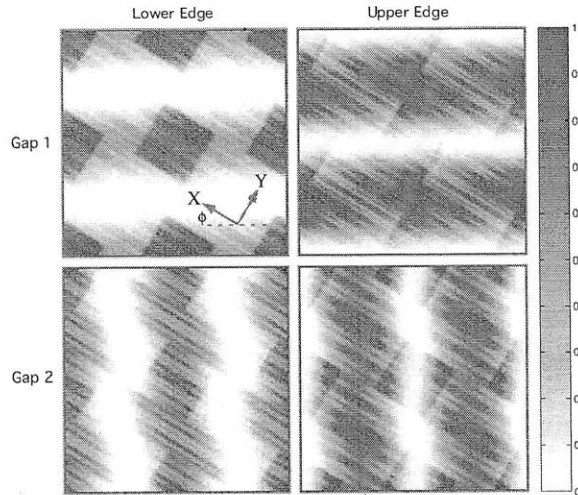


Figure B-5: Electric field energy distribution of the band edge states of gaps 1 and 2 in Fig. B-2. Although they have a complex small-scale structure, the large-scale variation is easily understood in terms of the structure of the superspace.

in the quasicrystal eigenmodes, when plotted in the superspace. In particular, Fig. B-5 displays the electric-field energy distribution of the band-edge states just above and below gaps 1 and 2 of Fig. B-2. Very similar to an ordinary two-dimensional photonic crystal, the bands just below the gaps are peaked in the dielectric squares, whereas the upper-edge bands have a nodal plane in these squares. If the same fields were plotted only in the physical coordinate space, the position of the peaks and nodes would vary between adjacent layers and this global pattern (including the relationship between the two gaps) might not be apparent. In contrast to a two-dimensional photonic crystal, on the other hand, the quasicrystalline field pattern has fractal oscillations in the superspace.

As the wavevector \mathbf{k} varies, most of the bands in the spectrum of Fig. B-2 are flat, except for certain modes (highlighted in blue) which appear to cross the band gaps relatively quickly, as shown in Fig. B-6. These are the spurious modes, as explained below.

In fact, a simple argument shows that, in the limit of infinite resolution, the physical spectrum cannot depend on \mathbf{k} , and hence any strongly \mathbf{k} -dependent band must be a numerical artifact. First, ω cannot depend on the components of \mathbf{k} in the unphysical directions Y , because the Maxwell operator of Eq. (B.1) has no \mathbf{y} -derivatives (equivalently, any phase oscillations in \mathbf{y} commute with the operator). Second, ω cannot depend on the components of \mathbf{k} in the physical directions X , either. The reason is that, from Bloch's theorem, \mathbf{k} and $\mathbf{k} + \mathbf{G}$ give the same eigensolutions for any reciprocal lattice vector \mathbf{G} , and the projections




Figure B-6: Enlarged view of the Fibonacci spectrum (Fig. B-2) showing a gap with a spurious band crossing it. Insets show the magnetic field $|H_z|$ for the spurious band at various k_x —the localization of this mode around the X -parallel edges of the dielectric indicate that this is a discretization artifact.

of the reciprocal lattice vectors are dense in X for a quasicrystal.

These “spurious” bands that appear arise from the discretization of the dielectric interfaces parallel to the slice direction. Because the slice is at an irrational angle, it will never align precisely with a uniform grid, resulting in inevitable staircasing effects at the boundary. With ordinary electromagnetic simulations, these staircasing effects can degrade the accuracy [150], but here the lack of derivatives perpendicular to the slice allows spurious modes to appear along these staircased edges (there is no frequency penalty to being localized perpendicular to the slice). Indeed, if one looks at the field patterns for the spurious modes as a function of \mathbf{k}_d (shown in Fig. B-6), one sees that the field intensity is peaked along the slice-parallel dielectric interfaces. Because they are localized to these interfaces and are therefore dominated by the unphysical staircasing, the spurious modes behave quite differently from the “real” solutions and are easily distinguished qualitatively and quantitatively (e.g via their \mathbf{k} -dependence). Most importantly, as the resolution is increased, the number of spurious modes in a given gap does not increase like all of the other bands, because the thickness of the staircased interface region decreases proportional to the resolution. This makes the gaps in the band structure obvious: here, they are the only frequency ranges for which the number of eigenvalues does not increase with resolution. Equivalently, as noted above, the contribution of the spurious bands to the density of states is asymptotically negligible as resolution is increased.




Figure B-7: Dielectric for the Fibonacci chain with $\epsilon = 2.56$ (light blue), and a defect—an additional $\epsilon_d = 8.0$ layer, shown in gray.

B.3.2 Defect Modes

Much of the interest in quasicrystal band gaps, similar to the analogous case of band gaps in periodic structures, centers around the possibility of localized states: by introducing a defect in the structure, e.g. by changing the thickness of a single layer, one can create exponentially localized states in the gap [27, 89]. In periodic systems, because such defects break the periodicity, they necessitate a larger computational cell, or supercell, that contains many unit cells. In quasicrystal systems once the gaps are known, on the other hand, defect states are arguably *easier* to compute than the gaps of the infinite structure, because an exponentially localized defect mode can be computed accurately with a traditional supercell and the infinite quasicrystal per se need not be included. Nevertheless, the superspace approach allows one to compute defect modes using the same higher-dimensional unit cell, which demonstrates the flexibility of this approach and provides an interesting (but not obviously superior) alternative to traditional supercells for defect states.

Ideally, if one had infinite spatial resolution, a defect in the crystal would be introduced as a very thin perturbation parallel to the slice direction. As the thickness of this perturbation goes to zero, it intersects the physical slice at greater and greater intervals in the physical space, corresponding to localized defects that are separated by arbitrarily large distances. In practice, of course, the thickness of the perturbation is limited by the spatial resolution, but one can still obtain defects that are very widely separated—since the associated defect modes are exponentially localized, the coupling between the defects is negligible.




Figure B-8: Varying the defect epsilon for resolutions 50 (blue) and 100 (red). The thickness of the defect is fixed to 0.02 lattice constants. The number of spurious modes increases with the resolution, the true defect state being the lowest of these modes.

In other words, one effectively has a very large supercell calculation, but expressed in only the unit cell of the higher-dimensional lattice.

As an example, we changed an $\varepsilon = 2.56$ layer to $\varepsilon = \varepsilon_d$ at one place in the Fibonacci quasicrystal. The corresponding superspace dielectric function is shown in Fig. B-7, where the defect is introduced as a thin ($0.02a$) strip of ε_d parallel to the slice direction. We compute the band structure as a function of the defect dielectric constant ε_d , varying it from the normal dielectric $\varepsilon_d = 2.56$ up to $\varepsilon_d = 11$. The thickness of the defect in the unphysical direction was fixed to be ≈ 0.02 . The reason for this is that the defect layer must be greater than one pixel thick in the Y directions in order to avoid staircasing effects in the spectrum. The resulting eigenvalues as a function of ε_d are shown in Fig. B-8 for two different spatial resolutions of 50 (blue) and 100 (red) pixels/ a . When the resolution is 50 the defect is only one pixel thick, the discretization effects might be expected to be large, although the frequency is within about 2% of the higher-resolution calculation. At the higher resolution, the frequency of the mode is converging (it is within 0.3% of a resolution-200 calculation, not shown). However, at the higher resolution there is a second, spurious mode due to the finite thickness (2 pixels) of the defect layer—this spurious mode is easily identified when the field is plotted Fig. B-9 (bottom), because it has a sign oscillation perpendicular to the slice (which would be disallowed if we could make the slice infinitesimally thin).

The defect modes for the resolution 100 are plotted in Fig. B-9 for both the real and the spurious modes, versus the physical coordinate (x) and also in the superspace unit cell




Figure B-9: Semi-log plots of the magnetic field magnitude H_z for the lowest (top) and highest (bottom) defect state for the configuration shown in Fig. B-7. *Insets:* Two-dimensional superspace visualizations of the defect states. Note the additional node in the lower figure (corresponding to an unphysical oscillation).

(insets). When plotted versus the physical coordinate x on a semilog scale, we see that the modes are exponentially localized as expected. The defect mode appears at multiple x values (every $\sim 20a$ on average) because the defect has a finite thickness—the physical slice intersects it infinitely many times (quasiperiodically), as discussed above. The spurious mode (bottom panel) is also exponentially localized; it has a sign oscillation perpendicular to the slice direction (inset) which causes it to have additional phase differences between the different defects.

Nevertheless, as emphasized above, we feel that the main advantages of the superspace approach are for studying the gaps and modes of the infinite, defect-free quasicrystal rather than for localized defect modes.

B.3.3 Continuously Varying the Cut Angle

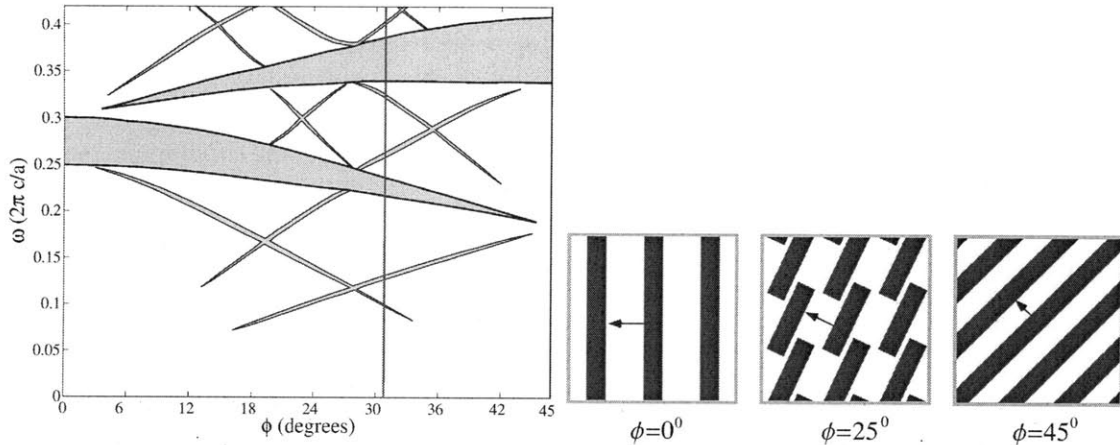


Figure B-10: Projected band structure vs. cut angle ϕ , showing different one-dimensional quasicrystal realizations. The vertical red line indicates the spectrum when the slope is the golden ratio τ (the spectra of ϕ and $\pi - \phi$ are equivalent)

The cut-and-project construction of quasicrystals provides a natural way to parameterize a family of periodic and quasiperiodic structures, via the cut angle ϕ . It is interesting to observe how the spectrum and gaps then vary with ϕ .

As ϕ is varied continuously from 0° to 45° , the structures vary from period a to quasiperiodic lattices (for $\tan \phi$ irrational) to long-period structures ($\tan \phi$ rational with a large denominator) to a period $a\sqrt{2}$ crystal. As we change ϕ , we rotate the objects in the unit cell, so that they are always extruded along the y direction with a length equal to the projection of the unit cell onto y [$a(\sin \phi + \cos \phi)$], corresponding the usual cut-and-project construction [235]. In this case, the spectrum varies continuously with ϕ , where the rational $\tan \phi$ correspond to “rational approximants” of the nearby irrational $\tan \phi$ [513, 552]. For a general unit cell with a rational $\tan \phi$, the physical spectrum might depend on the slice offset y and hence different from the total superspace spectrum, but this is not the case for dielectric structures like the one here, which satisfy a “closeness” condition [513] (the edges of the dielectric rods overlap when projected onto the Y direction). This makes the structure y -independent even for rational slices [513] The resulting structures are shown in the bottom panel of Fig. B-10 for three values of ϕ .

The corresponding photonic band gaps are shown in the top panel of Fig. B-10, as a continuous function of ϕ . Only the largest gaps are shown, of course, since we are unable to resolve the fractal structure to arbitrary resolution. As might be expected, there are

isolated large gaps at $\phi = 0^\circ$ and $\phi = 45^\circ$ corresponding to the simple $ABAB\dots$ periodic structures at those angles (with period a and $a/\sqrt{2}$, respectively, the latter resulting from two layers per unit cell). The $\phi = 45^\circ$ gap is at a higher frequency because of its shorter period, but interestingly it is not continuously connected to the $\phi = 0^\circ$ gap.

The reason for this is that the two gaps are dominated by different superspace reciprocal lattice vectors: $(1, 0) \cdot 2\pi/a$ for $\phi = 0^\circ$, and $(1, 1) \cdot 2\pi/a$ for $\phi = 45^\circ$. (In fact, it is possible to calculate, to first order, the locations of the gaps using the dynamic structure factor $S(\mathbf{k}, \omega)$ obtained from the projection of the superspace lattice [421].) For intermediate angles, a number of smaller gaps open and then close. If we were able to show the spectrum with higher resolution, we would expect to see increasing numbers of these smaller gaps opening, leading to the well-known fractal structure that arises e.g. for the Fibonacci crystal.

B.4 Defining Irrational Slice

In this section, we give an explicit derivation of the fact that an “irrational” slice densely fills the superspace unit cell, or rather a definition of the necessary conditions to be an “irrational” slice. These concepts are widely used in the quasicrystal literature, but a precise definition seems hard to find (one commonly requires that all of the Miller indices have incommensurate ratios, but this condition is stronger than necessary).

Without loss of generality, we can consider the unit cell in the superspace $Z = \mathbb{R}^n$ to be the unit cube (related to any lattice by an affine transformation) with lattice vectors along the coordinate directions. The physical slice X is d -dimensional, and it will be convenient to write the coordinates of a vector \mathbf{z} as $\mathbf{z} = (s_1, \dots, s_d, t_1, \dots, t_{n-d}) = (\mathbf{s}, \mathbf{t})$. By taking every coordinate modulo 1, we can map X to a set \bar{X} consisting of X 's intersection with each unit cell. We wish to show necessary and sufficient conditions for \bar{X} to densely fill the unit cell.

Assuming that the slice is not orthogonal to any of the coordinate axes (as otherwise it would clearly not densely fill the unit cell), we can parameterize the points \mathbf{z} of X so that the last $n - d$ coordinates (t_1, \dots, t_{n-d}) are written as a linear function $\mathbf{t}(s_1, \dots, s_d) \equiv \mathbf{t}(\mathbf{s})$ of the first d coordinates.

Consider the set T in \mathbb{R}^{n-d} formed by the $\mathbf{t}(\mathbf{s})$ coordinates of X when the components of \mathbf{s} take on integer values. This is a subset of X , and the corresponding set \bar{T} formed




Figure B-11: Schematic showing the superspace slice X (left) and the projected slice modulo 1 into the unit cell \bar{X} (right), along with the intersection \bar{T} of \bar{X} with the $s = 0$ hyperplane.

by taking $\mathbf{t} \in T$ modulo 1 is a subset of \bar{X} . The key fact is that \bar{X} is dense in the n -dimensional unit cell if and only if \bar{T} is dense in the $(n - d)$ -dimensional unit cell, and this is the case that we will analyze. This equivalence follows from the fact that \bar{X} is simply \bar{T} translated continuously along the slice directions (every point in \bar{X} is related to a point in \bar{T} by a simple projection). The set T is a lattice in \mathbb{R}^{n-d} consisting of all integer linear combinations of the basis vectors $\mathbf{t}_k = \mathbf{t}(s_j = \delta_{jk})$, since $\mathbf{t}(\mathbf{s})$ is a linear function.

For each basis vector \mathbf{t}_k , it is a well-known fact [392] that if it consists of m incommensurate irrational components, the set of integer multiples $\ell\mathbf{t}_k$ modulo 1 will densely fill an m -dimensional slice of the unit cell. More precisely, write $\mathbf{t}_k = \sum_{j=1\dots m_k} \alpha_k^j \mathbf{b}_k^j + \mathbf{q}_k$, where the \mathbf{b}_k^j and \mathbf{q}_k have purely rational components and the $\{\alpha_j\}$ are incommensurate irrational numbers, and m_k is therefore the number of incommensurate irrational components of \mathbf{t}_k . Then the set of integer multiples of \mathbf{t}_k modulo 1 densely fills an m_k -dimensional slice of the unit cell of \mathbb{R}^{n-d} . The basis vectors of this slice are precisely the vectors \mathbf{b}_k^j , which are rational and therefore commensurate with the basis vectors of \mathbb{R}^{n-d} , while the vector \mathbf{q}_k is simply a rational shift. This slice therefore cuts the unit cell of \mathbb{R}^{n-d} a finite number of times.

The set \bar{T} is then obtained as the direct sum of these dense slices for all $n - d$ vectors \mathbf{t}_k . This is then dense if and only if the set of vectors $\{\mathbf{b}_k^j\}_{k=1\dots d}^{j=1\dots m_k}$ spans \mathbb{R}^{n-d} . In other words, an “irrational slice,” which densely fills the unit cell, is one in which there are $n - d$ independent incommensurate slice components as defined above.

B.5 Concluding Remarks

We have presented a numerical approach to computing the spectra of photonic quasicrystals by directly solving Maxwell's equations extended to a periodic unit cell in higher dimensions, allowing us to exploit Bloch's theorem and other attractive properties of computations for periodic structures. In doing so, we extended the conceptual approach of cut-and-project techniques, which were developed as a way to *construct* quasicrystals, into a way to *simulate* quasicrystals. Compared to traditional supercell techniques, this allows us to capture the entire infinite aperiodic quasicrystal in a single finite computational cell, albeit at only a finite resolution. In this way, the single convergence parameter of spatial resolution replaces the combination of resolution and supercell size in traditional calculations, in some sense uniformly sampling the infinite quasicrystal. The resulting computations, applied to the test case of a Fibonacci quasicrystal, display the unique features of quasicrystals in an unusual fashion, in terms of higher-dimensional band structures and visualization techniques. This technique also allows defects and variation of cut angle (continuously varying between periodic and aperiodic structures) in a straightforward way.

In future work, we plan to apply this approach to modeling higher-dimensional quasicrystal structures, such as the Penrose [517] and 2D Fibonacci tilings [305], where computing the spectrum is currently more challenging using existing supercell techniques. To make a higher-dimensional superspace calculation practical, one must use iterative eigensolver methods [18, 249] rather than the simple dense-matrix techniques employed for our test case. Iterative techniques are most efficient for computing a few eigenvalues at a time, and so it will be useful to employ iterative methods designed to compute “interior” eigenvalues [18, 249], allowing one to search directly for large gaps without computing the lower-lying modes. Alternatively, numerical techniques have been developed, based on filter-diagonalization methods, to directly extract the spectrum of many eigenvalues without computing the corresponding eigenvectors [333].

Bibliography

- [1] Della Villa A., S. Enoch, G. Tayeb, V. Pierro, V. Galdi, and F. Capolino. Bang gap formation and multiple scattering in photonic quasicrystals with Penrose-type lattice. *Phys. Rev. Lett.*, 94:183903, 2005.
- [2] E. Abraham, W. J. Firth, and J. Carr. Self-oscillation and chaos in nonlinear Fabry-Perot resonators with finite response time. *Physics Lett. A*, pages 47–51, 1982.
- [3] G.P. Agrawal. *Applications of Nonlinear Fiber Optics*,. Optics and Photonics. Academic Press, San Diego, CA, 2001.
- [4] G. D. Aguanno, M. Centini, M. Scalora, C. Sibilìa, M. Bertolotti, M. J. Bloemer, and C. M. Bowden. Generalized coupled-mode theory for $\chi^{(2)}$ interactions in finite multi-layered structures. *J. Opt. Soc. Am. B*, 19:2111–2122, 2002.
- [5] G. D. Aguanno, M. Centini, M. Scalora, C. Sibilìa, Y. Dumeige, P. Vidavovic, J. A. Levenson, M. J. Bloemer, C. M. Bowden, J. W. Haus, and M. Bertolotti. Photonic band edge effects in finite structures and applications to $\chi^{(2)}$ interactions. *Phys. Rev. E*, 64:016609, 2001.
- [6] Yoshihiro Akahane, Takashi Asano, Bong-Shik Song, and Susumu Noda). Investigation of high-q channel drop filters using donor-type defects in two-dimensional photonic crystal slabs. *Appl. Phys. Lett.*, 83(8):1512–1514, 2003.
- [7] Vilson R. Almeida, Carlos A. Barrios, Roberto R. Panepucci, and Michal Lipson. All-optical control of light on a silicon chip. *Nature*, 431:1081–1084, 2004.
- [8] A. M. Ampere. *Mathematique Des Phenomenes Electro-Dynamiques, Uniquement Deduite de L'experience*. A. Blanchard, 1958.
- [9] A. Andornico, J. Claudon, J. M. Gerard, V. Berger, and G. Leo. Integrated terahertz source based on three-wave mixing of whispering-gallery mode. *Opt. Lett.*, 33(21):2416–2418, 2008.
- [10] A. André, L.-M. Duan, and M. D. Lukin. Coherent atom interactions mediated by dark-state polaritons. *Phys. Rev. Lett.*, 88(24):243602, 2002.
- [11] Lucio C. Andreani, Giovanni Panzarini, and Jean-Michel Gerard. *Phys. Rev. B*, 60:13276, 1999.
- [12] Dimitris G. Angelakis, Marcelo Franca Santos, and Sougato Bose. Photon-blockade-induced mott transitions and xy spin models in coupled cavity arrays. *Phys. Rev. A*, 76(3):031805, 2007.

- [13] Mauro Antezza, Lev P. Pitaevskii, Sandro Stringari, and Vitaly B. Svetovoy. *Phys. Rev. Lett.*, 97:223203, 2006.
- [14] E. Arimondo. *Progress in Optics*, volume 35, page 259. Elsevier Science, Amsterdam, 1997.
- [15] J. A. Armstrong, N. loembergen, J. Ducuing, and P. S. Pershan. Interactions between light waves in a nonlinear dielectric. *Phys. Rev.*, 127:1918–1939, 1962.
- [16] A. Ashkin, G. D. Boyd, and J. M. Dziedzic. Resonant optical second harmonic generation and mixing. *IEEE J. Quantum Electron.*, 2:109–124, 1966.
- [17] Y. H. Avetisyan. Cavity-enhanced terahertz region difference frequency generation in surface-emitting geometry. *Proc. SPIE*, 3795:501, 1999.
- [18] Zhaojun Bai, James Demmel, Jack Dongarra, Axel Ruhe, and Henk Van Der Vorst. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, Philadelphia, 2000.
- [19] A. V. Balakin, V. A. Bushuev, B. I. Mantsyzov, I. A. Ozheredov, E. V. Petrov, and A. P. Shkurinov. Enhancement of sum frequency generation near the photonic band edge under the quasiphase matching condition. *Phys. Rev. E*, 63:046609, 2001.
- [20] Constantine A. Balanis. *Antenna Theory: Analysis and Design*. Wiley, New York, 2nd edition, 1996.
- [21] R. Balian and B. Duplantier. Electromagnetic waves near perfect conductors. II. Casimir effect. *Ann. Phys.*, 112:165–208, 1978.
- [22] Phillip Ball. Fundamental physics: feel the force. *Nature*, 447:772–774, 2007.
- [23] Gan Bao, Lawrence Cowsar, and Wen Masters, editors. *Mathematical Modeling in Optical Science*, volume 22 of *Frontiers in Applied Mathematics*, chapter 7. SIAM, 2001.
- [24] P. E. Barclay, K. Srinivasan, and O. Painter. Nonlinear response of silicon photonic crystal microresonators excited via an integrated waveguide and fiber taper. *Opt. Express*, 13:801–820, 2005.
- [25] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, j. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, Philadelphia, 2nd edition, 1994.
- [26] M. Bayer, T.L. Reinecke, F. Weidner, A. Larionov, A. McDonald, and A. Forchel. *Phys. Rev. Lett.*, 86:3168, 2001.
- [27] M. Bayinding, E. Cubukcu, I. Bulu, and E. Ozbay. Photonic band gaps and localization in two-dimensional metallic quasicrystals. *Europhys. Lett.*, 56(1):41–46, 2001.
- [28] Matthew Beard, Gordon M. Turner, and Charles A. Schmuttenmaer. Terahertz spectroscopy. *J. Phys. Chem. B*, 106(29):7146–7159, 1997.

- [29] R.G. Beausoleil, W.J. Munro, and T.P. Spiller. *ArXiv:quant-ph*, page 0302109 v3, 2005.
- [30] Mikhail A. Belkin, Jonathan A. Fan, Sahand Hormoz, Federico Capasso, Suraj P. Khanna, Mohamed Lachab, A. G. Davies, and Edmund H. Linfield. Terahertz quantum cascade lasers with copper metal-metal waveguides operating up to 178 k. *Opt. Express*, 16(5):3242–3248, 2008.
- [31] C.H. Bennett, P.W. Shor, J.A. Smolin, and A.V. Thapliyal. *Phys. Rev. Lett.*, 85:3081, 1999.
- [32] V. S. Bentsen, R. Herikstad, S. Skriudalen, I. Brevik, and J. S. Hoye. Calculation of the Casimir force between similar and dissimilar metal plates at finite temperature. *J. Phys. A: Math. Gen.*, 38(43), 2005.
- [33] V Berger. Second-harmonic generation in monolithic cavities. *J. Opt. Soc. Am. B*, 14:1351, 1997.
- [34] Lennart Bergstrom. Hamaker constants of inorganic materials. *Adv. Colloid and Interface Science*, 70:125–169, 1997.
- [35] P. Bermel, J.D. Joannopoulos, Y. Fink, P.A. Lane, and C. Tapalian. *Phys. Rev. B*, 69:035316, 2004.
- [36] P. Bermel, A. Rodriguez, J. D. Joannopoulos, and M. Soljacic. Tailoring optical nonlinearities via the Purcell effect. *Phys. Rev. Lett.*, 99:053601, 2007.
- [37] P. Bermel, Alejandro Rodriguez, Steven G. Johnson, J. D. Joannopoulos, and Marin Soljačić. Single-photon all-optical switching using waveguide-cavity quantum electrodynamics. *Phys. Rev. A*, 74:043818, 2006.
- [38] J. C. Berquist, H. Hemmati, and W. M. Itano. High power second harmonic generation of 257 nm radiation in an external ring cavity. *Opt. Commun.*, 43:437–442, 1982.
- [39] D .S. Bethune. Optical harmonic generation and mixing in multilayer media: analysis using optical transfer matrix techniques. *J. Opt. Soc. Am. B*, 6:910–916, 1989.
- [40] V. B. Bezerra, G. L. Klimchitskaya, and C. Romero. Surface roughness contribution to the Casimir interaction between an isolated atom and a cavity wall. *Phys. Rev. A*, 61:022115, 2000.
- [41] G. C. Bhar and U. Chatterjee. Analyses of phase-matching for noncollinear three-wave mixing in uniaxial crystals. *Japanese J. App. Phys.*, 29(6):1103–1107, 1990.
- [42] M. Bieler. Thz generation from resonant excitation of semiconductor nanostructures: Investigation of second-order nonlinear optical effects. *IEEE J. Sel. Top. Quant. Elec.*, 14(2):458–469, 2008.
- [43] K.M. Birnbaum, A. Boca, R. Miller, A.D. Boozer, T.E. Northup, and H.J. Kimble. *Nature*, 436:87, 2005.
- [44] K.-J. Boller, Atac Imamoglu, and Steven E. Harris. *Phys. Rev. Lett.*, 66:2593, 1991.

- [45] Marc Bonnet. *Boundary Integral Equation Methods for Solids and Fluids*. Wiley, Chichester, England, 1999.
- [46] M. Bordag. Casimir effect for a sphere and a cylinder in front of a plane and corrections to the proximity force theorem. *Phys. Rev. D*, 73:125018, 2006.
- [47] M. Bordag, B. Geyer, G. L. Klimchitskaya, and V. M. Mostepanenko. Casimir force at both nonzero temperature and finite conductivity. *Phys. Rev. Lett.*, 85(3):503, 2000.
- [48] M. Bordag, U. Mohideen, and V. M. Mostepanenko. New developments in the Casimir effect. *Phys. Rep.*, 353:1–205, 2001.
- [49] M. Bostrom and Bo E. Sernelius. Thermal effects on the Casimir force in the 0.1-0.5 μm range. *Phys. Rev. Lett.*, 84(20):4757, 2000.
- [50] J. P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover, New York, 2nd edition, 2001.
- [51] Robert W. Boyd. *Nonlinear Optics*. Academic Press, California, 1992.
- [52] Robert W. Boyd. *Nonlinear Optics*. Academic Press, San Diego, 1992.
- [53] T. H. Boyer. Quantum electrodynamic zero-point energy of a conducting spherical shell and the casimir model for a charged particle. *Phys. Rev.*, 174:1764, 1968.
- [54] Timothy H. Boyer. Van der Waals forces and zero-point energy for dielectric and permeable materials. *Phys. Rev. A*, 9:2078–2084, 1974.
- [55] J. Bravo-Abad, S. Fan, S. G. Johnson, J. D. Joannopoulos, and M. Soljacic. Modeling nonlinear optical phenomena in nanophotonics. *J. Lightwave Tech.*, 25:2539–2546, 2007.
- [56] J. Bravo-Abad, A. W. Rodriguez, J. D. Joannopoulos, P. T. Rakich, S. G. Johnson, and M. Soljacic. Efficient low-power terahertz generation via on-chip triply-resonant nonlinear frequency mixing. *Appl. Phys. Lett.*, 96:101110, 2010.
- [57] Jorge Bravo-Abad, Alejandro W. Rodriguez, Peter Bermel, Steven G. Johnson, J. D. Joannopoulos, and Marin Soljačić. Enhanced nonlinear optics in photonic-crystal nanocavities. *Opt. Express*, 15(24):16161–16176, 2007.
- [58] F. Brechet, P. Roy, J. Marcou, and D. Pagnoux. Singlemode propagation into depressed-core-index photonic-bandgap fibre designed for zero-dispersion propagation at short wavelengths. *Electron. Lett.*, 36(6):514–515, 2000.
- [59] G. Bressi, G. Carugno, R. Onofrio, and G. Ruoso. Measurement of the Casimir force between parallel metallic surfaces. *Phys. Rev. Lett.*, 88:041804, 2002.
- [60] I. Brevick and G. Einevoll. Casimir force on a solid ball when $\varepsilon(\omega)\mu(\omega) = 1$. *Phys. Rev. D*, 37:2977, 1987.
- [61] I. Brevik, J. B. Aarseth, J. S. Høye, and K. A. Milton. Temperature dependence of the Casimir effect. *Phys. Rev. E*, 71(056101), 2005.

- [62] I. Brevik, E. K. Dahl, and G. O. Myhr. Casimir force on a micrometer sphere in a dip: proposal of an experiment. *J. Phys. A: Mathematical and General*, 38:49–56, 2005.
- [63] I. Brevik and M. Lygren. Casimir effect for a perfectly conducting wedge. *Ann. Phys.*, 51:157, 1996.
- [64] Iver Brevik and Simen A. Ellingsen. Comment on “casimir force acting on magnetodielectric bodies embedded in media”. *Phys. Rev. A*, 79(2):027801, 2009.
- [65] M. Brieger, H. Busener, A. Hese, F. V. Moers, and A. Renn. Enhancement of single frequency shg in a passive ring resonator. *Opt. Commun.*, 38:423–426, 1981.
- [66] X. Brokmann, L. Coolen, M. Dahan, and J. P. Hermier. *Phys. Rev. Lett.*, 93:107403, 2004.
- [67] M. Brown-Hayes, D. A. R. Davit, F. D. Mazzitelli, W. J. Kim, and R. Onofrio. Towards a precision measurement of the Casimir force in a cylinder-plane geometry. *Phys. Rev. A*, 72:052102, 2005.
- [68] S. Y. Buhmann and Dirk-Gunnar Welsch. Dispersion forces in macroscopic quantum electrodynamics. *Prog. Quant. Elec.*, 31(2):51–130, 2007.
- [69] Aurel Bulgac and Andreas Wirzba. Casimir interaction among objects immersed in a fermionic environment. *Phys. Rev. Lett.*, 87(12):120404, 2001.
- [70] I. B. Burgess, A. W. Rodriguez, M. W. McCutcheon, J. Bravo-Abad, Y. Zhang, S. G. Johnson, and M. Loncar. Difference-frequency generation with quantum-limited efficiency in triply-resonant nonlinear cavities. *OE*, 17(11):9241–9251, 2009.
- [71] R. Büscher and T. Emig. Nonperturbative approach to Casimir interactions in periodic geometries. *Phys. Rev. A*, 69:062101, 2004.
- [72] R. Büscher and T. Emig. Geometry and spectrum of Casimir forces. *Phys. Rev. Lett.*, 94:133901, 2005.
- [73] Antoine Canaguier-Durand, Paulo A. Maia Neto, Ines Cavero-Pelaez, Astrid Lambrecht, and Serge Reynaud. Casimir interaction between plane and spherical metallic surfaces. *Phys. Rev. Lett.*, 102(23):230404, 2009.
- [74] Federico Capasso, Jeremy N. Munday, Davide Iannuzzi, and H. B. Chan. Casimir forces and quantum electrodynamical torques: Physics and nanomechanics. *IEEE J. Selected Topics in Quant. Elec.*, 13(2):400–415, 2007.
- [75] V. P. Carey, G. Cheng, C. Grigoropoulos, M. Kaviani, and A. Majumdar. A review of heat transfer physics. *Nanoscale Micro. Thermophys. Eng.*, 12(1):1–60, 2006.
- [76] Tal Carmon and Kerry J. Vahala. Visible continuous emission from a silica microphotonic device by third-harmonic generation. *Nature*, 3:430–435, 2007.
- [77] H. B. G. Casimir. On the attraction between two perfectly conducting plates. *Proc. K. Ned. Akad. Wet.*, 51:793–795, 1948.
- [78] H. B. G. Casimir and D. Polder. The influence of retardation on the London-Van der Waals forces. *Phys. Rev.*, 13(4):360–372, 1948.

- [79] R. M. Cavalcanti. Casimir force on a piston. *Phys. Rev. D*, 69:065015, 2004.
- [80] E. Centeno and D. Felbacq. *Phys. Rev. B*, 62:R7683, 2000.
- [81] E. Centeno and D. Felbacq. Optical bistability in finite-size nonlinear bidimensional photonic crystals doped by a microcavity. *Phys. Rev. B*, 62:R7683–R7686, 2000.
- [82] H. B. Chan, V. A. Aksyuk, R. N. Kleinman, D. J. Bishop, and F. Capasso. Nonlinear micromechanical Casimir oscillator. *Phys. Rev. Lett.*, 87:211801, 2001.
- [83] H. B. Chan, V. A. Aksyuk, R. N. Kleinman, D. J. Bishop, and F. Capasso. Quantum mechanical actuation of microelectromechanical systems by the Casimir force. *Science*, 291:1941–1944, 2001.
- [84] H. B. Chan, Y. Bao, J. Zou, R. A. Cirelli, F. Klemens, W. M. Mansfield, and C. S. Pai. Measurement of the Casimir force between a gold sphere and a Silicon surface with a nontrench array. *Phys. Rev. Lett.*, 101:030401, 2008.
- [85] Y. S. Chan and Z. Y. Liu. Photonic band gaps in two dimensional photonic quasicrystals. *Phys. Rev. Lett.*, 80(5):956–959, 1998.
- [86] F. Chen, G. L. Klimchitskaya, V. M. Mostepanenko, and U. Mohideen. Demonstration of the difference in the casimir force for samples with different charge-carrier densities. *Phys. Rev. Lett.*, 97(17):170402, 2006.
- [87] F. Chen, G. L. Klimchitskaya, V. M. Mostepanenko, and U. Mohideen. Control of the Casimir force by the modification of dielectric properties of light. *Phys. Rev. B*, 76:035338, 2007.
- [88] F. Chen, U. Mohideen, G. L. Klimchitskaya, and V. M. Mostepanenko. Demonstration of the lateral Casimir force. *Phys. Rev. Lett.*, 88:101801, 2002.
- [89] Samuel S. M. Cheng, Lie-Ming Li, C. T. Chan, and Z. Q. Zhang. Defect and transmission properties of two-dimensional quasiperiodic photonic band-gap systems. *Phys. Rev. B*, 59(6):4091–4099, 1999.
- [90] W. C. Chew, J. M. Jin, E. Michelssen, and J. M. Song. Fast solution methods in electromagnetics. *IEEE Trans. Ant. Prop.*, 45(3):533–543, 1997.
- [91] Weng Cho Chew, Jin Jian-Ming, Eric Michielssen, and Song Jiming. *Fast and Efficient Algorithms in Computational Electromagnetics*. Artech, Norwood, MA, 2001.
- [92] R. Y. Chiao and J. Boyce. Superluminality, pareticity, and Earnshaw’s theorem in media with inverted populations. *Phys. Rev. Lett.*, 73:3383, 1994.
- [93] Andreas Christ and Hans L. Hartnagel. Three-dimensional finite-difference method for the analysis of microwave-device embedding. *IEEE Trans. Microwave Theory Tech.*, 35(8):688–696, 1987.
- [94] Bosko M. Cirjanic. A practical field plotter using a hybrid analog computer. *IEEE Trans. Ind. Electron.*, IECI-18(1):16–20, 1971.

- [95] R. Coccioli, M. Boroditsky, K. W. Kim, Y. Rahmat-Samii, and E. Yablonovitch. Smallest possible electromagnetic mode volume in a dielectric cavity. *IEE Proc. Optoelec.*, 145(6):391–397, 1998.
- [96] C. Cohen-Tannoudji, B. Din, and F. Laloë. *Quantum Mechanics*. Hermann, Paris, 1977.
- [97] Claude Cohen-Tannoudji, Bernard Diu, and Franck Laloë. *Quantum Mechanics*. John Wiley and Sons, New York, 1977.
- [98] M. J. Collet and R. B. Levien. Two-photon loss model of intracavity second-harmonic generation. *Phys. Rev. A*, 43(9):5068–5073, 1990.
- [99] W. L. Collett, C. A. Ventrice, and S. M. Mahajan. Electromagnetic wave technique to determine radiation torque on micromachines driven by light. *Appl. Phys. Lett.*, 82(16):2730–2733, 2003.
- [100] Robert E. Collin. *Field Theory of Guided Waves*. IEEE Press, New York, NY, second edition, 1991.
- [101] A. R. Cowan and J. F. Young. Optical bistability involving photonic crystal microcavities and fano line shapes. *Phys. Rev. E*, 68:046606, 2003.
- [102] A. R. Cowan and Jeff F. Young. Mode matching for second-harmonic generation in photonic crystal waveguides. *Phys. Rev. E*, 65:085106, 2002.
- [103] L. Dal Negro, M. Stolfi, Y. Yi, J. Michel, X. Duan, L. C. Kimerling, J. LeBlanc, and J. Haavisto. Photon band gap properties and omnidirectional reflectance in si/siO₂ thue-morse quasicrystals. *Appl. Phys. Lett.*, 84(25):5186–5188, 2004.
- [104] Luca Dal Negro, Claudio J. Oton, Zeno Gaburro, Lorenzo Pavesi, Patrick Johnson, Ad Lagendijk, Roberto Righini, Marcello Colocci, and Diedrik Wiersma. Light transport through the band-edge states of Fibonacci quasicrystals. *Phys. Rev. Lett.*, 90(5):055501, 2003.
- [105] D. A. Dalvit, P. A. Mia Neto, A. Lambrecht, and S. Reynaud. Proving quantum-vacuum geometrical effects with cold atoms. *Phys. Rev. Lett.*, 100(040405), 2008.
- [106] D. A. R. Dalvit, F. C. Lombardo, F. D. Mazzitelli, and R. Onofrio. Casimir force between eccentric cylinders. *Europhys. Lett.*, 67:517, 2004.
- [107] D. A. R. Dalvit, F. C. Lombardo, F. D. Mazzitelli, and R. Onofrio. Exact Casimir interaction between eccentric cylinders. *Phys. Rev. A*, 74:020101(R), 2006.
- [108] C. de Lange and T. Janssen. Electrons in incommensurate crystals: Spectrum and localization. *Phys. Rev. B*, 28(1):195–209, 1983.
- [109] R. S. Decca, F. Fischbach, G. L. Klimchitskaya, D. E. Krause, D. Lopez, and V. M. Mostepanenko. Improved tests of extra-dimensional physics and thermal quantum field theory from new Casimir force measurements. *Phys. Rev. D*, 68:116003, 2003.
- [110] R. S. Decca, D. Lopez, E. Fischbach, G. L. Klimchitskaya, D. E. Krause, and V. M. Mostepanenko. Tests of new physics from precise measurements of the Casimir pressure between two gold-coated spheres. *Phys. Rev. D*, 75:077101, 2007.

- [111] R. S. Decca, D. Lopez, E. Fischbach, and D. E. Krause. Measurement of the Casimir force between dissimilar metals. *Phys. Rev. Lett.*, 91:050402, 2003.
- [112] A. Della Villa, S. Enoch, G. Tayeb, V. Pierro, and V. Galdi. Localized modes in photonic quasicrystals with Penrose-type lattice. *Opt. Express*, 14(21):10021–10027, 2006.
- [113] Alessandro Della Villa, Vincenzo Galdi, Filippo Capolino, Vincenzo Pierro, Stefan Enoch, and Gerard Tayeb. A comparative study of representative categories of EGB dielectric quasi-crystals. *IEEE Ant. Wireless Prop. Lett.*, 5:331–333, 2006.
- [114] W. Demtroder. *Laser Spectroscopy*. Springer, 3 edition, 2003.
- [115] Parag B. Deotare, Murray W. McCutcheon, Ian W. Frank, Mughees Khan, and Marko Loncar. High quality factor photonic crystal nanobeam cavities. *Appl. Phys. Lett.*, 94:121106, 2009.
- [116] B. V. Derjaguin, I. I. Abrikosova, and E. M. Lifshitz. Direct measurement of molecular attraction between solids separated by a narrow gap. *Q. Rev. Chem. Soc.*, 10:295–329, 1956.
- [117] Ivan H. Deutsch, Raymond Y. Chiao, and John C. Garrison. Two-photon bound states: The diphoton bullet in dispersive self-focusing media. *Phys. Rev. A*, 47(4):3330–3336, 1993.
- [118] Andrea Di Falco, Claudio Conti, and Gaetano Assanto. Impedance matching in photonic crystal microcavities for second-harmonic generation. *Opt. Lett.*, 31:250, 2006.
- [119] P. A. M. Dirac. *Principles of Quantum Mechanics*. Clarendon, Oxford, 1982.
- [120] A. Ditkowski, K. Dridi, and J. S. Hesthaven. Convergent cartesian grid methods for maxwell’s equations in complex geometries. *J. Comp. Phys.*, 170:39–80, 2001.
- [121] G. J. Dixon, C. E. Tanner, and C. E. Wieman. 432-nm source based on efficient second-harmonic generation of gaalas diode-laser radiation in a self-locking external resonant cavity. *Opt. Lett.*, 14:731–733, 1989.
- [122] Babette Dobrich, Maarten DeKieviet, and Holger Gies. Scalar Casimir-Polder forces for uniaxial corrugations. *Phys. Rev. D*, 78:125022, 2008.
- [123] Tatyana V. Dolgova, Anton I. Maidykovski, Michail G. Martemyanov, Andrey A. Fedyanin, Oleg A. Aktsipetrov, Gerd Marowsky, Vladimir A. Yakovlev, Giorgio Mattei, Narumi Ohta, and Seiichiro Nakabayashi. Giant optical second-harmonic generation in single and coupled microcavities formed from one-dimensional photonic crystals. *J. Opt. Soc. Am. B*, 19:2129, 2002.
- [124] P.G. Drazin and R.S. Johnson. *Solitons: an Introduction*. Cambridge University Press, Cambridge, England, 1989.
- [125] P. D. Drummond, K. J. McNeil, and D. F. Walls. Non-equilibrium transitions in sub/second harmonic generation i: Semiclassical theory. *Optica Acta.*, 27(3):321–335, 1980.

- [126] Yannick Dumeige and Patrice Feron. Whispering-gallery-mode analysis of phase-matched doubly resonant second-harmonic generation. *Phys. Rev. A*, 74:063804, 2006.
- [127] Laurent Duraffourg and Philippe Andreucci. Casimir force between doped silicon slabs. *Phys. Lett. A*, 359:406–411, 2006.
- [128] Gupta S. Dutta and Jose Jolly. Third harmonic generation in layered media in presence of optical bistability of the fundamental. *Pranama-journal of Physics*, 50:239, 1988.
- [129] I. E. Dzyaloshinskiĭ, E. M. Lifshitz, and L. P. Pitaevskiĭ. The general theory of van der Waals forces. *Adv. Phys.*, 10(38):165–209, 1961.
- [130] W. Eckhardt. First and second fluctuation-dissipation-theorem in electromagnetic fluctuation theory. *Opt. Comm.*, 41(5):305–309, 1981.
- [131] W. Eckhardt. Macroscopic theory of electromagnetic fluctuations and stationary radiative heat transfer. *Phys. Rev. A*, 29(4):1991–2003, 1984.
- [132] T. Ederth. Template-stripped gold surfaces with 0.4-nm rms roughness suitable for force measurements: Application to the Casimir force in the 20–100-nm range. *Phys. Rev. A*, 62:062104, 2000.
- [133] M. D. Eisaman, A. Andre, F. Massou, M. Fleishhauer, A. S. Zibrov, and M. D. Lukin. Electromagnetically-induced transparency with tunable single-photon pulses. *Nature*, 438:837–841, 2005.
- [134] R. Eisenschitz, F. London, and Z. Physik. *H. Hettema, Quant. Chem.*, 60:491, 1930.
- [135] T. Emig. Casimir forces: An exact approach for periodically deformed objects. *Europhys. Lett.*, 62:466, 2003.
- [136] T. Emig and R. Büscher. Towards a theory of molecular forces between deformed media. *Nucl. Phys. B*, 696:468, 2004.
- [137] T. Emig, N. Graham, R. L. Jaffe, and M. Kardar. Casimir forces between compact objects: the scalar case. *Phys. Rev. D*, 77:025005, 2008.
- [138] T. Emig, A. Hanke, R. Golestanian, and M. Kardar. Probing the strong boundary shape dependence of the Casimir force. *Phys. Rev. Lett.*, 87:260402, 2001.
- [139] T. Emig, A. Hanke, R. Golestanian, and M. Kardar. Probing the strong boundary shape dependence of the Casimir force. *Phys. Rev. Lett.*, 87:260402, 2001.
- [140] T. Emig, A. Hanke, R. Golestanian, and M. Kardar. Normal and lateral Casimir forces between deformed plates. *Phys. Rev. A*, 67:022114, 2003.
- [141] T. Emig, R. L. Jaffe, M. Kardar, and A. Scardicchio. Casimir interaction between a plate and a cylinder. *Phys. Rev. Lett.*, 96:080403, 2006.
- [142] Thornsten Emig. Casimir-force-driven ratchets. *Phys. Rev. Lett.*, 98:160801, 2007.
- [143] Thornsten Emig, N. Graham, R. L. Jaffe, and Mehran Kardar. Casimir forces between arbitrary compact objects. *Phys. Rev. Lett.*, 99:170403, 2007.

- [144] Dirk Englund, David Fattal, Edo Waks, Glenn Solomon, Bingyang Zhang, Toshihiro Nakaoka, Yasuhiko Arakawa, Yoshihisa Yamamoto, and Jelena Vuckovic. *Phys. Rev. Lett.*, 95:013904, 2005.
- [145] Li Fan, Hsu Ta-Chen, Mahmoud Fallahi, James T. Murray, Robert Bedford, Yushi Kaneda, Jorg Hader, Aramais R. XZakharian, Jerome Moloney, Stephan W. Koch, and Wolfgang Stolz. Tunable watt-level blue-green vertical-external-cavity surface-emitting lasers by intracavity frequency doubling. *Appl. Phys. Lett.*, 88:2251117, 2006.
- [146] Shanhui Fan, Steven G. Johnson, J. D. Joannopoulos, C. Manolatou, and H. A. Haus. Waveguide branches in photonic crystals. *J. Opt. Soc. Am. B*, 18(2):162–165, 2001.
- [147] Shanhui Fan, P. R. Villeneuve, and J. D. Joannopoulos. Theoretical investigation of fabrication-related disorder on the properties of photonic crystals. *J. Appl. Phys.*, 78:1415–1418, 1995.
- [148] Shanhui Fan, Pierre R. Villeneuve, J. D. Joannopoulos, and H. A. Haus. Channel drop tunneling through localized states. *Phys. Rev. Lett.*, 80(5):960–963, 1998.
- [149] Andrei Faraon, Ilya Fushman, Dirk Englund, and Jelena Vučkovic. Photon blockade in a photonic crystal cavity with a strongly coupled quantum dot. *Advanced Optical Concepts in Quantum Computing, Memory, and Communication II*, 7225(1):72250D, 2009.
- [150] Ardavan Farjadpour, David Roundy, Alejandro Rodriguez, M. Ibanescu, Peter Bermel, Jeff Burr, J. D. Joannopoulos, and Steven G. Johnson. Improving accuracy by subpixel smoothing in the finite-difference time domain. *Opt. Lett.*, 31:2972–2974, 2006.
- [151] Ardavan Farjadpour, David Roundy, Alejandro Rodriguez, Mihai Ibanescu, Peter Bermel, John D. Joannopoulos, Steven G. Johnson, and G.W. Burr. *Opt. Lett.*, 31:2972, 2006.
- [152] A. A. Feiler, L. Bergstrom, and M. W. Rutland. Superlubricity using repulsive van der Waals forces. *Langmuir*, 24:2274–2276, 2008.
- [153] F. S. Felber and J. H. Marburger. Theory of nonresonant multistable optical devices. *Appl. Phys. Lett.*, 28(12):731–733, 1976.
- [154] Zhifang Feng, Xiangdong Zhang, Yiquan Wang, Zhi-Yuan Li, Bingying Cheng, and Dao-Zhong Zhang. Negative refraction and imaging using 12-fold-symmetry quasicrystals. *Phys. Rev. Lett.*, 94(24):247402, 2005.
- [155] Y. Fink, D. J. Ripin, S. Fan, C. Chen, J. D. Joannopoulos, and E. L. Thomas. Guiding optical light in air using an all-dielectric structure. *J. Lightwave Tech.*, 17:2039–2041, 1999.
- [156] Y. Fink, J. N. Winn, S. Fan, C. Chen, J. Michel, J. D. Joannopoulos, and E. L. Thomas. A dielectric omnidirectional reflector. *Science*, 282:1679–1682, 1998.
- [157] L. H. Ford. Spectrum of the Casimir effect and the Lifshitz theory. *Phys. Rev. A*, 48:2962, 1993.

- [158] J. S. Foresi, P. R. Villeneuve, J. Ferrera, E. R. Thoen, G. Steinmeyer, S. Fan, J. D. Joannopoulos, L. C. Kimerling, Henry I. Smith, and E. P. Ippen. Photonic-bandgap microcavities in optical waveguides. *Nature*, 390:143–145, 1997.
- [159] Barak Freedman, Guy Bartal, Mordechai Segev, Ron Lifshitz, Demetrios N. Christodoulides, and Jason W. Fleischer. Wave and defect dynamics in nonlinear photonic quasicrystals. *Nature*, 440:1166–1169, April 2006.
- [160] W. R. Frei and H. T. Johnson. Finite-element analysis of disorder effects in photonic crystals. *Phys. Rev. B*, 70:165116, 2004.
- [161] D. G. Friend and M. L. Hiber. Thermophysical property standard reference data from NIST. *Int. J. Thermophys.*, 15(6), 1994.
- [162] J. N. Fuchs, A. Recati, and W. Zwerger. Oscillating Casimir force between impurities and one-dimensional Fermi liquids. *Phys. Rev. A*, 75:043615, 2007.
- [163] J. U. Fürst, D. V. Strekalov, D. Elser, M. Lassen, U. L. Andersen, C. Marquardt, and G. Leuchs. Naturally phase-matched second-harmonic generation in a whispering-gallery-mode resonator. *Phys. Rev. Lett.*, 104(15):153901, 2010.
- [164] M. Gadenne. *Optical Properties of Nanostructured Random Media*. Topics Appl. Phys. Springer, Berlin, 2002.
- [165] Robert C. Gauthier and Khaled Mnaymneh. FDTD analysis of 12-fold photonic quasicrystal central pattern localized states. *Opt. Comm.*, 264:78–88, 2006.
- [166] Robert C. Gauthier and Khaled Mnaymneh. Photonic band gaps properties of 12-fold quasicrystal determined through fdtd analysis. *Opt. Express*, 13(6):1985–1998, 2005.
- [167] W. Gellermann, M. Kohmoto, B. Sutherland, and P.C. Taylor. Localization of light waves in fibonacci dielectric multilayers. *Phys. Rev. Lett.*, 72(5):633–636, 1994.
- [168] C. Genet, A. Lambrecht, P. Maia Neto, and S. Reynaud. The Casimir force between rough metallic plates. *Europhys. Lett.*, 62:484, 2003.
- [169] C. Genet, A. Lambrecht, and S. Reynaud. The Casimir effect in the nanoworld. *Eur. Phys. J. Special Topics*, 160:183–193, 2008.
- [170] Cyriaque Genet, Astrid Lambrecht, and Serge Reynaud. Casimir force and the quantum theory of lossy optical cavities. *Phys. Rev. A*, 67(4):043811, 2003.
- [171] A. C. Genz and A. A. Malik. An imbedded family of fully symmetric numerical integration rules. *SIAM J. Numer. Anal.*, 20(3):580–588, 1983.
- [172] A. I. Gerguson and M .H. Dunn. Intracavity second harmonic generation in continuous-wave dye lasers. *IEEE J. Quantum Electron.*, 13:751–756, 1977.
- [173] Christopher C. Gerry. Generation of optical macroscopic quantum superposition states via state reduction with a Mach-Zehnder interferometer containing a Kerr medium. *Phys. Rev. A*, 59(5):4095–4098, 1998.
- [174] Klaus M. Gheri, Wieland Alge, and Philippe Grangier. *Phys. Rev. A*, 60:R2673, 1999.

- [175] H. M. Gibbs. *Optical Bistability: Controlling Light with Light*. Academic Press, Orlando, FL, 1985.
- [176] H.M. Gibbs. *Optical Bistability: Controlling Light with Light*. Academic Press, Orlando, FL, 1985.
- [177] H. Gies, K. Langfeld, and L. Moyaerts. Casimir effect on the worldline. *J. High Energy Phys.*, 6:018, 2003.
- [178] Holger Gies and Klaus Klingmuller. Casimir edge effects. *Phys. Rev. Lett.*, 97:220405, 2006.
- [179] Holger Gies and Klaus Klingmuller. Casimir effect for curved geometries: Proximity-force-approximation validity limits. *Phys. Rev. Lett.*, 96:220401, 2006.
- [180] Holger Gies and Klaus Klingmuller. Worldline algorithms for Casimir configurations. *Phys. Rev. D*, 74:045002, 2006.
- [181] C. Godreche. Indexing the diffraction spectrum of a non-pisot self-similar structure. *Phys. Rev. B*, 45(1):176–185, 1992.
- [182] Israel Gohberg, Sehmour Goldberg, and Marinus A. Kaashoek. *Basic Classes of Linear Operators*. Birkhäuser, Basel, 2000.
- [183] Nigel Goldenfeld. *Lectures on phase transitions and the renormalization group*. Perseus Books, 1992.
- [184] Ramin Golestanian. Casimir-lifshitz interaction between dielectrics of arbitrary geometry: A dielectric contrast perturbation theory. *Phys. Rev. A*, 80(1):012519, 2009.
- [185] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, Maryland, 3rd edition, 1996.
- [186] N. Graham, M. Quandt, and H. Weigel. *Spectral Methods in Quantum Field Theory*. Springer-Verlag, Berlin, 2009.
- [187] Andrew D. Greentree, John A. Vaccaro, Sebastian R. de Echaniz, Alan V. Durrant, and Jon P. Marangos. *J. Opt. B: Quantum Semiclass. Opt.*, 2:252, 2000.
- [188] Lars Grüner-Nielsen, Stig Nissen Knudsen, Bent Edvold, Torben Veng, Dorte Magnussen, C. Christian Larsen, and Hans Damsgaard. Dispersion compensating fibers. *Optical Fiber Tech.*, 6:164–180, 2000.
- [189] A. G. GRushin and A. Cortijo. Tunable Casimir repulsion with three dimensional topological insulators. *arXiv:1002.3481*, 2010.
- [190] K. Grygiel and P. Szlachetka. Chaos in second-harmonic generation of light. the case of a strain of pulses. *Opt. Comm.*, 91:241–246, 1992.
- [191] Wolfgang Hackbush and Birkhauser Verlag. *Integral Equations: Theory and Numerical Treatment*. Birkhauser Verlag, Basel, Switzerland, 1995.
- [192] J Hald. Second harmonic generation in an external ring cavity with a brewster-cut nonlinear crystal: theoretical considerations. *Optics Communications*, 197:169, 2001.

- [193] B. Ham, P. Hemmer, and M. Shahriar. *Opt. Commun.*, 144:227, 1997.
- [194] R. E. Hamam, M. Ibanescu, E. J. Reed, P. Bermel, S. G. Johnson, E. Ippen, J. D. Joannopoulos, and M. Soljacic. Purcell effect in nonlinear photonic structures: a coupled mode theory analysis. *Opt. Express*, 12(11):2102–2116, 2008.
- [195] B. W. Harris, F. Chen, and Mohideen U. Precision measurement of the Casimir force using gold surfaces. *Phys. Rev. A*, 62:052109, 2000.
- [196] Steven E. Harris and Lene Hau. *Phys. Rev. Lett.*, 82:4611, 1999.
- [197] Masashi Hase, Hiroshi Miyazaki, Mitsuru Egashira, Norio Shinya, Kenji Kojima, and Shin-ichi Uchida. Photonic band gap and anisotropic structures in transmission spectra of two-dimensional fivefold and eightfold symmetric quasiperiodic photonic crystals. *Phys. Rev. B*, 66:214205–214212, 2002.
- [198] H. Hashemi, A. W. Rodriguez, J. D. Joannopoulos, M. Soljacic, and S. G. Johnson. Nonlinear harmonic generation and devices in doubly resonant kerr cavities. *Phys. Rev. A*, 79(1):013812, 2009.
- [199] Naoki Hashizume, Makoto Ohashi, Takashi Kondo, and Ryoichi Ito. Optical harmonic generation in multilayered structures: a comprehensive analysis. *J. Opt. Soc. Am. B*, 12:1894–1904, 1995.
- [200] Lene V. Hau, Steven E. Harris, Zachary Dutton, and C.H. Behroozi. *Nature*, 397:594, 1999.
- [201] H. Haus and Y. Lai. *J. Lightwave Tech.*, 9:754, 1991.
- [202] H. A. Haus. *Waves and Fields in Optoelectronics*. Prentice-Hall, Englewood Cliffs, NJ, 1984. Ch. 7.
- [203] H. A. Haus. *Waves and Fields in Optoelectronics*. Prentice-Hall, Englewood Cliffs, NJ, 1984.
- [204] J. Hebling, A. G. Stepanov, G. Almasi, B. Bartal, and J. Kuhl. Tunable thz pulse generation by optical rectification of ultrashort laser pulses with tilted pulse fronts. *Appl. Phys. B*, 78(5):593–599, 2004.
- [205] Eugene Hecht. *Optics*. Addison-Wesley, Reading, MA, 1998.
- [206] C. Henkel and K. Joulain. Casimir force between designed materials: What is possible and what not. *Europhys. Lett.*, 72:929–935, 2005.
- [207] M. Henseler, A. Wirzba, and T. Guhr. *Ann. Phys.*, 258:286, 1997.
- [208] M. P. Hertzberg, R. L. Jaffe, M. Kardar, and A. Scardicchio. Attractive Casimir forces in a closed geometry. *Phys. Rev. Lett.*, 95:250402, 2005.
- [209] Mark Hertzberg. The Casimir piston. Notes on 2d Casimir piston.
- [210] Mark P. Hertzberg, R. L. Jaffe, Kardar M., and A. Scardicchio. Casimir forces in a piston geometry at zero and finite temperatures. *arXiv:quant-ph/0705.0139*, 2007.

- [211] K. M. Ho, C. T. Chan, C. M. Soukoulis, R. Biswas, and M. Sigalas. Photonic band gaps in three dimensions: new layer-by-layer periodic structures. *Solid State Comm.*, 89:413–416, 1994.
- [212] C.J. Hood, T.W. Lynn, A.C. Doherty, A.S. Parkins, and H.J. Kimble. *Science*, 287:1447, 2000.
- [213] J. S. Hoye, I. Brevik, J. B. Aarseth, and K. A. Milton. What is the temperature dependence of the Casimir effect? *J. Phys. A: Math. Gen.*, 39(20):6031–6038, 2006.
- [214] Wei-Ping Huang. Coupled-mode theory for optical waveguides: an overview. *J. Opt. Soc. Am. A*, 11(3):963–983, 1994.
- [215] X. Q. Huang, S. S. Juang, R. W. Peng, and A. Hu. Perfect transmission and self-similar optical transmission spectra in symmetric Fibonacci-class multilayers. *Phys. Rev. B*, 63:245104–245112, 2001.
- [216] R.G. Hunsperger. *Integrated Optics*. Springer-Verlag, Berlin, 5 edition, 2002.
- [217] V. Hushwater. Repulsive Casimir force as a result of vacuum radiation pressure. *Am. J. Phys.*, 65(5):381–384, 1996.
- [218] D. Iannuzzi, M. Lisanti, and F. Capasso. Effect of hydrogen-switchable mirrors on the Casimir force. *Proc. Nat. Ac. Sci. USA*, 101:4019–4023, 2004.
- [219] Davide Iannuzzi and Federico Capasso. Comment on “repulsive Casimir forces”. *Phys. Rev. Lett.*, 91(2):029101, 2003.
- [220] Davide Iannuzzi, Mariangela Lisanti, Jeremy Munday, and Federico Capasso. The design of long range quantum electrodynamical forces and torques between macroscopic bodies. *Solid State Comm.*, 135:618–626, 2005.
- [221] Davide Iannuzzi, Mariangela Lisanti, Jeremy N. Munday, and Federico Capasso. Quantum fluctuations in the presence of thin metallic films and anisotropic materials. *J. Phys. A: Math. Gen.*, 39:6445–6454, 2006.
- [222] M. Ibanescu, Y. Fink, S. Fan, E. L. Thomas, and J. D. Joannopoulos. An all-dielectric coaxial waveguide. *Science*, 289:415–419, 2000.
- [223] M. Ibanescu, S. G. Johnson, D. Roundy, C. Luo, Y. Fink, and J. D. Joannopoulos. Anomalous dispersion relations by symmetry breaking in axially uniform waveguides. *Phys. Rev. Lett.*, 92(6):063903, 2004.
- [224] Mihai Ibanescu, Steven G. Johnson, Marin Soljačić, J. D. Joannopoulos, Yoel Fink, Ori Weisberg, Torkel D. Engeness, Steven A. Jacobs, and M. Skorobogatiy. Analysis of mode structure in hollow dielectric waveguide fibers. *Phys. Rev. E*, 67:046608, 2003.
- [225] Vladimir S. Ilchenko, Anatoliy A. Savchenkov, Andrey B. Matsko, and Lute Maleki. Nonlinear optics and crystalline whispering gallery mode cavities. *Phys. Rev. Lett.*, 92(4):043903, 2004.
- [226] Atac Imamoglu, Holger Schmidt, G. Woods, and M. Deutsch. *Phys. Rev. Lett.*, 79:1467, 1997.

- [227] G. Imeshev, M. E. Fermann, K. L. Vodopyanov, M. M. Fejer, X. Xu, J. S. Harris, D. Bliss, and D. Lynch. High-power source of thz radiation based on orientation-patterned gaas pumped by a fiber laser. *Opt. Express*, 14(10):4439–4444, 2006.
- [228] Yoseph Imry. Casimir zero-point radiation pressure. *Phys. Rev. Lett.*, 95:080404, 2005.
- [229] F. Intravaia and A. Lambrecht. Surface plasmon modes and the Casimir energy. *Phys. Rev. Lett.*, 94:110404, 2005.
- [230] J. N. Israelachvili. *Intermolecular and Surface Forces*. Academic Press, London, 1991.
- [231] J. D. Jackson. *Classical Electrodynamics*. Wiley, New York, third edition, 1998.
- [232] M. T. Jaekel and S. Reynaud. *J. Physique I*, 1:1395, 1991.
- [233] R. L. Jaffe. Casimir effect and the quantum vacuum. *Phys. Rev. D*, 72(2):021301, 2005.
- [234] R. L. Jaffe and A. Scardicchio. Casimir effect and geometric optics. *Phys. Rev. Lett.*, 92:070402, 2004.
- [235] C. Janot. *Quasicrystals*. Clarendon Press, 1992.
- [236] Li Jia-Hua, Luo Jin-Ming, Yang Wen-Xing, and Zhan Zhi-Ming. Efficient three-wave mixing in a three-level atomic medium with an assisting microwave driven field. *Chinese Phys.*, 15:132–136, 2006.
- [237] Jianming Jin. *The Finite Element Method in Electromagnetics*. Wiley, New York, second edition, 2002.
- [238] J. D. Joannopoulos, S. Fan, A. Mekis, and S. G. Johnson. *Photonic Crystals and Light Localization in the 21st Century*, volume 563 of *Series C: Mathematical and Physical Sciences*. Kluwer, Dordrecht, The Netherlands, 2001.
- [239] J. D. Joannopoulos, R. D. Meade, and J. N. Winn. *Photonic Crystals: Molding the Flow of Light*. Princeton Univ. Press, 1995.
- [240] J.D. Joannopoulos, R.D. Meade, and J.N. Winn. *Photonic Crystals: Molding the Flow of Light*. Princeton, Princeton, NJ, 1995.
- [241] John D. Joannopoulos, Steven G. Johnson, Joshua N. Winn, and Robert D. Meade. *Photonic Crystals: Molding the Flow of Light*. Princeton University Press, second edition, February 2008.
- [242] Sajeev John. *J. Opt. A*, 3:S103, 2001.
- [243] Sajeev John and Tran Quang. *Phys. Rev. A*, 50:1764, 1994.
- [244] Sajeev John and Tran Quang. *Phys. Rev. Lett.*, 76:2484, 1996.
- [245] Steven G. Johnson, Peter Bienstman, M. Skorobogatiy, Mihai Ibanescu, Eleftherios Lidorikis, and J. D. Joannopoulos. Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals. *Phys. Rev. E*, 66:066608, 2002.

- [246] Steven G. Johnson, Shanhui Fan, Attila Mekis, and J. D. Joannopoulos. Multipole-cancellation mechanism for high- q cavities in the absence of a complete photonic band gap. *Appl. Phys. Lett.*, 78(22):3388–3390, 2001.
- [247] Steven G. Johnson, Mihai Ibanescu, M. Skorobogatiy, Ori Weisberg, Torkel D. Engeness, Marin Soljačić, Steven A. Jacobs, J. D. Joannopoulos, and Yoel Fink. Low-loss asymptotically single-mode propagation in large-core OmniGuide fibers. *Opt. Express*, 9(13):748–779, 2001.
- [248] Steven G. Johnson and J. D. Joannopoulos. Three-dimensionally periodic dielectric layered structure with omnidirectional photonic band gap. *Appl. Phys. Lett.*, 77:3490–3492, 2000.
- [249] Steven G. Johnson and J. D. Joannopoulos. Block-iterative frequency-domain methods for Maxwell’s equations in a planewave basis. *Opt. Express*, 8(3):173–190, 2001.
- [250] Steven G. Johnson and J. D. Joannopoulos. *Photonic Crystals: The Road from Theory to Practice*. Kluwer, Boston, 2002.
- [251] Steven G. Johnson and John D. Joannopoulos. *Photonic Crystals: The Road from Theory to Practice*. Kluwer Academic, Boston, 2002.
- [252] Steven G. Johnson, Attila Mekis, Shanhui Fan, and J. D. Joannopoulos. Molding the flow of light. *Computing Sci. Eng.*, 3(6):38–47, 2001.
- [253] Steven G. Johnson, M. L. Povinelli, M. Soljacic, A. Karalis, S. Jacobs, and J. D. Joannopoulos. Roughness losses and volume-current methods in photonic-crystal waveguides. *Appl. Phys. B: Lasers Opt.*, 81(2):283–293, 2005.
- [254] M. A. Kaliteevski, S. Brand, R. A. Abram, T. F. Krauss, R. M. De La Rue, and P. Millar. Two-dimensional penrose-tiled photonic quasicrystals: Diffraction of light and fractal density of modes. *J. Mod. Opt.*, 47:1771–1778, 2000.
- [255] M. A. Kaliteevski, S. Brand, R. A. Abram, T. F. Krauss, P. Millar, and R. De La Rue. Diffraction and transmission of light in low-refractive index Penrose-tiled photonic quasicrystals. *J. Phys: Cond. Mat.*, 13:10459–10470, 2001.
- [256] M. A. Kaliteevski, J. Manzanares Martines, D. Cassagne, and J. P. Albert. *Phys. Status Solidi A*, 195:612, 2003.
- [257] Mehran Kardar. *Statistical Physics of Particles*. Cambridge University Press, Cambridge, UK, 2007.
- [258] Walter Karplus. *Analog simulation: solution of field problems*. McGraw-Hill, 1958.
- [259] O. Kenneth, I. Klich, A. Mann, and M. Revzen. Repulsive Casimir forces. *Phys. Rev. Lett.*, 89(3):033001, 2002.
- [260] Oded Kenneth and Israel Klich. Opposites attract: A theorem about the Casimir force. *Phys. Rev. Lett.*, 97:160401, 2006.
- [261] Oded Kenneth and Israel Klich. Casimir forces in a T-operator approach. *Phys. Rev. B*, 78:014103, 2008.

- [262] Sun-Kyung Kim, Jee-Hye Lee, Se-Heon Kim, In-Kang Hwang, and Yong-Hee Lee. Photonic quasicrystal single-cell cavity mode. *Appl. Phys. Lett.*, 86:031101, 2005.
- [263] H. J. Kimble, M. Dagenais, and L. Mandel. Photon antibunching in resonance fluorescence. *Phys. Rev. Lett.*, 39(11):691–695, 1977.
- [264] T. J. Kippenberg, S. M. Spillane, and K. J. Vahala. Kerr-nonlinearity optical parametric oscillation in an ultrahigh- q toroid microcavity. *Phys. Rev. Lett.*, 93(8):083904, 2004.
- [265] G. Kirchoff. On the motion of electricity in wires. *Philos. Mag.*, 13:393–412, 1857.
- [266] Lawrence A. Klein and Calvin T. Swift. An improved model for the dielectric constant of sea water at microwave frequencies. *IEEE Trans. Ant. and Prop.*, 25(1):104–111, 1977.
- [267] D. Kleppner. *Phys. Rev. Lett.*, 47:233, 1981.
- [268] G. L. Klimchitskaya, U. Mohideen, and V. M. Mostapanenko. The Casimir force between real materials: experiment and theory. *Rev. Mod. Phys.*, 81(4):1827–1885, 2009.
- [269] G. L. Klimchitskaya, S. I. Zanette, and A. O. Caride. Lateral projection as a possible explanation of the nontrivial boundary dependence of the Casimir force. *Phys. Rev. A*, 63:014101, 2001.
- [270] J. C. Knight, J. Broeng, T. A. Birks, and P. St.-J. Russell. Photonic band gap guidance in optical fibers. *Science*, 282:1476–1478, 1998.
- [271] Karl Koch and Gerald T. Moore. Singly resonant cavity-enhanced frequency tripling. *J. Opt. Soc. Am. B*, 16:448, 1999.
- [272] R. Kohler, A. Tredicucci, H. E. Beltram, F. Beere, E. H. Linfield, A. G. Davies, D. A. Ritchie, R. C. Iotti, and F. Rossi. Terahertz semiconductor-heterostructure laser. *Nature*, 417:156–159, 2002.
- [273] Mahito Kohmoto, Leo P. Kadanoff, and Chao Tang. Localization problem in one dimension: Mapping and escape. *Phys. Rev. Lett.*, 50(23):1870–1872, 1983.
- [274] Mahito Kohmoto, Bill Sutherland, and K. Iguchi. Localization in optics: Quasiperiodic media. *Phys. Rev. Lett.*, 58(23):2436–2438, 1987.
- [275] W. J. Kozlovsky, W. P. Risk, W. Lenth, B. G. Kim, G. L. Bona, H. Jaeckel, and D. J. Webb. Blue light generation by resonator-enhanced frequency doubling of an extended-cavity diode laser. *Appl. Phys. Lett.*, 65:525–527, 1994.
- [276] S. Kraiem, F. Hassen, H. Maaref, X. Marie, and E. Vaneelle. *Opt. Mat.*, 17:305, 2001.
- [277] U. Kuhl and H. J. Stockmann. Microwave realization of the Hofstadter butterfly. *Phys. Rev. Lett.*, 80(15):3232, 1997.
- [278] Karl S. Kunz and Raymond J. Luebbers. *The Finite-Difference Time-Domain Method for Electromagnetics*. CRC Press, Boca Raton, 1993.

- [279] A. Lambrecht, I. Pirozhenko, L. Duraffourg, and Ph. Andreucci. The casimir effect for silicon and gold slabs. *EPL*, 77(4):44006, 2007.
- [280] Astrid Lambrecht, Paulo A. Maia Neto, and Serge Reynaud. The Casimir effect within scattering theory. *New Journal of Physics*, 8(243), 2006.
- [281] Astrid Lambrecht, Paulo A. Maia Neto, and Serge Reynaud. Casimir energy and geometry: beyond the proximity force approximation. *J. Phys. A: Math. Thero.*, 41:164004, 2008.
- [282] P. Lambropoulos, G. M. Nikolopoulos, T. R. Nielsen, and S. Bay. *Rep. Prog. Phys.*, 63:455, 2000.
- [283] S. K. Lamoreaux. Demonstration of the Casimir force in the 0.6 to 6 μm range. *Phys. Rev. Lett.*, 78:5–8, 1997.
- [284] S. K. Lamoreaux. Demonstration of the Casimir force in the 0.6 to 6 μm range. *Phys. Rev. Lett.*, 78:5, 1997.
- [285] L. D. Landau and E. M. Lifshitz. *Electrodynamics of Continuous Media*. Pergamon Press, Oxford, 1960.
- [286] L. D. Landau and E. M. Lifshitz. *Quantum Mechanics*. Butterworth-Heinemann, Oxford, 3rd edition, 1977.
- [287] L. D. Landau, E. M. Lifshitz, and L. P. Pitaevskiĭ. *Statistical Physics Part 2*, volume 9. Pergamon, Oxford, 1960.
- [288] L. D. Landay, E. M. Lifshitz, and L. P. Pitaevskiĭ. *Electrodynamics of Continuous Media*, volume 8. Pergamon Press, Oxford, 1960.
- [289] Wah Tung Lau and Shanhui Fan. Creating large bandwidth line defects by embedding dielectric waveguides into photonic crystal slabs. *Appl. Phys. Lett.*, 81:3915–3917, 2002.
- [290] A. N. Lazarchik. Bragg fiber lightguides. *Radiotekhnika i elektronika*, 1:36–43, 1988.
- [291] Alexandra Ledermann, Ludovico Cademartiri, Martin Hermatschweiler, Costanza Toninelli, Geoffrey A. Ozin, Diederik S. Weirsma, Martin Wegener, and Georg Von Freymann. Three-dimensional silicon inverse photonic quasicrystals for infrared wavelengths. *Nature Mat.*, 5:942–945, December 2006.
- [292] S. Lee and W. M Sigmund. Repulsive van der waals forces for silica and alumina. *J. Colloid Interface Sci.*, 243:365–369, 2001.
- [293] S. Lee and W. M Sigmund. AFM study of repulsive van der waals forces between teflon AF thin film and silica or alumina. *J. Colloids Surf. A*, 204:43–50, 2002.
- [294] Seung-woo Lee and Wolfgang M. Sigmund. Repulsive van der Waals forces for Silica and Alumina. *J. Colloid Interface Science*, 243(2):365–369, 2001.
- [295] Y. S Lee, T. Meade, V. Perlin, H. Winful, T. B. Norris, and A. Galvanauskas. Generation of narrow-band terahertz radiation via optical rectification of femtosecond pulses in periodically poled lithium niobate. *Appl. Phys. Lett.*, 76:2505–2507, 2000.

- [296] G. Lenz, J. Zimmermann, T. Katsufuji, M. E. Lines, H. Y. Hwang, S. Splter, R. E. Slusher, S.-W. Cheong, J. S. Sanghera, and I. D. Aggarwal. *Opt. Lett.*, 25:254, 2000.
- [297] Ulf Leonhardt and Thomas G. Philbin. Quantum levitation by left-handed metamaterials. *New J. of Phys.*, 9(254):1–11, 2007.
- [298] Michael Levin, Alexander P. McCauley, Alejandro W. Rodriguez, Homer H. T. Reid, and Steven G. Johnson. **Casimir repulsion between metallic objects in vacuum.** *arXiv:1003.3487*, 2010.
- [299] Hao Li and Mehran Kardar. Fluctuation-induced forces between manifolds immersed in correlated fluids. *Phys. Rev. A*, 46(10):6490–6500, 1992.
- [300] Hongqiang Li, Bingying Cheng, and Daozhong Zhang. Two-dimensional disordered photonic crystals with an average periodic lattice. *Phys. Rev. B*, 56(17):10734–10736, 1997.
- [301] Jing Li, Degang Zhao, and Zhengyou Liu. Zero- \bar{n} photonic band gap in a quasiperiodic stacking of positive and negative refractive index materials. *Phys. Lett. A*, 332:461–468, 2004.
- [302] E. M. Lifshitz. *Dokl. Akad. Nauk. SSSR*, 100:879, 1955.
- [303] E. M. Lifshitz. *Sov. Phys. JETP*, 2:73, 1956.
- [304] E. M. Lifshitz and Lev P. Pitaevskii. *Statistical Physics: Part 2*. Pergamon, Oxford, 1980.
- [305] R. Lifshitz. The square fibonacci tiling. *J. of Alloys and Compounds*, 342:186–190, 2000.
- [306] Ron Lifshitz, Andy Arie, and Alon Bahabad. Photonic quasicrystals for nonlinear optical frequency conversion. *Phys. Rev. Lett.*, 95:133901, 2005.
- [307] M. Lisanti, D. Iannuzzi, and F. Capasso. Observation of the skin-depth effect on the Casimir force between metallic surfaces. *Proc. Nat. Ac. Sci. USA*, 102:11989–11992, 2005.
- [308] Marco Liscidini and Lucio .A Andreani. Highly efficient second-harmonic generation in doubly resonant planar microcavities. *Appl. Phys. Lett.*, 85:1883, 2004.
- [309] Marco Liscidini and Lucio .A Andreani. Second-harmonic generation in doubly resonant microcavities with periodic dielectric mirrors. *Phys. Rev. E*, 73:016613, 2006.
- [310] Natasha M. Litchinitser, A.K. Abeeluck, C. Headley, and B.J. Eggleton. *Opt. Lett.*, 27:1592, 2002.
- [311] Mengkun Liu, Nian Ji, Zhifang Lin, and Chuim S. T. Radiation torque on a birefringent sphere caused by an electromagnetic wave. *Phys. Rev. E*, 72:056610, 2005.
- [312] Tzu-Ming Liu, Cheng-Ta Yu, and Chi-Kuang Sun. 2 ghz repetition-rate femtosecond blue sources by second-harmonic generation in a resonantly enhanced cavity. *Appl. Phys. Lett.*, 86:061112, 2005.

- [313] Seth Lloyd. Quantum information matters. *Science*, 319(5867):1209–1211, 2008.
- [314] H.-K. Lo and H.F. Chau. *Science*, 283:2050, 1999.
- [315] Peter Lodahl, A. Floris van Driel, Ivan S. Nikolaev, Arie Irman, Karin Overgaag, Danil Vanmaekelbergh, and Willem L. Vos. *Nature*, 430:654, 2004.
- [316] Marco Lončar, Tomoyuki Yoshie, Axel Scherer, Pawan Gogna, and Yueming Qiu. Low-threshold photonic crystal laser. *Appl. Phys. Lett.*, 81(15):2680–2682, 2002.
- [317] Marko Lončar, Dušan Nedeljković, Theodor Doll, Jelena Vučković, Axel Scherer, and Thomas P. Pearsall. Waveguiding in planar photonic crystals. *Appl. Phys. Lett.*, 77(13):1937–1939, 2000.
- [318] F. London. The general theory of molecular forces. *Trans. Faraday Soc.*, 33:8–26, 1937.
- [319] Jian Ping Lu and Joseph L. Birman. Electronic structure of a quasiperiodic system. *Phys. Rev. B*, 36(8):4471–4474, 1987.
- [320] Wentao Lu, M. Rose, K. Pance, and S. Sridhar. Quantum resonances and decay of a chaotic fractal repeller observed using microwaves. *Phys. Rev. Lett.*, 82(26):5233, 1999.
- [321] M. D. Lukin and A. Imamoglu. Controlling photons using electromagnetically-induced transparency. *Nature*, 413:273–276, 2001.
- [322] Chiyuan Luo, A. Narayanaswamy, G. Ghen, and J. D. Joannopoulos. Thermal radiation from photonic crystals: A direct calculation. *Phys. Rev. Lett.*, 93(21):213905, 2004.
- [323] H. Mabuchi, M. Armen, B. Lev, M. Loncar, J. Vuckovic, H.J. Kimble, J. Preskill, M. Roukes, and A. Scherer. *Quantum Information and Computation*, 1:7, 2001.
- [324] Bjorn Maes, Peter Bienstman, and Roel Baets. Modeling second-harmonic generation by use of mode expansion. *J. Opt. Soc. Am. B*, 22:1378, 2005.
- [325] J. Mahanty and B. W. Ninham. *Dispersion forces*. Academic London, 1976.
- [326] J. Mahanty and B. W. Ninham. *Dispersion Forces*. Academic Press, London, 1976.
- [327] P. A. Maia Neto, A. Lambrecht, and S. Reynaud. Roughness correction to the Casimir force: Beyond the proximity force approximation. *Europhys. Lett.*, 69:924, 2005.
- [328] Martin Maldovan and Edwin L. Thomas. Diamond-structured photonic crystals. *Nature Materials*, 3:593–600, 2004.
- [329] A. M. Malvezzi, G. Vecchi, M. Patrini, G. Guizzetti, L. C. Andreani, F. Romanato, L. Businaro, E. Di Fabrizio, A. Passaseo, and M. De Vittorio. Resonant second-harmonic generation in a gas photonic crystal waveguide. *Phys. Rev. B*, 68:161306, 2003.
- [330] Weining Man, Mischa Megens, Paul J. Steinhardt, and P. M. Chaikin. Experimental measurement of the photonic crystal properties of icosahedral quasicrystals. *Nature*, 436:993–996, August 2005.

- [331] V. A. Mandelshtam and H. S. Taylor. Harmonic inversion of time signals and its applications. *J. Chem. Phys.*, 107(17):6756–6769, 1997. See erratum [332].
- [332] V. A. Mandelshtam and H. S. Taylor. Erratum: “Harmonic inversion of time signals and its applications”. *J. Chem. Phys.*, 109:4128, 1998.
- [333] Vladimir A. Mandelshtam and Arnold Neumaier. Further generalization and numerical implementation of pseudo-time Schrödinger equations for quantum scattering calculations. *J. Theoretical Comput. Chem.*, 1(1):1–15, 2002.
- [334] Valery N. Marachevsky. Casimir interaction of two plates inside a cylinder. *Phys. Rev. D*, 75:085019, 2007.
- [335] J. Marcou, F. Brechet, and Ph. Roy. Design of weakly guiding bragg fibres for chromatic dispersion shifting towards short wavelengths. *J. Opt. A: Pure Appl. Opt.*, 3:S144–S153, 2001.
- [336] D. Marcuse. *Theory of Dielectric Optical Waveguides*. Academic Press, San Diego, second edition, 1991.
- [337] M. G. Martemyanov, E. M. Kim, T. V. Dolgova, A. A. Fedyanin, O. A. Aktsipetrov, and G. Marowsky. Third harmonic generation in silicon photonic crystals and microcavities. *Phys. Rev. B*, 70:073311, 2004.
- [338] R. J. Martin, N. A. Masnari, and J. E. Rowe. Analog representation of Poisson’s equation in two dimensions. *IRE Trans. Electronic Comp.*, EC-9(4):490–496, 1960.
- [339] R. Matloob, A. Keshavarz, and D. Sedighi. Casimir effect for two lossy dispersive dielectric slabs. *Phys. Rev. A*, 60(5):3410–3420, 1999.
- [340] Reza Matloob. Casimir effect between two conducting plates. *Phys. Rev. A*, 60(5):3421–3428, 1999.
- [341] A. B. Matsko, D. V. Strekalov, and N. Yu. Sensitivity of terahertz photonic receivers. *Phys. Rev. A*, 77(4):043812, 2008.
- [342] F. D. Mazitelli, D. A. Dalvit, and F. C. Lobardo. Exact zero-point interaction energy between cylinders. *New Journal of Physics*, 8(240):1–21, 2006.
- [343] A. P. McCauley, M. T. H. Reid, P. Bermel, J. D. Joannopoulos, Steven G. Johnson, et al. Exact results for casimir forces in chiral metamaterials. *In preparation.*, 2010.
- [344] Alexander P. McCauley, Alejandro W. Rodriguez, J. D. Joannopoulos, and Steven G. Johnson. Casimir forces in the time domain: Applications. *Phys. Rev. A*, 81:012119, 2010.
- [345] Gail McConnell, Allister I. Ferguson, and Nigel Langford. Cavity-augmented frequency tripling of a continuous wave mode-locked laser. *J. Phys. D: Appl. Phys.*, 34:2408, 2001.
- [346] M. W. McCutcheon, D. E. Chang, Y. Zhang, M. D. Lukin, and M. Loncar. Broadband spectral control of single photon sources using a nonlinear photonic crystal cavity. *arXiv:0903.4706*, 2010.

- [347] K. S. Mendelson. The story of c . *Am. J. Phys.*, 74:995–997, 2006.
- [348] A. Messiah. *Quantum Mechanics: Vol. II*. Wiley, New York, 1976. Ch. 17.
- [349] A. Meurk, P. F. Luckham, and L. Bergstrom. Direct measurement of repulsive and attractive van der Waals forces between inorganic materials. *Langmuir*, 13:3896–3899, 1997.
- [350] D.A.B. Miller, S.D. Smith, and B.S. Wherrett. *Opt. Comm.*, 35:221, 1980.
- [351] A. Milling, Paul Mulvaney, and Ian Larson. Direct measurement of repulsive van der Waals interactions using an atomic force microscope. *J. Colloid and Interface Science*, 180(2):460–465, 1996.
- [352] P. W. Milonni. *The Quantum Vacuum: An Introduction to Quantum Electrodynamics*. Academic Press, San Diego, 1993.
- [353] P. W. Milonni, D. F. V. James, and H. Fearn. Comment on "superluminality, par-electricity, and Earnshaw's theorem in media with inverted populations". *Phys. Rev. Lett.*, 75(17):3194, 1995.
- [354] K. Milton, J. Wagner, P. Parashar, and I. Brevi. Casimir energy, dispersion and the Lifshitz formula. *arXiv:1001.4163*, 2010.
- [355] K. A. Milton. *The Casimir Effect: Physical Manifestations of Zero-Point Energy*. Singapore: World Scientific, 2001.
- [356] K. A. Milton, L. L. DeRaad Jr., and J. Schwinger. Casimir self-stress on a perfectly conducting spherical shell. *Ann. Phys.*, 115:388, 1978.
- [357] Kimball A. Milton. Zero-point energy in bag models. *Phys. Rev. D*, 22(6):1441–1443, 1980.
- [358] Kimball A. Milton. The Casimir effect: recent controversies and progress. *Journal of Physics A: Mathematical and General*, 37:R209–R277, 2004.
- [359] Kimball A. Milton, P. Parashar, J. Wagner, and C. Pelaez. Multiple scattering Casimir force calculations: layered and corrugated materials, wedges, and Casimir-Polder forces. *J. Vac. Sci. Tech. B*, 28(3):C4A8–C4A16, 2010.
- [360] Kimball A. Milton, Prachi Parashar, and Jef Wagner. Exact results for Casimir interactions between dielectric bodies: the weak-coupling or Van der Waals limit. *Phys. Rev. Lett.*, 101:160402, 2008.
- [361] S. F. Mingaleev and Y. S. Kivshar. Nonlinear transmission and light localization in photonic crystal waveguides. *J. Opt. Soc. Am. B*, 19:2241–2249, 2002.
- [362] MirFaez Miri and Ramin Golestanian. A frustrated nanomechanical device powered by the lateral Casimir force. *APL*, 92:113103, 2008.
- [363] Khaled Mnaymneh and Robert C. Gauthier. Mode localization and band-gap formation in defect-free photonic quasicrystals. *Opt. Express*, 15(8):5089–5099, 2007.

- [364] W. L. Mochan and C. Villarreal. Casimir effect for arbitrary materials: contributions within and beyond the light cone. *New Journal of Physics*, 8(242):1–21, 2006.
- [365] U. Mohideen and A. Roy. Precision measurement of the Casimir force from 0.1 to 0.9 μm . *Phys. Rev. Lett.*, 81:4549, 1998.
- [366] Gerald T. Moore, Karl Koch, and E. C. Cheung. Optical parametric oscillation with intracavity second-harmonic generation. *Optics Communications*, 113:463, 1995.
- [367] Y. A. Morozov, I. S. Nefedov, V. Y. Aleshkin, and I. V. Krasnikova. Terahertz oscillator based on nonlinear frequency conversion in a double vertical cavity. *Semiconductors*, 39(1):113, 2005.
- [368] Ralph Morrison. *Grounding and Shielding*. John Wiley and Sons, 2007.
- [369] V. M. Mostepanenko, R. S. Decca, E. Fischbach, G. L. Klimchitskaya, D. E. Krause, and D. Lopez. Stronger constraints on non-Newtonian gravity from the Casimir effect. *J. Phys. A: Math. Theor.*, 41:164054, 2008.
- [370] V. M. Mostepanenko and N. N Trunov. *The Casimir Effect and its Applications*. Clarendon Press, Oxford, 1997.
- [371] Xiaodong Mu, Yujie J. Ding, Haeyeon Yang, and Gregory J. Salamo. Cavity-enhanced and quasiphase-matched mutli-order reflection-second-harmonic generation from gaas/alas and gaas/algaas multilayers. *Appl. Phys. Lett.*, 79:569, 2001.
- [372] J. N. Munday and Federico Capasso. Measurement of the Casimir-Lifshitz force in fluids: the effect of electrostatic forces and Debye screening. *Phys. Rev. A*, 78:032109, 2008.
- [373] J. N. Munday, D. Iannuzzi, Y. Barash, and F. Capasso. Torque induced on birefringent plates by quantum fluctuations. *Phys. Rev. A*, 71:042102, 2005.
- [374] Jeremy Munday, Federico Capasso, and V. Adrian Parsegia. Measured long-range repulsive casimir-lifshitz forces. *Nature*, 457:170–173, 2009.
- [375] Jeremy N. Munday and Federico Capasso. Precision measurement of the Casimir-Lifshitz force in a fluid. *Phys. Rev. A*, 75:060102(R), 2007.
- [376] Arvind Narayanaswamy and Gang Chen. Thermal near-field radiative transfer between two spheres. *Phys. Rev. B*, 77(7):075125, 2008.
- [377] Kae Nemoto and W. J. Munro. Nearly deterministic linear optical controlled-NOT gate. *Phys. Rev. Lett.*, 93:250502, 2004.
- [378] M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, England, 2000.
- [379] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [380] W. D. Niven, editor. *The Scientific Papers of James Clerk Maxwell*, volume 1. Dover Publications, New York, 1952.

- [381] S. Noda, A. Chutinan, and M. Imada. Trapping and emission of photons by a single defect in a photonic bandgap structure. *Nature*, 407:608–610, 2000.
- [382] S. Noda, M. Fujita, and T. Asano. Spontaneous-emission control by photonic crystals and nanocavities. *Nature Photonics*, 1:449–458, 2007.
- [383] Andrew H. Norton and C. Martijn de Sterke. Optimal poling of nonlinear photonic crystals for frequency conversion. *Opt. Lett.*, 28:188, 2002.
- [384] Vance D. Norum, Marvin Adelberg, and Robert L. Farrenkopf. *Analog simulation of particle trajectories in fluid flow*. ACM, New York, NY, USA, 1962.
- [385] M. Notomi, H. Suzuki, T. Tamamura, and K. Edagawa. Lasing action due to the two-dimensional quasiperiodicity of photonic quasicrystals with Penrose lattice. *Phys. Rev. Lett.*, 92(12):123906, 2004.
- [386] Masaya Notomi, Akihiko Shinya, Satoshi Mitsugi, Goh Kira, Eiichi Kuramochi, and Takasumi Tanabe. *Opt. Express*, 13:2678, 2005.
- [387] Masaya Notomi, Akihiko Shinya, Satoshi Mitsugi, Goh Kira, Eiichi Kuramochi, and Takasumi Tanabe. Optical bistable switching action of Si high- q photonic-crystal nanocavities. *Opt. Express*, 13(7):2678–2687, 2005.
- [388] J. M. Obrecht, R. J. Wild, M. Antezza, L. P. Pitaevskii, S. Stringari, and E. A. Cornell. Measurement of the temperature dependence of the Casimir-Polder force. *Phys. Rev. Lett.*, 98(6):063201, 2007.
- [389] Ronald J. Onega. Analog computer solution of the electrodiffusion equation for a simple membrane. *Am. J. Phys.*, 40:390–394, 1971.
- [390] Roberto Onofrio. Casimir forces and non-Newtonian gravitation. *New J. Phys.*, 8:237, 2006.
- [391] Stellan Ostlund, Rahul Pandit, David Rand, Hans Scheller, and Eric Siggia. One-dimensional Schrödinger equation with an almost periodic potential. *Phys. Rev. Lett.*, 50(23):1873–1876, 1983.
- [392] Edward Ott. *Chaos in Dynamical Systems*. Cambridge University Press, 2nd edition, 2002.
- [393] Z. Y. Ou and H. J. Kimble. Enhanced conversion efficiency for harmonic generation with double resonance. *Opt. Lett.*, 18:1053–1055, 1993.
- [394] R. Ozeri, N. Katz, J. Steinhauer, E. Rowen, and N. Davidson. Three-wave mixing of bogoliubov quasiparticles in a bose-einstein condensate. *Phys. Rev. Lett.*, 90(17):170401, 2003.
- [395] O. Painter, J. Vučković, and A. Scherer. Defect modes of a two-dimensional photonic crystal in an optically thin dielectric slab. *J. Opt. Soc. Am. B*, 16(2):275–285, 1999.
- [396] G. Palasantzas. Self-affine roughness influence on the Casimir effect. *J. Appl. Phys.*, 97:126104, 2005.

- [397] Nicolae Panouiu, Mayank Bahl, and Jr. Richard Osgood. All-optical tunability of a nonlinear photonic crystal channel drop filter. *Opt. Express*, 12(8):1605–1610, 2004.
- [398] Alberto Parini, Gaetano Bellanca, Stefano Trillo, Matteo Conforti, Andrea Locatelli, and Costantino De Angelis. Self-pulsing and bistability in nonlinear Bragg gratings. *J. Opt. Soc. Am. B*, 24:2229–2237, 2007.
- [399] G. J. Parker, M. D. B. Charlton, M. E. Zoorob, J. J. Baumberg, M. C. Netti, and T. Lee. Highly engineered mesoporous structures for optical processing. *Phil. Trans. Soc. A*, 364:189–199, 2006.
- [400] Adrian V. Parsegian. *Van der Waals Forces: A Handbook for Biologists, Chemists, Engineers, and Physicists*. Cambridge University Press, NY, 2006.
- [401] V. A. Parsegian and George H. Weiss. Dielectric anisotropy and the van der waals interaction between bulk media. *J. Adhesion*, 3(4):259–267, 1972.
- [402] R. Paschotta, K. Fiedler, P. Kurz, and J. Mlynek. Nonlinear mode coupling in doubly resonant frequency doublers. *Appl. Phys. Lett.*, 58:117, 1994.
- [403] S. Pasquali and A. C. Maggs. Fluctuation-induced interactions between dielectrics in general geometries. *J. Chem. Phys.*, 129:014703, 2008.
- [404] S. Pasquali and A. C. Maggs. Numerical studies of Lifshitz interactions between dielectrics. *Phys. Rev. A (R)*, 79:020102(R), 2009.
- [405] S. Pearl, H. Lotem, and Y. Shimony. Optimization of laser intracavity second-harmonic generation by a linear dispersion element. *J. Opt. Soc. Am. B*, 16:1705, 1999.
- [406] M. A. Persaud, J. M. Tolchard, and A. I. Ferguson. Efficient generation of picosecond pulses at 243 nm. *IEEE J. Quantum Electron.*, 26:1253–1258, 1990.
- [407] R. Piessens, E. de Doncker-Kapenga, C. Uberhuber, and D. Hahaner. *QUADPACK: A Subroutine Package for Automatic Integration*. Springer-Verlag, Berlin, 1983.
- [408] L. P. Pitaevskii. Comment on "Casimir force acting on magnetodielectric bodies in embedded in media". *Phys. Rev. A*, 73:047801, 2006.
- [409] G. Plunien, B. Muller, and W. Greiner. The Casimir effect. *Phys. Rep.*, 134(87):87–193, 1986.
- [410] Mircea Popescu. Prediction of the electromagnetic torque in synchronous machines through Maxwell stress harmonic filter (HFT) method. *Electrical Engineering*, 89:117–125, 2006.
- [411] M. L. Povinelli, S. G. Johnson, and J. D. Joannopoulos. Slow-light, band-edge waveguides for tunable time delays. *Opt. Express*, 13:7145–7159, 2005.
- [412] M. L. Povinelli, Steven G. Johnson, Shanhui Fan, and J. D. Joannopoulos. Emulation of two-dimensional photonic crystal defect modes in a photonic crystal with a three-dimensional photonic band gap. *Phys. Rev. B*, 64:075313, 2001.
- [413] David M. Pozar. *Microwave Engineering 3rd Edition*. John Wiley and Sons, 2004.

- [414] John Preskill. Quantum computation lecture notes. <http://www.theory.caltech.edu/people/preskill/ph229/>, 2004.
- [415] E. M. Purcell. Spontaneous emission probabilities at radio frequencies. *Phys. Rev.*, 69:681–686, 1946.
- [416] E.M. Purcell. *Phys. Rev.*, 69:681, 1946.
- [417] David M. Pustai, Ahmed Sharkawy, Shouyuan Shi, and Dennis W. Prather. Tunable photonic crystal microcavities. *Appl. Opt.*, 41(26):5574–5579, 2002.
- [418] X. C. Zhang Q. Wu, M. Litz. Broadband detection capability of znTe electro-optic field detectors. *Appl. Phys. Lett.*, 68:2924–2926, 1996.
- [419] Minghao Qi, Eleftherios Lidorikis, Peter T. Rakich, Steven G. Johnson, J. D. Joannopoulos, Erich P. Ippen, and Henry I. Smith. A three-dimensional optical photonic crystal with designed point defects. *Nature*, 429:538–542, 2004.
- [420] Minghao Qi and H. Smith. Fabrication of 3d layered photonic crystal. private communications, 2003.
- [421] M. Quilichini. Phonon excitations in quasicrystals. *Reviews of Modern Physics*, 69(1):277–314, 1997.
- [422] Christian Raabe and Dirk-Gunnar Welsch. Casimir force acting on magnetodielectric bodies embedded in media. *Phys. Rev. A*, 71(1):013814, 2005.
- [423] Christian Raabe and Dirk-Gunnar Welsch. QED in arbitrary linear media: amplifying media. *arXiv:0710.2867*, 2007.
- [424] Christian Raabe and Dirk-Gunnar Welsch. Reply to “comment on ‘casimir force acting on magnetodielectric bodies embedded in media’ ”. *Phys. Rev. A*, 80(6):067801, 2009.
- [425] S. J. Rahi, Thornsten Emig, Robert L. Jaffe, and Mehran Kardar. Casimir forces between cylinders and plates. *Phys. Rev. A*, 78:012104, 2008.
- [426] S. J. Rahi, Thornsten Emig, Robert L. Jaffe, and Mehran Kardar. Scattering theory approach to electrodynamic Casimir forces. *Phys. Rev. D*, 80:085021, 2009.
- [427] S. J. Rahi, Mehran Kardar, and Thornsten Emig. Constraints on stable equilibria with fluctuation-induced forces. *arXiv:quant-ph/0911.5364v1*, 2009.
- [428] S. J. Rahi, Alejandro W. Rodriguez, Thorsten Emig, Robert L. Jaffe, Steven G. Johnson, and Mehran Kardar. Nonmonotonic effects of parallel sidewalls on casimir forces between cylinders. *Phys. Rev. A*, 77(3):030101(R), 2008.
- [429] Sahand Jamal Rahi and Saad Zaheer. Stable levitation and alignment of compact objects by casimir spring forces. *Phys. Rev. Lett.*, 104(7):070405, 2010.
- [430] S. M. Rao and N. Balakrishnan. Computational electromagnetics. *Current Science*, 77(10):1343–1347, 1999.
- [431] H. Razmi and S. M. Modarresi. Casimir torque for a perfectly conducting wedge: a canonical quantum field theoretical approach. *Int. J. Theor. Phys.*, 44(2):229–234, 2005.

- [432] S. Rebic, S.M. Tan, A.S. Parkins, and D.F. Walls. *J. Opt. B: Quantum Semiclass. Opt.*, 1:490, 1999.
- [433] Evan J. Reed, Michael R. Armstrong, Ki-Yong Kim, and James H. Glownia. Atomic-scale time and space resolution of terahertz frequency acoustic waves. *Phys. Rev. Lett.*, 101:014302, 2008.
- [434] M. T. Homer Reid, Alejandro W. Rodriguez, Jacob White, and Steven G. Johnson. Efficient computation of three-dimensional Casimir forces. *Phys. Rev. Lett.*, 103(4):040401, 2009.
- [435] Frederick Reif. *Fundamentals of statistical and thermal physics*. McGraw-Hill Series in Fundamentals of Physics, 1965.
- [436] F. F. Ren, R. Li, C. Cheng, Y. X. Fan, J. P. Ding, and H. T. Wang. Low-threshold and high-efficiency optical parametric oscillator using a one-dimensional single-defect photonic crystal with quadratic nonlinearity. *Phys. Rev. B*, 73:033104, 2006.
- [437] Fang-Fang. Ren, Rui Li, Chen Cheng, and Hui-Tian Wang. Giant enhancement of second harmonic generation in a finite photonic crystal with a single defect and dual-localized modes. *Phys. Rev. B*, 70:245109, 2004.
- [438] S. Reynaud, P. A. Maia Neto, and A. Lambrecht. Casimir energy and geometry: beyond the proximity force approximation. *J. Phys. A: Math. Theor.*, 41:164004, 2008.
- [439] Hannes Risken. *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer-Verlag, Heidelberg, New York, 1996.
- [440] R. B. Rodrigues, P. A. Maia Neto, A. Lambrecht, and S. Reynaud. Vacuum-induced torque between corrugated metallic plates. *Europhys. Lett.*, 75(5):822–828, 2006.
- [441] Robson B. Rodrigues, Paulo A. Maia Neto, Astrid Lambrecht, and Serge Reynaud. Lateral Casimir force beyond the proximity-force approximation. *Phys. Rev. Lett.*, 96:100402, 2006.
- [442] Alejandro Rodriguez, M. Ibanescu, David Iannuzzi, Federico Capasso, J. D. Joannopoulos, and Steven G. Johnson. Computation and visualization of Casimir forces in arbitrary geometries: Non-monotonic lateral-wall forces and failure of proximity force approximations. *Phys. Rev. Lett.*, 99(8):080401, 2007.
- [443] Alejandro Rodriguez, M. Ibanescu, David Iannuzzi, J. D. Joannopoulos, and Steven G. Johnson. Virtual photons in imaginary time: computing Casimir forces in arbitrary geometries via standard numerical electromagnetism. *Phys. Rev. A*, 76(3):032106, 2007.
- [444] Alejandro Rodriguez, Mihai Ibanescu, D. J. Joannopoulos, and G. Steven Johnson. Disorder-immune confinement of light in photonic-crystal cavities. *Opt. Lett.*, 30(23):3192, 2005.
- [445] Alejandro Rodriguez, Marin Soljačić, J. D. Joannopoulos, and Steven G. Johnson. $\chi^{(2)}$ and $\chi^{(3)}$ harmonic generation at a critical power in inhomogeneous doubly resonant cavities. *Opt. Express*, 15(12):7303–7318, 2007.

- [446] Alejandro W. Rodriguez, J. D. Joannopoulos, and Steven G. Johnson. Repulsive and attractive Casimir forces in a glide-symmetric geometry. *Phys. Rev. A*, 77(6):062107, 2008.
- [447] Alejandro W. Rodriguez, Alexander P. McCauley, Yehuda Avniel, and Steven G. Johnson. Computation and visualization of photonic quasicrystal spectra via Bloch's theorem. *Phys. Rev. B*, 77(10):104201, 2008.
- [448] Alejandro W. Rodriguez, Alexander P. McCauley, J. D. Joannopoulos, and Steven G. Johnson. Theoretical ingredients of a Casimir analog computer. *Submitted for publication*, 2008.
- [449] Alejandro W. Rodriguez, Alexander P. McCauley, J. D. Joannopoulos, and Steven G. Johnson. Casimir forces in the time domain: Theory. *Phys. Rev. A*, 80(1):012115, 2009.
- [450] Alejandro W. Rodriguez, Alexander P. McCauley, David Woolf, Federico Capasso, J. D. Joannopoulos, and Steven G. Johnson. Non-touching nanoparticle dimer clusters bound by repulsive and attractive Casimir forces. *Phys. Rev. Lett.*, 104(16):160402, 2010.
- [451] Alejandro W. Rodriguez, Alexander P. McCauley, David Woolf, Federico Capasso, J. D. Joannopoulos, and Steven G. Johnson. Non-touching nanoparticle dimer clusters bound by repulsive and attractive Casimir forces. *Phys. Rev. Lett.*, 104(16):160402, 2010.
- [452] Alejandro W. Rodriguez, Jeremy Munday, Diego Davlit, Federico Capasso, J. D. Joannopoulos, and Steven G. Johnson. Stable suspension and dispersion-induced transition from repulsive Casimir forces between fluid-separated eccentric cylinders. *Phys. Rev. Lett.*, 101(19):190404, 2008.
- [453] Alejandro W. Rodriguez, David Woolf, Alexander P. McCauley, Federico Capasso, and Steven G. Johnson. Achieving a strongly temperature-dependent Casimir effect. *arXiv:1004.2733*, 2010.
- [454] Pablo Rodriguez-Lopez, Sahand Jamal Rahi, and Thorsten Emig. Three-body Casimir effects and nonmonotonic forces. *Phys. Rev. A*, 80(2):022519, 2009.
- [455] J. Romero-Vivas, D. N. Chigrin, A. V. Lavrinenko, and C. M. Sotomayor Torres. Photonic quasicrystals for application in WDM systems. *Phys. Stat. Sol.*, 202(6):997–1001, 2005.
- [456] F. S. S. Rosa. On the possibility of Casimir repulsion using metamaterials. *J. Phys.: Conf. Series*, 161:012039, 2009.
- [457] F. S. S. Rosa, D. A. R. Dalvit, and P. Milonni. Casimir interactions for anisotropic magnetodielectric metamaterials. *Phys. Rev. A*, 78:032117, 2008.
- [458] F. S. S. Rosa, D. A. R. Dalvit, and P. W. Milonni. Casimir-Lifshitz theory and metamaterials. *Phys. Rev. Lett.*, 100:183602, 2008.

- [459] David Roundy and J. D. Joannopoulos. Photonic crystal structure with square symmetry within each layer and a three-dimensional band gap. *Appl. Phys. Lett.*, 82(22):3835–3837, 2003.
- [460] A. Roy, C. Y. Lin, and U. Mohideen. Improved precision measurement of the Casimir force. *Phys. Rev. D*, 60:111101(R), 1999.
- [461] A. Roy and U. Mohideen. Demonstration of the nontrivial boundary dependence of the Casimir force. *Phys. Rev. Lett.*, 82:4380, 1999.
- [462] Zhichao Ruan, Georgios Veronis, Konstanti L. Vodopyanov, Marty M. Fejer, and Shanhui Fan. Enhancement of optics-to-thz conversion efficiency by metallic slot waveguides. *Opt. Express*, 17(16):13502–13515, 2009.
- [463] W. Rudin. *Real and Complex Analysis*. McGraw-Hill, New York, 1966.
- [464] H. Y. Ruy, M. Notomi, and Y. H. Lee. High-quality-factor and small-mode-volume hexapole modes in photonic-crystal slab microcavities. *Appl. Phys. Lett.*, 83:4294–4296, 2003.
- [465] S. M. Rytov, V. I. Tatarskii, and Yu. A. Kravtsov. *Principles of Statistical Radiophysics II: Correlation Theory of Random Processes*. Springer-Verlag, 1989.
- [466] H.Y. Ryu and M. Notomi. *Opt. Lett.*, 28:2390, 2003.
- [467] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, Reading, MA, 1994.
- [468] J. J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, Reading, MA, revised edition, 1994.
- [469] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*. Wiley, New York 1991.
- [470] Luigi Scaccabarozzi, M. M. Fejer, Yijie Huo, Shanhui Fan, Xiaojun Yu, and James S. Harris. Enhanced second-harmonic generation in algaas/al_xo_y tightly confining waveguides and resonant cavities. *OL*, 31(24):3626–3630, 2006.
- [471] A. Scardicchio and R. L. Jaffe. Casimir effects: An optical approach I. foundations and examples. *Nuclear Physics B.*, 704(3):552–582, 2005.
- [472] J. E. Schaar, K. L. Vodopyanov, and M. M. Fejer. Intracavity terahertz-wave generation in a synchronously pumped optical parametric oscillator using quasi-phase-matched gaas. *Opt. Lett.*, 32(10):1284–1286, 2007.
- [473] M. Schaden and L. Spruch. Infinity-free semiclassical evaluation of Casimir effects. *Phys. Rev. A*, 58:935–953, 1998.
- [474] Stephan Schiller. *Principles and Applications of Optical Monolithic Total-Internal-Reflection Resonators*. PhD thesis, Stanford University, Stanford, CA, March 1993.
- [475] Holger Schmidt and Atac Imamoglu. *Opt. Lett.*, 21:1936, 1996.
- [476] K Schneider, Stephan Schiller, and Jürgen Mlynek. 1.1-w single-frequency 532-nm radiation by second-harmonic generation of a miniature nd:yag ring laser. *Opt. Lett.*, 21:1999–2001, 1996.

- [477] Pierre Scotto, Pere Colet, and Maxi San Miguel. All-optical image processing with cavity type ii second-harmonic generation. *Opt. Lett.*, 28:1695, 2003.
- [478] Marlan O. Scully and M. Suhail Zubairy. *Quantum Optics*. Cambridge University Press, Cambridge, England, 1997.
- [479] Marlan O. Scully and Suhail Zubairy. *Quantum Optics*. Cambridge University Press, Cambridge, UK, 1997.
- [480] R. Sedmik, I. Vasiljevich, and M. Tajmar. Detailed parametric study of Casimir forces in the Casimir Polder approximation for nontrivial 3d geometries. *J. Computer-Aided Mat. Des.*, 14(1):119–132, 2007.
- [481] F. Michael Serry, Dirk Walliser, and Maclay G. Jordan. The role of the Casimir effect in the static deflection of and stiction of membrane strips in microelectromechanical systems MEMS. *J. Appl. Phys.*, 84:2501, 1998.
- [482] S. S. Shamailov, A. S. Parkins, M. J. Collet, and H. J. Carmichael. Multi-photon blockade and dressing of the dressed state. *Opt. Comm.*, 283(5):766–772, 2010.
- [483] F. Shanhui, Wonjoo Suh, and J. D. Joannopoulos. Temporal coupled-mode theory for the fano resonance in optical resonators. *JOSA A*, 20(3):569–572, 2003.
- [484] Cheng-Gang Shao, Ai-Hong Tong, and Jun Luo. Casimir torque between birefringent plates. *Phys. Rev. A*, 72:022102, 2005.
- [485] Cheng-Gang Shao, Da-Lei Zheng, and Jun Luo. Repulsive Casimir effect between anisotropic dielectric and permeable plates. *Phys. Rev. A*, 74:012103, 2006.
- [486] Jung-Tsung Shen and Shanhui Fan. *Phys. Rev. Lett.*, 95:213001, 2005.
- [487] Jung-Tsung Shen and Shanhui Fan. Strongly correlated two-photon transport in a one-dimensional waveguide coupled to a two-level system. *Phys. Rev. Lett.*, 98(15):153003, 2007.
- [488] Jung-Tsung Shen and Shanhui Fan. Theory of single-photon transport in a single-mode waveguide. i. coupling to a cavity containing a two-level atom. *Phys. Rev. A*, 79(2):023837, 2009.
- [489] Jung-Tsung Shen and Shanhui Fan. Theory of single-photon transport in a single-mode waveguide. ii. coupling to a whispering-gallery resonator containing a two-level atom. *Phys. Rev. A*, 79(2):023838, 2009.
- [490] Jung-Tsung Shen, M. L. Povinelli, Sunil Sandhu, and Shanhui Fan. Stopping single photons in one-dimensional circuit quantum electrodynamics systems. *Phys. Rev. B*, 75(3):035320, 2007.
- [491] Yury Sherkunov. Van der Waals interaction of excited media. *Phys. Rev. A*, 72:052703, 2005.
- [492] Anthony Siegman. *Lasers*. University Science Books, Mill Valley, CA, 1986.
- [493] Anthony E. Siegman. *Lasers*. University Science Books, Sausalito, CA, 1986.

- [494] S. Singh. *Nonlinear Optical Materials*, volume 3 of *Handbook of laser science and technology*. CRC Press, 1986.
- [495] M. Skorobogatiy, S. A. Jacobs, Steven G. Johnson, and Y. Fink. Geometric variations in high index-contrast waveguides, coupled mode theory in curvilinear coordinates. *Opt. Express*, 10(21):1227–1243, 2002.
- [496] Arlee V. Smit. Group-velocity-matched three-wave mixing in birefringent crystals. *Opt. Lett.*, 26(10):719–721, 2001.
- [497] David D. Smith, George Fischer, Robert W. Boyd, and Don A. Gregory. Cancellation of photoinduced absorption in metal nanoparticle composites through a counterintuitive consequence of local field effects. *J. Opt. Soc. Am. B*, 14:1625, 1997.
- [498] R. G. Smith. Theory of intracavity optical second-harmonic generation. *IEEE J. Quantum Electron.*, 6:215–223, 1970.
- [499] A. W. Snyder and J. D. Love. *Optical Waveguide Theory*. Chapman and Hall, London, 1983.
- [500] M. Soljacic and J. D. Joannopoulos. Enhancement of nonlinear effects using photonic crystals. *Nature Mater.*, 3:211–219, 2004.
- [501] Marin Soljačić, Mihai Ibanescu, Steven G. Johnson, Yoel Fink, and J. D. Joannopoulos. Optimal bistable switching in non-linear photonic crystals. *Phys. Rev. E Rapid Commun.*, 66:055601(R), 2002.
- [502] Marin Soljacic, Mihai Ibanescu, Steven G. Johnson, Yoel Fink, and John D. Joannopoulos. *Phys. Rev. E*, 66:055601, 2002.
- [503] Marin Soljacic and John D. Joannopoulos. *Nat. Mater.*, 3:211, 2004.
- [504] Marin Soljačić, Steven G. Johnson, Shanhui Fan, Mihai Ibanescu, Erich Ippen, and J. D. Joannopoulos. Photonic-crystal slow-light enhancement of non-linear phase sensitivity. *J. Opt. Soc. Am. B*, 19:2052–2059, 2002.
- [505] Marin Soljacic, Elefterios Lidorikis, John D. Joannopoulos, and Lene V. Hau. *Appl. Phys. Lett.*, 86:171101, 2005.
- [506] Marin Soljacic, C. Luo, J. D. Joannopoulos, and S. Fan. Nonlinear photonic crystal microdevices for optical integration. *Opt. Lett.*, 28:637–639, 2003.
- [507] M. Soljacioc, M. Imabenscu, S. G. Johnson, J. D. Joannopoulos, and Y. Fink. Optical bistability in axially modulated omniguide fibers. *Opt. Lett.*, 28:516–518, 2003.
- [508] Bong-Shik Song, Susumu Noda, Takashi Asano, and Yoshihiro Akahane. Ultra-high- q photonic double-heterostructure nanocavity. *Nature Materials*, 4(3):207–210, 2005.
- [509] H. S. Sözüer and J. P. Dowling. Photonic band calculations for woodpile structures. *J. Mod. Opt.*, 41(2):231–239, 1994.
- [510] S.M. Spillane, T.J. Kippenberg, K.J. Vahala, K.W. Goh, E. Wilcut, and H.J. Kimble. *Phys. Rev. A*, 71:013817, 2005.

- [511] K. Srinivasan, P. E. Barclay, O. Painter, J. X. Chen, A. Y. Cho, and C. Gmachl. Experimental demonstration of high quality factor photonic crystal microcavity. *Appl. Phys. Lett.*, 83:1915–1917, 2003.
- [512] D. H. Staelin, A. W. Morgenthaler, and J. A. Kong. *Electromagnetic Waves*. Prentice Hall, Englewood Cliffs, NJ, 1994.
- [513] Walter Steurer and Torsten Haibach, editors. *Physical Properties of Quasicrystals*, chapter 2. Springer, 1999.
- [514] Walter Steurer and Daniel Sutter-Widmer. Photonic and phononic quasicrystals. *J. Phys. D*, 40:R229–R247, 2007.
- [515] John Strikwerda. *Finite Difference Schemes and Partial Differential Equations*. Wadsworth and Brooks/Cole, Pacific Grove, CA, 1989.
- [516] Steven H. Strogatz. *Nonlinear Dynamics and Chaos*. Westview Press, Boulder, CO, 1994.
- [517] J. B. Suck, M. Schreiber, and P. Haussler, editors. *Quasicrystals*, chapter 2. Springer, 2004.
- [518] Yoshimasa Sugimoto, Yu Tanaka, Naoki Ikeda, Yusui Nakamura, and Kiyoshi Asakawa. Low propagation loss of 0.76 db/mm in GaAs-based single-line-defect two-dimensional photonic crystal slab waveguides up to 1 cm in length. *Opt. Express*, 12(6):1090–1096, 2004.
- [519] Wonjoo Suh, Zheng Wang, and Shanhui Fan. *IEEE J. Quantum Electron.*, 40:1511, 2004.
- [520] Wonjoo Suh, Zheng Wang, and Shanhui Fan. Temporal coupled-mode theory and the presence of non-orthogonal modes in lossless multimode cavities. *IEEE J. Quantum Electron.*, 40(10):1511–1518, 2004.
- [521] Michael Tabor. *Chaos and Integrability in Nonlinear Dynamics: An Introduction*. Wiley, New York, 1989.
- [522] Allen Taflove and Susan C. Hagness. *Computational Electrodynamics*. Artech House, Norwood, MA, 2nd edition, 2000.
- [523] Allen Taflove and Susan C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. Artech, Norwood, MA, 2000.
- [524] M. Tajmar. Finite element simulation of Casimir forces in arbitrary geometries. *Intl. J. Mod. Phys. C*, 15(10):1387–1395, 2004.
- [525] D. V. Talapin, A. L. Rogach, A. Kornowski, M. Haase, and H. Weller. *Nano Lett.*, 1:207, 2001.
- [526] T. Tanabe, M. Notomi, S. Mitsugi, A. Shinya, and E. Kuramochi. Fast bistable all-optical switch and memory on a silicon photonic crystal on-chip. *Opt. Lett.*, 30:2575–2577, 2005.

- [527] Richard S. Tasgal and Boris A. Band Y. B. Malomed. Gap solitons in a medium with third-harmonic generation. *Phys. Rev. E*, 72:016624, 2005.
- [528] T. Tiedgje, E. Yablonovitch, G. D. Cody, and B. G. Brooks. Limiting efficiency of silicon solar cells. *IEEE Trans. Elec. Dev.*, 31(5):711–716, 1984.
- [529] M. S. Tomaš. Casimir force in absorbing multilayers. *Phys. Rev. A*, 66:052103, 2002.
- [530] M. Tonouchi. Cutting-edge terahertz technology. *Nature*, 1(97), 2007.
- [531] Lloyd N. Trefethen and David Bau. *Numerical linear algebra*. SIAM, Philadelphia, 1st edition, 1997.
- [532] U. Trottenberg, C. Ooseterlee, and A. Schüller. *Multigrid*. Academic Press, London, 2001.
- [533] J. Trull, R. Vilaseca, J. Martorell, and R. Corbalan. Second-harmonic generation in local modes of a truncated periodic structure. *Opt. Lett.*, 20:1746–1748, 1995.
- [534] Q.A. Turchette, C.J. Hood, W. Lange, H. Mabuchi, and H.J. Kimble. *Phys. Rev. Lett.*, 75:4710, 1995.
- [535] W. Turner. Terahertz quantum-cascade-laser source based on intracavity difference-frequency generation. *Nature Photonics*, 1:288–292, 2007.
- [536] A.V. Turukhin, V.S. Sudarshanam, M.S. Shahriar, J.A. Musser, B.S. Ham, and P.R. Hemmer. *Phys. Rev. Lett.*, 88:023602, 2002.
- [537] Kerry J. Vahala. Optical microcavities. *Nature*, 424:839–846, 2003.
- [538] P. H. G. M. van Blokland and J. T. G. Overbeek. Van der Waals forces between objects covered with a chromium layer. *J. Chem. Soc. Far. Trans. I*, 74:2637–2651, 1978.
- [539] S. J. Van Enk. The Casimir effect in dielectrics: A numerical approach. *Journal of Modern Optics*, 42(2):321–338, 1995.
- [540] S. J. van Enk. Casimir torque between dielectrics. *Phys. Rev. A*, 52(4):2569–2575, 1995.
- [541] S.J. van Enk, J.J. Cirac, and P. Zoller. *Phys. Rev. Lett.*, 78:4293, 1997.
- [542] D. W. Vernooy, V. S. Ilchenko, H. Mabuchi, E. W. Streed, and H. J. Kimble. High-q measurements of fused-silica microspheres in the near infrared. *Opt. Lett.*, 23(4):247–249, 1998.
- [543] A. Villeneuve, C. C. Yang, G. I. Stegeman, C. Lin, and H. Lin. Nonlinear refractive-index and two-photon absorption near half the band gap in algaas. *Appl. Phys. Lett.*, 62:2465–2467, 1993.
- [544] P. R. Villeneuve, Shanhui Fan, Steven G. Johnson, and J. D. Joannopoulos. Three-dimensional photon confinement in photonic crystals of low-dimensional periodicity. *IEE Proc. Optoelec.*, 145(6):384–390, 1998.

- [545] Pierre R. Villeneuve, Shanhui Fan, and J. D. Joannopoulos. Microcavities in photonic crystals: Mode symmetry, tunability, and coupling efficiency. *Phys. Rev. B*, 54:7837–7842, 1996.
- [546] Y. A. Vlasov, Xiang-Zheng-Bo, J. C. Sturm, and D. J. Norris. On-chip natural assembly of silicon photonic bandgap crystals. *Nature*, 414:289–293, 2001.
- [547] K. L. Vodopyanov, M. M. Fejer, X. Xu, J. S. Harris, Y. S. Lee, W. C. Hurlbut, V. G. Kozlov, D. Bliss, and C. Lynch. Terahertz-wave generation in quasi-phase-matched gaas. *Appl. Phys. Lett.*, 89(14):141119, 2006.
- [548] John L. Volakis, Arindam Chatterjee, and Leo C. Kempel. *Finite Element Method Electromagnetics: Antennas, Microwave Circuits, and Scattering Applications*. IEEE Press, New York, 2001.
- [549] Jelena Vučković, Marco Lončar, Hideo Mabuchi, and Axel Scherer. Design of photonic crystal microcavities for cavity qed. *Phys. Rev. E*, 65:016608, 2002.
- [550] A. Wade, G. Fedorov, D. Smirnov, S. Kumar, Q. Williams, Hu, and J. L. Reno. Magnetic-field-assisted terahertz quantum cascade laser operating up to 225 K. *Nature Photonics*, 3:41–45, 2008.
- [551] D.F. Walls and G.J. Milburn. *Quantum Optics*. Springer-Verlag, Berlin, 1994.
- [552] K. Wang, S. David, A. Chelnokov, and J. M. Lourtioz. Photonic band gaps in quasicrystal-related approximant structures. *J. Mod. Opt.*, 50(13):2095–2105, 2003.
- [553] Yiquan Wang, Cheng Bingying, and Daozhong Zhang. The density of states in quasiperiodic photonic crystals. *J. Phys.: Cond. Mat.*, 15:7675–7680, 2003.
- [554] M. R. Watts, S. G. Johnson, H. A. Haus, and J. D. Joannopoulos. Electromagnetic cavity with arbitrary q and small modal volume without a complete photonic bandgap. *Opt. Lett.*, 27(20):1785–1787, 2002.
- [555] Alexej Weber and Holder Gies. Nonmonotonic thermal Casimir force from geometry-temperature interplay. *arXiv:1003.0430*, 2010.
- [556] M.J. Werner and A. Imamoglu. *Phys. Rev. A*, 61:011801(R), 1999.
- [557] Karen Windmeier Wetz. Repulsive Casimir force between dielectric planes. *arXiv:quant-ph/0109134*, 2001.
- [558] D. S. Wiersma, R. Sapienza, S. Mujumdar, M. Colocci, M. Ghulinyan, and L. Pavesi. Optics of nanostructured dielectrics. *J. Opt. A: Pure and Appl. Opt.*, 7:S190–S197, 2005.
- [559] D. L. Wilcox, M. Berg, T. Bernat, D. Kellerman, and J. K. Cochran. *Hollow and Solid Spheres and Microspheres: Science and Technology Associated with Their Fabrication and Application*, volume 372. Society Symposium Proceedings, 1995.
- [560] Benjamin Williams, Sushil Kumar, Qing Hu, and John Reno. Operation of terahertz quantum-cascade lasers at 164 k in pulsed mode and at 117 k in continuous-wave mode. *Opt. Express*, 13(9):3331–3339, 2005.

- [561] Benjamin S. Williams. Terahertz quantum-cascade lasers. *Nature Photonics*, 1:517–525, 2007.
- [562] Chee Wei Wong, Peter T. Rakich, Steven G. Johnson, Minhai Qi, Henry I. Smith, Erich P. Ippen, Lionel C. Kimerling, Yongbae Jeon, George Barbastathis, and Sang-Gook Kim. Strain-tunable silicon photonic band gap microcavities in optical waveguides. *Appl. Phys. Lett.*, 84:1242–1245, 2004.
- [563] Ling-An Wu, Min Xiao, and H. J. Kimble. Squeezed states of light from an optical parametric oscillator. *JOSA-B*, 4:1465–1476, 1987.
- [564] Ping Xie, Zhao-Qing Zhang, and Xiangdong Zhang. Gap solitons and soliton trains in finite-sized two-dimensional periodic and quasiperiodic photonic crystals. *Phys. Rev. E*, 67:026607, 2003.
- [565] Jie Xiong and Weng Cho Chew. Efficient evaluation of casimir force in z-invariant geometries by integral equation methods. *Appl. Phys. Lett.*, 95:154102, 2009.
- [566] Jie L. Xiong, Mei Song Tong, Phillip Atkins, and Weng Cho Chew. Efficient evaluation of casimir force in arbitrary three-dimensional geometries by integral equation methods. *Physics Letters A*, 374(25):2517 – 2520, 2010.
- [567] Q. F. Xu and M. Lipson. Carrier-induced optical bistability in silicon ring resonators. *Opt. Lett.*, 31:341–343, 2006.
- [568] Qianfan Xu and Michal Lipson. Carrier-induced optical bistability in silicon ring resonators. *Opt. Lett.*, 31(3):341–343, 2005.
- [569] Yong Xu, Reginald K. Lee, and Amnon Yariv. Quantum analysis and the classical analysis of spontaneous emission in a microcavity. *Phys. Rev. A*, 61(3):033807, 2000.
- [570] E. Yablonovitch. Inhibited spontaneous emission in solid-state physics and electronics. *Phys. Rev. Lett.*, 58:2059–2062, 1987.
- [571] Yoshihisa Yamamoto and Atac Imamoglu. *Mesoscopic Quantum Optics*. John Wiley and Sons, Inc., New York, 1999.
- [572] Alexey Yamilov and Hui Cao. Highest-quality modes in disordered photonic crystals. *Phys. Rev. A*, 69(3):031803, 2004.
- [573] V. A. Yampol’skii, Sergey Savel’ev, Z. A. Mayselis, S. S. Apostolov, and Franco Nori. Anomalous temperature dependence of the Casimir force for thin metal films. *Phys. Rev. Lett.*, 101(9):096803, 2008.
- [574] Xiaodong Yang and Wong Chee Wei. Coupled-mode theory for stimulated Raman scattering in high- q/v_m silicon photonic band gap defect cavity lasers. *Opt. Express*, 15(8):4763–4781, 2006.
- [575] M. F. Yanik, S. Fan, and M. Soljacic. High-contrast all-optical bistable switching in photonic crystal microcavities. *Appl. Phys. Lett.*, 83:2739–2741, 2003.
- [576] Mehmet F. Yanik, Shanhui Fan, Marin Soljačić, J. D. Joannopoulos, and Yanik. All-optical transistor action with bistable switching in a photonic crystal cross-waveguide geometry. *Opt. Lett.*, 68:2506, 2004.

- [577] V. Yannopoulos and N. V. Vitanov. First-principles study of casimir repulsion in metamaterials. *Phys. Rev. Lett.*, 103:120401, 2009.
- [578] A. Yariv. *Optical Electronics in Modern Communication*. Oxford Univ. Press, Oxford, 5th edition, 1997.
- [579] A Yariv. *Quantum Electronics, 3rd ed.* Wiley, New York, 1998.
- [580] A. Yariv, Y. Xu, R. K. Lee, and A. Scherer. Coupled-resonator optical waveguide: a proposal and analysis. *Opt. Lett.*, 24:711–713, 1999.
- [581] Kiyotoshi Yasumoto. *Electromagnetic Theory and Applications for Photonic Crystals*. CRC Press, 2005.
- [582] M. S. Yeganeh, J. Qi, J. P. Culver, A. G. Yodh, and M. C. Tamargo. Three-wave-mixing spectroscopy of znse/gaas(001) heterointerfaces. *Phys. Rev. B*, 49(16):11196–11209, 1994.
- [583] K.-L. Yeh, M. C. Hoffmann, J. Hebling, and Keith A. Nelson. Generation of 10 mu j ultrashort terahertz pulses by optical rectification. *Applied Physics Letters*, 90(17):171121, 2007.
- [584] P. Yeh, A. Yariv, and E. Marom. Theory of Bragg fiber. *J. Opt. Soc. Am.*, 68:1196–1201, 1978.
- [585] Hiroyuki Yokoyama and kikuo Ujihara, editors. *Spontaneous emission and laser oscillation in microcavities*. CRC Press, Boca Raton, Florida, 1995.
- [586] T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H.M. Gibbs, G. Rupper, C. Ell, O.B. Shchekin, and D.G. Deppe. *Nature*, 432:200, 2004.
- [587] T. Yoshie, J. Vuckovic, A. Scherer, H. Chen, and D. Deppe. High quality two-dimensional photonic crystal slab cavities. *Appl. Phys. Lett.*, 79:4289–4291, 2001.
- [588] J. L. Young and R. O Nelson. A summary and systematic analysis of FDTD algorithms for linearly dispersive media. *IEEE Antennas and Propagation Magazine*, 43(1):61–126, 2001.
- [589] I. A. Zagorodnov, R. Schuhmann, and T. Weiland. A uniformly stable conformal fdtd-method in cartesian grids. *Int. J. Numer. Model*, 16:127–141, 2003.
- [590] Saad Zaheer, Alejandro W. Rodriguez, Steven G. Johnson, and Robert L. Jaffe. Optical-approximation analysis of sidewall-spacing effects on the force between two squares with parallel sidewalls. *Phys. Rev. A*, 76(6):063816, 2007.
- [591] J. Y. Zhang, H. L. Tam, W. H. Wong, Y. B. Pun, J. B. Xia, and K. W. Cheah. Isotropic photonic bandgap in Penrose textured metallic microcavity. *Sol. State Comm.*, 138:247–249, 2006.
- [592] Xiangdong Zhang, Zhao-Qing Zhang, and C. T. Chan. Absolute photonic band gaps in 12-fold symmetric photonic quasicrystals. *Phys. Rev. B*, 63:081105–081108, 2001.

- [593] Z. S. Zhang, B. Zhang, J. Xu, Z. J. Yang, Z. X. Qin, T. J. Yu, and D. P. Yu. Effects of symmetry of GaN-based two-dimensional photonic crystal with quasicrystal lattices on enhancement of surface light extraction. *Appl. Phys. Lett.*, 88:171103, 2006.
- [594] R. Zhao, T. Koschny, E. N. Economou, and C. M. Soukoulis. Comparison of chiral metamaterial designs for repulsive casimir force. *arXiv:0911.2019*, 2009.
- [595] R. Zhao, J. Zhou, E. N. Economou, and C. M. Soukoulis. Repulsive casimir force in chiral metamaterials. *Phys. Rev. Lett.*, 103:103602, 2009.
- [596] Yu Zhu and Andreas C. Cangellaris. *Multigrid Finite Element Methods for Electromagnetic Field Modelling*. John Wiley and Sons, Hooboke, NJ, 2006.
- [597] V. M. Zhuravlev. Exactly integrable model of three-wave mixing in an inhomogeneous nonlinear medium. *Soviet Phys. JETP*, 61(4):264–269, 1995.
- [598] E. S. Zijlstra and T. Janssen. Non-spiky density of states of an icosahedral quasicrystal. *Europhys. Lett.*, 52(5):578–583, 2000.
- [599] Richard W. Ziolkowski. *Metamaterials: Physics and Engineering Explorations*. John and Wiley and Sons, Inc., Hoboken, NJ, 2006.
- [600] M. E. Zoorob, M. D. B. Charlton, G. J. Parker, J. J. Baumberg, and M. C. Netti. Complete and absolute photonic bandgaps in highly symmetric photonic quasicrystals embedded in low refractive index materials. *Mat. Science and Eng.*, 74:168–174, 2000.
- [601] M. E. Zoorob, M. D.B. Charlton, G. J. Parker, J. J. Baumberg, and M. C. Netti. Complete photonic bandgaps in 12-fold symmetric quasicrystals. *Nature*, 404:740–743, 2000.
- [602] I I. Zootoverkh, Kravtsov N V., and E G Lariontsev. Enhancement of the efficiency of second-harmonic generation in microlaser. *Quantum Electronics*, 30:565, 2000.