

MIT Open Access Articles

Numerical Methods for Computing Casimir Interactions

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Johnson, Steven G. "Numerical Methods for Computing Casimir Interactions." In *Casimir Physics*, edited by Diego Dalvit, Peter Milonni, David Roberts, and Felipe da Rosa, 834:175–218. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.

As Published: http://dx.doi.org/10.1007/978-3-642-20288-9_6

Publisher: Springer

Persistent URL: <http://hdl.handle.net/1721.1/115587>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Chapter 7

Numerical methods for computing Casimir interactions

Steven G. Johnson

Abstract

We review several different approaches for computing Casimir forces and related fluctuation-induced interactions between bodies of arbitrary shapes and materials. The relationships between this problem and well known computational techniques from classical electromagnetism are emphasized. We also review the basic principles of standard computational methods, categorizing them according to three criteria—choice of problem, basis, and solution technique—that can be used to classify proposals for the Casimir problem as well. In this way, mature classical methods can be exploited to model Casimir physics, with a few important modifications.

7.1 Introduction

Thanks to the ubiquity of powerful, general-purpose computers, large-scale numerical calculations have become an important part of every field of science and engineering, enabling quantitative predictions, analysis, and design of ever more complex systems. There are a wide variety of different approaches to such calculations, and there is no single “best” method for all circumstances—not only are some methods better suited to particular situations than to others, but there are also often severe trade-offs between generality/simplicity and theoretical efficiency. Even in relatively mature areas like computational classical electromagnetism (EM), a variety of techniques spanning a broad range of sophistication and generality remain in widespread use (and new variations are continually developed) [1–8]. Semi-analytical approaches also remain important, especially perturbative techniques to decom-

Steven G. Johnson
Department of Mathematics, Massachusetts Institute of Technology, Cambridge MA 02139
e-mail: stevenj@math.mit.edu

pose problems containing widely differing length scales (the most challenging situation for brute-force numerics). Nevertheless, many commonalities and guiding principles can be identified that seem to apply to a range of numerical techniques.

Until a few years ago, Casimir forces and other EM fluctuation-induced interactions occupied an unusual position in this tableau. Realistic, general numerical methods to solve for Casimir forces were simply unavailable; solutions were limited to special high-symmetry geometries (and often to special materials like perfect metals) that are amenable to analytical and semi-analytical approaches. This is not to say that there were not, *in principle*, decades-old theoretical frameworks capable of describing fluctuations for arbitrary geometries and materials, but practical techniques for *evaluating* these theoretical descriptions on a computer have only been demonstrated in the last few years [9–27]. In almost all cases, these approaches turn out to be closely related to computational methods from *classical* EM, which is fortunate because it means that Casimir computations can exploit decades of progress in computational classical EM once the relationship between the problems becomes clear. The long delay in developing numerical methods for Casimir interactions, from the time the phenomenon was first proposed in 1948 [28], can be explained by three factors. First, accurate measurements of Casimir forces were first reported only in 1997 [29] and experimental interest in complex Casimir geometries and materials has only recently experienced dramatic growth due to the progress in fabricating nanoscale mechanical devices. Second, even the simplest numerical prediction of a single force requires the equivalent of a large number of classical EM simulations, a barrier to casual numerical experimentation. Third, there have historically been many equivalent theoretical formulations of Casimir forces, but some formulations are much more amenable to computational solution than others, and these formulations are often couched in a language that is opaque to researchers from classical computational EM.

This purpose of this review is to survey the available and proposed numerical techniques for evaluating Casimir forces, energies, torques, and related interactions, emphasizing their relationships to standard classical-EM methods. Our goal is not to identify a “best” method, but rather to illuminate the strengths and weaknesses of each approach, highlighting the conclusions that can be gleaned from the classical experience. We will review an intellectual framework in which to evaluate different numerical techniques, comparing them along several axes for which quasi-independent choices of approach can be made. We will also emphasize a few key departures of Casimir problems from ordinary classical EM, such as the necessity of imaginary- or complex-frequency solutions of Maxwell’s equations and the need for wide-bandwidth analyses, that impact the adaptation of off-the-shelf computational methods.

7.2 Characterization of numerical methods: Three axes

Numerical methods from distinct groups or research papers often differ in several ways simultaneously, complicating the task of directly comparing or even describing them. In order to organize one’s understanding of numerical approaches, it is useful to break them down along three axes of comparison, representing (roughly) independent choices in the design of a method:

- What **problem** does the method solve—even within a single area such as classical EM, there are several conceptually different *questions* that one can ask and several *ways of asking* them that lead to different categories of methods.
- What **basis** is used to express the unknowns—how the *infinite* number of unknowns in the exact partial differential equation (PDE) or integral equation are reduced to a *finite* number of unknowns for solution on a computer.
- What **solution technique** is used to determine these unknowns—even with the same equations and the same unknowns, there are vast differences among the types of direct, sparse, and iterative methods that can be used to attack the problem, and the efficient application of a particular solution technique to a particular problem is sometimes a research task unto itself.

In this section, we briefly summarize the available problems, basis choices, and solution techniques for Casimir problems. In subsequent sections, we then discuss in more detail the specific approaches that have currently been demonstrated or proposed.

7.2.1 Posing Casimir problems

In classical EM, there are several types of problems that are typically posed [6, appendix D], such as computing source-free time-harmonic eigensolutions $\mathbf{E}, \mathbf{H} \sim e^{-i\omega t}$ and eigenfrequencies ω , computing time-harmonic fields resulting from a time-harmonic current source $\mathbf{J} \sim e^{-i\omega t}$, or computing the time-dependent fields created by an arbitrary time-dependent source $\mathbf{J}(t)$ starting at $t = 0$. Although these are all closely mathematically related, and in some sense the solution of one problem can give solutions to the other problems, they lead to very different types of numerical simulations.

In a similar way, despite the fact that different formulations of Casimir-interaction problems are ultimately mathematically equivalent (although the equivalencies are often far from obvious)—and are usually answering the same conceptual question, such as what is the force or interaction energy for some geometry—each one leads most naturally to distinct classes of computational methods. Here, we exclude formulations such as proximity-force (“parallel-

plate”) approximations [30–32], pairwise summation of Casimir–Polder forces (valid in the dilute-gas limit) [33–35], and ray optics [36–39], that are useful in special cases but represent uncontrolled approximations if they are applied to arbitrary geometries. Although at some point the distinctions are blurred by the mathematical equivalencies, we can crudely categorize the approaches as:

- Computing the eigenfrequencies ω_n and summing the zero-point energy $\sum_n \frac{\hbar\omega_n}{2}$ [28, 40]. See Sec. 7.3.
- Integrating the mean energy density or force density (stress tensor), by evaluating field correlation functions $\langle E_i E_j \rangle_\omega$ and $\langle H_i H_j \rangle_\omega$ in terms of the classical EM Green’s functions at ω via the fluctuation–dissipation theorem [11–13, 19, 20, 23, 27]. See Sec. 7.5.
- Evaluating a path-integral expression for the interaction energy (or its derivative), constrained by the boundary conditions—usually, portions of the path integrals are performed analytically to express the problem in terms of classical scattering matrices or Green’s functions at each ω [9, 10, 14–18, 21, 22, 24–26]. See Sec. 7.6.

In each case, the result must be summed/integrated over all frequencies ω to obtain the physical result (corresponding to thermodynamic/quantum fluctuations at all frequencies). The relationship of the problem to *causal* Green’s functions (fields appear *after* currents) means that the integrand is analytic for $\text{Im}\omega \geq 0$ [41]. As a consequence, there is a choice of *contours* of ω integration in the upper-half complex plane, which is surprisingly important—it turns out that the integrands are wildly oscillatory on the real- ω axis and require accurate integration over a huge bandwidth, whereas the integrands are much better-behaved along the imaginary- ω axis (“Wick-rotated” or “Matsubara” frequencies). This means that Casimir calculations almost always involve classical EM problems evaluated at *complex or imaginary frequencies*, as is discussed further in Sec. 7.4. The nonzero-temperature case, where the integral over imaginary frequencies becomes a sum (numerically equivalent to a trapezoidal-rule approximation), is discussed in Sec. 7.8.

There is also another way to categorize the problem to be solved: whether one is solving a partial differential equation (**PDE**) or an **integral equation**. In a PDE, one has *volumetric unknowns*: fields or other functions at every point in space, related to one another *locally* by derivatives and so on. In an integral equation, one typically has *surface unknowns*: the fields or currents on the *boundaries* between piecewise-homogeneous regions, related to one another *non-locally* by the Green’s functions of the homogeneous regions (typically known analytically) [1, 3] (described further in Sec. 7.5.3). The key point is to take advantage of the common situation in which one has piecewise-constant materials, yielding a *surface integral equation*. (There are also *volume integral equations* for inhomogeneous media [42], as well as hybrid integral/PDE approaches [1], but these are less common.) There are other hybrid approaches such as *eigenmode expansion* [43–45], also called *rigorous*

coupled-wave analysis (RCWA) [46,47] or a *cross-section* method [48]: a structure is broken up along one direction into piecewise-constant cross-sections, and the unknown fields at the *interfaces* between cross-sections are propagated in the uniform sections via the *eigenmodes* of those cross-sections (computed analytically or numerically by solving the PDE in the cross-section). Eigenmode expansion is most advantageous for geometries in which the cross-section is constant over substantial regions, just as integral-equation methods are most advantageous to exploit large homogeneous regions.

7.2.2 Choices of basis

Casimir problems, for the most part, reduce to solving classical EM linear PDEs or integral equations where the unknowns reside in an infinite-dimensional vector space of functions. To *discretize* the problem approximately into a finite number N of unknowns, these unknown functions must be expanded in some finite *basis* (that converges to the exact solution as $N \rightarrow \infty$). There are three typical types of basis:

- **Finite differences** [2, 49, 50] (**FD**): approximate a function $f(x)$ by its values on some uniform grid with spacing Δx , approximate derivatives by some difference expression [e.g. second-order center differences $f'(x) \approx \frac{f(x+\Delta x) - f(x-\Delta x)}{2\Delta x} + O(\Delta x^2)$] and integrals by summations (e.g. a trapezoidal rule).
- **Finite-element methods** [1, 3, 4, 7] (**FEM**): divide space into geometric *elements* (e.g. triangles/tetrahedra), and expand an unknown $f(x)$ in a simple *localized* basis expansion for each element (typically, low-degree polynomials) with some continuity constraints. (FD methods are viewable as special cases of FEMs for uniform grids.) For an integral-equation approach, where the unknowns are functions on surfaces, the same idea is typically called a **boundary-element method (BEM)** [1, 3, 7, 51, 52].¹
- **Spectral methods** [53]: expand functions in a non-localized complete basis, truncated to a finite number of terms. Most commonly, Fourier series or related expansions are used (cosine series, Fourier–Bessel series, spherical or spheroidal harmonics, Chebyshev polynomials, *etc.*).

Finite differences have the advantage of simplicity of implementation and analysis, and the disadvantages of uniform spatial resolution and relatively low-order convergence (errors typically $\sim \Delta x^2$ [2] or even $\sim \Delta x$ in the presence of discontinuous materials unless special techniques are used [54, 55]). FEMs can have nonuniform spatial resolution to resolve disparate feature

¹ The name *method of moments* is also commonly applied to BEM techniques for EM. However, this terminology is somewhat ambiguous, and can refer more generally to Galerkin or other weighted-residual methods (and historically referred to monomial test functions, yielding statistical “moments”) [53].

sizes in the same problem, at a price of much greater complexity of implementation and solution techniques, and can have high-order convergence at the price of using complicated curved elements and high-order basis functions. Spectral methods can have very high-order or possibly exponential (“spectral”) convergence rates [53] that can even suit them to analytical solution—hence, spectral methods were the dominant technique before the computer era and are typically the first class of methods that appear in any field, such as in Mie’s classic solution of wave scattering from a sphere [56]. However, exponential convergence is usually obtained only if all discontinuities and singularities are taken explicitly into account in the basis [53]. With discontinuous materials, this is typically only practical for very smooth, high-symmetry geometries like spheres, cylinders, and so on; the use of a generic Fourier/spectral basis for arbitrary geometries reduces to a brute-force method that is sometimes very convenient [57], but may have unremarkable convergence rates [53, 57, 58]. BEMs require the most complicated implementation techniques, because any nontrivial change to the Green’s functions of the homogeneous regions (e.g. a change in dimensionality, boundary conditions, or material types) involves tricky changes to the singular-integration methods required to assemble the matrix [59–61] and to the fast-solver methods mentioned in Sec. 7.2.3.

Given FEM/BEM or spectral basis functions $b_n(x)$ and a linear equation $\hat{A}u(x) = v(x)$ for an unknown function u in terms of a linear differential/integral operator \hat{A} , there are two common ways [53] to obtain a finite set of N equations to determine the N unknown coefficients c_n in $u(x) \approx \sum_n c_n b_n(x)$. One is a **collocation** method: require that $(\hat{A}u - v)|_{x_n} = 0$ be satisfied at N collocation points x_n . The other is a **Galerkin** method: require that $\langle b_k, \hat{A}u - v \rangle = 0$ be satisfied for $k = 1, \dots, N$, where $\langle \cdot, \cdot \rangle$ is some inner product on the function space. Both approaches result in an $N \times N$ matrix equation of the form $\mathbf{A}\mathbf{u} = \mathbf{v}$. A Galerkin method has the useful property that if \hat{A} is Hermitian and/or definite then the matrix $A_{kn} = \langle b_k, \hat{A}b_n \rangle$ has the same properties.

The specific situation of vector-valued unknowns in EM creates additional considerations for the basis functions. In order to obtain center-difference approximations for all the field components, FD methods for EM typically use a staggered **Yee grid** [2, 49], in which each component of the EM fields is offset onto its own $\frac{\Delta x}{2}$ -shifted grid. In FEMs for EM, in order to maintain the appropriate continuity conditions for curl or divergence operators, one uses special classes of *vector-valued* basis functions such as *Niçedjçelec* elements [7, 62]. In BEMs for EM, vector-valued **RWG** (Rao, Wilton, and Glisson) basis functions [63] (or generalizations thereof [64]) are used in order to enforce a physical continuity condition on surface currents (to preclude accumulation of charge at element edges); see also Fig. 7.3 in Sec. 7.5.3. A spectral integral-equation method for EM with cylindrical or spherical scatterers is sometimes called a **multipole-expansion** method [5], since the

obvious spectral basis is equivalent to expanding the scattered fields in terms of multipole moments.

7.2.3 Solution techniques for linear equations

Given a particular problem and basis choice, one at the end obtains some $N \times N$ set of linear equations $A\mathbf{x} = \mathbf{b}$ to solve (or possibly eigenequations $A\mathbf{x} = \lambda B\mathbf{x}$).² Note also that a *single* Casimir-force calculation requires the solution of *many* such equations, at the very least for an integral over frequencies (see Sec. 7.4). There are essentially three ways to solve such a set of equations:

- **Dense-direct** solvers: solve $A\mathbf{x} = \mathbf{b}$ using direct matrix-factorization methods (e.g. Gaussian elimination),³ requiring $O(N^2)$ storage and $O(N^3)$ time [65].
- **Sparse-direct** solvers [66]: if A is *sparse* (mostly zero entries), use similar direct matrix-factorization methods, but cleverly re-arranged in an attempt to preserve the sparsity. Time and storage depend strongly on the *sparsity pattern* of A (the pattern of nonzero entries).
- **Iterative methods** [65, 67, 68]: repeatedly improve a guess for the solution \mathbf{x} (usually starting with a random or zero guess), only referencing A via repeated matrix–vector multiplies. Time depends strongly on the properties of A and the iterative technique, but typically requires only $O(N)$ storage. Exploits any fast way [ideally $O(N)$ or $O(N \log N)$] to multiply A by any arbitrary vector.

If the number N of degrees of freedom is small, i.e. if the basis converges rapidly for a given geometry, dense-direct methods are simple, quick, and headache-free (and have a standard state-of-the-art implementation in the free LAPACK library [69]). For example, $N = 1000$ problems can be solved in under a second on any modern computer with a few megabytes of memory. Up to $N \sim 10^4$ is reasonably feasible, but $N = 10^5$ requires almost 100 GB of memory and days of computation time without a large parallel computer. This makes dense-direct solvers the method of choice in simple geometries with a rapidly converging spectral basis, or with BEM integral-equation methods for basic shapes that can be accurately described by a few thousand triangular panels, but they rapidly become impractical for larger problems involving many and/or complex objects (or for moderate-size PDE problems even in two dimensions).

² This applies equally well, if somewhat indirectly, to the path-integral expressions of Sec. 7.6 where one evaluates a log determinant or a trace of an inverse, since this is done using either eigenvalues or the same matrix factorizations that are used to solve $A\mathbf{x} = \mathbf{b}$.

³ Technically, all eigensolvers for $N > 4$ are necessarily iterative, but modern dense-eigensolver techniques employ direct factorizations as steps of the process [65].

In PDE methods with a localized (FD or FEM) basis, the matrices A have a special property: they are *sparse* (mostly zero). The locality of the operators in a typical PDE means that each grid point or element directly interacts only with a bounded number of neighbors, in which case A has only $O(N)$ nonzero entries and can be stored with $O(N)$ memory. The process of solving $A\mathbf{x} = \mathbf{b}$, e.g. computing the LU factorization $A = LU$ by Gaussian elimination [65], unfortunately, ordinarily destroys this sparsity: the resulting L and U triangular matrices are generally not sparse. However, the pattern of nonzero entries that arises from a PDE is not random, and it turns out that clever re-orderings of the rows and columns during factorization can partially preserve sparsity for typical patterns; this insight leads to *sparse-direct* solvers [66], available via many free-software packages implementing different sparsity-preserving heuristics and other variations [68]. The sparsity pattern of A depends on the dimensionality of the problem, which determines the number of neighbors a given element interacts with. For meshes/grids having nearest-neighbor interactions, a sparse-direct solver typically requires $O(N)$ time and storage in 1d (where the matrices are band-diagonal), $O(N^{3/2})$ time with $O(N \log N)$ storage in 2d, and $O(N^2)$ time with $O(N^{4/3})$ storage in 3d [66, 70]. The practical upshot is that sparse-direct methods work well for 1d and 2d PDEs, but can grow to be impractical in 3d. For BEM and spectral methods, the interactions are not localized and the matrices are not sparse, so sparse-direct methods are not directly applicable (but see below for an indirect technique).

For the largest-scale problems, or for problems lacking a sparse A , the remaining possibility is an *iterative* method. In these methods, one need only supply a fast way to multiply A by an arbitrary vector y , and the trick is to use this Ay operation on a clever sequence of vectors in such a way as to make an arbitrary initial guess x_0 converge as rapidly as possible to the solution x , ideally using only $O(N)$ storage. Many such techniques have been developed [65, 67, 68]. The most favorable situation for $Ax = b$ occurs when A is Hermitian positive-definite, in which case an ideal Krylov method called the *conjugate-gradient* method can be applied, with excellent guaranteed convergence properties [65, 67], and fortunately this is precisely the case that usually arises for the imaginary-frequency Casimir methods below. There are two wrinkles that require special attention, however. First, one must have a fast way to compute Ay . If A is sparse (as for PDE and FD methods), then only $O(N)$ nonzero entries of A need be stored (as above) and Ay can be computed in $O(N)$ operations. In a spectral method, A is generally dense, but for spectral PDE methods there are often fast $O(N \log N)$ techniques to compute Ay using only $O(N)$ storage (A is stored implicitly), based on fast Fourier transform (FFT) algorithms [53, 57]. In a BEM, where A is again dense, a variety of sophisticated methods that require only $O(N \log N)$ computation time and $O(N)$ storage to compute Ay (again storing A implicitly) have been developed [1, 3, 7, 71], beginning with the pioneering fast-multipole method (FMM) [72]. These fast BEMs exploit the localized basis and the

decaying, convolutional nature of the Green’s function to approximate long-range interactions (to any desired accuracy). FMMs can be viewed as an approximate factorizations into sparse matrices, at which point sparse-direct methods are also applicable [73]. A second wrinkle is that the convergence rates of iterative methods depend on the condition number of A (the ratio of largest and smallest singular values) [65,67], and condition numbers generally worsen as the ratio of the largest and smallest lengthscales in the problem increases. To combat this, users of iterative methods employ **preconditioning** techniques: instead of solving $Ax = b$, one solves $KAx = Kb$ or similar, where the preconditioning matrix K is some crude *approximate inverse* for A (but much simpler to compute than A^{-1} !) such that the condition number of KA is reduced [67]. The difficulty with this approach is that good preconditioners tend to be highly problem-dependent, although a variety of useful approaches such as incomplete factorization and coarse-grid/multigrid approximations have been identified [65,67]. The upshot is that, while the largest-scale solvers almost invariably use iterative techniques, for any given class of physical problems it sometimes takes significant research before the iterative approach becomes well-optimized.

7.3 The impracticality of eigenmode summations

Perhaps the simplest way to express the Casimir energy, at zero temperature, is as a sum of zero-point energies of all oscillating EM modes in the system:

$$U = \sum_{\omega} \frac{\hbar\omega}{2}, \quad (7.1)$$

where ω is the frequency of the mode ($\sim e^{-i\omega t}$) [28,74]. That is, when the electromagnetic field is quantized into photons with energy $\hbar\omega$, it turns out that the vacuum state in the absence of photons is not empty, but rather has the energy equivalent of “half a photon” in each mode. The computational strategy is then straightforward, in principle: compute the EM eigenfrequencies ω in the problem by some numerical method (many techniques are available for computing eigenfrequencies [1,57]) and sum them to obtain U . Forces are then given by the derivative of U with respect to changes in the geometry, which could be approximated by finite differences or differentiated analytically with a Hellman–Feynman technique [75] (more generally, derivatives of any computed quantity can be computed efficiently by an adjoint method [76]).

Of course, U in eq. (7.1) has the disadvantage of being formally infinite, but this is actually a minor problem in practice: as soon as one discretizes the problem into a finite number of degrees of freedom (e.g., a finite number of grid points), the number of eigenfrequencies becomes finite (with the upper

bound representing a Nyquist-like frequency of the grid). This is the numerical analogue [12] of analytical regularization techniques that are applied to truncate the same sum in analytical computations [28]. (These regularizations do not affect energy *differences* or forces for rigid-body motions.) Matters are also somewhat subtle for dissipative or open systems [77]. But the most serious problem is that, even in the lossless case, this sum is badly behaved: even when one differentiates with separation a to obtain a finite force $F = -\frac{\hbar}{2} \sum \frac{d\omega}{da}$, the summand is wildly oscillatory and includes substantial contributions from essentially *every* frequency, which mostly cancel to leave a tiny result [12, 78]. Numerically, therefore, one must ostensibly compute *all* of the modes, to high precision, which requires $O(N^3)$ time and $O(N^2)$ storage (for a dense-direct eigensolver [65]) given N degrees of freedom. This is possible in simple 1d problems [12, 40], but is impractical as a general approach.

Because of the mathematical equivalence of the different approaches to the Casimir problem, the mode-summation method is sometimes useful as a starting point to derive alternative formulations, but the end result is invariably quite different in spirit from computing the eigenfrequencies one by one and summing them. For example, if one has a function $z(\omega)$ whose roots are the eigenfrequencies, then one can equivalently write U , via the residue theorem of complex analysis, as $U = \frac{1}{2\pi i} \oint_C \frac{\hbar\omega}{2} \frac{d(\ln z(\omega))}{d\omega} d\omega$, where C is any closed contour in the complex- ω plane that encloses the roots [79]. However, finding functions whose roots are the eigenfrequencies naturally points towards Green's functions (to relate different boundary conditions), and the contour choices typically involve Wick rotation as in Sec. 7.4, so this approach leads directly to imaginary-frequency scattering-matrix techniques as in Sec. 7.6 [18]. A similar contour integral arises from a zeta-function regularization of (7.1) [80].

7.4 The complex-frequency plane and contour choices

In order to better understand the frequency integration/summation in Casimir problems, it is illustrative to examine the analytical formula for the simple case of two perfect-metal plates in vacuum separated by a distance a , in which case it can be derived in a variety of ways that the attractive force F is given by [81]:

$$\begin{aligned}
 F &= \frac{\hbar}{\pi^2 c^3} \operatorname{Re} \left[\int_0^\infty d\omega \int_1^\infty dp \frac{p^2 \omega^3}{e^{2ip(\omega+i0^+)a/c} - 1} \right] \\
 &= \operatorname{Re} \left[\int_0^\infty f(\omega) d\omega \right] = \operatorname{Im} \left[\int_0^\infty f(i\xi) d\xi \right] = \frac{\hbar c}{240a^4},
 \end{aligned} \tag{7.2}$$

where $f(\omega)$ is the contribution of each frequency ω to the force and p is related to the plate-parallel momentum of the contributing modes/fluctuations. In this special case, the entire integral can be performed analytically, but for parallel plates of some finite permittivity ε the generalization (the *Lifshitz formula* [81]) must be integrated numerically. In practice, however, the formula and its generalizations are never integrated in the form at left—instead, one uses the technique of contour integration from complex analysis to *Wick rotate* the integral to imaginary frequencies $\omega = i\xi$, integrating over ξ . (In fact, the formula is typically derived *starting* in imaginary frequencies, via a Matsubara approach [81].) In this section, we review why a trick of this sort is both *possible* and *essential* in numerical computations for all of the methods described below.

Wick rotation is always *possible* as a consequence of causality. It turns out that the frequency contributions $f(\omega)$ for arbitrary materials and geometries, for all of the different formulations of the Casimir force below, are ultimately expressed in terms of classical EM *Green's functions* at ω : the EM fields in response to time-harmonic currents $\mathbf{J} \sim e^{-i\omega t}$. As a consequence of the causality of Maxwell's equations and physical materials—EM fields always arise *after* the source currents, not before—it mathematically follows that the Green's functions must be analytic functions (no poles or other singularities) when $\text{Im } \omega > 0$ (the upper-half complex plane) [41]. Poles in the Green's function correspond to eigenfrequencies or resonances of the source-free Maxwell's equations, and must lie at $\text{Im } \omega \leq 0$ for any physical system with dissipative materials (with the poles approaching $\text{Im } \omega = 0^-$ in the idealized lossless limit). [One can easily see explicitly that this is true for the $f(\omega)$ above: the poles result from a vanishing denominator in the p integrand, which only occurs for purely real ω corresponding to the real-frequency modes trapped between two perfect-metal plates.] As an elementary consequence of complex analysis, this analyticity means that the $\int d\omega$ can be arbitrarily deformed to any contour in the upper-half complex- ω plane without changing the integration result.

Wick rotation is *essential* for computation because the frequency contributions $f(\omega)$ to the force (or interaction energy or other related quantities) are extremely ill-behaved near to the real- ω axis: they are wildly oscillatory and slowly decaying. For example, the magnitude and phase of the function $f(\omega)$ are plotted in the complex ω plane in Fig. 7.1, where the p integral was evaluated numerically with a high-order Clenshaw–Curtis quadrature scheme [82]. Merely evaluating $f(\omega)$ along the real- ω axis is difficult because of singularities (which ultimately reduce the integral to a summation over eigenfrequency-contributions as in Sec. 7.3); in physical materials with dissipation, the real- ω axis is non-singular but is still badly behaved because of poles (lossy modes) located just below the axis. Along any contour parallel to the real- ω axis, the integrand is oscillatory (as can be seen from the phase plot) and non-decaying (as can be seen from the magnitude plot): formally, just as with the infinite summation over eigenmodes in Sec. 7.3, one must




Fig. 7.1 Contributions $f(\omega)$ to the Casimir force, from each fluctuation/mode frequency ω , for two perfect-metal plates with separation a , in the complex- ω plane. *Left*: magnitude $|f(\omega)|$. *Right*: phase $\angle f(\omega)$. [The magnitude is truncated at $10^3\hbar/a^3$, as it diverges towards the real- ω axis, and some numerical artifacts (rapid oscillations) are visible near the real- ω axis in the phase due to difficulty in evaluating $f(\omega)$.] The key point is that $f(\omega)$ is badly behaved (oscillatory and non-decaying) along contours parallel to the real- ω axis, whereas $f(\omega)$ is nicely behaved (non-oscillatory and exponentially decaying) along contours parallel to the imaginary- ω axis.

integrate over an infinite bandwidth, regularized in some way (e.g. by the Nyquist frequency placing an upper bound on ω for a finite grid), where the oscillations almost entirely cancel to leave a tiny remainder (the force). (Any physical materials must cease to polarize as $\omega \rightarrow \infty$ where the susceptibility vanishes [41], which will make the force contributions eventually vanish as $\omega \rightarrow \infty$ even in 1d, but a very wide-bandwidth oscillatory integral is still required.) This is a disaster for any numerical method—even when one is only integrating an analytical expression such as the Lifshitz formula, mere roundoff errors are a severe difficulty for real ω . Along the imaginary- ω axis, on the other hand (or any sufficiently vertical contour), $f(\omega)$ is exponentially decaying and mostly non-oscillatory—an ideal situation for numerical integration.

Therefore, in order for classical EM solvers to be used for Casimir problems, they need to be adapted to solve Maxwell’s equations at complex or imaginary ω . Although this sounds strange at first, the frequency-domain problem actually becomes numerically *easier* in every way at imaginary ω ; this is discussed in more detail in Sec. 7.5.1.2. In fact, one can even identify an exact mathematical equivalence between a particular complex- ω contour and a *real*-frequency system where an *artificial dissipation* has been introduced, as discussed in Sec. 7.5.5 below—using this trick, one can actually use classical EM solvers with no modification at all, as long as they handle dissipative media. In any case, one needs an integral over frequencies to compute a physically meaningful quantity, which means that solvers and material models, not to mention any physical intuition used for guidance, must

be valid for more than just a narrow real- ω bandwidth (unlike most problems in classical EM).

Numerically, it should be pointed out that the $f(i\xi)$ integrand is smooth and exponentially decaying, and so the ξ integral can be approximated to high accuracy by an exponentially convergent quadrature (numerical integration) scheme using evaluations at relatively few points ξ . For example, one can use Gauss–Laguerre quadrature [83], Gaussian quadrature with an appropriate change of variables [84], or Clenshaw–Curtis quadrature with an appropriate change of variables [82].

7.5 Mean energy/force densities and the fluctuation–dissipation theorem

Another, equivalent, viewpoint on Casimir interactions is that they arise from geometry-dependent fluctuations of the electromagnetic fields \mathbf{E} and \mathbf{H} , which on average have some nonzero energy density and exert a force. If we can compute these average fields, we can integrate the resulting energy density, stress tensors, and so on, to obtain the energy, force, or other quantities of interest. The good news is that there is a simple expression for those fluctuations in terms of the **fluctuation–dissipation** theorem of statistical physics: the correlation function of the fields is related to the corresponding *classical* Green’s function [81]. Ultimately, this means that any standard classical EM technique to compute Green’s functions (fields from currents) can be applied to compute Casimir forces, with the caveat that the techniques must be slightly modified to work at imaginary or complex frequencies as described below.

7.5.1 Background

The temperature- T correlation function for the fluctuating electric field at a given frequency ω is given by [81]:

$$\langle E_j(\mathbf{x})E_k(\mathbf{x}') \rangle_\omega = -\frac{\hbar}{\pi} \text{Im} [\omega^2 G_{jk}^E(\omega; \mathbf{x}, \mathbf{x}')] \coth(\hbar\omega/2k_{\text{B}}T), \quad (7.3)$$

where $G_{jk}^E = (\mathbf{G}_k^E)_j$ is the classical dyadic “photon” Green’s function, proportional⁴ to the relationship between an electric-dipole current in the k direction at \mathbf{x}' to the electric field at \mathbf{x} , and solves

⁴ The electric field $\mathbf{E}(\mathbf{x})$ from a dipole current $\mathbf{J} = \delta^3(\mathbf{x} - \mathbf{x}')\hat{e}_k e^{-i\omega t}$ is $\mathbf{E}(\mathbf{x}) = i\omega \mathbf{G}_k^E(\omega, \mathbf{x}, \mathbf{x}')e^{-i\omega t}$.

$$[\nabla \times \mu(\omega, \mathbf{x})^{-1} \nabla \times - \omega^2 \varepsilon(\omega, \mathbf{x})] \mathbf{G}_k^E(\omega, \mathbf{x}, \mathbf{x}') = \delta^3(\mathbf{x} - \mathbf{x}') \hat{e}_k, \quad (7.4)$$

where ε is the electric permittivity tensor, μ is the magnetic permeability tensor, and \hat{e}_k is a unit vector in direction k . Similarly, the magnetic-field correlation function is

$$\langle H_j(\mathbf{x}) H_k(\mathbf{x}') \rangle_\omega = -\frac{\hbar}{\pi} \text{Im} [\omega^2 G_{jk}^H(\omega; \mathbf{x}, \mathbf{x}')] \coth(\hbar\omega/2k_{\text{B}}T). \quad (7.5)$$

The magnetic Green's function \mathbf{G}^H can be defined in two essentially equivalent ways. The first is as derivatives $\frac{1}{\omega^2 \mu(\mathbf{x})} \nabla \times \mathbf{G}^E \times \nabla' \frac{1}{\mu(\mathbf{x}')}$ of the electric Green's function $G_{jk}^E(\mathbf{x}, \mathbf{x}')$, where ∇ and ∇' denote derivatives with respect to \mathbf{x} and \mathbf{x}' (∇' acting to the left), respectively [81]. The second way to define \mathbf{G}^H is proportional to the magnetic field in response to a magnetic-dipole current, analogous to (7.4):

$$[\nabla \times \varepsilon(\omega, \mathbf{x})^{-1} \nabla \times - \omega^2 \mu(\omega, \mathbf{x})] \mathbf{G}_k^H(\omega, \mathbf{x}, \mathbf{x}') = \delta^3(\mathbf{x} - \mathbf{x}') \hat{e}_k, \quad (7.6)$$

which can be more convenient for numerical calculation [13]. These two definitions are related [86] by $\mathbf{G}^H = \frac{1}{\omega^2 \mu(\mathbf{x})} \nabla \times \mathbf{G}^E \times \nabla' \frac{1}{\mu(\mathbf{x}')} - \frac{1}{\omega^2 \mu(\mathbf{x}')} \delta(\mathbf{x} - \mathbf{x}') I$ (with I being the 3×3 identity matrix),⁵ where the second (diagonal) term has no effect on energy differences or forces below and is therefore irrelevant. Now, these equations are rather nasty along the real- ω axis: not only will there be poles in \mathbf{G} just below the axis corresponding to lossy modes, but in the limit where the dissipative losses vanish (ε and μ become real), the combination of the poles approaching the real axis with the Im in the correlation function results in a delta function at each pole⁶ and integrals of the correlation functions turn into sums over modes as in Sec. 7.3. However, the saving grace, as pointed out in Sec. 7.4, is that Green's functions are *causal*, allowing us to transform any integral over all real fluctuation frequencies into an integral over *imaginary* fluctuation frequencies $\omega = i\xi$. The \coth factor has poles that alter this picture, but we will eliminate those for now by considering only the $T = 0^+$ case where $\coth(+\infty) = 1$, returning to nonzero temperatures in Sec. 7.8.

⁵ This can be seen more explicitly by substituting $\mathbf{G}^H = \frac{1}{\omega^2} \frac{1}{\mu} \nabla \times \mathbf{G}^E \times \nabla' \frac{1}{\mu'} - \frac{1}{\omega^2 \mu'} \delta$ into (7.6), with δ denoting $\delta(\mathbf{x} - \mathbf{x}') I$ and μ or μ' denoting $\mu(\mathbf{x})$ or $\mu(\mathbf{x}')$, respectively. In particular, $[\nabla \times \frac{1}{\varepsilon} \nabla \times - \omega^2 \mu] (\frac{1}{\omega^2} \frac{1}{\mu} \nabla \times \mathbf{G}^E \times \nabla' \frac{1}{\mu'} - \frac{1}{\omega^2 \mu'} \delta)$ yields $\nabla \times [\frac{1}{\omega^2 \varepsilon} \nabla \times \frac{1}{\mu} \nabla \times \mathbf{G}^E - \mathbf{G}^E] \times \nabla' \frac{1}{\mu'} - \nabla \times \frac{1}{\omega^2 \mu' \varepsilon} \nabla \times \delta + \delta$, which via (7.4) gives $+\nabla \times \frac{1}{\omega^2 \varepsilon} \delta \times \nabla' \frac{1}{\mu'} - \nabla \times \frac{1}{\omega^2 \mu' \varepsilon} \nabla \times \delta + \delta = \delta$ as desired, where in the last step we have used the fact that $\delta \times \nabla' = \nabla \times \delta$ [since $\nabla \times$ is antisymmetric under transposition and $\nabla' \delta(\mathbf{x} - \mathbf{x}') = -\nabla \delta(\mathbf{x} - \mathbf{x}')$].

⁶ This follows from the standard identity that the limit $\text{Im}[(x + i0^+)^{-1}]$, viewed as a distribution, yields $-\pi\delta(x)$ [85].

7.5.1.1 Energy density

In particular, to compute the Casimir energy U , we merely integrate the classical energy density in the EM field [41] over all positions and all fluctuation frequencies, Wick-rotated to an integral over imaginary frequencies, resulting in the expression:

$$U = \int_0^\infty d\xi \int \frac{1}{2} \left[\frac{d(\xi\varepsilon)}{d\xi} \langle |\mathbf{E}|^2 \rangle_{i\xi} + \frac{d(\xi\mu)}{d\xi} \langle |\mathbf{H}|^2 \rangle_{i\xi} \right] d^3\mathbf{x}, \quad (7.7)$$

where we have simplified to the case of isotropic materials (scalar ε and μ). At thermodynamic equilibrium, this expression remains valid even for arbitrary dissipative/dispersive media thanks to a direct equivalence with a path-integral expression [87], which is not obvious from the classical viewpoint in which the energy density is usually only derived in the approximation of negligible absorption [41]. (Thanks to the relationship between the Green's function and the local density of states [88], there is also a direct equivalence between this energy integral and eigenmode summation [12].) In the common case where μ has negligible frequency dependence (magnetic responses are usually negligible at the short wavelengths where Casimir forces are important, so that $\mu \approx \mu_0$), we can use the identity⁷ that $\int \varepsilon |\mathbf{E}|^2 = \int \mu |\mathbf{H}|^2$ for fields at any given frequency [6] to simplify this expression to [12]:

$$U = \int_0^\infty d\xi \int \frac{1}{2\xi} \frac{d(\xi^2\varepsilon)}{d\xi} \langle |\mathbf{E}|^2 \rangle_{i\xi} d^3\mathbf{x}. \quad (7.8)$$

Here, the zero-temperature imaginary-frequency mean-square electric field is given by:

$$\langle |\mathbf{E}(\mathbf{x})|^2 \rangle_{i\xi} = \frac{\hbar}{\pi} \xi^2 \text{tr} \mathbf{G}^E(i\xi; \mathbf{x}, \mathbf{x}), \quad (7.9)$$

where tr denotes the trace $\sum_j G_{jj}^E$ and the Im has disappeared compared to (7.3) because $\mathbf{G}^E(i\xi)$ is real and the Im cancels the i in $d\omega \rightarrow i d\xi$.

Equation (7.13) may at first strike one as odd, because one is evaluating the Green's function (field) at \mathbf{x} from a source at \mathbf{x} , which is formally infinite. This is yet another instance of the formal infinities that appear in Casimir problems, similar to the infinite sum over modes in Sec. 7.3. In practice, this is not a problem either analytically or numerically. Analytically, one typi-

⁷ Lest the application of this field identity appear too glib, we can also obtain the same equality directly from the Green's functions in the correlation functions. We have $\int \mu \langle |\mathbf{H}|^2 \rangle = \frac{\hbar}{\pi} \text{tr} \int \xi^2 \mu \mathbf{G}^H(\mathbf{x}, \mathbf{x})$, and from the identity after (7.6) we know that $\xi^2 \mu \mathbf{G}^H = -\nabla \times \mathbf{G} \times \nabla' \frac{1}{\mu'} + \delta$. However, because $\nabla \times$ is self-adjoint [6], we can integrate by parts to move $\nabla \times$ from the first argument/index of \mathbf{G}^E to the second, obtaining $-\mathbf{G}^E \times \nabla' \frac{1}{\mu'} \times \nabla' = \xi^2 \varepsilon' \mathbf{G}^E - \delta$ from the first term under the integral. (Here, we employ the fact that \mathbf{G}^E is real-symmetric at imaginary $\omega = i\xi$, from Sec. 7.5.1.2, to apply (7.10) to the second index/argument instead of the first.) This cancels the other delta from $\xi^2 \mu \mathbf{G}^H$ and leaves $\xi^2 \varepsilon \mathbf{G}^E$, giving $\varepsilon \langle |\mathbf{E}|^2 \rangle$ as desired.

cally regularizes the problem by subtracting off the vacuum Green's function (equivalent to only looking at the *portion* of the fields at \mathbf{x} which are reflected off of inhomogeneities in ε or μ) [81]. Numerically, in an FD or FEM method with a finite grid, the Green's function is everywhere finite (the grid is its own regularization) [12]. In a BEM, the Green's function is explicitly written as a sum of the vacuum field and scattered fields, so the former can again be subtracted analytically [12]. As in Sec. 7.3, these regularizations do not affect physically observable quantities such as forces or energy differences, assuming rigid-body motion.

7.5.1.2 The remarkable imaginary-frequency Green's function

This imaginary-frequency Green's function is actually a remarkably nice object. Wick-rotating eq. (7.4), it satisfies:

$$[\nabla \times \mu(i\xi, \mathbf{x})^{-1} \nabla \times + \xi^2 \varepsilon(i\xi, \mathbf{x})] \mathbf{G}_k^E(i\xi, \mathbf{x}, \mathbf{x}') = \delta^3(\mathbf{x} - \mathbf{x}') \hat{e}_k. \quad (7.10)$$

Because of causality, it turns out that ε and μ are strictly real-symmetric and positive-definite (in the absence of gain) along the imaginary-frequency axis, even for dissipative/dispersive materials [41]. Furthermore, the operator $\nabla \times \mu^{-1} \nabla \times$ is real-symmetric positive-semidefinite for a positive-definite real-symmetric μ [6]. Thus, the entire bracketed operator $[\dots]$ in eq. (7.10) is *real-symmetric positive-definite* for $\xi > 0$, which lends itself to some of the best numerical solution techniques (Cholesky decomposition [65], tridiagonal QR [65], conjugate gradients [65, 67], and Rayleigh-quotient methods [68]). (This definiteness is also another way of seeing the lack of poles or oscillations for $\omega = i\xi$.) It follows that the integral operator whose kernel is \mathbf{G}^E , i.e. the inverse of the $[\dots]$ operator in eq. (7.10), is also real-symmetric positive-definite, which is equally useful for integral-equation methods.

In vacuum, the 3d real- ω Green's function $\sim e^{i\omega|\mathbf{x}-\mathbf{x}'|/c}/|\mathbf{x}-\mathbf{x}'|$ [41] is Wick-rotated to $\sim e^{-\xi|\mathbf{x}-\mathbf{x}'|/c}/|\mathbf{x}-\mathbf{x}'|$, an *exponentially decaying, non-oscillatory* function. This is yet another way of understanding why, for $\omega = i\xi$, there are no interference effects and hence no “modes” (poles in \mathbf{G}), and integrands tend to be non-oscillatory and exponentially decaying (as $\xi \rightarrow \infty$, \mathbf{G} becomes exponentially short-ranged and does not “see” the interacting objects, cutting off the force contributions). (It also means, unfortunately, that a lot of the most interesting phenomena in classical EM, which stem from interference effects and resonances, may have very limited consequences for Casimir interactions.)

One other property we should mention is that the operator becomes semidefinite for $\xi = 0$, with a nullspace encompassing any static field distribution ($\nabla\phi$ for any scalar ϕ). This corresponds to the well-known singularity of Maxwell's equations at zero frequency [89, 90], where the electric and magnetic fields decouple [41]. Since we eventually integrate over ξ , the




Fig. 7.2 Schematic depiction of two objects whose Casimir interaction is desired. One computational method involves integrating a mean stress tensor around some closed surface (dashed red line) surrounding an object, yielding the force on that object.

measure-zero contribution from $\xi = 0$ does not actually matter, and one can use a quadrature scheme that avoids evaluating $\xi = 0$. However, in the nonzero-temperature case of Sec. 7.8 one obtains a sum over discrete- ξ contributions, in which case the zero-frequency term is explicitly present. In this case, $\xi = 0$ can be interpreted if necessary as the limit $\xi \rightarrow 0^+$ (which can be obtained accurately in several ways, e.g. by Richardson extrapolation [91], although some solvers need special care to be accurate at low frequency [89,90]); note, however, that there has been some controversy about the zero-frequency contribution in the unphysical limit of perfect/dissipationless metals [92].

7.5.1.3 Stress tensor

In practice, one often wants to know the Casimir force (or torque) on an object rather than the energy density. In this case, instead of integrating an electromagnetic energy density over the volume, one can integrate an electromagnetic *stress tensor* over a surface enclosing the object in question, schematically depicted in Fig. 7.2 [81]. The mean stress tensor for the Casimir force is [81]:

$$\begin{aligned} \langle T_{jk}(\mathbf{x}) \rangle_\omega = \varepsilon(\mathbf{x}, \omega) & \left[\langle E_j(\mathbf{x}) E_k(\mathbf{x}) \rangle_\omega - \frac{\delta_{jk}}{2} \sum_\ell \langle E_\ell(\mathbf{x})^2 \rangle_\omega \right] \\ & + \mu(\mathbf{x}, \omega) \left[\langle H_j(\mathbf{x}) H_k(\mathbf{x}) \rangle - \frac{\delta_{jk}}{2} \sum_\ell \langle H_\ell(\mathbf{x})^2 \rangle_\omega \right]. \end{aligned} \quad (7.11)$$

As above, the field correlation functions are expressed in terms of the classical Green's function, and the integral of the contributions over all ω is Wick-rotated to imaginary frequencies $\omega = i\xi$:

$$\mathbf{F} = \int_0^\infty d\xi \oint\!\!\!\!\!\oint_{\text{surface}} \langle \mathbf{T}(\mathbf{x}) \rangle_{i\xi} \cdot d\mathbf{S}, \quad (7.12)$$

with (zero-temperature) correlation functions

$$\langle E_j(\mathbf{x})E_k(\mathbf{x}) \rangle_{i\xi} = \frac{\hbar}{\pi} \xi^2 G_{jk}^E(i\xi; \mathbf{x}, \mathbf{x}), \quad (7.13)$$

$$\langle H_j(\mathbf{x})H_k(\mathbf{x}) \rangle_{i\xi} = \frac{\hbar}{\pi} \xi^2 G_{jk}^H(i\xi; \mathbf{x}, \mathbf{x}), \quad (7.14)$$

corresponding to the fields on the stress-integration surface in response to currents placed on that surface. To compute a Casimir torque around an origin \mathbf{r} , one instead uses $(\mathbf{x} - \mathbf{r}) \times \langle \mathbf{T}(\mathbf{x}) \rangle_{i\xi} \cdot d\mathbf{S}$ [93].

The derivation of this stress tensor (7.11) is not as straightforward as it might at first appear. If the stress-integration surface lies entirely in vacuum $\varepsilon \approx \varepsilon_0$ and $\mu \approx \mu_0$, then one can interpret (7.11) as merely the ordinary EM stress tensor from the microscopic Maxwell equations [41], albeit integrated over fluctuations. If the stress-integration surface lies in a dispersive/dissipative medium such as a fluid, however, then the classical EM stress tensor is well known to be problematic [41] and (7.11) may seem superficially incorrect. However, it turns out that these problems disappear in the context of thermodynamic equilibrium, where a more careful free-energy derivation of the Casimir force from fluctuations indeed results in equation (7.11) [81, 94],⁸ which has also proved consistent with experiments [95, 96]. Note also that, while (7.11) assumes the special case of isotropic media at \mathbf{x} , it can still be used to evaluate the force on objects made of anisotropic materials, as long as the stress-integration surface lies in an isotropic medium (e.g. vacuum or most fluids).

This formulation is especially important for methods that use an iterative solver for the Green's functions as discussed below, because it only requires solving for the response to currents on the stress-integration surface, rather than currents at every point in space to integrate the energy density, greatly reducing the number of right-hand sides to be solved for the linear equation (7.10) [12]; additional reductions in the number of right-hand sides are described in Sec. 7.5.6.

7.5.2 Finite-difference frequency-domain (FDFD)

In order to determine the Casimir energy or force, one evaluates the Green's function by solving eq. (7.10) and then integrates the appropriate energy/force density over volume/surface and over the imaginary frequency ξ . The central numerical problem is then the determination of the Green's function by solving a set of linear equations corresponding to (7.10), and probably the simplest approach is based on a finite-difference (FD) basis: space is divided

⁸ If compressibility of the fluid and the density-dependence of ε is not neglected, then there is an additional $\partial\varepsilon/\partial\rho$ term in (7.11) resulting from fluctuations in the density ρ [81].

into a uniform grid with some resolution Δx , derivatives are turned into differences, and one solves (7.10) by some method for every desired right-hand side. In classical EM (typically finding the field from a given current at real ω), this is known as a finite-difference frequency-domain (FDFD) method, and has been widely used for many years [5, 49].

For example, in one dimension for z -directed currents/fields with $\mu = 1$, (7.10) becomes $\left(-\frac{d^2}{dx^2} + \xi^2\right)\varepsilon G_{zz}^E = \delta(x-x')\hat{z}$. If we approximate $G_{zz}^E(n\Delta x, x') \approx G_n$, then the corresponding finite-difference equation, with a standard center-difference approximation for d^2/dx^2 [50], is

$$-\frac{G_{n+1} - 2G_n + G_{n-1}}{\Delta x^2} + \xi^2 \varepsilon_n G_n = \frac{\delta_{nn'}}{\Delta x}, \quad (7.15)$$

replacing the $\delta(x-x')$ with a discrete equivalent at n' . Equation (7.15) is a tridiagonal system of equations for the unknowns G_n . More generally, of course, one has derivatives in the y and z directions and three unknown \mathbf{G} (or \mathbf{E}) components to determine at each grid point. As mentioned in Sec. 7.2.2, it turns out that accurate center-difference approximations for the $\nabla \times \nabla \times$ operator in three dimensions are better suited to a “staggered” grid, called a *Yee* grid [2, 49], in which different field components are discretized at points slightly offset in space: e.g., $E_x([n_x + \frac{1}{2}]\Delta x, n_y\Delta y, n_z\Delta z)$, $E_y(n_x\Delta x, [n_y + \frac{1}{2}]\Delta y, n_z\Delta z)$, and $E_z(n_x\Delta x, n_y\Delta y, [n_z + \frac{1}{2}]\Delta z)$ for the \mathbf{E} field components. Note that any arbitrary frequency dependence of ε is trivial to include, because in frequency domain one is solving each ξ separately, and a perfect electric conductor is simply the $\varepsilon(i\xi) \rightarrow \infty$ limit.

One must, of course, somehow truncate the computational domain to a finite region of space in order to obtain a finite number N of degrees of freedom. There are many reasonable ways to do this because Casimir interactions are rapidly decaying in space (force $\sim 1/a^{d+1}$ or faster with distance a in d dimensions, at least for zero temperature). One could simply terminate the domain with Dirichlet or periodic boundary conditions, for example, and if the cell boundaries are far enough away from the objects of interest then the boundary effects will be negligible (quite different from classical EM problems at real ω !) [12]. In classical EM, one commonly uses the more sophisticated approach of a **perfectly matched layer** (PML), an artificial reflectionless absorbing material placed adjacent to the boundaries of the computational domain to eliminate outgoing waves [2]. Mathematically, a PML in a direction x is equivalent to a complex “coordinate stretching” $\frac{d}{dx} \rightarrow (1 + i\sigma/\omega)^{-1} \frac{d}{dx}$ for an artificial PML “conductivity” $\sigma(x) > 0$ [2, 98–100], where the $1/\omega$ factor is introduced to give an equal attenuation rate at all frequencies. However, at imaginary frequencies $\omega = i\xi$, a PML therefore results simply in a *real* coordinate stretching $\frac{d}{dx} \rightarrow (1 + \sigma/\xi)^{-1} \frac{d}{dx}$: in a PDE with decaying, non-oscillatory solutions (such as the imaginary- ω Maxwell equations), it is well known that a reasonable approach to truncating infinite domains is to perform a (real) coordinate transformation that compresses space far away

where the solution is small [53]. A convenience of Maxwell’s equations is that any coordinate transformation (real or complex) can be converted merely into a change of ε and μ [101], so any PML can be expressed simply as a change of materials while keeping the same PDE and discretization [99, 100].

Such a center-difference scheme is nominally second-order accurate, with discretization errors that vanish as $O(\Delta x^2)$ [2, 50]. One can also construct higher-order difference approximations (based on more grid points per difference). As a practical matter, however, the accuracy is limited instead by the treatment of material interfaces where ε changes discontinuously. If no special allowance is made for these interfaces, the method still converges, but its convergence rate is reduced by the discontinuity to $O(\Delta x)$ [54, 55, 102, 103] (unless one has \mathbf{E} polarization completely parallel to all interfaces so that there is no field discontinuity). There are various schemes to restore second-order (or higher) accuracy by employing specialized FD equations at the interfaces [54, 142], but an especially simple scheme involves unmodified FD equations with modified materials: it turns out that, if the discontinuous ε is smoothed in a particular way (to avoid introducing first-order errors by the smoothing itself), then second-order accuracy can be restored [55, 102, 103].⁹

Given the FD equations, one must then choose a solution technique to solve the resulting linear equations $A\mathbf{x} = \mathbf{b}$, where \mathbf{x} is the Green’s function (or field), \mathbf{b} is the delta-function (or current) right-hand side, and A is the discretized $\nabla \times \mu^{-1} \nabla \times + \xi^2 \varepsilon$ operator. Note that A is a real-symmetric positive-definite matrix at imaginary frequencies, as discussed in Sec. 7.5.1.2. Because A is sparse [only $O(N)$ nonzero entries], one can utilize a sparse-direct Cholesky factorization $A = R^T R$ (R is upper-triangular) [66] (for which many software packages are available [68]). Given this factorization, any right-hand side can be solved quickly by backsubstitution, so one can quickly sum the energy density over all grid points (essentially computing the trace of A^{-1}) to find the Casimir energy, or alternatively sum the stress tensor over a stress-integration surface to find the force. Precisely such a sparse-direct FD method for the Casimir energy was suggested by Pasquali and Maggs [23], albeit derived by a path-integral log det expression that is mathematically equivalent to summing the energy density (7.8) [87]. The alternative is an iterative technique, and in this case A ’s Hermitian definiteness means that an ideal Krylov method, the conjugate-gradient method [65, 67] can be employed [12]. The conjugate-gradient method requires $O(N)$ storage and time per iteration, and in the absence of preconditioning requires a number of iterations in d dimensions proportional to the diameter $O(N^{1/d})$ of the grid for each right-hand side [104]. The stress-tensor approach reduces the number of right-hand sides to be solved compared to energy-density in-

⁹ Even if the ε discontinuities are dealt with in this way, however, one may still fail to obtain second-order accuracy if the geometry contains sharp corners, which limit the accuracy to $O(\Delta x^p)$ for some $1 < p < 2$ [102]. This is an instance of *Darboux’s principle*: the convergence rate of a numerical method is generally limited by the strongest singularity in the solution that has not been explicitly compensated for [53].

tegration: one only needs to evaluate the Green's function for sources on a stress-integration surface, which has $O(N^{\frac{d-1}{d}})$ points in d dimensions. This gives a total time complexity of $O(N) \cdot O(N^{1/d}) \cdot O(N^{\frac{d-1}{d}}) = O(N^2)$ for an unpreconditioned iterative method; an ideal multigrid preconditioner can in principle reduce the number of iterations to $O(1)$ [4,105] (when N is increased by improving spatial resolution), yielding an $O(N^{2-\frac{1}{d}})$ time complexity. Substantial further improvements are obtained by realizing that one does not, in fact, need to sum over every point on the stress-integration surface, instead switching to a different spatial integration scheme described in Sec. 7.5.6.

7.5.3 Boundary-element methods (BEMs)

In some sense, a volume discretization such as an FD method is too general: in most physical situations, the medium is piecewise-constant, and one might want to take advantage of this fact. In particular, for the basic problem of finding the field in response to a current source at a given frequency, one can instead use a surface integral-equation approach: the unknowns are *surface currents* on the interfaces between homogeneous materials, and one solves for the surface currents so that the total field (source + surface currents) satisfies the appropriate boundary conditions at the interfaces [1,3,7]. For example, in the case of a perfect electric conductor, the surface-current unknowns can be the physical electric currents \mathbf{J} at the interface, and the boundary condition is that of vanishing tangential \mathbf{E} field.¹⁰ In the case of permeable media ε and μ , the physical (bound) currents are volumetric within the medium [e.g., the electric bound current is $\mathbf{J} = -i\omega(\varepsilon - \varepsilon_0)\mathbf{E}$], not surface currents [41]. However, it turns out that one can introduce *fictitious* surface electric and magnetic currents at all interfaces to provide enough degrees of freedom to satisfy the boundary condition of continuous tangential \mathbf{E} and \mathbf{H} , and thus to fully solve Maxwell's equations. The application of this *equivalence principle*¹¹ to obtain surface integral equations for BEM is known as the PMCHW approach (Poggio, Miller, Chang, Harrington, and Wu) [110–112]. In either case, one has surface (electric and/or magnetic) currents \mathbf{J}_s , plus an external current source \mathbf{J} [e.g., the right-hand side of (7.4)], so one can express the \mathbf{E} or \mathbf{H} field at any point \mathbf{x} as a convolution of $\mathbf{J} + \mathbf{J}_s$ with the *analytically known* Green's function $\mathbf{G}_0(\mathbf{x} - \mathbf{x}')$ of the corresponding *homogeneous* medium at \mathbf{x} . In BEM, one expresses \mathbf{J}_s , in turn, as a sum of *localized* basis functions \mathbf{b}_k

¹⁰ This is known as an *electric-field integral equation* (EFIE); one can also express the equations for perfect conductors in terms of boundary conditions enforced on *magnetic* fields (MFIE) or some linear *combination* of the two (CFIE), and the most effective formulation is still a matter of debate [90].

¹¹ The idea of solving scattering problems by introducing fictitious boundary currents had its origins [106–109] many years before its application to BEM by Harrington [110] and subsequent refinements.




Fig. 7.3 Example triangular mesh of the surfaces of two objects for a BEM solver [9]. Associated with each edge k is an “RWG” basis function \mathbf{b}_k [63], schematically represented in the inset, which vanishes outside the adjacent two triangles.

associated with some discrete mesh approximation of the surface. For example, Fig. 7.3 depicts a standard triangular-type mesh of two objects, where there is a localized basis function \mathbf{b}_k (inset) associated with each *edge* of this mesh such that \mathbf{b}_k is nonzero only on the adjacent two triangles [63]; this is the RWG basis mentioned in Sec. 7.2.2. Abstractly, the resulting equations for the fields could then be written in the following form:

$$\text{field}(\mathbf{x}) = \mathbf{G}_0 * (\mathbf{J} + \mathbf{J}_s) = \mathbf{G}_0 * \mathbf{J} + \sum_{k=1}^N \mathbf{G}_0 * \mathbf{b}_k c_k, \quad (7.16)$$

where $\mathbf{G}_0 *$ denotes convolution with the (dyadic) analytical Green’s function of the homogeneous medium at \mathbf{x} , and c_k are the unknown coefficients of each basis function.¹² (More generally, $\mathbf{G}_0 * \mathbf{J}$ could be replaced by any arbitrary incident field, regardless of how it is created.) In a Galerkin method (see Sec. 7.2.2), one obtains N equations for the N unknowns c_k by taking the inner product of both sides of this equation (substituted into the appropriate boundary condition) with the same basis functions \mathbf{b}_j (since they work just as well as a basis for the tangential field as for the tangential surface currents). This ultimately results in a set of linear equations $\mathbf{A}\mathbf{c} = \mathbf{d}$, where the matrix \mathbf{A} multiplying the unknown coefficients c_k is given by

¹² Technically, only currents from surfaces bordering the medium of \mathbf{x} contribute to this sum.

$$A_{jk} = \iint \bar{\mathbf{b}}_j(\mathbf{x}) \cdot \mathbf{G}_0(\mathbf{x} - \mathbf{x}') \cdot \mathbf{b}_k(\mathbf{x}') d^2\mathbf{x} d^2\mathbf{x}'. \quad (7.17)$$

[For the case of a perfect conductor with a vanishing tangential \mathbf{E} , the right-hand-side \mathbf{d} is given by $d_j = -\langle \mathbf{b}_j, \mathbf{G}_0 * \mathbf{J} \rangle = -\iint \bar{\mathbf{b}}_j(\mathbf{x}) \cdot \mathbf{G}_0(\mathbf{x} - \mathbf{x}') \cdot \mathbf{J}(\mathbf{x}') d^2\mathbf{x} d^2\mathbf{x}'$.] One then solves the linear system for the unknown coefficients c_k , and hence for the unknown surface currents \mathbf{J}_s . Implementing this technique is nontrivial because the A_{jk} integrands (7.17) are singular for $j = k$ or for j adjacent to k , necessitating specialized quadrature techniques for a given form of \mathbf{G}_0 [60,61], but substantial guidance from the past decades of literature on the subject is available.

Given these currents, one can then evaluate the electric or magnetic field at any point \mathbf{x} , not just on the surface, by evaluating eq. (7.16) at that point. In particular, one can evaluate the field correlation functions via the fluctuation-dissipation theorem (7.3): $\langle E_j(\mathbf{x}) E_k(\mathbf{x}) \rangle_\omega$ is given in terms of the electric field in the j direction at \mathbf{x} from a delta-function current \mathbf{J} in the k direction at \mathbf{x} . Of course, as noted previously, this is infinite because the $\mathbf{G}_0 * \mathbf{J}$ term (the field from the delta function) blows up at \mathbf{x} , but in the Casimir case one is only interested in the *change* of the correlation functions due to the geometry—so, one can use the standard trick [81] of subtracting the vacuum contribution $\mathbf{G}_0 * \mathbf{J}$ and only computing the surface-current contribution $\mathbf{G}_0 * \mathbf{J}_s$ to the field at \mathbf{x} . In this way, one can compute the stress tensor, the energy density, and so on, as desired.

As explained in Sec. 7.4, the integral of contributions over all frequencies is best performed at imaginary frequencies, so all of the above must use $\omega = i\xi$. This only has the effect of Wick-rotating the homogeneous-medium dyadic Green's function \mathbf{G}_0 to the $\sim e^{-\xi|\mathbf{x}-\mathbf{x}'|}/|\mathbf{x}-\mathbf{x}'|$ imaginary-frequency Green's function. This makes the problem *easier*, in principle. First, the exponential decay cuts off long-range interactions, making fast-solver techniques (see Sec. 7.2.3) potentially even more effective. Second, the matrix A is now real-symmetric and positive-definite, which allows the use of more efficient linear solvers as noted previously. Fortunately, the $1/|\mathbf{x}-\mathbf{x}'|$ singularity of \mathbf{G}_0 is the same at real and imaginary frequencies, allowing existing techniques for the integration of (7.17) to be leveraged.

At first glance, this approach seems most straightforwardly applicable to the stress-tensor technique, as suggested in Ref. [12]: one uses the BEM solution to evaluate the mean stress tensor $\langle \mathbf{T} \rangle$ on any integration surface around a body, integrating via some quadrature technique to obtain the force. If one uses a dense-direct solver (when N is not too big), the Cholesky factorization of A can be computed once for a given ξ and then many right-hand sides can be solved quickly via backsubstitution [65] in order to integrate $\langle \mathbf{T} \rangle$ over the stress-integration surface. Precisely such a dense-direct BEM stress-tensor method was recently demonstrated to compute Casimir forces in two and three dimensions [19,20]. As described in Sec. 7.2.3, fast-solver techniques can be applied to multiply A by a vector in $O(N \log N)$ time

with $O(N)$ storage; given a good preconditioner, this implies that an iterative method such as conjugate-gradient (applicable since A is real-symmetric positive-definite) could find $\langle \mathbf{T} \rangle$ at a single \mathbf{x} and ξ in $O(N \log N)$ time. The remaining question is the number of points required for the surface integral of $\langle \mathbf{T} \rangle$, which depends on *why* one is increasing N : either to increase accuracy for a fixed geometry or to increase the complexity of the geometry for a fixed accuracy. In the former case, the smoothness of $\langle \mathbf{T} \rangle$ in \mathbf{x} means that exponentially convergent quadrature techniques are applicable, which converge much faster than the (polynomial) BEM basis for the surface currents, so that ultimately the number of stress-quadrature points¹³ should be independent of N and the overall complexity becomes $O(N \log N)$. In the latter case, for a fixed accuracy and increasingly complex geometry (or smaller feature sizes), it appears likely that the number of stress-quadrature points will increase with N , but detailed studies of this scaling are not yet available.

It turns out that this BEM approach is closely related to the BEM path-integral approach described in Sec. 7.6.3. Both approaches end up solving linear equations with exactly the same matrix A of eq. (7.17), with the same degrees of freedom. The path-integral approach shows, however, that this same matrix can be applied to compute the Casimir interaction energy as well as the force, with comparable computational cost for dense solvers. Moreover, as explained below, expressing the force in terms of the derivative of the path-integral energy results in a trace expression that is conceptually equivalent to integrating a stress tensor over the surface of an object, where the number of “quadrature points” is now exactly equal to N . An unanswered question, at this point, is whether a fast solver can be more efficiently (or more easily) exploited in the stress-tensor approach or in the path-integral approach.

7.5.4 Other possibilities: FEM and spectral methods

There are of course, many other frequency-domain techniques from classical EM that could potentially be used to solve for the Green’s function and hence the energy/force density. For example, one could use spectral integral-equation methods, such as multipole expansions for spheres and cylinders [5], to compute responses to currents, although the advantages of this approach compared to the spectral path-integral approach in Sec. 7.6.2 are unclear.

¹³ Numeric integration (*quadrature*) approximates an integral $\int f(x)dx$ by a sum $\sum_i f(x_i)w_i$ over quadrature points x_i with weights w_i . There are many techniques for the selection of these points and weights, and in general one can obtain an error that decreases exponentially fast with the number of points for analytic integrands [53,83,84,113]. Multi-dimensional quadrature, sometimes called *cubature*, should be used to integrate the stress tensor over a 2d surface, and numerous schemes have been developed for low-dimensional cubature [114,115] (including methods that adaptively place more quadrature points where they are most needed [116]). For spherical integration surfaces (or surfaces that can be smoothly mapped to spheres), specialized methods are available [117,118].

One can also solve the PDE formulation of the Green’s function (7.10) using a finite-element (FEM) approach with some general mesh; in principle, existing FEM techniques from classical EM [1, 3, 4, 7] are straightforwardly applicable. One subtlety that arises in FEM methods with a nonuniform resolution is the regularization, however [12]. In principle, as mentioned above, one needs to subtract the vacuum Green’s function contribution from the field correlation functions in order to get a physical result [since the vacuum Green’s function $\mathbf{G}(\mathbf{x}, \mathbf{x}')$ diverges as $\mathbf{x}' \rightarrow \mathbf{x}$, although the divergence is cut off by the finite mesh resolution]. With a uniform mesh, this vacuum contribution is the same everywhere in the mesh and hence automatically integrates to zero in the force (when the stress tensor is integrated over a closed surface or the energy is differentiated). For a nonuniform mesh, however, the vacuum contribution varies at different points in space with different resolution, so some “manual” regularization seems to be required (e.g., subtracting a calculation with the same mesh but removing the objects). These possibilities currently remain to be explored for Casimir physics.

7.5.5 *Finite-difference time-domain (FDTD) methods*

Casimir effects are fundamentally broad-bandwidth, integrating contributions of fluctuations at all frequencies (real or imaginary), although the imaginary-frequency response is dominated by a limited range of imaginary frequencies. In classical EM, when a broad-bandwidth response is desired, such as a transmission or reflection spectrum from some structure, there is a well-known alternative to computing the contributions at each frequency separately—instead, one can simulate the same problem in *time*, Fourier-transforming the response to a short pulse excitation in order to obtain the broad-bandwidth response in a single *time-domain* simulation [6, 119]. The same ideas are applicable to the Casimir problem, with a few twists, yielding a practical method [13, 27] that allows Casimir calculations to exploit off-the-shelf time-domain solvers implementing the standard *finite-difference time-domain* (FDTD) method [2]. There are two key components of this approach [13]: first, converting the frequency integral to a time integral and, second, finding a time-domain equivalent of the complex-frequency idea from Sec. 7.4.

As reviewed above, the mean fluctuations in the fields, such as $\langle E^2(\mathbf{x}) \rangle_\omega$, can be expressed in terms of the fields at \mathbf{x} from a frequency- ω current at \mathbf{x} . If, instead of a frequency- ω current, one uses a current with $\delta(t)$ time dependence, it follows by linearity of (7.4) that the Fourier transform of the resulting fields must yield exactly the same $\langle E^2(\mathbf{x}) \rangle_\omega$. Roughly, the procedure could be expressed as follows: First, we compute some function $\Gamma(t)$ of the time-domain fields from a sequence of simulations with $\delta(t)$ sources, e.g. where $\Gamma(t)$ is the result of spatially integrating the fields making up the mean

stress tensor $\langle \mathbf{T}(\mathbf{x}) \rangle$ [noting that each point \mathbf{x} involves several separate $\delta(t)$ -response simulations]. Second, we Fourier transform $\Gamma(t)$ to obtain $\tilde{\Gamma}(\omega)$. Third, we obtain the force (or energy, etcetera) by integrating $\int \tilde{\Gamma}(\omega) \tilde{g}(\omega) d\omega$ with appropriate frequency-weighting factor $\tilde{g}(\omega)$ (which may come from the frequency dependence of ε in $\langle \mathbf{T} \rangle$, a Jacobian factor from below, etcetera). At this point, however, it is clear that the Fourier transform of Γ was entirely unnecessary: because of the unitarity of the Fourier transform (the Plancherel theorem), $\int \tilde{\Gamma}(\omega) \tilde{g}(\omega) d\omega = \int \Gamma(t) g(-t) dt$. That is, we can compute the force (or energy, etcetera) by starting with $\delta(t)$ sources and simply integrating the response $\Gamma(t)$ in time (accumulated as the simulation progresses) multiplied by some (precomputed, geometry-independent) kernel $g(t)$ (which depends on temperature if the coth factor is included for $T > 0$). The details of this process, for the case of the stress tensor, are described in Refs. [13, 27].

Although it turns out to be possible to carry out this time-integration process as-is, we again find that a transformation into the complex-frequency plane is desirable for practical computation (here, to reduce the required simulation time) [13]. Transforming the *frequency* in a *time*-domain method, however, requires an indirect approach. The central observation is that, in the equation (7.4) for the electric-field Green's function \mathbf{G}^E , the frequency only appears explicitly in the $\omega^2 \varepsilon$ term, together with ε . So, any transformation of ω can equivalently be viewed as a transformation of ε . In particular suppose that we wish to make some transformation $\omega \rightarrow \omega(\xi)$ to obtain an ω in the upper-half complex plane, where ξ is a real parameter (e.g. $\omega = i\xi$ for a Wick rotation). Equivalently, we can view this as a calculation at a *real* frequency ξ for a transformed *complex material*: $\omega^2 \varepsilon(\omega, \mathbf{x}) \rightarrow \xi^2 \varepsilon_c(\xi, \mathbf{x})$ where the transformed material is [13, 120]

$$\varepsilon_c(\xi, \mathbf{x}) = \frac{\omega^2(\xi)}{\xi^2} \varepsilon(\omega(\xi), \mathbf{x}). \quad (7.18)$$

For example, a Wick rotation $\omega \rightarrow i\xi$ is equivalent to operating at a real frequency ξ with a material $\varepsilon(\omega) \rightarrow -\varepsilon(i\xi)$. However, at this point we run into a problem: multiplying ε by -1 yields exponentially growing solutions at negative frequencies [13, 120], and this will inevitably lead to exponential blowup in a time-domain simulation (which cannot avoid exciting negative frequencies, if only from roundoff noise). In order to obtain a useful time-domain simulation, we must choose a contour $\omega(\xi)$ that yields a *causal, dissipative* material ε_c , and one such choice is $\omega(\xi) = \xi \sqrt{1 + i\sigma/\xi}$ for any constant $\sigma > 0$ [13, 120]. This yields $\varepsilon_c = (1 + i\sigma/\omega)\varepsilon$, where the $i\sigma/\omega$ term behaves exactly like an artificial *conductivity* added everywhere in space. In the frequency-domain picture, we would say from Sec. 7.4 that this $\omega(\xi)$ contour will improve the computation by moving away from the real- ω axis, transforming the frequency integrand into something exponentially decaying and less oscillatory. In the time-domain picture, the σ term adds a *dissipation* everywhere in space that causes $\Gamma(t)$ to *decay exponentially in time*, allow-

ing us to truncate the simulation after a short time. As long as we include the appropriate Jacobian factor $\frac{d\omega}{d\xi}$ in our frequency integral, absorbing it into $g(t)$, we will obtain the *same result* in a much shorter time. The computational details of this transformation are described in Refs. [13, 27]. More generally, this equivalence between the Casimir force and a relatively narrow-bandwidth real-frequency response of a dissipative system potentially opens other avenues for the understanding of Casimir physics [120].

The end result is a computational method for the Casimir force in which one takes an off-the-shelf time-domain solver (real time/frequency), adds an artificial conductivity σ everywhere, and then accumulates the response $\Gamma(t)$ to short pulses multiplied by a precomputed (geometry independent) kernel $g(t)$. The most common time-domain simulation technique in classical EM is the FDTD method [2]. Essentially, FDTD works by taking the same spatial Yee discretization as in the FDFD method above, and then also discretizing time with some time step Δt . The fields are then marched through time in steps of Δt , where each time step requires $O(N)$ work for N spatial grid points. Because the complex- ω contour is implemented entirely as a choice of materials ε_c , existing FDTD software can be used without modification to compute Casimir forces, and one can exploit powerful existing software implementing parallel calculations, various dimensionalities and symmetries, general dispersive and anisotropic materials, PML absorbing boundaries, and techniques for accurate handling of discontinuous materials. One such FDTD package is available as free/open-source software from our group [119], and we have included built-in facilities to compute Casimir forces [121].

7.5.6 Accelerating FD convergence

Finally, we should mention a few techniques that accelerate the convergence and reduce the computational cost of the finite-difference approaches. These techniques are not *necessary* for convergence, but they are simple to implement and provide significant efficiency benefits.

The simplest technique is extrapolation in Δx : since the convergence rate of the error with the spatial resolution Δx is generally known *a priori*, one can fit the results computed at two or more resolutions in order to extrapolate to $\Delta x \rightarrow 0$. The generalization of this approach is known as *Richardson extrapolation* [91], and it can essentially increase the convergence order cheaply, e.g., improving $O(\Delta x)$ to $O(\Delta x^2)$ [123].

Second, suppose one is computing the force between two objects A and B surrounded by a homogeneous medium. If one of the objects, say B , is removed, then (in principle) there should be no net remaining force on A . However, because of discretization asymmetry, a computation with A alone will sometimes still give a small net force, which converges to zero as $\Delta x \rightarrow 0$. If this “error” force is subtracted from the A – B force calculation, it turns out

that the net error is reduced. More generally, the error is greatly reduced if one computes the A – B force and then subtracts the “error” forces for A alone and for B alone, tripling the number of computations but greatly reducing the resolution that is required for an accurate result [12].

Third, when integrating the stress tensor $\langle \mathbf{T}(\mathbf{x}) \rangle_{i\xi}$ over \mathbf{x} to obtain the net force (7.12), the most straightforward technique in FD is to simply sum over all the grid points on the integration surface—recall that each point \mathbf{x} requires a linear solve (a different right-hand side) in frequency domain, or alternatively a separate time-domain simulation (a separate current pulse). This is wasteful, however, because $\langle \mathbf{T}(\mathbf{x}) \rangle_{i\xi}$ is conceptually smoothly varying in space—if one could evaluate it at arbitrary points \mathbf{x} (as is possible in the BEM approach), an exponentially convergent quadrature scheme could be exploited to obtain an accurate integral with just a few \mathbf{x} ’s. This is not directly possible in an FD method, but one can employ a related approach. If the integration surface is a box aligned with the grid, one can expand the fields on each side of the box in a cosine series (a discrete cosine transform, or DCT, since space is discrete)—this generally converges rapidly, so only a small number terms from each side are required for an accurate integration. But instead of putting in point sources, obtaining the responses, and expanding the response in a cosine series, it is equivalent (by linearity) to put in cosine sources directly instead of point sources. [Mathematically, we are exploiting the fact that a delta function can be expanded in any orthonormal basis $b_n(\mathbf{x})$ over the surface, such as a cosine series, via: $\delta(\mathbf{x} - \mathbf{x}') = \sum_n \bar{b}_n(\mathbf{x}') b_n(\mathbf{x})$. Substituting this into the right-hand side of (7.10), each $b_n(\mathbf{x})$ acts like a current source and $\bar{b}_n(\mathbf{x}')$ scales the result, which is eventually integrated over \mathbf{x}' .] The details of this process and its convergence rate are described in Ref. [27], but the consequence is that many fewer linear systems (fewer right-hand sides) need be solved (either in frequency or time domain) than if one solved for the stress tensor at each point individually.

7.6 Path integrals and scattering matrices

Another formulation of Casimir interactions is to use a derivation based on path integrals. Although the path-integral derivation itself is a bit unusual from the perspective of classical EM, and there are several slightly different variations on this idea in the literature, the end result is straightforward: Casimir energies and forces are expressed in terms of log determinants and traces of classical scattering matrices [10, 14–18, 21, 22, 24–26], or similarly the interaction matrices (7.17) that arise in BEM formulations [9]. Here, we omit the details of the derivations and focus mainly on the common case of piecewise-homogeneous materials, emphasizing the relationship of the resulting method to surface-integral equations from classical EM via the approach in Ref. [9].

Path integrals relate the Casimir interaction energy U of a given configuration to a functional integral over all possible vector-potential fields \mathbf{A} . Assuming piecewise-homogeneous materials, the constraint that the fields in this path integral must satisfy the appropriate boundary conditions can be expressed in terms of auxiliary fields \mathbf{J} at the interfaces (a sort of Lagrange multiplier) [124].¹⁴ At this point, the original functional integral over \mathbf{A} can be performed analytically, resulting in an energy expression involving a functional integral $Z(\xi)$ over only the auxiliary fields \mathbf{J} at each imaginary frequency ξ , of the form (at zero temperature):

$$U = -\frac{\hbar c}{2\pi} \int_0^\infty \log \det \frac{Z(\xi)}{Z_\infty(\xi)} d\xi, \quad (7.19)$$

$$Z(\xi) = \int \mathcal{D}\mathbf{J} e^{-\frac{1}{2} \iint d^2\mathbf{x} \iint d^2\mathbf{x}' \mathbf{J}(\mathbf{x}) \cdot \mathbf{G}_\xi(\mathbf{x}-\mathbf{x}') \cdot \mathbf{J}(\mathbf{x}')}. \quad (7.20)$$

Here, Z_∞ denotes Z when the objects are at infinite separation (non-interacting), regularizing U to just the (finite) interaction energy (See also the chapter by S. J. Rahi et al. in this volume for additional discussion of path integrals and Casimir interactions.) In the case of perfect electric conductors in vacuum, \mathbf{J} can be interpreted as a surface current on each conductor (enforcing the vanishing tangential \mathbf{E} field), and \mathbf{G}_ξ is the vacuum Green's function in the medium outside the conductors [9]. For permeable media (finite ε and μ), it turns out that a formulation closely related to the standard PMCHW integral-equation model (see Sec. 7.5.3) can be obtained: \mathbf{J} represents fictitious surface electric and magnetic currents on each interface (derived from the continuity of the tangential \mathbf{E} and \mathbf{H} fields), with \mathbf{G}_ξ again being a homogeneous Green's function (with one Z factor for each contiguous homogeneous region) [125]. Alternatively, because there is a direct correspondence between surface currents and the outgoing/scattered fields from a given interface, "currents" \mathbf{J} can be replaced by scattered fields, again related at different points \mathbf{x} and \mathbf{x}' by the Green's function of the homogeneous medium; this is typically derived directly from a T-matrix formalism [14, 22, 25, 26]. Here, we will focus on the surface-current viewpoint, which is more common in the classical-EM integral-equation community.

The path integral (7.20) is somewhat exotic in classical EM, but it quickly reduces to a manageable expression once an approximate (finite) basis \mathbf{b}_k is chosen for the currents \mathbf{J} . Expanding in this basis, $\mathbf{J} \approx \sum c_k \mathbf{b}_k(\mathbf{x})$ and the functional integral $\mathcal{D}\mathbf{J}$ is replaced by an ordinary integral over the basis coefficients $dc_1 \cdots dc_N$. Equation (7.20) is then a Gaussian integral that can be performed analytically to obtain $Z(\xi) = \# / \sqrt{\det A(\xi)}$ for a proportionality constant $\#$ [9], where $A_{jk} = \int \bar{\mathbf{b}}_j \cdot \mathbf{G}_\xi \cdot \mathbf{b}_k$ is essentially the same as the

¹⁴ Alternatively, the path integral can be performed directly in \mathbf{A} , resulting in an expression equivalent to the sum over energy density in Sec. 7.5 [87] and which in an FD discretization reduces in the same way to repeated solution of the Green's-function diagonal at every point in space [23].

BEM matrix (7.17), albeit here in an arbitrary basis. In the log det of (7.19), proportionality constants and exponents cancel, leaving:

$$U = +\frac{\hbar c}{2\pi} \int_0^\infty \log \det [A_\infty(\xi)^{-1} A(\xi)] d\xi. \quad (7.21)$$

Just as in Sec. 7.17, the use of a real-symmetric positive-definite homogeneous Green's function \mathbf{G}_ξ at imaginary frequencies means that $A(\xi)$ is also real-symmetric and positive-definite, ensuring positive real eigenvalues and hence a real log det. Several further simplifications are possible, even before choosing a particular basis. For example, let \mathbf{p} be the position of some object for which the force \mathbf{F} is desired. The components F_i of the force (in direction p_i) can then be expressed directly as a trace [9, 126]:

$$F_i = -\frac{dU}{dp_i} = -\frac{\hbar c}{2\pi} \int_0^\infty \text{tr} \left(A^{-1} \frac{\partial A}{\partial p_i} \right) d\xi. \quad (7.22)$$

Equivalently, this trace is the sum of eigenvalues λ of the generalized eigenproblem $\frac{\partial A}{\partial p_i} \mathbf{v} = \lambda A \mathbf{v}$; again, these λ are real because A is real-symmetric positive-definite and $\partial A / \partial p_i$ is real-symmetric. (If dense-direct solvers are used, computing $A^{-1} \frac{\partial A}{\partial p_i}$ via Cholesky factorization is much more efficient than computing eigenvalues, however [65].) The matrix A can be further block-decomposed in the usual case where one is computing the interactions among two or more disjoint objects (with disjoint surface currents \mathbf{J}). For example, suppose that one has two objects 1 and 2, in which case one can write

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{12}^T & A_{22} \end{pmatrix}, \quad (7.23)$$

where A_{11} and A_{22} couple currents on each object to other currents on the same object, and A_{12} and $A_{12}^T = A_{21}$ couple currents on object 1 to object 2 and vice versa. In the limit of infinite separation for A_∞ , one obtains $A_{12} \rightarrow 0$ while A_{11} and A_{22} are unchanged, and one can simplify the log det in (7.21) to

$$\log \det [A_\infty(\xi)^{-1} A(\xi)] = \log \det [I - A_{22}^{-1} A_{12}^T A_{11}^{-1} A_{21}]. \quad (7.24)$$

Computationally, only A_{12} depends on the relative positions of the objects, and this simplification immediately allows several computations to be re-used if the energy or force is computed for multiple relative positions.

7.6.1 Monte-Carlo path integration

Before we continue, it should be noted that there also exists a fundamentally different approach for evaluating a path-integral Casimir formulation. Instead of reducing the problem to surface/scattering unknowns and analytically in-

tegrating Z to obtain a matrix log det expression, it is possible to retain the original path-integral expression, in terms of a functional integral over vector potentials \mathbf{A} in the volume, and perform this functional integral numerically via Monte-Carlo methods [127, 128]. This reduces to a Monte-Carlo integration of an action over all possible closed-loop paths (“worldlines”), discretized into some number of points per path. Because this technique is so different from typical classical EM computations, it is difficult to directly compare with the other approaches in this review. Evaluating its computational requirements involves a statistical analysis of the scaling of the necessary number of paths and number of points per path with the desired accuracy and the complexity of the geometry [12], which is not currently available. A difficulty with this technique is that it has currently only been formulated for scalar fields with Dirichlet boundary conditions, not for the true Casimir force of vector electromagnetism.

7.6.2 Spectral methods

One choice of basis functions \mathbf{b}_k for the path-integral expressions above is a spectral basis, and (mirroring the history of integral equations in classical EM) this was the first approach applied in the Casimir problem. With cylindrical objects, for example, the natural spectral basis is a Fourier-series $e^{im\phi}$ in the angular direction ϕ . For planar surfaces the natural choice is a Fourier transform, for spheres it is spherical harmonics $Y_{\ell m}$ (or their vector equivalents [41]), and for spheroids there are spheroidal harmonics [129]. Equivalently, instead of thinking of surface currents expanded in a Fourier-like basis, one can think of the scattered fields from each object expanded in the corresponding Fourier-like basis (e.g. plane, cylindrical, or spherical waves), in which case A relates the incoming to outgoing/scattered waves for each object; this has been called a “scattering-matrix” or “T-matrix” method and is the source of many pioneering results for Casimir interactions of non-planar geometries [14, 22, 25, 26]. Even for nonspherical/spheroidal objects, one can expand the scattered waves in vector spherical harmonics [22], and a variety of numerical techniques have been developed to relate a spherical-harmonic basis to the boundary conditions on nonspherical surfaces [58]. These spectral scattering methods have their roots in many classical techniques for EM scattering problems [56, 122] (See also the chapters by S. J. Rahi et al. and A. Lambrecht et al. in this volume for additional discussions of scattering techniques and Casimir interactions.) Here, we will use the surface-current viewpoint rather than the equivalent scattered-wave viewpoint.

Many simplifications occur in the interaction matrix A of (7.23) for geometries with highly symmetrical objects and a corresponding spectral basis [22]. Consider, for example, the case of spherical objects, with surface currents expressed in a vector spherical-harmonic basis (spherical harmonics for two po-

larizations [22]). In the interaction matrix $A_{jk} = \iint \mathbf{b}_j(\mathbf{x}) \cdot \mathbf{G}(\mathbf{x} - \mathbf{x}') \cdot \mathbf{b}_k(\mathbf{x}')$, the convolution $\int G(\mathbf{x} - \mathbf{x}') \cdot \mathbf{b}_k(\mathbf{x}')$ of a Green's function \mathbf{G} with \mathbf{b}_k is known analytically: it is just the outgoing spherical wave produced by a spherical-harmonic current. If \mathbf{b}_j is another spherical-harmonic current on the same sphere, then the orthogonality of the spherical harmonics means that the \mathbf{x} integral of $\mathbf{b}_j(\mathbf{x})$ against the spherical wave is zero unless $j = k$. Thus, the self-interaction blocks A_{11} and A_{22} of (7.23), with an appropriate normalization, are simply identity matrices. The A_{12} entries are given by the coupling of a spherical wave from \mathbf{b}_k on sphere 2 with a spherical-harmonic basis function \mathbf{b}_j on sphere 1, but again this integral can be expressed analytically, albeit as an infinite series: the spherical wave from sphere 2 can be re-expressed in the basis of spherical waves centered on sphere 1 via known translation identities of spherical waves, and as a result A_{12} takes the form of a “translation matrix” [22]. Furthermore, if there are only two spheres in the problem, then their spherical harmonics can be expressed with respect to a common z axis passing through the centers of the spheres, and a $Y_{\ell m}$ on sphere 1 will only couple with a $Y_{\ell' m'}$ on sphere 2 if $m = m'$, greatly reducing the number of nonzero matrix elements. Related identities are available for coupling cylindrical waves around different origins, expanding spherical/cylindrical waves in terms of planewaves for coupling to planar surfaces, and so on [22].

As was noted in Sec. 7.2.2, such a spectral basis can converge exponentially fast if there are no singularities (e.g. corners) that were not accounted for analytically, and the method can even lend itself to analytical study. Especially for cylinders and spheres, the method is simple to implement and allows rapid exploration of many configurations; the corresponding classical “multipole methods” are common in classical EM for cases where such shapes are of particular interest [5]. On the other hand, as the objects become less and less similar to the “natural” shape for a given basis (e.g. less spherical for spherical harmonics), especially objects with corners or cusps, the spectral basis converges more slowly [58]. Even for the interaction between two spheres or a sphere and a plate, as the two surfaces approach one another the multipole expansion will converge more slowly [25, 130, 140]—conceptually, a spherical-harmonic basis has uniform angular resolution all over the sphere, whereas for two near-touching surfaces one would rather have more resolution in the regions where the surfaces are close (e.g. by using a nonuniform BEM mesh). This exponential convergence of a spectral (spherical harmonic [22]) Casimir calculation is depicted in Fig. 7.4 for the case of the Casimir interaction energy U between two gold spheres of radius $R = 1 \mu\text{m}$, for various surface-to-surface separations a . The error $\Delta U/U$ decreases exponentially with the maximum spherical-harmonic order ℓ [corresponding to $N = 4\ell(\ell+2)$ degrees of freedom for two spheres], but the exponential rate slows as a/R decreases. (On the other hand, for small a/R a perturbative expansion or extrapolation may become applicable [130].)




Fig. 7.4 Fractional error $\Delta U/U$ in the Casimir interaction energy U between two gold spheres of radius $R = 1\mu\text{m}$, for various surface-surface separations a , as a function of the maximum spherical-harmonic order ℓ of the spectral path-integral (scattering-matrix/multipole) method. The error converges exponentially with ℓ , but the exponential rate slows as a/R shrinks. (Calculations thanks to A. Rodriguez.)

7.6.3 Boundary-element methods (BEMs)

In a BEM, one meshes the interfaces, say into triangles, and uses a set of localized basis functions \mathbf{b}_k as discussed in Sec. 7.5.3. In this case, the interaction matrix A that arises in the path-integral formulation is exactly the same as the interaction matrix that arises in classical BEM methods (albeit at an imaginary frequency), and is the same as the matrix A that arises in a BEM stress-tensor approach as described in Sec. 7.5.3. The main difference, compared to the stress-tensor approach, lies in how one *uses* the matrix A : instead of solving a sequence of linear equations to find the mean stress tensor $\langle \mathbf{T} \rangle$ at various points on a surface around an object, one computes $\log \det A$ or $\text{tr} \left[A^{-1} \frac{\partial A}{\partial p_i} \right]$ to obtain the energy (7.21) or force (7.22). We have demonstrated this approach for several three-dimensional geometries, such as the crossed capsules of Fig. 7.3 [9].

If one is using dense-matrix techniques, the advantage of this approach over the stress-tensor technique seems clear [9]: it avoids the complication of picking a stress-integration surface and an appropriate surface-integration technique, and allows the size of the linear system to be easily reduced via blocking as in eq. (7.24). The situation is less clear as one moves to larger and larger problems, in which dense-matrix solvers become impractical and one

requires an iterative method. In that case, computing $\text{tr} \left[A^{-1} \frac{\partial A}{\partial p_i} \right]$ straightforwardly requires N linear systems to be solved; if each linear system can be solved in $O(N \log N)$ time with a fast solver (as discussed in Secs. 7.2.3 and 7.5.3), then the overall complexity is $O(N^2 \log N)$ [with $O(N)$ storage], whereas it is possible that the stress-tensor surface integral may require fewer than N solves. On the other hand, there may be more efficient ways to compute the trace (or log det) via low-rank approximations: for example, if the trace (or log det) is dominated by a small number of extremal eigenvalues, then these eigenvalues can be computed by an iterative method [68] with the equivalent of $\ll N$ linear solves. The real-symmetric property of A , as usual, means that the most favorable iterative methods can be employed, such as a Lanczos or Rayleigh-quotient method [68]. Another possibility might be sparse-direct solvers via a fast-multipole decomposition [73]. The most efficient use of a fast $O(N \log N)$ BEM solver in Casimir problems, whether by stress-tensor or path-integral methods, remains an open question (and the answer may well be problem-dependent).

In the BEM approach with localized basis functions, the $\text{tr} \left[A^{-1} \frac{\partial A}{\partial p_i} \right]$ expression for the force corresponds to a sum of a diagonal components for each surface element, and in the exact limit of infinite resolution (infinitesimal elements) this becomes an integral over the object surfaces. Expressing the force as a surface integral of a quantity related to Green's-function diagonals is obviously reminiscent of the stress-tensor integration from Sec. 7.5.1.3, and it turns out that one can prove an exact equivalence using only vector calculus [125]. (At least one previous author has already shown the algebraic equivalence of the stress tensor and the derivative of the path-integral energy for forces between periodic plates [131].)

7.6.4 Hybrid BEM/spectral methods

It is possible, and sometimes very useful, to employ a hybrid of the BEM and spectral techniques in the previous two sections. One can discretize a surface using boundary elements, and use this discretization to solve for the scattering matrix A_{kk} of each object in a spectral basis such as spherical waves. That is, for any given incident spherical wave, the outgoing field can be computed with BEM via (7.16) and then decomposed into outgoing spherical waves to obtain one row/column of A_{kk} at a time; alternatively, the multipole decomposition of the outgoing wave can be computed directly from the multipole moments of the excited surface currents \mathbf{J}_s [41]. This approach appears to be especially attractive when one has complicated objects, for which a localized BEM basis works well to express the boundary conditions, but the interactions are only to be computed at relatively large separations where the Casimir interaction is dominated by a few low-order multipole moments.




Fig. 7.5 Schematic problem for which eigenmode-expansion is well suited: the interaction between two corrugated surfaces, with period Λ . The Casimir problem reduces to computing the reflection matrices R_1 and R_2 for each individual surface, in a planewave basis. Eigenmode expansion works by expanding the field in each cross-section (dashed lines) in the basis of eigenmodes of a z -invariant structure with that cross-section, and then matching boundary conditions whenever the cross-section changes.

One can perform the BEM computation once per object and re-use the resulting scattering matrix many times via the analytical translation matrices, allowing one to efficiently compute interactions for many rearrangements of the same objects and/or for “dilute” media consisting of many copies of the same objects [132]. (Essentially, this could be viewed as a form of low-rank approximation of the BEM matrix, capturing the essential details relevant to moderate-range Casimir interactions in a much smaller matrix.) Such a hybrid approach is less attractive for closer separations, however, in which the increasing number of relevant multipole moments will eventually lead to an impractically large matrix to be computed.

7.6.5 Eigenmode-expansion/RCWA methods

Consider the case of the interaction between two corrugated surfaces depicted in Fig 7.5, separated in the z direction. From the scattering-matrix viewpoint, it is natural to consider scattering off of each object by planewaves. In this case, the self-interaction matrices A_{11}^{-1} and A_{22}^{-1} can be re-expressed in terms of reflection matrices R_1 and R_2 for each surface, relating the amplitudes of incident waves at some plane (dashed line) above each surface to the reflected (specular and non-specular) planewave amplitudes. The matrices A_{12} and A_{21} are replaced by a diagonal matrix $D_{12} = D_{21}^T$ that relates the planewave amplitudes at the planes for objects 1 and 2, separated by a distance a —at real frequencies, this would be a phase factor, but at imaginary frequencies it is an exponential decay as discussed below. This results in the following expression for the Casimir interaction energy:

$$U = \frac{\hbar c}{2\pi} \int_0^\infty \log \det [I - R_2 D_{12} R_1 D_{12}] d\xi. \quad (7.25)$$

Alternatively, instead of viewing it as a special case of the T-matrix/scattering-matrix idea [22], the same expression can be derived starting from an eigenmode-summation approach [18].

The problem then reduces to computing the scattering of an incident planewave off of a corrugated surface, with the scattered field decomposed into outgoing planewaves. For this problem, one could use any of the tools of computational EM (such as BEM, FD, and so on), but there is a notable method that is often well-suited to the case of periodic surfaces, especially periodic surfaces with piecewise-constant cross-sections¹⁵ (as in object 2 of Fig. 7.5). This method is called *eigenmode expansion* [43–45] or *rigorous coupled-wave analysis* (RCWA) [46, 47], or alternatively a *cross-section* method [48]. RCWA has a long history because it is closely tied to semi-analytical methods to study waveguides with slowly/weakly varying cross-sections [48, 133]. An analogous method was recently applied to Casimir problems [18]. In RCWA, one computes reflection and scattering matrices at a given frequency ω along some direction z by expanding the fields at each z in the basis of the eigenmodes of the cross-section at that z (waves with z dependence $e^{i\beta z}$ at a given ω , where β is called the *propagation constant* of the mode). Along regions of uniform cross-section, the z dependence $e^{i\beta z}$ of each mode is known analytically and no computation is required (the mode amplitudes are multiplied by a diagonal propagation matrix D). Regions of continuously varying cross-section are approximated by breaking them up into a finite number of constant-cross-section layers (as in object 1 of Fig. 7.5). At any z where the cross-section changes, a change of basis is performed by matching boundary conditions (the xy components of the fields must be continuous), yielding a *transfer matrix* at that interface. All these transfer and propagation matrices can then be combined to compute scattering/reflection matrices for an entire structure.

The main difference here from classical RCWA computations is that the modes are computed at imaginary frequencies ξ . As in Sec. 7.5.1.2, this actually simplifies the problem. At an imaginary frequency $\omega = i\xi$, the modes of a given cross-section $\varepsilon(i\xi, x, y)$ and $\mu(i\xi, x, y)$ with z dependence $e^{i\beta z} = e^{-\gamma z}$ ($\gamma = -i\beta$) satisfy the eigenequation (for isotropic materials) [134, 135]:

¹⁵ See also the chapter by A. Lambrecht et al. in this volume for additional discussion of Casimir interactions among periodic structures.

$$\begin{aligned} & \begin{pmatrix} \xi\varepsilon + \nabla_{xy} \times \frac{1}{\xi\mu} \nabla_{xy} \times & \\ & \xi\mu + \nabla_{xy} \times \frac{1}{\xi\varepsilon} \nabla_{xy} \times \end{pmatrix} \begin{pmatrix} \mathbf{E}_{xy} \\ \mathbf{H}_{xy} \end{pmatrix} \\ &= \gamma \begin{pmatrix} & & 1 \\ & -1 & \\ -1 & & \\ 1 & & \end{pmatrix} \begin{pmatrix} \mathbf{E}_{xy} \\ \mathbf{H}_{xy} \end{pmatrix}, \quad (7.26) \end{aligned}$$

where the xy subscript indicates a two-component vector with xy (transverse) components. The operators on both the left- and right-hand sides are real-symmetric, while the operator on the left-hand side is positive-definite, and as a result the eigenvalues γ are purely real. This means that the propagation constants β are purely imaginary (all of the imaginary-frequency modes are evanescent in z), and the analogues of incoming/outgoing waves are those that are exponentially decaying towards/away from the surface. Moreover, the numerical problem of solving for these eigenmodes in a given cross-section reduces to a positive-definite generalized eigenvalue problem (a definite matrix pencil [69]), to which the most desirable numerical solvers apply [68,69] (unlike the classical real- ω problem in which there are both propagating and evanescent modes because the problem is indefinite [134,135]). For homogeneous cross-sections (as in the space between the two objects), the solutions are simply planewaves of the form $e^{ik_x x + ik_y y - \gamma z + \xi t}$, where for vacuum $\gamma = \pm \sqrt{|\mathbf{k}_{xy}|^2 + \xi^2/c^2}$.

For sufficiently simple cross-sections, especially in two-dimensional or axisymmetric geometries, it is possible to solve for the modes analytically and hence obtain the scattering matrices, and this is how the technique was first applied to the Casimir problem [18]. For more general geometries, one can solve for the modes numerically by a variety of techniques, such as by a transfer-matrix method in two dimensions [43] or by a planewave expansion (in the xy cross-section) in three dimensions [46]. Of course, one truncates to a finite number of modes via some cutoff $|\gamma|$ (which follows automatically from discretizing the cross-section in a finite grid, for example), and convergence is obtained in the limit as this cutoff increases. Given a basis of eigenmodes with some cutoff, the process of constructing the scattering/reflection matrices is thoroughly discussed elsewhere [43–47], so we do not review it here.

The strength of RCWA is that regions of uniform cross-section are handled with at most a 2d discretization of the cross-section, independent of the thickness of the region, so very thick or very thin layers can be solved efficiently. The main limitation of RCWA methods is that the transfer matrices (and the resulting reflection matrices R_1 and R_2) are dense $N \times N$ matrices, where N is the number of modes required for convergence. If N is large, as in complicated three-dimensional structures, the problem can quickly become impractical because of the $O(N^2)$ storage and $O(N^3)$ computation requirements. The most favorable case is that of periodic structures with relatively simple unit cells, in which case the problem can be reduced to that of computing the modes of each periodic unit cell (with Bloch-periodic boundary

conditions) as discussed below, and RCWA can then be quite practical even in three dimensions. Non-periodic structures, such as compact objects, can be handled by perfectly matched layer (PML) absorbing boundaries [44], albeit at greater computational cost because of the increased cross-section size.

7.7 Periodicity and other symmetries

In this section, we briefly discuss the issue of periodicity and other symmetries, which can be exploited to greatly reduce the computational effort in Casimir calculations just as for classical EM calculations.

If a structure is periodic in the x direction with period Λ , as in Fig. 7.5, the problem simplifies considerably because one can reduce the computation to a single unit cell of thickness Λ . In particular, one imposes Bloch-periodic boundary conditions—the fields at $x = \Lambda$ equal the fields at $x = 0$ multiplied by a phase factor $e^{ik_x\Lambda}$ —and computes the Casimir energy or force for each Bloch wavevector k_x separately, then integrates the result over k_x via $\int_{-\pi/\Lambda}^{\pi/\Lambda}(\dots)dk_x$. This can be derived in a variety of ways, for example by applying Bloch’s theorem [6] to decompose the eigenmodes into Bloch-wave solutions for each k_x , or by expanding the delta functions of the fluctuation–dissipation approach in a Fourier series [12]. More generally, for any periodic unit cell, one can perform the Casimir energy/force computation for the unit cell with Bloch periodic boundaries and then integrate the Bloch wavevector \mathbf{k} over the irreducible Brillouin zone (multiplied by the volume ratio of the Brillouin zone and the irreducible Brillouin zone).

The specific case of continuous translational symmetry, say in the x direction, corresponds $\Lambda \rightarrow 0$ and one must integrate over all k_x (the Brillouin zone is infinite). Certain additional simplifications apply in the case of a perfect-metal structure with continuous translational symmetry, in which case the fields decompose into two polarizations and the k integration can be performed implicitly [12].

Rotational symmetry can be handled similarly: the fields can be decomposed into fields with $e^{im\phi}$ angular dependence, and the total force or energy is the sum over all integers m of the contributions for each m [27]. More generally, the Casimir contributions can be decomposed into a sum of contributions from irreducible representations of the symmetry group of the structure (e.g. all eigenmodes can be classified into these representations [136, 141]); translational and rotational symmetries are merely special cases. As another example, in a structure with a mirror symmetry one could sum even- and odd-symmetry contributions (in fact, this is the underlying reason for the TE/TM polarization decomposition in two dimensions [6]).

7.8 Nonzero-temperature corrections

In the preceding sections, we discussed only the computation of Casimir interactions at zero temperature $T = 0^+$. However, the modification of any imaginary-frequency expression for a Casimir interaction from $T = 0$ to $T > 0$ is almost trivial: one simply performs a sum instead of an integral. If the $T = 0$ interaction (energy, force, etc.) is expressed as an integral $\int_0^\infty C(\xi) d\xi$ of some contributions $C(\xi)$ at each imaginary frequency ξ , then the $T > 0$ interaction is well known to be simply [81]:

$$\int_0^\infty C(\xi) d\xi \rightarrow \frac{2\pi k_B T}{\hbar} \sum'_{n=0} C\left(\frac{2\pi k_B T}{\hbar} n\right), \quad (7.27)$$

where k_B is Boltzmann's constant and \sum' indicates a sum with weight $\frac{1}{2}$ for the $n = 0$ term. The frequencies $\xi_n = 2\pi k_B T n / \hbar$ are known as Matsubara frequencies, and the corresponding (imaginary) Matsubara wavelengths are $\lambda_n = 2\pi / \xi_n = \lambda_T / n$ where $\lambda_T = \hbar / k_B T$. The conversion of the $T = 0$ integral into a summation can be derived in a variety of ways, most directly by considering thermodynamics in the Matsubara formalism [81]. Physically, this arises from the $\coth(\hbar\omega/2kT)$ Bose–Einstein distribution factor that appears in the fluctuation–dissipation expressions (7.3) for nonzero temperatures. When the contour integration is performed over ω , the \coth introduces poles at $\hbar\omega/2kT = i\pi n$ that convert the integral into a sum via the residue theorem (with the $n = 0$ residue having half weight because it lies on the real- ω axis) [77]. As explained in Sec. 7.5.1.2, some care must be applied in evaluating the $n = 0$ term because of the well known singularity of Maxwell's equations at $\omega = 0$ (where the \mathbf{E} and \mathbf{H} fields decouple), and one may need to take the limit $\xi \rightarrow 0^+$ (although there is some controversy in the unphysical case of perfect metals [92])

Mathematically, the sum of eq. (7.27) is exactly the same as a trapezoidal-rule approximation for the $T = 0$ integral, with equally spaced abscissas $\Delta\xi = 2\pi/\lambda_T$ [53, 91, 137]. Thanks to the $O(\Delta\xi^2)$ convergence of the trapezoidal rule [53], this means that the $T > 0$ result is quite close to the $T = 0$ result unless $C(\xi)$ varies rapidly on the scale of $2\pi/\lambda_T$. In particular, suppose that $C(\xi)$ varies on a scale $2\pi/a$, corresponding to some lengthscale a in the problem (typically from a surface–surface separation). In that case, assuming $C(\xi)$ has nonzero slope¹⁶ at $\xi = 0^+$ (typical for interactions between realistic metal surfaces), then the nonzero- T correction should be of order $O(a^2/\lambda_T^2)$. At room temperature ($T = 300$ K), $\lambda_T \approx 7.6 \mu\text{m}$, and the temperature corrections to Casimir interactions are typically negligible for submicron separations [81, 138]. On the other hand, it is possible that careful material and geometry choices may lead to larger temperature effects [137].

¹⁶ If $C(\xi)$ has zero slope at $\xi = 0^+$, then the trapezoidal rule differs from the integral by $O(\Delta\xi^4)$ or less, depending upon which derivative is nonzero at $\xi = 0^+$ [53].

There is also the possibility of interesting effects in nonequilibrium situations (objects at different temperatures) [97, 139], but such situations are beyond the scope of this review.

7.9 Concluding remarks

The area of numerical Casimir computations remains rich with opportunities. Relatively few geometry and material combinations have as yet been explored, and thus many newly answerable questions remain regarding the ways in which Casimir phenomena can be modified by exploiting the degrees of freedom available in modern nanofabrication. In the regime of computational techniques, while several effective methods have already been proposed and demonstrated, the parallels with computational electromagnetism lead us to anticipate ongoing improvements and developments for some time to come. The same parallels also caution against any absolute “rankings” of the different approaches, as different numerical techniques have always exhibited unique strengths and weaknesses in both theory and practice. And because computer time is typically much less expensive than programmer time, there is something to be said for methods that may be theoretically suboptimal but are easy to implement (or are available off-the-shelf) for very general geometries and materials. Nor is the value of analytical and semi-analytical techniques diminished, but rather these approaches are freed from the tedium of hand computation to focus on more fundamental questions.

Acknowledgements

This work was supported in part by the Army Research Office through the ISN under contract W911NF-07-D-0004, by the MIT Ferry Fund, and by the Defense Advanced Research Projects Agency (DARPA) under contract N66001-09-1-2070-DOD. We are especially grateful to our students, A. W. Rodriguez, A. P. McCauley, and H. Reid for their creativity and energy in pursuing Casimir simulations. We are also grateful to our colleagues F. Capasso, D. Dalvit, T. Emig, R. Jaffe, J. D. Joannopoulos, M. Kardar, M. Levin, M. Lončar, J. Munday, S. J. Rahi, and J. White, for their many suggestions over the years.

References

1. Chew, W.C., Jian-Ming, J., Michielssen, E., Jiming, S.: *Fast and Efficient Algorithms in Computational Electromagnetics*. Artech, Norwood, MA (2001)
2. Taflov, A., Hagness, S.C.: *Computational Electrodynamics: The Finite-Difference Time-Domain Method*. Artech, Norwood, MA (2000)
3. Volakis, J.L., Chatterjee, A., Kempel, L.C.: *Finite Element Method Electromagnetics: Antennas, Microwave Circuits, and Scattering Applications*. IEEE Press, New York (2001)
4. Zhu, Y., Cangellaris, A.C.: *Multigrid Finite Element Methods for Electromagnetic Field Modelling*. John Wiley and Sons, Hoboken, NJ (2006)
5. Yasumoto, K. (ed.): *Electromagnetic Theory and Applications for Photonic Crystals*. CRC Press (2005)
6. Joannopoulos, J.D., Johnson, S.G., Winn, J.N., Meade, R.D.: *Photonic Crystals: Molding the Flow of Light*, 2nd edn. Princeton University Press (2008)
7. Jin, J.: *The Finite Element Method in Electromagnetics*, 2nd edn. Wiley, New York (2002)
8. Rao, S.M., Balakrishnan, N.: *Computational electromagnetics*. *Current Science* **77**(10), 1343–1347 (1999)
9. Reid, M.T.H., Rodriguez, A.W., White, J., Johnson, S.G.: Efficient computation of three-dimensional Casimir forces. *Phys. Rev. Lett.* **103**(4), 040401 (2009)
10. Reynaud, S., Maia Neto, P.A., Lambrecht, A.: Casimir energy and geometry: beyond the proximity force approximation. *J. Phys. A: Math. Theor.* **41**, 164004 (2008)
11. Rodriguez, A., Ibanescu, M., Iannuzzi, D., Capasso, F., Joannopoulos, J.D., Johnson, S.G.: 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)
12. Rodriguez, A., Ibanescu, M., Iannuzzi, D., Joannopoulos, J.D., Johnson, S.G.: Virtual photons in imaginary time: computing Casimir forces in arbitrary geometries via standard numerical electromagnetism. *Phys. Rev. A* **76**(3), 032106 (2007)
13. Rodriguez, A.W., McCauley, A.P., Joannopoulos, J.D., Johnson, S.G.: Casimir forces in the time domain: Theory. *Phys. Rev. A* **80**(1), 012115 (2009)
14. Emig, T., Graham, N., Jaffe, R.L., Kardar, M.: Casimir forces between arbitrary compact objects. *Phys. Rev. Lett.* **99**, 170403 (2007)
15. Emig, T., Hanke, A., Golestanian, R., Kardar, M.: Probing the strong boundary shape dependence of the Casimir force. *Phys. Rev. Lett.* **87**, 260402 (2001)
16. Emig, T., Jaffe, R.L., Kardar, M., Scardicchio, A.: Casimir interaction between a plate and a cylinder. *Phys. Rev. Lett.* **96**, 080403 (2006)
17. Lambrecht, A., Maia Neto, P.A., Reynaud, S.: The Casimir effect within scattering theory. *New Journal of Physics* **8**(243) (2008)
18. Lambrecht, A., Marachevsky, V.I.: Casimir interactions of dielectric gratings. *Phys. Rev. Lett.* **101**, 160403 (2008)
19. Xiong, J.L., Chew, W.C.: Efficient evaluation of Casimir force in z -invariant geometries by integral equation methods. *Appl. Phys. Lett.* **95**, 154102 (2009)
20. Xiong, J.L., Tong, M.S., Atkins, P., Chew, W.C.: Efficient evaluation of Casimir force in arbitrary three-dimensional geometries by integral equation methods. *Phys. Lett. A* **374**(25) 2517–2520 (2010)
21. Rahi, S.J., Emig, T., Jaffe, R.L., Kardar, M.: Casimir forces between cylinders and plates. *Phys. Rev. A* **78**, 012104 (2008)
22. Rahi, S.J., Emig, T., Jaffe, R.L., Kardar, M.: Scattering theory approach to electrodynamic Casimir forces. *Phys. Rev. D* **80**, 085021 (2009)
23. Pasquali, S., Maggs, A.C.: Numerical studies of Lifshitz interactions between dielectrics. *Phys. Rev. A. (R)* **79**, 020102(R) (2009)

24. Maia Neto, P.A., Lambrecht, A., Reynaud, S.: Roughness correction to the Casimir force: Beyond the proximity force approximation. *Europhys. Lett.* **69**, 924–930 (2005)
25. Maia Neto, P.A., Lambrecht, A., Reynaud, S.: Casimir energy between a plane and a sphere in electromagnetic vacuum. *Phys. Rev. A* **78**, 012,115 (2008)
26. Kenneth, O., Klich, I.: Casimir forces in a T-operator approach. *Phys. Rev. B* **78**, 014,103 (2008)
27. McCauley, A.P., Rodriguez, A.W., Joannopoulos, J.D., Johnson, S.G.: Casimir forces in the time domain: Applications. *Phys. Rev. A* **81**, 012,119 (2010)
28. Casimir, H.B.G.: On the attraction between two perfectly conducting plates. *Proc. K. Ned. Akad. Wet.* **51**, 793–795 (1948)
29. Lamoreaux, S.K.: Demonstration of the Casimir force in the 0.6 to 6 μ m range. *Phys. Rev. Lett.* **78**, 5–8 (1997)
30. Derjaguin, B.V., Abrikosova, I.I., Lifshitz, E.M.: Direct measurement of molecular attraction between solids separated by a narrow gap. *Q. Rev. Chem. Soc.* **10**, 295–329 (1956)
31. Bordag, M.: Casimir effect for a sphere and a cylinder in front of a plane and corrections to the proximity force theorem. *Phys. Rev. D* **73**, 125,018 (2006)
32. Bordag, M., Mohideen, U., Mostepanenko, V.M.: New developments in the Casimir effect. *Phys. Rep.* **353**, 1–205 (2001)
33. Casimir, H.B.G., Polder, D.: The influence of retardation on the London-Van der Waals forces. *Phys. Rev.* **13**(4), 360–372 (1948)
34. Sedmik, R., Vasiljevich, I., Tajmar, M.: Detailed parametric study of Casimir forces in the Casimir Polder approximation for nontrivial 3d geometries. *J. Computer-Aided Materials Design* **14**(1), 119–132 (2007)
35. Tajmar, M.: Finite element simulation of Casimir forces in arbitrary geometries. *Intl. J. Mod. Phys. C* **15**(10), 1387–1395 (2004)
36. Jaffe, R.L., Scardicchio, A.: Casimir effect and geometric optics. *Phys. Rev. Lett.* **92**, 070,402 (2004)
37. Hertzberg, M.P., Jaffe, R.L., M., K., Scardicchio, A.: Casimir forces in a piston geometry at zero and finite temperatures. *Phys. Rev. D* **76**, 045,016 (2007)
38. Zaheer, S., Rodriguez, A.W., Johnson, S.G., Jaffe, R.L.: Optical-approximation analysis of sidewall-spacing effects on the force between two squares with parallel sidewalls. *Phys. Rev. A* **76**(6), 063,816 (2007)
39. Scardicchio, A., Jaffe, R.L.: Casimir effects: An optical approach I. foundations and examples. *Nuclear Physics B.* **704**(3), 552–582 (2005)
40. Van Enk, S.J.: The Casimir effect in dielectrics: A numerical approach. *Journal of Modern Optics* **42**(2), 321–338 (1995)
41. Jackson, J.D.: *Classical Electrodynamics*, 3rd edn. Wiley, New York (1998)
42. Schaubert, D.H., Wilton, D.R., Glisson, A.W.: A tetrahedral modeling method for electromagnetic scattering by arbitrarily shaped inhomogeneous dielectric bodies. *IEEE Trans. Antennas Propagat.* **32**, 77–85 (1984)
43. Bienstman, P., Baets, R.: Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers. *Optical and Quantum Electron.* **33**(4–5), 327–341 (2001)
44. Bienstman, P., Baets, R.: Advanced boundary conditions for eigenmode expansion models. *Optical and Quantum Electron.* **34**, 523–540 (2002)
45. Willems, J., Haes, J., Baets, R.: The bidirectional mode expansion method for two dimensional waveguides. *Optical and Quantum Electron.* **27**(10), 995–1007 (1995)
46. Moharam, M.G., Grann, E.B., Pommet, D.A., Gaylord, T.K.: Formulation for stable and efficient implementation of the rigorous coupled-wave analysis of binary gratings. *J. Opt. Soc. Am. A* **12**, 1068–1076 (1995)
47. Moharam, M.G., Pommet, D.A., Grann, E.B., Gaylord, T.K.: Stable implementation of the rigorous coupled-wave analysis for surface-relief gratings: enhanced transmittance matrix approach. *J. Opt. Soc. Am. A* **12**, 1077–1086 (1995)

48. Katsenelenbaum, B.Z., Mercader del Río, L., Pereyaslavets, M., Sorolla Ayzá, M., Thumm, M.: *Theory of Nonuniform Waveguides: The Cross-Section Method*. Inst. of Electrical Engineers, London (1998)
49. Christ, A., Hartnagel, H.L.: Three-dimensional finite-difference method for the analysis of microwave-device embedding. *IEEE Trans. Microwave Theory Tech.* **35**(8), 688–696 (1987)
50. Strikwerda, J.: *Finite Difference Schemes and Partial Differential Equations*. Wadsworth and Brooks/Cole, Pacific Grove, CA (1989)
51. Bonnet, M.: *Boundary Integral Equation Methods for Solids and Fluids*. Wiley, Chichester, England (1999)
52. Hackbusch, W., Verlag, B.: *Integral Equations: Theory and Numerical Treatment*. Birkhauser Verlag, Basel, Switzerland (1995)
53. Boyd, J.P.: *Chebyshev and Fourier Spectral Methods*, 2nd edn. Dover, New York (2001)
54. Ditkowski, A., Dridi, K., Hesthaven, J.S.: Convergent Cartesian grid methods for Maxwell’s equations in complex geometries. *J. Comp. Phys.* **170**, 39–80 (2001)
55. Oskooi, A.F., Kottke, C., Johnson, S.G.: Accurate finite-difference time-domain simulation of anisotropic media by subpixel smoothing. *Optics Letters* **34**, 2778–2780 (2009)
56. Stratton, J.A.: *Electromagnetic Theory*. McGraw-Hill, New York (1941)
57. Johnson, S.G., Joannopoulos, J.D.: Block-iterative frequency-domain methods for Maxwell’s equations in a planewave basis. *Opt. Express* **8**(3), 173–190 (2001)
58. Kuo, S.H., Tidor, B., White, J.: A meshless, spectrally accurate, integral equation solver for molecular surface electrostatics. *ACM J. Emerging Technologies in Computer Systems* **4**(2), 1–30 (2008)
59. Sladek, V., Sladek, J.: *Singular Integrals in Boundary Element Methods*. WIT Press, Southampton, UK (1998)
60. Taylor, D.J.: Accurate and efficient numerical integration of weakly singular integrals in Galerkin EFIE solutions. *IEEE Trans. Antennas Propagat.* **51**, 1630–1637 (2003)
61. Tong, M.S., Chew, W.C.: Super-hyper singularity treatment for solving 3d electric field integral equations. *Microwave Opt. Tech. Lett.* **49**, 1383–1388 (2006)
62. Nédélec, J.C.: Mixed finite elements in \mathbb{R}^3 . *Numerische Mathematik* **35**, 315–341 (1980)
63. Rao, S.M., Wilton, D.R., Glisson, A.W.: Electromagnetic scattering by surfaces of arbitrary shape. *IEEE Trans. Antennas Propagat.* **30**, 409–418 (1982)
64. Cai, W., Yu, Y., Yuan, X.C.: Singularity treatment and high-order RWG basis functions for integral equations of electromagnetic scattering. *Intl. J. Numer. Meth. Eng.* **53**, 31–47 (2002)
65. Trefethen, L.N., Bau, D.: *Numerical Linear Algebra*. SIAM, Philadelphia (1997)
66. Davis, T.A.: *Direct Methods for Sparse Linear Systems*. SIAM, Philadelphia (2006)
67. Barrett, R., Berry, M., Chan, T.F., Demmel, J., Donato, J., Dongarra, j., Eijkhout, V., Pozo, R., Romine, C., Van der Vorst, H.: *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, 2nd edn. SIAM, Philadelphia (1994)
68. Bai, Z., Demmel, J., Dongarra, J., Ruhe, A., Van Der Vorst, H.: *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, Philadelphia (2000)
69. Anderson, E., Bai, Z., Bischof, C., Blackford, S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., Sorensen, D.: *LAPACK Users’ Guide*, 3rd edn. SIAM, Philadelphia (1999)
70. Duff, I.S., Erisman, A.M., Reid, J.K.: On George’s nested dissection method. *SIAM J. Numer. Anal.* **13**, 686–695 (1976)
71. Phillips, J.R., White, J.K.: A precorrected-FFT method for electrostatic analysis of complicated 3-D structures. *IEEE Trans. Computer-Aided Design* **16**, 1059–1072 (1997)

72. Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *J. Comp. Phys.* **73**, 325–348 (1987)
73. Greengard, L., Gueyffier, D., Martinsson, P.G., Rokhlin, V.: Fast direct solvers for integral equations in complex three-dimensional domains. *Acta Numerica* **18**, 243–275 (2009)
74. Milonni, P.W.: *The Quantum Vacuum: An Introduction to Quantum Electrodynamics*. Academic Press, San Diego (1993)
75. Cohen-Tannoudji, C., Din, B., Laloë, F.: *Quantum Mechanics*. Hermann, Paris (1977)
76. Strang, G.: *Computational Science and Engineering*. Wellesley-Cambridge Press, Wellesley, MA (2007)
77. Lamoreaux, S.K.: The Casimir force: background, experiments, and applications. *Rep. Prog. Phys.* **68**(201–236) (2005)
78. Ford, L.H.: Spectrum of the Casimir effect and the Lifshitz theory. *Phys. Rev. A* **48**, 2962–2967 (1993)
79. Nesterenko, V.V., Pirozhenko, I.G.: Simple method for calculating the Casimir energy for a sphere. *Phys. Rev. D* **57**, 1284–1290 (1998)
80. Cognola, G., Elizalde, E., Kirsten, K.: Casimir energies for spherically symmetric cavities. *J. Phys. A* **34**, 7311–7327 (2001)
81. Lifshitz, E.M., Pitaevskii, L.P.: *Statistical Physics: Part 2*. Pergamon, Oxford (1980)
82. Boyd, J.P.: Exponentially convergent Fourier–Chebyshev quadrature schemes on bounded and infinite intervals. *J. Scientific Computing* **2**, 99–109 (1987)
83. Stroud, A.H., Secrest, D.: *Gaussian Quadrature Formulas*. Prentice-Hall, Englewood Cliffs, NJ (1966)
84. Piessens, R., de Doncker-Kapenga, E., Uberhuber, C., Hahaner, D.: *QUADPACK: A Subroutine Package for Automatic Integration*. Springer-Verlag, Berlin (1983)
85. Gel'fand, I.M., Shilov, G.E.: *Generalized Functions*. Academic Press, New York (1964)
86. Buhmann, S.Y., Scheel, S.: Macroscopic quantum electrodynamics and duality. *Phys. Rev. Lett.* **102**, 140,404 (2009)
87. Milton, K.A., Wagner, J., Parashar, P., Brevek, I.: Casimir energy, dispersion, and the Lifshitz formula. *Phys. Rev. D* **81** 065,007 (2010)
88. Economou, E.N.: *Green's Functions in Quantum Physics*, 3rd edn. Springer, Heidelberg (2006)
89. Zhao, J.S., Chew, W.C.: Integral equation solution of Maxwell's equations from zero frequency to microwave frequencies. *IEEE Trans. Antennas Propagat.* **48**, 1635–1645 (2000)
90. Epstein, C.L., Greengard, L.: Debye sources and the numerical solution of the time harmonic Maxwell equations. *Commun. Pure Appl. Math.* **63**, 413–463 (2009)
91. Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edn. Cambridge Univ. Press (1992)
92. Hoye, J.S., Brevik, I., Aarseth, J.B., Milton, K.A.: What is the temperature dependence of the Casimir effect? *J. Phys. A: Math. Gen.* **39**(20), 6031–6038 (2006)
93. Rodriguez, A.W., Munday, J., Davlit, D., Capasso, F., Joannopoulos, J.D., Johnson, S.G.: Stable suspension and dispersion-induced transition from repulsive Casimir forces between fluid-separated eccentric cylinders. *Phys. Rev. Lett.* **101**(19), 190,404 (2008)
94. Pitaevskii, L.P.: Comment on "Casimir force acting on magnetodielectric bodies in embedded in media". *Phys. Rev. A* **73**, 047,801 (2006)
95. Munday, J., Capasso, F., Parsegian, V.A.: Measured long-range repulsive casimir-lifshitz forces. *Nature* **457**, 170–173 (2009)
96. Munday, J.N., Capasso, F.: Precision measurement of the Casimir-Lifshitz force in a fluid. *Phys. Rev. A* **75**, 060,102(R) (2007)
97. Najafi, A., Golestanian, R.: Forces induced by nonequilibrium fluctuations: The Soret–Casimir effect. *Europhys. Lett.* **68**, 776–782 (2004)

98. Chew, W.C., Weedon, W.H.: A 3d perfectly matched medium from modified Maxwell's equations with stretched coordinates. *Microwave and Optical Tech. Lett.* **7**(13), 599–604 (1994)
99. Zhao, L., Cangellaris, A.C.: A general approach for the development of unsplit-field time-domain implementations of perfectly matched layers for FDTD grid truncation. *IEEE Microwave and Guided Wave Lett.* **6**(5), 209–211 (1996)
100. Teixeira, F.L., Chew, W.C.: General closed-form PML constitutive tensors to match arbitrary bianisotropic and dispersive linear media. *IEEE Microwave and Guided Wave Lett.* **8**(6), 223–225 (1998)
101. Ward, A.J., Pendry, J.B.: Refraction and geometry in Maxwell's equations. *J. Mod. Opt.* **43**(4), 773–793 (1996)
102. Farjadpour, A., Roundy, D., Rodriguez, A., Ibanescu, M., Bermel, P., Joannopoulos, J.D., Johnson, S.G., Burr, G.: Improving accuracy by subpixel smoothing in FDTD. *Optics Letters* **31**, 2972–2974 (2006)
103. Kottke, C., Farjadpour, A., Johnson, S.G.: Perturbation theory for anisotropic dielectric interfaces, and application to sub-pixel smoothing of discretized numerical methods. *Phys. Rev. E* **77**, 036,611 (2008)
104. Golub, G.H., Loan, C.F.V.: *Matrix Computations*, 3rd edn. Johns Hopkins Univ. Press, Baltimore, MD (1996)
105. Trottenberg, U., Oosterlee, C., Schüller, A.: *Multigrid*. Academic Press, London (2001)
106. Rengarajan, S.R., Rahmat-Samii, Y.: The field equivalence principle: Illustration of the establishment of the non-intuitive null fields. *IEEE Antennas Propagation Magazine* **42**, 122–128 (2000)
107. Schelkunoff, S.A.: Some equivalence theorems of electromagnetic waves. *Bell Syst. Tech. J.* **15**, 92–112 (1936)
108. Stratton, J.A., Chu, L.J.: Diffraction theory of electromagnetic waves. *Phys. Rev.* **56**, 99–107 (1939)
109. Love, A.E.H.: The integration of equations of propagation of electric waves. *Phil. Trans. Roy. Soc. London A* **197**, 1–45 (1901)
110. Harrington, R.F.: Boundary integral formulations for homogeneous material bodies. *J. Electromagnetic Waves Appl.* **3**, 1–15 (1989)
111. Umashankar, K., Taflov, A., Rao, S.: Electromagnetic scattering by arbitrary shaped three-dimensional homogeneous lossy dielectric objects. *IEEE Trans. Antennas Propagat.* **34**, 758–766 (1986)
112. Medgyesi-Mitschang, L.N., Putnam, J.M., Gedera, M.B.: Generalized method of moments for three-dimensional penetrable scatterers. *J. Opt. Soc. Am. A* **11**, 1383–1398 (1994)
113. Trefethen, L.N.: Is Gauss quadrature better than clenshaw–curtis. *SIAM Review* **50**, 67–87 (2008)
114. Cools, R.: Advances in multidimensional integration. *J. Comput. Appl. Math* **149**, 1–12 (2002)
115. Cools, R.: An encyclopaedia of cubature formulas. *J. Complexity* **19**, 445–453 (2003)
116. Berntsen, J., Espelid, T.O., Genz, A.: An adaptive algorithm for the approximate calculation of multiple integrals. *ACM Trans. Math. Soft.* **17**, 437–451 (1991)
117. Atkinson, K., Sommariva, A.: Quadrature over the sphere. *Elec. Trans. Num. Anal.* **20**, 104–118 (2005)
118. Le Gia, Q.T., Mhaskar, H.N.: Localized linear polynomial operators and quadrature formulas on the sphere. *SIAM J. Num. Anal.* **47**, 440–466 (2008)
119. Oskooi, A.F., Roundy, D., Ibanescu, M., Bermel, P., Joannopoulos, J.D., Johnson, S.G.: MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method. *Comp. Phys. Comm.* **181**, 687–702 (2010)

120. Rodriguez, A.W., McCauley, A.P., Joannopoulos, J.D., Johnson, S.G.: Theoretical ingredients of a Casimir analog computer. *Proc. Nat. Acad. Sciences* **107**, 9531–9536 (2010)
121. McCauley, A.P., Rodriguez, A.W., Johnson, S.G.: Casimir Meep wiki. <http://ab-initio.mit.edu/meep/casimir>
122. Waterman, P.C.: The T -matrix revisited. *J. Opt. Soc. Am. A* **24**(8), 2257–2267 (2007)
123. Werner, G.R., Cary, J.R.: A stable FDTD algorithm for non-diagonal anisotropic dielectrics. *J. Comp. Phys.* **226**, 1085–1101 (2007)
124. Li, H., Kardar, M.: Fluctuation-induced forces between rough surfaces. *Phys. Rev. Lett.* **67**, 3275–3278 (1991)
125. Reid, H.: Casimir interactions among dielectric media: An explicitly surface-only formulation (2009). Manuscript in preparation.
126. Emig, T.: Casimir forces: An exact approach for periodically deformed objects. *Europhys. Lett.* **62**, 466 (2003)
127. Gies, H., Klingmuller, K.: Worldline algorithms for Casimir configurations. *Phys. Rev. D* **74**, 045,002 (2006)
128. Gies, H., Langfeld, K., Moyaerts, L.: Casimir effect on the worldline. *J. High Energy Phys.* **6**, 018 (2003)
129. Emig, T., Graham, N., Jaffe, R.L., Kardar, M.: Orientation dependence of casimir forces. *Phys. Rev. A* **79**, 054,901 (2009)
130. Emig, T.: Fluctuation-induced quantum interactions between compact objects and a plane mirror. *J. Stat. Mech.* p. P04007 (2008)
131. Bimonte, G.: Scattering approach to Casimir forces and radiative heat transfer for nanostructured surfaces out of thermal equilibrium. *Phys. Rev. A* **80**, 042,102 (2009)
132. McCauley, A.P., Zhao, R., Reid, M.T.H., Rodriguez, A.W., Zhao, J., Rosa, F.S.S., Joannopoulos, J.D., Dalvit, D.A.R., Soukoulis, C.M., Johnson, S.G.: Microstructure effects for Casimir forces in chiral metamaterials. arXiv:1006.5489 (2010)
133. Marcuse, D.: *Theory of Dielectric Optical Waveguides*, 2nd edn. Academic Press, San Diego (1991)
134. Johnson, S.G., Bienstman, P., Skorobogatiy, M., Ibanescu, M., Lidorikis, E., Joannopoulos, J.D.: Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals. *Phys. Rev. E* **66**, 066,608 (2002)
135. Skorobogatiy, M., Yang, J.: *Fundamentals of Photonic Crystal Guiding*. Cambridge University Press (2009)
136. Inui, T., Tanabe, Y., Onodera, Y.: *Group Theory and Its Applications in Physics*. Springer, Heidelberg (1996)
137. Rodriguez, A.W., Woolf, D., McCauley, A.P., Capasso, F., Joannopoulos, J.D., Johnson, S.G.: Achieving a strongly temperature-dependent Casimir effect. *Phys. Rev. Lett.* **105** 060,401 (2010).
138. Milton, K.A.: The Casimir effect: recent controversies and progress. *Journal of Physics A: Mathematical and General* **37**, R209–R277 (2004)
139. Obrecht, J.M., Wild, R.J., Antezza, M., Pitaevskii, L.P., Stringari, S., Cornell, E.A.: Measurement of the temperature dependence of the Casimir-Polder force. *Phys. Rev. Lett.* **98**(6), 063,201 (2007)
140. Emig, T., Jaffe, R.L.: Casimir forces between arbitrary compact objects. *J. Phys. A* **41**, 164,001 (2008)
141. Tinkham, M.: *Group Theory and Quantum Mechanics*. Dover, New York (2003)
142. Zhao, S.: High order matched interface and boundary methods for the Helmholtz equation in media with arbitrarily curved interfaces. *J. Comp. Phys.* **229**, 3155–3170 (2010)