Geometric Manipulation of Light: From Nonlinear Optics to Invisibility Cloaks

by

Hila Hashemi

B.S., University of California, Berkeley

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2012

Author .....
Department of Mathematics
January 9, 2012

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

Accepted by .....
Michel Goemans
Leighton Family Professor of Applied Mathematics

Geometric Manipulation of Light:
From Nonlinear Optics to Invisibility Cloaks

by

Hila Hashemi

Submitted to the Department of Mathematics
on January 9, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this work, we study two different manipulations of electromagnetic waves governed by macroscopic Maxwell's equations. One is frequency conversion of such waves using small intrinsic material nonlinearities. We study conversion of an input signal at frequency ω_1 to frequency ω_k due to second or third harmonic generation or four-wave mixing using coupled-mode theory. Using this framework, we show there is a critical input power at which maximum frequency conversion is possible. We study in depth the case of third harmonic generation, its solutions, and their stability analysis. Based on the dynamics of the system, we propose a regime of parameters that 100%-efficient frequency conversion is possible and propose a way of exciting this solution. We also look at same analysis for the case of degenerate four-wave mixing and come up with 2d and 3d designs of a device that exhibits high-efficiency second-harmonic generation.

Second, we consider proposals for invisibility cloaks to change the path of electromagnetic waves in a certain way so that the object appears invisible at a certain frequency or a range of frequencies. Transformation-based invisibility cloaks make use of the coordinate invariance of Maxwell's Equations and require complex material configuration ε and μ in the cloak. We study the practical limitations of cloaking as a function of the size of the object being cloaked. Specifically, we study the bandwidth, loss, and scattering limitations of cloaking as the object gets larger and show that cloaking of objects many times larger than the wavelength in size becomes practically impossible.

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

Acknowledgments

I'd like to start with thanking my thesis advisor Prof. Steven G. Johnson without whose support this thesis would have not been possible. I thank Steven for all his dedication, patience, caring, and passion for teaching. Not only he is an exceptional scientist and researcher, he is also an amazing teacher willing to spend lots of his energy and time on teaching you what he knows. And this combination of a brilliant research mind and an exceptional teacher is what any graduate student would wish in a thesis advisor and I feel very lucky and grateful to have had Steven as my thesis advisor. In addition, a characteristic that Steven has which is not so easy to find, is how much he cares about his students, their success, and their level of happiness with their graduate education. Working with Steven, he constantly gave me the feeling that this Ph.D was all about me and not a small part of his group research.

I would also like to thank Prof. John D. Joannopoulos for all his support, valuable advise in terms of research details, broader picture and prospective of my projects, and his support for the larger Ab-Initio group that I am a member of. Without his vision, support, and exceptional scientific and emotional intelligence, the group would have not been standing so strong today.

Many thanks to Prof. Laurent Demanet for being on my thesis committee and the time and effort that he has put into it.

This thesis would have not been possible without great work of my collaborators and their valuable inputs to the projects. Thank you Alejandro W. Rodriguez, Zhuanfang Bi, David M. Ramirez, Dr. Alexander P. McCauley, Dr. Ardavan F. Oskooi, Prof. Marin Soljacic, Dr. Baile Zhang, Prof. Cheng Wei Qiu, and Dr. David Duchesne.

Thank you to all the SGJ and JDJ group members for making being a member of the group such an amazing experience. I have learned a lot from you and have enjoyed working with you. Special thanks to Dr. Alexander P. McCauley for insightful discussions and for being an amazing offciemate and Prof. Peter Bermel and Dr. Ardaval F. Oskooi for all the computer and technical help.

And thanks to all the individuals that have inspired me at different stages of my life and education and gave me the motivation, inspiration, and valuable advise that helped me through my educational path and the important decisions I made. Special thanks to my math teacher at Foothill Collage, Brian Stanley, my high school physics teachers Mrs. Azadeh and Mrs. Fatemi, and Prof. Ioana Dumitriu.

I'd also like to thank the Sidney-Pacific Graduate Community for making Sidney-Pacific my home, and the community my family in the past 5 years. Special thanks to the housemasters: Prof. and Mrs. Mark and Prof. Kim and Dr. Tang for opening the door to their warm and lovely families and providing me with the warmth and support that my immediate family was sometimes too far away to provide.

And of course, my graduate experience at MIT would have not been so amazing without the wonderful friends I have been very lucky to have during my time here. Thank you Sepehr Shahshahani, Dr. Alexander P. McCauley, Ramis Movassagh, Dr. Pouyan M. Ghaemi, Dr. S. J. Rahi, Dr. Maissam Barkeshli, Dr. Ardavan F. Oskooi, Herve Martin-Rivas, Dr. Danial Lashkari, Dr. Christiana Athanasiou, Dr. Nabil Iqbal, Dr. Ali Hosseini, Dr. Gilda Shayan, Nicole Tariverdian, Dr. Pedram Hekmati, Dr. Kasia Gora, Matthieu Varagnat, Dr. Aristidis Karalis, Dr. Alejadro W. Rodriguez, Dr. Nicolas Pinto, Dr. Nicolas Poilvert, and Dr. Karen Lee.

And thanks to the MIT Ballroom Dance Team for providing me with such an exceptional and efficient outlet for stress and an amazing group of friends in the past two years. Special thanks to my coach Armin Kappacher and my partner Tiago Souza: you have made significant contribution to my experience and life at MIT.

And very special thanks to my family for their unconditional love and support, for trusting me with my decisions and respecting my choices. Without their support in every way they could, I would have not been able to pull through years of studying and making my personal development and education my first priority. Thanks to my sister Haleh Hashemi for all the energy and inspiration, for always being there for me, and doing so much for me. Thanks to my loving husband for all his help and advise, support and patience, and all the inspiration and motivation he has provided me with along the way.

To my loving beautiful parents:

Lida Nahavandian & Jamal Hashemi.

I know no way to thank you enough for my all.

Contents

List of Figures

Chapter 1

Introduction

The field of photonics has been evolving quickly in recent years. For the longest time it was limited to the regime of ray optics for the visible and infrared regimes, in which the devices and objects involved have length scales much larger than the wavelength. However, in recent years, new technologies and man's ability to fabricate devices in the sub-wavelength scale and artificial materials with controlled properties have opened door to many interesting optical phenomena and manipulations of light that were not possible in the past [82, 165].

Today, researchers are able to fabricate delicate optical devices in sub-wavelength scale as small as tens of nano-meters with control and precision over their design and properties that would allow us to predict and control the behavior of light in such devices such as optical waveguides and cavities [56, 69, 77, 91, 95, 132, 137, 138], photonic crystals and slabs [88, 106, 117, 128, 161, 162, 178] and optical fibers [28, 90, 96] to name a few. At the same time, there has been significant advancement in the field of meta-materials: using our ability in sub-wavelength fabrications, we are able to construct artificial materials that exhibit electromagnetic properties which natural materials don't possess [182, 185]. Besides, we are able to modify these properties continuously by varying the geometric parameters. This gives us access to devices and materials with varying parameters and properties so that we are no longer limited to homogeneous materials found in nature.

In addition to fabrication and engineering capabilities, one can not underestimate

the power that computational methods have given us in designing such systems and devices [17, 33, 35, 58, 81, 100, 167, 192, 219]. With fast computers and efficient algorithms, we are now capable of simulating such systems and devices accurately and of studying the behavior of light in such systems to better than experimental accuracy.

All these advancements and capabilities have given rise to many new fields, applications, and physical and optical phenomena and have given us power to control the behavior of light in many different ways: from optical switches and filters, LEDs, optical transistors and lasers, second and third harmonic generation, to "super lenses" and invisibility cloaks [14, 44, 50, 114, 123, 129, 132, 133, 172, 175, 186, 187, 209]

However, despite all the fabrication and computational possibilities and capabilities, it still remains a challenge to come up with a device and a geometry that corresponds to a specific optical process or phenomena that one wishes to create. Maxwell's equations govern electromagnetic waves. These equations are a set of partial differential equations that we need to solve to determine the behavior of light in different geometries. Although these equations are linear, they are not solvable analytically except for few very simple geometries. Solving Maxwell's equations in complex systems and geometries requires fast algorithms and powerful computers. In addition, although these are linear PDE's, they are however very nonlinear in their geometry dependence. In other words, it is far from trivial to come up with a geometry and device design given a particular wave behavior and process.

Given our engineering and computational capabilities today, important questions become: what new phenomena are possible, what useful devices do these phenomena enable, and what are the fundamental limiting factors in practical realization of interesting theoretical phenomena. Answering these questions involves coming up with novel processes and manipulations of light that were not possible or thought of until today, modeling such processes and systems and studying their dynamics, coming up with an appropriate device design among all the fabrication tools and devices available, and finally doing the fabrication and experiments, and

potentially advancing the fabrication method to something that can be commercialized and widely used depending on applications of the process. It is equally important to study the limitations of such processes both in terms of fundamentals of physics such as causality, conservation of energy implications, and also limitations that inevitable material losses and our current fabrication capabilities impose. And this is where the field of optics stands today, trying to come up with novel design and applications of optical processes through thinking through the questions mentioned above.

One such an optical process is the process of frequency conversion: converting the light at one frequency to light at another frequency. Such conversion is done through use of intrinsic material nonlinearities that allow light units called photons of different frequencies interact with each other and create new photons at new frequencies, through processes such as: second and third harmonic generation, difference and sum frequency generation, and degenerate four-wave mixing [9, 10, 23, 48, 52, 97, 114, 116, 133, 134, 166, 168, 183, 193], two of which are shown schematically in Fig. 1-1. Such processes can be used in many applications: theoretically, they can act as a light source at the generated frequency for frequencies that are otherwise not easy to generate [12, 19, 42, 62, 99, 118, 139, 156, 193] besides other applications such as optical imaging, sensing, and image processing in medical applications [25, 59, 177], retardation measurements in optical elements [26] and spectroscopy [134]. The material nonlinearities that create these processes can be found in nature. However, optical nonlinearities are usually very weak in natural materials. For example, to induce only a 0.1% change in the refractive index in silica glass, one would need to apply an electromagnetic field with intensity of about 50 kW/μm^2, which could easily cut through solid steel, much less melt glass. As a result, if just left to nature, such frequency conversions occur at a very slow rate resulting in very small conversion efficiencies; a 10% conversion efficiency would be an impressively high efficiency for such processes [112]. In addition, macroscopic devices typically require relatively high powers ($\sim$ Watts) [147, 172]. The efficiency and high power requirements remain the main challenges in the field. In this thesis

we look closely at these processes, analyze the physics at a fundamental level decoupled from any specific geometry, and propose a general strategy and specific geometries that overcome the efficiency and high power requirement challenges. Another optical phenomena is the very popular idea of invisibility cloaks; an idea that man has dreamed of from the time of Greek myths of Perseus to today's science-fiction books and movies such as Harry Potter, Star Trek, or the Invisible Man. Today, this long-time dream may appear closer to reality thanks to availability and progressive advancement of the meta-materials. As Pendry showed mathematically in 2006 [155], if one picked the cloak material with certain electromagnetic properties, one could modify the path of light and make it go around the object so that it appears perfectly invisible. Subsequently, it was shown that one could approximate these properties in reality using metamaterials, and experimental demonstrations of partial "cloaking" of small objects have given hope that some form of practical cloaking might be around the corner [32, 50, 60, 63, 93, 107, 120, 122, 125, 126, 175, 184, 199, 215]. However, fundamentals of physics impose severe limitations on practical cloaking. In the second part of this thesis, we focus on the fundamental and practical limitations of transformation-based invisibility cloaks. In particular, we show for the first time that there are practical limitations that grow with the size of the object being cloaked, leading to an intrinsic difficulty in scaling up experimental cloaking demonstrations from wavelength-scale objects to larger objects.

1.1 Frequency conversion

In this thesis, we study different forms of frequency conversion, also known as harmonic generation, in which an input signal at a certain frequency ω_{in} is converted to a signal at a different frequency ω_{out}. Different techniques and processes can generate different ω_{out}. For example, one can use a Kerr ($\chi^{(3)}$) nonlinearity, which allows interactions between three photons to generate $\omega_{\text{out}} = 3\omega_{\text{in}}$, or "third-harmonic generation" (THG). Similarly, a $\chi^{(2)}$ nonlinearity

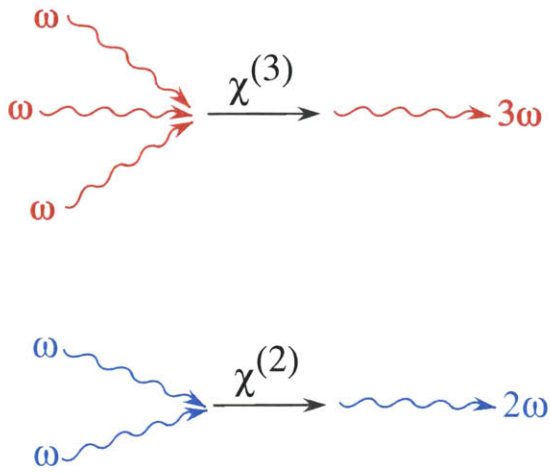

Figure 1-1: *Top:* Third harmonic generation: three photons at frequency ω are added through a $\chi^{(3)}$ (Kerr) nonlinearity and generate a photon at triple frequency 3ω. *Bottom*: similarly, a $\chi^{(2)}$ nonlinearity allows interaction between two photons at frequency ω and generates a photon at double frequency 2ω.

allows interaction between two photons and results in doubling the frequency of the input signal; this is "second-harmonic generation" (SHG). THG and SHG are shown schematically in Fig. 1-1. Similar processes can result in frequency summation of two input signals (sum-frequency generation, SFG) or difference-frequency generation (DFG). Regardless of the relationship of the output signal with the input signal(s), all these processes use some nonlinearity in the medium they occur in to allow the photons of the input signal(s) interact with each other and create new photons at the output frequency.

What allows such processes is intrinsic material nonlinearities found in nature. Such nonlinearities allow photons to interact and generate new photons. However, as discussed earlier, such natural nonlinearities are extremely weak and pose a challenge to most practical applications of such processes. These challenges resulting from the weakness of nonlinearities come in two ways. (i) Conversion efficiency: due to weak interactions, the conversion rate is very low and results in very low conversion efficiencies (less than 10%); most of the input signal is outputted unconverted. (ii) High power requirements: in order to get any sensible conversion, even with low efficiencies reported above, high input powers are required in order to

enhance nonlinear interactions. These challenges limit the applicability of such processes to real applications and remain the main challenges in this field of optics.

Due to broad interest and wide range of applications mentioned earlier, a lot of work has been done to come up with different methods of doing frequency conversion. One common approach has been in context of waveguides: an input signal travels in a waveguide made of nonlinear material and some of the photons in the signal get converted while traveling down the waveguide and the two signals at ω_{in} and ω_{out} co-propagate in the waveguide [5, 6, 10, 38, 48, 129, 144, 153]. This method however, suffers from extremely low efficiency and high required power, the reason being that the method does not do much in enhancing nonlinear interactions; unless the waveguide is very long to allow long time exposure of the signal to nonlinearities or the power is high to create high intensity fields in the nonlinear material, the conversion efficiency stays very low.

Another common approach is to use a single-mode cavity at the pump frequency [8, 9, 15, 20, 37, 42, 44, 61, 62, 65, 99, 121, 134, 139, 140, 156, 172, 174, 183] or the harmonic frequency [41]. Using an optical cavity with resonance at one of these frequencies allows confinement of that mode in the cavity and results in enhancement of the nonlinear interactions. However since the cavity supports only one of the two modes involved in the process strongly, it still requires high powers (of order of Watt [147, 172]) and/or amplification in the cavity.

Yet, another approach would be to use doubly resonant optical cavities; an optical cavity that is resonant at both ω_{in} and ω_{out} so both modes are confined in the nonlinear cavity for long times. This would provide both the spatial and temporal confinement at both modes promising higher efficiencies and lower input powers required. In the case of second harmonic generation, there are some works that looked at doubly-resonant cavities mostly only in the low-efficiency regime [12, 47, 119, 127, 150, 221] and few that looked at high efficiency regime [45, 147, 203]. Although it has been shown before that frequency conversion can be done for arbitrary low input powers, that would require very long cavity lifetimes and hence very small band-width. In the case of third harmonic generation,

however, we are not aware of any work done prior to the projects that form a large portion of this thesis that studies such process in doubly resonant cavities.

This thesis is concerned with addressing these two main limiting factors in harmonic generation processes, namely the efficiency and high input power requirements. It is an interesting problem to think of ways to enhance these nonlinear interactions, increase the conversion efficiency while keeping the power low. Two ways to think about it are to enhance such interactions by temporal and spatial confinement of the modes. Temporal confinement means trapping the fields inside the nonlinear medium for a long time so the photons have more time to interact and this should increase the conversion efficiency. And spatial confinement means to trap the fields in a small volume where their intensity is high and therefore there's more interaction; this should help with keeping the input power required low while maintaining high field intensities in the cavity, enhanced nonlinear interactions, and so high conversion efficiencies.

Given our objectives, a nonlinear doubly resonant optical cavity that supports all the operating and generated frequencies and modes seems to be a natural choice as it will provide us with a nonlinear medium in which the nonlinear interaction are enhanced by 1) trapping the fields at the input and output frequencies for a long time and 2) localizing them with high intensity in a small modal volume. Imagine an input channel from which an input signal at frequency ω_1 can couple to the cavity and an output channel where the generated fields at frequency ω_k can escape the cavity. Now if one designed this whole system carefully and picked all the parameters starting from the field lifetimes in the cavity, their coupling strength to the input and output channels, and the power of the input signal optimally, one would expect an optimal conversion efficiency. In this thesis we look closely at these processes: starting from their mathematical modeling, coming up with ways to achieve high efficiency conversion, and designing devices that can demonstrate such processes.

To model such a system mathematically and analyze its properties and behavior, we take advantage of the weakness of nonlinearities and model this nonlinear process in

the form of time-invariant first order perturbation theory. More specifically, we use coupled-mode theory to model these processes. Coupled-mode theory is a simple model of this physical phenomena based on few simple assumptions: 1) the nonlinearities are weak, 2) All other interactions and nonlinear processes that result in photons 'with frequencies different than the input and output frequencies are weak and negligible, 3) The materials and geometries are time-invariant and some general properties such as conservation of energy and time-reversal invariance.

As we'll see in detail, based on these simple assumptions and first principle, we develop a framework, a set of ordinary differential equations based on the system parameters such as field lifetimes in the cavity, coupling strengths of the fields in the nonlinear medium, and input power, that describe this rather complex process of frequency conversion. As it was shown in a prior work by my collaborator Alejandro W. Rodriguez [169], this system achieves 100% efficient frequency conversion for a "critical" input power in the case of 2nd and 3rd harmonic generation. In the same work, he also showed that this required "critical" power can be made arbitrary small by either increasing the mode lifetimes in the cavity or making the modal volume of the modes small while maintaining a reasonable bandwidth. However, this early analysis left out any consideration of the nonlinear dynamics of the system, and simplified the problem by omitting a key nonlinear process that always occurs in $\chi^{(3)}$ medium: at the same time that frequency conversion is occurring, the light also shifts the resonant frequencies of the cavities (possibly driving the system out of resonance). I rectified these limitations by developing a more complete analysis [71] of 3rd harmonic generation, the 100% efficient solution, and its stability. I analyzed all the criteria and system parameters where the complete conversion solution is stable and achievable in practice. Finally, given the dynamics of the system, I introduced different possibilities of achieving and exciting the complete conversion solution. In addition to our complete frequency conversion focus, I also found that the system possesses a variety of nonlinear phenomena and dynamics including interesting bifurcations, limit cycles, multi-stable solutions, and hysteresis behavior. These complex dynamics not only are interesting on their own,

but they could also be used in many different applications other than frequency conversion such as optical switches (using hysteresis and multi-stable behavior) and optical clocks (using the limit cycles).

Next, I co-supervised a UROP student whom we led through a parallel study and analysis for the case of degenerate four-wave mixing (DFWM). Such a process also uses Kerr nonlinearities ($\chi^{(3)}$), but to convert input signals at frequencies ω_0 and $\omega_m = \omega_0 - \Delta\omega$ to generate a signal at $\omega_p = \omega_0 + \Delta\omega$. Similar to the case of 3rd harmonic generation, we observe interesting nonlinear dynamics including bifurcations, multi-stability, and limit cycles. In this case however, 100% conversion efficiency is not possible due to a "quantum limit" (which can be derived classically as well as from photon considerations) and the maximum efficiency possible is given by $\omega_p/2\omega_0$ (for $\omega_p < 2\omega_0$).

Next, we look at design ideas for devices that posses the characteristics we are looking for and are in the parameter regimes where achieving the high-efficiency solutions are possible in practice. The design process involves choosing a type of optical cavity that can have resonances at the frequencies involved. For example, in the case of 2nd and 3rd harmonic generation, we propose using of a ring resonator as the cavity. Since the spacing between the operating and harmonic frequencies is large, it is relatively easy to design a ring resonator that supports both frequencies and is phase-matched. In the case of degenerate four-wave mixing, we normally look at small frequency spacing. Therefore a ring resonator would hardly work. In that case, it is easier to use waveguide cavities with a relatively large band-gap that include all the operating and harmonic frequencies. After choosing of a specific type of optical cavity, we twig its parameters so that it supports the desired modes. This cavity is coupled to an input and output channel (waveguides). The coupling strength is what determines the field lifetimes in the cavity, parameters very important in controlling the conversion efficiency and the critical input power. We show that, in order to obtain optimal efficiency, or even to attain stable solutions at all in the case of THG, a novel design of a ring with separate input-output waveguides, one of which is designed to have a low-frequency cutoff, is required in

order to achieve the desired coupling parameters. Although my initial considerations were for THG, we found that a similar design was desirable for SHG and we carried out the design of a ring-resonator SHG structure in collaboration with a visiting student, Zhuanfang Bi, who performed full nonlinear Maxwell calculations to verify the predictions of coupled-mode theory and demonstrated $\sim 90\%$ SHG efficiency in the presence of losses, ten times higher than what has been achieved in previous experimental designs. Finally, we propose ideas and discuss challenges of designing the same system for 3rd harmonic generation and give a brief design idea for the degenerate four-wave mixing (which is currently being pursued by another student).

1.2 Invisibility cloaks

Invisibility "cloaking" refers to the idea of making an object appear invisible, or at least greatly reducing its scattering cross-section, by surrounding it with appropriate materials. There has been large interest in cloaking of objects; because it is scientifically very interesting, it's a popular and exciting idea to the public and has shown up in myths, stories, and movies, ancient and recent, and it can have numerous applications if achieved at a practical level, especially for military purposes.

In particular, we consider the most ambitious goal, that of true invisibility: to make an object undetectable by any observer of light in a given frequency range, in which the observer is capable of distinguishing any alteration in the phase or amplitude of light from sources at any location. In contrast, all current "stealth" technology attacks much more modest problems. For example, stealth aircraft are primarily designed to reduce radar backscattering (e.g. by radar-absorbing materials), under the assumption that the observer is co-located with the radar source [27, 189], but still casts a "shadow" behind the aircraft. Camouflage clothing not only fails to eliminate the shadow cast by a backlit person, but also causes light bouncing off the person to experience a different time delay than it would when scattering off of the background—it relies on the inability of the observer to distinguish small time

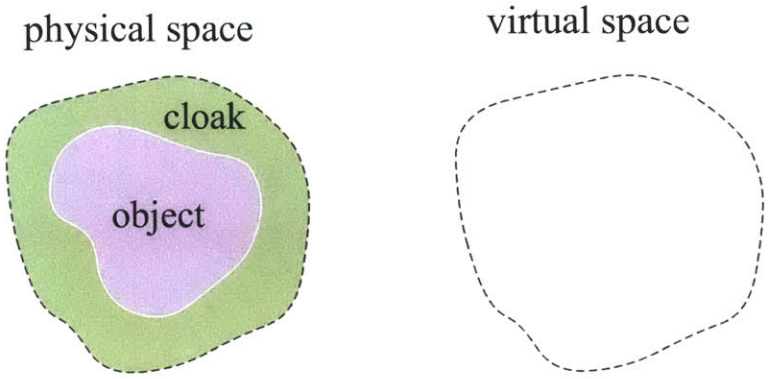

Figure 1-2: A transformation-based isolated-object cloak. The cloak transforms the physical space on the left where the object and the cloak are to an empty virtual space shown on the right.

delays combined with imperfect knowledge of the background and other limitations of human observers. One can also consider "active" cloaking in which the cloaked object is covered with emitters that radiate signals to mimic the background and/or to destructively interfere with the scattered waves, but such devices also intrinsically suffer from inevitable time delays in their response [135].

In 2006, a theoretical cloaking proposal by Pendry [155] attracted extensive popular and research interest. Pendry's idea was based on the coordinate invariance properties of Maxwell's equations that we will discuss in chapter 2. In short, a coordinate transformation in Maxwell's equations is equivalent to a change in materials' electromagnetic properties ϵ and μ. We can utilize this property to cloak an object: surround the object with appropriate material with ϵ and μ that are equivalent to a coordinate transformation of the physical space to a virtual space in which the object space does not exist (has been mapped to a single point), as depicted in Fig. 1-2. In the physical space, these materials make any incident wave go around the object, as shown in Fig. 1-3, exiting the cloak on the exact trajectory (and with the exact phase and amplitude) the wave would have if there were no object or cloak. Therefore, to an observer, it will look like the object was never there. However, unfortunately, cloaking of isolated objects turns out to be severely restricted, in that speed-of-light/causality constraints intrinsically limit perfect cloaking to an infinitesimal bandwidth [135, 155]. An easy way to see this limitation

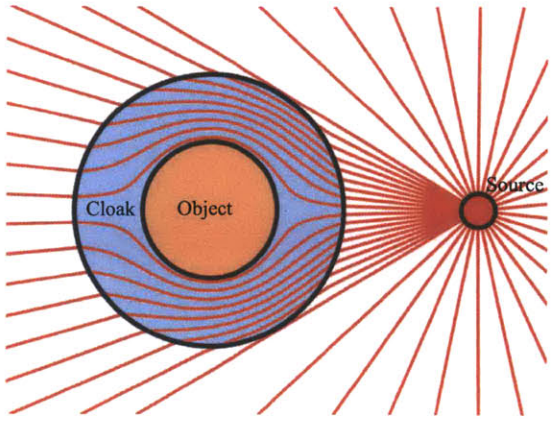

Figure 1-3: A demonstration of how a transformation-based cloak modifies the path of light so that it does not hit and reflect from the object. A point source is placed near the object/cloak system where the cloak is designed such that the rays incident to the cloak travel around the object. This picture is copied and slightly modified from [155].

is the following: in order for the object to appear invisible, the incident wave must travel around the object (in the cloak) in the same period of time as it would have traveled in a straight line if the object was not there. Therefore, the beam must travel a longer path in the same time. If in vacuum, this implies the incident beam must travel faster than the speed of light c. And this is not possible over a non-zero bandwidth.

An alternative with no intrinsic bandwidth limitations is ground-plane cloaking [113], in which the goal is to make an object sitting on a reflective surface indistinguishable from the bare surface as shown in Fig. 1-4. This again uses the coordinate-transformation idea, choosing cloak materials that map the object "into the ground," but the practical realization is much easier because the transformation is now non-singular. The easiest case is if the ground is a perfect absorber ("black" ground), in which case one would only need a black cloak that absorbs all incident waves. The harder and more interesting case is one in which the ground is a good reflector, in which case the cloak needs to mimic reflection off of the ground as if the object weren't there. There have been theoretical investigation of several variations on the underlying idea of ground-plane cloaking [49, 105, 205, 214] and considerable fundamental theoretical interest in isolated-object cloaking [7, 24, 30, 31, 39, 66, 67, 76,

80, 92, 102, 109, 110, 124, 143, 163, 170, 176, 196, 206, 212, 216, 218, 220]. In terms of experiments, there have been several experimental demonstrations of ground-plane cloaking [32, 50, 60, 63, 107, 120, 125, 126, 199, 215] and single-frequency cloaking of small isolated objects [93, 122, 175, 184] including some three dimensional cloaks.

However, despite all the attention and scientific investment that the field of transformation-based cloaking has acquired, there had been almost no attention to limitations, both theoretical and practical, that can severely restrict what can be done in practice and what practical applications the current experiments can have.

For example, although several theoretical simulations included absorption loss [24, 29, 39, 66, 80, 92, 115, 216], the first work we are aware of to suggest a trade off between absorption tolerance and object size was a numerical calculation by a colleague [212] following a suggestion by our group. Similarly, in the experiments, there seems to have been no discussion of or concerns over practicality of such experiments and whether such experiments are scalable to cloaking of larger objects. The objects cloaked in experiments were at most few wavelengths in size and in fact are mostly sub-wavelength in size. Although cloaking of objects no matter how small is a very exciting and valuable scientific achievement, one can not ignore the question of whether such experiments are scalable since for truly resolving an object of interest, one would use an incident beam with wavelength smaller that the length-scale of the object. And therefore the interesting regime with potential application is when the object is many wavelengths in size.

In the second part of this thesis, we focus on analyzing different practical limitations of transformation-based cloaking, both in the case of isolated-object cloaking and ground-plane cloaks. We study how absorption losses, scattering due to random imperfections in the cloak, and bandwidth of the cloaking scale with the size of the object. First, we illustrate these limitations with an idealized one-dimensional (1d) system in which cloaking is much simpler than in three dimensions (3d)—only one incident wave need be considered—but in which the same limitations appear. Just based on this simple model, we find that basic physical principles imply that cloaking of human-scale objects is challenging even at radio frequencies (RF), while

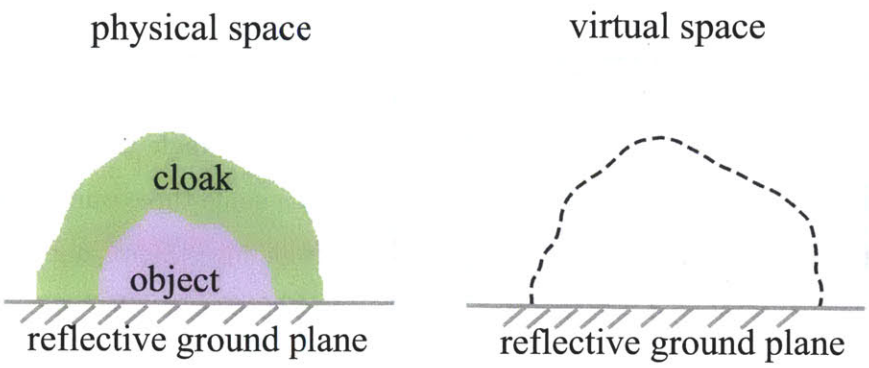

Figure 1-4: An example of a ground-plane cloak. The cloak makes the object and cloak look like a reflective sheet in an empty space.

cloaking such objects at much shorter (e.g. visible) wavelengths is rendered impractical by the delay–loss product: the absorption losses have to asymptotically vanish as the object gets larger, and given the materials available, the required properties becomes impractical. Despite the simplicity of this analysis, we arrive at fundamental criteria that may help guide future research on the frontiers of cloaking phenomena. Indeed, one of the conclusions of our early analysis was that cloaking would become easier for objects immersed in a fluid (although ultimately the same scaling limitations apply), and two subsequent experiments demonstrated cloaking of mm-scale objects at visible wavelengths in fluids [32, 215]. For example, it has been suggested that gain could be used to compensate for absorption loss (but not other imperfections) in the cloak [66], but a corollary of our results still requires that such compensation must become increasingly exact as the object diameter increases (and in any case gain is susceptible to nonlinear saturation).

Generalizing this simple one-dimensional argument, we then study three dimensional cloaks both in the isolated-object case and also ground-plane cloaking and give rigorous proofs of limitations for arbitrary cloaking transformations in three dimensions. In both cases, our key starting point is the assumption that the attainable refractive indices are bounded. We show that (for both isolated-object and ground-plane cloaks), the cloak thickness must scale with the object size and hence any losses per unit volume (including both absorption and scattering from

imperfections) must scale inversely with the object size. In addition to scattering and loss, systematic imperfections (such as an overall shift in the indices or an overall neglect of anisotropy in favor of approximate isotropic materials [113]) must also vanish inversely with object thickness, since such systematic errors produce a worst-case phase shift in the reflected field proportional to the imperfection and the thickness of the cloak (the path length, which scales with the object). (For oblique angles of incidence, such phase shifts can cause a lateral shift in the reflected beam [214], analogous to a Goos-Hänchen shift.) From this, if one requires a bounded reduction in the scattering cross-section, it follows that the loss (due to absorption or other imperfections) per unit volume must scale inversely with the object thickness, and we quantify this scaling more precisely in the case of absorption loss and scattering from disorder. In the case of isolated object cloaking, there is an additional limitation: the singularity of the cloaking transformation (which maps an object to a single point) corresponds to very extreme material responses (e.g. vanishing effective indices) at the inner surface of a perfect cloak [213]. Independent of the limitations on losses, we show that if the attainable refractive indices are bounded below, then the cloak is necessarily imperfect: it reduces the object cross section by a bounded fraction, even for otherwise lossless and perfect materials. And finally we look closely at the bandwidth of isolated-object cloaking and how it changes depending on the size of the object. Although perfect cloaking of an isolated object is not possible over a nonzero bandwidth, it is still worth investigating how fast the scattering cross-section increases once we move away from the operating frequency ω_{op} and how this rate of scattering cross-section increase (decrease in the cloak performance) scales with the size of the object being cloaked. We show that if one requires a bound on the scattering cross-section, the bandwidth of cloaking decreases inversely proportional to the square root of the radius of the object.

Chapter 2

Fundamentals: Maxwell's Equations and Properties

Light is an electromagnetic wave and therefore is governed by a wave equation. As derived by James Clerk Maxwell in four publications in 1861–1865, the governing equations are a set of partial differential equations written both in the microscopic and macroscopic forms.

All the macroscopic electromagnetic waves are governed by Maxwell's equations:

$$\begin{aligned} \nabla \cdot \mathbf{B} &= 0 \quad & \nabla \times \mathbf{E} + \frac{\partial \mathbf{B}}{\partial t} &= 0 \\ \nabla \cdot \mathbf{D} &= \rho \quad & \nabla \times \mathbf{H} - \frac{\partial \mathbf{D}}{\partial t} &= \mathbf{J} \end{aligned} \tag{2.1}$$

in SI units, where $\mathbf{E}$ and $\mathbf{H}$ are the macroscopic electric and magnetic fields, $\mathbf{D}$ and $\mathbf{B}$ are the displacement and magnetic induction fields, and ρ and $\mathbf{J}$ are the free charge and current densities. The displacement fields $\mathbf{D}$ (or $\mathbf{B}$) are related to $\mathbf{E}$ (or $\mathbf{H}$) via power series describing the polarizability of different materials in powers of $\mathbf{E}$ (or $\mathbf{H}$). Physically, when subject to an applied field, the charges in a material experience forces that polarize the atoms into microscopic dipoles, leading to a macroscopic polarization density $\mathbf{P}$. To lowest-order, the polarization is linear in the applied field: $\mathbf{P} = \varepsilon_0 \chi_e \mathbf{E}$ where χ_e is the electric susceptibility of the material (related to the polarizability of the atoms) and $\epsilon_0 \approx 8.854 \times 10^{-12}$ Farad/m is the

vacuum permittivity. The displacement field $\mathbf{D}$ is then related to $\mathbf{E}$ by

$$\mathbf{D} = \varepsilon_0 \mathbf{E} + \mathbf{P} = \varepsilon_0(1 + \chi_e)\mathbf{E} = \varepsilon \mathbf{E},$$

defining the material's electric permittivity $\varepsilon = \varepsilon_0(1 + \chi_e)$. To higher orders, the polarization can be expanded in a power series in terms of the electric field. The series would relate the components the displacement field D_i to components of the electric field E_i via:

$$D_i/\varepsilon_0 = \Sigma_j \varepsilon_{ij} E_j + \Sigma_{j,k} \chi^{(2)}_{ijk} E_j E_k + \Sigma_{ijkl} \chi^{(3)} E_j E_k E_l + O(E^4) \tag{2.2}$$

Most electromagnetic systems operate in the linear regime, where the higher order terms $\chi^{(2)}$ and $\chi^{(3)}$ are negligible. However, in nonlinear systems and when one is interested in higher order effects, one goes one step further and looks at the first non-zero term in the series. One would naturally think that the next non-zero term would be the $\chi^{(2)}$ term, which is true for some materials and we will study applications of such nonlinearities in this thesis a little bit, namely second harmonic generation where two photons at frequency ω_1 generate a photon at the double frequency $2\omega_1$. However, in many common materials (including silicon and glass), this first nonlinear term is zero. This arises from inversion symmetry, in which the atom structure is indistinguishable under the coordinate transformation $\mathbf{x} \to -\mathbf{x}$. If a material has inversion symmetry, then flipping $\mathbf{E} \to -\mathbf{E}$ must flip $\mathbf{P} \to -\mathbf{P}$ and hence the $\chi^{(2)}$ coefficient (or the coefficient of any even power of $\mathbf{E}$ in the series) must be zero. Therefore, the third term in the series, $\chi^{(3)}$, otherwise known as a Kerr nonlinearity, becomes the dominant nonlinear process in many natural materials and applications.

There is a similar power series for $\mathbf{B} = \mu \mathbf{H}$, but at infrared and optical frequencies most materials have negligible magnetic response, so they act as linear materials with $\mu \approx \mu_0$, the vacuum permeability $4\pi \times 10^{-7}$ Henry/m. The velocity of light (electromagnetic waves) in vacuum is $c = 1/\sqrt{\varepsilon_0 \mu_0}$ and the phase velocity in a medium is $1/\sqrt{\varepsilon \mu}$. The *refractive index* (familiar from Snell's law of optics) is

$n = \sqrt{\varepsilon\mu}/\sqrt{\varepsilon_0\mu_0} \approx \sqrt{\varepsilon/\varepsilon_0}$ (corresponding to a phase velocity c/n). The relative permittivity $\varepsilon/\varepsilon_0$ is also known as the *dielectric constant*. (In many theoretical settings we choose units where ε_0 and μ_0 are both 1.)

2.1 Nonlinear processes

Let's take a closer look at what type of processes this third-order nonlinearity can create: suppose a signal of frequency $\omega_1 (\mathbf{E} = \mathbf{E}_1 e^{i\omega_1 t} + \mathbf{E}_1^* e^{-i\omega_1 t})$ is subject to such a nonlinearity. The $\chi^{(3)} E^3$ terms then immediately lead to $(e^{\pm i\omega_1 t})^3$ terms and hence fields at the *third-harmonic* frequency $3\omega_1$. There is also *self-phase modulation*: E^3 terms of the form $(e^{\pm i\omega_1 t})^2 e^{\mp i\omega_1 t} = e^{\pm i\omega_1 t}$ that act as a nonlinear change in the refractive index at ω_1, where the change is proportional to the field intensity E^2 (not cubed). Once the THG process has occurred, however, so that $\mathbf{E} = \mathbf{E}_1 e^{i\omega_1 t} + \mathbf{E}_1^* e^{-i\omega_1 t} + \mathbf{E}_3 e^{i3\omega_1 t} + \mathbf{E}_3^* e^{-i3\omega_1 t}$ for some amplitudes $\mathbf{E}_1$ and $\mathbf{E}_3$, the $\chi^{(3)}$ terms will yield additional frequencies. There will be higher harmonics $9\omega_1$ and so on, but one will also see *down-conversion* from $\chi^{(3)} (e^{\mp i\omega_1 t})^2 e^{\pm i3\omega_1 t} = e^{\pm i\omega_1 t}$, which combine ω_1 and $3\omega_1$ waves to produce more ω_1 waves. More generally, whenever waves at three frequencies $\omega_{1,2,3}$ are present, one obtains E^3 terms at frequencies $\pm\omega_1 \pm \omega_2 \pm \omega_3$, a process known as *four-wave mixing* (FWM). *Degenerate four-wave mixing* (DFWM) is when two of the frequencies are the same (of which down-conversion is a special case). Finally, another special case of DFWM is *cross-phase modulation* (XPM): when two frequencies $\omega_{1,2}$ produce terms at frequency $\omega_1 - \omega_1 + \omega_2 = \omega_2$, which acts like a change in the refractive index at ω_2 proportional the field intensity E_1^2 at ω_1. Self- and cross-phase modulation complicate the design of resonant nonlinear devices, because they cause the resonant frequencies of optical microcavities to shift as the field intensity changes; this is discussed in detail in the subsequent chapters.

In the case of a $\chi^{(2)}$ nonlinearity, there are similar *second-harmonic generation* (SHG: $\omega_1 \to 2\omega_1$), *sum-frequency generation* (SFG: $\omega_{1,2} \to \omega_1 + \omega_2$), and *difference-frequency generation* (DFG: $\omega_{1,2} \to \omega_1 - \omega_2$ processes, but matters are

simplified compared to $\chi^{(3)}$ because there is no self/cross-phase modulation. (Note that SHG combined with DHG produces a down-conversion process in which $2\omega_1$ waves are converted back down to ω_1 waves.)

In this thesis, we will closely look at the $\chi^{(3)}$ effect (i.e. Kerr nonlinearity) and its applications in 3rd harmonic generation and degenerate four-wave mixing. We will also briefly look at a system that uses $\chi^{(2)}$ nonlinearities to generate 2nd harmonic modes. But for the cloaking section, we only look at linear materials (where we will set ε_0 and μ_0 to 1 for convenience).

2.2 Linear equations in frequency domain

In the approximation of a linear, non-dispersive material, one can write:

$$\begin{aligned} \mathbf{D}(\mathbf{x},t) &= \varepsilon(\mathbf{x})\mathbf{E}(\mathbf{x},t) \\ \mathbf{B}(\mathbf{x},t) &= \mu(\mathbf{x})\mathbf{H}(\mathbf{x},t) \end{aligned} \tag{2.3}$$

where $= 4\pi \times 10^{-7}$ Henry/m is the vacuum permeability. In the case of anisotropic materials, ε and μ are 3×3 matrices (rank-2 tensors). Substituting these in Eq. (2.1) gives us:

$$\begin{aligned} \nabla \times \mathbf{E}(\mathbf{x},t) + \mu(\mathbf{x})\frac{\partial \mathbf{H}(\mathbf{r},t)}{\partial t} &= 0 \\ \nabla \times \mathbf{H}(\mathbf{x},t) - \varepsilon(\mathbf{x})\frac{\partial \mathbf{E}(\mathbf{x},t)}{\partial t} &= \mathbf{J}(\mathbf{x},t) \end{aligned} \tag{2.4}$$

Because these equations are linear and time-invariant, one can Fourier-transform in time to consider only a single frequency ω at a time:

$$\mathbf{H}(\mathbf{x},t) = \mathbf{H}(\mathbf{x})e^{-i\omega t}, \tag{2.5}$$

$$\mathbf{E}(\mathbf{x},t) = \mathbf{E}(\mathbf{x})e^{-i\omega t}, \tag{2.6}$$

$$\mathbf{J}(\mathbf{x},t) = \mathbf{J}(\mathbf{x})e^{-i\omega t}. \tag{2.7}$$

[Often, instead of talking about the frequency ω, it is convenient to instead use the (vacuum) *wavelength* $\lambda = 2\pi c/\omega$.] Plugging this into Eq. (2.4) gives the two curl equations in the form:

$$\nabla \times \mathbf{E}(\mathbf{x}) - i\omega\mu(\mathbf{x})\mathbf{H}(\mathbf{x}) = 0 \tag{2.8}$$

$$\nabla \times \mathbf{H}(\mathbf{x}) + i\omega\varepsilon(\mathbf{x})\mathbf{E}(\mathbf{x}) = \mathbf{J}(\mathbf{x}). \tag{2.9}$$

(More generally, dispersive materials, ε and μ will depend upon ω.) In the case of source-free ($\mathbf{J} = 0$) waves, with very little algebra, one can write these equations in form of generalized eigenvalue problems for $\mathbf{E}$ or $\mathbf{H}$. In particular, one obtains:

$$\nabla \times \frac{1}{\mu(\mathbf{x})}\nabla \times \mathbf{E} = \omega^2\varepsilon(\mathbf{x})\mathbf{E}, \tag{2.10}$$

$$\nabla \times \frac{1}{\varepsilon(\mathbf{x})}\nabla \times \mathbf{H} = \omega^2\mu(\mathbf{x})\mathbf{H}. \tag{2.11}$$

We will be using Maxwell's equations in this form both for the harmonic generation problem (in order to describe the linear resonant modes before nonlinear interactions are added as a small perturbation), and more generally in the form of a frequency-domain scattering problem (for some incident field or a field created by a $\mathbf{J}$ outside the object) in the cloaking sections.

2.2.1 Scaling properties

The frequency-domain linear Maxwell's equations have interesting scaling and invariance properties that are used throughout this thesis. Therefore, we do a review of these properties here.

For example, rescaling the problem by a compression/expansion factor s so that $\mathbf{x}' = s\mathbf{x}$, one can get back to Eq. (2.11) by rescaling: $\varepsilon'(\mathbf{x}) = \varepsilon(\mathbf{x}/s)$ and $\omega' = \omega/s$, $\mathbf{H}'(\mathbf{x}) = \mathbf{H}(\mathbf{x}'/s)$ and $\mathbf{E}'(\mathbf{x}) = \mathbf{E}(\mathbf{x}'/s)$. In other words, if we rescale the geometry by a factor s and the wavelength by s, the solutions are the same, just scaled by the same factor. This is a very useful property as it allows us to solve or design a system in an arbitrary or convenient lengthscale and then just rescale the solutions

and the results afterwards if needed. One caveat to this is material dispersion: generally, ε varies significantly with wavelength, so to obtain the same solutions at different lengthscales one may need to choose different materials so as to have the same ε (and μ).

2.3 Coordinate transformation

In a more general form, Maxwell's equations also possess a form of invariance under coordinate transformations. More specifically, it has been shown that a coordinate transformation of Maxwell's equations is equivalent to a change in material properties ε and μ [155]. This has many implications and applications both in making solving of Maxwell's equations easier both analytically and computationally and also in using this invariance to create very interesting physical phenomena and optical devices such as invisibility cloaks.

For example, if you want to design a geometry that warps light in a particular way (e.g. bending it, focusing it, and so on), one instead comes up with a coordinate transformation that captures the desired warping, and arrives at materials and geometry that mimic this transformation. This design technique has come to be known as *transformation optics*. Even though the resulting geometries and materials may be very complicated, the solution of Maxwell's equation is often trivial because it is merely a remapping of the solution in vacuum or some similarly simple system; the challenge is instead coming up with metamaterials with the requisite properties (or finding transformations that lead to attainable materials). Invisibility cloaks are but one dramatic application of this technique.

Other applications of these transformation properties arise in electromagnetic simulations. Suppose we have a code that can solve Maxwell's equations in Cartesian coordinates, but we have a system more efficiently described in cylindrical coordinates (e.g. because it is axisymmetric) that we would like to solve. Then we could convert the system to Cartesian coordinates by changing the materials properly and use our available code to solve it. Another well-known computational

application of this coordinate invariance property is the perfectly matched layer (PML), in which one uses a transformation to *complex* coordinates near the boundary of the computational cell to obtain an artificial absorbing material that absorbs all the outgoing waves. Coordinate transformations can even be used in semi-analytical techniques to express the effect of small perturbations to dielectric interfaces (e.g. bumps or surface roughness) in a tractable way [98].

In particular, consider a differentiable coordinate transformation $\mathbf{x} \to \mathbf{x}'(\mathbf{x})$. Let $\mathcal{J}$ be the 3×3 Jacobian matrix characterizing this transformation:

$$\mathcal{J}_{ij} = \frac{\partial x_i'}{\partial x_j}$$

Then, Ward and Pendry [200] showed that Maxwell's equations take on the same form in the transformed coordinate system, with ∇ replaced by ∇', if one makes the following transformations:

$$\mathbf{E}' = (\mathcal{J}^T)^{-1}\mathbf{E} \tag{2.12}$$

$$\mathbf{H}' = (\mathcal{J}^T)^{-1}\mathbf{H} \tag{2.13}$$

$$\varepsilon' = \frac{\mathcal{J}\varepsilon\mathcal{J}^T}{\det \mathcal{J}} \tag{2.14}$$

$$\mu' = \frac{\mathcal{J}\mu\mathcal{J}^T}{\det \mathcal{J}} \tag{2.15}$$

$$\mathbf{J}' = \frac{\mathcal{J}\mathbf{J}}{\det \mathcal{J}} \tag{2.16}$$

$$\rho' = \frac{\rho}{\det \mathcal{J}} \tag{2.17}$$

where $\mathcal{J}^T$ is the transpose. (Although the original derivation was somewhat cumbersome and was limited to isotropic ε and μ, a general one-page derivation making use of Einstein/repeated-index notation can be found in later work [98].) Transformation-based cloaking is all based on this simple invariance property of Maxwell's equations. It is worth noting that even if we start with isotropic ε and μ, we may very well end up with anisotropic material after the transformation as we will see in cloaking application. This fact has consequences for both computational

and experimental purposes: to simulate such systems using numerical methods we would need to have codes that support anisotropic materials, and in experiments we would need to be able to fabricate such anisotropic materials.

Chapter 3

Coupled-Mode Theory

Even in the linear regime, Maxwell's equations are almost never solvable in complex geometries without resorting to computer simulations, and matters become even more difficult when nonlinear interactions are included. Although computational models are powerful and indispensable tools, it is useful to also have analytical insight into the *general* behaviors of a particular class of systems, rather than the solutions for very specific problems, in order to develop a real understanding and to guide the design of devices to achieve specific behaviors. A key way to obtain this understanding is to exploit that many of the most interesting problems contain small parameters and weak interactions: weak nonlinearities, slow loss mechanisms, and resonant cavities that couple weakly (escape slowly) into input/output channels like waveguides. The existence of small parameters allows us to apply *perturbative* techniques, in which these small parameters are treated analytically. For example, this allows us to solve the linear Maxwell's equations first (perhaps with a computer) and then analytically insert the effects of nonlinearities. For systems of resonant cavities (which trap light for a long time in a small volume) and waveguides (which transport light to and from the cavities), it allows us to solve for the harmonic modes of the isolated components (cavities and waveguides) first, and then analytically predict what will happen when then are combined—-a very powerful approach because we can analytically explore many possible combinations and choose the best one.

Generally speaking, there are two classes of perturbation methods one can use. One approach would be to start with the the exact equations involved, in this case Maxwell's equations, and apply perturbation methods to them to expand the solution in the small parameter and obtain explicit partial differential equations and formulas. An alternative and somewhat lesser-known approach, however, is to take a step back and to apply as *little* knowledge of the physics as possible: we try to write the most general possible "coupled-mode" equations relating the amplitudes of the eigenmodes in the cavity and waveguides, and then constrain them using fundamental principles like conservation of energy. This latter approach, sometimes called *temporal coupled-mode theory* (TCMT), allows one to describe a set of universal behaviors for all systems in a certain class (for example, all ways of coupling a single waveguide channel to a single resonant cavity), often independently of whether we are talking about electromagnetism, sound waves, or even quantum mechanics. In this work, we use *both* approaches. First we use the latter approach to model the system using temporal coupled-mode theory in the form of a set of ODEs, in order to describe the general possible behaviors of any nonlinear frequency-conversion device in a certain broad class, and then we use the former approach and apply perturbation methods to Maxwell's equations to find the exact coupling coefficients (to plug into the coupled-mode equations) in terms of the solutions of linear electromagnetism (provided by the computer).

TCMT has a long history and appears in many ways in different fields. In quantum mechanics, it is essentially what is known as Breit-Wigner scattering theory [36]. In electromagnetism, it originated in the theory of microwave devices due to work by Haus and others [73]. Similar ideas appear in describing mechanical resonant devices. Because TCMT derives a universal description for a certain class of devices, the equations obtained appear over and over again in the literature, but are not always derived at the same level of generality; for example, much of the work on nonlinear optical cavities superficially derives the equations only for light bouncing back and forth between mirrors on a macroscopic scale [147, 174], even though the same equations are much more general (and apply equally well to nanophotonic

devices where ray-optic descriptions are invalid).

In this chapter, we review a basic derivation of the TCMT equations for a few key example systems, starting with a linear case and moving on to a nonlinear case relevant to the third-harmonic generation discussed in the next chapter. For further information, see e.g. [73, 83]. As we will see, one obtains a description of all possible behaviors of a given class of coupled-resonator devices, parametrized by a small number of unknowns like the frequencies and lifetimes of the cavity modes which must be supplied from separate calculations for a specific geometry.

3.1 The two-port linear case

Consider a simple model problem, depicted schematically in Fig. 3-1(a): a (linear) resonant cavity with a mode at a frequency ω_0 that decays slowly into two input/output channels (left and right) with given exponential decay rates.

There are many examples of such a situation. For instance, consider the (one-dimensional) circumstance of light bouncing back and forth between two parallel partially reflecting mirrors, as shown in Fig. 3-1(b), called a Fabry-Perot cavity (or sometimes an "etalon"). Such a cavity will support resonant modes at certain frequencies (corresponding round-trip phases that are multiples of 2π) that decay slowly to both the left and the right. In one dimension (or for large flat mirrors in 3d) the light escaping (or entering) at left and right will only do so at normal incidence, corresponding to single-mode channels at left and right. In nanophotonics, microcavities can be formed in many ways, such as by ring resonators (waveguides along a circular path), and input-output channels might consist of single-mode waveguides.

Now we wish to analyze circumstances like those in Fig. 3-1(a) in general, including (but not limited to) specific realizations like the one in Fig. 3-1(b). We will first derive the coupled-mode equations and will then analyze the resulting possible behaviors in the next two subsections.

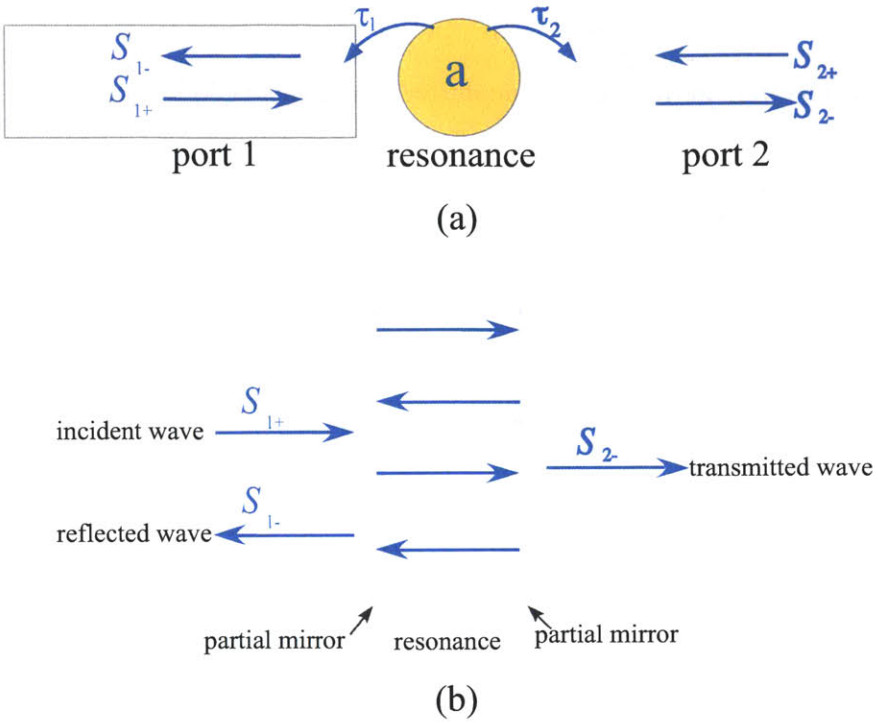

Figure 3-1: *a)* A schematic illustration of a two-port linear resonant system. A resonant cavity is coupled to two ports with lifetimes τ_1 and τ_2. *b)*: A Fabry-Paerot cavity: two parallel partial mirrors form a resonant cavity in the space between them. This is an example of a two-port linear system that is modeled using coupled-mode theory in this section.

3.1.1 Derivation of coupled-mode equations

First, what are the degrees of freedom in such a system? We are assuming a system with a single resonant mode, and such a mode is described by a single complex amplitude a. Second, we assumed that the input/output channels were single-mode also, so they are described again by complex amplitudes of the incoming and outgoing waves at each port. We denote these amplitudes by s_{1+} and s_{2+} for the incoming-wave amplitudes at ports 1 and 2, respectively, and by s_{1-} and s_{2-} for the outgoing wave amplitudes. Next, we have to choose a normalization for these amplitudes, and it is convenient to normalize a so that $|a|^2$ is the *energy* stored in the cavity by a mode of amplitude a, while we normalize $s_{k\pm}$ so that $|s_{k\pm}|^2$ is the incoming or outgoing *power* in the waveguide modes. The phase choice of these

amplitudes is still arbitrary and will be set below in order to put the equations in a simple form.

The most central assumption of TCMT is that of resonance: one or more trapped modes in the cavity that decay slowly (compared to the frequency) into the surrounding radiative/loss channels. We assume that the mode in the cavity is a resonant mode. A resonant mode is one that decays very slowly and that can be modeled as an exponential decay with a lifetime much larger than the period $1/\omega_0$. (Mathematically, a resonance is a pole in the scattering matrix at a point in the complex plane slightly below the real axis. The resonant "mode" is really a saddle-point approximation to the local field pattern, really a integral over all real frequencies superimposing a continuum of radiation modes, which is accurate when the pole is close to the real axis, corresponding to decay lifetimes that are much larger than the optical period of the resonance.) To start with, let us consider the simple case of a resonant mode that has been excited and then allowed to decay with no input sources of energy ($s_{k+} = 0$). In this case, by assumption we will have $a = Ae^{i\omega_0 t - t/\tau}$, an exponentially decaying sinusoid with some amplitude A, resonant frequency ω_0, and lifetime τ. This satisfies an ODE:

$$\frac{da}{dt} = -i\omega_0 a - a/\tau = -i\omega_0 a - a/\tau_1 - a/\tau_2,$$

where we have split the decay rate $1/\tau = 1/\tau_1 + 1/\tau_2$ into separate decay rates to the left and right channels. How can this equation be modified in the presence of input waves, i.e. nonzero s_{k+}?

The second key assumption is the linear time invariance. Under this assumption, the most general possible modification of the a equation by s_{k+} is of the form:

$$\frac{da}{dt} = -i\omega_0 a - a/\tau_1 - a/\tau_2 + \hat{\nu}_1 * s_{1+} + \hat{\nu}_2 * s_{2+},$$

where s_{k+} is convolved ($*$) with some $\hat{\nu}_k(t)$. We could also Fourier-transform this to

frequency domain, yielding

$$-i\omega a = -i\omega_0 a - a/\tau_1 - a/\tau_2 + \nu_1(\omega)s_{1+} + \nu_2(\omega)s_{2+},$$

where $\nu_k(\omega)$ is a frequency-dependent coefficient. However, since this is a resonant system, sources at frequencies very different from ω_0 will not be relevant, so it is a good approximation to replace $\nu_k(\omega)$ with $\nu_k(\omega_0) = \nu_k$. (More formally, we could Taylor expand ν_k around ω_0 and we would find that the other terms are all higher-order in $1/\tau$ after the analysis below shows that $\nu_k \sim \sqrt{1/\tau}$.) The corresponding time-domain equation now has simply a constant coefficient for each s_{k+} instead of a convolution:

$$\frac{da}{dt} = -i\omega_0 a - a/\tau_1 - a/\tau_2 + \nu_1 s_{1+} + \nu_2 s_{2+}. \tag{3.1}$$

Similarly, the most general possible linear time-invariant equation for the outgoing wave amplitudes s_{k-}, again using resonance to replace convolutions with a constant coefficient corresponding to the resonant frequency, is of the form:

$$s_{k-} = \varsigma_k s_{k+} + \gamma_k a, \tag{3.2}$$

in terms of unknown coefficients ς_k and γ_k.

These ODEs are still too general to be useful, because of all of the unknown parameters ν_k, ς_k, and γ_k. However, it turns out that *all* of these unknowns are determined by τ_k once one assumes that two additional physical principles apply: conservation of energy and time-reversal symmetry (true in many wave problems, including electromagnetism). (Energy dissipation will be added in later, under the assumption that it is slow.)

First, using conservation of energy, one can find the γ_k coefficients. Consider the simplified case where the cavity is decoupled from the second waveguide, $\tau_2 \to \infty$, and suppose no input power, $s_{k+} = 0$. Then all the energy from the cavity, $|a|^2$ must

decay into the outgoing power $|s_{1-}|^2$:

$$-\frac{d|a|^2}{dt} = \frac{2}{\tau_1}|a|^2 = |s_{1-}|^2 = |\gamma_1|^2|a|^2. \tag{3.3}$$

and so $|\gamma_1|^2 = 2/\tau_1$. Since one can pick the phase of s_{1-} arbitrarily, one can then choose $\gamma_1 = \sqrt{2/\tau_1}$. By symmetry of the problem, $\gamma_2 = \sqrt{2/\tau_2}$.

The constants ς_k and nu_k can be determined by time-reversal symmetry, which means that if you run a solution backwards in time then it also solves the same equations. (Alternatively, one could use the principle of reciprocity to obtain the same result [cite Suh & Fan].) Time-reversing the decaying solution ($s_{k+} = 0$) from the previous paragraph, we obtain a new solution $a(t) = A(0)e^{-i\omega_0 t + t/\tau}$ with input fields $s_{k+} = \sqrt{2/\tau_k}a$ and zero output ($s_{k-} = 0$). Plugging this into Equation (3.2) gives us $\varsigma_k = -1$. Finally, to determine ν_1, again take the case of $\tau - 2 \to \infty$ and plug in $A(t)$ into Eq. (3.1) which would give $\nu_1\sqrt{2/\tau_1}A = 2A/\tau_1$. Therefore, $\nu_k = \sqrt{2/\tau_k}$.

Hence, the temporal coupled mode equations for the case of a two-port linear singly-mode resonant cavity coupled to two waveguides is:

$$\frac{da}{dt} = \dot{a} = -i\omega_0 a + \sum_k \left(-a/\tau_k + \sqrt{\frac{2}{\tau_{k+}}} s_{k+} \right), \tag{3.4}$$

$$s_{k-} = -s_{k+} + \sqrt{2/\tau_k}a. \tag{3.5}$$

3.1.2 Results

Given the coupled-mode Eqs. (3.4–3.5), one can find the transmission spectrum of any weakly-coupled waveguide-cavity-waveguide system. Assuming an input s_{1+} at frequency ω and no sources at port $s_{2+} = 0$ and solving the Fourier-transformed equations, one obtains the transmission spectrum:

$$T(\omega) = \frac{|s_{2-}|^2}{|s_{1+}|^2} = \frac{\frac{4}{\tau_1\tau_2}}{(\omega - \omega_o)^2 + (\frac{1}{\tau_1} + \frac{1}{\tau_2})^2} \tag{3.6}$$

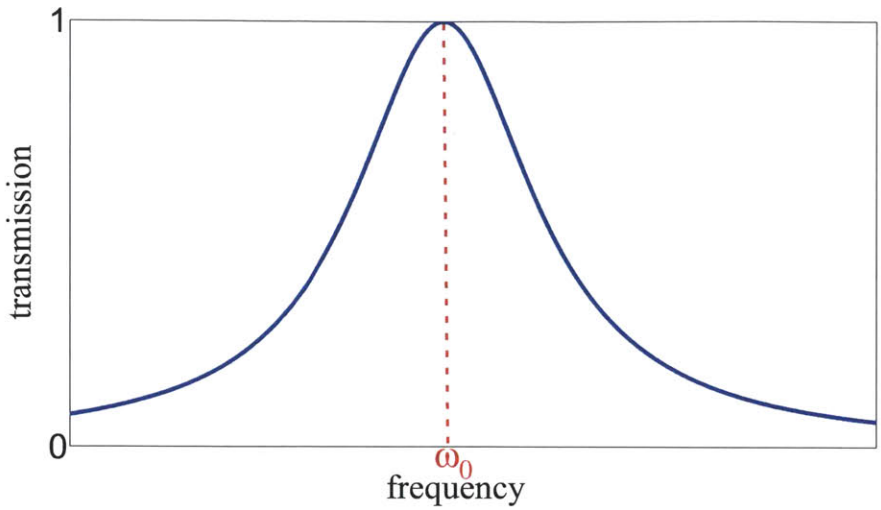

Figure 3-2: Transmission spectrum of a two-port linear system as in Eq. (3.6) for a case where $\tau_1 = \tau_2$. Complete transmission occurs at $\omega = \omega_0$

which is a Lorentzian peak with a maximum at $\omega = \omega_0$ as shown in Fig. 3-2.

For the case of $\tau_1 = \tau_2$, the peak is at 100% transmission, a remarkable and somewhat counter-intuitive result. Consider, for example, the two-mirror example of Fig. 3-1(b). Intuitively, there will always be *some* reflection of input light from the left mirror, and yet here we are predicting 100% transmission (0% reflection) whenever the two mirrors have equal reflectivity. In fact, if the left mirror is initially *less* reflective than the right mirror, *increasing* its reflectivity will cause the reflection of light from the left to *decrease*. The explanation comes from Eq. (3.5): there are two terms, one corresponding to direct reflection and one corresponding to light leaking *back* out of the cavity, and for $\tau_1 = \tau_2$ these terms exactly cancel on resonance: perfect destructive interference between the direct reflection and the cavity leakage. Amazingly, this perfect cancellation *must* occur whenever $\tau_1 = \tau_2$, regardless of the specific geometry or the physical mechanism used for the mirrors. (This situation is somewhat related to the concept of "impedance matching" in electrical engineering.)

3.1.3 Dissipation loss

Of course, real cavities inevitably have some additional dissipation loss mechanism, if only from material absorption (and also possibly radiative scattering into other directions as in the example of two mirrors with some surface roughness). As long as these losses are small (loss rates much less than ω_0), they can be incorporated into the same framework. Most simply, one could add an additional "port" 3 (or "e" for "external") corresponding to the losses.

Technically, the addition of a $1/\tau_e$ loss term might modify the other coupling coefficients, but such modifications will be a higher-order effect since it will go as the product of powers of $1/\tau$ or $\sqrt{1/\tau}$.

3.2 Nonlinear doubly-resonant cavity

Now, in preparation for the results of the following chapter, we will review the TCMT for doubly-resonant nonlinear cavity with third-harmonic generation, as derived in Ref. 169. Here, we will consider a cavity with *two* modes at resonant frequencies ω_1^{cav} and $\omega_3^{\text{cav}} \approx 3\omega_1^{\text{cav}}$, with corresponding amplitudes $a_{1,3}$. The "fundamental" frequency ω_1^{cav} is coupled to a single-mode channel with incoming/outgoing amplitudes $s_{1\pm}$, and the "harmonic" frequency ω_3^{cav} is coupled to a single-mode channel with incoming/outgoing amplitudes $s_{3\pm}$, as depicted in Fig. 3-3. (Physically, the s_1 and s_3 modes may simply be modes propagating at different frequencies in the same waveguide, but here we depict them as separate "channels" for clarity.)

In a linear system, this would satisfy equations very similar to those in the previous section, in which the modes at the two frequencies are completely decoupled. However, we will now modify this situation by assuming that there is a $\chi^{(3)}$-type process that introduces terms $\sim a^3$ which allows the two modes to nonlinearly couple. We will see below that this allows both an up-conversion process (third harmonic generation) as well as a down-conversion process (degenerate four-wave mixing) from ω_1^{cav} to ω_3^{cav} and back. The conversion rate will depend upon the input

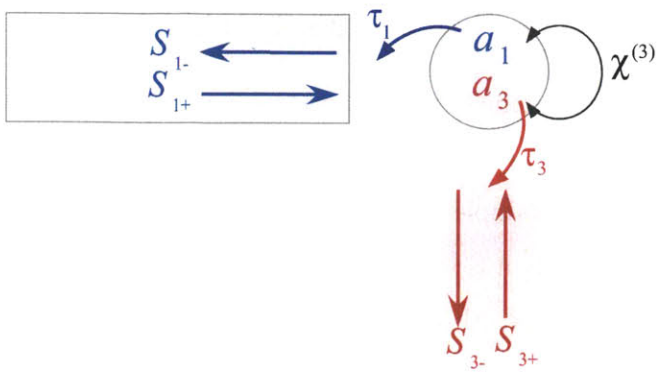

Figure 3-3: A doubly-resonant nonlinear cavity coupled to two waveguides. The $\chi^{(3)}$ nonlinearity is used for 3rd harmonic generation.

power, which determines the strength of the nonlinear terms. Intuitively, this is conceptually somewhat similar to the two-port system of the previous section, in which $1/\tau_2$ leakage to the output port is replaced by a nonlinear leakage to s_{3-}. Therefore one might guess that a similar "impedance matching" phenomenon might occur: at a certain critical input power, the nonlinear leakage rate will be exactly matched to the $1/\tau_1$ linear escape rate and 100% conversion efficiency will occur. In practice, the equations below are much more complicated than the ones in the previous section, but it turns out that this intuition is remarkably accurate: 100% conversion efficiency is possible at a certain critical input power.

3.2.1 Derivation of coupled-mode equations

One can derive the TCMT equations for Fig. 3-3 by starting with the linear TCMT equations and adding the most general possible cubic nonlinear terms, subject to a couple of assumptions. First, one assumes that the nonlinearities are weak, so that both higher-order nonlinear terms (e.g. sextic) and modifications of the $1/\tau$ or $\sqrt{1/\tau}$ pre-existing linear terms are higher-order (products of small nonlinearities and slow decay rates) and can be neglected. Second, one makes the "rotating-wave approximation": that only nonlinear processes coupling resonant modes are

significant, whereas nonlinear processes coupling to frequencies where there are no resonances are neglected. So, for example, THG processes that generate $3\omega_3^{\text{cav}}$ from the second mode are negligible since there is no resonance at $3\omega_3^{\text{cav}}$. Using the rotating-wave approximation, one considers all cubic combinations of $a_{1,3}$ and $a_{1,3}^*$ terms and discards those that lie at the wrong frequencies, arriving at the following TCMT equations [169]:

$$\begin{aligned} \dot{a_1} = & \left[i\omega_1^{\text{cav}} \left(1 - \alpha_{11}|a_1|^2 - \alpha_{13}|a_3|^2\right) - \frac{1}{\tau_1} \right] a_1 \\ & - i\omega_1\beta_1(a_1^*)^2 a_3 + \sqrt{\frac{2}{\tau_{s,1}}} s_{1+} \end{aligned} \tag{3.7}$$

$$\begin{aligned} \dot{a_3} = & \left[i\omega_3^{\text{cav}}(1 - \alpha_{31}|a_1|^2 - \alpha_{33}|a_3|^2) - \frac{1}{\tau_3} \right] a_3 \\ & - i\omega_3\beta_3 a_1^3 \end{aligned} \tag{3.8}$$

Here the α and β coefficients are geometry/material-dependent constants that express the strength of various nonlinear effects for the given modes. ω_1 and ω_3 are frequencies to be explained below; think of them for now as arbitrary scale factors. The α_{ij} terms correspond to self- and cross-phase modulation effects: they clearly give rise to effective frequency shifts in the two modes. The β_i terms characterize the energy transfer between the modes: the β_3 term describes frequency up-conversion (THG) and the β_1 term describes down-conversion (DFWM).

One can use conservation of energy to relate β_1 and β_3. Consider a system with no inflow or outflow of energy: $s_{k+} = s_{k-} = 0$ and $\tau_1, \tau_3 \to \infty$. Then conservation of energy implies:

$$\begin{aligned} & \frac{d}{dt}\left(|a_1|^2 + |a_3|^2\right) = 0 \\ & a_1\dot{a}_1^* + \dot{a}_1 a_1^* + a_3\dot{a}_3^* + \dot{a}_3 a_3^* = 0 \\ & - i\omega_1\beta_1 a_1^* 3a_3 + c.c. + i\omega_3\beta_3^* a_1^{*3} a_3 + c.c. = 0 \end{aligned} \tag{3.9}$$

and so one obtains $\omega_1\beta_1 = \omega_3\beta_3^*$.

3.2.2 Coupling coefficients

The nonlinear coupling coefficients α and β are not further constrained by TCMT or general principles. In the TCMT framework, we will consider the general possible behaviors for different values of α and β (and τ, etcetera). However, eventually one would like to make predictions for specific device designs. Resonance frequencies and lifetimes can be calculated by simulations of the linear system, and in principle one could compute α and β parameters by simulating the nonlinear system. However, because nonlinear effects in electromagnetism are weak one can instead compute α and β semi-analytically, by first solving for the modes of the linear system and then applying perturbative techniques to Maxwell's equations to express the lowest-order nonlinear corrections in terms of the linear modes. This was carried out in [169], and we review the results here.

When a dielectric structure is perturbed by a small $\delta\varepsilon$, a standard result of perturbation theory states that the corresponding change $\delta\omega$ in an eigenfrequency ω to first order is [83]:

$$\frac{\delta\omega}{\omega} = -\frac{1}{2}\frac{\int d^3\mathbf{x}\delta\epsilon|\mathbf{E}|^2}{\int d^3x\epsilon|\mathbf{E}|^2} = -\frac{1}{2}\frac{\int d^3\mathbf{x}\mathbf{E}^*\cdot\delta\mathbf{P}}{\int d^3\mathbf{x}\varepsilon|\mathbf{E}|^2} \tag{3.10}$$

where E is unperturbed electric field and $\delta\mathbf{P} = \delta\epsilon\mathbf{E}$ is the change in the polarization density due to the index change. The first-order effect of a nonlinearity is computed by substituting the nonlinear polarization $\mathbf{P}$ into this equation. The only thing that requires some care is that nonlinear terms must be computed using the *real* fields, not the $\sim e^{-i\omega t}$ complex fields that are normally used for convenience in the frequency domain.

First, one computes the nonlinear first-order frequency perturbation due to the total field $\mathbf{E}$. In case of Kerr nonlinearity, $P = \epsilon\chi^{(3)}|\mathbf{E}|^2E$, where $\mathbf{E} = \mathrm{Re}[\mathbf{E}_1e^{i\omega_1 t} + \mathbf{E}_3e^{i\omega_3 t}] + (\mathrm{complexconjugate})$. Eliminating all terms except those

at the ω_1 and ω_3 frequencies, one obtains:

$$\frac{\delta\omega_1}{\omega_1} = -\frac{1}{8}\left(\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[|\mathbf{E}_1\cdot\mathbf{E}_1|^2+2|\mathbf{E}_1\cdot\mathbf{E}_1^*|^2+2(\mathbf{E}_1\cdot\mathbf{E}_1^*)(\mathbf{E}_3\cdot\mathbf{E}_3^*)\right]}{\int d^3\mathbf{x}\epsilon|\mathbf{E}_1|^2}\right. \tag{3.11}$$

$$\left.+\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[2|\mathbf{E}_1\cdot\mathbf{E}_3|^2+2|\mathbf{E}_1\cdot\mathbf{E}_3^*|^2+(\mathbf{E}_1^*\cdot\mathbf{E}_1^*)(\mathbf{E}_1^*\cdot\mathbf{E}_3)\right]}{\int d^3\mathbf{x}\epsilon|\mathbf{E}_1|^2}\right) \tag{3.12}$$

$$\frac{\delta\omega_3}{\omega_3} = -\frac{1}{8}\left(\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[|\mathbf{E}_3\cdot\mathbf{E}_3|^2+2|\mathbf{E}_3\cdot\mathbf{E}_3^*|^2+2(\mathbf{E}_1\cdot\mathbf{E}_1^*)(\mathbf{E}_3.\mathbf{E}_3^*)\right]}{\int d^3\mathbf{x}\epsilon|\mathbf{E}_3|^2}\right. \tag{3.13}$$

$$\left.+\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[2|\mathbf{E}_1\cdot\mathbf{E}_3|^2+2|\mathbf{E}_1\cdot\mathbf{E}_3^*|^2+(\mathbf{E}_1\cdot\mathbf{E}_1)(\mathbf{E}_1\cdot\mathbf{E}_3^*)\right]}{\int d^3\mathbf{x}\epsilon|\mathbf{E}_3|^2}\right) \tag{3.14}$$

Then by setting $\omega_k \to \omega_k + \delta\omega_k$ in Eq. (3.4) and comparing to Eqs. (3.7–3.8), one finds the coupling coefficients:

$$\alpha_{ii} = \frac{1}{8}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[|\mathbf{E}_i\cdot\mathbf{E}_i|^2+2|\mathbf{E}_i\cdot\mathbf{E}_i^*|^2\right]}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_i|^2\right]^2} \tag{3.15}$$

$$\alpha_{13} = \frac{1}{4}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[(\mathbf{E}_1\cdot\mathbf{E}_1^*)(\mathbf{E}_3\cdot\mathbf{E}_3^*)+|\mathbf{E}_1\cdot\mathbf{E}_3|^2+|\mathbf{E}_1\cdot\mathbf{E}_3^*|^2\right]}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_1|^2\right]^2\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_3|^2\right]^2} \tag{3.16}$$

$$\alpha_{31} = \alpha_{13} \tag{3.17}$$

$$\beta_1 = \frac{3}{8}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}(\mathbf{E}_1^*\cdot\mathbf{E}_1^*)(\mathbf{E}_1^*\cdot\mathbf{E}_3)}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_1|^2\right]^{3/2}\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_3|^2\right]^{1/2}} \tag{3.18}$$

$$\beta_3 = \frac{1}{8}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}(\mathbf{E}_1\cdot\mathbf{E}_1)(\mathbf{E}_1\cdot\mathbf{E}_3^*)}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_1|^2\right]^{3/2}\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_3|^2\right]^{1/2}} \tag{3.19}$$

There are three different α_{ij} parameters (two SPM coefficients α_{11} and α_{33} and one XPM coefficient $\alpha_{13} = \alpha_{31}$). All three values are different, in general, but are determined by similar integrals of the field patterns, produce similar frequency-shifting phenomena, and all scale inversely with some measure of the modal volume V. Therefore, in order to limit the parameter space analyzed, in the following chapter we consider the simplified case where all three frequency-shifting terms have the same strength $\alpha_{ij} = \alpha$.

Note that when $\omega_3 = 3\omega_1$, we have $\omega_1\beta_1 = \omega_3\beta_3^*$ as required by conservation of energy in the previous section.

3.2.3 Results

As is reviewed in Sec. 4.2, the nonlinear system also exhibits "full transmission" very similar to the linear system. "Full transmission" in this case would mean complete conversion so that all the input power at ω_1 is converted to the ω_3 mode and coupled back to the output port. It turns out that there is a "critical power" s_{1+}^{crit} at which this complete conversion occurs. This is the power at which the rates of down-conversion and up-conversion processes and the coupling of the modes to the ports are matched exactly so that destructive interference occurs and the reflected signal into the input channel vanishes, $s_{1-} = 0$. Again, this may seem counter-intuitive at first as one may expect the conversion efficiency increase with the input power. However, one must remember that the conversion efficiency is a result of two competing processes in the cavity: up-conversion and down-conversion and so higher input power and energy in the cavity does not necessarily mean higher efficiency since the down-conversion process may easily dominate and decrease the efficiency. Therefore, the optimal input power is the one that gives an optimal balance between all the processes involved. We will review the exact expression for this "critical power" in Sec. 4.2.

However, as is discussed in detail in the next chapter, there is a complication due to the α (SPM and XPM) terms. The simplest case is when all the α terms are (artificially) set to zero. Then one sets $\omega_1^{\text{cav}} = \omega_1$ and $\omega_3^{\text{cav}} = \omega_3 = 3\omega_1$ to the THG frequency, and increases the input power until the critical power is reached. However, when $\alpha \neq 0$, as the power is increased towards the critical power then the cavity frequencies *shift out of resonance*, spoiling THG. In order to combat this in the next chapter, we will choose $\omega_3 = 3\omega_1$ to be the desired cavity frequencies *at the critical power* (including the α terms), while ω_1^{cav} and ω_3^{cav} will be the *linear* (low-power) cavity frequencies which are *pre-shifted* from $\omega_{1,3}$ by amounts exactly chosen so that the modes shift into resonance as the input reaches the critical power. These frequency shifts greatly complicate the nonlinear dynamics of the problem, a difficulty that was neglected in previous work and which we address in

the next chapter.

Chapter 4

3rd Harmonic Generation

In this chapter, we study 3rd harmonic generation in detail. We use temporal coupled-mode theory to model this complex nonlinear process. Once we have the ODE's governing the system (independent of the exact geometry), we do a full analysis of the dynamics of this nonlinear system, its solutions and their stabilities. We show there exists a solution of 100% efficient conversion and that it is stable for a range of system parameters. Finally, we propose ways to excite this complete conversion solution in practice. This work was published in Ref. 71.

4.1 Introduction

Consider a waveguide coupled to a doubly resonant cavity with two resonant frequencies $\omega_1^{\text{cav}} = \omega_1$ and $\omega_3^{\text{cav}} = \omega_3 = 3\omega_1$ (below, we will shift ω_k^{cav} to differ slightly from ω_k), and corresponding lifetimes τ_1 and τ_3 describing their radiation rates into the waveguide (or quality factors $Q_k = \omega_k \tau_k/2$). In addition, these modes are coupled to one another via the Kerr nonlinearity. This system is depicted schematically in Fig. 4-1.

In a Kerr ($\chi^{(3)}$) medium, there is a change in the refractive index proportional to the square of the electric field; for an oscillating field at a frequency ω, this results in a shift in the index at the same frequency (self-phase modulation, SPM), generation of power at the third-harmonic frequency 3ω, and also other effects when

multiple frequencies are present [cross-phase modulation (XPM) and four-wave mixing (FWM)] [18]. When the field is confined in a cavity, restricting to a small modal volume V for a long time given by the quality factor Q (a lifetime in units of the optical period) [84], such nonlinear effects are enhanced by both the increased field strength for the same input power and by the frequency sensitivity inherent in resonant effects (since the fractional bandwidth is $1/Q$). This enhancement is exploited, for example, in nonlinear harmonic and sum-frequency generation, most commonly for $\chi^{(2)}$ effects where the change in index is proportional to the electric field (which requires a non-centrosymmetric material) [18]. One can further enhance harmonic generation by using a cavity with *two* resonant modes, one at the source frequency and one at the harmonic frequency [13, 45, 47, 119, 127, 147, 150, 203]. In this case, one must also take into account a nonlinear downconversion process that competes with harmonic generation [45, 147, 203], but it turns out to be theoretically possible to obtain 100% harmonic conversion for either $\chi^{(2)}$ ($\omega \rightarrow 2\omega$) or $\chi^{(3)}$ ($\omega \rightarrow 3\omega$) nonlinearities at a specific "critical" input power P_{crit} (both in a one-dimensional model of propagating waves for $\chi^{(2)}$ nonlinearities [173] and also in a more general coupled-mode model for either $\chi^{(2)}$ or $\chi^{(3)}$ nonlinearities [169]). In particular, my collaborators studied the harmonic-generation and downconversion processes in a broad class of model systems depicted in Fig. 4-1: a single input channel (e.g. a waveguide) is coupled to a nonlinear cavity with two resonant frequencies, where both reflected and harmonic fields are emitted back into the input channel. In this case, 100% harmonic generation at a critical power P_{crit} proportional to V/Q^3 for $\chi^{(2)}$ and V/Q^2 for $\chi^{(3)}$ was predicted [169]. However, they only looked at the steady-state solution of the system and not its dynamics or stability. Moreover, in the $\chi^{(3)}$ case there can also be an SPM/XPM effect that shifts the cavity frequencies out of resonance and spoils the harmonic-generation effect. In this thesis, we consider both of these effects, describe how to compensate for SPM/XPM, and demonstrate the different regimes of stability in such $\chi^{(3)}$ doubly resonant systems. We show that the parameters and the initial conditions must be chosen within certain regimes to obtain a stable steady state with high

conversion efficiency.

In other regimes, we demonstrate radically different behaviors: not only low-efficiency steady states, but also limit-cycle solutions where the efficiency oscillates slowly with a repetition period of many thousands of optical cycles. With infrared light, these limit cycles form a kind of optical oscillator/clock with a period in the hundreds of GHz or THz (and possibly lower, depending on the cavity parameters). Previously, limit-cycle/self-pulsing behaviors have been observed in a number of other nonlinear optical systems, such as: doubly-resonant $\chi^{(2)}$ cavities coupled by second-harmonic generation [45]; bistable multimode Kerr cavities with time-delayed nonlinearities [1]; nonresonant distributed feedback in Bragg gratings [149]; and a number of nonlinear lasing devices [180]. However, the system considered in this work seems unusually simple, especially among $\chi^{(3)}$ systems, in that it only requires two modes and an instantaneous Kerr nonlinearity, with a constant-frequency input source, to attain self-pulsing, and partly as a consequence of this simplicity the precise self-pulsing solution is quite insensitive to the initial conditions. In other nonlinear optical systems where self-pulsing was observed, other authors have also observed chaotic solutions in certain regimes. Here, we did not observe chaos for any of the parameter regimes considered, where the input was a constant-frequency source, but it is possible that chaotic solutions may be excited by an appropriate pulsed input as in the $\chi^{(2)}$ case [45].

Another interesting phenomenon that can occur in nonlinear systems is multistability, where there are multiple possible steady-state solutions that one can switch among by varying the initial conditions. In Kerr ($\chi^{(3)}$) media, an important example of this phenomenon is bistable transmission through nonlinear cavities: for transmission through a *single*-mode cavity, output can switch discontinuously between a high-transmission and a low-transmission state in a hysteresis effect that results from SPM [186]. For example, if one turns on the power gradually from zero the system stays in the low-transmission state, but if the power is increased further and then decreased to the original level, the system can be switched to the high-transmission state. This effect, which has been observed experimentally [145],

can be used for all-optical logic, switching, rectification, and many other functions [186]. In a cavity with multiple closely-spaced resonances, where the nonlinearity is strong enough to shift one cavity mode's frequency to another's, the same SPM phenomenon can lead to more than two stable solutions [53]. Here, we demonstrate a much richer variety of multistable phenomena in the doubly-resonant case for widely-separated cavity frequencies coupled by harmonic generation in addition to SPM—not only can there be more than two stable states, but the transitions between them can exhibit complicated oscillatory behaviors as the initial conditions are varied, and there are also Hopf bifurcations into self-pulsing solutions.

4.2 100% efficient solution and critical power

We begin by reviewing the results of Ref. 169 in the absence of SPM/XPM effects and considering only steady-state effects without concern for stability. This is an important first step, in which it is established that a 100%-efficiency steady-state solution exists, after which we must then analyze whether this solution is maintainable in practice.

Figure 4-2 shows the steady-state conversion efficiency ($|s_{3-}|^2/|s_{1+}|^2$) versus input power of light that is incident on the cavity at ω_1^{cav}, for $Q_1 = 1000$, $Q_3 = 3000$, $\beta_1 = (4.55985 - 0.7244i) \times 10^{-5}$ in dimensionless units (i.e. assuming negligible self- and cross-phase modulation so that $\alpha = 0$), and not considering the stability of the steady state. This plot was made by solving the coupled mode equations for the case $\alpha_{ij} = 0$ for different values of input power. And the plot is in agreement with FDTD simulations of the system. As shown by the solid red curve, as one increases the input power, the efficiency increases, peaking at 100% conversion for a critical input power $P_{\text{crit}} = |s_{1+}^{\text{crit}}|$. Going back the the coupled-mode equations, one can find this critical power in terms of the known system parameters [169]. Requiring $s_{3-} = 0$, one can use Equation (3.5) to solve a_1 in terms of s_{1+}. Plugging it into Equation (3.7) (with setting $\alpha_{ij} = 0$) and assuming steady state ($\dot{a}_k = 0$), one can

then solve for the critical power [169]:

$$|s_{1+}^{\text{crit}}| = \left(\frac{4}{|\omega_1\beta_1|^2\tau_1^3\tau_3}\right)^{1/4}. \tag{4.1}$$

The efficiency decreases if the power is too low (in the linear regime) or too high (dominated by down-conversion). The critical input power $|s_{1+}^{\text{crit}}|^2$ scales as V/Q^2, so one can in principle obtain very low-power efficient harmonic conversion by increasing Q and/or decreasing V. Including absorption or other losses decreases the peak efficiency, but does not otherwise qualitatively change this solution as we will discuss later. In addition, the harmonic conversion efficiency goes to zero if the input power is too low or too high. The former is quite intuitive, however, understanding the latter takes some more understanding of the process. It is well known that in order to get 100% transmission from an input port to an output port, the coupling rates to the two ports must be matched so that they exactly cancel the back-reflected wave. Here, the input coupling rate is $1/Q_1$ whereas the output coupling rate is determined by the nonlinear coupling (down conversion terms) β's and $1/Q_3$. And so if not matched correctly, too week or too strong non-linear coupling would give rise to back reflection instead of conversion. In other words, the critical power is the power at which there's perfect destructive interference for the reflected wave. And the dependence of this critical power on the coupling parameters becomes apparent by having β_k and τ_k in Eq. (4.1).

The system obtains a 100% efficient solution at which all the input signal is converted to the harmonic mode and coupled back to the output channel. However, this is only the beginning of the story. There are may questions and concerns that have to be addressed about this system and the complete conversion solution before we can even think of practical applications and designs. For one thing, we have come to this conclusion ignoring the self- and cross-phase modulation effects. Such effects shift the frequency of the system and take the signal out of resonance with the cavity. And this would affect the coupling of the input signal with the cavity and the conversion rate. Normally, the frequencies must be in a $1/Q$ window of each

other for couplings to occur. We address this issue in the next section. Second, there is the question of how many solutions the system has. This 100% efficient solution is one solution that we have found to the system. But this is a multi-dimensional complex system and as we will see later, can have up to five steady-state solutions. In order to utilize this 100% solution, it is important to understand the whole system and all the possible steady-state solutions which we will study in Sec. 4.4. In addition to existence of other solutions, stability becomes an important question in this highly non-linear dynamical system. If this 100% efficient solution is an unstable one, its existence means nothing for practical purposes as there is no way in practice to maintain the system in this steady-state solution. Even if the solution is stable, there are still stability questions to be addressed. For example, if there are more than one stable solutions, then exciting the right one (100% efficient one) can be quite challenging as we will see since the system can as well fall into other stable solutions that don't have high efficiency.

4.3 Self- and cross-phase modulation

There are two effects that we did not previously analyze in detail, however, which can degrade this promising solution: nonlinear frequency shifts and instability. Here, we first consider frequency shifts, which arise whenever $\alpha \neq 0$, and consider stability in the next section. The problem with the α terms is that efficient harmonic conversion depends on the cavities being tuned to harmonic frequencies $\omega_3 = 3\omega_1$; a nonlinear shift in the cavity frequencies due to self/cross-phase modulation will spoil this resonance. More specifically, there is a relative change of $1 - \alpha_{11}|a_1|^2 - \alpha_{13}|a_3|^2$ in ω_1 and $1 - \alpha_{13}|a_1|^2 - \alpha_{33}|a_3|^2$ in ω_3 as Eqs. (3.7–3.8) suggest. In principle, there is a straightforward solution to this problem, as depicted in Fig. 4-3. Originally (for $\alpha = 0$), the cavity was designed to have the frequency ω_1 in the linear regime, but with $\alpha \neq 0$ the effective cavity frequency ω_1^{NL} (including self/cross-phase modulation terms) is shifted away from the design frequency as shown by the blue line. Instead, we can simply design the linear cavity to have a frequency ω_1^{cav} slightly different

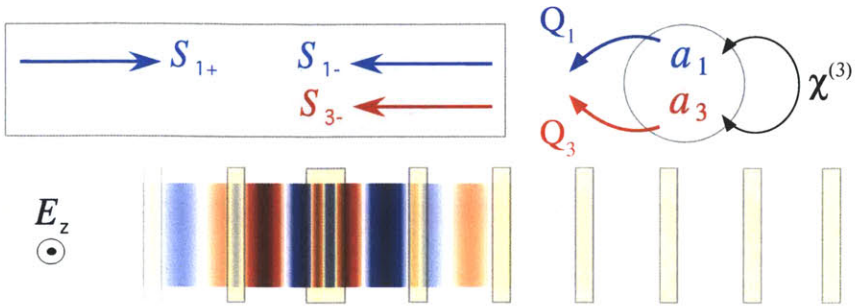

Figure 4-1: *Top:* Schematic of general scheme for third-harmonic generation, and dynamical variables for coupled-mode equations: a single input/output channel (with incoming/outgoing field amplitudes $s_\pm$) is coupled to a resonant cavity with two modes at frequencies ω_1 and $3\omega_1$ (and corresponding amplitudes a_1 and a_3). The two resonant modes are nonlinearly coupled by a Kerr ($\chi^{(3)}$) nonlinearity. *Bottom:* An example realization [169], in one dimension, using a semi-infinite quarter-wave stack of dielectric layers with a doubled-layer defect (resonant cavity) that is coupled to incident plane waves; the electric field of a steady-state $3\omega_1$ solution is shown as blue/white/red for negative/zero/positive.

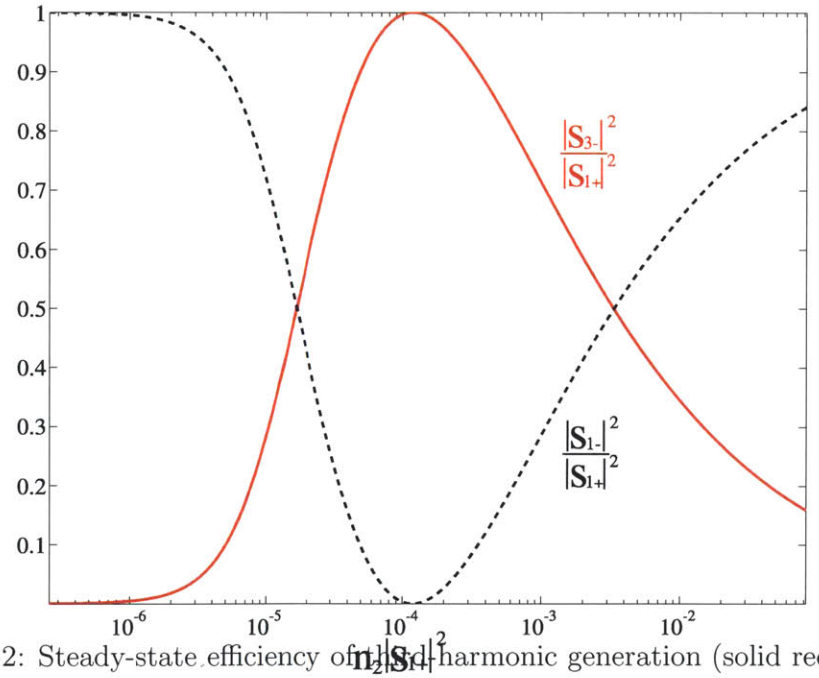

Figure 4-2: Steady-state efficiency of third-harmonic generation (solid red line) from Ref. 169, for $\alpha = 0$ (no self-phase modulation), as a function of input power $|s_{1+}|^2$ scaled by the Kerr coefficient $n_2 = 3\chi^{(3)}/4c\varepsilon$. The reflected power at the incident frequency ω_1 is shown as a dashed black line. There is a critical power where the efficiency of harmonic generation is 100. The parameters used in this plot are $Q_1 = 1000$, $Q_3 = 3000$, $\beta_1 = (4.55985 - 0.7244i) \times 10^{-5}$ in dimensionless units o f $\chi^{(3)}/V\epsilon$.

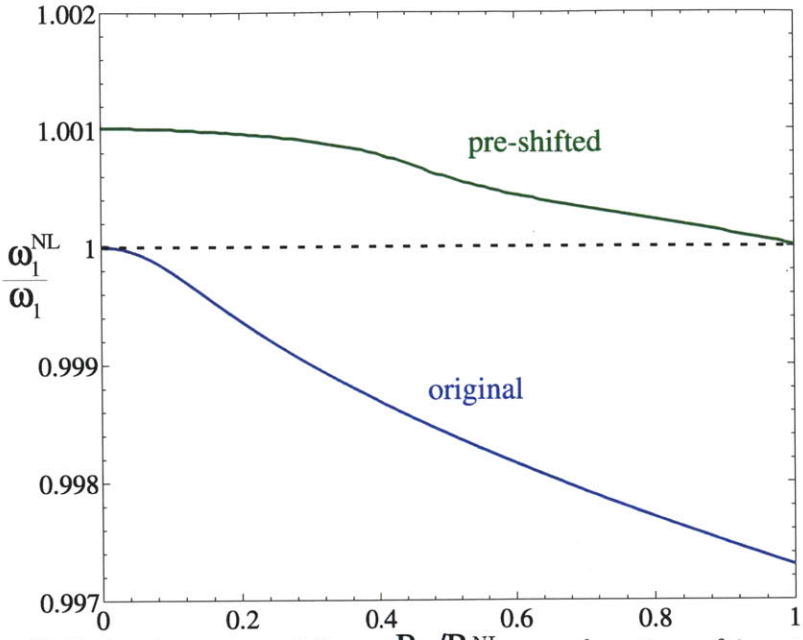

Figure 4-3: Shift in the resonant frequency ω_1^{NL} as a function of input power, due to self/cross-phase modulation. (There is an identical shift in ω_3^{NL}.) If the cavity is designed so that the linear ($P_{\mathrm{in}} \to 0$) frequencies are harmonics, the nonlinearity pushes the system out of resonance (lower blue line) as the power increases to the critical power for 100% efficiency. This is corrected by pre-shifting the cavity frequencies (upper green line) so that the nonlinear frequency shift pushes the modes into resonance at P_{crit}.

from the operating frequency ω_1, so that self/cross-phase modulation shifts ω_1^{NL} exactly to ω_1 at the critical input power, as depicted by the green line in Fig. 4-3. Exactly the same strategy is used for ω_3^{NL}, by pre-shifting ω_3^{cav}.

More precisely, to compute the required amount of pre-shifting, we examine the coupled-mode equations Eqs. (3.7–3.8). First, we solve for the critical power P_{crit} assuming $\alpha = 0$ as we did in Sec. 4.2, and obtain the corresponding critical cavity fields using Eq. (3.5):

$$\left|a_1^{\mathrm{crit}}\right|^2 = \left(\frac{1}{\omega_1^2 \left|\beta_1\right|^2 \tau_3 \tau_{1,s}}\right)^{1/2} \tag{4.2}$$

$$\left|a_3^{\mathrm{crit}}\right|^2 = \left(\frac{\omega_3 \beta_3 \tau_3}{(\omega_1 \beta_1 \tau_{1,s})^3}\right)^{1/2}. \tag{4.3}$$

Then, we substitute these critical fields into the coupled-mode equations for $\alpha \neq 0$,

and solve for the new cavity frequencies ω_k^{cav} so as to cancel the α terms and make the a_k^{crit} solutions still valid. This yields the following transformation of the cavity frequencies:

$$\omega_1^{\text{cav}} = \frac{\omega_1}{1 - \alpha_{11}|a_1^{\text{crit}}|^2 - \alpha_{13}|a_3^{\text{crit}}|^2}$$
$$\omega_3^{\text{cav}} = \frac{\omega_3}{1 - \alpha_{13}|a_1^{\text{crit}}|^2 - \alpha_{33}|a_3^{\text{crit}}|^2}.$$

By inspection, when substituted into Eqs. (3.7–3.8) at the critical power, these yield the same steady-state solution as for $\alpha = 0$. (There are two other appearances of ω_1 and ω_3 in the coupled-mode equations, in the β_k terms, but we need not change these frequencies because that is a higher-order effect, and the derivation of the coupled-mode equations considered only first-order terms in $\chi^{(3)}$.) There are some limitations on how much the pre-shifting the cavity can help overcome the self- and cross-phase modulations. As we see in Figure 4-3, the ω_{NL} gradually drops from ω_{cav} to ω_1 as we increase the power and reaches exactly at the desired ω_1 and ω_3 at the critical power. However, the discrepancy between ω_1 and ω_{cav} can not be more than $1/Q_1$. Because if it is, then the input signal won't couple to the cavity at the very beginning. The other limitation is the materials; if the frequency shifting is any higher than 1%, there'll be material meltdown. Now depending on how high Q_1 is, one of the two limitations above is the limiting factor and should be taken into account when designing the system.

4.4 Solutions and their stability analysis

Now it remains to find all the possible solutions to the system for a set of parameters and do their stability analysis. First, we reduce the space of parameters by doing some scaling and invariance analysis. The nonlinear dynamics turn out to depend only on four dimensionless parameters: $\tau_3/\tau_1 = Q_3/3Q_1$, α_{11}/β_1, α_{33}/β_1, and $\alpha_{13}/\beta_1 = \alpha_{31}/\beta_1$. The overall scale of Q, α, etcetera, merely determines the absolute scale for the power requirements: it is clear from the equations that multiplying all

α and β coefficients by an overall constant K can be compensated by dividing all a and s amplitudes by $\sqrt{K}$ [which happens automatically for s at the critical power by Eq. (4.1)]; the case of scaling $\tau_{1,3}$ by an overall constant is more subtle and is considered below. As mentioned above, for simplicity we take $\alpha_{11} = \alpha_{33} = \alpha_{13} = \alpha_{31} = \alpha$. Therefore, in the subsequent sections we will analyze the asymptotic efficiency as a function of τ_3/τ_1 and α/β_1.

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

To understand the dynamics and stability of the nonlinear coupled-mode equations, we apply the standard technique of identifying fixed points of the equations and analyzing the stability of the linearized equations around each fixed point [191].

By a "fixed point," we mean a steady-state solution corresponding to an input frequency ω_1 ($s_{1+} \sim e^{-i\omega_1 t}$) and hence $a_1(t) = A_1 e^{-i\omega_1 t}$ and $a_3(t) = A_3 e^{-i3\omega_1 t}$ for some unknown constants A_1 and A_3. [An input frequency ω_1 can also generate higher harmonics, such as $9\omega_1$ or $5\omega_1$, but these are negligible: both because they are higher-order effects ($\sim [\chi^{(3)}]^2$, and all such terms were dropped in deriving the coupled-mode equations), and because we assume there is no resonant mode present at those frequencies.] By substituting this steady-state form into Eqs. (3.7–3.8), one obtains two coupled polynomial equations whose roots are the fixed points. Given the steady-state solutions (the roots), their stability is determined by linearizing the original equations around these points to a first-order linear equation of the form $d\mathbf{x}/dt = A\mathbf{x}$; a stable solution is one for which the eigenvalues of A have negative real parts (leading to solutions that decay exponentially towards the fixed point) [191]. To study the stability of these fixed points or in fact any other fixed point a_{k0}, we consider small perturbations $a_{k0} \to a_{k0} + \delta a_k$ about the fixed points:

$$\frac{d\delta a_1}{dt} = -\frac{1}{\tau_{s,1}}(\delta a_1 + a_{10}) + i\omega_1(\alpha_{11}\,|a_{10}|^2 + \alpha_{13}\,|a_{30}|^2 - \alpha_{11}\,|a_{10}+\delta a_1|^2 - \alpha_{13}\,|a_{30}+\delta a_3|^2)(a_{10}+\delta_{a10})$$

$$-i\omega_1\beta_1(a_{10}^* + \delta a_1^*)^2(\delta a_3 + a_{30}) + \sqrt{\frac{2}{\tau_{s,1}}}s_{1+}$$

$$\frac{d\delta a_3}{dt} = -\frac{1}{\tau_{s,3}}(\delta a_3 + a_{30}) + i\omega_3(\alpha_{33}\,|a_{30}|^2 + \alpha_{13}\,|a_{10}|^2 - \alpha_{33}\,|a_{30}+\delta a_3|^2 - \alpha_{13}\,|a_{10}+\delta a_1|^2)(a_{30}+\delta_{a30})$$

$$-i\omega_3\beta_3(\delta a_1 + a_{10})^3$$

We linearize the equations above about the fixed points ($\delta a_1 = \delta a_3 = 0$). Taking the linearized equations and their complex conjugates gives a 4×4 linear system of ODE's:

$$\frac{d}{dt}\begin{pmatrix} \delta a_1 \\ \delta a_1^* \\ \delta a_3 \\ \delta a_3^* \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}\begin{pmatrix} \delta a_1 \\ \delta a_1^* \\ \delta a_3 \\ \delta a_3^* \end{pmatrix}$$

Where,

$$a_{11} = -\frac{1}{\tau_{s,1}}, \quad a_{12} = -i\omega_1\alpha_{11}a_{10}^2 - 2i\omega_1\beta_1 a_{10}^* a_{30}$$

$$a_{13} = -i\omega\beta_1 a_{10}^{*2} - i\omega_1\alpha_{13}a_{30}^* a_{10}, \qquad a_{14} = -i\omega_1\alpha_{13}a_{30}a_{10}$$

$$a_{21} = a_{12}^*, \quad a_{22} = a_{11}^*, \quad a_{23} = a_{14}^*, \quad a_{24} = a_{13}^*$$

$$a_{31} = -i\omega_3\alpha_{13}a_{10}^*a_{30} - i3\omega_3\beta_3 a_{10}^2, \quad a_{32} = -i\omega_3\ \alpha_{13}a_{10}a_{30},$$

$$a_{33} = -\frac{1}{\tau_3}, \quad a_{34} = -i\omega_3\alpha_{33}a_{30}^2$$

$$a_{41} = a_{32}^*, \quad a_{42} = a_{31}^* \quad a_{43} = a_{34}^*, \quad a_{44} = a_{33}^*$$

The steady state solutions are stable only if all the four eigenvalues of the matrix above have negative real parts. A very simple Matlab code can tell us whether the system is stable or not.

We already know one of the fixed points from the previous section, the 100% efficiency solution. Let's first study the stability of this specific solution of interest: We found the "critical" field amplitudes $|a_k|$ in Eqs. (4.2–4.3). Let us first consider the case of $\alpha_{ij} = 0$. It can be verified easily, by a simple function in Matlab, that this system is stable for $\tau_1 > \tau_3$ and unstable otherwise. Luckily and interestingly, however, it turns out the unstable solutions are limit cycles with very high average efficiency. Therefore, as long as we are fixable about having small oscillations in the output power, we can still obtain high efficiency conversion even for the regime $\tau_1 < \tau_3$. This alternative is very important in practice since the regime $\tau_1 < \tau_3$ is much easier to design and obtain (specially in the case of ring resonators).

Although the 100% efficient solution is stable for $\tau_1 > \tau_3$, for practical purposes it is important to study how "robust" this stability is. In other words, for small perturbations and discrepancies in the input power or frequency, how the stability of this solution and its efficiency change. In Fig. 4-4, we plot the efficiency for different input power and frequency detunings. We observe that the system maintains a high-efficiency stable solution around the origin (critical power, zero frequency discrepancy). We see a sudden drop in the efficiency for larger power detunings (the blue area); however this appears for power detunings of about 50% which is quite large. This is not a concern for practical purposes. And therefore, we can reasonably conclude that the 100% efficient solution is robust.

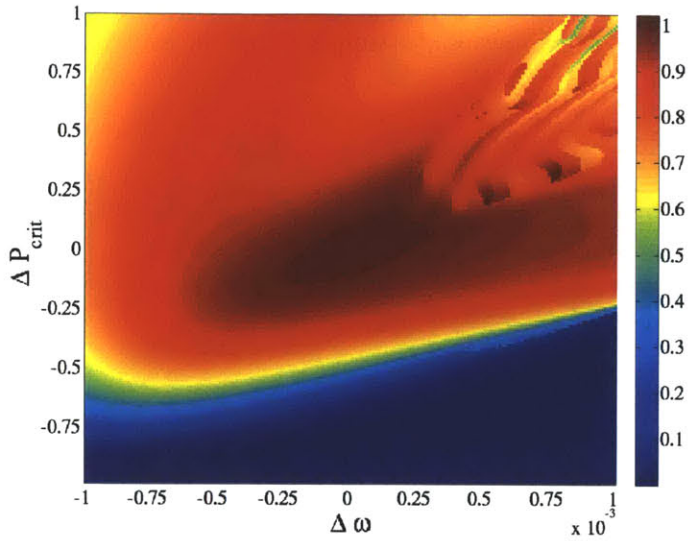

Figure 4-4: Efficiency vs. power and frequency detuning in the case where the system is stable at the origin.

In the case where $\alpha_{ij} \neq 0$, there are similar dynamics. The 100% efficient solution is the only stable solution when $\alpha_{ij} = 0$. However, additional stable solutions with different efficiencies show up depending on the ratio α/β for non-zero α_{ij}. In such a regime, the system could be singly, doubly-, or multi-stable. This on one hand makes the dynamics of the system more interesting to study, but practically makes reaching the 100% efficient solution harder since there are other stable solutions in the phase space.

But to fully characterize the system one would like to know all of the fixed points (both stable and unstable). We solved these polynomial equations using Mathematica, which is able to compute all of the roots, but some transformations were required to put the equations into a solvable form.

To find these solvable polynomials, we enforce the steady state condition on Eqs. (3.7–3.8) to get a set of coupled polynomial equations, whose roots we compute using Mathematica. However, some transformations were required in order to put the equations into a solvable form.

In particular, we eliminated the complex conjugations by writing $A_k = r_k e^{i\phi_k}$ and

assuming (without loss of generality) that s_{1+} is real. Multiplying Eq. (3.7) by $e^{-i\phi_1}$ and Eq. (3.8) by $e^{-i3\phi_1}$ allows us to simply solve the system for $e^{-i\phi_1}$ and $e^{i(\phi_3-3\phi_1)}$:

$$
\begin{aligned}
e^{i\phi_3} &= \frac{i\omega_3\beta_3 r_1^3}{r_3\left[i(\omega_3^{\mathrm{cav}}(1-\alpha(r_1^2+r_3^2))-\omega_3)-\frac{1}{\tau_3}\right]} \\
e^{i(\phi 3-3\phi_1)} &= \frac{r_1}{s_{1+}}\left[\frac{-\omega_1\omega_3\beta_1\beta_3 r_1^4}{i[\omega_3^{\mathrm{cav}}-\omega_3-\omega_3^{\mathrm{cav}}\alpha(r_1^2+r_3^2)]-1/\tau_3}+i(\omega_1^{\mathrm{cav}}-\omega_1+\omega_1^{\mathrm{cav}}\alpha(r_1^2+r_3^2))+1/\tau_1\right]
\end{aligned}
\tag{4.4}
$$

Requiring the magnitude of these two quantities to be unity yields two polynomials in $x = r_1^2$ and $y = r_3^2$, which Mathematica can handle:

$$
\begin{aligned}
0 &= \omega_3^2\beta_3^2 x^3 - \left[1/\tau_3^2+\left[w_3^{\mathrm{cav}}(1-\alpha(x+y))-\omega_3\right]^2\right]^2 y \\
0 &= 1/\tau_3^2+(\omega_3^{\mathrm{cav}}(1-\alpha(x+y))-\omega_3)^2-\frac{x}{s_{1+}^2}\left[(-\omega_1\omega_3\beta_1\beta_3 x^2-\frac{1}{\tau_1\tau_3}+\right. \\
&\quad (\omega_3^{\mathrm{cav}}(1-\alpha(x+y\ ))-\omega_3)(\omega_1^{\mathrm{cav}}(1-\alpha(x+y))-\omega_1))^2+ \\
&\quad \left.\left(\frac{\omega_3^{\mathrm{cav}}(1-\alpha(x+y))-\omega_3}{\tau_1}+\frac{\omega_1^{\mathrm{cav}}(1-\alpha(x+y))-\omega_1}{\tau_3}\right)^2\right]
\end{aligned}
\tag{4.5}
$$

The resulting polynomial is of an artificially high degree, resulting in spurious roots, but the physical solutions are easily identified by the fact that x and y must be real and non-negative. And using Mathematica's "Solve" function, we can find roots of this set of polynomials. Each computation would take about a minute to finish. However, we should also note that this root-finding process is highly sensitive to roundoff error [160], independent of the physical stability of the solutions, but we dealt with that problem by employing 50 decimal places of precision.

As mentioned above, the dynamics are independent of the overall scale of $\tau_{1,3}$, and depend only on the ratio τ_3/τ_1. This can be seen from the equations for $A_{1,3}$, in which the $\omega_{1,3}$ oscillation has been removed. In these equations, if we multiply τ_1 and τ_3 by an overall constant factor K, the $A_{1,3}$ equations are invariant if we rescale $A_1 \to A_1/\sqrt{K}$, $A_3 \to A_3/\sqrt{K}$, rescale time $t \to Kt$, and rescale the input $s_{1+} \to s_{1+}/K$ [which happens automatically for the critical power by Eq. (4.1)]:

$$\frac{a_1\sqrt{k}}{kdt} = \left[i\frac{\omega_1^{\text{cav}}}{k}\left(1 - k\alpha_{11}\frac{|a_1|^2}{k} - k.\alpha_{13}\frac{|a_3|^2}{k}\right) - \frac{1}{k.\tau_1}\right]\frac{a_1}{\sqrt{k}} - i\frac{\omega_1}{k}k.\beta_1\frac{(a_1^*)^2 a_3}{k^{3/2}} + \sqrt{\frac{2}{k\tau_{s,1}}}\frac{s_{1+}}{k}$$

$$\rightarrow \frac{1}{k^{3/2}}[\dot{a_1}] = \frac{1}{k^{3/2}}\left[\left[i\omega_1^{\text{cav}}\left(1 - \alpha_{11}|a_1|^2 - \alpha_{13}|a_3|^2\right) - \frac{1}{\tau_1}\right]a_1 - i\omega_1\beta_1(a_1^*)^2 a_3 + \sqrt{\frac{2}{\tau_{s,1}}}s_{1+}\right]$$

$$\rightarrow \dot{a_1} = \left[i\omega_1^{\text{cav}}\left(1 - \alpha_{11}|a_1|^2 - \alpha_{13}|a_3|^2\right) - \frac{1}{\tau_1}\right]a_1 - i\omega_1\beta_1(a_1^*)^2 a_3 + \sqrt{\frac{2}{\tau_{s,1}}}s_{1+} \tag{4.6}$$

Note also that the conversion efficiency $|s_{3-}/s_{1+}|^2 = (2/\tau_3)|A_3/s_{1+}|^2$ is also invariant under this rescaling by K. That is, the powers and the timescales of the dynamics change if you change the lifetimes, unsurprisingly, but the steady states, stability, etcetera (as investigated in the next section) are unaltered.

The results of this fixed-point and stability analysis are shown in Fig. 4-5 as a "phase diagram" of the system as a function of the relative lifetimes $\tau_3/\tau_1 = 3Q_3/Q_1$ and the relative strength of self-phase-modulation vs. four-wave mixing α/β_1. Our original 100%-efficiency solution is always present, but is only stable for $\tau_3 < \tau_1$ and becomes unstable for $\tau_3 > \tau_1$. The transition point, $\tau_3 = \tau_1$, corresponds to equal energy $|a_1|^2 = |a_3|^2$ in the fundamental and harmonic mode at the critical input power. The unstable region corresponds to $|a_3|^2 > |a_1|^2$ (and the down-conversion term is stronger than the up-conversion term)—intuitively, this solution is unstable because, if any perturbation causes the energy in the harmonic mode to decrease, there is not enough pumping from up-conversion to bring it back to the 100%-efficiency solution. Conversely, in the stable $|a_3|^2 < |a_1|^2$ ($\tau_3 < \tau_1$) regime, the higher-energy fundamental mode is being directly pumped by the input and can recover from perturbations. Furthermore, as α/β_1 increases, additional lower-efficiency stable solutions are introduced, resulting in regimes with two (doubly stable) and three (triply stable) stable fixed points. These different regimes are explored in more detail via bifurcation diagrams below, and the excitation of the different stable solutions is considered in the next section.

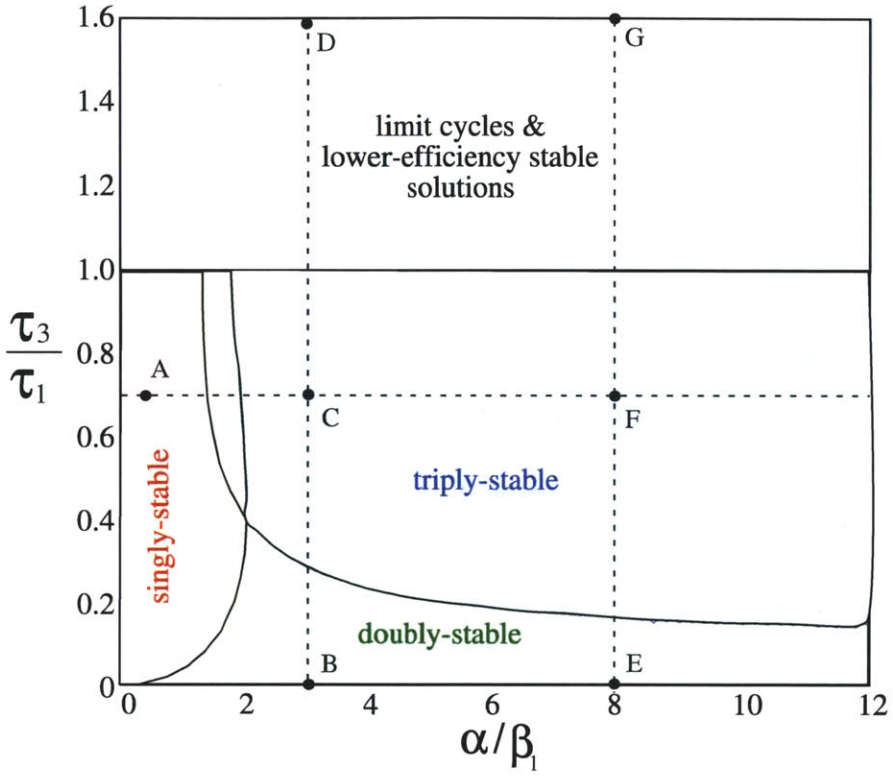

Figure 4-5: Phase diagram of the nonlinear dynamics of the doubly-resonant nonlinear harmonic generation system from Fig. 4-1 as a function of the relative cavity lifetimes ($\tau_3/\tau_1 = 3Q_3/Q_1$) and the relative strength of SPM/XPM vs. harmonic generation (α/β_1) for input power equal to the critical power for 100% efficiency. For $\tau_3 < \tau_1$ there is always one *stable* 100%-efficiency solution, and for nonzero α the system may have additional stable solutions. For $\tau_3 > \tau_1$ the 100%-efficiency solution becomes unstable, but there are limit cycles and lower-efficiency stable solutions. Various typical points A–G in each region are labeled for reference in the subsequent figures.

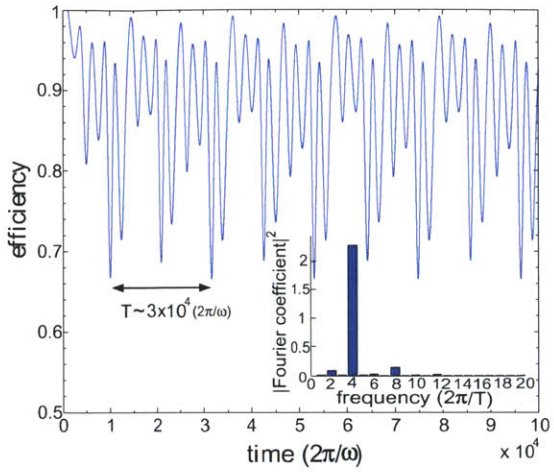

Figure 4-6: An example of a limit-cycle solution, with a periodically oscillating harmonic-generation efficiency as a function of time, corresponding to point D in Fig. 4-5. Perturbations in the initial conditions produce only phase shifts in the asymptotic cycle. Here, the limit cycle has a period of around 3×10^4 optical cycles. *Inset:* Square of Fourier amplitudes (arbitrary units) for each harmonic component of the limit cycle in the Fourier-series expansion of the $|A_3|$.

For $\tau_3 > \tau_1$, the 100%-efficiency solution is unstable, but there are lower-efficiency steady-state solutions and also another interesting phenomenon: limit cycles. A limit cycle is a stable oscillating-efficiency solution, one example of which (corresponding to point D in Fig. 4-5) is plotted as a function time in Fig. 4-6. (In general, the existence of limit cycles is difficult to establish analytically [191], but the phenomenon is clear in the numerical solutions as a periodic oscillation insensitive to the initial conditions). In fact, as we shall see below, these limit cycles result from a "Hopf bifurcation," which is a transition from a stable fixed point to an unstable fixed point and a limit cycle [190]. In this example at point D, the efficiency oscillates between roughly 66% and nearly 100%, with a period of several thousand optical cycles. As a consequence of the time scaling described in the last paragraph of the previous section, the period of such limit cycles is proportional to the τ's. If the frequency ω_1 were 1.55 μm, for a Q_1 of 500 optical cycles, this limit cycle would have a frequency of around 70 GHz, forming an interesting type of optical "clock" or oscillator. Furthermore, the oscillation is not sinusoidal and contains several higher harmonics as shown in the inset of Fig. 4-6; the dominant

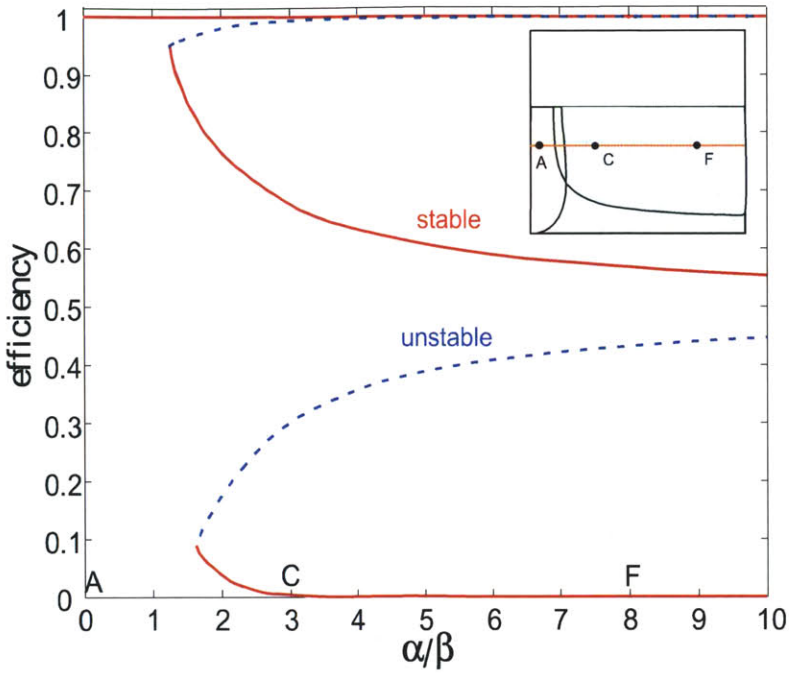

Figure 4-7: Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of α/β_1 for a fixed $\tau_3/\tau_1 = 0.7$, corresponding to the line ACF in Fig. 4-5 (see inset). The input power is the critical power P_{crit}, so there is always a 100%-efficiency stable solution, but as α/β_1 increases new stable and unstable solutions appear at lower efficiencies.

frequency component in this case is the fourth harmonic ($\sim$ 280 GHz), but different points in the phase diagram yield limit cycles with different balances of Fourier components.

4.5 Bifurcations and the non-linear wonderland

To better understand the phase diagram of Fig. 4-5, it is useful to plot the efficiencies of both the stable and unstable solutions as a function of various parameters. Several of these bifurcation diagrams (in which new fixed points typically appear in stable/unstable pairs) are shown in Figs. 4-7–4-9. To begin with, Figs. 4-7 and 4-8 correspond to lines connecting the labeled points ACF, BCD, and ECG, respectively, in Fig. 4-5, showing how the stability changes as a function of α/β_1 and τ_3/τ_1. Figure 4-7 shows how first one then two new stable fixed points

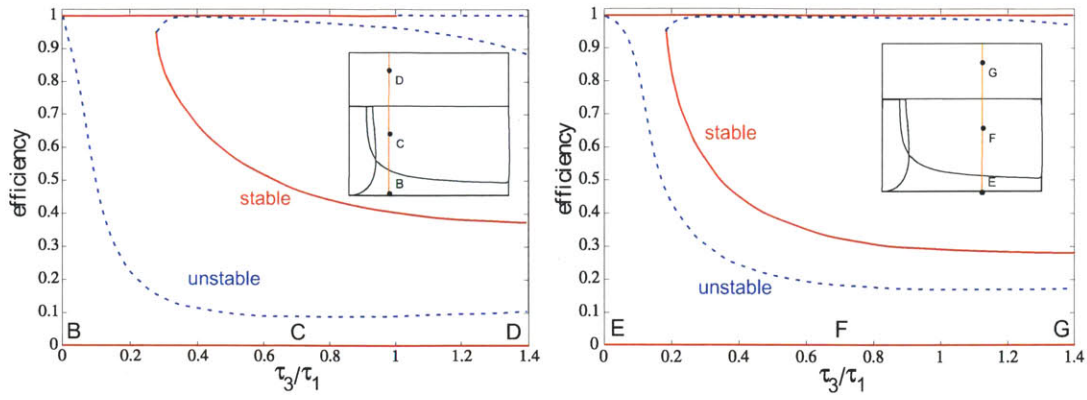

Figure 4-8: Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of τ_3/τ_1 for a fixed $\alpha/\beta_1 = 3$ (left) or $= 8$ (right), corresponding to the lines BCD or EFG, respectively, in Fig. 4-5 (see insets). The input power is the critical power P_{crit}, so there is always a 100%-efficiency steady-state solution, but it becomes unstable for $\tau_3 > \tau_1$ (a Hopf bifurcation leading to limit cycles as in Fig. 4-6).

appear as α/β_1 is increased, one approaching zero efficiency and the other closer to 50%. Along with these two stable solutions appear two unstable solutions (dashed lines). (A similar looking plot, albeit inverted, can be found in Ref. 53 for SPM-coupled closely-spaced resonances.) In particular, the fact that one of the unstable solutions approaches the 100%-efficiency stable solution causes the latter to have a smaller and smaller basin of attraction as α/β_1 increases, making it harder to excite as described in the next section. The next two plots, in Fig. 4-8, both show the solutions with respect to changes in τ_3/τ_1 at two different values of α/β_1. They demonstrate that at $\tau_1 = \tau_3$, a Hopf bifurcation occurs where the 100%-efficiency solution becomes unstable for $\tau_3 \geq \tau_1$ and limit cycles appear, intuitively seeming to "bounce between" the two nearby unstable fixed points. (The actual phase space is higher dimensional, however, so the limit cycles are not constrained to lie strictly between the efficiencies of the two unstable solutions.) It is worth to note that the remaining nonzero-efficiency stable solution (which appears at a nonzero τ_3/τ_1) becomes less efficient as τ_3/τ_1 increases.

The above analysis and results were for the steady-state-solutions when operating at the critical input power to obtain a 100%-efficiency solution. However, one can, of

course, operate with a different input power—although no other input power will yield a 100%-efficient steady-state solution, different input powers may still be useful because (as noted above and in the next section) the 100%-efficiency solution may be unstable or practically unattainable. Figure 4-9(left) is the bifurcation diagram with respect to the input power $P_{\text{in}}/P_{\text{crit}}$ at fixed α/β_1 and fixed τ_3/τ_1, corresponding to point C in Fig. 4-5. This power bifurcation diagram displays a number of interesting features, with the steady-state solutions transitioning several times from stable to unstable and vice versa. As we will see in the next section, the stability transitions in the uppermost branch are actually supercritical (reversible) Hopf bifurcations to/from limit cycles. Near the critical power, there is only a small region of stability of the near-100%-efficiency solution, as shown in the inset of Fig. 4-9(left). In contrast, the lower-efficiency stable solutions have much larger stable regions of the curve while still maintaining efficiencies greater than 70% at low powers comparable to $P_{\text{crit}} \sim V/Q^2$, which suggests that they may be attractive regimes for practical operation when α/β_1 is not small. This is further explored in the next section, and also by Fig. 4-9(right) which shows the bifurcation diagram along the line ACF in Fig. 4-5 [similar to Fig. 4-7], but at 135% of the critical input power. For this higher power, the system becomes at most doubly stable as α/β_1 is increased, and the higher-efficiency stable solution becomes surprisingly close to 100% as $\alpha/\beta_1 \to 0$.

4.6 Excitation of the high-efficiency solution

One remaining concern in any multistable system is how to excite the desired solution—depending on the initial conditions, the system may fall into different stable solutions, and simply turning on the source at the critical input power may result in an undesired low-efficiency solution. If α/β is small enough, of course, then from Fig. 4-5 the high-efficiency solution is the only stable solution and the system must inevitably end up in this state no matter how the critical power is turned on. Many interesting physical systems will correspond to larger values of α/β,

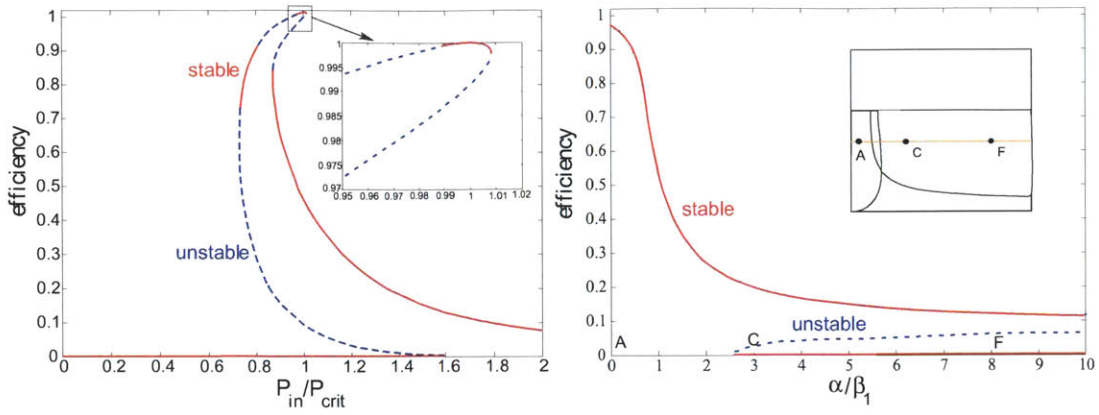

Figure 4-9: *Left:* Bifurcation diagram showing the harmonic-generation efficiency of the stable (solid red lines) and unstable (dashed blue lines) steady-state solutions as a function of input power $P_{\text{in}}/P_{\text{crit}}$ at fixed $\alpha/\beta_1 = 3$ and $\tau_3/\tau_1 = 0.7$, corresponding to point C in Fig. 4-5; the inset shows an enlarged view of the high-efficiency solutions. *Right:* Bifurcation diagram as a function of α/β_1 for fixed $P_{\text{in}}/P_{\text{crit}} = 1.35$ and fixed $\tau_3/\tau_1 = 0.7$; in this case, because it is not at the critical power, there are no 100%-efficiency solutions.

however [169], and in this case the excitation problem is complicated by the existence of other stable solutions. Moreover, the basins of attraction of each stable solution may be very complicated in the phase space, as illustrated by Fig. 4-10, where varying the initial cavity amplitudes $A_{1,3}$ from the 100%-efficiency solution causes the steady state to oscillate in a complicated way between the three stable solutions (at point C in Fig. 4-5). We have investigated several solutions to this excitation problem, and found an "adiabatic" excitation technique that reliably produces the high-efficiency solution without unreasonable sensitivity to the precise excitation conditions.

First, we considered a simple technique similar to the one described in Ref. 186 for exciting different solutions of a bistable filter: as shown in Fig. 4-11, we "turn on" the input power by superimposing a gradual exponential turn-on (asymptoting to $P_1 = P_{\text{crit}}$) with a Gaussian pulse of amplitude P_0 and width δT. The function of the initial pulse is to "kick" the system into the desired stable solution. We computed the eventual steady-state efficiency (after all transient effects have disappeared) as a function of the pulse amplitude P_0 at point C in Fig. 4-5, where

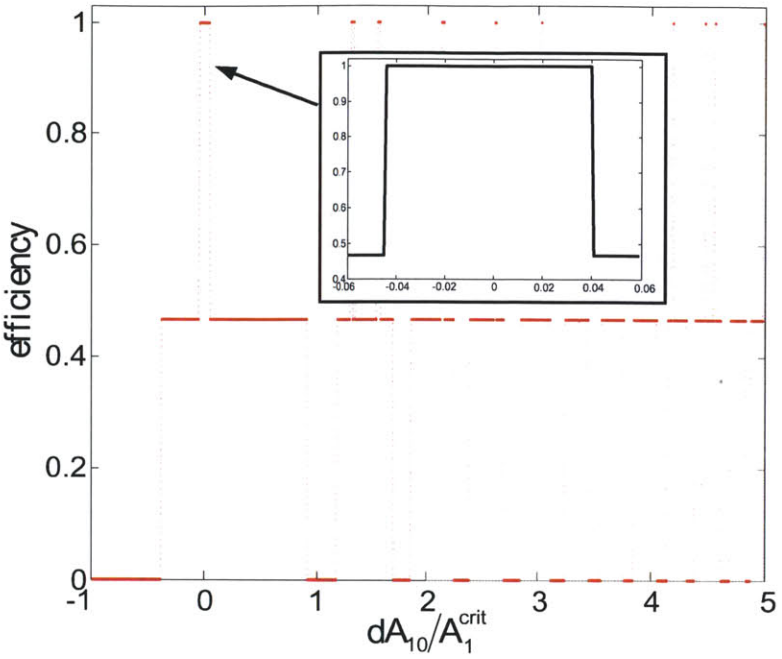

Figure 4-10: Asymptotic steady-state efficiency at point C (triply-stable) in the phase diagram (Fig. 4-5), with the initial conditions perturbed from the 100%-efficiency stable solution. The initial amplitudes A_{10} and A_{30} are perturbed by δA_{10} and δA_{30}, respectively, with $\delta A_{10}/A_1^{\text{crit}} = \delta A_{30}/A_3^{\text{crit}}$. The oscillation of the steady-state efficiency with the perturbation strength is an indication of the complexity of the phase space and the shapes of the basins of attraction of each fixed point.

there are three stable solutions. The results are shown in Fig. 4-12(left), and indeed we see that all three stable solutions from point C in Fig. 4-7: one at near-zero efficiency, one at around 47% efficiency, and one at 100% efficiency. Unfortunately, the 100% efficiency solution is obviously rather difficult to excite, since it occurs for only a very narrow range of P_0 values. One approach to dealing with this challenge is to relax the requirement of 100% efficiency (which will never be obtained in practice anyway due to losses), and operate at a power $P_1 < P_{\text{crit}}$. In particular, Fig. 4-9(left) shows that there is a much larger stable region for $P_1 \approx 0.8P_{\text{crit}}$ with efficiency around 90%, leading one to suspect that this solution may be easier to excite than the 100%-efficiency solution at $P_1 = P_{\text{crit}}$. This is indeed the case, as is shown in Fig. 4-12(right), plotting efficiency vs. P_0 at point C with $P_1 \approx 0.8P_{\text{crit}}$. In this case, there are only two stable solutions, consistent with Fig. 4-9(left), and there are much wider ranges of P_0 that attain the high-efficiency ($\approx 90\%$) solution.

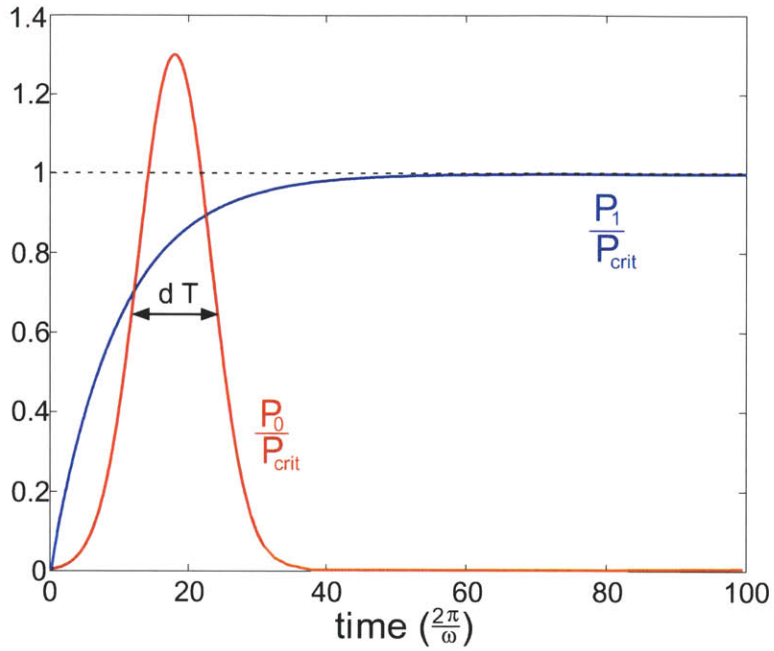

Figure 4-11: One way of exciting the system into a controlled stable solution: the input power is the sum of an exponential turn-on (the blue curve, P_1) and a Gaussian pulse with amplitude P_0 and width δT. The amplitude P_0 is altered to control which stable solution the system ends up in.

There are also many other ways to excite the high-efficiency solution (or whatever steady-state solution is desired). For example, because the cavity is initially detuned from the input frequency, as described in chapter 2, much of the initial pulse power is actually reflected during the transient period, and a more efficient solution would vary the pulse frequency in time to match the cavity frequency as it detunes. One can also, of course, vary the initial pulse width or shape, and by optimizing the pulse shape one may obtain a more robust solution .

In particular, one can devise a different (constant-frequency) input pulse shape that robustly excites the high-efficiency solution, insensitive to small changes in the initial conditions, by examining the power-bifurcation diagram in Fig. 4-9(left) in more detail. First, we observe that input powers $\gtrsim 1.45 P_{\text{crit}}$ have only one stable solution, meaning that this stable solution is excited regardless of the initial conditions or the manner in which the input power is turned on. Then, if we slowly decrease the input power, the solution must "adiabatically" follow this stable

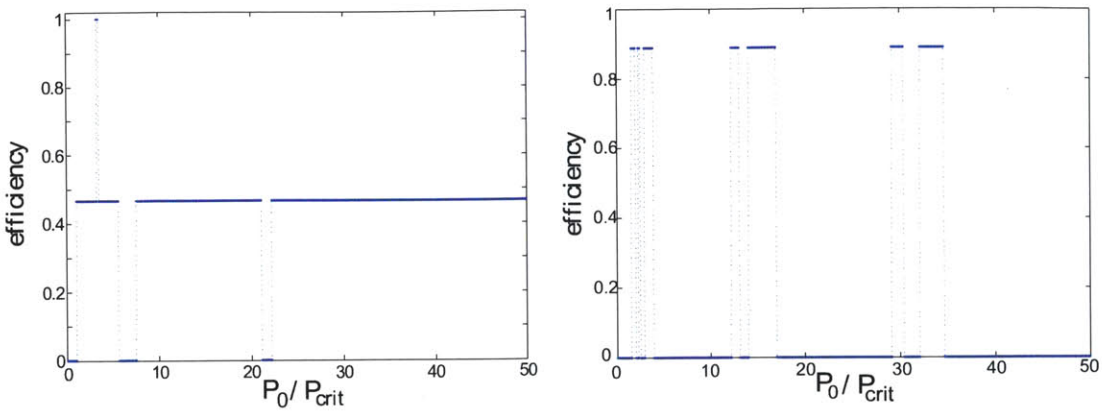

Figure 4-12: *Left:* Steady-state efficiency at point C in Fig. 4-5 as a function of the transient input-pulse power P_0 from Fig. 4-11, showing how all three stable solutions can be excited by an appropriate input-pulse amplitude. *Right:* Same, but for an asymptotic input power $P_1 \approx 0.8P_{\text{crit}}$, for which the maximum efficiency is $\approx 90\%$ from Fig. 4-9(right), but is easier to excite.

solution in the bifurcation diagram until a power $\approx 0.95P_{\text{crit}}$ is reached, at which point that stable solution disappears. In fact, by inspection of Fig. 4-9(left), at that point there are *no* stable solutions, and solution jumps into a limit cycle. If the power is further decreased, a high-efficiency stable solution reappears and the system must drop into this steady state (being the only stable solution at that point). This process of gradually decreasing the power is depicted in Fig. 4-13(left), where the instantaneous "efficiency" is plotted as a function of input power, as the input power is slowly decreased. (The efficiency can exceed unity,because we are plotting instantaneous output vs. input power, and in the limit-cycle self-pulsing solution the output power is concentrated into pulses whose peak can naturally exceed the average input or output power.) Already, this is an attractive way to excite a high-efficiency ($> 90\%$) solution, because it is insensitive to the precise manner in which we change the power as long as it is changed slowly enough—this rate is determined by the lifetime of the cavity, and since this lifetime is likely to be sub-nanosecond in practice, it is not difficult to change the power "slowly" on that timescale. However, we can do even better, once we attain this high-efficiency steady state, by then *increasing* the power adiabatically. As we increase the power, starting from the high-efficiency steady-state solution below the critical power, the

system first enters limit-cycle solutions when the power becomes large enough that the stable solution disappears in Fig. 4-9(left). As we increase the power further, however, we observe that these limit cycles *always* converge adiabatically into the 100%-efficiency solution when $P \to P_{\text{crit}}$. This process is shown in Fig. 4-13(right). What is happening is actually a supercritical Hopf bifurcation at the two points where the upper branch changes between stable and unstable: this is a reversible transition between a stable solution and a limit cycle (initially small oscillations, growing larger and larger away from the transition). This is precisely what we observe in Fig. 4-13, in which the limit cycle amplitudes become smaller and smaller as the stable solutions on either side of the upper branch are approached, leading to the observed reversible transitions between the two. The important fact is that, in this way, by first decreasing and then increasing the power to P_{crit}, one always obtains the 100%-efficiency solution regardless of the precise details of how the power is varied (as long as it is "slow" on the timescale of the cavity lifetime).

4.7 Losses

So far all the analysis we have done assume no losses. In practice, however, there will be different losses such as radiation, material absorption, two-photon absorption, and thermal effects. For example, never all the energy/fields in the cavity couple fully and directly to the waveguides; some of it will radiate into free space. In the case of a ring resonator, the smaller the ring-resonator, the higher radiation losses. In addition the closer the waveguide channels are to the ring, the higher radiation and refection losses are. As we discussed, that could be accounted for as a radiation lifetime $1/\tau_{e,k}$. That would add a small term $1/\tau_{e,k}$ to our coupled-mode equations. More specifically, in Equations (3.7–3.8), we no longer assume $\tau_k = \tau_{s,k}$. Instead $1/\tau_k = 1/\tau_{s,k} + 1/\tau_{e,k}$. This would change the dynamics of the system slightly. And as we have observed the system and specially the 100% solution can be very sensitive to small changes in the system parameters and therefore it is important to study how inclusion of different losses affects our

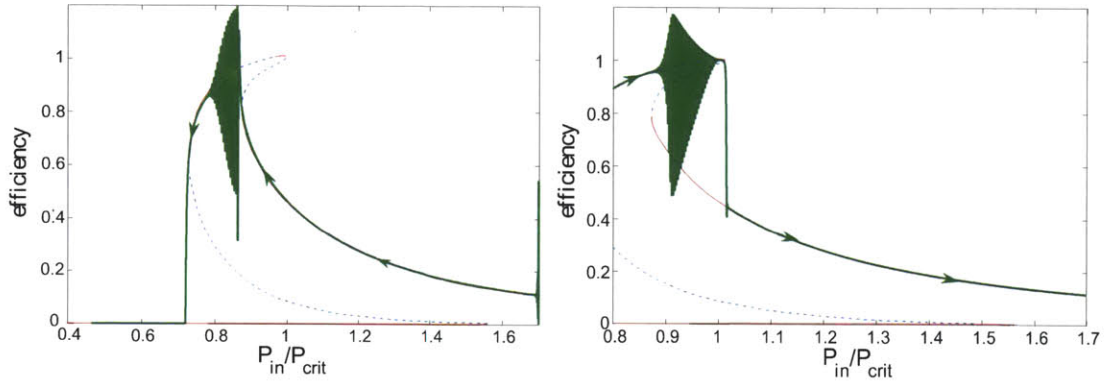

Figure 4-13: *Left:* Black line with arrows indicates instantaneous "efficiency" (harmonic output power / input power) as the input power is slowly decreased, starting at a power $\approx 1.7P_{\text{crit}}$. For comparison, Fig. 4-9(left) is superimposed as solid-red and dashed-blue lines. The solution "adiabatically" follows a steady state until the steady state becomes unstable, at which point it enters limit cycles, and then returns to a high-efficiency steady state, and finally goes drops to a low-efficiency steady-state if the power is further decreased. *Right:* Similar, but here the power is *increased* starting at the high-efficiency steady state solution for $P < P_{\text{crit}}$. In this case, it again enters limit cycles, but then it returns to a high-efficiency steady-state solution as the power is further increased, eventually reaching the 100%-efficiency stable solution. If the power is further increased, it drops discontinuously to the remaining lower-efficiency steady-state stable solution.

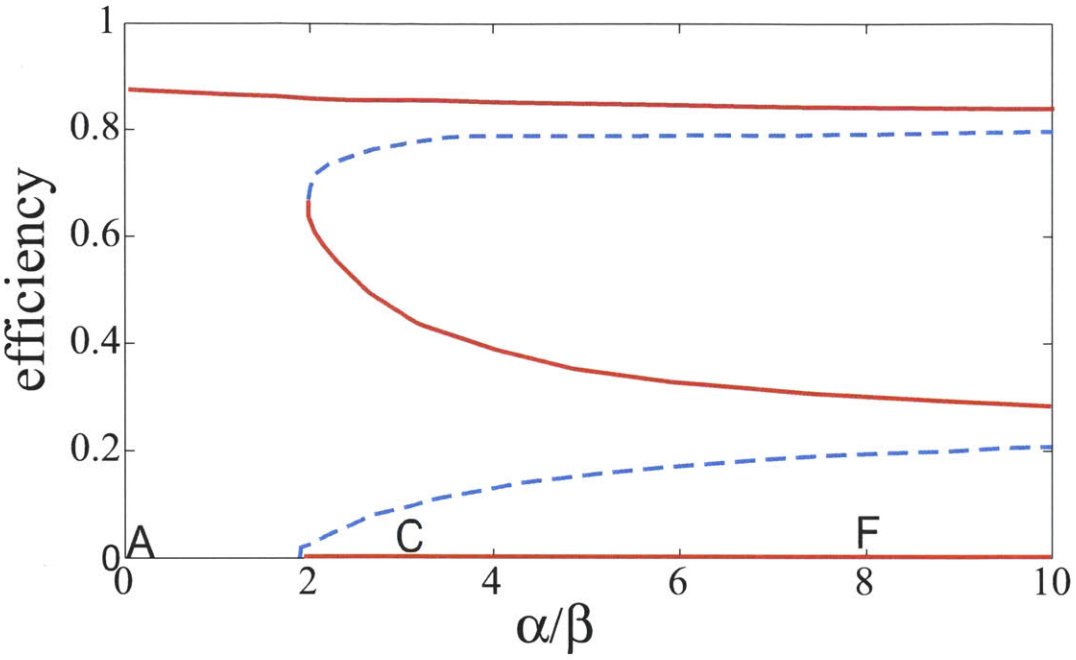

Figure 4-14: Effect of two-photon absorption on the conversion efficiency. This is calculation of the bifurcation diagram Fig. 4-7 with the difference that example two-photon absorption is included in the calculations. The qualitative behavior of the system is the same as before; only the efficiency of the high-efficiency solution decreases as α gets larger.

analysis. However, we reproduced the bifurcation diagrams after including linear losses and observed that the qualitative behavior of the system did not change; only the efficiency of the high-efficiency solution was shifted below 100%.

Similarly, we checked whether two-photon absorption makes any qualitative changes in the system dynamics. In the case of two-photon absorption, the losses are proportional to field energy in the cavity. This can be represented in the coupled-mode equations by adding an imaginary part to the α coefficients. For illustration, we added an imaginary part of magnitude $1/(20\tau_1)$ to the alpha coefficient and recalculated the bifurcation diagram of Fig. 4-7. The results are shown in Fig. 4-14: the solutions and their dynamics are very similar to the lossless case except that the efficiency of the high-efficiency solution decreases as α gets larger.

Therefore, as observed, although the losses decrease the overall conversion efficiency

(as expected), they do not make fundamental changes to the dynamics of the system and their stability,

Chapter 5

Analysis for Degenerate Four-Wave Mixing in Kerr Nonlinearities

Here, we look at a slightly different problem but with very similar dynamics to THG in chapter 4, namely degenerate four-wave mixing. More specifically, in a medium with Kerr ($\chi^{(3)}$) nonlinearity, we use two input signals at frequencies ω_0 and ω_m to generate a signal at frequency ω_p where the three frequencies are related in the following way:

$$\omega_m = \omega_0 - \Delta \tag{5.1}$$

$$\omega_p = \omega_0 + \Delta, \tag{5.2}$$

Unlike THG, this allows frequency conversion between *nearby* frequencies, for example to convert one channel into another in frequency-multiplexed optical communications, although we will see that the previous THG case arises as the special case $\Delta = 2\omega_0$. The analytical process is conceptually very similar to that of the THG case, however, and I co-supervised a UROP student, David Ramirez, in carrying it out and making the necessary adaptations. This work was published in Ref. 166.

In what follows, we extend the previous work on SHG, DFG and THG in resonant cavities to the case of DFWM in $\chi^{(3)}$ media. Previous work has studied FWM in

the context of optical fibers [68,78,94,111,179] and even matter waves [40], as well as demonstrating the use of FWM in applications such as phase conjugation [4,51,201,210] and generation of two-photon coherent states [111,181,211]. While there has been recent experimental work on intra-cavity FWM in $\chi^{(3)}$ media (degenerate or otherwise) [2,3,11,21,46,54,55,75,151,152,197,198], we are not aware of any detailed studies of the underlying theoretical phenomena in general cavities. As we shall see, DFWM in triply-resonant cavities shares many qualitative features with SHG, DFG, and THG, including the existence of critical powers at which optimal conversion efficiency is achieved as well as interesting nonlinear phenomena such as limit cycles and multistability. As in DFG, and unlike SHG or THG, there exist Manley-Rowe limitations on the overall conversion efficiency. In Sec. 5.2, we discuss the corresponding relations governing four-wave mixing and illustrate their implications for conversion efficiency. These relations can be obtained classically through temporal coupled-mode theory [73,74], but they are more easily motivated and understood from a quantum perspective [22,202].

We begin apply the coupled-mode formalism to the case of DFWM in a triply-resonant cavity in Sec. 5.1, to obtain the coupled-mode equations of motion as well as explicit expressions for the nonlinear coupling coefficients. We then briefly discuss general properties of the conversion process in Sec. 5.2 and, using the standard Manley-Rowe relations and simple photon-counting arguments, obtain limits on the maximal efficiency of the system. In Secs. 5.3.1 and 5.3.2, we analyze the stability and dynamics of the solutions to the coupled-mode equations obtained in chapter 3, neglecting SPM and XPM effects, and demonstrate the existence of the maximal conversion efficiencies obtained in Sec. 5.2. Finally, in Sec. 5.3.3, we briefly consider the effects of SPM and XPM using a simple model to illustrate the qualitative behavior of the system; in particular, we demonstrate the existence of stable, maximal efficiency solutions even including SPM and XPM effects.

5.1 Coupled-mode equations

Similar to before, we use coupled-mode theory to describe the system in terms of ODE's:

$$\frac{da_0}{dt} = \left[i\omega_0(1 - \alpha_{00}|a_0|^2 - \alpha_{0m}|a_m|^2 - \alpha_{0p}|a_p|^2) - \frac{1}{\tau_0}\right] a_0 - i\omega_0\beta_0 a_0^* a_m a_p + \sqrt{\frac{2}{\tau_{s,0}}} s_{0,+} \tag{5.3}$$

$$\frac{da_m}{dt} = \left[i\omega_m(1 - \alpha_{m0}|a_0|^2 - \alpha_{mm}|a_m|^2 - \alpha_{mp}|a_p|^2) - \frac{1}{\tau_m}\right] a_m - i\omega_m\beta_m a_0^2 a_p^* + \sqrt{\frac{2}{\tau_{s,m}}} s_{m,+} \tag{5.4}$$

$$\frac{da_p}{dt} = \left[i\omega_p(1 - \alpha_{p0}|a_0|^2 - \alpha_{pm}|a_m|^2 - \alpha_{pp}|a_p|^2) - \frac{1}{\tau_p}\right] a_p - i\omega_p\beta_p a_0^2 a_m^* \tag{5.5}$$

$$s_{k,-} = \sqrt{\frac{2}{\tau_{s,k}}} a_k - s_{k,+}. \tag{5.6}$$

where as before, τ_p, τ_m and τ_0 are the corresponding cavity lifetimes of the three modes and the α_{ij}'s and β's are the coupling coefficients with α's describing the self- and cross-phase modulation effects and β terms, the energy transfer between different modes. Using the same method as for 3rd harmonic generation, we find the coupling coefficients to be:

$$
\begin{aligned}
\beta_0 &= \frac{1}{8}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[(\mathbf{E}_0^*\cdot\mathbf{E}_0^*)(\mathbf{E}_m\cdot\mathbf{E}_p)+2(\mathbf{E}_0^*\cdot\mathbf{E}_m)(\mathbf{E}_0^*\cdot\mathbf{E}_p)\right]}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_0|^2\right]\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_m|^2\right]^{1/2}\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_p|^2\right]^{1/2}} \\
\beta_m &= \beta_p = \frac{1}{2}\beta_0^* \\
\alpha_{jj} &= \frac{1}{8}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[|\mathbf{E}_j\cdot\mathbf{E}_j^*|^2+|\mathbf{E}_j\cdot\mathbf{E}_j|^2\right]}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_j|^2\right]^2} \\
\alpha_{jk} &= \frac{1}{8}\frac{\int d^3\mathbf{x}\epsilon_0\chi^{(3)}\left[|\mathbf{E}_j\cdot\mathbf{E}_k|^2+|\mathbf{E}_j\cdot\mathbf{E}_k^*|^2+|\mathbf{E}_j|^2|\mathbf{E}_k|^2\right]}{\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_j|^2\right]\left[\int d^3\mathbf{x}\epsilon|\mathbf{E}_k|^2\right]} \\
\alpha_{kj} &= \alpha_{jk}.
\end{aligned}
\tag{5.7}
$$

It is worth noting few details about these equations and the system we are looking at: Since the system operates with Kerr nonlinearities, we still observe self- and cross-phase modulation effects and therefore frequency shifting. As we will see later however, dealing with the frequency shifting in this case is easier than in the case of 3rd harmonic generation. Moreover, in this system, we have two input signals, one at frequency ω_0 which we will refer to as the input *pump* and one at ω_m that we refer to as the input *signal* to generate the *shifted signal* at ω_p.

Conservation laws give us relations between frequencies and coupling coefficients. Namely, conservation of frequency gives us $2\omega_0 = \omega_m + \omega_p$. In addition, we can again deduce relations between β's using conservation of energy. Requiring $\frac{d}{dt}(l|a_0|^2 + |a_p|^2 + |a_m|^2) = 0$ in an isolated cavity with no coupling to waveguides gives us:

$$
\omega_0\beta_0^* = \omega_m\beta_m + \omega_p\beta_p. \tag{5.8}
$$

This system operates in two very different regimes: 1) $\Delta\omega < \omega_0$, and 2) $\Delta\omega > \omega_0$.

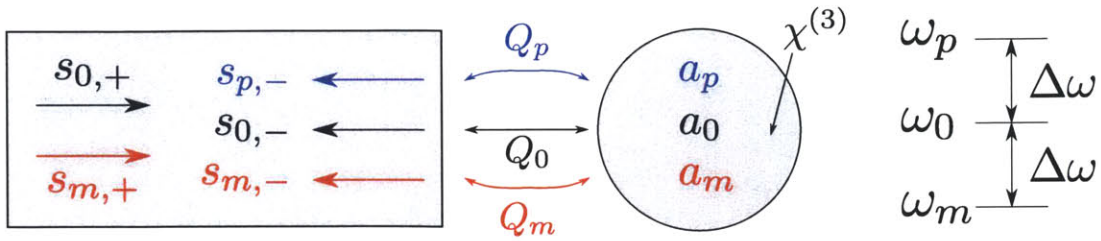

Figure 5-1: (*Left*) Schematic for *degenerate* four-wave mixing involving a coupled waveguide-cavity system. Dynamical variables for coupled-mode equations represent: a single input (output) channel (with incoming/outgoing field amplitudes $s_{\pm}$) coupled to a resonant cavity with three modes at frequencies ω_0, $\omega_m = \omega_0 - \Delta\omega$ and $\omega_p = \omega_0 + \Delta\omega$ (and corresponding amplitudes a_0, a_m and a_p). The three resonant modes are nonlinearly coupled by a Kerr ($\chi^{(3)}$) nonlinearity. (*Right*) Diagram illustrating the relationship between the three resonant frequencies.

The two different regimes give rise to different dynamics. Most importantly, they allow different maximum conversion efficiencies.

5.2 Maximum efficiency: Quantum-limited vs. complete

As described below, the DFWM process we consider here exhibits drastically different behavior depending on the ratio of $\Delta\omega$ to ω_0. In particular, there exist at least two distinct regimes of operation, corresponding to quantum-limited ($|\Delta\omega| < \omega_0$) and complete ($\Delta\omega \geq \omega_0$) conversion. It turns out that, although our coupled-mode formalism is entirely classical, the *same* behaviors can be more easily understood by considering photon interactions in a quantum picture. Although this system is, of course, described by the general Manley-Rowe relations, which can be derived from both classical [18, 73, 74] and quantum [22, 202] arguments similar to those here, it is useful to review a basic picture of such limits and their physical consequences for the specific case of intra-cavity DFWM.

Our focus here is the up-conversion process (or interaction) corresponding to taking input light at frequencies ω_0 and ω_m and generating output light at frequency ω_p. Therefore, an appropriate figure of merit is the ratio of the output power in the ω_p

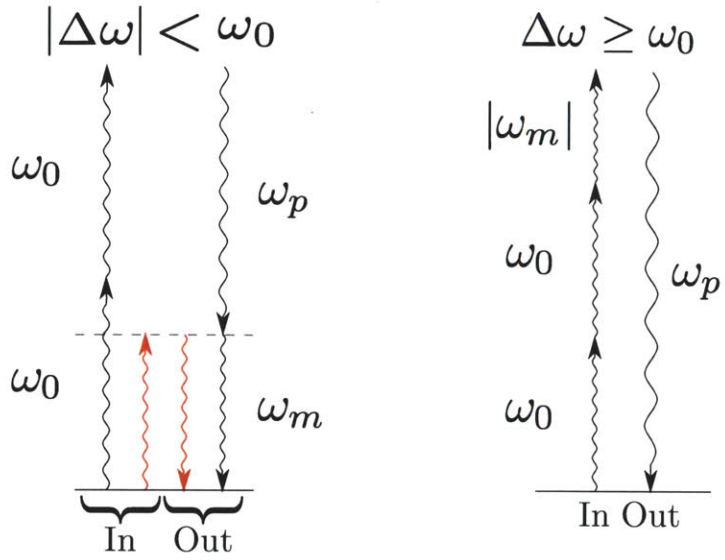

Figure 5-2: Diagram of nonlinear up-conversion process involving input light at ω_0 and ω_m and output light at ω_p and ω_m. The conversion efficiency of DFWM is determined by $\Delta\omega$, and photon energy conservation consideration (see text), leading to at least two different regimes of operation: (*Left:*) for $|\Delta\omega| < \omega_0$, two ω_0 pump photons and an signal ω_m photon are converted into two ω_m signal photons and an ω_p photon. The input ω_m photon is only necessary to initiate the conversion process and emerges unchanged after the interaction (indicted by red). (*Right:*) for $\Delta\omega \geq \omega_0$, two incoming ω_0 and a single ω_m photon are combined to produce an ω_p photon. In contrast to the previous regime, the ω_m photon is energetically needed to produce the ω_p photon.

mode to the total input power, which we define as the absolute efficiency

$$\eta = |s_{p,-}|^2/(|s_{0,+}|^2 + |s_{m,+}|^2).$$

As described in the chapter 3, the coupled-mode equations [Eqs. (5.3–5.6)] follow from very general and purely classical considerations. The same considerations yield relationships between the frequencies and coupling coefficients of the problem, such as frequency conservation ($\omega_m + \omega_p = 2\omega_0$) and energy conservation ($\omega_m\beta_m + \omega_p\beta_p = \omega_0\beta_0^*$). Additional conservation rules that are perhaps best understood from quantum arguments, such as photon energy ($\hbar|\omega|$) conservation and standard $\chi^{(3)}$ selection rules [18], also play a substantial role in the physics of nonlinear frequency conversion. In the case of the DFWM up-conversion process considered here, $\chi^{(3)}$ selection rules imply that nonlinear interactions can only be initiated if there exist at least three input photons: $2\omega_0$ photons and one ω_m photon.

In the $|\Delta\omega| < \omega_0$ regime, there are at least two important features that can be understood from the above relations: First, depletion of the signal input power ($s_{m,+}$) is impossible, leading to a conversion efficiency $\eta < 1$. Second, in order to maximize the total conversion efficiency, one desires $s_{m,+}$ to be as small as possible. These features can be understood by considering a simple picture of the nonlinear photon–photon interaction, as follows. From the DFWM $\chi^{(3)}$ selection rule [18], it follows that the creation of an ω_p photon is accompanied by the destruction of two ω_0 photons and one ω_m photon. The latter, along with photon energy conservation, leads to the process considered in Fig. 5-2 (left), in which two ω_0 photons and an ω_m photon interact to yield two ω_m photons and an ω_p photon. From the figure, and since $2\omega_0 > \omega_p$, one can see that the incident ω_m photon (depicted in red) is merely required by the $\chi^{(3)}$ selection rule to initiate the interaction, and emerges unmodified, accompanied by an ω_p photon and an additional ω_m photon. Thus, it is clear that the input ω_m photon does not actively participate in the energy transfer and therefore merely reduces the maximum possible conversion efficiency. This implies that one desires a minimal input signal power to initiate the up-conversion. Effectively, the incident ω_m photons are amplified by the conversion process (a similar amplification effect is a crucial component in other nonlinear interactions, such as OPAs in $\chi^{(2)}$ media [8, 18, 34]). In addition, it is clear that complete depletion of the signal photons, i.e. $s_{m,-} = 0$, is not possible for non-zero $s_{m,+}$, and therefore the conversion efficiency must be less than 100% (since the total input power is conserved). No such restriction is placed on $s_{0,-}$, and therefore we expect that maximal efficiency will be obtained for arbitrarily low signal power and complete depletion of the pump power, i.e. $s_{0,-} = 0$.

Based on these arguments, we can predict the maximal efficiency of the conversion process by considering the ratio of the energy of the output ω_p photon ($\hbar\omega_p$) to the energy of the three input photons [$\hbar(2\omega_0 + \omega_m)$]. Since the ω_m photons can be provided with arbitrarily low amplitude, we therefore expect maximal efficiency to be achieved upon neglecting their contribution, i.e. we predict a maximal efficiency

of:

$$\eta_{\max}(|\Delta\omega < \omega_0|) = \frac{\hbar\omega_p}{2\hbar\omega_0} = \frac{\omega_p}{2\omega_0}. \tag{5.9}$$

Note that this efficiency depends only on the ratio of $\Delta\omega$ to ω_0 and $\hbar$ cancels, so it should appear in the classical limit as well. As we shall see in Sec. 5.3.1, this prediction is verified analytically by examining the steady-state solution of our coupled-mode equations.

In the $\Delta\omega \geq \omega_0$ regime, the conversion process is fundamentally different and, in particular, complete depletion of the ω_m and ω_0 photons is possible, leading to 100% conversion efficiency. Basically, because $\omega_p > 2\omega_0$ in this case, no additional photons are required to satisfy photon energy conservation, yielding the nonlinear interaction process depicted in Fig. 5-2 (right), where two input ω_0 photons and an ω_m photon combine to produce an ω_p photon. Note that now the input ω_m photon actively participates in the energy transfer, in contrast to the $|\Delta\omega| < \omega_0$ regime, leading to a maximal conversion efficiency occurring when $s_{0,+}$ and $s_{m,+}$ are both non-zero. Furthermore, since ω_p is now the only product of the interaction, we expect that complete depletion of both the pump and signal powers, $s_{0,-} = s_{m,-} = 0$, should be possible, leading to 100% conversion efficiency. As before, this can also be quantified by comparing the ratio of the output energy ($\hbar\omega_p$) to the input energy $[\hbar(2\omega_0 + |\omega_m|)]$ (note that now the energy of ω_m photon is $\hbar|\omega_m|$), and the result follows from the fact that $2\omega_0 + |\omega_m| = \omega_p$. Again, we shall see in Sec. 5.3.2, this prediction is validated analytically and directly from the coupled-mode equations, yielding also the critical input powers at which 100% conversion is achieved.

In this section, we made a number of predictions based on very general arguments relying on a quantum interpretation of the nonlinear interactions, allowing us to obtain predictions of maximal conversion efficiency. Our final results, of course, contained no factors of $\hbar$ and it is therefore not surprising that we recover the same results (albeit with more detail, e.g. predictions of the values of critical powers) in the ensuing analysis of the purely-classical coupled-mode equations. Nevertheless,

the heuristic quantum picture of Fig. 5-2 has the virtue of being simple and revealing, while the classical derivation is more complicated (although more quantitative). Similar quantum arguments have also proven useful in other contexts, such as in many problems involving classical radiation [79], or the recently-studied problem of optical bonding/anti-bonding in waveguide structures [159].

5.3 Coupled-mode analysis

In order to gain a simple understanding of the system, we shall first consider frequency conversion in the absence of self- and cross-phase modulation, i.e. $\alpha_{jk} = 0$. The nonzero-α case will be considered in Sec. 5.3.3. Section 5.3.1 focuses on the $|\Delta\omega| < \omega_0$ regime, whereas Sec. 5.3.2 focuses on the $\Delta\omega \geq \omega_0$ regime. In both cases, we describe the solutions to the coupled-mode equations [Eqs. (5.3–5.5))] in the steady state, including the stability of these solutions and their dependence on the cavity parameters.

5.3.1 $|\Delta\omega| < \omega_0$ regime: Limited conversion

Although the analysis in this section is general, for the purposes of plotting results we choose the specific parameters: $\alpha_{jk} = 0$, $\tau_0 = \tau_m = \tau_p = 100/\omega_0$, $\beta = 10^{-4}$, and $\Delta\omega = 0.05\omega_0$. The qualitative results remain unchanged as these parameters are varied, provided that the Q are large enough such that mode overlap is minimal as required by CMT. The influence of varying these parameters is discussed further at the end of the section.

To understand the stability and dynamics of the nonlinear coupled-mode equations in the quantum-limited regime, we apply the standard technique of identifying the fixed points of Eqs. (5.3–5.5) and analyzing the stability of the linearized equations around each fixed point [191]. A fixed point is given by a steady-state solution where the mode amplitudes vary as $a_k(t) = A_k e^{i\omega_k t}$, with the A_k being unknown constants. Plugging this steady-state ansatz into Eqs. (5.3–5.5), we obtain three coupled polynomial equations in the parameters $A_0, A_m, A_p, s_{0,+}$, and $s_{m,+}$. These

polynomials were solved using Mathematica to obtain the mode energies $|A_k|^2$, which are then used to calculate the efficiency $\eta = |s_{p,-}|^2/(|s_{0,+}|^2 + |s_{m,+}|^2)$. The phases of the A_k can be easily determined from the steady-state equations of motion; A_0 and A_m acquire the phases of $s_{0,+}$ and $s_{m,+}$ respectively, while the phase of A_p is that of $\beta_p A_0^2 A_m^*$ rotated by $\pi/2$. Without loss of generality, $s_{0,+}$ and $s_{m,+}$ can be chosen to be real.

In general, this system has either one or three solutions, only one of which is ever stable. The stability and efficiency of this solution are shown in Fig. 5-3 for the specific parameters mentioned above. We observe that maximal conversion efficiency is obtained in the limit as input signal power $s_{m,+}$ is reduced to zero, consistent with the discussion in the previous section. To obtain the maximum efficiency and the corresponding critical input powers, complete depletion of the pump (ω_0) photon is required, i.e. $s_{0,-} = 0$ (note that one cannot require depletion of the signal photon, for the reasons discussed in the previous section). We find that the maximum efficiency $\eta_{\max}$ is obtained at $|s_{0,+}^{\text{crit}}|^2 = P_0$ as $|s_{m,+}|^2 \to 0$, where:

$$\begin{aligned} P_0 &= \frac{4}{\tau_0 |\beta_0| \sqrt{\tau_m \tau_p |\omega_m \omega_p|}} \\ \eta_{\max} &= \frac{\omega_p}{2\omega_0} = \frac{1}{2}\left(1 + \frac{\Delta\omega}{\omega_0}\right), \end{aligned} \tag{5.10}$$

Note that Eq. (5.10) is identical to the value predicted in the previous section. In the important case of narrow-band conversion, $|\Delta\omega| \ll \omega_0$, the maximum efficiency is approximately 50%. (however, this is relative to the pump power— compared to the input signal alone, the output signal is amplified to an arbitrary degree). If $Q_0 \sim Q_m \sim Q_p$, then, as in THG, the critical power scales as V/Q^2, where V is the modal volume (recall that $\beta \sim 1/V$).

As $\Delta\omega \to \omega_0$, the maximum efficiency approaches unity, i.e. 100% conversion can be achieved in the limit. This limit is reminiscent of second-harmonic generation, since $\omega_p = 2\omega_0$. However, the interaction process is fundamentally different from the standard ($\chi^{(2)}$) SHG in a number of ways. First, one is converting DC ($\omega_m \approx 0$) light and ω_0 pump light into $2\omega_0$. Second, the stability of this solution (described

below) is quite different from that of SHG [45, 64, 171]. Finally, the critical power in this case, P_0, diverges as $1/\sqrt{1-(\Delta\omega/\omega_0)^2}$ for $\Delta\omega$ near ω_0. However, $\Delta\omega$ close but not equal to ω_0 yields a reasonable P_0: for example, $\Delta\omega = 0.95\omega_0$ yields efficiency $\eta = 0.975$ with a critical power roughly three times the critical power for $\Delta\omega$ near zero. Because this near-"SHG" situation involves coupling resonances at very different frequency scales, it is reminiscent of using $\chi^{(2)}$ DFG to produce THz from infrared [23].

Eq. (5.10) are only valid in the limit $|s_{m,+}|^2 \to 0$, which is ideal from an efficiency perspective. However, it is interesting to consider the system for non-infinitesimal $s_{m,+}$, in which case we solve for the input power that yields a stable solution with maximal efficiency for a given $s_{m,+}$. We denote this input power by $P_c(|s_{m,+}|^2) = |s_{0,+}^{\text{crit}}|^2 + |s_{m,+}|^2$, where $|s_{0,+}^{\text{crit}}|^2$ (a function of $|s_{m,+}|^2$) is defined to be the pump power required to achieve maximum, stable conversion efficiency for a given signal power $|s_{m,+}|^2$. As seen in Fig. 5-3, this efficiency is always $\leq \eta_{\max}$, and $P_c \to P_0$ as $s_{m,+} \to 0$. In the non-zero $|s_{m,+}|^2$ regime, P_c does not correspond to complete depletion of the pump. Requiring pump depletion ($s_{0,-} = 0$) for a given signal power $|s_{m,+}|^2$ yields two pump powers, which we label $P_\pm(|s_{m,+}|^2)$. $P_+(|s_{m,+}|^2)$ does indeed provide a solution with maximal efficiency, however this solution is always unstable. As seen from Fig. 5-3, only for small signal power $s_{m,+}$ does depletion of the pump lead to maximal efficiency.

In general, to obtain the largest efficiency while retaining stability, one would aim to operate with low signal power $|s_{m,+}|^2$ and use a pump power near the critical power P_0 given in Eq. (5.10). However, it is interesting to consider the unstable solutions, because they turn out to be related to limit cycles. As mentioned above, the system contains either one or three steady-state solutions for given input powers. Fig. 5-4 plots these stable and unstable solutions as a function of pump power $|s_{0,+}|^2$ at fixed signal power $|s_{m,+}|^2 = 0.1P_0$, corresponding to the horizontal dashed line in Fig. 5-3. For low input pump power $|s_{0,+}|^2$, the system has a single steady-state solution; as the pump power is increased, the system experiences a bifurcation yielding two unstable solutions. As mentioned above, the higher efficiency solution emerging

from the bifurcation achieves a maximum corresponding at $|s_{0,+}|^2 = P_+$, coinciding with complete depletion of the pump ($s_{0,-} = 0$), but this maximal efficiency solution is always unstable; note that there may be a stable solution at $|s_{0,+}|^2 = P_+$, but the stable solution will have a lower efficiency than the maximal, unstable solution, as shown in Fig. 5-4. Furthermore, the original stable solution eventually becomes unstable as the pump power is increased (this can occur before or after the bifurcation, depending on the system parameters); this onset of instability coincides with the onset of limit cycles, stable oscillating-efficiency solutions. An example of these limit cycles are shown in Fig. 5-4, where the green dashed lines indicate the bounds of the oscillations and the solid green line gives the average. The limit cycles are plotted as a function of time in the inset of Fig. 5-4. The limit cycles shown here were obtained by numerically time-evolving the coupled-mode equations. In general, we find that these limit cycles oscillate with a period proportional to τ_p.

Figures 5-3–5-4 describe a system corresponding to a particular set of values for the parameters $\Delta\omega$ and τ_k. Qualitatively, the most important features of the figures remain largely unchanged as these parameters are varied. Basically, there exist at most three solutions to the coupled-mode equations, one of which has a finite region of stability as a function of $s_{0,+}$ and $s_{m,+}$, with the general shape that is shown in Fig. 5-3, and two others that are always unstable and bifurcate at a finite $s_{0,+}$. There are however, some differences to note: First, as $\Delta\omega$ increases from 0, the maximum steady-state efficiency also increases, asymptoting to $\eta = 1$ as $\Delta\omega \to \omega_0$. This was obtained analytically and is quantified in Eq. (5.10). Unfortunately, we find that as $\Delta\omega$ increases, the region of instability in Fig. 5-3 also increases, and furthermore, the conversion efficiency at finite $s_{m,+}$ also drops off more rapidly. (In particular, we observe in the "SHG" limit of $\Delta\omega \to \omega_0$, the system becomes largely unstable except for very low signal powers.) These tendencies are depicted in Fig. 5-5, which plots $P_c(|s_{m,+}|^2)$ and the corresponding conversion efficiency for different values of $\Delta\omega$. The kinks observed in the plots of P_c are due to the discontinuity in the slope of the P_c curve as it reaches the region of instability, corresponding to the point U in Fig. 5-3.

Varying τ_k does not affect the maximum possible efficiency and also leaves Fig. 5-3 qualitatively unchanged, changing only the scale of the critical input power P_0. The stability of the system however, does depend on the relative lifetimes of the cavity modes. In particular, the stability depends largely on the ratio τ_0/τ_p, and decreases weakly as τ_m increases with respect to either τ_0 or τ_p. This makes sense since, as argued in Sec. 5.2, the ω_m photons do not actively participate in the energy transfer. (A similar dependence on the ratio of the lifetimes was also observed in the case of THG.) More quantitatively, we follow the position of the point U (the point where P_c reaches the region of instability) as the τ_k are varied. Assuming equal modal lifetimes ($\tau_0 = \tau_m = \tau_p$ as in Fig. 5-3), we find that U lies at critical input powers $|s_{0,+}|^2 \approx 1.28P_0$ and $|s_{m,+}|^2 \approx 0.35P_0$. Increasing τ_0/τ_p, from 1 to 10, we find that U moves to $|s_{0,+}|^2 \approx 10P_0$ and $|s_{m,+}|^2 \approx 4.75P_0$. However, if we instead keep $\tau_0 = \tau_p$ and increase τ_m such that $\tau_m/\tau_0 = \tau_m/\tau_p = 10$, U moves only to $|s_{0,+}|^2 \approx 1.05P_0$ and $|s_{m,+}|^2 \approx 0.27P_0$. Note that, as mentioned previously, maximal stable conversion efficiency is obtained for low signal power $|s_{m,+}|^2$ and input power $|s_{0,+}|^2$ near the critical power P_0, regardless of τ_k. We note that rescaling β simply scales the input power and therefore changing β does not affect the dynamics.

Thus far, we have focused on the up-conversion process: taking input light at frequencies ω_0 and ω_m and generating output light at frequency $\omega_p > \omega_0$. However, it suffices to consider the above system when $\Delta\omega < 0$ to understand the physics of the alternative, down-conversion process: taking input light at frequencies ω_0 and ω_p and generating output light at frequency ω_m. For $\Delta\omega < 0$, we effectively have $\omega_m \leftrightarrow \omega_p$. In this regime, all of the above analysis holds, and in particular, the maximal efficiency, given by Eq. (5.10), is obtained as $|s_{m,+}|^2 \to 0$ with $|s_{0,+}|^2 = P_0$. Similarly, the stability of the solutions follow similar trends to those outlined above.

5.3.2 $\Delta\omega \geq \omega_0$ regime: Complete conversion

When $\Delta\omega$ is larger than ω_0, we argued in Sec. 5.2 that the system is capable of complete conversion, i.e. $\eta = 1$. In this section, we demonstrate the existence of a critical steady-state solution to the classical coupled mode equations with complete

conversion and analyze the stability of this critical solution, as well as relate DFWM to THG we studied in detain in chapter 4.

As in the previous section, we consider the equations of motion Eqs. (5.3–5.5) in the steady state. To obtain the critical solution, we again require depletion of the pump power, i.e. $s_{0,-} = 0$. However, as argued in Sec. 5.2, complete depletion of the signal, $s_{m,-} = 0$ must also occur. Recall from Sec. 5.2 that complete ω_m depletion is possible in the $\Delta\omega \geq \omega_0$ regime since the up-conversion process does not produce ω_m photons (see Fig. 5-2). Imposing the depletion constraints on the steady-state equations of motion yields the following critical cavity energies $|a_k^{\text{crit}}|^2$:

$$|a_0^{\text{crit}}|^2 = \frac{1}{|\beta_m|\sqrt{\tau_m \tau_p |\omega_m \omega_p|}}, \tag{5.11}$$

$$|a_m^{\text{crit}}|^2 = \frac{\tau_m |\omega_m|}{2\tau_0 \omega_0} |a_0^{\text{crit}}|^2, \tag{5.12}$$

$$|a_p^{\text{crit}}|^2 = \frac{\tau_p \omega_p}{2\tau_0 \omega_0} |a_0^{\text{crit}}|^2, \tag{5.13}$$

which lead to the following critical powers:

$$|s_{0,+}^{\text{crit}}|^2 = P_0 \tag{5.14}$$

$$|s_{m,+}^{\text{crit}}|^2 = \frac{|\omega_m|}{2\omega_0} P_0, \tag{5.15}$$

where P_0 is given by Eq. (5.10). Solving for the corresponding output signal $|s_{p,-}|^2$, the output power is indeed 100% of the input power, as required by energy conservation. (In contrast, the assumption that $s_{0,-} = s_{m,-} = 0$ in the $|\Delta\omega| < \omega_0$ case yields no solution). Note that the critical signal power $|s_{m,+}^{\text{crit}}|^2$ is now non-zero, due to the fact that the energy from the signal ω_m photons is necessary to produce the output ω_p photons. This is in contrast with the $|\Delta\omega| < \omega_0$ regime where maximal conversion efficiency was only achieved in the limit as input signal power $|s_{m,+}|^2$ decreased to zero. The critical pump and signal powers, with the corresponding maximum efficiency η, are plotted versus $\Delta\omega$ in Fig. 5-6 for both $\Delta\omega$ regimes.

As may be noted from Fig. 5-6, there are two particular values of $\Delta\omega$ that warrant special attention when $\Delta\omega \geq \omega_0$. The first case, when $\Delta\omega = \omega_0$, the "SHG" case, was discussed in the previous section. The second case is when $\Delta\omega = 2\omega_0$. In this case, $\omega_m = -\omega_0$ and $\omega_p = 3\omega_0$, reminiscent of third-harmonic generation (THG). In fact, this case of DFWM corresponds exactly to $\chi^{(3)}$ THG, and thus $\Delta\omega > \omega_0$ strictly generalizes our previous THG analysis in chapter 4. To see this, some care must be taken to adjust the coupling coefficients β_k given in Eq. (5.10) to properly implement the rotating wave approximation; since $\omega_m = -\omega_0$, we have that $a_m = a_0^*$, and thus $\beta_0 \to \beta_0 + \beta_m^*$ and $\beta_m \to \beta_m + \beta_0^*$. This results in $\beta_0 = \beta_m^* = 3\beta_p^*$, exactly as shown in chapter 3. Furthermore, we have $|s_{0,+}^{\text{crit}}|^2 = |s_{m,+}^{\text{crit}}|^2 = P_0$ (note that this differs by a factor of two from Eq. (5.15), due to the adjusted β_k values); upon requiring that $\tau_0 = \tau_m$, this recovers the critical power previously obtained for THG. Note that the correspondence between $\Delta\omega = 2\omega_0$ and $\chi^{(3)}$ THG is exact, whereas the $\Delta\omega = \omega_0$ limit has little in common with $\chi^{(2)}$ SHG as discussed above. The existence of an $s_{0,-} = s_{m,-} = 0$ solution having demonstrated the existence of critical powers where 100% conversion can be achieved, we are now interested in characterizing the system at this critical power by studying all of the fixed points. These fixed points were obtained using Mathematica as in the previous section, and their stability was determined via linear stability analysis as before. For the critical input power, the steady-state equations of motion yield three solutions; however, in contrast to the $|\Delta\omega| < \omega_0$ regime, there exists multistability when $\Delta\omega \geq \omega_0$. Similar to the case of THG ($\Delta\omega = \omega_0$), the system is either singly stable, doubly stable, or unstable, depending on the values of the mode lifetimes τ_k (see Fig. 5-7). In this $\Delta\omega > \omega_0$ regime, the stability of the solutions does not depend on $\Delta\omega$, again in contrast with the quantum-limited regime. Unlike the $|\Delta\omega| < \omega_0$ regime, the value of τ_m now plays a significant role in the stability of the solutions.

5.3.3 Self- and Cross-Phase Modulation ($\alpha \neq 0$)

Finally, we briefly consider the effects of SPM and XPM. This corresponds to taking the coefficients α_{jk} to be non-zero; as mentioned above, for simplicity, we take all

the coefficients to be equal, i.e. $\alpha_{jk} = \alpha$ for all j, k. As we showed in Sec. 4.3, one simple way to overcome this difficulty is to pre-shift the cavity resonant frequencies so as to compensate for the SPM/XPM effects when operating near the critical input power. Unfortunately, this will inevitably affect the stability analysis obtained in the $\alpha = 0$ case as we observed in chapter 4, and therefore a new analysis that includes SPM/XPM effects must be performed. In this remainder of this section, we only analyze the stability of the maximal-efficiency solutions obtained in Secs. 5.3.1 and 5.3.2, and in particular, we find that 100% photon-conversion efficiency can be obtained in this case as well.

The change in cavity frequency due to SPM/XPM can be accounted for by a pre-shifting technique described in Sec. 4.3. In particular, the α terms in Eqs. (5.3–5.5) act to shift the cavity resonant frequencies from $\omega_k^{\text{cav}} \to \omega_k^{\text{NL}}$, spoiling the frequency-conservation relations necessary for efficient nonlinear frequency conversion as well as detuning the resonances from the input light. However, one can simply design the cavity frequencies to be resonant at the shifted frequencies, i.e. $\omega_k^{\text{cav}} = \omega_k^{\text{NL}}$, for a given steady-state solution. For the critical solutions corresponding to 100% photon-conversion efficiency, this implies that the new cavity frequencies will be given by:

$$\omega_0^{\text{cav}} = \frac{\omega_0}{1 - \alpha(|a_0^{\text{crit}}|^2 + |a_m^{\text{crit}}|^2 + |a_p^{\text{crit}}|^2)} \tag{5.16}$$

$$\omega_m^{\text{cav}} = \frac{\omega_m}{1 - \alpha(|a_0^{\text{crit}}|^2 + |a_m^{\text{crit}}|^2 + |a_p^{\text{crit}}|^2)} \tag{5.17}$$

$$\omega_p^{\text{cav}} = \frac{\omega_p}{1 - \alpha(|a_0^{\text{crit}}|^2 + |a_m^{\text{crit}}|^2 + |a_p^{\text{crit}}|^2)}., \tag{5.18}$$

where $|a_k^{\text{crit}}|^2$ are the energies of the modes at critical power. For cavities resonances

ω_k^{cav}, the new equations of motion are given by:

$$\frac{da_0}{dt} = \left[i\omega_0^{\text{cav}}(1 - \alpha_{00}|a_0|^2 - \alpha_{0m}|a_m|^2 - \alpha_{0p}|a_p|^2) - \frac{1}{\tau_0}\right] a_0 - i\omega_0\beta_0 a_0^* a_m a_p + \sqrt{\frac{2}{\tau_{s,0}}} s_{0,+} \tag{5.19}$$

$$\frac{da_m}{dt} = \left[i\omega_m^{\text{cav}}(1 - \alpha_{m0}|a_0|^2 - \alpha_{mm}|a_m|^2 - \alpha_{mp}|a_p|^2) - \frac{1}{\tau_m}\right] a_m - i\omega_m\beta_m a_0^2 a_p^* + \sqrt{\frac{2}{\tau_{s,m}}} s_{m,+} \tag{5.20}$$

$$\frac{da_p}{dt} = \left[i\omega_p^{\text{cav}}(1 - \alpha_{p0}|a_0|^2 - \alpha_{pm}|a_m|^2 - \alpha_{pp}|a_p|^2) - \frac{1}{\tau_p}\right] a_p - i\omega_p\beta_p a_0^2 a_m^*, \tag{5.21}$$

Note that the frequencies ω_k multiplying the β_k terms do not need to be shifted, since the terms introduced by such a shifting will be higher order in $\chi^{(3)}$. By inspection, we observe that the solutions obtained in Secs. 5.3.1 and 5.3.2 at critical input power a_k^{crit} are also solutions of Eqs. (5.19–5.20), but as explained above, their stability may change. Using the results from Secs. 5.3.1 and 5.3.2, we now study the stability properties of these solutions in the two $\Delta\omega$ regimes.

We first consider the $\Delta\omega \leq \omega_0$ regime. As in Sec. 5.3.1, we restrict our analysis to a specific parameter regime ($\tau_0 = \tau_m = \tau_p = 100/\omega_0$, $\beta = 10^{-4}$, and $\Delta\omega = 0.05\omega_0$) for simplicity, although our qualitative conclusions apply to other parameter ranges. As discussed above in Sec. 5.3.1, the maximal efficiency is obtained for input light with $|s_{0,+}|^2 = P_0$ as $|s_{m,+}|^2 \to 0$. Since one must always pump with finite $|s_{m,+}|^2$, and there are no analytic solutions in this case, we solve for the field energies $|a_k^{\text{crit}}|^2$ numerically at a small $|s_{m,+}|^2$ and for $|s_{0,+}|^2 = P_0$ in the case of $\alpha = 0$ in order to compute the shifted frequencies Eqs. (5.16–5.18). This allows us to solve the coupled-mode equations Eqs. (5.19–5.20) and therefore obtain the steady-state field amplitudes and phases. As in chapter 4, the inclusion of self- and cross-phase modulation introduces new steady-state solutions absent in the $\alpha = 0$ case, and the stability of the old and new solutions are then examined again via a linear stability

analysis, as in Sec. 5.3.1. In particular, we find that the inclusion of SPM/XPM does not destroy the stability of the maximal efficiency solution in the $\alpha = 0$ case studied in Sec. 5.3.1, and in fact creates additional stable solutions, as shown in Fig. 5-8. A similar analysis can be performed in the $\Delta\omega > \omega_0$ regime, where it is possible to obtain the analytic form of the maximal efficiency solutions [Eqs. (5.11–5.13) in Sec. 5.3.2]. We find that, as in the previous regime, the presence of α introduces additional stable solutions, while retaining the original 100% efficiency $\alpha = 0$ solution, over finite regions of the parameter space.

Therefore, the inclusion of these effects is not prohibitive for 100% nonlinear frequency conversion although predicting which parameter regimes allow for such conversion will depend on the system under question. As in chapter 4, the presence of multiple stable solutions means that the manner in which the source is initiated will determine which solution is excited, but a simple initialization procedure similar to that in Sec. 4.6 should be possible to excite the maximal-efficiency solution.

5.4 conclusions

By exploiting a simple but rigorous coupled-mode theory framework, we have demonstrated the possibility of achieving highly-efficient (low-power) DWFM in triply-resonant cavities, similar to our previous work in SHG and THG [71, 169]. We conclude that there are two main regimes of operation, determined by the ratio of $\Delta\omega$ to ω_0. In particular, whereas the maximal efficiency obtainable in the $\Delta\omega \leq \omega_0$ regime, corresponding to conversion between closely-spaced resonances, is bounded above by a quantum-limited process, there is no such bound when $\Delta\omega > \omega_0$. In both regimes, a suitable choice of system parameters leads to stable, maximal-efficiency nonlinear frequency conversion, even in the presence of SPM and XPM effects. We remark that all of the results obtained in this paper correspond to the idealized case of lossless interactions, since the main focus of the paper is in examining the basic considerations involved in operating with these systems rather than predicting results for specific experimentally-relevant systems. Nevertheless, based on our

previous experience with SHG and THG, we expect that linear and nonlinear losses, e.g. coming from radiation or material absorption, will only act to slightly decrease the overall conversion efficiency and will not affect the qualitative predictions here.

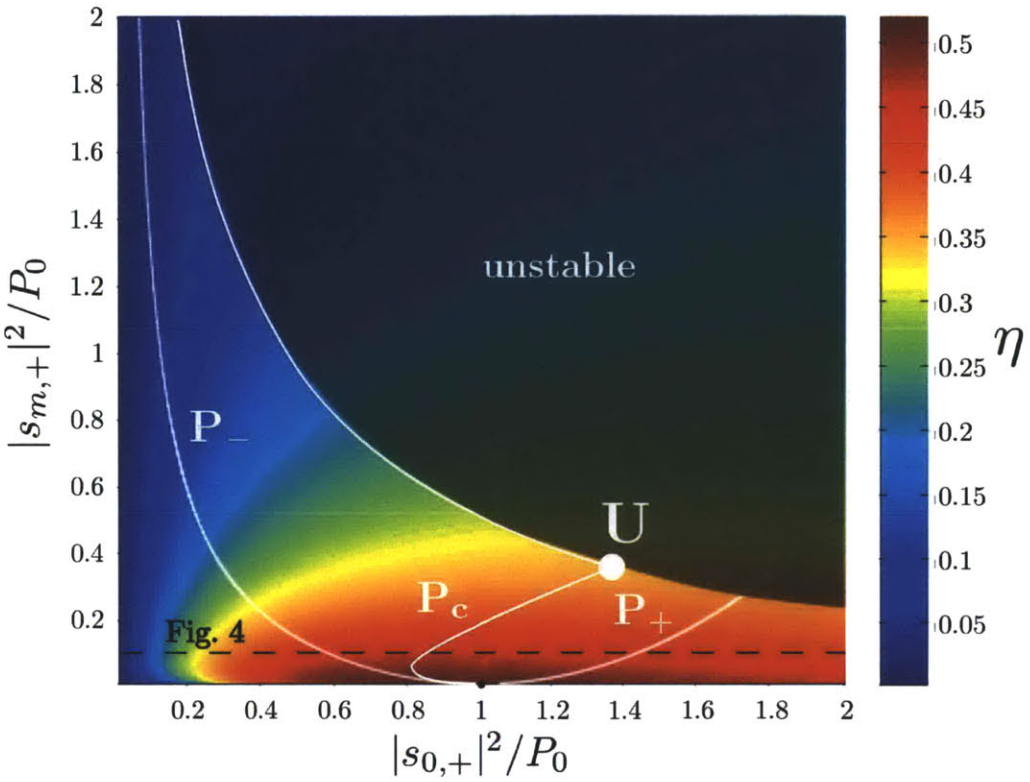

Figure 5-3: Color plot of the steady-state conversion efficiency $\eta = |s_{p,-}|^2/(|s_{0,+}|^2 + |s_{m,+}|^2)$ as a function of input power $|s_{0,+}|^2$ and $|s_{m,+}|^2$, for a system consisting of $\Delta\omega = 0.05\omega_0$, $\beta = 10^{-4}$, and $\tau_0 = \tau_p = \tau_m = 100$. Both powers are normalized by the critical power $P_c(|s_{m,+}|^2 \to 0) = P_0 = 2/\tau_0|\beta|\sqrt{\tau_m\tau_p|\omega_m\omega_p|}$ (black dot). The shaded region indicates that the solution is unstable. The curves $P_\pm$ indicate the powers at which depletion of the ω_0 input light is achieved, i.e. $s_{0,-} = 0$; the critical power $P_c(|s_{m,+}|^2)$ is defined as the total input power that yields the highest *stable* efficiency for any given $|s_{m,+}|^2$. The dash line is the cross-section shown in Fig. 5-4.

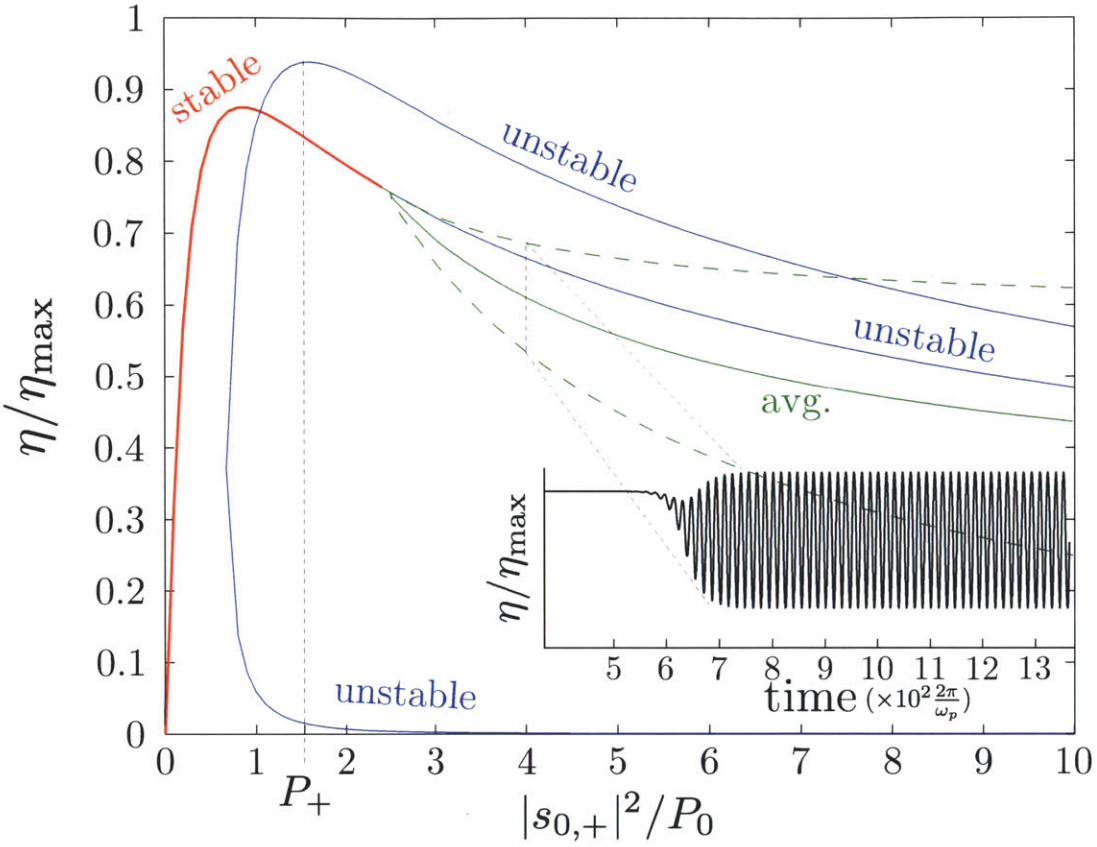

Figure 5-4: Bifurcation diagram of the steady-state efficiency η, normalized by the quantum-limited maximum efficiency $\eta_{\max} = \frac{1}{2}\omega_p/\omega_0$, as a function of $|s_{0,+}|^2$, normalized by P_0, for signal power $|s_{m,+}|^2 = 0.1P_0$ (indicated by the black dashed line of Fig. 5-3). Red/blue correspond to a stable/unstable solution (note that the two bifurcating solutions are always unstable). The green dashed line illustrates the bounds of the limit cycles obtained from time domain simulations, where the solid green line yields the average over the cycle. (*Inset:*) Efficiency as a function of time in units of the period $T_p = 2\pi/\omega_p$ in a regime where there exists a limit cycle.

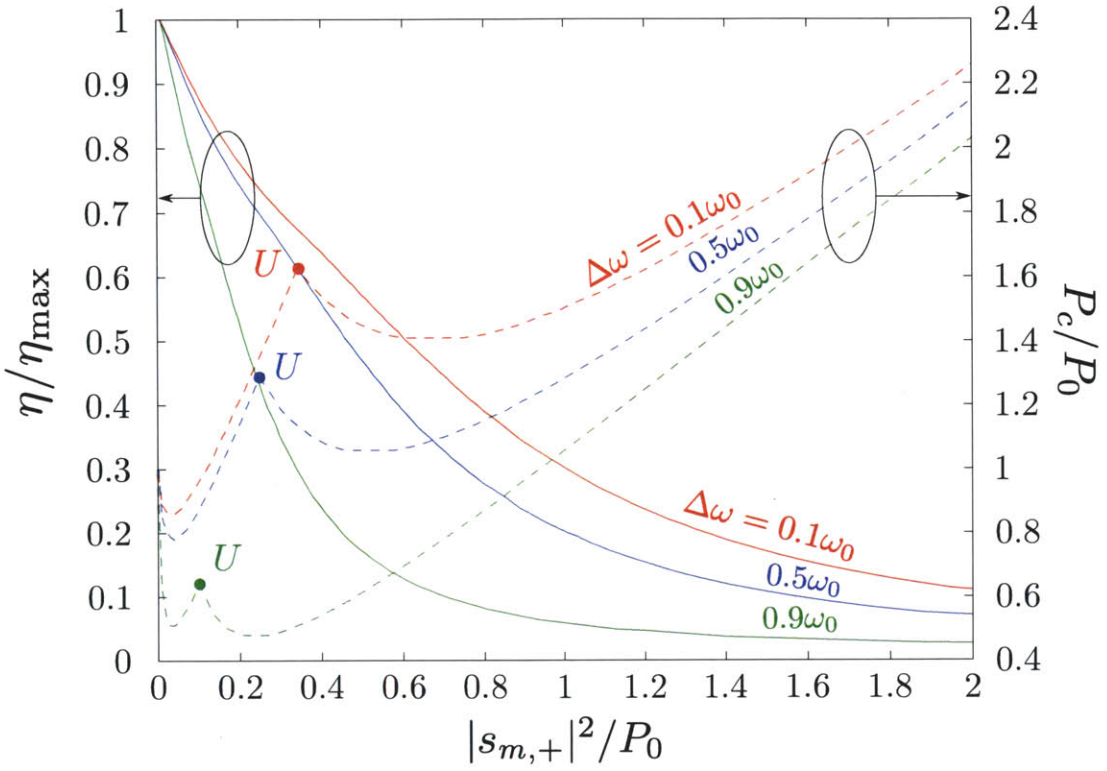

Figure 5-5: Plot of the steady-state efficiency η (solid lines) along the critical solution [total input power $P_c(|s_{m,+}|^2) = |s_{m,+}|^2 + |s^c_{0,+}|^2$ that yields the maximum efficiency for a given $|s_{m,+}|^2$ (solid white curve of Fig. 5-3)] and the value of P_c (dashed lines) as a function of $|s_{m,+}|^2$, normalized by P_0, for three different values of $\Delta\omega$: $0.1\omega_0$ (red), $0.5\omega_0$ (blue), and $0.9\omega_0$ (green). The kinks in the P_c curves correspond to the point U where P_c reaches the region of instability (see Fig. 5-3). The coupling lifetimes τ and coefficient β of the system are equivalent to those of Fig. 5-3.

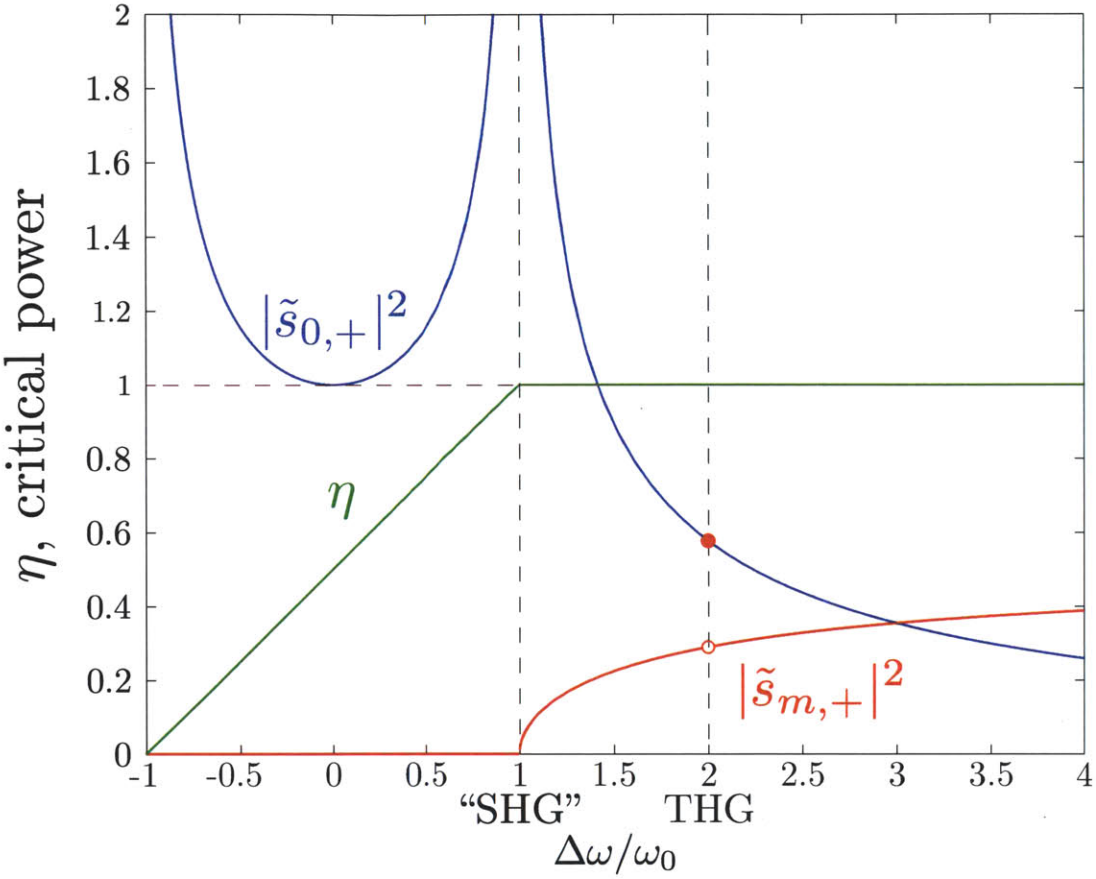

Figure 5-6: Plot of the critical powers $|\tilde{s}_{0,+}|^2$ (blue), $|\tilde{s}_{m,+}|^2$ (red), and maximum steady-state efficiency η (green) as a function of $\Delta\omega/\omega_0$ (the tilde over the critical powers indicates that the values have been rescaled by the factor $4/\tau_0|\beta_0|\sqrt{\tau_m\tau_p\omega_0^2}$). The vertical dashed lines at $\Delta\omega = \omega_0$ and $\Delta\omega = 2\omega_0$ indicate special degenerate regimes, corresponding to "second harmonic generation" (SHG) and third harmonic generation (THG). (Note the discontinuity in $|\tilde{s}_{m,+}|^2$ located at $\Delta\omega = 2\omega_0$, explained in the text).

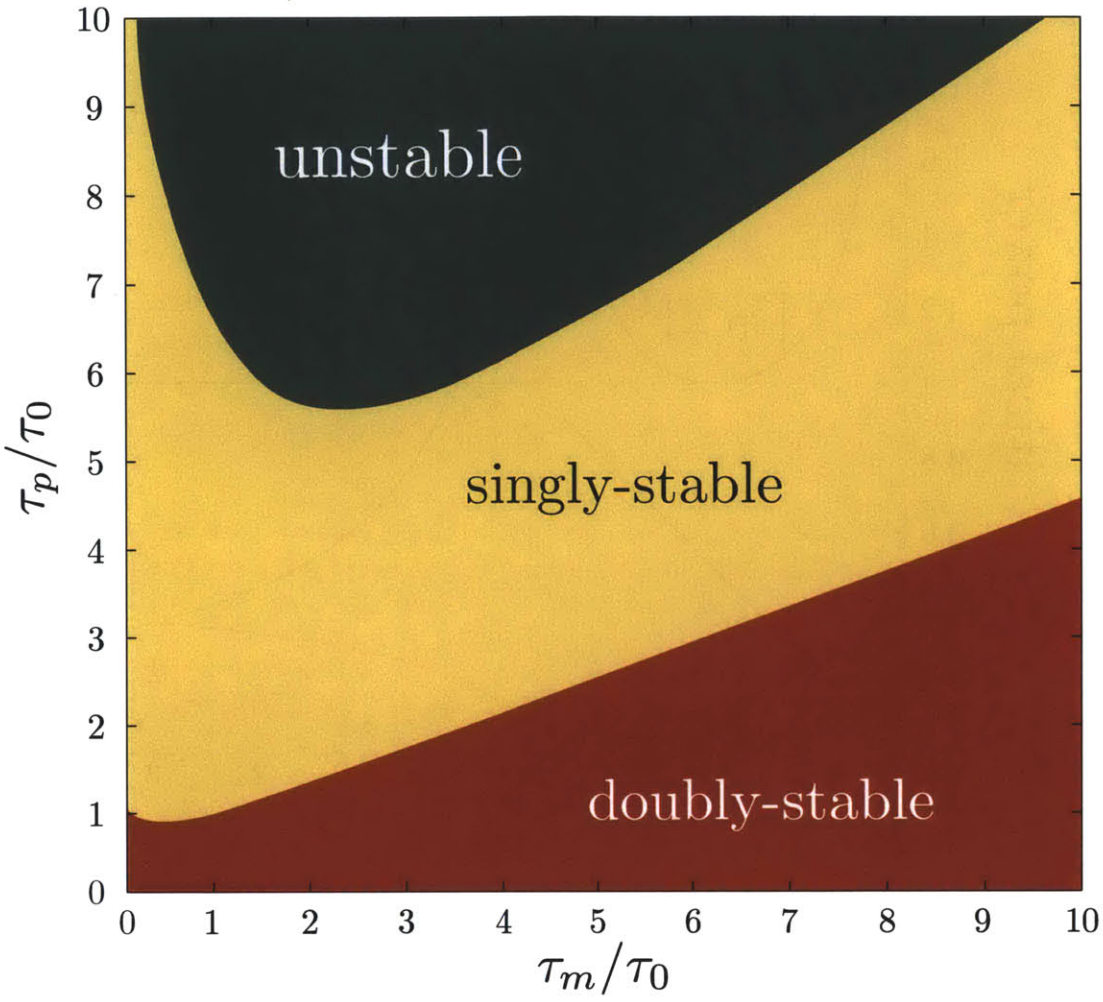

Figure 5-7: Stability contours (number of stable solutions) as a function of modal lifetimes τ_m and τ_p, normalized by τ_0, pumping at the critical input powers $|s_{0,+}^{\text{crit}}|^2$ and $|s_{m,+}^{\text{crit}}|^2$. The stability in the $\Delta\omega \geq \omega_0$ regime is independent of the value of $\Delta\omega$.

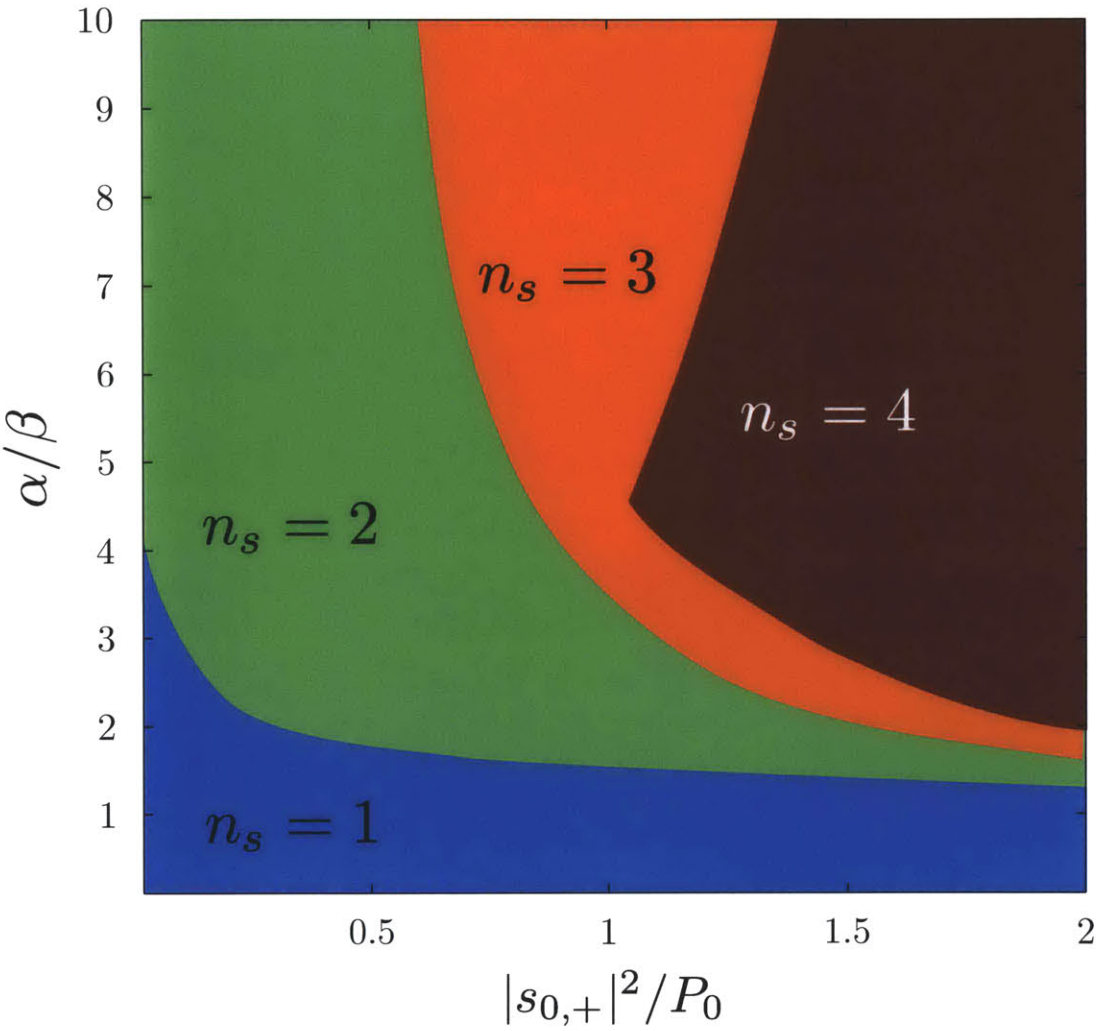

Figure 5-8: Contour plot of number of stable solutions (n_s) as a function of α/β and $|s_{0,+}|^2/P_0$, for input pump power $|s_{m,+}|^2 = 0.1P_0$, and for the system described in the text.

Chapter 6

Designing Waveguide–Cavity Systems for High-efficiency Frequency Conversion

In the previous chapters, we have performed a theoretical analysis of frequency conversion in resonant cavities coupled to input and output channel(s). We used temporal coupled-mode theory to model such systems and used the model to do an analysis of these complex dynamical systems, find their solutions and check their stabilities. We obtained results telling us for what system parameter regimes, there exists a stable high efficiency solution and how we can excite such a solution in practice.

The next step to realization of these processes is to design actual systems that have such properties and parameters we are looking for and can operate in experiments and demonstrate such high efficiency nonlinear processes. The design process starts with selecting a type of optical cavity that supports all the involved modes and frequencies and is made of nonlinear material. Next is to design the cavity and set its parameters and geometry so that it is frequency matched with correct wave vectors. Next is to decide on the input/output channel(s): whether to use one or two separate channels and design the parameters so that they support the frequencies and wave vectors they need to couple to from the cavity modes.

Once we have the cavity designed, we can compute the coupling coefficients between the modes. And once we have all the pieces of the system together, we can place the waveguide(s) at appropriate distance from the cavity so that it gives us the cavity lifetimes of the modes that we are looking for and that fit the parameter regime we are looking at. Note that we have little control over what the mode coupling coefficients come out to be. The best we can do is to choose different mode profiles with spatial configuration that we think would give stronger overlap between the modes. However, it is not easy to predict how the exact value of the coupling coefficients would change by changing different cavity parameters. On the other hand, we have a lot more control on the values of the field lifetimes in the cavity; we can directly control them by changing the distance between the cavity and the waveguide(s). As we will discuss in more detail, this gives us a good reason to use two separate waveguides as the input and output channels so we have full control over the cavity lifetimes independently and therefore have more parameters we can play with to get the system into the desired regime.

In the case of $\chi^{(3)}$ processes such as third harmonic generation (THG) and DFWM, there is an additional complexity in the design process: the self- and cross-phase modulation shift the cavity frequencies out of resonance. Therefore as proposed in chapter 4, we pre-shift the cavity so that it falls into resonance at the desired frequency at the critical power. We will discuss how this should be done when designing the cavity.

In this chapter we go in detail in the design process and ideas. In the end, we give a demonstration of a 2d system that exhibits second harmonic generation. Simulations of the system including losses give 89% conversion efficiency. My initial work on this problem established many of the basic design parameters of the problem, including the technique of a dual-coupled ring-resonator system with frequency-cutoff waveguides to independently control the Q's at the fundamental and harmonic frequencies, as well as the computational techniques. We then found that an even simpler realization could be attained for $\chi^{(2)}$ second-harmonic generation (SHG) where no pre-shifting is required, and calculations and design for

the SHG case were carried out by a visiting student, Zhuanfang Bi, in collaboration with me and our other co-authors. These SHG calculations are included in the second part of this chapter and were recently submitted for publication [16].

6.1 Cavity design

The design process starts with selecting a type of optical devise we would wish to use as the non-linear optical resonator where the frequency conversion occurs. This cavity must be made of non-linear medium ($\chi^{(2)}$ or $\chi^{(3)}$) and must support the modes and frequencies involved in the frequency conversion process.

There have been several different optical resonators proposed for the purpose of frequency conversion including ring resonators [19, 112, 207] and complete-gap photonic crystals [194]. Ring resonators have the advantage that they are easy to frequency match specially in applications where the spacing between the involved frequencies is large; It is relatively easy to design a ring resonator that supports modes both at frequency ω_1 and its double $\omega_2 = 2\omega_1$ or its triples $\omega_3 = 3\omega_1$. On the other hand, they have the disadvantage that there is always a trade-off between the modal volume (size of the ring) and the radiation losses Q_e: the smaller the modal volume, the higher losses the resonator acquires. Complete-gap photonic crystals, however, do not suffer from this loss-modal volume tradeoff but they are difficult to frequency match for processes in which the spacing between the frequencies involved is large such as in second and third harmonic generation processes: it is difficult to design such cavities that have complete band gaps both at the incident frequency and its double/triple. On the other hand, for degenerate four-wave mixing in the regime where $\Delta\omega$ is small, designing a photonic crystal that has a complete band gap that contains all the three frequencies ω_0, ω_m, and ω_p is easy. Given these facts, it is more sensible to choose a ring resonator as the nonlinear optical cavity for the second and third harmonic generation as shown in Figs. 6-3 and 6-1 and a complete-band gap photonic crystal for the four-wave mixing Fig. 6-2.

Once we have chosen the cavity type we will use, the next step is to design the

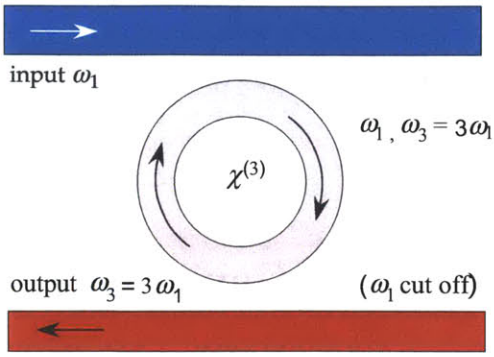

Figure 6-1: Schematic ring-resonator waveguide-cavity system: input light from a waveguide supporting a propagating mode of frequency ω_1 (input power $|s_{1+}|^2$) is coupled to a ring-resonator cavity mode of frequency ω_1, converted to a cavity mode of three times the frequency $\omega_3 = 3\omega_1$ by a nonlinear $\chi^{(3)}$ process, and coupled out by another waveguide supporting a propagating mode of frequency ω_3 (the waveguide does not support a propagating ω_1 mode).

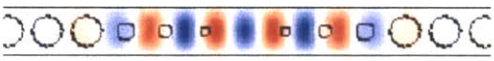

Figure 6-2: Example photonic-crystal cavity system for DFWM in 2d, where the photonic crystal consists of a periodic sequence of air holes in a dielectric waveguide [84]. Calculations performed by T. Alcorn, a UROP student I am helping to supervise.

cavity and frequency-match it. To do this, we use a FDTD simulation of Maxwell's equation in the linear system that gives us the resonant modes of the cavity. Then we use either an optimizer or a zero finder algorithm to change the geometrical parameters of the cavity so that the cavity is frequency-matched and quasi-phase matched. Once we have the cavity modes, we can easily compute the field coupling coefficients by integrating the fields given by time-domain simulations of the system at steady state. We go in detail in the design process of such a ring resonator in 2d for the purpose of second harmonic generation in Sec. 6.3

6.2 Input/output channel(s): waveguide design

Next is to design the input/output channel waveguide(s) and couple them to the cavity. There are several important criteria and aspects of waveguide design:

- Obviously, the waveguides have to guide the frequencies/modes that they carry.
- The waveguide must be phase-matched to the cavity modes in order to optimize the coupling between them.
- The waveguides and their positions relative to the resonant cavity determines the mode quality factors Q_1 and Q_k which affect several tradeoffs:
  - The (fractional) bandwidth of conversion is $1/\max(Q_1, Q_k)$.
  - The critical power is proportional to $1/Q_1^2 Q_2$ in the case of 2nd harmonic generation and $1/Q^2$ in the case of THG and FWM.
  - The sensitivity to perturbations in the structure is determined by $\min(Q_1, Q_k)$.
- The design of the input/output waveguide(s) is critical to ensure that the system supports only a single incoming and a single outgoing wave at both ω_1 and ω_2—additional channels will lower the efficiency (unless they have much larger coupling Q) [83].

Given all these requirements the input/output channel(s) have to satisfy, it becomes a natural choice to us two separate waveguides as separate input and output channels. This is specially helpful since it gives us independent control over Q_1 and Q_k. If we used only one waveguide instead, it is well known that higher-frequency modes will have higher Q, so coupling the ring to a single waveguide would yield $Q_k \gg Q_1$ and so we would be very limited on where in the parameter space and what efficiency regimes we can get the system in. In particular, in the case of 3rd harmonic generation, as we saw in chapter 4, the complete conversion solution is unstable for $Q_3 > Q_1$.

However, the last requirement above adds a new twist to the waveguide design: we have to ensure that each waveguide only supports one of the frequencies/modes and

not the other one: the input channel must only carry the ω_1 and the output channel must only couple to the $\omega - k$ mode. The former is already satisfied modulo small errors. As we discussed before given same waveguide-cavity spacing, the quality factor of the lower frequency Q_1 will be much smaller than Q_k. Therefore the tiny leakage of the Q_k mode in the input channel is negligible specially if the coupling to the output channel is much stronger (higher quality factor). The latter however, we need to take care of, for exact same reason: if we do nothing, the coupling Q_1 to the output channel will be comparable to the coupling to the input channel. But leakage of the input signal to the output channel is a loss and would bring down the conversion efficiency significantly. To do this, we design the output waveguide so that it has a *low-frequency cuttoff* $> \omega_1$.

Once we have the waveguides phase- and frequency-matched separately, then it remains to decide on their position (distance relative to the resonant cavity. The distance of each is determined by what Q we are looking for.) However, to minimize the losses, the coupling quality factor Q_s must be much smaller than the radiation quality factor Q_e of the mode. Otherwise, we will loose too much energy due to radiation.

Once we have everything in place, we know all the system parameters and can therefore calculate the critical input power that would give us the stable high efficiency solution. Note that while designing the system, we should have been aware of the different system regimes where the maximum efficiency solution is stable and relatively easy to excite. This seems to be much easier in the case of 2nd harmonic generation since the dynamics are simple and the high efficiency solution is either stable or falls in high-average-efficiency limit cycles.

Finally, we are ready for a full time-domain FDTD simulation of the nonlinear system. We add the nonlinearities to the resonant cavity, put a source at the critical power in the input channel, and hope to get high efficiency frequency conversion and an output signal at ω_k coming out of the output waveguide.

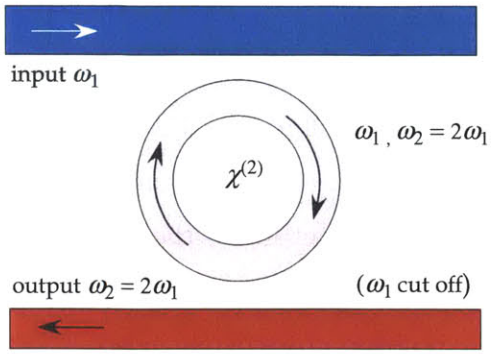

Figure 6-3: Schematic ring-resonator waveguide-cavity system: input light from a waveguide supporting a propagating mode of frequency ω_1 (input power $|s_{1+}|^2$) is coupled to a ring-resonator cavity mode of frequency ω_1, converted to a cavity mode of twice the frequency $\omega_2 = 2\omega_1$ by a nonlinear $\chi^{(2)}$ process, and coupled out by another waveguide supporting a propagating mode of frequency ω_2 (the waveguide does not support a propagating ω_1 mode).

6.3 Computational methods: a $2d$ design for a 2nd harmonic generation process

In order to develop the harmonic-generation design, we needed to compute micro-cavity modes, frequencies, and lifetimes (Q), as well as waveguide dispersion relations. The final design was evaluated both semi-analytically with coupled-mode theory (CMT) and with a full nonlinear Maxwell simulation. We began by studying a two-dimensional (2d) model system, and continued to full 3d calculations. The computational methods for these calculations are described here.

The basic cavity design is that of a ring resonator coupled with two adjacent waveguides, as depicted in Fig. 6-3. To begin with, we studied the isolated cavities, uncoupled to any waveguide. Since the isolated waveguide is axisymmetric, it can be modeled in cylindrical coordinates. We did so using a free finite-difference time-domain (FDTD) software package (Meep) [146]. The simulation cell is surrounded by a perfectly matched layer (PML) absorbing boundary region. The use of cylindrical coordinates in this simulation reduces the 2d problems to a 1d problem, thereby reducing simulation times significantly. Another advantage is that, the angular dependence of the fields in systems with continuous rotational

symmetry can be given by the angular momentum parameter (input variable) m, which is easy to control. To begin with, we inserted a broad Gaussian pulse in the structure in order to excite all of the (TM polarized) modes within a chosen bandwidth and with a fixed m; we then re-ran the simulation with a narrow-band source around each mode and outputted the corresponding fields at the end. The resonance frequency and lifetime Q_{rad} were obtained by Harminv, which is a free program to solve the problem of harmonic inversion [130]. The waveguide modes were computed using an iterative eigenmode solver in a planewave basis, using a freely available software package (MPB) [85].

The combined waveguide-cavity system, with waveguides adjacent to the ring resonator, is not axisymmetric and requires a full 2d FDTD calculation. Bringing in the waveguides, there are two decay mechanisms for the modes in this cavity: the mode can decay into the adjacent waveguides, and it can radiate into the surrounding air. The total dimensionless decay rate $1/Q_{tot}$ can be written as the sum of two decay rates: $1/Q_{tot} = 1/Q_w + 1/Q_{rad}$, where $1/Q_w$ and $1/Q_{rad}$ are the waveguide and radiative decay rates, respectively. We obtain Q_{tot} from a filter-diagonalization analysis of the field decay in FDTD [146]. However, we also need to know Q_w and Q_{rad} individually, both of which are modified for different ring–waveguide separations. Therefore, for each separation, we computed the linear transmission spectrum in FDTD. Then, comparing with the transmission equation obtained from coupled mode theory [83],

$$T(\omega) = \frac{|s_-|^2}{|s_+|^2} = \frac{\omega_0^2(/Q_w - 1/Q_{rad})^2 + 4(\omega - \omega_0)^2}{\omega_0^2(1/Q_w + 1/Q_{rad})^2 + 4(\omega - \omega_0)^2}, \tag{6.1}$$

one can solve for both Q_{rad} and Q_w given $T(\omega_0)$ (the minimum T) and Q_{tot}.

Given the waveguide modes, frequencies, and Q values, the resulting nonlinear system can be modeled semi-analytically using the coupled-mode equations of Ref. [169]. A key parameter of this model is the overlap integral β_1 between the fundamental (ω_1) and harmonic ($\omega_2 = 2\omega_1$) modes, which is computed from the

following integral:

$$\beta_1 = \frac{1}{4}\frac{\int \mathrm{d}^3x \sum_{ijk} \epsilon_0 \chi^{(2)}_{ijk} \left(E^*_{1i}(E_{2j}E^*_{1k} + E^*_{1j}E_{2k})\right)}{\left(\int \mathrm{d}^3x\, \epsilon |E_1|^2\right)\left(\int \mathrm{d}^3x\, \epsilon |E_2|^2\right)^{1/2}}, \tag{6.2}$$

This coupling coefficient can be obtained by applying perturbation theory to Maxwell's equations in the presence of a $\chi^{(2)}$ nonlinearity, as explained in Ref. [169] and very similar to our analysis in chapter 3. β_1 and $beta_2$ coefficients are related to one another via conservation of energy: $\omega_1\beta_1 = \omega_2\beta_2^*$.

CMT makes several approximations: it assumes that cavity-waveguide and cavity-radiation coupling is weak (high Q), it neglects nonlinear coupling to modes not at $2\omega_1$ (the rotating-wave approximation), it assumes that the input waveguide couples only to a single direction mode of propagation around the ring (clockwise or counterclockwise), and correspondingly that each ring mode couples out to only a single direction of propagation in the waveguides (despite the fact that we use waveguides that are not identical to the ring structure for reasons described below). To ensure that the full complexities of the nonlinear Maxwell equations in this geometry were accurately captured by the coupled-mode equations, we also performed fully nonlinear FDTD simulations (where the nonlinear constitutive equations are solved by Padé approximates [146]). In particular, we excited the input waveguide with a continuous plane wave of frequency ω_1, and computed the SHG power escaping through the output waveguide at ω_2. The nonlinear FDTD calculations were then checked against the CMT predictions, using the frequencies, decay rates, and β_1 coefficient computed from a set of linear FDTD simulations.

In the absence of losses or reflections, 100% conversion is achieved in the steady state. The critical power (including the possibility of losses,

$$\left|s^{crit}_{1+}\right|^2 = \frac{\omega_1 Q_1}{2\left|\beta_1\right|^2 Q_{w,2} Q^3_{w,1}},$$

can then be related to the various linear modal parameters, and ends up scaling as V/Q^3. Hence, one can in principle obtain very low-power efficient harmonic

conversion by increasing Q and/or decreasing V (i.e., maximize β). The efficiency decreases if the power is either too low (in the linear regime) or too high (dominated by down-conversion) compared to $\left|s_{1+}^{crit}\right|^2$.

6.3.1 Ring-resonator design

Previously, achieving efficient SHG in waveguides or Fabry–Perot etalons required techniques to obtain "phase-matching" of the fundamental and harmonic modes [57]. Phase matching is a selection rule arising from the approximate translational symmetry for propagation over long uniform regions, according to which the fundamental and harmonic modes must have the same phase velocities in order to couple efficiently. In microcavities where the fields are confined to within a few wavelengths, such a constraint is instead replaced by selection rules resulting from symmetry considerations, which determine whether the overlap integral in Eq. (6.2) is nonzero. In our geometry, involving cylindrically symmetric cavities, the fields can be chosen to have azimuthal dependence $\sim e^{im\phi}$, determined by the conserved angular momentum "quantum number" $m \in \mathbb{Z}$; by simple inspection of Eq. (6.2) for $\chi^{(2)}_{zzz}$ coupling, this leads to the requirement $m_2 = 2m_1$, where m_1 and m_2 are the corresponding quantum numbers of the fundamental and harmonic modes, respectively. Because m is constrained to integer values, perturbing the cavity parameters does not alter the m of a given mode, so the $m_2 = 2m_1$ condition is easy to satisfy and robust. On the other hand, perturbing the cavity parameters does change the frequencies of the modes at given m values, so the key difficulty is to find modes at a given pair of m's that satisfy $\omega_2 = 2\omega_1$.

For any given geometry and given m, we compute the resonant modes of the linear system by FDTD in cylindrical coordinates, which reduces to a 1d problem in r for a given m. We then compute the modes at m_1 and $2m_1$ for various choices of m_1, and vary the geometry until we find a pair of modes with $\omega_2 = 2\omega_1$. (More precisely, the finite bandwidth means that there is some tolerance $\sim 1/Q$ on the frequency mismatch $\Delta\omega = \omega_2 - 2\omega_1 \lesssim \omega_2/Q$.) We began exploring the the simplest possible space of designs: we considered a fixed ring width a and varied the inner radius R,

as depicted in the right inset of Fig. 6-4. However finding modes that satisfy both phase- and frequency-matching requirements for this single-ring design, at chosen index contrast, turned out to be unfeasible. In particular, Fig. 6-4 shows the frequency mismatch $\Delta\omega$ (units of $2\pi c/a$) corresponding to two phase-matched modes with $m_1 = 15$ and $m_2 = 30$, as a function of the ring radius R. As observed, $\Delta\omega$ is always positive and large $\gtrsim 0.04$ $(2\pi c/a)$ relative to the mode bandwidths ($Q \sim 1000$s) considered here, and we obtained similar results for other values of m_1 and other pairs of modes. We therefore abandoned the single-ring design and instead considered a double-ring structure, depicted schematically in the left inset of Fig. 6-4. The reason to consider such a structure is based on the fact that, for the single ring resonator, the $\omega_2 = 2\omega_1$ condition was most nearly matched by a pair of modes where the ω_1 mode was approximately even-symmetrical through the middle of the waveguide and the ω_2 mode was approximately odd-symmetrical. We therefore expect that an air groove introduced within the waveguide will affect the ω_1 mode more than the ω_2 mode (which has a node in the waveguide), thereby allowing us to shift ω_1 into $\omega_2/2$. Furthermore, by introducing the air groove off-center in the waveguide, we break the approximate mirror symmetry of the waveguide (which is only slightly broken by the small curvature), and thereby allow the ω_1 mode and ω_2 modes to have a much larger overlap integral. (Exactly even and odd modes would not couple at all, a fact we return to in 3d.) Thus, this structure yields two additional degrees of freedom, the widths of the inner and outer rings, w_1 and w_2 respectively, with which it is possible to shift ω_1 relative to ω_2. After a relatively small search through the design space, we settled on a structure having $w_1 = 0.45a$ and $w_2 = 0.35a$, where a denotes the total width of the double-ring structure including the air groove of width $a - w_1 - w_2 = 0.2a$ (labeled in the left inset of Fig. 6-4). For this configuration, we find that, for the same azimuthal modes $m_1 = 15$ and $m_2 = 30$ as above, $\Delta\omega$ switches sign as R increases and is zero at $R \approx 4.51a$, as shown by the red line of Fig. 6-4. Using the field patterns of

these modes, we also computed the coupling coefficient β_1, which we found to be

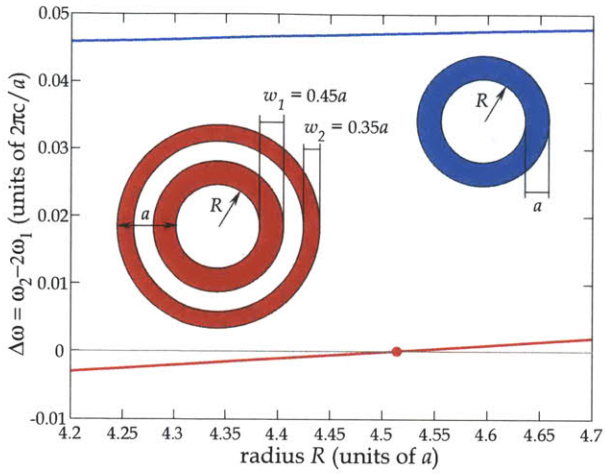

Figure 6-4: Plot of the frequency difference $\Delta\omega = \omega_2 - 2\omega_1$ (units of $2\pi c/a$) of two $LiNbO_3$ ring-resonator modes of frequencies ω_1 and ω_2, and azimuthal momentum $m_1 = 15$ and $m_2 = 30$, respectively, corresponding to two different ring-resonator geometries (insets), as a function of inner radius R. The blue and red lines correspond to the single-ring (right inset) and double-ring resonators.

$$\beta_1 \approx -0.000017848693063 - 0.000026297992627i \text{ (in units of } \chi^{(2)}/a^{3/2}).$$

As described below, the final ring parameters are slightly perturbed by the presence of the input/output waveguides, so our design procedure was to design the isolated ring, design the coupling waveguides as described below, and then tweak the ring design to restore the $\omega_2 = 2\omega_1$ condition. We obtained a final ring radius of $R = 4.585$ for frequencies $\omega_1 = 0.277172 \cdot 2\pi c/a$ (vacuum wavelength $3.6a$) and $\omega_2 = 0.554344 \cdot 2\pi c/a$ (vacuum wavelength $1.8a$).

6.3.2 Input/output coupling waveguides

Given a ring resonator with appropriate modes as described above, we must then design the coupling to adjacent waveguides so that they have the desired coupling lifetimes $Q_{w,1}$ and $Q_{w,2}$ at ω_1 and ω_2. In order to obtain small radiation losses, these coupling lifetimes should be much smaller than the radiative lifetimes $Q_{rad,1}$ and $Q_{rad,2}$ (approximately 10^5 and 10^7 for the isolated ring), the so-called "over-coupled" regime. For this reason, and also to obtain a reasonable bandwidth of conversion, and to limit computation times, we chose to work with $Q_w \sim 10^4$.

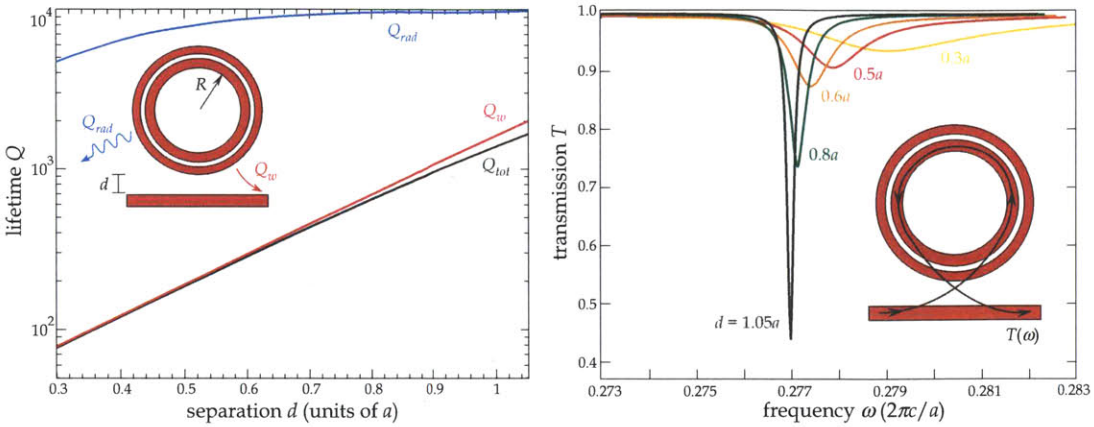

Figure 6-5: (Left:) Semilog plot of the radiative (Q_{rad}), waveguide-coupling (Q_w), and total (Q_{tot}) lifetimes of the ω_1 mode of Fig. 6-6, as a function of the ring-waveguide separation d_1. (Right:) Corresponding transmission spectrum at various separations.

Furthermore, we don't want to have to bring the waveguide too close to the ring, which would require a high computational resolution and might also induce additional radiative scattering.

In order to obtain Q_w values that are not too large for moderate ring–waveguide separations, we phase-match the waveguide mode to the ring-resonator mode. Conceptually one designs the waveguide to have a phase velocity equal to $m/r\omega$, the phase velocity of the ring mode, but this condition is ambiguous because the "phase velocity" of the ring mode varies with radius r. For a large R, the difference between R and $R + a$ is negligible and so one can simply use a waveguide of the same width as the ring [204]. In our case, however, R/a is too small for this to be a good approximation and an identical waveguide is not optimal. Instead, we varied the waveguide width to minimize Q_w for a given ring–waveguide separation d. Furthermore, this allowed us to have good coupling between the double-ring structure and a simple dielectric waveguide with no air groove, as well as with the cutoff waveguide described below. For example, for the ω_1 mode we found that this procedure corresponded to an optimal phase velocity $m/r\omega$ with $r = R + 1.3a$. Fixing the waveguide width and varying the ring–waveguide gap d, we obtain the plot of the $Q_{w,1}$ and $Q_{rad,1}$ in Fig. 6-7(a), which illustrates two effects. F First, the coupling Q_w decreases exponentially with d, thanks to the exponentially increasing

overlap of the evanescent tails of the waveguide and cavity modes. Second, although for large d the losses Q_{rad} asymptotes to a constant given by the radiation loss of the isolated ring, for sufficiently small d the radiation losses increase due to scattering of the cavity mode from the waveguide. As explained in Sec. 6.3, we obtain Q_{rad} and Q_w from a combination of mode decay and transmission simulations, and several of these transmission spectra are shown in Fig. 6-7(b). When $Q_{rad} \gg Q_w$, the linear ring–waveguide system approaches an all-pass filter with 100% transmission (but a resonant delay), while as Q_w approaches Q_{rad} for large d the radiation loss increases and one observes a resonant dip in the transmission.

In principle, we could use a single waveguide to couple ω_1 into the ring resonator and ω_2 out after SHG. However, for the same ring–waveguide separation d the ω_2 mode will normally have $Q_{w,2} \gg Q_{w,1}$ because the evanescent tails of a waveguide mode decay more rapidly at higher frequencies. (Furthermore, if the waveguide is optimized to couple ω_1, the phase matching will not be optimized at ω_2.) For example, if we designed the ring to couple to a single dielectric waveguide at ω_1 with a gap $d_1 = 0.5a$, we would obtain $Q_{w,1} = 191$ and $Q_{2,w} = 155,000$, and the difference increases for larger d_1. Instead, therefore, we couple ω_2 out with a *second* waveguide that is optimized to couple at that frequency, on the other side of the ring. This introduces a new problem, however: we must prevent the second waveguide from coupling ω_1, since CMT predicts that the introduction of multiple ω_1 channels will lower the attainable conversion efficiency significantly. We solve this problem by designing the second waveguide to have a *low-frequency cutoff* $> \omega_1$, so that ω_1 cannot couple out via that channel. In 3d below, a low-frequency cutoff is introduced simply by the waveguide substrate; in our 2d model system here we obtain a cutoff by the simple expedient of placing the ω_2 waveguide next to a perfect electric conductor (PEC).

By this procedure, we obtain an input waveguide of width $0.5a$ and an output waveguide of width $0.35a$ (adjacent to PEC), with dispersion relations shown in Fig. 6-5. Note the cutoff in the ω_2 mode, as desired. In Fig. 6-6, we show the field distribution of the two modes in the ring resonator coupled with these two

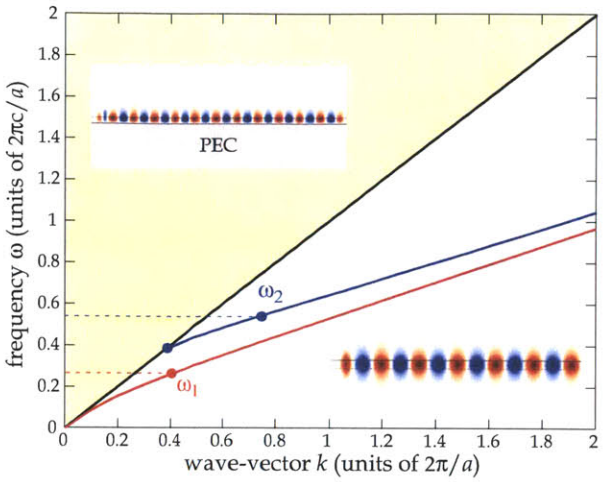

Figure 6-6: Band diagram or frequency ω (units of $2\pi c/a$) as a function of wave-vector k (units of $2\pi/a$), corresponding to the fundamental (red line) and second-order (blue line) modes of two two different $LiNbO_3$ waveguides of thickness $w_1 = 0.5a$ and $w_2 = 0.35a$, respectively. Here, a denotes the thickness of the double-ring resonator of Fig. 6-4. The right and left insets show the E_z field profile (blue/white/red denote positive/zero/negative amplitude) of two different modes, with frequencies $\omega_1 = 0.277(2\pi c/a)$ and $\omega_2 = 2\omega_1$, and corresponding wave-vectors $k_1 = 0.39(2\pi/a)$ and $k_2 = 2k_1$, respectively.

waveguides; for illustration. Note that the ω_2 mode has negligible leakage into the upper waveguide because the coupling Q in that direction is so much larger. Also, note that each ring mode (propagating counter-clockwise) couples primarily to waveguide modes traveling in the same direction. This is critical in order to mimic the theoretically optimal situation as described in Ref. [169]: ω_1 must enter the resonator from a single channel and exit in a single channel for ω_1 and in a single channel for ω_2. For the nonlinear simulations below, we use larger values of d_1 and d_2 in order to obtain a lower critical power. In particular, we use $d_1 = 1.05a$ and $d_2 = 0.7a$, obtaining $Q_{w,1} = 2000$ and $Q_{w,2} = 8992$.

6.3.3 Nonlinear characterization and SHG efficiency

Given these parameters of the linear resonator system, CMT can predict the behavior of the nonlinear system when a $\chi^{(2)}$ is introduced. In particular, it predicts that 82% SHG efficiency should be obtained at a certain critical power ($< 100\%$

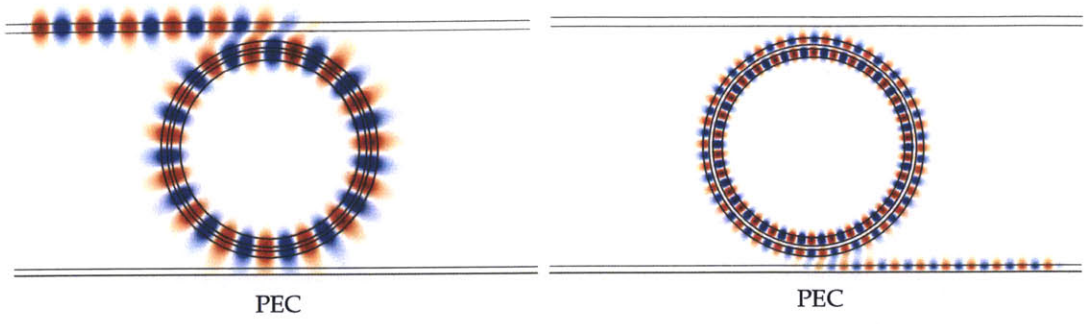

Figure 6-7: E_z field snapshot of two double-ring (Fig. 6-4) resonator modes propagating counter-clockwise, with frequencies $\omega_1 = 0.277(2\pi c/a)$ (left) and $\omega_2 = 2\omega_1$ (right) and azimuthal momentum $m_1 = 15$ and $m_2 = 30$ (effective $k_1 = 0.39(2\pi/a)$ and $k_2 = 2k_1$). The ring resonator is side-coupled to two adjacent waveguides, separated by a distance $d_1 = d_2 = 0.5a$, supporting phase-matched propagating modes at ω_1 (top waveguide) and ω_2 (bottom waveguide).

because of radiations from the finite Q_{rad}/Q_w). However, as noted in Sec. 6.3, CMT makes many approximations with respect to the full Maxwell equations, and while each of these approximations seems justified in the present case, it is desirable to validate the CMT predictions against a full nonlinear simulation as described in Sec. 6.3.

In Fig. 6-8(left), we plot the SHG conversion efficiency (output to the lower waveguide) versus the input power at ω_1, incident from the upper-left port, as computed by both CMT and by nonlinear FDTD (run long enough to reach steady state from zero initial fields). The nonlinear FDTD results agree well with the CMT: the FDTD efficiency peaks at the predicted critical power with a maximum efficiency of 78%, which is reasonable agreement especially considering that it is difficult to determine the resonator Q value from the transmission fits with more than a few percent accuracy. A snapshot of the nonlinear FDTD simulation is shown in Fig. 6-8(lower-right), in which both the ω_1 input and the ω_2 output are visible (with a complicated superposition of the two modes in the cavity). Most of the 22% of unconverted power is lost to radiation (visible in the plot) due to the finite Q_{rad}/Q_w, but the imperfections represented by these losses also give rise to a few percent of the ω_1 power escaping into the upper-right port and $\approx 2\%$ of the ω_2

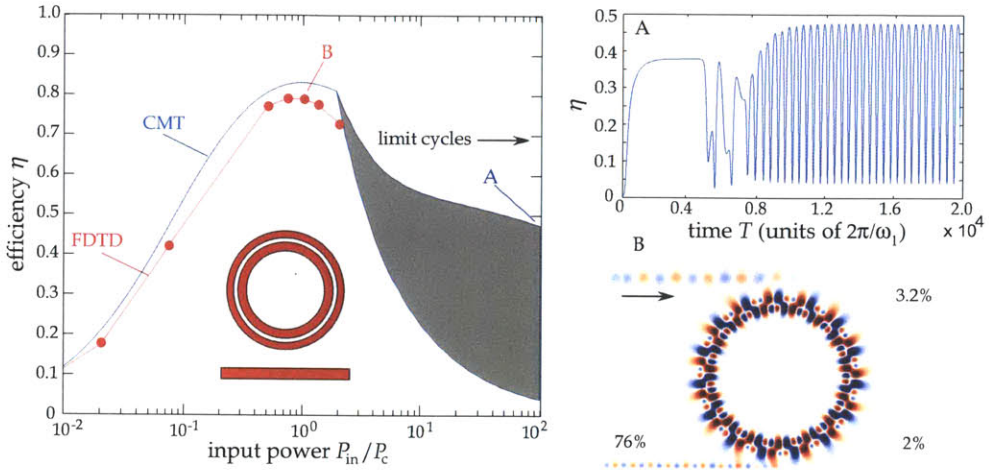

Figure 6-8: (Left:) Plot of SHG efficiency $\eta = P_{SH}/P_{in}$ versus P_{in}, for the double-ring resonator system of Fig. 6-6 with waveguide-separations $d_1 = 1.05a$ and $d_2 = 0.7a$, obtained both via FDTD simulations (red circles) and CMT (blue line). The gray region denotes the presence of instabilities that lead to limit-cycle behavior. (Right:) An example of limit cycle at point B and the manifestation of nonlinear conversion processing at point A (the efficiency peak) shown in (Left).

power escaping to the lower-left port.

Another intriguing prediction of CMT for indra-cavity SHG is that, once the input power is significantly larger than the critical power, the steady-state solution is replaced by a "self-pulsing" solution in which a constant input power at ω_1 produces an oscillating output power at ω_2 [45]. These "limit cycles" occur in the shaded region of Fig. 6-8(left), and because the system never reaches a steady state simply taking the efficiency at the end of the simulation gives a somewhat noisy value as seen in this plot from the CMT data. A plot of efficiency vs. time, from the CMT for $P_{\text{in}} = 2P_c$, is shown in Fig. 6-8(upper-right), and the limit cycles are clearly visible (after an initial transient).

When it comes to realistic numbers, we are most interested in the 3d results of the next section, but it is still interesting to estimate the critical power for the 2d system in experimental units, assuming a vacuum wavelength $\lambda_1 = 1.55\,\mu$m for ω_1 (and hence $a = 0.277 \times 1.55 = 0.43\,\mu$m). To obtain a power in W from a 2d calculation, we must assume a certain finite thickness in the z direction, and in this case we assume 200 nm (good vertical confinement), and obtain a critical input

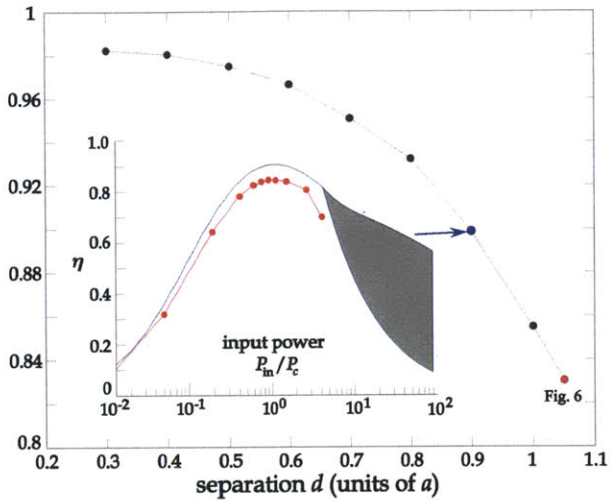

Figure 6-9: Maximum efficiency vs. separation between input waveguide and ring resonator. (inset: Conversion efficiency from CMT and FDTD in the case $d_1 = 0.9a$)

power of 800 mW for $LiNbO_3$ at these Q values.

The $< 100\%$ efficiency was due to the finite Q_{rad}/Q_w and in particular the limiting factor is the loss at ω_1 ($Q_{rad,1}/Q_{w,1} = 9847/1998 = 4.9$), so we can obtain higher efficiency by making $Q_{w,1}$ smaller. This is accomplished by bringing the input waveguide closer to the ring resonator. Substituting the Q values from Fig. 6-7 into CMT, we predict an increase in efficiency with decreasing d_1 as shown by the black curve in Fig. 6-9. The inset of Fig. 6-9 shows the results in both CMT and FDTD for a separation $d_1 = 0.9a$ (for which $Q_{rad}/Q_w = 9616/1043 = 9.2$) and obtain 90% and 85% efficiency from CMT and FDTD, respectively.

6.4 3D Design

We now apply the same basic principles that we validated in the 2d example to a more realistic 3d design. In 3d, we will consider SHG from $\lambda_1 = 1.55\,\mu$m in a GaAs film bonded to an SiO_2 substrate. As in 2d, we will design a ring resonator to have resonances at ω_1 and $\omega_2 = 2\omega_1$, which is side-coupled to two waveguides for input and output, as shown schematically in Fig. 6-10. Also as in 2d, we will design the output waveguide to have a low-frequency cutoff $> \omega_1$ so that we can independently

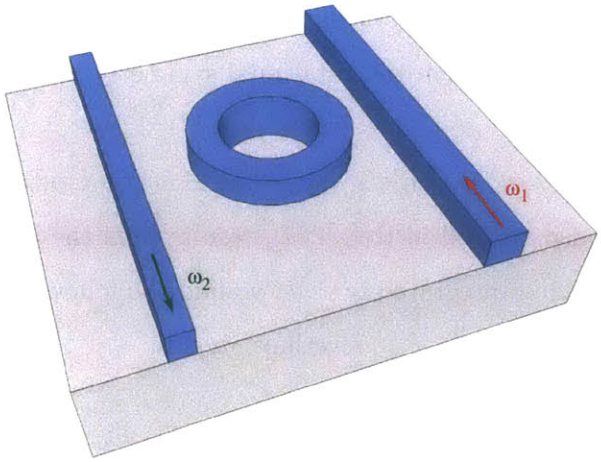

Figure 6-10: Schematic diagram of 3d ring-resonator waveguide-cavity system

control the coupling Q_w for the two frequencies via the respective ring–waveguide gaps. However, there are also several key differences from 2d. First, we include realistic material dispersion of the GaAs: its index is 3.374 for ω_1 and 3.692 for ω_2 [148], and these large indices relative to our 2d structure allow a smaller-radius ring for the same Q. GaAs has a large *off-diagonal* $\chi^{(2)}$ coefficient ≈ 180 pm/V [18].

For purely off-diagonal components $\chi^{(2)}_{xyz}$, the expression for β_1 in Eq. (6.2) is slightly more complicated than those involving diagonal $\chi^{(2)}$ components (coupling modes with the same polarization). In this off-diagonal case, it is instructive to explicitly write down the coupling coefficients. Moreover, since our geometry has cylindrical symmetry, it is useful to write down Eq. (6.2) in cylindrical coordinates. Specifically, if we write the overlap integral in the numerator $\sim 4(E_{1x}E_{1y}E_{2z} + E_{1x}E_{1z}E_{2y} + E_{1y}E_{1z}E_{2c})$ in polar coordinates, using the coordinate transformations $E_x = E_r\cos(\theta) - E_\theta\sin(\theta)$ and $E_y = E_r\sin(\theta) + E_\theta\cos(\theta)$ we find coupling terms of the form:

$$E_{1x}E_{1y} = (E_{1r}^2 - E_{1\theta}^2)\sin(2\theta)/2 + E_{1r}E_{1\theta}\cos(2\theta) \tag{6.3}$$

$$\begin{aligned} E_{1x}E_{2y} + E_{2x}E_{1y} &= 2(E_{1r}E_{2r} - E_{1\theta}E_{2\theta})\sin(\theta)\cos(\theta) \\ &\quad + (E_{1r}E_{2\theta} + E_{2r}E_{1\theta})[\cos^2(\theta) - \sin^2(\theta)]. \end{aligned} \tag{6.4}$$

If we also assume that the fields that are rotationally symmetric functions multiplied by $e^{im_{1,2}\theta}$ (from the rotational symmetry of the ring), and we write $\cos(\theta) = (e^{i\theta} + e^{-i\theta})/2$ and $\sin(\theta) = (e^{i\theta} - e^{-i\theta})/2i$, we obtain an overlap function whose integral in θ is zero unless $m_2 = 2m_1 \pm 2$; this modification to the previous selection rule, also derived in Ref. [207], results from the additional $e^{\pm i2\theta}$ terms in the integral. Having integrated over θ, we find the following simplified cylindrical overlap integral:

$$\beta_1 = \frac{1}{2}\frac{\int r\mathrm{d}r\mathrm{d}z\,\epsilon_0\chi^{(2)}_{xyz}\left[E_{1r}E_{1\theta}E_{2z} + (E_{1r}E_{2\theta} + E_{2r}E_{1\theta})E_{1z}\right]}{\left(\int \mathrm{d}^3\mathbf{x}\,\epsilon|E_1|^2\right)\left(\int \mathrm{d}^3\mathbf{x}\,\epsilon|E_2|^2\right)^{1/2}} \pm \frac{i}{4}\frac{\int r\mathrm{d}r\mathrm{d}z\,\epsilon_0\chi^{(2)}_{xyz}\left[2(E_{1r}E_{2r} - E_{1\theta}E_{2\theta})E_{1z} + (E_{1r}^2 - E_{1\theta}^2)E_{2z}\right]}{\left(\mathrm{d}^3\mathbf{x}\,\epsilon|E_1|^2\right)\left(\int \mathrm{d}^3\mathbf{x}\,\epsilon|E_2|^2\right)^{1/2}}, \quad (6.5)$$

While the modes in 3d are not purely polarized, because of the near mirror symmetries (both laterally and vertically) for strongly confined modes, in the center of the waveguide they are mostly TE-like (E in-plane) or TM-like (E out-of-plane), and therefore it can be convenient to describe the modes as TE-like or TM-like. (Even for purely symmetric waveguides the modes are only purely polarized in the mid-planes, outside of which they have other components [84]). On the other hand, this terminology can be misleading, because, for example, TM-like modes often have significant in-plane components and hence there can be significant coupling between two TM-like modes. More explicitly, the overlap integral in the numerator of Eq. (6.5) has nonzero $\sim E_{1\theta}^2 E_{2z}$ terms that couple two TM-like modes (the symmetry allows for a large overlap if E_{2z} is even because $E_{1\theta}$ is squared). (As mentioned in the concluding remarks, preliminary work suggests that the overlap can be significantly improved by optimizing over a wider parameter space to consider additional modes, and at the same time one should obtain modes with higher radiative Q_{rad} to reduce losses.) The ability to couple modes having the same polarization is a dramatic departure from 2d and belies some of the conventional wisdom on this subject.

As a consequence of all of these changes, it turns out that a single ring (no air

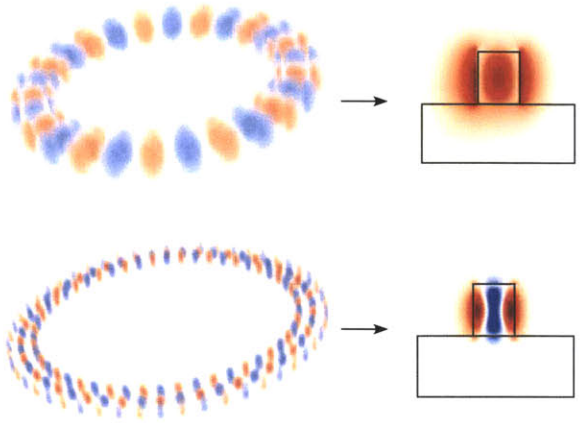

Figure 6-11: Field distribution (left) and corresponding lateral cross-section (right) for the ω_1 (top) and ω_2 (bottom) modes.

groove) is sufficient to obtain the desired two modes with both matched frequencies and excellent mode overlap (large β_1). Here, we chose a conventional TE-like to TM-like design that we found by searching through a small space of parameters. We chose a square $0.8a \times a$ cross-section of the ring waveguide, with an inner radius of R. The key factor in the ring design is the modified selection rule from the overlap integral of Eq. (6.5). For several choices of m_1, we varied the ring radius R and looked for $\omega_2 = 2\omega_1$ pairs of modes at each of the three possible m_2 values. We found a suitable pair of modes for $R = 5.07a$, $m_1 = 16$, and $m_2 = 34$, for which $\omega_1 = 0.244 \cdot 2\pi c/a$, corresponding to $a = 0.244 \times 1.55\,\mu\text{m} = 378$ nm. The resulting field patterns are shown in Fig. 6-11. As mentioned above, in 3d we can use the substrate itself to induce the required low-frequency cutoff rather than an unphysical perfect metal. We now describe these differences and the resulting 3d design in more detail.

The fields in Fig. 6-11 are especially attractive because they satisfy a second, *approximate*, selection rule. Because the waveguide width is small compared to the ring radius, the modes closely resemble those of a straight waveguide with the same cross-section, and this cross-section has a mirror symmetry plane bisecting the waveguide perpendicular to the substrate. In a straight waveguide, therefore, all

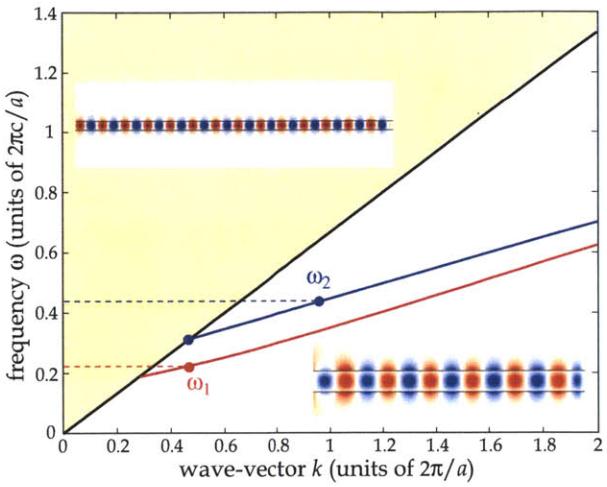

Figure 6-12: Left: Band diagram of the 3d waveguides, Right: Cross section of the field distribution (upper: Ez; lower: Ey)

modes would have fields patterns that are either symmetric or anti-symmetric (even or odd) with respect to this mirror. In a ring, the curvature breaks the symmetry but the modes are still *nearly* even or odd with respect to this midplane. In the coupling integral, even or odd ω_1 modes are squared to become even, and hence can only couple to even-symmetry ω_2 modes. Therefore, we should consider only ω_2 modes that have nearly even symmetry such as the one in Fig. 6-11, as nearly odd ω_2 modes would have nearly zero β_1. As a figure of merit, we can compare the overlap integral β_1 to a "perfect" overlap integral β_0 in which we assume that the fields $= 1$ inside the ring and $= 0$ outside, and we find that $|\beta_1/\beta_0| \approx 0.06$ for our modes, indicating reasonably good overlap. (Better overlaps should be easily achieved.)

For input coupling, we use an $0.8a \times a$ waveguide identical to the ring cross-section. For output coupling, we reduce the width to $0.15a = 57\,\text{nm}$ so that the waveguide has a low-frequency cutoff $> \omega_1$. The corresponding dispersion relations are shown in Fig. 6-12. (Note that the substrate causes a cutoff in both waveguides, but the cutoff is $< \omega_1$ for the input waveguide.)

Since we use separate input/output waveguides, it should be possible to independently control their coupling Q's by varying the waveguide–ring separations. The radiation Q values for the isolated ring are $Q_{rad,1} = 7129$ and

$Q_{rad,2} \approx 1.136 \times 10^7$. So, if we choose $Q_{w,1} = 1000$ and $Q_{w,2} = 2000$, CMT predicts a peak of 88% efficiency at a critical power of 18 mW. As mentioned above, significantly lower critical powers can be easily achieved with further optimization of the structure (larger overlap integrals), and indeed preliminary work suggests the possibility of almost an order of magnitude improvement.

6.5 Summary and remarks on the results

In this chapter, we demonstrated that a simple two-port, two-mode CMT can accurately capture the complexities involved in a full nonlinear Maxwell system involving rings, losses, and multiple output ports. We have also presented proof-of-concept designs for ring-resonator intra-cavity SHG at high efficiencies and low powers, illustrating the care that is required to obtain appropriate modes and symmetries and in the design of the input/output coupling. However, these designs could be altered in many ways depending upon the needs of a particular experiment. For example, lower powers could be achieved by going to larger Q_w values, at the expense of bandwidth and sensitivity, while increasing the ring radius R in order to prevent radiation loss from increasing. Conversion from 10.6 μm to 5.3 μm is especially attractive, both because of the paucity of sources at 5 μm and also because a 10$\times$ increase in length scales should simplify fabrication. As can be seen from the comparison of the 2d and 3d designs, the specific parameters of the ring design depend very strongly on the materials (refractive indices and dispersion), the form of the $\chi^{(2)}$ susceptibility, and on the details of the vertical confinement. So, the the specific parameter choices here are far from universal, but the basic design principles, especially the selection rules, the role of the different Q values, and the advantage of cutoffs for separate input/output coupling, will remain. A key practical concern in any intra-cavity frequency-conversion design such as this one is the sensitivity to fabrication imperfections, which will slightly shift both ω_1 and ω_2. Any overall shift in the frequencies can be compensated by a tunable laser source for the input, in order to match ω_1. However, another tuning parameter is

required if imperfections spoil the $\omega_2 = 2\omega_1$ condition. Fortunately, we found in our designs that varying a single parameter, in our case the radius R, was sufficient to bring the frequencies into alignment. Although R cannot easily be changed post-fabrication, other dynamically tunable parameters should play a similar role. For example, strain-induced deformation of the cavity or strain-induced birefringence should affect the ω_1 and ω_2 modes differently and hence be capable of correcting small errors in $\omega_2 - 2\omega_1$.

Chapter 7

Review of Transformation-Based Invisibility Cloaking

For the rest of the thesis, we focus on transformation-based cloaking and its limitations. We will argue that the problem of cloaking becomes intrinsically more difficult as the size of the object to be cloaked increases compared to the wavelength, and is ultimately limited by fundamental considerations involving the delay–bandwidth and delay–loss products, even for ground-plane cloaks [113, 120, 217] where bandwidth is not limited by causality constraints. To begin with, in this chapter, we will first review the basic ideas underlying transformation-based cloaking, as proposed by [113, 155].

Transformation-based cloaking uses the coordinate-invariance property of Maxwell's equations we discussed in chapter 2, to make objects appear invisible. As discussed in chapter 2, a coordinate transformation of space $x \rightarrow x'$ is equivalent to a material change/transformation $\varepsilon \rightarrow \varepsilon'$. Therefore, by an appropriate choice of cloak material ϵ and μ, one can transform the physical space to a virtual space in which the part of space where the object is does not exist. A simple example is shown in Fig. 7-1

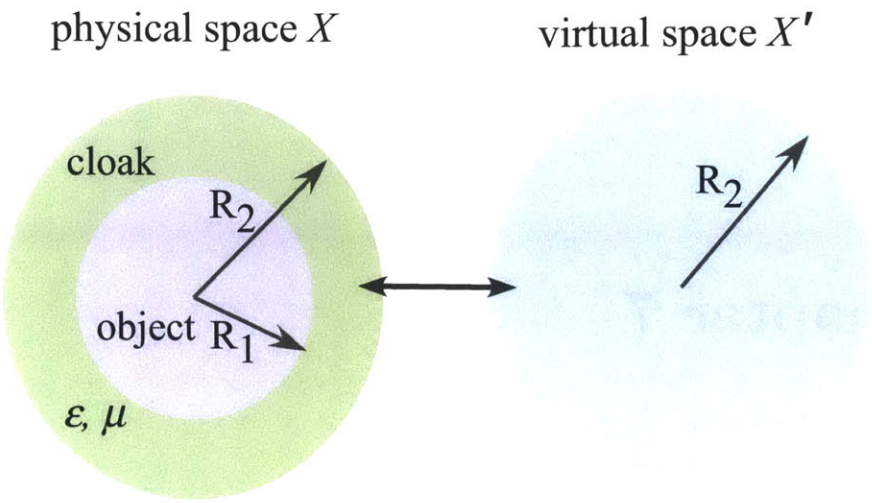

Figure 7-1: A demonstration of how a transformation-based cloak transforms the space of a spherical object and cloak to empty space.

where a spherical object is cloaked using a simple linear transformation:

$$r = R_1 + r' \frac{R_2 - R_1}{R_2}, \tag{7.1}$$

$$\theta = \theta', \tag{7.2}$$

$$\phi = \phi' \tag{7.3}$$

which gives result to a cloak configuration of:

$$\varepsilon_r = \mu_r = \frac{R_2}{R_2 - R_1} \frac{(r - R_1)^2}{r^2} \tag{7.4}$$

$$\varepsilon_\theta = \mu_\theta = \frac{R_2}{R_2 - R_1} \tag{7.5}$$

$$\varepsilon_\phi = \mu_\theta = \frac{R_2}{R_2 - R_1} \tag{7.6}$$

Mathematically, the key point is that the transformation is the identity for $r > R_2$, so the solutions outside the cloak are *identical to those in virtual space* where there is no object and hence no scattering. Physically, when an incident wave hits the cloak, the cloak materials cause the wave to bend around the object, with no reflections at the cloak interface thanks to perfect "impedance matching," and exit

the cloak exactly on their original trajectory, as depicted in Fig. 1-3. After a little thought about these ray trajectories in the cloak, one can immediately see why perfect cloaking in air/vacuum is impossible over nonzero bandwidth, as pointed out by Pendry [155]. Because rays traveling around the object propagate over a longer distance than they would in the absence of the object, but must reach the opposite side in the same amount of time, they would have to have velocity $> c$. This is possible at a single frequency because then only the phase velocity must exceed c, but is impossible over a nonzero bandwidth since the group velocity is bounded above by c. This can also interpreted as a causality constraint [135]. In chapter 10, we extend these ideas to prove a causality-induced limitation on the bandwidth even for imperfect cloaking (and incidentally give an alternative proof that perfect cloaking is impossible over a nonzero bandwidth even for idealized materials).

Another important point to pay attention to is the singularity of the transformation; the coordinate transformation is singular at the origin in X', corresponding to the inner surface of the cloak in X. This implies extreme material properties ε and μ as the inner surface of the cloak is approached. Therefore, analysis of such a cloak requires additional care at the singular point [213]. The same is true for numerical methods applied to such a cloaking system. In practice also, although having ε and μ values of exactly zero is unlikely—not only will there generally be absorption loss that gives a nonzero imaginary part, but manufacturing imperfections most likely mean that the real parts of ε and μ will not pass through zero simultaneously. As we show in chapter 10, if one sets a nonzero lower bound on the attainable refractive index then the attainable reduction in scattering cross-section is bounded even at a single frequency.

7.1 Ground-plane cloaking

It was subsequently proposed by Li and Pendry [113] that such bandwidth limitations and extreme material requirements are removed for a ground-plane

cloak, in which an object is hidden by a coordinate transformation mapping it into a ground plane or substrate as illustrated in Fig. 7-2. This type of cloak does not require a singular transformation and also has a smaller range of ε and μ in the cloak and so appears more realistic. In addition, causality constraints do not apply to ground-plane cloaks, because the reflected wave travels a shorter distance in the presence of the cloak and hence does not need a speed $> c$ to simulate absence of the object. Although this design makes cloaking easier in both theory and practice, we argue in the next two chapters that even ground-plane cloaking is subject to delay–bandwidth/size and delay–loss limitations that become more stringent as the size of the cloaked object increases.

The most challenging case to consider is that of a perfectly reflecting "ground," since in that case the cloak must not absorb any incident light, but must instead reflect it with a time delay that mimics the reflection in the absence of the object. The opposite extreme is a perfectly absorbing "black" ground, in which case cloaking is trivially attained by a black (perfectly absorbing) cloak. (In the intermediate case of a partially absorbing ground, the coordinate transformation would map the object deep enough into the "ground" in virtual space so as to make it invisible in virtual space, and hence in physical space.)

Consider an object in a volume V_o on a reflective ground plane, surrounded by a homogeneous isotropic ambient medium with permittivity ε_a and permeability μ_a. This object is cloaked by choosing the materials ε and μ in a surrounding volume V_c to mimic a coordinate transformation, with Jacobian $\mathcal{J}$, mapping the physical space X to a virtual space X' in which the object is absent ($\mathcal{J}_{ij} = \partial x_i / \partial x'_j$), as shown in Fig. 7-2. This is achieved by $\varepsilon = \varepsilon_a \mathcal{J}\mathcal{J}^T / \det \mathcal{J}$ and $\mu = \mu_a \mathcal{J}\mathcal{J}^T / \det \mathcal{J}$ (for isotropic ε_a, μ_a) [155]. (The surface of the object in X is mapped to the ground plane in X', and so the inner surface of the cloak must be reflective like the ground plane.)

The lack of singularity in this coordinate transformation makes realization of the corresponding material properties by metamaterials considerably easier than for isolated-object cloaks. An especially favorable situation arises for a thick cloak in

which a gradual, smoothly varying coordinate transformation is chosen, so that ε and μ are slowly varying in X (and are never large compared to the ambient medium). In this case, Li and Pendry pointed out [113] that for a sufficiently slowly varying μ, the permeability μ can be approximately commuted with $\nabla \times$ in Maxwell's equations so as to combine it with ε. In this way, the challenging problem of attaining a effective magnetic response from non-magnetic constituent materials is avoided, and the slowly varying ε can be achieved simply via a nanoporous material with a gradually varying air-fill fraction [113]. Most of the attempted experimental realizations of ground-plane cloaking have followed this fabrication strategy. Another fabrication strategy recently demonstrated was to cloak a specially chosen ambient *fluid* for which an index-matched naturally occurring birefringent crystalline material exists that closely approximates the desired cloak ε (albeit for a 2d cloak)—in that way, manufacturing imperfections associated with synthetic metamaterials are avoided and very low-loss cloaks are attained in order to cloak an object $\sim 1000\times$ the wavelength (mm-scale) [32, 215].

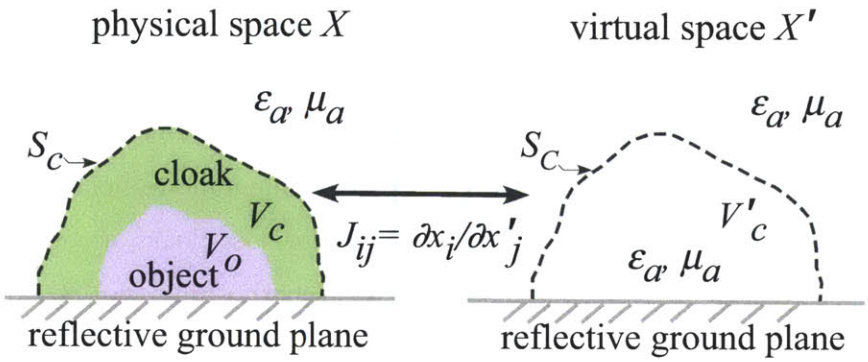

Figure 7-2: Schematic of a general cloaking problem: an object in a volume V_o sitting on a reflective ground is cloaked by choosing the materials ε and μ in a surrounding volume V_c to mimic a coordinate transformation, with Jacobian $\mathcal{J}$, mapping the physical space X to a virtual space X' in which the object is mapped into the ground and V_c is mapped into the entire $V_c' = V_c \cup V_o$ volume with the homogeneous ambient-space properties ε_a and μ_a. S_c denotes the outer surface of the cloak (identical in X and X').

Chapter 8

Introduction to limitations of cloaking: a 1D model problem

In order to understand the limitations of ground-plane cloaking, we first consider the simplest possible circumstance: a 1d cloak to hide an object of thickness h on top of a substrate (e.g. a conducting plane) in vacuum. This problem is conceptually much easier than general cloaking, in that only a single incident (and reflected) wave need be considered. In contrast, even two-dimensional cloaking is far more complex: not only would the object need to be cloaked from incident waves at all angles, but for incident waves parallel to the ground the cloaking problem becomes more similar to that of cloaking an isolated object—with the associated causality constraints—as the height of the object increases. Since 1d cloaking appears to be so much easier, any fundamental limitations that arise in this case should intuitively apply even more strongly in 2d and 3d, and we show explicitly that this is the case in chapter 9. This work was published in Ref. 72.

In this idealized 1d case, the cloak consists of some arbitrary materials in a region of thickness d on top of the object, as depicted in Fig. 8-1. We assume that the ground plane reflects light with negligible loss in the bandwidth of interest (in the trivial case of a black ground plane, one would merely need a black cloak). The function of the cloak is now simple: the cloak must reflect incident waves with a delay equal to the time (and phase) delay $\tau_0 \geq 2(h+d)/c$ that the reflected wave would incur in

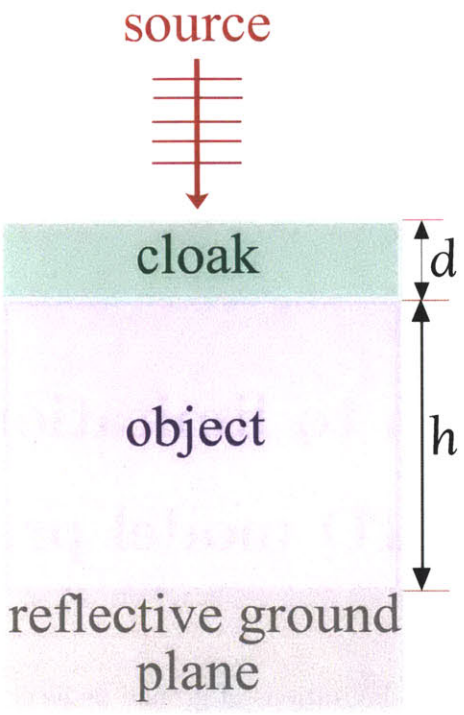

Figure 8-1: A 1d ground-plane cloak

the absence of the cloaked object. A similar delay must also be achieved in 2d/3d cloaking for beams at any angle—the cloak must simulate the delay from bouncing "through" the object off the ground plane, and in fact the required delay increases for more oblique incidence (longer paths through the object). To be more precise, suppose that the reflected wave from the bare ground plane, at a height $h + d$, has a phase $\phi(\omega) \approx \phi(\omega_0) + \phi'(\omega_0)(\omega - \omega_0)$ near some frequency ω_0, where the derivative $\phi'(\omega_0) = \tau_0$ is the time delay [108]. There are two cases. First, the *phase-delay case*: if the bandwidth is narrow, so that $\phi'(\omega_0)(\omega - \omega_0)$ can be neglected, then the cloak merely need achieve the correct phase $\phi(\omega_0)$, but this imposes a bandwidth constraint: the delay–bandwidth product $\tau_0 \Delta\omega$ must be small. (This corresponds to incident pulses of duration $\sim 1/\Delta\omega \gg \tau_0$.) Second, the *time-delay case*: if the delay–bandwidth product is not small, then $\phi'(\omega_0)(\omega - \omega_0)$ cannot be neglected and the cloak must achieve a true time delay τ_0 (an ω-dependent phase). This raises two additional difficulties. First, it is well known that the achievable delay–bandwidth

product in finite-size passive linear systems is limited [108, 136, 196]. Second, a long dwell time in the cloak means that loss in the cloak must be small. We deal with each of these requirements below.

8.1 Delay–bandwidth limitation

The achievable time delay τ_0 in a passive linear system (unlike time-varying active devices [208]) is limited: for a given bandwidth $\Delta\omega$ and diameter d of the delay region, the maximum delay is proportional to $d/\Delta\omega$. The scaling of delay with bandwidth is known as the delay–bandwidth product limitation [108], and in the case of a single resonant filter the upper bound on $\tau_0\Delta\omega$ is of order unity as a consequence of the Fourier uncertainty relation [108, 158]. To obtain a bandwidth much larger than $1/\tau_0$, one can chain multiple filters into a slow-light delay line, or even forgo slow light and use propagation through a long region—in any case, the maximum delay is proportional to the diameter of the region. A more careful analysis for slow-light delay lines yields a delay–bandwidth limit of $\tau_0\Delta\omega/\omega \lesssim (n-1)2d/c$, where n is the effective index in the delay region [196], and a more optimistic bound of $n(n-1)2d/c$ was derived under more general assumptions [136]. As a consequence of this and $\tau_0 > 2h/c$, the cloak thickness d must grow proportional to h:

$$d \gtrsim \frac{h}{n(n-1)} \frac{\Delta\omega}{\omega}. \tag{8.1}$$

[This is probably optimistic in the wide-bandwidth regime where slow light is difficult to utilize; for a non-slow cloak of thickness d, where the time delay is simply $2dn/c$ and must be $> 2(h+d)/c$, one obtains a minimum thickness $d > h/(n-1)$.] One can relax this tradeoff if a larger n can be obtained, but large indices of refraction (arising from resonances) are associated with narrow bandwidths and/or large losses [79].

8.2 Delay–loss limitation

In the time-delay regime, a larger object for a given bandwidth means that the incident wave needs to spend more time in the cloak, which will tend to increase losses due to absorption and imperfections. The loss per time γ is proportional to $\gamma \sim \omega \operatorname{Im} n / \operatorname{Re} n$ for light confined mostly in a given index n [84]. To maintain effective invisibility, the loss incurred in the cloak must be small: one must have $\tau_0 \gamma \ll 1$ for negligible absorption. But, since $\tau_0 > 2(h+d)/c$, this implies the following limitation on the loss tangent:

$$\frac{\operatorname{Im} n}{\operatorname{Re} n} \ll \frac{1}{4\pi} \frac{\lambda}{h+d}. \tag{8.2}$$

That is, less and less loss can be tolerated for larger objects relative to the vacuum wavelength λ.

In the *phase-delay* regime, the dwell time inside the cloak can be independent of h, in which case the loss tolerance does *not* decrease as h/λ increases, at the expense of greatly reduced bandwidth.

8.3 Interface reflections

A low-loss cloak achieving the requisite time delay is useless if there is substantial reflection off the surface of the cloak itself. In 1d, eliminating reflections reduces to the problem of impedance-matching the cloak with vacuum [188]. In the transformation-optics approach, impedance-matching is attained automatically: the 1d transformation results in a cloak material that has a both a permittivity ε and a permeability μ (for polarizations transverse to the surface normal), such that the impedance $\sqrt{\mu/\varepsilon}$ exactly equals that of vacuum [200]. Alternatively, if the material is varying slowly enough, this μ can be approximately commuted with the curls in Maxwell's equations to combine it with ε into an index $n = \sqrt{\varepsilon\mu/\varepsilon_0\mu_0}$. This is equivalent to an anti-reflective (AR) coating formed by a slowly varying n (in the "adiabatic" limit of slow variation the reflection vanishes [87]).

This means that a homogeneous medium cannot be used for the cloak. To obtain a μ, metamaterials employing subwavelength metallic resonances are typically used [154, 175], whereas a continuously varying n is typically achieved with a microporous structure whose porosity is gradually varied [50, 60, 101, 199]. In either case, the loss limit in the previous section must then include fabrication disorder and surface roughness in addition to absorption.

8.4 Examples and results

Let us take some real-world examples of cloaking applications and study what practical limitations one would face even for an idealized 1d ground-plane cloak. For microwave frequencies, consider cloaking a vehicle of height $\approx 2\,\mathrm{m}$ from a radar of wavelength $\approx 1.25\,\mathrm{cm}$ (24 GHz). Using a time-delay cloak of thickness 10 cm over a 10% bandwidth, Eqs. (8.1) and (8.2) imply an effective index of $\gtrsim 1.4$ and a loss tangent of $\ll 4.7 \times 10^{-4}$. (Operating in the phase-delay regime would imply a fractional bandwidth of $< 10^{-4}$.) [Although one might expect a cloak of thickness $h/20$ to need $n = 20$ for the requisite delay, Eq. (8.1) assumes that slow-light/resonances are used to exploit the narrow $\Delta\omega$.] To time-delay cloak the same object at visible frequencies with a 10 cm cloak, aiming for 10% bandwidth around 575 nm, we would again need $n \gtrsim 1.4$, but with a loss tangent $\ll 2 \times 10^{-8}$. (In the phase-delay regime, the bandwidth would be only 0.013 pm.)

Although such low losses may seem attainable, e.g. with oxides, even in 1d a microstructured medium is required for impedance-matching as described above, and in 2d and 3d even more complicated metamaterials seem necessary [50, 60, 101, 113, 120, 199] (anisotropy requirements can be minimized via quasiconformal transformations [113], although discarding anisotropy incurs a lateral shift in reflected beams [214]). For a cloak at ≈ 10 GHz, an experimental absorption loss tangent $\approx 10^{-3}$ was obtained [175] for a wavelength-scale object; this is already too lossy for 1d cloaking a meter-scale object, from above. A ground-plane cloak can use non-resonant micro-structures that may be lower-loss [113, 120], but for an

object that stands many wavelengths above the ground, the problem of cloaking against oblique waves seems to approach isolated-object cloaking. Nevertheless, we cannot say that the loss bounds from the 1d cloaking problem are definitely unattainable for cloaking meter-scale objects at microwave frequencies, although it appears challenging. On the other hand, loss tangents $\ll 10^{-8}$ seem impossibly small for any metamaterial with metallic constituents at infrared or visible frequencies. Even if ground-plane cloaking permits the use of purely dielectric constituents, such a loss tangent appears almost unattainable when scattering from fabrication disorder and nonzero gradients (non-adiabaticity) is included, since the requisite gradients (especially for cloaks not too much bigger than the cloaked object) seem to imply constituent materials with large index contrasts (oxide/air or larger) [50, 60, 199]. For comparison, a waveguide with a loss tangent of 10^{-9} at $1\,\mu$m wavelength would correspond to decay lengths of ~ 1 km—orders of magnitude better than the cm-scale decay lengths typically achieved at infrared frequencies in geometries (such as strip waveguides) with wavelength-scale geometric components, and a metamaterial requires components much smaller than the wavelength (thus many more surfaces). Visible-wavelength cloaking, therefore, seems restricted to cloaking objects that are many orders of magnitude smaller than meter scales.

These calculations demonstrate the difficulty of cloaking objects much larger than the wavelength when the ambient medium is air/vacuum. The problem may become easier if the ambient medium itself is lossy, such as for an object immersed in water or inside a lossy waveguide. In that case, the loss of the cloak need only be comparable to that of the surrounding medium. The delay–bandwidth constraint remains, however: the cloak thickness must grow proportional to that of the object being cloaked, for a fixed bandwidth. However, if the velocity of light in the ambient medium is $< c$, the causality constraint on wide-bandwidth cloaking of isolated objects [135, 155] is lifted. A possible example can be found in Ref. 185, which cloaked an "object" (a place where two surfaces touched) roughly 100 wavelengths in diameter (this "diameter" was indirectly measured and may not be comparable to the diameters of objects used in other cloaking problems), but did so

in a waveguide between two metallic surfaces. Such a waveguide has a group velocity $< c$, eliminating the causality constraint, and may also have non-negligible absorption loss. In addition, the cloaking region in Ref. 185 was achieved by curving the surfaces smoothly, which allows a smooth variation of the effective index without microstructured media—it seems plausible that such a cloak has at least 100 times lower absorption/scattering loss than was present in metamaterial cloaks. Although we presented the basic delay–loss and delay–bandwidth/size tradeoffs in the context of an idealized 1d cloaking problem, we will show in next sections that similar tradeoffs must apply even more strongly to the more difficult problem of cloaking objects in 2d and 3d, especially without a ground plane. In fact, recent numerical experiments have shown that similar tradeoffs (loss tolerance scaling inversely with diameter) arise for three-dimensional cloaking of a perfectly conducting sphere of diameter h [212]. In Fig. 8-2, we show the loss threshold vs. h for this 3d cloak when the (single-ω) reduction in scattering cross section is fixed to $1/100$, for cloak thickness $d = h/12$. The scattering is calculated with a transfer-matrix method in a spherical-wave basis [163, 212]. Not only does the maximum loss scale exactly inversely with the diameter, but the constant factor in this relationship is consistent with the requirement Eq. (8.2) that the loss be much smaller than $\lambda/2\pi(h + d)$ (shaded region) derived from the much simpler 1d model (using a path length h instead of $2h$ since there is no ground plane to double the optical path).

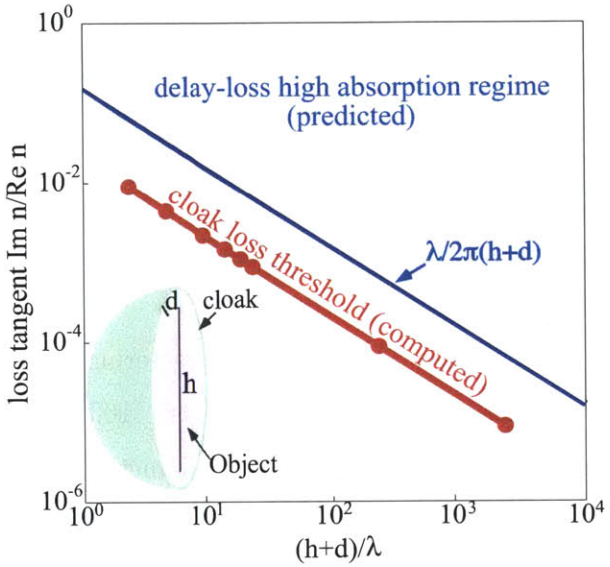

Figure 8-2: Maximum cloak loss tangent versus diameter h for cloaking a perfectly conducting sphere, for cloak of thickness $d = h/12$. Shaded area is the regime of high absorption predicted by the simple 1d model of Eq. (8.2). The red curve, data from Ref. 212, shows the maximum loss tangent to obtain 99% reduction in the scattering cross section using a Pendry-type cloak.

Chapter 9

Limitations of Ground-Plane Cloaking: 3D Case

The previous chapter proposed a general scaling limitation for ground-plane cloaking—cloaking gets harder as the object gets bigger—derived in a simple one-dimensional model, independent of how the cloak is designed. In this chapter, we derive the same result in three dimensions (and of course an analogous proof applies to two dimensions as well), assuming that the cloak is designed by the principles of transformation optics. This work was published in Ref. 70.

In particular, we derive the limitations of cloaking under the following two practical requirements:

- The attainable refractive index contrast $\sqrt{\varepsilon\mu}/\sqrt{\varepsilon_a\mu_a}$ (the eigenvalues of $\mathcal{J}\mathcal{J}^T/\det\mathcal{J}$) is bounded above by B and below by b.
- The scattering cross-section (nonzero due to imperfections in the cloak) for any incident wave is bounded above by some fraction f of the geometric cross-section s_g.

Given these two constraints, we derive the following relations between the difficulty of cloaking and the size of the object to be cloaked:

- The thickness of the cloak must scale with the object thickness (divided by B).

- The allowed imperfections (e.g. disorder or absorption) must scale at most inversely with the object thickness.

9.1 Cloak thickness

The volume V_c is given by $\int_{V_c'} |\det \mathcal{J}| dx'dy'dz'$, and therefore constraints on V_c/V_c' immediately follow from two facts. First, $|\det \mathcal{J}|$ can be bounded due to the bound B on the index change above. Second, an even tighter bound follows from the fact that the outer surface S_c of the cloak is invariant under the coordinate transformation. In particular, defining $\mathcal{J}_x$ and $\mathcal{J}_y$ as the first two columns of $\mathcal{J}$ and denoting singular values of $\mathcal{J}$ by σ_i (σ_i^2 are eigenvalues of $\mathcal{J}\mathcal{J}^T$ [195]), referring to the cross sections $A'(z')$ defined in Fig. 9-1, we show the following sequence of bounds:

$$\begin{aligned} V_c &= \iiint_{V_c'} |\det \mathcal{J}|\, dx'dy'dz' \\ &\geq (\min \sigma_i) \int_0^{z_0} dz' \iint_{A'(z')} |\mathcal{J}_x \times \mathcal{J}_y|\, dx'dy' \\ &= (\min \sigma_i) \int_0^{z_0} A(z')\, dz' \\ &\geq (\min \sigma_i) \int_0^{z_0} A'(z')\, dz' = (\min \sigma_i) V_c' \\ &\geq V_c'/B. \end{aligned} \tag{9.1}$$

The first inequality in Eq. (9.1) follows from $|\det \mathcal{J}| = \|\mathcal{J}^T(\mathcal{J}_x \times \mathcal{J}_y)\|$, and $\|\mathcal{J}^T\mathbf{u}\| \geq (\min \sigma_i)\|\mathbf{u}\|$ for any vector $\mathbf{u}$ [195]. The $x'y'$ integral in the second line of Eq. (9.1) is simply the area $A(z')$ that $A'(z')$ maps to, and $A(z') \geq A'(z')$ because the outer boundary (solid dots in Fig. 9-1) is identical in A and A' and the flat surface A' is the minimal area for this boundary. Finally, the last step in Eq. (9.1) stems from elementary properties of singular values: the eigenvalues of $\mathcal{J}\mathcal{J}^T/|\det \mathcal{J}|$ are simply $b \leq \sigma_i^2/\sigma_1\sigma_2\sigma_3 \leq B$ [195], and algebraic manipulation of this inequality yields $B^{-1} \leq \sigma_i \leq b^{-1}$. Thus, we have shown that $V_c \geq V_c'/B$, but

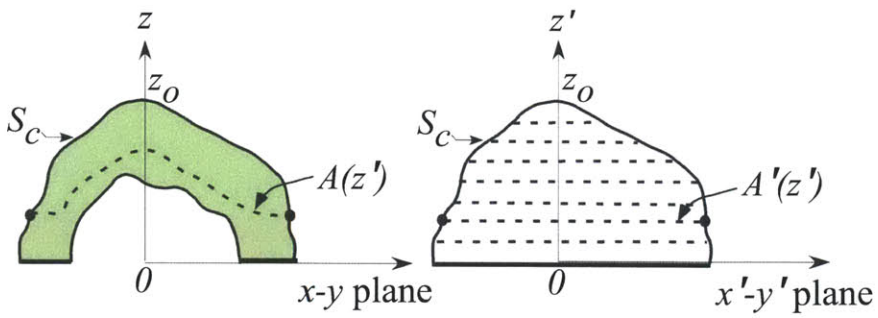

Figure 9-1: The cloaked volume V_c' in virtual space can be divided into flat cross-sections $A'(z')$ for each $z' \in [0, z_0]$. These are mapped to curved surfaces $A(z')$ in X. The invariance of the outer surface S_c means that the boundaries (solid dots) of $A(z')$ and $A'(z')$ coincide, and hence $A(z') \geq A'(z')$.

since the outer surface S_c is invariant it follows that the thickness of the cloak must scale proportional to the object thickness divided by B.

9.2 Cloak losses

We will analyze losses due to imperfections via perturbation theory: we first obtain the fields in a perfect cloak (ε and μ given exactly by the transformation law), and then consider the lowest-order absorption or scattering in the presence of small imperfections. Suppose the object is illuminated by an incident planewave with electric-field amplitude E_0, which means that the total field (incident + reflected) in the ambient medium and in V_c' has amplitude $\leq 2E_0$. For a perfect cloak, the fields in the cloak V_c are simply given by $(\mathcal{J}^T)^{-1}$ multiplied by the reflected planewave in virtual space X' [155], and hence the field amplitude E in the cloak is $\leq 2E_0/(\min \sigma_i) \leq 2E_0 B$ (using the bounds on the singular values σ_i from above). We must consider the worst-case losses for arbitrary incident waves; since this is bounded below by the loss from any particular incident wave, it is convenient to consider glancing-angle p-polarized planewaves where the field is a constant E_0 everywhere in V_c' and $E \geq E_0 b$ in the cloak (from the bound on $\max \sigma_i$).

In assuming that the incident wave has a constant amplitude E_0 in virtual space, we are implicitly assuming that the ambient medium is lossless (as for vacuum) or at

least that its losses are negligible over the lengthscale of the cloaked object. The case of a more lossy ambient medium is discussed in Sec. 9.2.3.

9.2.1 Absorption

An absorption imperfection is a small deviation $\Delta \operatorname{Im} \varepsilon$ in the imaginary part of ε compared with $\varepsilon_a \mathcal{J}\mathcal{J}^T / \det \mathcal{J}$. (Similarly for μ, but it suffices to consider electric absorption here.) This gives a change $\frac{\omega}{2} \operatorname{Re} \int_{V_c} \mathbf{E}^* \Delta \operatorname{Im} \varepsilon \mathbf{E}$ in the time-average absorbed power at a frequency ω [79]. To lowest order in $\Delta\varepsilon$, we can take $\mathbf{E}$ to be the field in the perfect cloak, and suppose for simplicity that the absorption is isotropic ($\Delta \operatorname{Im} \varepsilon$ is a scalar), and therefore the worst-case change in the absorbed power is $\geq \frac{\omega}{2} E_0^2 b^2 |\int_{V_c} \Delta \operatorname{Im} \varepsilon|$. Combined with our initial requirement on the scattering cross-section, the change in the absorbed power must be $\leq f s_g I_0$ where $I_0 = \frac{1}{2} E_0^2 \sqrt{\varepsilon_a / \mu_a}$ is the incident intensity, and we obtain the following bound:

$$\text{mean } \Delta \operatorname{Im} \varepsilon \leq \frac{f\sqrt{\varepsilon_a/\mu_a}}{\omega b^2} \frac{s_g}{V_c} \leq \frac{f\sqrt{\varepsilon_a/\mu_a} B}{\omega b^2} \frac{s_g}{V_c'}, \tag{9.2}$$

(using the V_c inequality above). The ratio s_g/V_c' scales as 1/thickness, so this means that the mean $\Delta \operatorname{Im} \varepsilon$ scales inversely with the thickness of the object to be cloaked. (The $1/\omega$ dependence means that this can be interpreted as a bound on the conductivity.)

Equation (9.2) is a necessary condition on the loss, but is too optimistic to be a sufficient condition. For example, suppose that the ambient medium is lossy, so that $\Delta \operatorname{Im} \varepsilon$ can have either sign depending on whether the cloak is more or less lossy than the ambient medium (or alternatively, $\Delta \operatorname{Im} \varepsilon < 0$ could come from gain). Equation (9.2) is satisfied if the more-lossy and less-lossy regions of the cloak average to zero, but in fact this will not result in a zero scattering cross-section for arbitrary incident waves. For example, a narrow incident beam (rather than a planewave) will interrogate the loss in some regions of the cloak more than others, and even a planewave at a different angle will create a standing-wave pattern that has higher field intensity in some regions—in these cases, any delicate cancellation in

the absorption will be destroyed. Instead, we can derive a sufficient condition on the loss by bounding the change in absorption above rather than below. In particular, $|\frac{\omega}{2}\operatorname{Re}\int_{V_c}\mathbf{E}^*\Delta\operatorname{Im}\varepsilon\mathbf{E}| \leq \frac{\omega}{2}V_c(2E_0B)^2\max|\Delta\operatorname{Im}\varepsilon|$, and thus it is sufficient for

$$\max|\Delta\operatorname{Im}\varepsilon| \leq \frac{f\sqrt{\varepsilon_a/\mu_a}}{4\omega B^2}\frac{s_g}{V_c} \leq \frac{f\sqrt{\varepsilon_a/\mu_a}}{4\omega B}\frac{s_g}{V_c'}. \tag{9.3}$$

This is, perhaps, stronger than strictly necessary; we conjecture that a weaker sufficient condition exists that replaces $\max|\Delta\operatorname{Im}\varepsilon|$ by an average of $\Delta\operatorname{Im}\varepsilon$ in the smallest region that can be interrogated by an incident wave (i.e. some wavelength-scale region). Regardless, both the sufficient and necessary conditions on the absorption imperfections scale inversely with the object thickness. This is true regardless of whether the ambient medium is lossy, and also means that any gain-based compensation of absorption must become increasingly exact for larger objects.

9.2.2 Random imperfections

Small random imperfections can be thought of as scatterers distributed randomly throughout V_c with some polarizability α (dipole moment $\mathbf{p} = \alpha\mathbf{E}$) [89]. For example, a small change $\Delta\varepsilon$ in a small region δV corresponds to $\alpha = \Delta\varepsilon\delta V$. Computing α for surface roughness is more involved, but is conceptually similar [89]. If these imperfections are uncorrelated, then on average the scattered power is simply the mean dipole radiation from each scatterer multiplied by the density d_α of scatterers (the radiation from different scatterers is incoherent, so interference terms average to zero) [79]. This radiation is most easily computed by transforming each scatterer back to virtual space X', where the polarizability is $\alpha' = (\mathcal{J}^T)^{-1}\alpha\mathcal{J}^{-1}\det\mathcal{J}$, so that $|\alpha'| \geq |\alpha|/B$ from above. The radiated power of a point source in virtual space (homogeneous above a ground plane) varies with distance and orientation above the ground plane due to the image dipole source below the ground plane, but the worst-case (over all incident waves) average (over all scatterer positions) scattered power is proportional (with a constant factor of

order unity) to the radiated power of a point source in the homogeneous medium, $(\alpha' E_0)^2 \omega^4 \mu_a \sqrt{\varepsilon_a \mu_a}/12\pi$ [79]. Multiplying this by the number $d_\alpha V_c$ of scatterers and comparing to the requirement on the worst-case loss, one finds that $\alpha^2 d_\alpha$ is bounded above by a quantity proportional to $B^3 f s_g / V_c'$, which again scales inversely with the thickness of the cloaked object.

Note that, while gain could conceivably be used to compensate for absorption loss, it does not seem applicable to scattering from imperfections.

9.2.3 Lossy ambient media

In the beginning of Sec. 9.2 above, we assumed the incident wave, in virtual space (no object) has roughly constant amplitude across all of V_c'. This assumption, however must be modified in an ambient medium with significant absorption losses, e.g. for cloaking in a lossy fluid. In particular, suppose that the incident wave has a decay length L in the ambient medium (i.e. it decays as $\sim e^{-\text{distance}/L}$). As long as the object diameter is $\ll L$, we can neglect the ambient losses and treat the incident wave as above, but when the diameter becomes $\gtrsim L$ the analysis must be modified. Intuitively, at this point an observer can only "see" a portion of diameter $\sim L$ of the object anyway, so effectively the object diameter is bounded and we would no longer expect the cloaking difficulty to increase with object size. More precisely, Eq. (9.2) is modified because the absorption loss $\frac{\omega}{2} E_0^2 b^2 |\int_{V_c} \Delta \operatorname{Im} \varepsilon|$ of an incident planewave now includes an $e^{-\text{distance}/L}$ factor in the integrand. This means that, while the integral is over all of V_c, only a volume $\sim s_g L$ contributes significantly, and one obtains $\sim s_g L \operatorname{mean} \Delta \operatorname{Im} \varepsilon$ instead of $\sim V_c \operatorname{mean} \Delta \operatorname{Im} \varepsilon$. The s_g factors in Eq. (9.2) now cancel and the bound on $\operatorname{mean} \Delta \operatorname{Im} \varepsilon$ becomes independent of diameter as expected. Similarly for the integral in Eq. (9.3) and for the number of scatterers ($\sim d_\alpha s_g L$ instead of $d_\alpha V_c$) in Sec. 9.2.2.

Chapter 10

Limitations of Isolated-Object Cloaking: 3D Case

In the previous chapter, we studied the limitations of ground-plane cloaking. We showed that the absorption losses and scattering off of imperfections scale with the size of the object being cloaked. Here, we look at the case of isolated-object cloaking. We arrive at similar results as in the case of ground-plane cloaking. Namely, we show that the losses scale with the size of the object. However, isolated-object cloaking faces two additional limitations: if we require any bounds on the material properties b, then perfect cloaking is impossible even at one frequency; the object can not be mapped to a single point and become fully invisible and there is a lower bound on the effective surface area of the object in the virtual space. This work was published in Ref. 70. Second, although perfect cloaking over a non-zero bandwidth is impossible by causality, we show that even in non-perfect cloaking, the bandwidth of cloaking decreases with the size of the object; we are currently preparing a manuscript for publication on these bandwidth limitations. Consider the problem of cloaking an isolated object of volume V_o in a homogeneous isotropic ambient medium using a transformation-based cloak of volume V_c surrounding the object. As above, the cloak material is determined from a coordinate mapping with Jacobian $\mathcal{J}$, via the equations $\varepsilon = \varepsilon_a \mathcal{J}\mathcal{J}^T/\det\mathcal{J}$ and $\mu = \mu_a \mathcal{J}\mathcal{J}^T/\det\mathcal{J}$. Now, however, there is no ground plane, and so the coordinate

mapping instead attempts to shrink the object: it maps the physical space X to a virtual space X' in which the object volume V_o is mapped to a smaller volume V_o', as shown in Fig. 10-1. (As for ground-plane cloaking, the outer cloak surface S_c is invariant, since the transformation is the identity outside of the cloak.) Perfect cloaking corresponds to the case in which V_o' is a single point, but we will show that this is not possible if the index contrast (eigenvalues of $\mathcal{J}\mathcal{J}^T/\det\mathcal{J}$) is bounded below by $b > 0$ as above. We will also show results analogous to our results for ground-plane cloaking: if the index contrast is bounded above by B, then the thickness of the cloak must scale with the object diameter, and correspondingly the losses (from absorption or imperfection) in the cloak must decrease with the object diameter.

Before developing our general results, however, we will begin with a specific illustrative example which demonstrates these scalings: an adaptation of Pendry's linear-scaling cloak design [155] to a nonzero V_o'.

10.1 Example: A spherical linear-scaling cloak

Suppose that the cloaked object is a sphere of radius R_1 and the cloak has outer radius R_2. We shrink the object to a sphere of radius R_1' with the transformation $r' = R_1' + (r - R_1)(R_2 - R_1')/(R_2 - R_1)$ in spherical coordinates. This leads to the following transformed materials in the cloak region (applying the general spherical-coordinate version of the transformation [155]):

$$\begin{aligned} \varepsilon_\theta/\varepsilon_a = &\mu_\theta/\mu_a = \varepsilon_\phi/\varepsilon_a = \mu_\phi/\mu_a = \frac{R_2 - R_1'}{R_2 - R_1}, \\ \varepsilon_r/\varepsilon_a = &\mu_r/\mu_a = \frac{R_2 - R_1}{R_2 - R_1'} \frac{\left[R_1' + \frac{r - R_1}{R_2 - R_1}(R_2 - R_1')\right]^2}{r^2} \end{aligned} \tag{10.1}$$

At the inner cloak surface $r = R_1$, the radial components of ε and μ simplify to $\frac{R_2 - R_1}{R_2 - R_1'}(R_1'/R_1)^2$, which vanishes for $R_1' = 0$. If we impose a lower bound b on the singular values of $\varepsilon\mu/\varepsilon_a\mu_a$, however, then R_1' cannot vanish. In particular, if $R_1 \gg R_1'$ then $(R_1'/R_1)^2 \sim b$, and hence the area reduction $S_o'/S_o \sim b$. Below, we

show in general (for arbitrary non-spherical transformations) that a $b > 0$ condition imposes a lower bound on the area reduction.

Even if $b = 0$, there is still a relationship between the index-contrast upper bound B and the cloak thickness. Suppose $R_1' \ll R_1$ (i.e. it is an effective cloak). Then the tangential components of ε and μ are $\approx R_2/(R_2 - R_1) \leq B$, from which it immediately follows that the cloak thickness $R_2 - R_1$ is bounded by $R_2/B \sim R_1/B$. This is identical to what we found for ground-plane cloaking, above: cloak thickness scales with object thickness divided by B. Below, we will generalize this thickness scaling to arbitrary non-spherical cloaks.

10.2 General limits on cloaking cross section

In this section, we show in general that the S_o', the effective surface area of the cloaked object, must be $\geq S_o b^2$. That is, if $b > 0$, then the area S_o' as seen by an observer must scale proportional to the object area, so that the cross section is reduced by a bounded factor. (For objects much larger than the wavelength, the scattering cross section is proportional to the geometric cross section S_o'.) From the spherical example above, it might be possible to derive an even tighter $\sim b$ bound on S_o', but the $S_o b^2$ bound here is sufficient to demonstrate the scaling with S_o.

Our basic approach is to write down S_o in terms of an integral over S_o', and then to bound the integral:

$$S_o = \oiint_{S_o'} |\mathcal{J}_u \times \mathcal{J}_v| \, dA'(u', v'), \tag{10.2}$$

where $(u, v) \leftrightarrow (u', v')$ is some coordinate system of the surfaces, $dA'(u', v')$ is the area element in S_o', and $\mathcal{J}_u$ and $\mathcal{J}_v$ are columns of the Jacobian matrix as in Sec. 9.1. We must then bound $|\mathcal{J}_u \times \mathcal{J}_v|$ similar to our previous analysis, but this is conceptually complicated somewhat by the fact that the coordinate system here is non-Cartesian, meaning that $\mathcal{J}$ no longer has the same bounds. To circumvent this difficulty, we apply a standard trick from differential geometry [43]: we cover the

surface S'_o with small overlapping neighborhoods N'_k (a locally finite open covering [43]) where the surface is locally approximately flat, in which case we can define a local Cartesian coordinate system and use the Cartesian bounds on $\mathcal{J}$ (ultimately taking the limit of infinitesimal neighborhoods so that the local flatness becomes exact). To combine these local results, one uses a partition of unity [43]: a set of "bump" functions $p'_k(u', v')$ (nonzero only on N'_k) such that $\sum_k p'_k = 1$ on S'_o. Thus, we obtain:

$$S_o = \sum_k \iint_{N'_k} |\mathcal{J}_u \times \mathcal{J}_v| \, p'_k dA'. \tag{10.3}$$

Now, in the limit of infinitesimal neighborhoods N'_k, we can freely treat each integral as being over a local Cartesian coordinate system, in which case $\mathcal{J}$ is the ordinary Cartesian Jacobian matrix. Now, we use two facts. First, we know $|\det \mathcal{J}| = \sigma_1\sigma_2\sigma_3 \leq (\max \sigma_i)^2(\min \sigma_i)$. Second, similar to Sec. 9.1, $|\det \mathcal{J}| = \|\mathcal{J}^T(\mathcal{J}_u \times \mathcal{J}_v)\| \geq (\min \sigma_i)|\mathcal{J}_u \times \mathcal{J}_v|$ from general properties of singular values [195]. Combining these two inequalities, we find $|\mathcal{J}_u \times \mathcal{J}_v| \leq (\max \sigma_i)^2 = 1/b^2$. Substituting this bound into the integral above, we finally obtain:

$$\begin{aligned} S_o &\leq \sum_k \iint_{N'_k} p'_k dA'/b^2 \\ &= \oiint_{S'_o} \left(\sum_k p'_k \right) dA'/b^2 = \oiint_{S'_o} (1) d\ A'/b^2 \\ &= S'_o/b^2, \end{aligned} \tag{10.4}$$

the desired inequality. (Note that if the object contains corners where S'_o is not locally flat, that does not affect this analysis since those corner regions have zero measure in the integral. The bounds on $\mathcal{J}$ mean that corners in S_o must be mapped to corners in S'_o and vice versa, and the integrand is finite.)

We suspect that a tighter bound, proportional to b instead of b^2, can be proved by taking into account the fact that the coordinate transformation must leave S_c invariant. In the spherically symmetrical example above, the purely radial nature of

the coordinate transformation caused it to have at most one factor of $1/b$ in $|\det \mathcal{J}|$ (only one eigenvalue of ε and μ vanishes for $R_1' = 0$) and not two, leading to an $S_o \sim S_o'/b$ dependence. A similar single factor of b in the general case would give a $1/b$ scaling in the inequalities above by replacing $(\max \sigma_i)^2$ with $\max \sigma_i$.

10.3 General scaling of cloak thickness and loss

A simple linear scaling of the cloak volume, as mentioned at the beginning of Sec. 9.1, follows immediately from the bounds on $|\det \mathcal{J}|$:

$$V_c = \iiint_{V_c'} |\det \mathcal{J}|\, dx'dy'dz' \geq (\min \sigma_i)^3 V_c' \geq V_c'/B^3 \tag{10.5}$$

As for ground-plane cloaking, we can define a mean thickness V_c/S_c of the cloak. For a useful cloak, $V_o' \ll V_o$ and hence $V_c' \approx V_c + V_o$. Thus, $V_c'/S_c' = V_c'/S_c$ is a mean total (cloak + object) diameter. Therefore, we have just demonstrated the inequality $V_c/S_c \geq (V_c'/S_c')/B^3$ (to lowest order in V_o'/V_o), which means that the mean cloak thickness V_c/S_c must scale proportional to the object diameter.

The spherical Pendry example above leads us to suspect that a tighter bound $\sim 1/B$ can be derived. Similar to the ground-plane example in Sec. 9.1, we expect that the key point is that we have not yet taken into account the constraint $S_c' = S_c$ on the cloaking transformation. However, our main goal here is to demonstrate the scaling of cloak thickness with object thickness, not to fine-tune the constant factor. Now that we know that cloak thickness must scale with object thickness, an analysis of losses similar to that in the ground-plane case above must apply. As the object becomes thicker, incident rays travel for a longer distance through the cloak, and hence for a fixed tolerance on the cross section the losses per unit distance must shrink proportional to the object diameter. Hence absorption losses (or rather, the *difference* between the cloak absorption in X' and the absorption of the ambient medium) and imperfections must scale inversely with the diameter. As we pointed out in chapter 8 [72], precisely such an inverse scaling of absorption tolerance with

diameter has been demonstrated numerically [213] for spherical cloaks, and our work now shows that this relationship is general. As suggested by some authors [66], gain could be used to compensate for absorption (but not disorder), but as discussed above our results imply that this compensation must become more and more exact as the object diameter increases.

For example, let us explicitly consider absorption imperfections for the idealized case of $b = 0$: that is, suppose we are able to map V_o to a single point $V_o' = 0$, but have a finite B and are still concerned with imperfect materials. In this case, an incident planewave of amplitude E_0 corresponds to a planewave of constant amplitude in the cloak for perfect materials, or to lowest order in the imperfections for imperfect materials. Then, similar to our analysis in Sec. 9.2.1, the change in absorbed power is bounded below by $\frac{\omega}{2}E_0^2|\int_{V_c}\Delta\,\mathrm{Im}\,\varepsilon/(\max\sigma_i)^2|$. Since we require that this change also be bounded above by the incident intensity ($\sim E_0^2$) multiplied by some fraction f of the geometric cross-section s_g, we obtain a necessary condition on the absorption imperfection as in Sec. 9.2.1. In particular it follows that an averaged absorption imperfection $|\int_{V_c}\Delta\,\mathrm{Im}\,\varepsilon/(\max\sigma_i)^2|/V_c$ must scale proportional to $s_g/V_c \sim s_g B^3/V_c'$, where V_c'/s_g is proportional to the diameter. In Sec. 9.2.1, we further used $\max\sigma_i \leq 1/b$ and pulled out a b^2 factor, but that is not appropriate when $b = 0$, so instead we must leave a $1/(\max\sigma_i)^2$ weight factor (which is only zero at the inner surface of the cloak) in the average of -$\Delta\,\mathrm{Im}\,\varepsilon$. As in Sec. 9.2.1, this is not a sufficient condition because the observer need not use planewaves—for interrogating the cloak with a focused beam, a stronger condition must apply, in which $\Delta\,\mathrm{Im}\,\varepsilon$ within a small (wavelength-scale) volume must go to zero as diameter increases. If $b > 0$, this analysis is only slightly modified in principle (although the precise expression becomes much more complicated): the small scattered field (assuming $V_o' \ll V_o$) from the nonzero $V_o' > 0$ modifies the field in X' by a small amount over most of V_c' (except immediately adjacent to V_o'), which should only change the proportionality of the $\mathrm{Im}\,\Delta\varepsilon$ scaling by a small factor.

10.4 Bandwidth limitations and scaling

A major obstacle with cloaking of isolated objects is the bandwidth [155]. As discussed by Pendry [155] and reviewed in chapter 7, perfect cloaking of an isolated objects over a non-zero bandwidth is impossible, a result that also follows from causality [135]. Here we want to generalize this question, considering not only the zero bandwidth of perfect cloaking, but also the nonzero bandwidth of *imperfect* cloaking, focusing on how the dispersive behavior of materials affect the performance of an isolated-object cloak and how this leads to bandwidth limitations that scale with the size of the object. More specifically, we study how the performance of the cloak over a fixed bandwidth $\Delta\omega$ around a design frequency ω_{op} gets worse as R increases regardless of how well the cloak performs at ω_{op}. To quantify the performance of the cloak, we look at the average total scattering cross-section σ_{scat} divided by the geometric cross-section s_g over some bandwidth $\Delta\omega$ and its behavior as the size of the object R increases.

Assuming:

- The total scattering cross section attains a minimum at the design frequency. In other words, the cloak is designed to optimize invisibility at the design frequency ω_{op}.
- We are interested in the performance of the cloak over a fixed bandwidth $\Delta\omega$.
- We are using best materials available and the widest range of n practically possible: bounds b and B are fixed independent of R as in the previous sections.

We show that:

- An average total scattering cross-section over bandwidth $\Delta\omega$ equals the total scattering cross-section evaluated at a single complex frequency $\omega_{\text{op}} + \Delta\omega$.
- Going to complex frequency $\omega_{\text{op}} + \Delta\omega$ is equivalent of introducing two kinds of losses: homogeneous losses throughout the space, and losses due to dispersion. The latter imposes bandwidth limitations on the cloak performance.

- The average total scattering cross-section divided by the geometric cross section over bandwidth $\Delta\omega$ increases proportional to the size of the object R.

- The increase in the average scattering cross-section divided by the geometric cross section due to non-zero bandwidth/dispersion is proportional to the size of the object.

If, instead of fixing the bandwidth, we fix the desired cloaking performance (the ratio of scattering cross-section to geometric cross-section), we show that the cloaking bandwidth must decrease as the object size increases. In addition to general proofs of these statements, we also illustrate them with numerical calculations for a specific spherical-cloak example, performed in collaboration with Prof. Cheng-Wei Qiu of NUS.

10.4.1 Frequency average of the scattering cross-section

To quantify the cloak performance over a finite bandwidth, we look at the frequency average of the total scattering cross section σ_{scat} (including both scattering and absorption) window centered at ω_{op} with bandwidth $\Delta\omega$. It is convenient to define this average via a Lorentzian window:

$$\text{mean}\,\sigma_{\text{scat}} = \int_{-\infty}^{\infty} \sigma_{\text{scat}} \cdot \frac{\Delta\omega/\pi}{(\omega-\omega_{\text{op}})^2 + \Delta\omega^2} d\omega \tag{10.6}$$

The key to a simple analysis of this average cross section is to transform the problem so that we no longer need to consider the cross section at many frequencies at once. Instead, using techniques from complex analysis, it turns out that we can express $\text{mean}\,\sigma_{\text{scat}}$ in terms of a scattering problem at a *single complex* frequency, which we then show to be equivalent to a *single real*-frequency problem with an *artificial loss*.

If σ_{scat} were an analytic function in the upper-half complex-frequency plane ($\text{Im}\,\omega > 0$), then Eq. (10.6) could be evaluated by means of a contour integration. By the residue theorem, closing the integration contour in the upper-half plane by an infinite-radius semicircle, we would obtain the residue of the single pole of the

Lorentzian at $\omega = \omega_{\rm op} + i\Delta\omega$. Unfortunately, $\sigma_{\rm scat}$ is not an analytic function in general, but it is *related* to an analytic function by the optical theorem: $\sigma_{\rm scat}(\omega) = \mathrm{Im}\, f(\omega)$, where f is proportional to the complex amplitude of the forward-scattered wave. (We review a proof of the optical theorem for complex frequency in Sec. 10.4.2.) Thanks to causality, $f(\omega)$ must be an analytic function in the upper-half plane (just like any other linear-response function), and hence Eq. (10.6) becomes:

$$\begin{aligned} \text{mean}\,\sigma_{\rm scat} &= \mathrm{Im}\int_{-\infty}^{\infty} f \frac{\Delta\ \omega/\pi}{(\omega-\omega_{\rm op})^2+\Delta\omega^2} d\omega \\ &= \mathrm{Im}\, f(\omega_{\rm op}+i\Delta\omega) \\ &= \text{“}\sigma_{\rm scat}(\omega_{\rm op}+i\Delta\omega),\text{”} \end{aligned} \tag{10.7}$$

where we have put the last line in quotation marks as it turns out that $\mathrm{Im}\, f$ at a complex frequency is not quite a conventional scattering cross section, as discussed in detail in Sec. 10.4.2.

The equivalence of frequency average of the scattering cross section and its value at the complex frequency $\omega_{\rm op} + \Delta\omega$ leads to a particularly simple way to analyze the effect of frequency averaging, because using a complex frequency is mathematically equivalent to using complex (lossy) materials at a real frequency. In particular, consider Maxwell's equations in the frequency domain:

$$\nabla \times \mathbf{E} = -i\omega\mu\mathbf{H} + \mathbf{K} \tag{10.8}$$

$$\nabla \times \mathbf{H} = +i\omega\epsilon\mathbf{E} + \mathbf{J} \tag{10.9}$$

where J and K are amplitudes of single frequency electric and magnetic currents respectively. The first terms on the right-hand sides of these equations for a single complex frequency become $i(\omega_{\rm op}+i\Delta\omega)\epsilon(\omega_{\rm op}+i\Delta\omega)\mathbf{E} = i\left[\epsilon(\omega_{\rm op}+i\Delta\omega).(1+i\frac{\Delta\omega}{\omega_{\rm op}})\right]\omega_{\rm op}$, which looks exactly like the term for a real frequency $\omega_{\rm op}$ at a modified permittivity $\epsilon(\omega_{\rm op}+i\Delta\omega).(1+i\frac{\Delta\omega}{\omega_{\rm op}})$. (The μ also works in a similar manner.) Therefore

operating at a complex frequency corresponds to two changes in the material: (i) it adds an absorption to the whole space due to multiplication of ϵ by $1 + i\Delta\omega/\omega_{\text{op}}$, and (2) we evaluate ϵ at $(\omega_{\text{op}} + i\Delta\omega)$ rather than at the real frequency ω_{op}.

We discuss the consequences these changes in the materials have to the cloaking performance in the next section, but the basic idea is that, since frequency averaging corresponds to a single frequency with modified lossy materials, we can quote the results from the previous section: as the object diameter increases, the consequences of these "effective" losses become worse and worse proportional to the diameter. It follows that the only way to maintain even imperfect cloaking is to decrease the bandwidth as the diameter increases.

10.4.2 The optical theorem and analytic continuation to complex ω

The key our understanding of bandwidth-averaged cross sections is our ability to analytically continue the problem into the complex frequency plane. In order to fully understand the consequences of such a technique, however, it is useful to review a derivation of the optical theorem, and in particular that the total cross section is the imaginary part of a causal function (analytic in the upper half plane). Moreover, we need to explain precisely how the analytic continuation of this function leads to a scattering problem at complex frequency, in order to analyze the implications for cloaking.

The optical theorem [79] has a long history, going all the way back to 1871 [141], and is closely related to energy conservation: energy that is scattered or absorbed must be correspondingly removed from the forward-transmitted wave (i.e., there must be a shadow). The first version of the theorem was stated by Lord Rayleigh in the context of scattering by small particles, and it was later combined with the Kramers–Kronig relations to derive a relationship between the index of refraction and absorption measuremnts [141]. Generalizations of the theorem to arbitrary scatterers appeared first in quantum mechanics using scattering-matrix formalisms,

and in electromagnetism shortly thereafter [141]. The theorem is widely used in electromagnetic, acoustic, and quantum scattering theory [141], as well as in quantum field theory [157]. In quantum field theory, the optical theorem is important in order to elucidate the analytic properties of scattering cross section [157]. Similar in spirit to this paper, the optical theorem is also used in the "ITEP sum rules" of quantum chromodynamics to relate integrals of scattering cross sections multiplied by Lorentzian windows (or powers thereof) to scattering amplitudes at single points in momentum space via contour integration [157]. On another related note, contour integration of the imaginary part of Green's functions (and other scattering amplitudes) is also a key technique for computing Casimir interactions in quantum field theory [86], where the relationship between the imaginary part of the Green's function and the field fluctuation statistics (the fluctuation–dissipation theorem) is, like the optical theorem, derived from energy-conservation considerations [104].

Consider any scatterer (in a linear time-invariant system), described by some change $\Delta\varepsilon$ and $\Delta\mu$, within a finite volume, in the permittivity and the permeability compared to the ambient medium. Let the incident field (in the absence of the scatterer) be described by a six-component vector field $\psi_{\text{inc}} = \begin{pmatrix} E_{\text{inc}} \\ H_{\text{inc}} \end{pmatrix}$. In the presence of the scatterer, this is modified to a new total field $\psi = \psi_{\text{inc}} + \psi_{\text{scat}}$, the sum of the incident and scattered fields. The charge perturbations in the scatterer are described by electric and magnetic polarization currents $\mathbf{J} = -i\omega\Delta\varepsilon\mathbf{E}$ and $\mathbf{K} = -i\omega\Delta\mu\mathbf{H}$, which can be combined into a six-component current field $\xi = \begin{pmatrix} J \\ K \end{pmatrix}$. Abstractly, we can write $\xi = A\psi_{\text{inc}}$ for some linear operator A relating the incident fields to the induced currents (the solution to a volume-integral equation), and causality (currents come after fields) implies that A is an analytic function in the upper-half complex-ω plane [103].

Physically, the scattered field is the field produced by these oscillating induced currents ξ in the scatterer, and the interactions of the currents and fields provide a simple way to characterized the absorbed and scattered powers. Then the total

absorbed power $P_{textabs}$ equals the total time-average (ψ) incoming Poynting flux, which by Poynting's theorem can be shown to be the time-average work done by the total fields on the induced currents:

$$P_{\text{abs}} = \frac{1}{2}\operatorname{Re}\int(\mathbf{E}^* \cdot \mathbf{J} + \mathbf{H}^* \cdot \mathbf{K}) = \frac{1}{2}\operatorname{Re}\langle\psi, \xi\rangle = \frac{1}{2}\operatorname{Re}\langle\psi_{\text{inc}}, \xi\rangle + \frac{1}{2}\operatorname{Re}\langle\psi_{\text{scat}}, \xi\rangle,$$

where we have defined an inner product $\langle\cdots|\cdots\rangle$ in bra-ket notation. Similarly, the scattered power is the work done by the currents on the scattered field: $-\frac{1}{2}\operatorname{Re}\langle\psi_{\text{scat}}, \xi\rangle$. Therefore, the total (absorbed + scattered) power is:

$$\begin{aligned} P_{\text{tot}} = P_{\text{abs}} + P_{\text{scat}} &= \left(\frac{1}{2}\operatorname{Re}\langle\psi_{\text{inc}}, \xi\rangle + \frac{1}{2}\operatorname{Re}\langle\psi_{\text{scat}}, \xi\rangle\right) - \frac{1}{2}\operatorname{Re}\langle\psi_{\text{scat}}, \xi\rangle \\ &= \frac{1}{2}\operatorname{Re}\langle\psi_{\text{inc}}, \xi\rangle \\ &= \frac{1}{2}\operatorname{Re}\langle\psi_{\text{inc}}, A\psi_{\text{inc}}\rangle \\ &= \operatorname{Im} f(\omega), \end{aligned} \tag{10.10}$$

where we have defined $f(\omega) = \frac{i}{2}\langle\psi_{\text{inc}}, A\psi_{\text{inc}}\rangle$, a causal function since A represents a causal linear response. The interpretation of f as a forward-scattering amplitude is actually irrelevant to our proof, but it follows from the fact that $f(\omega) = \frac{\omega}{2}\langle\psi_{\text{inc}}, i\xi/\omega\rangle$, which is simply the ψ_{inc} component (a Fourier component for an incident planewave) of the field $\sim i\xi/\omega$ corresponding to the currents ξ.

The key question, for our purposes, is to understand what f looks like at a complex frequency. Suppose we have homogeneous, lossless ambient medium with an incident planewave $\psi_{\text{inc}} = \psi_0 e^{\omega x/c}$ propagating in the $+x$ direction for a constant amplitude ψ_0. Then

$$f(\omega) = \frac{i}{2}\left\langle\psi_0, e^{-i\omega x/c}A(\omega)e^{+i\omega x/c}\psi_0\right\rangle, \tag{10.11}$$

moving all the x dependence to the right-hand side of the inner product. At a complex frequency $\omega \to \omega + i\gamma$ in the upper-half plane ($\gamma > 0$), $A(\omega + i\gamma)$ is analytic by causality while the $e^{\mp i\omega x/c \pm \gamma x/c}$ terms are exponentially growing and decaying

with x, respectively. In order to perform contour integrations of $f(\omega)$ multiplied by a Lorentzian, it is not enough for f to be analytic, however: it must also be bounded as $|\omega| \to \infty$ so that we can close the contour above. Intuitively, this occurs because the the exponential growth of $e^{+\gamma x/c}$ cancels the exponential decay of $e^{-\gamma x/c}$. Mathematically, as $|\omega| \to \infty$ the $\Delta\varepsilon$ and $\Delta\mu$ must $\to 0$ (since all susceptibilities vanish in the limit of infinite frequency [79, 103]), and so A simplifies: to lowest order for weak scatterers (i.e., the first Born approximation), $\xi \approx A\psi_{\text{inc}} \approx \begin{pmatrix} -i\omega\Delta\varepsilon E_{\text{inc}} \\ -i\omega\Delta\mu H_{\text{inc}} \end{pmatrix}$, in which case the exponential factors exactly cancel. We will use a similar procedure, below, to analyze the effects of small imperfections in the cloak due to material dispersion.

10.5 Consequences of dispersion on frequency-averaged scattering

Combining the previous two sections, we can now relate the frequency-averaged scattering cross-section, for an incident planewave (say in the x direction) to the solution of a single scattering problem at a single complex frequency:

$$\text{mean}\,\sigma_{\text{scat}} = \frac{1}{2}\,\text{Re}\left\langle \psi_0, e^{-i\omega x/c} A(\omega) e^{+i\omega x/c}\psi_0 \right\rangle, \tag{10.12}$$

evaluated at $\omega = \omega_{\text{op}} + i\Delta\omega$. The $A(\omega)e^{+i\omega x/c}\psi_0$ represents the induced currents ξ, for the materials evaluated at the complex ω, in response to an exponentially decaying incident planewave (multiplied by a constant amplitude ψ_0). As in Sec. 9.2.1, it is convenient to imagine solving this problem in *virtual* space, where an exact cloak would be replaced by no scatterers at all and an imperfect cloak is replaced by weak scatterers.

Furthermore, as noted above, solving the scattering problem at a complex frequency is equivalent to solving it at the real frequency ω_{op} with complex materials

$\epsilon(\omega_{\rm op}+i\Delta\omega).(1+i\frac{\Delta\omega}{\omega_{\rm op}})$ and $\mu(\omega_{\rm op}+i\Delta\omega).(1+i\frac{\Delta\omega}{\omega_{\rm op}})$. These materials differ from those at $\omega_{\rm op}$ in two ways: they are multiplied by an artificial absorption $1+i\frac{\Delta\omega}{\omega_{\rm op}}$, and also material dispersion means that $\epsilon(\omega_{\rm op}+i\Delta\omega)\neq\epsilon(\omega_{\rm op})$ (and similarly for μ). We discuss each of these effects in turn below.

Irrelevance of artificial absorption

The first change in the materials is the multiplication of ϵ and μ by complex constant $(1+i\frac{\Delta\omega}{\omega_{\rm op}})$. But this change does not hurt the cloaking performance simply because if $\epsilon=\mathcal{J}\epsilon_a\mathcal{J}^T/\det\mathcal{J}$ is a valid transformation-based cloak, so is $\epsilon.(1+\frac{\Delta\omega}{\omega_{\rm op}})=\mathcal{J}\epsilon_a.(1+\frac{\Delta\omega}{\omega_{\rm op}})\mathcal{J}^T/\det\mathcal{J}$, except that this is now a transformation-based cloak for a lossy ambient medium $\epsilon_a.(1+i\frac{\Delta\omega}{\omega_{\rm op}})$, an similarly for μ. Therefore this change in the materials does not create any scattering.

Impact of material dispersion

The second change in the materials is that we need to evaluate ϵ and μ at the complex frequency $\omega_{\rm op}+i\Delta\omega$ instead of $\omega_{\rm op}$, and here the presence of material dispersion (unavoidable for any material other than vacuum) acts to spoil the cloak. In particular, as reviewed below, causality and other fundamental principles imply that $\mathrm{Im}\,\epsilon$ and $\mathrm{Im}\,\mu$ are both strictly *positive* $\omega_{\rm op}+i\Delta\omega$ for $\Delta\omega>0$ [103], corresponding to an unavoidable additional absorption in the complex-frequency scattering problem. Unlike the artificial absorption in the previous section, this is an absorption *defect*: ϵ (both Im and Re) differs from the cloaking transformation of the ambient medium by

$$\Delta\epsilon=\left[\epsilon(\omega_{\rm op}+i\Delta\omega)-\epsilon(\omega_{\rm op})\right](1+i\frac{\Delta\omega}{\omega_{\rm op}})$$

(similarly for μ).

We can analyze the effect of this effective absorption imperfection similarly to Sec. 9.2.1. As in Sec. 9.2.2, however, it convenient to transform the imperfections to virtual space where the scattering problem is easier: $\Delta\epsilon'\geq\Delta\epsilon/B$. We are only

interested in the case of a good cloak,[1] which means that the imperfections must be small, and therefore virtual space at the complex ω is nearly a homogeneous medium with small imperfections $\Delta\epsilon'$ (and $\Delta\mu'$) from material dispersion. In this case, as in Sec. 9.2.1, to lowest order the scattering current is simply $\mathbf{J}' = -i\omega_{\text{op}}\Delta\epsilon'\mathbf{E}_{\text{inc}}$ and similarly for $\mathbf{K}'$, or $\xi' \approx A\psi_{\text{inc}} \approx \begin{pmatrix} -i\omega\Delta\varepsilon' E_{\text{inc}} \\ -i\omega\Delta\mu' H_{\text{inc}} \end{pmatrix}$. Therefore we find that the $e^{\pm\Delta\omega x/c}$ exponential exactly cancel as in the previous section, leaving:

$$\text{mean}\,\sigma_{\text{scat}} \approx \frac{1}{2}\int_{V_c'} \left[|\mathbf{E}_0|^2\,\text{Im}\,\Delta\varepsilon' + |\mathbf{H}_0|^2\,\text{Im}\,\Delta\mu'\right] \tag{10.13}$$

This has two main consequences. First, $\sigma_{\text{scat}} > 0$ since $\text{Im}\,\Delta\varepsilon$ and $\text{Im}\,\Delta\mu$ are strictly positive as noted above. Even if the cloak is a perfect cloak at w_{op}, it is imperfect over a non-zero bandwidth. This is, therefore, an alternative proof that cloaking of isolated-objects over a non-zero bandwidth is impossible for physical, causal materials. Second, exactly as in Sec. 9.2.1, it immediately follows that $\text{mean}\,\sigma_{\text{scat}}$ grows $\sim V_c' \sim V_c \sim V_o$ and hence the frequency-averaged cloaking efficiency $\text{mean}\,\sigma_{\text{scat}}/s_g$ scales proportional to the mean diameter V_o/s_g. We further explore the consequences of this diameter scaling in Sec. 10.5.1.

Causality and materials at complex ω

Here, we review a known consequence of causality that is the key to our analysis above: for a passive medium, $\text{Im}\,\epsilon(\omega) > 0$ in the upper half plane $\text{Im}\,\omega > 0$ as proved in Ref. 103. A condensed proof of this fact is as follows: $\text{Im}\,\epsilon$ is analytic in the upper-half plane by causality, and is therefore a harmonic function in the upper-half plane, and so can obtains its minimum only on the boundary of its domain, except in the trivial case of vacuum where it is a constant function ($\text{Im}\,\epsilon = 0$). In particular, consider the upper-right quadrant of the complex-ω plane. Along the positive real axis, $\text{Im}\,\epsilon \geq 0$ for a passive material (in the absence of gain),

[1] Once the imperfections become large and the cloak becomes useless, all of these scaling relations will break down and one will simply have a σ_{scat} comparable to the geometric cross-section.

even for idealized lossless materials. Along the positive imaginary axis, $\operatorname{Im}\epsilon = 0$ since $\epsilon(-\omega) = \epsilon(\omega)^*$ for real-valued physical fields [79]. As $|\omega| \to \infty$ one must have $\operatorname{Im}\omega \to 0$. Hence the minimum of $\operatorname{Im}\epsilon$ along the boundary of the upper-right quadrant is zero and it is strictly positive in the interior.

For physical materials, $\operatorname{Im}\omega > 0$ along the positive real-ω axis except at $\omega = 0$ in order to satisfy the second law of thermodynamics, and this is the usual case in which the above statement is proved. However, it is also interesting to consider the idealized limit of lossless materials (such as a plasma model or a lossless resonance), in order to study bandwidth-averaged cloaking with idealized materials. The above proof works even in this case as long as one is a little cautious about the case of poles in ε lying exactly on the real-ω axis, in order to exclude the case where $\operatorname{Im}\epsilon$ diverges to $-\infty$ as the real axis is approached. In particular, we must restrict ourselves to idealized materials that are the limit of physical lossy materials as the losses go zero, so that we are taking the limit as poles approach the real axis from below. In this case, $\operatorname{Im}\epsilon$ in the upper-half plane is the limit of a positive quantity and hence cannot go to $-\infty$.

10.5.1 Scaling of cloaking bandwidth with diameter

As shown in the previous section, the fractional cross-section mean σ_{scat}/s_g over a bandwidth $\Delta\omega$ around ω_{op} must scale proportional to the diameter d, at least as long as mean $\sigma_{\text{scat}}/s_g \ll 1$ (i.e. until cloaking breaks down completely). There are two possible sources of this linear scaling, depending on whether mean σ_{scat} is limited by the cross-section at ω_{op} or whether it is limited by the bandwidth of a dip in the cross-section around ω_{op}. In the former case, the physical mechanism is simply that the losses at ω_{op} scale with diameter, as already shown earlier in this chapter. In the latter case, however, it leads to a new prediction: in a bandwidth-limited cloak, the bandwidth must narrow with object diameter.

First, let us consider the bandwidth scaling from generic dimensional considerations. Because of material dispersion, one expects good cloaking to only be possible in some limited bandwidth $\sim \Gamma$ around some design frequency ω_{op}, in which case small

mean σ_{scat}/s_g requires $\Delta\omega \ll G$. Then, if we Taylor-expand mean σ_{scat}/s_g in $\Delta\omega/G$, we would generically expect an expansion of the form:

$$\text{mean}\,\sigma_{\text{scat}}/s_g = \sigma_{\text{scat}}(\omega_{\text{op}})/s_g + \#\Delta\omega/\Gamma + O[(\Delta\omega/\Gamma)^2]$$

for some coefficient $\# > 0$.[2] For sufficiently small bandwidths $\Delta\omega \ll \Gamma\sigma_{\text{scat}}(\omega_{\text{op}})/s_g$, this is dominated by the scattering at ω_{op}, which scales linearly with diameter in the presence of imperfections as shown earlier in this chapter. On the other hand, for a good single-frequency cloak at ω_{op}, there is a regime $\Gamma\sigma_{\text{scat}}(\omega_{\text{op}})/s_g \ll \Delta\omega \ll \Gamma$ where the second term dominates, i.e. where material dispersion is the limiting factor. This term must also scale linearly with the diameter in order for mean σ_{scat}/s_g to scale linearly, and hence we can conclude that the bandwidth Γ of the cloak generally *scales inversely with diameter.*

This is best illustrated by a simple example. Consider the case where σ_{scat}/s_g achieves a minimum $f \ll 1$ at ω_{op} with a Lorentzian lineshape of width Γ, going to 1 at frequencies far from ω_{op}:

$$\frac{\sigma_{\text{scat}}(\omega)}{s_g} = 1 - \frac{\Gamma^2}{(\omega - \omega_{\text{op}})^2 + \Gamma^2}(1 - f). \tag{10.14}$$

The integral of Eq. (10.6) for this $\sigma_{\text{scat}}(\omega)$ can be evaluated analytically by the residue theorem to obtain

$$\text{mean}\,\sigma_{\text{scat}} = 1 - \frac{1 - f}{1 + \Delta\omega/\Gamma} \approx f + (1 - f)\Delta\omega/\Gamma + O[(\Delta\omega/\Gamma)^2].$$

Since this must scale linearly with diameter for any $f \ll 1$ and any $\Delta\omega \ll \Gamma$, it follows that both f and $1/\Gamma$ scale linearly with diameter until cloaking breaks down.

[2] One might imagine a circumstance in which # is exactly zero, but this seems to require a very unusual-shaped $\sigma_{\text{scat}}(\omega)$; we have not been able to explicitly construct any such function.

10.5.2 Numerical example

To illustrate and validate our predictions of this scaling of cloaking bandwidth with diameter, we collaborated with Prof. Cheng-Wei Qui of NUS to perform explicit numerical calculations of the scaling for an example spherical-cloaking problem. At an operating frequency ω_{op}, we use exactly a Pendry cloak as in chapter 7: a sphere of radius R_1 is surrounded by a cloak of radius $R_2 > R_1$ that is linearly mapped to an empty sphere ($R_1' = 0$), resulting in materials:

$$\epsilon_r(\omega_{\text{op}}) = \mu_r(\omega_{\text{op}}) = \frac{R_2 - R_1}{R_2}\left(\frac{r - R_1}{r}\right)^2$$

$$\epsilon_\theta(\omega_{\text{op}}) = \epsilon_\phi(\omega_{\text{op}}) = \mu_\theta(\omega_{\text{op}}) = \mu_\phi(\omega_{\text{op}}) = \frac{R_2}{R_2 - R_1}$$

in units where $\epsilon_0 = \mu_0 = 1$. We fixed $R_2 = 1.5R_1$, so that the cloaking material parameters are the same for all R_1, merely rescaled in space.

So, at ω_{op} we have a *perfect* cloak, and the only limiting factor is the bandwidth: we use idealized lossless materials but with causal dispersion relations that satisfy the Kramers–Kronig constraints. In particular, we use a combination of two limiting cases: a plasma model (a limit of a Drude model as losses go to zero) and a limit of lossless Lorentzian resonance (corresponding to a polarization field described by a lossless harmonic oscillator) at a frequency ω_0. Combined with the prescribed Pendry values at ω_{op} from above, this results in dispersion relations:

$$\epsilon_r(r,\omega) = \mu_r(r,\omega) = 1 - \left[1 - \frac{R_2}{R_2 - R_1}\left(\frac{r - R_1}{r}\right)^2\right]\frac{\omega_{\text{op}}^2}{\omega^2},$$

$$\epsilon_\theta(\omega,r) = \epsilon_\phi(\omega,r) = \mu_\theta(\omega,r) = \mu_\phi(\omega,r) = 1 + \frac{R_1}{R_2 - R_1}\frac{\omega_0^2 - \omega_{\text{op}}^2}{\omega_0^2 - \omega^2}.$$

The Lorentzian resonance frequency ω_0 can be chosen arbitrarily, and we used

$$\omega_0 = 2\omega_{\text{op}}.$$

This geometry was then simulated using a spectral (spherical-harmonic expansion) scattering-matrix method as described in Ref. 164. The continuously varying

anisotropic material parameters are approximated by a large number of piecewise-homogeneous isotropic layers [164]. The total scattering cross-section was computed over a range of frequencies for $R_1 = \lambda_{\text{op}}, 2\lambda_{\text{op}}, and 4\lambda_{\text{op}}$ (where $\lambda_{\text{op}} = 2\pi c/\omega_{\text{op}}$), and is plotted in Fig. 10-2. The results in Fig. 10-2 are converged with resolution (the number of spherical layers) to within a few percent accuracy. At ω_{op}, the cloak should theoretically be perfect, but we obtain a small nonzero σ/s_g ($< 10^{-3}$) due to the discretization errors, which vanishes with increasing resolution. As expected, the material dispersion prevents this from being a good cloak except at frequencies in a narrow bandwidth around ω_{op}, and this bandwidth becomes narrower as the diameter increases. Quantitatively, if we look at the bandwidth at a fixed σ/s_g of about 1/4 its maximum, we find that the bandwidths for $R_1 = 2\lambda_{\text{op}}$ and $R_1 = 4\lambda_{\text{op}}$ are $\approx 1/2.1$ and $1/4.1$ times the bandwidth for $R_1 = \lambda_{\text{op}}$, respectively, almost exactly the predicted linear scaling.

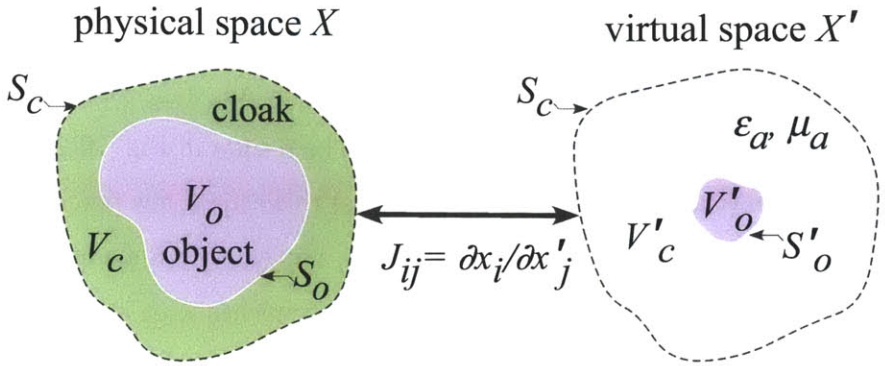

Figure 10-1: isolated-object cloak

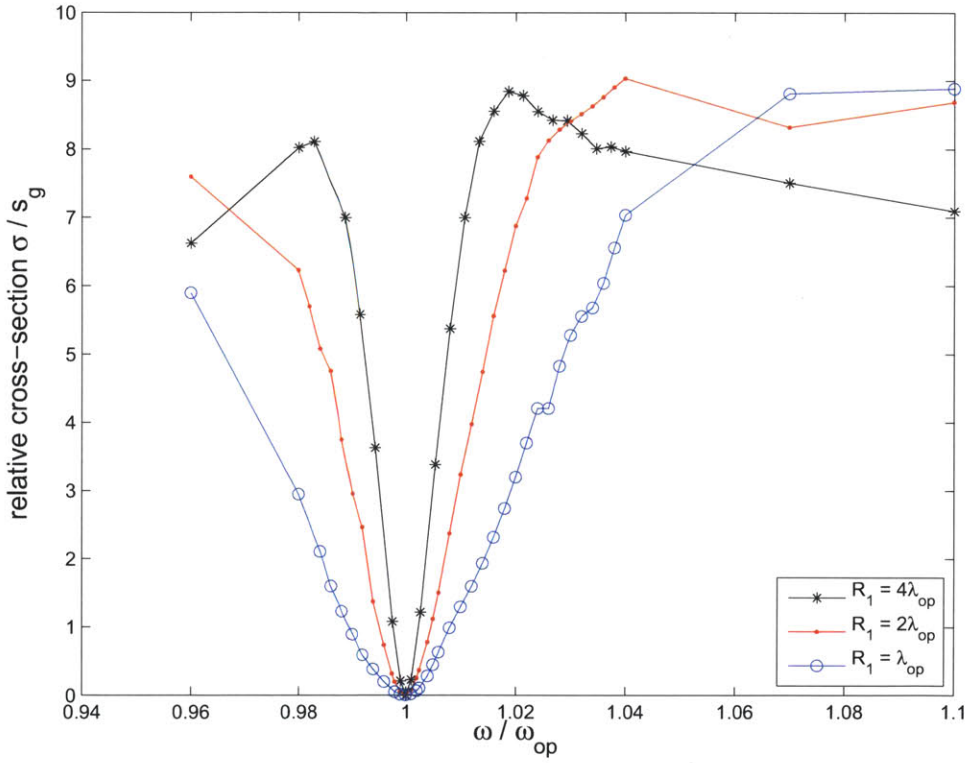

Figure 10-2: Relative cross-section versus frequency for a spherical cloak designed to be a perfect Pendry cloak at ω_{op} and showing the effects of material dispersion at other frequencies, computed by a spectral scattering-matrix method. As predicted, the cloaking bandwidth decreases linearly with the object radius, for three object radii relative to $\lambda_{\text{op}} = 2\pi c/\omega_{\text{op}}$.

Chapter 11

Concluding Remarks

11.1 Nonlinear optics

Starting from coupled-mode theory, we modeled systems involving nonlinear frequency conversion processes. We did a detailed analysis of the system dynamics and conversion efficiencies for the case of THG and DFWM and introduced the system-parameter regimes in which there exists a stable high-efficiency solution and introduced different ways of exciting such solutions.

We observed that these systems show very interesting non-linear behavior such as Hopf bifurcations, limit cycles, and multi-stability. Some of these behaviors could possibly have applications other than frequency conversion. for example, one could use the limit cycles to build optical cloaks or use bi-stable regimes as an optical switch. These are some of the future directions that can start from this work. As for the application of frequency conversion, we showed that such systems, if designed with correct system parameters, obtain efficiencies ten times larger than the highest reported in literature.

Next we proposed the design process of such systems. We introduced devices and their design process for SHG, THG, and DFWM. We went in detail to design a system exhibiting high-efficiency second harmonic generation. A next step in this direction is to achieve an actual design for the case of THG. This would require further care due to SPM/XPM effects. Another interesting idea is to explore

DFWM in a realistic geometry such as a photonic crystal slab coupled to an index-guided waveguide and study some of the dynamical effects arising from SPM/XPM; a project that we are currently working on. Obviously, the next step after deigning actual system is their fabrication.

Given the promise of high-efficiency frequency conversion and its wide applications, I believe this area of nonlinear optics has a lot of potential.

11.2 Cloaking

Generalizing from simple $1d$ thought experiment, we have proved scaling laws that point to an inherent practical difficulty (though not a mathematical impossibility) in scaling experimental ground-plane cloaking of small objects to larger ones. Furthermore, we showed that very similar analysis can be applied to cloaking of isolated objects—bounded index contrasts will imply a bounded reduction in the scattering cross section, a cloak thickness proportional to the object diameter, and imperfection tolerances that shrink with the object diameter. It might be possible to further generalize the results in this thesis to cloaks that are not derived from coordinate transformations (similar to the generality of our one-dimensional analysis from chapter 8 [72]).

However, the most serious constraint on isolated-object cloaking seems to be the bandwidth, which must be zero for perfect cloaking [135, 155]. Clearly, if perfect cloaking is possible (theoretically) at a single frequency, then imperfect cloaking (reduction of the cross section by a given factor) must persist over a nonzero bandwidth, and we showed that the bandwidth of such imperfect isolated-object cloaking scales with the object diameter from causality constraints. The same limitation does not apply to ground-plane cloaking, and the reason our proof does not translate to this case is that we *want* ground-plane cloaks to scatter, albeit in a way that mimics the ground-plane, so total scattering cross-section is not a useful figure of merit.

An alternative direction is to consider relaxations of the cloaking problem that

might prove more practical. In particular, it would be valuable to make precise the intuition that the cloaking problem becomes easier if the incident waves are restricted (e.g. to planewaves from a certain range of angles) and/or the observer is limited (e.g. only scattered waves at certain angles are visible, or only amplitude but not phase can be detected), since this is arguably the situation in most experiments. (For example, current "stealth" aircraft are designed in the radar regime mainly to reduce back-scattering only [142].) Another interesting possibility is to consider "cloaking" that attempts to make one object look like a different object of a similar size rather than making it invisible (although this approach is similar in spirit to ground-plane cloaking and may have similar limitations).

Bibliography

[1] E. Abraham, W. J. Firth, and J. Carr. Self-oscillation and chaos in nonlinear Fabry-Perot resonators with finite response time. *Physics Lett. A*, pages 47–51, 1982.

[2] P. P. Absil, J. V. Hryniewicz, B. E. Little, P. S. Cho, R. A. Wilson, L. G. Joneckis, and P. T. Ho. Wavelength conversion in gaas micro-ring resonators. *OL*, 25:554–556, 2000.

[3] I. H. Agha, Y. Okawachi, M. A. Foster, J. E. Sharping, and A. L. Gaeta. Four-wave parametric oscillations in disperion-compensated high-q silica microspheres. *PRA*, 76:043837, 2007.

[4] G. P. Agrawal. *Fiber-Optic Communication Systems.* Wiley, Hoboken, NJ, 2002.

[5] G. D. Aguanno, M. Centini, M. Scalora, C. Sibilia, M. Bertolotti, M. J. Bloemer, and C. M. Bowden. Generalized coupled-mode theory for $\chi^{(2)}$ interactions in finite multi-layered structures. *J. Opt. Soc. Am. B*, 19:2111–2122, 2002.

[6] G. D. Aguanno, M. Centini, M. Scalora, C. Sibilia, Y. Dumeige, P. Vidavovic, J. A. Levenson, M. J. Bloemer, C. M. Bowden, J. W. Haus, and M. Bertolotti. Photonic band edge effects in finite structures and applications to $\chi^{(2)}$ interactions. *Phys. Rev. E*, 64:016609, 2001.

[7] Christos Argyropoulos, Efthymios Kallos, and Yang Hao. Dispersive cylindrical cloaks under nonmonochromatic illumination. *Phys. Rev. E*, 81(1):01661, 2010.

[8] J. A. Armstrong, N. loembergen, J. Ducuing, and P. S. Pershan. Interactions between light waves in a nonlinear dielectric. *Phys. Rev.*, 127:1918–1939, 1962.

[9] A. Ashkin, G. D. Boyd, and J. M. Dziedzic. Resonant optical second harmonic generation and mixing. *IEEE J. Quantum Electron.*, 2:109–124, 1966.

[10] A .V. Balakin, V. A. Bushuev, B. I. Mantsyzov, I. A. Ozheredov, E. V. Petrov, and A. P. Shkurinov. Enhancement of sum frequency generation near the photonic band edge under the quasiphase matching condition. *Phys. Rev. E*, 63:046609, 2001.

[11] Guy Bartal, Ofer Manela, and Mordechai Segev. Spatial four wave mixing in nonlinear periodic structures. *Phys. Rev. Lett.*, 97(7):073906, Aug 2006.

[12] V. Berger. Second-harmonic generation in monolithic cavities. *J. Opt. Soc. Am. B*, 14(6):1351–1360, Jun 1997.

[13] V Berger. Second-harmonic generation in monolithic cavities. *J. Opt. Soc. Am. B*, 14:1351, 1997.

[14] P. Bermel, Alejandro Rodriguez, Steven G. Johnson, J. D. Joannopoulos, and Marin Soljačić. Single-photon all-optical switching using waveguide-cavity quantum electrodynamics. *Phys. Rev. A*, 74:043818, 2006.

[15] J. C. Berquist, H. Hemmati, and W. M. Itano. High power second harmonic generation of 257 nm radiation in an external ring cavity. *Opt. Commun.*, 43:437–442, 1982.

[16] Zhuanfang Bi, Alejandro W. Rodriguez, Hila Hashemi, David Duchesne, Marko Loncar, Ke-Ming Wang, and Steven G. Johnson. High-efficiency second-harmonic generation in doubly-resonant $\chi^{(2)}$ microring resonators. Submitted for publication, 2012.

[17] Peter Bienstman and R. Baets. Optical modelling of photonic crystals and VCSELs using eigenmode expansion and perfectly matched layers. *Optical and Quantum Electron.*, 33(4–5):327–341, 2001.

[18] Robert W. Boyd. *Nonlinear Optics.* Academic Press, California, 1992.

[19] Jorge Bravo-Abad, Alejandro W. Rodriguez, John D. Joannopoulos, Peter T. Rakich, Steven G. Johnson, and Marin Soljačić. Efficient low-power terahertz generation via on-chip triply-resonant nonlinear frequency mixing. *Applied Physics Letters*, 96:101110, 2010.

[20] M. Brieger, H. Busener, A. Hese, F. V. Moers, and A. Renn. Enhancement of single frequency shg in a passive ring resonator. *Opt. Commun.*, 38:423–426, 1981.

[21] D. H. Broaddus, M. A. Foster, I. H. Agha, J. T. Robinson, M. Lipson, and A. L. Gaeta. Silicon-wavelength-coupled high-q chalcogenide microspheres. *Opt. Express*, 17:5998–6003, 2009.

[22] J. Brown. Proof of the manley-rowe relations from quantum considerations. *EL*, 1:23–24, 1965.

[23] Ian B. Burgess, Alejandro W. Rodriguez, Murray W. McCutcheon, Jorge Bravo-Abad, Yinan Zhang, Steven G. Johnson, and Marko Lončar. Difference-frequency generation with quantum-limited efficiency in triply-resonant nonlinear cavities. *Optics Express*, 17:9241–9251, May 2009.

[24] Wenshan Cai, Uday K. Chettiar, Alexander V. Kildishev, and Vladimir M. Shalaev. Optical cloaking with metamaterial. *Nature Photonics*, 1:224–227, 2007.

[25] P. Campagnola and C.-Y. Dong. Second harmonic generation in microscopy: principles and applications to disease diagnosis. *Laser and Photonics Reviews*, 5(1):13–26, 2011.

[26] Stefano Cattaneo and Martti Kauranen. Application of second-harmonic generation to retardation measurements. *J. Opt. Soc. Am. B*, 20(3):11415–11421, 2003.

[27] B. Chambers and A. Tennant. Optimised design of jaumann radar absorbing materials using a genetic algorithm. *Radar, Sonar and Navigation, IEE Proceedings* -, 143(1):23 –30, feb 1996.

[28] J. K. Chandalia, B. J. Eggleton, R. S. Windeler, S. G. Kosinski, X. Liu, and C. Xu. Adiabatic coupling in tapered air-silica microstructured optical fiber. *IEEE Photon. Tech. Lett.*, 13(1):52–54, 2001.

[29] Hongsheng Chen, Bae-Ian Wu, Baile Zhang, and Jin Au Kong. Electromagnetic wave interactions with a metamaterial cloak. *Phys. Rev. Lett.*, 99(6):063903, 2007.

[30] Huanyang Chen and C. T. Chan. Time delays and energy transport velocities in three dimensional ideal cloaking devices. *J. Appl. Phys.*, 104:033113, 2008.

[31] Huanyang Chen, Zixian Liang, Peijun Yao, Xunya Jiang, Hongru Ma, and C. T. Chan. Extending the bandwidth of electromagnetic cloaks. *Phys. Rev. B*, 76(24):241104(R), 2007.

[32] Xianzhong Chen, Yu Luo, Jingjing Zhang, Kyle Jiang, John B. Pendry, and Shiang Zhang. Macroscopic invisibility cloaking of visible light. *Nature Communications*, 2(176), 2011.

[33] Weng Cho Chew, Jin Jian-Ming, Eric Michielssen, and Song Jiming. *Fast and Efficient Algorithms in Computational Electromagnetics*. Artech, Norwood, MA, 2001.

[34] S. K. Choi, R. D. Li, C. Kim, and P. Kumar. Traveling-wave optical parametric amplifier: investigation of its phase-sensitive and phase-insensitive gain response. *JOSA-B*, 14:1564–1575, 1997.

[35] Andreas Christ and Hans L. Hartnagel. Three-dimensional finite-difference method for the analysis of microwave-device embedding. *IEEE Trans. Microwave Theory Tech.*, 35(8):688–696, 1987.

[36] C. Cohen-Tannoudji, B. Din, and F. Laloë. *Quantum Mechanics*. Hermann, Paris, 1977.

[37] M. J. Collet and R. B. Levien. Two-photon loss model of intracavity second-harmonic generation. *Phys. Rev. A*, 43(9):5068–5073, 1990.

[38] A. R. Cowan and Jeff F. Young. Mode matching for second-harmonic generation in photonic crystal waveguides. *Phys. Rev. E*, 65:085106, 2002.

[39] Steven A. Cummer, Bogdan-Ioan Popa, David Schurig, David R. Smith, and John Pendry. Full-wave simulation of electromagnetic cloaking structures. *Phys. Rev. E*, 74(3):036621, 2006.

[40] L. Deng, E. W. Hagley, M. Trippenbach, Y. Band, P. S. Julienne, J. E. Simsarian, K. Helmerson, S. L. Rolston, and W. D. Phillips. Four-wave mixing with matter waves. *Nature*, 398:218–220, 1999.

[41] Andrea Di Falco, Claudio Conti, and Gaetano Assanto. Impedance matching in photonic crystal microcavities for second-harmonic generation. *Opt. Lett.*, 31:250, 2006.

[42] G. J. Dixon, C. E. Tanner, and C. E. Wieman. 432-nm source based on efficient second-harmonic generation of GaAlAs diode-laser radiation in a self-locking external resonant cavity. *Opt. Lett.*, 14:731–733, 1989.

[43] Manfredo P. do Carmo. *Riemannian Geometry.* Mathematics. Birkhäuser, Boston, MA, 1992.

[44] Tatyana V. Dolgova, Anton I. Maidykovski, Michail G. Martemyanov, Andrey A. Fedyanin, Oleg A. Aktsipetrov, Gerd Marowsky, Vladimir A. Yakovlev, Giorgio Mattei, Narumi Ohta, and Seiichiro Nakabayashi. Giant optical second-harmonic generation in single and coupled microcavities formed from one-dimensional photonic crystals. *J. Opt. Soc. Am. B*, 19:2129, 2002.

[45] P. D. Drummond, K. J. McNeil, and D. F. Walls. Non-equilibrium transitions in sub/second harmonic generation i: Semiclassical theory. *Optica Acta.*, 27(3):321–335, 1980.

[46] D. Duchesne, L. Razzari, M. Ferrera, R. Morandotti, S. Chu, B.E. Little, and D.J. Moss. Integrated optical hyper-parametric oscillator. In *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)*, pages 1 –3, 2010.

[47] Yannick Dumeige and Patrice Feron. Wispering-gallery-mode analysis of phase-matched doubly resonant second-harmonic generation. *Phys. Rev. A*, 74:063804, 2006.

[48] Gupta S. Dutta and Jose Jolly. Third harmonic generation in layered media in presence of optical bistability of the fundamental. *Pranama-journal of Physics*, 50:239, 1988.

[49] Tolga Ergin, Jad C. Halimeh, Nicolas Stenger, and Martin Wegener. Optical microscopy of 3d carpet cloaks:ray-tracing calculations. *Opt. Express*, 18(19):20535–20545, 2010.

[50] Tolga Ergin, Nicolas Stenger, Patrice Brenner, John B. Pendry, and Martin Wegener. Three-dimensional invisibility cloak at optical wavelength. *Science*, 328:337–339, 2010.

[51] J. Feinberg. Self-pumped, continuous-wave phase conjugator using internal reflection. *Opt. Lett.*, 7:486–488, 1982.

[52] M. M. Fejer. Nonlinear optical frequency conversion. *Phys. Today*, 47:25–32, 1994.

[53] F. S. Felber and J. H. Marburger. Theory of nonresonant multistable optical devices. *Appl. Phys. Lett.*, 28(12):731–733, 1976.

[54] M. Ferrera, D. Duchesne, L. Razzari, M. Peccianti, R. Morandotti, P. Cheben, S. Janz, D. X. Xu, B. E. Little, S. Chu, and D. J. Moss. Low power four wave mixing in an integrated, micro-ring resonator with q=1.2 million. *OE*, 17:14098–14103, 2009.

[55] M. Ferrera, L . Razzari, D. Duchesne, R. Morandotti, Z. Yang, M. Liscidini, J. E. Sipe, S. Chu, B. E. Little, and D. J. Moss. Low-power continuous-wave nonlinear optics in doped silica glass integrated waveguide structures. *Nature Photonics*, 2:737–740, 2008.

[56] Y. Fink, S. Fan, J. D. Joannopoulos, C Chen, and E. L. Thomas. Omnidirectional multilayer coating for enhanced optical waveguiding. *US Patent Application, Provisional*, October 1998. Ser. No. 60/104,153.

[57] A. Fiore, V. Berger, E. Rosencher, P. Bravetti, and J. Nagle. Phase matching using an isotropic nonlinear optical material. *Letters to Nature*, 391:463–466, 1997.

[58] W. R. Frei and H. T. Johnson. Finite-element analysis of disorder effects in photonic crystals. *Phys. Rev. B*, 70:165116, 2004.

[59] Canek Fuentes-Hernandez, Gabriel Ramos-Ortiz, Shuo-Yen Tseng, Michael P. Gaj, and Bernard Kippelen. Third-harmonic generation and its applications in optical image processing. *J. Mater. Chem.*, 19:7394–7401, 2009.

[60] Lucas H. Gabrielli, Jaime Cardenas, Carl B. Poitras, and Michal Lipson. Silicon nano-structure cloak operating at optical frequencies. *Nature Photonics*, 3:461–463, 2009.

[61] A. I. Gerguson and M .H. Dunn. Intracavity second harmonic generation in continuous-wave dye lasers. *IEEE J. Quantum Electron.*, 13:751–756, 1977.

[62] Christopher C. Gerry. Generation of optical macroscopic quantum superposition states via state reduction with a Mach-Zehnder interferometer containing a Kerr medium. *Phys. Rev. A*, 59(5):4095–4098, 1998.

[63] Magid Gharghi, Christopher Gladden, Thomas Zentgraf, Yongmin Liu, Xiaobo Yin, Jason Valentine, and Xiang Zhang. A carpet cloak for visible light. *Nano Letters*, 11(7):2825–2828, 2011.

[64] K. Grygiel and P. Szlatchetka. Chaos in second-harmonic generation of light. the case of a strain of pulses. *Opt. Comm.*, 91:241–246, 1992.

[65] J Hald. Second harmonic generation in an external ring cavity with a brewster-cut nonlinear cystal: theoretical considerations. *Optics Communications*, 197:169, 2001.

[66] Tiancheng Han, Cheng-Wei Qiu, Jiaming Hao, Xiaohong Tang, and Said Zouhdi. Gain-assisted transformation optics. *Opt. Express*, 19(9):8610–8615, 2011.

[67] Tiancheng Han, Cheng-Wei Qiu, and Xiaohong Tang. The general two-dimensional open-closed cloak with tunable inherent discontinuity and directional communication. *Appl. Phys. Lett.*, 97(124104), 2010.

[68] J. Hansryd, P. A. Andrekson, M. Westlund, J. Li, and P. O. Hedexkvist. Fiber-based optical parametric amplifiers and their applications. *IEEE J. Quantum Electron.*, 8:506–520, 2002.

[69] James A. Harrington. A review of IR transmitting, hollow waveguides. *Fiber Integr. Opt.*, 19:211–227, 2000.

[70] Hila Hashemi, A. Oskooi, J. D. Joannopoulos, and Steven G. Johnson. General scaling limitations of ground-plane and isolated-object cloaks. *Physical Review A*, 84:023815, 2011.

[71] Hila Hashemi, Alejandro W. Rodriguez, John D. Joannopoulos, Marin Soljačić, and Steven G. Johnson. **Nonlinear harmonic generation and devices in doubly-resonant Kerr cavities**. *Physical Review A*, 79(1):013812, 2008.

[72] Hila Hashemi, Baile Zhang, J. D. Joannopoulos, and Steven G. Johnson. Delay-bandwidth and delay-loss limitations for cloaking of large objects. *Phys. Rev. Lett.*, 104(25):253903, 2010.

[73] H. A. Haus. *Waves and Fields in Optoelectronics*. Prentice-Hall, Englewood Cliffs, NJ, 1984.

[74] H. A. Haus and W. Huang. Coupled-mode theory. *Proc. of the IEEE*, 79:1505–1518, 1991.

[75] J. E. Heebner, N. N. Lepeshkin, A. Schweinsberg, G. W. Wicks, R. W. Boyd, R. Grover, and P. T. Ho. Enhanced linear and nonlinear optical phase response of algaas microring resonators. *OL*, 29:769–771, 2004.

[76] Ying Huang, Yijun Feng, and Tian Jiang. Electromagnetic cloaking by layered structure of homogeneous isotropic materials. *Opt. Express*, 15(18):11133–11141, 2007.

[77] M. Ibanescu, Y. Fink, S. Fan, E. L. Thomas, and J. D. Joannopoulos. An all-dielectric coaxial waveguide. *Science*, 289:415–419, 2000.

[78] K. Inque. Four-wave-mixing in an optical fiber in the zero-dispersion wavelength region. *J. Lightwave Tech.*, 10:1553–1561, 1992.

[79] J. D. Jackson. *Classical Electrodynamics.* Wiley, New York, third edition, 1998.

[80] Wei Xiang Jiang, Tie Jun Cui, Guan Xia Yu, Xian Qi Lin, and Qiang Cheng. Arbitrarily elliptical-cylindrical invisible cloaking. *Journal of Physics D: Applied Physics*, 41:085504, 2008.

[81] Jianming Jin. *The Finite Element Method in Electromagnetics.* Wiley, New York, second edition, 2002.

[82] J. D. Joannopoulos, S. Fan, A. Mekis, and S. G. Johnson. *Photonic Crystals and Light Localization in the 21st Century*, volume 563 of *Series C: Mathematical and Physical Sciences*, chapter Novelties of Light with Photonic Crystals, pages 1–24. Kluwer, Dordrecht, The Netherlands, 2001.

[83] John D. Joannopoulos, Steven G. Johnson, Joshua N. Winn, and Rob ert D. Meade. *Photonic Crystals: Molding the Flow of Light.* Princeton University Press, second edition, February 2008.

[84] John D. Joannopoulos, Steven G. Johnson, Joshua N. Winn, and Robert D. Meade. *Photonic Crystals: Molding the Flow of Light.* Princeton University Press, second edition, February 2008.

[85] Steven Johnson and John Joannopoulos. Block-iterative frequency-domain methods for maxwell's equations in a planewave basis. *Opt. Express*, 8(3):173–190, 2001.

[86] Steven G. Johnson. Numerical methods for computing Casimir interactions. In Diego Dalvit, Peter Milonni, David Roberts, and Felipe da Rosa, editors, *Casimir Physics*, volume 834 of *Lecture Notes in Physics*, chapter 6, pages 175–218. Springer, Berlin, 2011.

[87] Steven G. Johnson, Peter Bienstman, M. A. Skorobogatiy, Mihai Ibanescu, Elefterios Lidorikis, and J. D. Joannopoulos. Adiabatic theorem and continuous coupled-mode theory for efficient taper transitions in photonic crystals. *Phys. Rev. E*, 66:066608, 2002.

[88] Steven G. Johnson and J. D. Joannopoulos. *Photonic Crystals: The Road from Theory to Practice*. Kluwer, Boston, 2002.

[89] Steven G. Johnson, M. L. Povinelli, M. Soljačić, A. Karalis, S. Jacobs, and J. D. Joannopoulos. Roughness losses and volume-current methods in photonic-crystal waveguides. *Appl. Phys. B*, 81:283–293, 2005.

[90] V. P. Kalosha and A. P. Khapalyuk. Mode birefringence in a single-mode elliptic optical fiber. *Sov. J. Quantum Electron.*, 13:109–111, 1983.

[91] V. P. Kalosha and A. P. Khapalyuk. Mode birefringence of a three-layer elliptic single-mode fiber waveguide. *Sov. J. Quantum Electron.*, 14:427–430, 1984.

[92] Boubacar Kanté, Andreé de Lustrac, Jean-Michel Lourtioz, and Shah Nawaz Burokur. Infrared cloaking based on the electric response of split ring resonators. *Opt. Express*, 16(12):9191–9198, 2008.

[93] Boubacar Kanté, Dylan Germain, and Andreé de Lustrac. Infrared cloaking based on the electric response of split ring resonators. *Phys. Rev. B*, 80(20):201104, 2009.

[94] M. Karisson. Four-wave-mixing in fibers with randomly varying zero-dispersion wavelength. *JOSA-B*, 15:2269–2275, 1998.

[95] T. Kawanishi and M. Izutsu. Coaxial periodic optical waveguide. *Opt. Express*, 7(1):10–22, 2000.

[96] J. C. Knight, J. Broeng, T. A. Birks, and P. St.-J. Russell. Photonic band gap guidance in optical fibers. *Science*, 282:1476–1478, 1998.

[97] Karl Koch and Gerald T. Moore. Singly resonant cavity-enhanced frequency tripling. *J. Opt. Soc. Am. B*, 16:448, 1999.

[98] Chris Kottke, Ardavan Farjadpour, and Steven G. Johnson. Perturbation theory for anisotropic dielectric interfaces, and application to sub-pixel smoothing of discretized numerical methods. *Phys. Rev. E*, 77:036611, 2008.

[99] W. J. Kozlovsky, W. P. Risk, W. Lenth, B. G. Kim, G. L. Bona, H. Jaeckel, and D. J. Webb. Blue light generation by resonator-enhanced frequency doubling of an extended-cavity diode laser. *Appl. Phys. Lett.*, 65:525–527, 1994.

[100] Karl S. Kunz and Raymond J. Luebbers. *The Finite-Difference Time-Domain Method for Electromagnetics*. CRC Press, Boca Raton, 1993.

[101] Mei-Ling Kuo, David J. Poxson, Yong Sung Kim, Frank W. Mont, Jong Kyu Kim, E. Fred Schubert, and Shawn-Yu Lin. Realization of a near-perfect antireflection coating for silicon solar energy utilization. *Opt. Lett.*, 33:2527–2529, 2008.

[102] Do-Hoon Kwon and Douglas H. Werner. Two-dimensional eccentric elliptic electromagnetic cloaks. *Appl. Phys. Lett.*, 92(1):013505, 2008.

[103] L. Landau, E. M. Lifshitz, and L. P. Pitaevskii. *Electrodynamics of Continuous Media.* Butterworth-Heinemann, Oxford, 2nd edition, 1984.

[104] L. D. Landau and E. M. Lifshitz. *Statistical Physics: Part 1.* Butterworth-Heinemann, Oxford, 3rd edition, 1980.

[105] N. I. Landy, N. Kundtz, and D. R. Smith. Designing three-dimensional transformation optical media using quasiconformal coordinate transformations. *Phys. Rev. Lett.*, 105:193902, 2010.

[106] Wah Tung Lau and Shanhui Fan. Creating large bandwidth line defects by embedding dielectric waveguides into photonic crystal slabs. *Appl. Phys. Lett.*, 81:3915–3917, 2002.

[107] J. H. Lee, J. Blair, V. A. Tamma, Q. Wu, S. J. Rhee, C. J. Summers, and W. Park. Direct visualization of optical frequency invisibility cloak based on silicon nanorod array. *Opt. Express*, 17:12922–12928, 2009.

[108] G. Lenz, B. J. Eggleton, C. K. Madsen, and R. E. Slusher. Optical delay lines based on optical filters. *IEEE J. Quantum Electron.*, 37(4):525–532, 2001.

[109] Ulf Leonhardt. Notes on conformal inisibility devices. *New Journal of Physics*, 8(7):118, 2006.

[110] Ulf Leonhardt and Tomáš Tyc. Broadband invisibility by non-euclidean cloaking. *Science*, 323(5910):110–112, 2009.

[111] M. D. Levenson, R. M. Shelby, A. Aspect, M. Reid, and D. F. Walls. Generation and detection of squeezed states of light by nondegenerate four-wave mixing in an optical fiber. *Phys. Rev. A*, 32:1550–1562, 1985.

[112] Jacob S. Levy, Mark A. Foster, Alexander L. Gaeta, and Michal Lipson. Harmonic generation in silicon nitride ring resonators. *Opt. Express*, 19(12):11415–11421, 2011.

[113] Jensen Li and J. B. Pendry. Hiding under the carpet: A new strategy for cloaking. *Phys. Rev. Lett.*, 101(20):203901, 2008.

[114] Rui Li, Jing Chen, Qian Xu, Fang-Fang Ren, Jianping Ding, and Hui-Tian Wang. Saturation effect and forward-dominant second harmonic generation in single-defect photonic crystals with dual localizations. *Opt. Lett.*, 31(22):3327, 2006.

[115] Wei Li, Jianguo Guan, Zhigang Sun, Wei Wang, and Qingjie Zhang. A near-perfect invisibility cloak constructed with homogeneous materials. *Opt. Express*, 17(26):23410–23416, 2009.

[116] Ron Lifshitz, Andy Arie, and Alon Bahabad. Photonic quasicrystals for nonlinear optical frequency conversion. *Phys. Rev. Lett.*, 95:133901, 2005.

[117] Shawn-Yu Lin, Edmund Chow, Vince Hietala, Pierre R. Villeneuve, and J. D. Joannopoulos. Experimental demonstration of guiding and bending of electromagnetic waves in a photonic crystal. *Science*, 282:274–276, 1998.

[118] Marco Liscidini and Lucio .A Andreani. Highly efficient second-harmonic generation in doubly resonant planar microcavities. *Appl. Phys. Lett.*, 85:1883, 2004.

[119] Marco Liscidini and Lucio .A Andreani. Second-harmonic generation in doubly resonant microcavities with periodic dielectric mirrors. *Phys. Rev. E*, 73:016613, 2006.

[120] R. Liu, C. Ji, J. J. Mock, J. Y. Chin, T. J. Cui, and D. R. Smith. Broadband ground-plane cloak. *Science*, 323:366–369, 2009.

[121] Tzu-Ming Liu, Cheng-Ta Yu, and Chi-Kuang Sun. 2 GHz repetition-rate femtosecond blue sources by second-harmonic generation in a resonantly enhanced cavity. *Appl. Phys. Lett.*, 86:061112, 2005.

[122] Xiu Liu, Chao Li, Kan Yao, Xiankun Meng, Wei Feng, Bingheng Wu, and Fang Li. Experimental verification of broadband invisibility using a cloak based on inductor-capacitor networks. *Appl. Phys. Lett.*, 95:191107, 2009.

[123] Marco Lončar, Tomoyuki Yoshie, Axel Scherer, Pawan Gogna, and Yueming Qiu. Low-threshold photonic crystal laser. *Appl. Phys. Lett.*, 81(15):2680–2682, 2002.

[124] Hua Ma. The open cloak. *Appl. Phys. Lett.*, 94:103501, 2009.

[125] Hui Feng Ma and Tie Jun Cui. Three-dimensional broadband ground-plane cloak made of metamaterials. *Nature Communications*, 1(124), 2010.

[126] Hui Feng Ma, Wei Xiang Jiang, Xin Mi Yang, Xiao Yang Zhou, and Tie Jun Cui. Compact-sized and broadband carpet cloak and free-space cloak. *Opt. Express*, 17:19947–19959, 2009.

[127] Bjorn Maes, Peter Bienstman, and Roel Baets. Modeling second-harmonic generation by use of mode expansion. *J. Opt. Soc. Am. B*, 22:1378, 2005.

[128] Martin Maldovan and Edwin L. Thomas. Diamond-structured photonic crystals. *Nature Materials*, 3:593–600, 2004.

[129] A. M. Malvezzi, G. Vecchi, M. Patrini, G. Guizzeti, L. C. Andreani, F. Romanato, L. Businaro, E. Di Fabrizio, A. Passaseo, and M. De Vittorio. Resonant second-harmonic generation in a GaAs photonic crystal waveguide. *Phys. Rev. B*, 68:161306, 2003.

[130] V. A. Mandelshtam and H. S. Taylor. Harmonic inversion of time signals and its applications. *J. Chem. Phys.*, 107(17):6756–6769, 1997. See erratum [131].

[131] V. A. Mandelshtam and H. S. Taylor. Erratum: "Harmonic inversion of time signals and its applications". *J. Chem. Phys.*, 109:4128, 1998.

[132] E. A. Marcatili and R. A. Schmeltzer. Hollow metallic and dielectric waveguides for long distance optical transmission and lasers. *Bell Syst. Tech. J.*, 43:1783–1809, 1964.

[133] Przemyslaw P. Markowicz, Hanifi Tiryaki, Haridas Pudavar, Paras N. Prasad, Nick N. Lepeshkin, and Robert W. Boyd. Dramatic enhancement of third-harmonic generation in three-dimensional photonic crystals. *Phys. Rev. Lett.*, 92(083903), 2004.

[134] Gail McConnell, Allister I. Ferguson, and Nigel Langford. Cavity-augmented frequency tripling of a continuous wave mode-locked laser. *J. Phys. D: Appl.Phys*, 34:2408, 2001.

[135] David A. B. Miller. On perfect cloaking. *Opt. Express*, 14(25):12457–12466, 2006.

[136] David A. B. Miller. Fundamental limit to linear one-dimensional slow light structures. *Phys. Rev. Lett.*, 99(4):203903, 2007.

[137] Mitsunobu Miyagi, Kazuhide Harada, and Shojiro Kawakami. Wave propagation and attenuation in the general class of circular hollow waveguides with uniform curvature. *IEEE Trans. Microwave Theory Tech.*, MTT-32(5):513–521, 1984.

[138] Mitsunobu Miyagi, Akihito Hongo, and Shojiro Kawakami. Transmission characteristics of dielectric-coated metallic waveguides for infrared transmission: slab waveguide model. *IEEE J. Quantum Electron.*, QE-19(2):136–145, 1983.

[139] Gerald T. Moore, Karl Koch, and E. C. Cheung. Optical parametric oscillation with intracavity second-harmonic generation. *Optics Communications*, 113:463, 1995.

[140] Xiaodong Mu, Yujie J. Ding, Haeyeon Yang, and Gregory J. Salamo. Cavity-enhanced and quasiphase-matched mutli-order reflection-second-harmonic generation from GaAs/AlAs and GaAs/AlGaAs multilayers. *Appl. Phys. Lett.*, 79:569, 2001.

[141] Roger G. Newton. Optical theorem and beyond. *Americal Journal of Physics*, 44(7), 1976.

[142] Leland M. Nicolai and Grant E. Carichner. *Fundamentals of aircraft and airship design.* American Institute of Aeronautics and Astronautics, Reston, VA, 2010.

[143] A. Nicolet, D. Zolla, and S. Guenneau. Electromagnetic analysis of cylindrical cloaks of an arbitrary cross section. *Opt. Lett.*, 33:1584–1586, 2008.

[144] Andrew H. Norton and C. Martijn de Sterke. Optimal poling of nonlinear photonic crystals for frequency conversion. *Opt. Lett.*, 28:188, 2002.

[145] Masaya Notomi, Akihiko Shinya, Satoshi Mitsugi, Goh Kira, Eiichi Kuramochi, and Takasumi Tanabe. Optical bistable switching action of Si high-q photonic-crystal nanocavities. *Opt. Express*, 13(7):2678–2687, 2005.

[146] A. F. Oskooi, D. Roundy, M. Ibanescu, P. Bermel, J. D. Joannopoulos, and S. G. Steven. Meep: A flexible free-software package for electromagnetic simulations by the fdtd method. *Comp. Phys. Comm.*, 181:687–702, 2010.

[147] Z. Y. Ou and H. J. Kimble. Enhanced conversion efficiency for harmonic generation with double resonance. *Opt. Lett.*, 18:1053–1055, 1993.

[148] Edward D. Palik, editor. *Handbook of optical constants of solids II.* Academic Press, 1991.

[149] Alberto Parini, Gaetano Bellanca, Stefant Trillo, Matteo Conforti, Andrea Locatelli, and Costantino De Angelis. Self-pulsing and bistability in nonlinear Bragg gratings. *J. Opt. Soc. Am. B*, 24:2229–2237, 2007.

[150] R. Paschotta, K. Fiedler, P. Kurz, and J. Mlynek. Nonlinear mode coupling in doubly resonant frequency doublers. *Appl. Phys. Lett.*, 58:117, 1994.

[151] Alessia Pasquazi, Raja Ahmad, Martin Rochette, Michael Lamont, Brent E. Little, Sai T. Chu, Roberto Morandotti, and David J. Moss. All-optical wavelength conversion in an integrated ring resonator. *Opt. Express*, 18:3858–3863, 2010.

[152] Alessia Pasquazi, Yongwoo Park, José Aza na, François Légaré, Roberto Morandotti, Brent E. Little, Sai T. Chu, and David J. Moss. Efficient wavelength conversion and net parametric gain via four wave mixing in a high index doped silica waveguide. *Opt. Express*, 18:7634–7641, 2010.

[153] S. Pearl, H. Lotem, and Y. Shimony. Optimization of laser intracavity second-harmonic generation by a linear dispersion element. *J. Opt. Soc. Am. B*, 16:1705, 1999.

[154] J. B. Pendry, A. J. Holden, D. J. Robbins, and W. J. Stewart. Magnetism from conductors and enhanced nonlinear phenomena. *IEEE Trans. Microwave Theory Tech.*, 47(11):2075–2084, 1999.

[155] J. B. Pendry, D. Schurig, and D. R. Smith. Controlling electromagnetic fields. *Science*, 312:1780–1782, 2006.

[156] M. A. Persaud, J .M. Tolchard, and A. I Ferguson. Efficient generation of picosecond pulses at 243 nm. *IEEE J. Quantum Electron.*, 26:1253–1258, 1990.

[157] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory.* Addison-Wesley Publishing Group, Reading, Massachusetts, 1995.

[158] Mark A. Pinsky. *Introduction to Fourier analysis and wavelets.* Brooks/Cole, 2002.

[159] M. L. Povinelli, M. Loncar, M. Ibanescu, E. J. Smythe, S. G. Johnson, F. Capasso, and J. D. Joannopoulos. Evanescent-wave bonding between optical waveguides. *Opt. Lett.*, 30:3042, 2005.

[160] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing.* Cambridge Univ. Press, second edition, 1992.

[161] Minghao Qi, Eleftherios Lidorikis, Peter T. Rakich, Steven G. Johnson, J. D. Joannopoulos, Erich P. Ippen, and Henry I. Smith. A three-dimensional optical photonic crystal with designed point defects. *Nature*, 429:538–542, 2004.

[162] Minghao Qi and H. Smith. Fabrication of 3d layered photonic crystal. private communications, 2003.

[163] Chengwei Qiu, Li Hu, Baile Zhang, Bae-Ian Wu, Steven G. Johnson, and John D. Joannopoulos. Spherical cloaking using nonlinear transformations for improved segmentation into concentric isotropic coatings. *Opt. Express*, 17(16):13467–13478, 2009.

[164] Chengwei Qiu, Li Hu, Baile Zhang, Bae-Ian Wu, Steven G. Johnson, and John D. Joannopoulos. Spherical cloaking using nonlinear transformations for improved segmentation into concentric isotropic coatings. *Optics Express*, 17:13467–13478, 2009.

[165] Marco Rahm, David Shurig, Daniel A. Roberts, Steven A. Cummer, David R. Smith, and John B. Pendry. Extending the bandwidth of electromagnetic cloaks. *Photonics and Nanostructures: Fundamentals and Applications*, 6(1):87–95, 2008.

[166] David Ramirez, Alejandro W. Rodriguez, Hila Hashemi, J.D. Joannopoulos, Marin Soljačić, and Steven G. Johnson. Degenerate four-wave mixing in triply-resonant Kerr cavities. *Physical Review A*, 83:033834, March 2011.

[167] S. M. Rao and N. Balakrishnan. Computational electromagnetics. *Current Science*, 77(10):1343–1347, 1999.

[168] Fang-Fang. Ren, Rui Li, Chen Cheng, and Hui-Tian Wang. Giant enhancement of second harmonic generation in a finite photonic crystal with a single defect and dual-localized modes. *Phys. Rev. B*, 70:245109, 2004.

[169] Alejandro Rodriguez, Marin Soljačić, J. D. Joannopulos, and Steven G. Johnson. $\chi^{(2)}$ and $\chi^{(3)}$ harmonic generation at a critical power in inhomogeneous doubly resonant cavities. *Opt. Express*, 15(12):7303–7318, 2007.

[170] Zhichao Ruan, Min Yan, and Min Qiu. Cylindrical invisibility cloak with simplified material parameters is inherently visible. *Phys. Rev. Lett.*, 99(11):113903, 2007.

[171] C. M. Savage and D. F. Walls. Optical chaos in second-harmonic generation. *Optica Acta*, 30:557–561, 1983.

[172] Luigi Scaccabarozzi, M. M. Fejer, Yijie Huo, Shanhui Fan, Xiaojun Yu, and James S. Harris. Enchanced second-harmonic generation in AlGaAs/Al_xo_y tightly confining waveguides and resonant cavities. *OL*, 31(24):3626–3630, 2006.

[173] Stephan Schiller. *Principles and Applications of Optical Monolithic Total-Internal-Reflection Resonators.* PhD thesis, Stanford University, Stanford, CA, March 1993.

[174] K Schneider, Stephan Schiller, and Jürgen Mlynek. 1.1-w single-frequency 532-nm radiation by second-harmonic generation of a miniature nd:yag ring laser. *Opt. Lett.*, 21:1999–2001, 1996.

[175] D. Schurig, J. J. Mock, B. J. Justice, S. A. Cummer, J. B. Pendry, A. F. Starr, and D. R. Smith. Metamaterial electromagnetic cloak at microwave frequencies. *Science*, 314:977–980, 2006.

[176] D. Schurig, J. B. Pendry, and D. R. Smith. Calculation of material properties and ray tracing in ransformation media. *Opt. Express*, 14(21):9794–9804, 2006.

[177] Pierre Scotto, Pere Colet, and Maxi San Miguel. All-optical image processing with cavity type ii second-harmonic generation. *Opt. Lett.*, 28:1695, 2003.

[178] D. N. Sharp, M. Campbell, E. R. Dedman, M. T. Harrison, R. G. Denning, and A. J. Turberfield. Photonic crystals for the visible spectrum by holographic lithography. *Optical and Quantum Electron.*, 34:3–12, 2002.

[179] N. Shibata, R. P. Braun, and R. G. Waarts. Phase-mismatch dependence of efficiency of wave generation through four-wave mixing in single-mode optical fiber. *IEEE J. Quantum Electron.*, QE-23(7):1205–1210, 1987.

[180] Anthony Siegman. *Lasers.* University Science Books, Mill Valley, CA, 1986.

[181] R. E. Slusher, L. W. Hollberg, B. Yurke, J. C. Mertz, and J. F. Valley. Observation of squeezed states generated by four-wave mixing in an optical cavity. *Phys. Rev. Lett.*, 55:2409–2014, 1985.

[182] D. R. Smith, J. P. Pendry, and M. C. Wiltshire. Metamateirals and negative refractive index. *Science*, 305(5685):788–792, 2004.

[183] R .G. Smith. Theory of intracavity optical second-harmonic generation. *IEEE J. Quantum Electron.*, 6:215–223, 1970.

[184] I. I. Smolyaninov, Y. J. Hung, and C. C. Davis. Two-dimensional metamaterial structure exhibiting reduced visibility at 500nm. *Opt. Lett.*, 33:1342–1344, 2008.

[185] I. I. Smolyaninov, Vera N. Smolyaninova, Alexander V. Kildishev, and Vladimir M. Shalaev. Anisotropic metamaterials emulated by tapered waveguides: Application to optical cloaking. *Phys. Rev. Lett.*, 102:213901, 2009.

[186] Marin Soljačić, Mihai Ibanescu, Steven G. Johnson, Yoel Fink, and J. D. Joannopoulos. Optimal bistable switching in non-linear photonic crystals. *Phys. Rev. E Rapid Commun.*, 66:055601(R), 2002.

[187] B. Z. Steinberg, A. Boag, and R. Lisitsin. Sensitivity analysis of narrowband photonic crystal filters and waveguides to structure variations and inaccuracy. *J. Opt. Soc. Am. A*, 20(1):138–146, 2003.

[188] Julius Adams Stratton. *Electromagnetic Theory.* McGraw-Hill, New York, 1941.

[189] H.C. Strifors and G.C. Gaunaurd. Scattering of electromagnetic pulses by simple-shaped targets with radar cross section modified by a dielectric coating. *Antennas and Propagation, IEEE Transactions on*, 46(9):1252 –1262, sep 1998.

[190] Steven H. Strogatz. *Nonlinear Dynamics and Chaos.* Westview Press, Boulder, CO, 1994.

[191] Michael Tabor. *Chaos and Integrability in Nonlinear Dynamics: An Introduction.* Wiley, New York, 1989.

[192] Allen Taflove and Susan C. Hagness. *Computational Electrodynamics: The Finite-Difference Time-Domain Method.* Artech, Norwood, MA, 2000.

[193] Richard S. Tasgal and Boris A. Band Y. B. Malomed. Gap solitons in a medium with third-harmonic generation. *Phys. Rev. E*, 72:016624, 2005.

[194] Isao Tomita, Masaki Asobe, Hiroyuki Suzuki, Junji Yumoto, and Yuzo Yoshikuni. Broadband quasi-phase-matched second-harmonic generation in a nonlinear photonic crystal. *Journal of Applied Physics*, 100(2):023120, 2006.

[195] Lloyd N. Trefethen and David Bau III. *Numerical Linear Algebra.* SIAM: Society for Industrial and Applied Mathematics, 1997.

[196] R. S. Tucker, P. C. Ku, and C. J. Chang-Hasnain. Delay-bandwidth product and storage density in slow-light optical buffers. *Electron. Lett.*, 41(4):208–209, 2005.

[197] A. C. Turner, M. A. Foster, A. L. Gaeta, and M. Lipson. Ultra-low power parametric frequency conversion in a silicon microring resonator. *OE*, 16:4881–4887, 2008.

[198] Kerry J. Vahala. Optical microcavities. *Nature*, 424:839–846, 2003.

[199] Jason Valentine, Jensen Li, Thomas Zentgraf, Guy Bartal, and Xiang Zhang. An optical cloak made of dielectrics. *Nature Materials*, 8(7):568–571, 2009.

[200] A. J. Ward and J. B. Pendry. Refraction and geometry in maxwell's equations. *J. Mod. Opt.*, 43(4):773–793, 1996.

[201] S. Watanabe, T. Naito, and T. Chikama. Compensation of chromatic dispersion in a single-mode fiber by optical-phase conjugation. *PTL*, 5:92–95, 1993.

[202] M. T. Weiss. Quantum derivation of energy relations analogous to those for nonlinear reactances. *Proc. IRE*, 45:1012–1013, 1957.

[203] Ling-An Wu, Min Xiao, and H. J. Kimble. Squeezed states of light from an optical parametric oscillator. *JOSA-B*, 4:1465–1476, 1987.

[204] Qianfan Xu and Michal Lipson. Carrier-induced optical bistability in silicon ring resonators. *Opt. Lett.*, 31(3):341–343, 2005.

[205] Xiaofei Xu, Yijun Feng, Yu Hao, Juming Zhao, and Tian Jiang. Infrared carpet cloak designed with uniform silicon grating structure. *Appl. Phys. Lett.*, 95(18), 2009.

[206] Min Yan, Zhichao Ruan, and Min Qiu. Cylindrical invisibility cloak with simplified material parameters is inherently visible. *Phys. Rev. Lett.*, 99(23):233901, 2007.

[207] Zhenshan Yang, Philip Chak, Alan D. Bristow, Henry M. van Driel, Rajiv Iyer, J. Stewart Aitchison, Arthur L. Smirl, and J. E. Sipe. Enhanced second-harmonic generation in algaas microring resonators. *Opt. Lett.*, 32(7):826–828, Apr 2007.

[208] Mehmet F. Yanik and Shanhui Fan. Dynamic photonic structures: Stopping, storage, and time reversal of light. *Studies in Applied Mathematics*, 115:233–254, 2005.

[209] Mehmet F. Yanik, Shanhui Fan, Marin Soljačić, , J. D. Joannopoulos, and Yanik. All-optical transistor action with bistable switching in a photonic crystal cross-waveguide geometry. *Opt. Lett.*, 68:2506, 2004.

[210] A. Yariv. Compensation for channel dispersion by nonlinear optical-phase-conjugation. *Opt. Lett.*, 4:52–54, 1978.

[211] H. P. Yuen and J. H. Shapiro. Generation and detection of 2-photon coherent states in degenerate four-wave mixing. *Opt. Lett.*, 4:334–336, 1979.

[212] B. L. Zhang, H. S. Chen, and B.-I. Wu. Practical limitations of an invisibility cloak. *Progress In Electromagnetics Research*, 97:407–416, 2009.

[213] Baile Zhang. *Study on Transformation-Based Invisibility Cloaks.* PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2009.

[214] Baile Zhang, Tucker Chan, and Bae-Ian Wu. Lateral shift makes ground-plane cloak detectable. *Phys. Rev. Lett.*, 104(233903), 2010.

[215] Baile Zhang, Yuan Luo, Xiaogang Liu, and George Barbastathis. Macroscopic invisibility cloak for visible light. *Phys. Rev. Lett.*, 106(3):033901, 2011.

[216] Baile Zhang and Bae-Ian Wu. Cylindrical cloaking at oblique incidence with optimized finite multilayer parameters. *Opt. Express*, 35(16):2681, 2010.

[217] Pu Zhang, Yi Jin, and Sailing He. Cloaking an object on a dielectric half-space. *Opt. Express*, 16(5):3161–3166, 2008.

[218] Pu Zhang, Yi Jin, and Sailing He. Obtaining a nonsingular two-dimensional cloak of complex shape from a perfect three-dimensional cloak. *Appl. Phys. Lett.*, 93:243502, 2008.

[219] Yu Zhu and Andreas C. Cangellaris. *Multigrid Finite Element Methods for Electromagnetic Field Modelling.* John Wiley and Sons, Hooboke, NJ, 2006.

[220] Ferederic Zolla, Sebastien Guenneau, Andre Nicolet, and J. B. Pendry. Electromagnetic analysis of cylindrical invisibility cloaks and mirage effect. *Opt. Lett.*, 32(9):1069–1071, 2007.

[221] I I. Zootoverkh, Kravtsov N V., and E G Lariontsev. Enhancement of the efficiency of second-harmonic generation in microlaser. *Quantum Electronics*, 30:565, 2000.