Dissipation in Deforming Chaotic Billiards

A thesis presented
by

Alexander Harvey Barnett

to
The Department of Physics
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of

Physics

Harvard University
Cambridge, Massachusetts
October 2000

Thesis advisor
Eric J. Heller

Author
Alexander Harvey Barnett

Dissipation in Deforming Chaotic Billiards

Abstract

Chaotic billiards (hard-walled cavities) in two or more dimensions are paradigm systems in the fields of classical and quantum chaos. We study the dissipation (irreversible heating) rate in such billiard systems due to general shape deformations which are periodic in time. We are motivated by older studies of one-body nuclear dissipation and by anticipated mesoscopic applications. We review the classical and quantum linear response theories of dissipation rate and demonstrate their correspondence in the semiclassical limit. In both pictures, heating is a result of stochastic energy spreading. The heating rate can be expressed as a frequency-dependent friction coefficient $\mu(\omega)$, which depends on billiard shape and deformation choice. We show that there is a special class of deformations for which μ vanishes as like a power law in the small-ω limit. Namely, for deformations which cause translations and dilations $\mu \sim \omega^4$ whereas for those which cause rotations $\mu \sim \omega^2$. This contrasts the generic case for which $\mu \sim \omega^0$. We show how a systematic treatment of this special class leads to an improved version of the 'wall formula' estimate for $\mu(0)$.

We show that the special nature of dilation (a new result) is semiclassically equivalent to a quasi-orthogonality relation between the (undeformed) billiard quantum eigenstates on the boundary. This quasi-orthogonality forms the heart of a 'scaling method' for the numerical calculation of quantum eigenstates, invented recently by Vergini and Saraceno. The scaling method is orders of magnitude more efficient than any other known billiard quantization method, however an adequate explanation for its success has been lacking until now. We explain the scaling method, its errors, and applications. We also present improvements to Heller's plane wave method.

Two smaller projects conclude the thesis. Firstly, we give a new formalism for quantum point contact (QPC) conductance in terms of scattering cross-section in the half-plane, of use in open mesoscopic and atomic systems. We derive the maximum conductance through a tunneling QPC coupled to a resonator. Secondly, we numerically model a novel design of coherent atom waveguide which uses the dipole force due to evanescent light fields leaking from an optical waveguide mounted on a substrate.

Contents

Citations to Previously Published Work

Large portions of Chapters 3 and 4, as well as some of Sections 2.3 and 6.1.2 have appeared in the following two papers:

> "Deformations and dilations of chaotic billiards: dissipation rate, and quasi-orthogonality of the boundary wavefunctions", A. H. Barnett, D. Cohen, and E. J. Heller, Phys. Rev. Lett. **85**, 1412 (2000), `nlin.CD/0003018`;
>
> "Rate of energy absorption for a driven chaotic cavity", A. H. Barnett, D. Cohen, and E. J. Heller, submitted to J. Phys. A, `nlin.CD/0006041`.

The numerical methods of Chapter 6 were used to calculate data appearing in the above papers and in the following:

> "Parametric evolution for a deformed cavity", D. Cohen, A. H. Barnett, W. Bies, and E. J. Heller, submitted to Phys. Rev. E, `nlin.CD/0008040`.

Chapter 7 appears in its entirety as

> "Mesoscopic scattering in the half-plane: how much conductance can you squeeze through a small hole?", A. H. Barnett, M. Blaauboer, A. Mody, and E. J. Heller, submitted to Phys. Rev. B, `cond-mat/0008279`.

Finally, most of Chapter 8 has been published as

> "Substrate-based atom waveguide using guided two-color evanescent light fields", A. H. Barnett, S. P. Smith, M. Olshanii, K. S. Johnson, A. W. Adams, M. Prentiss, Phys. Rev. A **61**, 023608 (2000), `physics/9907014`.

Electronic preprints (shown in `typewriter font`) are available on the Internet at the following URL:

> `http://arXiv.org`

Acknowledgments

Completing this doctoral work has been a wonderful and often overwhelming experience. It is hard to know whether it has been grappling with the physics itself which has been the real learning experience, or grappling with how to write a paper, give a coherent talk, work in a group, teach section, code intelligibly, recover a crashed hard drive, stay up until the birds start singing, and... stay, um... focussed.

I have been very privileged to have undoubtedly the most intuitive, smart and supportive advisor anyone could ask for, namely Rick Heller. Ever since I learned from him what an avoided crossing was (animated in full colour, naturally), I have been stimulated and excited by his constant flow of good ideas. Rick has an ability to cut through reams of algebra with a single visual explanation that I will always admire, and I have learned a great deal of physics from him. He has fostered certainly the most open, friendly, collaborative and least competitive research group in the theory wing of this Department. He has also known when (and how) to give me a little push in the forward direction when I needed it.

Throughout my six years, I was supported for many semesters by the National Science Foundation, and the Institute for Theoretical Atomic and Molecular Physics (ITAMP), through the generosity of my advisor. During my first year I thank the Kennedy Memorial Trust in London for providing my funding.

Rick's other students and post-docs, both past and present, comprise a superb research group. The ability to bounce ideas off so many excellent minds has been priceless. My most intense collaboration has been with Doron Cohen, whose clarity, persistence, ability to create new models, and ability to write a new publication every month, has taught me a lot. (I am still working on the publication per month aspect.) Michael Haggerty has similiarly influenced me, and I think I can safely say that everyone in the group has benefitted from his generosity and collaboration. In my first couple of years in the Heller group, Adam Lupu-Sax and I swapped many thoughts on numerical methods, while Lev Kaplan prevented me from getting too scarred by the new realm of quantum chaos. Jesse Hersch shared many good times and wacky science experiments with me, and introduced me to a world of electrostatically-levitated chalk particles. Areez Mody has been a companion and goldmine of mathematical knowledge since the year we arrived together. Scot Shaw has been my numerical methods partner-in-crime. Bill Bies was my buddy through thesis formatting hell. It has also been my pleasure to work with (and hang out with) Jiří Vaníček, Axel André, Miriam Blaauboer, Ragnar Fleischmann, Greg Fiete, Allison Kalben, Natasha Lepore, Michael Efroimsky (and his stories), Manny Tannenbaum (and his accents), Sheng Li and Stephan Filipov; also from the earlier days Maurizio Carioli, Allan Tameshtit, Jonathan Edwards, Kazuo Hirai (and his valves), Bill Hoston, Neepa Maitra, Martin Naraschewski and Steve Tomsovic (and his jazz).

Our group's secretary Carol Davis is surely the kindest, coolest and most witty person one could possibly hope to spend a lunch-break gossiping with. She has also helped me out many a time during those pre-presentation panics. I must also thank Mary Lampros for her unique blend of caring and total organisation. She will be sorely missed in this Department.

On the foreign front, I must thank Eduardo Vergini for fascinating email exchanges, and guiding me towards an understanding of the incredible numerical method he invented along with Marcos Saraceno. From the group of Mara Prentiss, Steve Smith and Maxim Ol'shanii were always there to fill me in on the practice and theory of atom optics.

There are countless others who have been there for me throughout my time as a

graduate student. From the world of physics, Meredith Betterton, Joseph Thywissen and Rosalba Perna stand out as both good, caring friends and providers of fascinating physics conversations. From other worlds, Erika Evasdottir and my housemates Vitaly Napadow and John Iversen have been wonderful influences and friends; they have lead me from deconstruction to drumming.

My fascination with the physical world is undoubtedly due to the influence of my father, Ross. He showed me the joys of high-dispersion prisms in sunlight at an early age, and let me spend most of my childhood building Tesla coils and playing with BASIC on our 6502 home computer. Even now, he proofreads my papers. My late mother Pat taught me a love of language, of art, of teaching, and of signs, and it is sad to realise that she cannot be present for this moment in my life. The same realisation holds true for my grandfather Oliver, who is highly responsible for my tinkering nature. My courageous sister Jess has kept me sane during my hardest times via the transatlantic phone system.

Finally, Liz Canner has been my guiding light and love over these last two years. She has seen my best and my worst, and provided support, hugs, and taken me places I never imagined. Even when my emotional and research brains became so hopelessly entwined that I dreamt that the two of us were overlapping eigenfunctions, she still loved me. And she even thought it was cute.

Dedicated to my father Ross,
my late mother Pat,
and my sister Jess.

Chapter 1

Introduction and summary

Structure of this thesis

This thesis falls naturally into four parts, which are relatively independent:

- Study of dissipation rate in deformed chaotic billiards (Chapters 2, 3 and 4),
- Improved numerical methods for quantization of billiards (Chapters 5 and 6),
- Half-plane scattering approach to mesoscopic conductance (Chapter 7), and
- Design of an atom waveguide using two-color evanescent light fields (Chapter 8).

The first two parts form the main body of the thesis, and they are both devoted to the study of billiard systems (hard-walled cavities enclosing a region of free space) in which the classical motion is chaotic. The quantum mechanics of such systems has become known as the field of 'quantum chaos'. The first part probably contains the most significant new physical results; this is reflected in the choice of thesis title. The second part can be viewed merely as a description of numerical quantum-mechanical calculations that play a supporting role in the first part. However, there will also turn out to be a surprising reciprocal connection, namely that results from the first part will provide a much-needed explanation for the success of a very efficient numerical technique in the second part. The intertwining of these two subject areas had turned out to be one of the most beautiful surprises in this body of research.

The third and fourth parts form essentially separate projects, and can therefore be read independently. However they do share with the rest of the thesis the common theme of wave mechanics: the third presents a new approach to the transport of quasiparticle waves in mesoscopic systems, and the fourth models confined electromagnetic waves to trap and guide atoms (which themselves can be treated as coherent matter waves).

The goals and subject matter of the four parts are sufficiently different to merit individual introductions and summaries, which now follow without further ado.

Chapters 2,3 and 4: Dissipation rate and deformations of chaotic billiards

The dynamics of a particle inside a cavity (billiard) in $d = 2$ or 3 dimensions is a major theme in studies of classical and quantum chaos [154, 91, 25]. Whereas the physics of

time-independent chaotic systems is extensively explored, less is known about the physics when such a system is 'driven' (time-dependent chaotic Hamiltonian). The main exceptions are the studies of the kicked rotator and related systems [75]. However, the rotator (with no kicks) is a $d = 1$ integrable system, whereas we are interested in chaotic ($d \geq 2$) cavities.

Driven cavities have been of special interest since the 1970s in studies of the so-called 'one-body' dissipation rate in vibrating nuclei [29, 120, 118, 106, 107]. A renewed interest in this problem is anticipated in the field of mesoscopic physics. Quantum dots [20, 65] can be regarded as small 2D cavities whose shape is controlled by electrical gates. Quasiparticle motion inside the dot can have long coherence (dephasing) times, and enable the semiclassical regime to be approached (many wavelengths across the system).

In Chapter 2 I give tutorial review of the theory of dissipation in general driven ergodic systems, which is quite a young field. The Hamiltonian is controlled by a single parameter x, whose time-dependence will be $x(t) = A\sin(\omega t)$ where A is the amplitude and ω is the driving frequency. In both the classical (Section 2.1) and quantum-mechanical (Section 2.2) pictures, dissipation is a result of *stochastic energy spreading*. Once this spreading is established, the pictures can be unified [46]. Irreversible growth of energy (heating) is then a result of biased diffusion (a random walk) in energy. I will confine myself to a regime where linear response theory (LRT) is valid. In the quantum case this is known as the Kubo formalism, although the language of the energy spreading picture appears different (I connect the two pictures in Section 2.2.5). The heating rate is given by

$$\frac{d}{dt}\langle \mathcal{H} \rangle \;=\; \mu(\omega) \cdot \tfrac{1}{2}(A\omega)^2, \tag{1.1}$$

where the 'friciton coefficient' $\mu(\omega)$ is related linearly[1] to $\tilde{C}_{\rm E}(\omega)$, the correlation spectrum of the time-dependent fluctuating quantity $-\partial\mathcal{H}/\partial x$. The latter (a generalized 'force' on the parameter x) specifies the random 'kicks' up or down in energy that particles receive. In the quantum case $\tilde{C}_{\rm E}(\omega)$ is the 'band profile' (average off-diagonal shape) of the matrix $\partial\mathcal{H}/\partial x$ in the energy representation. In Section 2.3 I demonstrate theoretically and numerically the semiclassical equivalence of the classical and quantum versions of $\tilde{C}_{\rm E}(\omega)$, and examine quantum effects beyond the band profile.

In Chapters 3 and 4 we specialize to a system of non-interacting particles inside a billiard whose walls are deformed by the parameter x. One can specify any deformation by a function $D(\mathbf{s})$, where $\mathbf{s}$ specifies location of a wall element on the boundary (surface) of the billiard, and $D(\mathbf{s})x$ is the resulting normal displacement of this wall element. We will be interested in low-frequency driving, meaning $\omega \ll 1/\tau_{\rm bl}$ the mean collision frequency. In the $\omega \to 0$ limit (uniform parameter velocity $\dot{x}$), the heating rate (1.1) is given by the dc friction $\mu(0)$ proportional to $\nu_{\rm E} = \tilde{C}_{\rm E}(\omega \to 0)$. An assumption of uncorrelated collisions (the white noise approximation or WNA) gives an estimate for $\nu_{\rm E}$, which in turn in $d = 3$ leads to the well-known 'wall formula' [29] (from the nuclear application),

$$\mu_{\rm E} \;=\; \frac{N}{\mathsf{V}} m v_{\rm E} \oint D(\mathbf{s})^2 d\mathbf{s}. \tag{1.2}$$

This (ω-independent) estimate of the friction applies to a microcanonical ensemble of N particles with speed $v_{\rm E}$ in a billiard volume V.

[1] The relation depends upon the initial energy distribution.

We analyze $\mu(\omega)$ numerically in 2D billiard shapes (generalized Sinai, and Bunimovich stadium). We believe this is the first study of frequency-dependent heating rate in billiard systems. The chief discovery (Section 3.3) is a class of deformations whose heating rate vanishes in the $\omega \to 0$ limit, like a power-law $\tilde{C}_{\rm E}(\omega) \sim \omega^{\gamma}$. This holds even for billiards with strong chaos, and goes completely against the WNA prediction. The class of 'special' deformations turns out to be just the class which preserves the billiard shape. For translations and dilations $\gamma = 4$ and for rotations $\gamma = 2$. This is to be compared to the case of a generic deformation, for which $\gamma = 0$ as the WNA prediction would predict. We give classical explanations for the power-laws (which rely on correlation on short timescales $\sim \tau_{\rm bl}$). Importantly, the special class is manifested in the quantum band profile too. Thus the special nature of dilation, believed to be new in the literature, corresponds to a quasi-orthogonality relation between eigenstates on the boundary, which in turn will be the key to the powerful numerical method of Chapter 6. We also discuss (Section 3.4) non-generic shape-dependent effects (such as marginally-stable orbits) which may alter the power-laws given above.

The goal of Chapter 4 is as follows: given a general deformation $D(\mathbf{s})$, in a given billiard shape, we seek an analytical estimate of $\nu_{\rm E}$ (and hence $\mu(0)$). It is an exact result that $\nu_{\rm E}$ is a quadratic form in the function space of $D(\mathbf{s})$. The WNA fails to take into account that $\nu_{\rm E}$ vanishes for special deformations (which form a linear subspace in $D(\mathbf{s})$). We show that how it is possible to systematically subtract (project out) the 'special' components of a general $D(\mathbf{s})$. Applying the WNA only to the remaining ('normal') component gives an improved estimate of $\nu_{\rm E}$. We analytically and numerically justify this projection procedure, and test the quality of the improved formula. The quality is limited by that of the WNA estimate of the 'normal' component, which relies on the assumption of strongly chaotic motion. However, in the generalized Sinai billiard the improved formula is found to perform much better than the original WNA.

Our work replaces all *ad hoc* corrections which had been introduced [29] to account for the intuitive result that translations and rotations cause no heating at $\omega = 0$. We thereby clear up some inconsistent habits in the nuclear community (Section 4.3.2). More significantly, the incorporation of the special nature of dilation is entirely new.

Note that the effect of interaction between the particles is not considered. If the mean free path for inter-particle collisions is large compared with the size of the cavity, then we expect that our analysis still applies. (However if the mean free path is much smaller, then we get into the hydrodynamic regime, where viscosity becomes the dominant dissipative effect).

This work was performed in collaboration with Doron Cohen. I also benefitted greatly from use of a classical billiard trajectory code written by Michael Haggerty.

Chapters 5 and 6: Improved billiard quantization methods

The rapid development of electronic computing machines in the last 50 years has had an impact on scientific research whose size is hard to grasp[2]. The interplay between

[2]Some idea of the rate of progress of this technology can be gained from such quotes as "Such a machine in the hands of a competent operator can produce 400 full-length products or 1,000 sums during an 8-hour working day"[42], referring to an electromechanical desk calculator typical for scientific use in the 1950s.

numerical simulations and theoretical models now plays a crucial role in most areas of physics, chemistry, engineering, and other sciences. However, this impact would have been be drastically reduced were it not for the parallel development of efficient numerical algorithms. For instance, the invention of two techniques alone—the diagonalization of dense matrices[81], and the Fast Fourier Transform[161]—has allowed scientists to handle hitherto undreamed-of problems on a daily basis.

Quantum chaos [91] is no exception: it has relied heavily on numerical solutions almost since its inception (as did its forebear, classical chaotic dynamics[154]). Billiards in $d = 2$ (or sometimes 3) dimensions have been popular systems for study because use of the free-space Green's function allows formulation as a boundary problem. Thus quantum eigenstates can be calculated at much higher energy than with the traditional (*e.g.* finite element) methods which cover the entire domain. High energies are so sought-after because most of the theoretical predictions involve the semiclassical limit $\hbar \to 0$. In this part of the thesis I will present new and efficient methods for finding these high-energy eigenstates.

The time-independent Schrodinger's equation in such a system is the Helmholtz wave equation,

$$(\nabla^2 + k^2)\psi(\mathbf{r}) \;=\; 0, \tag{1.3}$$

with certain boundary conditions. This problem is common to many other areas of physics and engineering (mesoscopic devices, acoustics, elastodynamics of thin plates, scalar electromagnetics and optics), and there has been some (but not that much) exchange of ideas between those communities and that of quantum chaos. If the boundary conditions are open then we have a scattering problem; if closed, an eigenvalue problem. I will be concerned only with the latter. In Section 5.1 I present a review, and categorize solution methods dependent on whether the basis does not (Class A) or does (Class B) depend on energy. Only Class B allows formulation as a boundary problem. The pioneering chaotic eigenstate studies of McDonald and Kaufman [145] and Berry and Wilkinson [28] used the Boundary Integral Method [36, 121] (BIM or BEM), while those of Heller [90, 91] used the Plane Wave Decomposition Method (PWDM, an original technique). Semiclassical quantization methods have also been developed based on boundary matching [23] or the surface of section [34]. All these methods are Class B, and all require an expensive search ('sweep' or 'hunt') in energy-space for zero-determinants of a matrix.

In Chapter 5 I present an original Class B sweep method which is a simplified version of Heller's PWDM. The problems of missing states and sensitivity to basis size choice and matching point density have been solved, and the efficieny increased. The coefficient vector $\mathbf{x}$ of the nearest eigenfunction to a given wavenumber k is given by the largest-eigenvalue (λ_1) solution to

$$[G(k) - \lambda F(k)]\,\mathbf{x} \;=\; \mathbf{0}, \tag{1.4}$$

where the matrix G takes the norm in the domain, and F takes the norm of the boundary condition error (the 'tension'). I show that G can be expressed entirely on the boundary, and discuss improved 'hunt' methods for zeros of λ_1^{-1} (which give the desired eigenwavenumbers k). However, a few diagonalizations of (1.4) are still required per state found.

The diagonalization of an 8-by-8 matrix was a weekend-long task[125].

Of much more significance is Chapter 6. Here I analyse the 'scaling method' of Vergini and Saraceno [195, 194], which despite being little-understood and little-used, is without doubt the most significant advance in numerical billiard quantization in the last 15 years. Eigenstates are given by the large-λ solutions of

$$\left[\frac{dF}{dk}(k) - \lambda F(k)\right] \mathbf{x} \;=\; \mathbf{0}, \tag{1.5}$$

however, through a certain choice of boundary weighting function an amazing property of F and dF/dk emerges: they are *quasi-diagonal* (have very small off-diagonal elements) in a basis of the exact eigenfunctions rescaled to all have the same wavenumber k. This allows (1.5) to return up to $N/10$ useful eigenfunctions for a *single* diagonalization, and entirely eliminates the need for 'hunt' procedures. Here N is the matrix size (semiclassical basis size). The relative efficiency over sweep methods is $\sim 10^3$ when there are several hundred wavelengths across the system, and moreover, increases further with increasing k and dimension d! Remarkably, no adequate explanation of the key quasi-diagonality property has been known until now. I give, for the first time, a semiclassical explanation in terms of the 'special' nature of the dilation deformation (from Chapter 3). I also correct errors in the original authors' derivation [195, 194] of higher-order tension terms.

Both chapters are presented as a practical 'how-to' guide to the diagonalization of d-dimensional billiards, and I hope they may be of use to other communities who solve the Helmholtz eigenproblem. I thoroughly analyse the various types of error in both the sweep and scaling methods, compare results from the two, and discuss the use of real and evanescent plane wave basis sets. For illustration, I use Bunimovich's stadium billiard (a shape known to be classically-chaotic [40]), in which evanescent basis sets have been pioneered by Vergini[194]. Currently the scaling method applies only to Dirichlet boundary conditions. Adequate basis sets for more general shapes is an area in dire need of future research. My work has involved deriving a collection of useful new formulae for boundary evaluation of domain integrals of Helmholtz solutions: these are presented in Appendix H.

Two applications of the scaling method are presented in this thesis. The first is the quantum band-profile calculations for Chapters 2–4. The second is an efficient evaluation of overlaps of eigenstates of a billiard with eigenstates of the same billiard deformed by various finite amounts (Section 6.4). The profiles of the resulting matrices can be viewed as local densities of states ('line shapes'), which are analysed in our publication [48]. The diagonalization of the deformed stadium billiard is believed to be new.

During this work, I have benefitted much from fruitful exchanges with Eduardo Vergini. I must thank Doron Cohen for first alerting me to the semiclassical estimation of the band profile of matrix elements on the boundary. Finally, Appendix H resulted from collaboration with Michael Haggerty.

Chapter 7: Quantum point contact conductance and scattering in the half-plane

This third part continues the theme of wave mechanics of non-interacting particles. However attention shifts from closed to open systems, namely the transport of quasiparticles (in a 2D electron gas) though a general two-terminal mesoscopic electronic device, or 'quantum point contact' (QPC) [20, 65].

We model the conductance of a QPC, in linear response. If the QPC is highly non-adiabatic or near to scatterers in the open reservoir regions, then the usual distinction between 'leads' and 'reservoirs' breaks down. This situation arises in the recent experimental work of Katine [112] and Topinka [192] in the Westervelt group here at Harvard, where open resonant and scattering geometries were studied. In such systems the Landauer formula [129, 130, 41, 55, 65] for the conductance (including spin degeneracy),

$$G \;=\; \frac{2e^2}{h}\mathrm{Tr}(t^\dagger t), \tag{1.6}$$

is no longer convenient because no conventional transverse 'lead' states exist between which to define the transmission matrix t. Rather, a technique based on scattering theory in the two-dimensional infinite half-plane is appropriate. We relate conductance to transmission *cross section*, defined as an effective collision size on the reflective boundary of a half-plane (reservoir) region. We also introduce a new half-plane radial basis of 'lead' states in which the usual Landauer formula is recovered.

The relation between the Landauer and the half-plane scattering formalism is expressed by

$$\int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\mathrm{T}}(k,\phi) \;=\; \lambda \mathrm{Tr}(t^\dagger t), \tag{1.7}$$

where $\sigma_{\mathrm{T}}(k,\phi)$ is the angle-dependent transmission cross section and λ the Fermi wavelength, which I derive for both hard-walled and soft-walled reflective potential barriers.

We analyse an idealized, highly non-adiabatic slit QPC system in the extreme quantum, intermediate, and semiclassical regimes. We derive the counterintuitive result (first due to Heller) that an arbitrarily small (tunneling) QPC can reach a p-wave channel conductance of $2e^2/h$ when coupled to a suitable resonant cavity. We also find that if two or more resonances coincide, the total conductance can in theory reach multiples of this value.

This leads to some thought-experiments on attempting to overcome the maximum conductance $2e^2/h$ per quantum channel. We also discuss reciprocity (left-to-right symmetry) of conductance, and the possibility of its breakdown in a proposed QPC (which could exhibit 'conductance' quantization) for atom waves[191]. We emphasizes the importance of the *thermal occupation* of states in phase space (as is usual in 2D electron systems), for reciprocity to exist. An analogous atomic QPC in 3D need not have this thermal occupation, thus in this system reciprocity can be broken.

This work has been in collaboration with Areez Mody and Miriam Blaauboer, and at the earlier stages many contributions were made by Adam Lupu-Sax. Joseph Thywissen, and professors Charlie Marcus and Daniel Fisher also contributed via stimulating discussions.

Chapter 8: Waveguides for neutral atoms using evanescent light fields

The fourth and final part is a self-contained proposal for a new design of coherent atom waveguide, using the forces exerted on atoms by near-resonant laser fields. As an independent project, it does not connect directly with other theoretical work in this thesis.

However it shares many common themes: the motion of atoms in the trapping potential is a 2D quantum bound mode problem (in a smooth potential), and the optical waveguide bound mode calculation is also similar to this same problem (dielectric constant playing the role of a negative potential). Because the optical 'potential' is not hard-walled, and because the fields are vector rather than scalar, the efficient methods of Chapter 5 or 6 do not apply; rather Finite Elements[56] will be used.

There has been much recent progress in the trapping and cooling of neutral atoms, opening up new areas of ultra-low energy and matter-wave physics [45, 51, 159]. Waveguides for such atoms are of great interest for atom optics, atom interferometery, and atom lithography. Multimode atom waveguides act as incoherent atom pipes that could trap atoms, transport them along complicated paths or between different environments, or deliver highly localized atom beams to a surface. Single-mode waveguides (or multimode guides populated only by atoms in the transverse ground-state) could be used for coherent atom optics and interferometry [151, 2], as well as a tool for one-dimensional physics such as boson-fermion duality [190, 152, 44] and low-dimensional Bose-Einstein condensation effects [115].

In Chapter 8 we propose a dipole-force linear waveguide which confines neutral atoms up to $\lambda/2$ above a microfabricated single-mode dielectric optical guide. The optical guide carries far blue-detuned light in the horizontally-polarized TE mode and far red-detuned light in the vertically-polarized TM mode, with both modes close to optical cut-off. A trapping minimum in the transverse plane is formed above the optical guide due to the differing evanescent decay lengths of the two modes. This design allows manufacture of mechanically stable atom-optical elements on a substrate.

We find that a rectangular optical guide of $0.8\,\mu$m by $0.2\,\mu$m carrying 6 mW of total laser power (detuning ± 15 nm about the D2 line) gives a trap depth of $200\,\mu$K for cesium atoms ($m_F = 0$), transverse oscillation frequencies of $f_x = 40$ kHz and $f_y = 160$ kHz, collection area $\sim 1\,\mu\text{m}^2$ and coherence time of 9 ms. The laser powers required are orders of magnitude less than those commonly needed for dipole traps. The large tranverse frequencies achieved allow the possibility of atomic single-mode occupation (hence coherent guiding) when fed from a source at cesium MOT temperature ($\approx 3\mu K$). We present design equations allowing optimal parameter choices to be made. We also discuss the effects of non-zero m_F, the D1 line, surface interactions, heating rate, the substrate refractive index n_s, and the limits on waveguide bending radius. It emerges that lowering n_s is the main goal if large trap depths are desired of order an optical wavelength from the guide surface.

As known in the engineering community, the optical bound mode problem is notoriously hard [56]. We calculate the full vector bound modes for an arbitrary guide shape using two-dimensional non-uniform finite elements in the frequency-domain, allowing us to optimize atom waveguide properties. We chose rectangular guide cross-sections for this optimization, for simplicity. There are many other shapes possible; the fabrication technique will be the determining factor on what is practical.

This work on atom waveguides, an admittedly far-fetched topic for a student of Rick Heller, was in fact a collaboration with the following members of the Prentiss Group: Steve Smith (my principal collaborator), Maxim Olshanii, Kent Johnson, Allan Adams (who introduced me to the problem), and Mara Prentiss. I also benefitted from discussions with Joseph Thywissen and Yilong Lu.

Chapter 2

Quantum and classical theories of ergodic dissipation

I review the linear response theories of quantum and classical dissipation in general systems. Notation will be introduced which is necessary for fully appreciating the three chapters which follow, in particular the significance of the correlation spectrum ('band profile') $\tilde{C}_{\mathrm{E}}(\omega)$ for finding heating rate. In the final part I discuss and demonstrate correspondence between the two.

A system is 'driven' by changing a parameter on which the Hamiltonian depends, in some prescribed function of time—the parameter is not a dynamical variable, unlike in some of the original studies. The Hamiltonian is assumed to be classically completely chaotic for all parameter values reached (this contrasts the nuclear dynamics example [118, 29, 158, 149]). Dissipation is defined as irreversible growth of average energy. Rather than requiring a 'bath' (for instance, harmonic oscillators in the quantum case [65], or the surrounding gas in the case of classical Brownian motion) which has infinite degrees of freedom, dissipation is a result of ergodicity in closed system with a small number of degrees of freedom (as little as 2). Dissipation can be thought of as a 'friction force' felt by the agency responsible for changing the parameter, which arises from ergodic motion in the system.

The theory of dissipation in an ergodic system has generally taken many forms appropriate to various subfields (for instance, nuclear dynamics [118]). Much of what I present in this chapter relates to the work of Cohen[46, 47], who has attempted to unify the field, outline a general theory of dissipation and its various regimes of applicability. In the picture of Cohen the quantum and classical languages are intertwined. The first two parts of this chapter (and the associated Appendix A) can be read as a tutorial introduction to this field. For this reason I will present the quantum and classical cases in separate languages, which only connect once the existence of *stochastic energy spreading* has been established. Therefore the following classical and quantum reviews can be read somewhat independently of one another, although the connections are many.

2.1 Review of classical dissipation in a general system

Here I review the mechanism for dissipation in a classical ergodic system whose Hamiltonian is dependent on an external parameter $x(t)$. An example model to keep in mind is that of a gas particle trapped inside a deformable chaotic cavity, where heating of the particle (increase in expectation value of the particle kinetic energy) can occur due to time-dependent deformation (driving). Note that the cavity wall is "cold", *i.e.* there are no thermal fluctuations in its velocity. Heating rate in the cavity case will be discussed in detail in Chapters 3 and 4.

The heating has two components: I will only be concerned with the *irreversible* part, which we will see is due to broadening of the distribution in energy. The *reversible* part, which vanishes for any cyclic parameter change, is due to solely to changes of phase space volume and corresponds to adiabatic heating found in the usual thermodynamic treatment of a classical gas. In terms of the *generalized force* acting on the parameter, the former corresponds to a friction force ($\propto \dot{x}^2$), and the latter an elastic force (gradient of a conservative potential). Note that the notion of irreversibility relies on the assumption that you are not allowed to "look" at the particle in order to decide how to vary the deformation (clearly in that case, any heating is 'reversible').

As is usual in classical statistical mechanics, a 'state' of the system will mean a distribution in phase space. The expectation of quantity is then defined over this distribution (this will sometimes be referred to as an 'ensemble average'). It is simplest to derive the energy spreading rate assuming an initial distribution which is microcanonical (uniform in a single energy shell). I will derive this for the case of the parameter changing slowly and non-periodically in time with velocity $\dot{x} = V$, and in the case of sinusoidal time-dependence $x(t) = A\sin(\omega t)$. Energy spreading leads to an energy diffusion equation. The heating rate can then be found for arbitrary initial distribution.

This field is quite young: Koonin and Randrup [120] first derived the dissipation rate in the context of one-body nuclear viscous forces, in 1977, using classical linear response theory. My presentation generally follows that of the wave of activity since 1990, chiefly Wilkinson [201] (following on the work of Ott [153]), Jarzynski [106, 107], and Cohen [47, 46]. I will not consider the interesting case of more than one parameter, in which Berry's phase effects arise (see [27]).

2.1.1 Adiabatic invariance of phase-space volume

For a d-dimensional system with Hamiltonian $\mathcal{H}(\mathbf{q}, \mathbf{p}; x)$, at a given fixed parameter value x the surfaces of $\mathcal{H} = E$ form shells in phase-space $(\mathbf{q}, \mathbf{p})$. At an energy E, the phase-space volume enclosed by the shell is given by $\Omega(E; x)$, and the *weight* of the shell

$$g(E; x) \equiv \frac{d\Omega}{dE}(E; x) = \int d\mathbf{q}d\mathbf{p}\ \delta(\mathcal{H}(\mathbf{q}, \mathbf{p}; x) - E) \tag{2.1}$$

can be thought of as the surface integral of the shell 'thickness' $1/|\nabla\mathcal{H}|$. Note that to connect with quantum mechanics, phase space volume could be measured in Planck-cell units ($d\mathbf{q}d\mathbf{p} \to d\mathbf{q}d\mathbf{p}/(2\pi\hbar)^d$), whereupon $g(E)$ is just the density of states. Following Jarzynski, I will not use these units here.

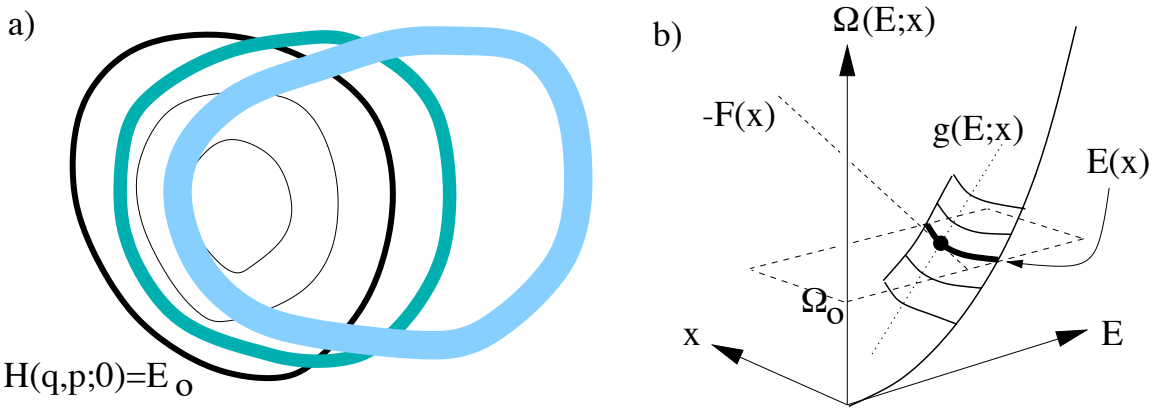

Figure 2.1: a) General picture: constant-energy contours in phase-space at initial time, and subsequent phase-space spreading about a moving energy shell. b) Invariance of phase-space volume Ω determines shell energy $E(x)$ in the adiabatic limit. This can be seen as contours of the surface $\Omega(E;x)$ in 3D. The elastic force $F(x)$ is the negative of the slope of $E(x)$. At any parameter x and energy E, the shell weight $g(E;x)$ is the tangent slope in the constant-x plane.

Ott[153] showed that in the adiabatic limit of slow time-dependent $x(t)$, an initial phase-space distribution on an energy shell which encloses phase-space volume Ω_0 remains on an energy shell for all future times. What energy $E(x)$ the shell has at a future time (when the parameter has value x) is given by the condition that the enclosed phase-space volume remains constant:

$$\Omega(E(x);x) = \Omega_0. \tag{2.2}$$

This is the *reversible* energy change. The gradient of this energy accounts for the elastic 'force' $F(x)$ on the parameter x,

$$F(x) \equiv -\frac{dE}{dx} = \left\langle -\frac{\partial \mathcal{H}}{\partial x}(\mathbf{q}(t),\mathbf{p}(t),x) \right\rangle_{\text{E}}. \tag{2.3}$$

The subscript E, wherever used, implies a microcanonical average at energy E. The second equality can be proved by differentiation of (2.2) with respect to x, with Ω written as a phase-space integral of the step function $\theta(E-\mathcal{H})$. As shown in Fig. 2.1b, the elastic force is the gradient of the tangent to the surface $\Omega(E;x)$ lying in a constant-Ω plane, that is, $F(x) = -- (\partial E/\partial x)_\Omega$. In the billiard case, $F(x)$ is zero when the deformation preserves the billiard volume.

Now this time-dependence of the shell's energy can be written in terms of the time-evolution of a probability density function. This is an effective equation for timescales longer than the ergodic (mixing) time, so that the distribution across the surface of an energy shell is assumed to have equilibrated (when viewed on a coarse-grained scale) between time-steps. This means that the distribution can be completely characterized by a function of energy alone: it is reduced to a one-dimensional time-dependent distribution $\rho(E,t)$. An equivalent representation is the distribution $\eta(\Omega,t)$ in the phase space volume variable Ω. Keep in

mind that there is a time-dependent (one-to-one) relationship between Ω and E: this can be visualized as constant-x slices of a surface shown in Fig. 2.1b, at the value $x = x(t)$. Likewise the probability densities are related through the time-dependent Jacobean,

$$\eta(\Omega(E; x(t)), t) = \frac{1}{g(E; x(t))}\rho(E, t), \tag{2.4}$$

where the implied dependences on energy and time are shown in full.

There is *no* evolution in the adiabatic limit when written in terms of $\eta(\Omega, t)$:

$$\dot{\eta}(\Omega, t) = 0. \tag{2.5}$$

The adiabatic invariance of Ω for each energy shell implies that the distribution is unchanging.

2.1.2 Energy spreading—constant velocity case

We now derive the next-order correction to the adiabatic limit, namely spreading in energy (or equivalently in Ω), which causes the *irreversible* heating. This was first found by Ott [153] using 'multiple-time-scale analysis'. To derive the spreading rate, we now assume constant parameter velocity $x = Vt$ over a time interval long enough to establish diffusive behaviour (following [47, 46]). The longer-time evolution for a general $x(t)$ can then be built from these short constant-velocity segments.

For any given particle trajectory $(\mathbf{q}(t), \mathbf{p}(t))$ (ensemble member launched at $(\mathbf{q}_0, \mathbf{p}_0)$ at $t = 0$) there is an associated stochastic time-dependent 'force' $-\frac{\partial \mathcal{H}}{\partial x}(\mathbf{q}(t), \mathbf{p}(t), x(t))$ on the parameter x. The time-dependence arises from that of the trajectory. In the billiard case, this is simply the impulses the particle exerts on the deforming part of the wall. We extract the fluctuation of this quantity about its average (which from (2.3) is just the conservative force $F(x)$), giving the definition

$$\mathcal{F}(t) \;\equiv\; -\frac{\partial \mathcal{H}}{\partial x}(\mathbf{q}(t), \mathbf{p}(t), x(t)) \;-\; F(x(t)). \tag{2.6}$$

The external work done in changing the parameter from 0 to x can be found by integrating force over distance to give $Q = -\int_0^x dx'[\mathcal{F}(t(x')) + F(x')]$, which must be converted to an increase in the particle energy. Writing this as a time integral, the particle's energy difference at time t from the adiabatic value $E(x(t))$ is

$$\mathcal{H}(t) - E(x(t)) = -V \int_0^t dt_1 \mathcal{F}(t_1). \tag{2.7}$$

Now $\mathcal{F}(t)$ has a finite correlation time $\tau_{\rm cl}$ which is similar to or less than the ergodic time. For $t \gg \tau_{\rm cl}$ then the above expression is simply the final energy change resulting from a 'random walk' of step size $\sim \tau_{\rm cl}$. Squaring this and taking a microcanonical average (at energy E) gives the energy variance

$$\langle [\mathcal{H}(t) - E(x(t))]^2 \rangle_{\rm E} \;=\; V^2 \int_0^t \int_0^t dt_1 dt_2 \, \langle \mathcal{F}(t_1)\mathcal{F}(t_2) \rangle_{\rm E} \tag{2.8}$$

$$\approx\; 2D_{\rm E} \cdot t \qquad \text{for } t \gg \tau_{\rm cl}. \tag{2.9}$$

This last approximation is good for $t \gg \tau_{\rm cl}$, where the double integral grows linearly in t (with fractional error from linearity dying like $\tau_{\rm cl}/t$). This is simply the variance of a random walk growing linearly in time. It can also be shown using a transform of variables to $t' = (t_2 + t_1)/2$ and $\tau = t_2 - t_1$, with the range of t' being $[0, t]$, and the range of τ can be taken to $[-\infty, \infty]$, as illustrated in Fig. 2.2a. We will write $D_{\rm E}$ in terms of $C_{\rm E}(\tau)$, the autocorrelation function of the fluctuating force. $C_{\rm E}(\tau)$ depends on both the classical motion at energy E and on the particular deformation chosen, and is defined by

$$C_{\rm E}(\tau) \equiv \langle \mathcal{F}(0)\mathcal{F}(\tau)\rangle_{\rm E} \tag{2.10}$$

$$= \langle \mathcal{F}(t')\mathcal{F}(t'+\tau)\rangle_{\rm E}, \qquad 0 < t' < t \tag{2.11}$$

where the second equality states the assumption that the autocorrelation does not change over the timescale t. This latter condition restricts Vt to be much smaller than δx_c^{cl}, the parametric change which changes the classical Hamiltonian (and hence the statistical properties of the system) significantly. In combination with the limit $t \gg \tau_{\rm cl}$, this gives

$$V \ll \delta x_c^{cl}/\tau_{\rm cl} \tag{2.12}$$

(this is the 'trivial slowness condition' of Cohen [46]).

We now have established diffusive energy spreading with a rate (given by substitution of (2.10) into (2.8)) of

$$D_{\rm E} = \tfrac{1}{2}\nu_{\rm E} V^2 \tag{2.13}$$

where $\nu_{\rm E}$ is the (zero-frequency) "noise intensity" of the fluctuating force, equal to the first moment of the autocorrelation function:

$$\nu_{\rm E} \equiv \int_{-\infty}^{\infty} d\tau C_{\rm E}(\tau) \equiv \tilde{C}_{\rm E}(\omega = 0). \tag{2.14}$$

The *correlation power spectrum* or *spectral density* of $\mathcal{F}(t)$ is called $\tilde{C}_{\rm E}(\omega)$, and is the fourier transform of the autocorrelation function. I use the convention

$$\tilde{C}_{\rm E}(\omega) \equiv \int_{-\infty}^{\infty} C_{\rm E}(\tau) e^{i\omega\tau} d\tau. \tag{2.15}$$

Now that diffusion is established for short time-steps t, the evolution over longer times with a general $x(t)$ can be composed of independent diffusive steps each operating on the probability distribution $\rho(E, t)$ given by the previous step. This 'memory-less' stochastic approximation is called *Markovian* [79]. Note that, in this picture, if $x(t)$ ever becomes comparable to δx_c^{cl} then the diffusion rate should be treated as parameter-dependent $D_{\rm E}(x(t))$.

2.1.3 Energy spreading—periodic driving case

We can build on the results of the previous section to see how things change if the parameteric driving is sinusoidal: $x(t) = A\sin(\omega t)$. We will find the new spreading rate $D_{\rm E}$ in the same fashion. For a general $x(t)$, Eq. (2.7) becomes

$$\mathcal{H}(t) - E(x(t)) = -\int_0^t dt_1 \dot{x}(t_1)\mathcal{F}(t_1), \tag{2.16}$$

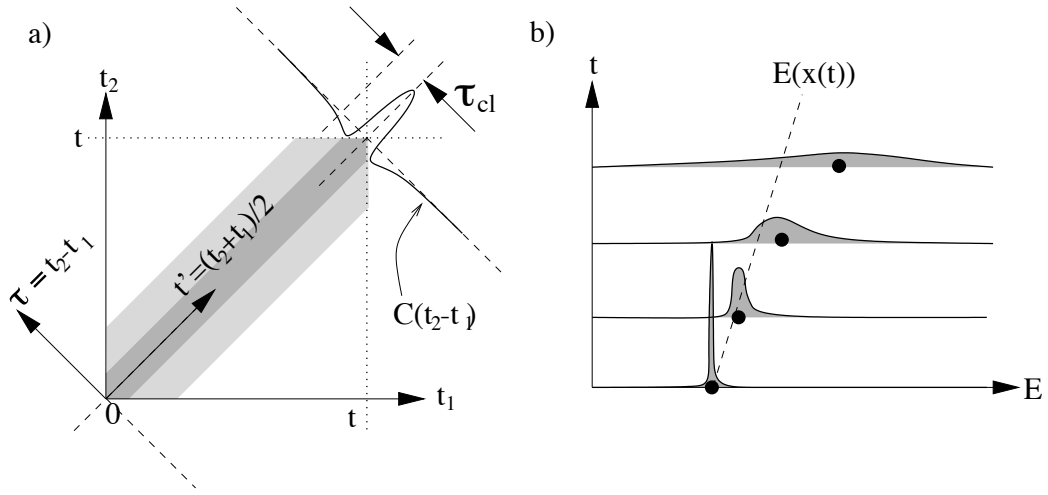

Figure 2.2: **a)** Plot of the microcanonical average of $\mathcal{F}(t_1)\mathcal{F}(t_2)$, showing how its double integral (Eq. 2.8) from time 0 to t can be rewritten as t times the first moment of the autocorrelation function $C_{\rm E}(\tau)$. $C(\tau)$ itself is shown shaded, and it is significant only for $|\tau| < \tau_{\rm cl}$. **b)** Evolution of $\rho(E,t)$ (shaded) showing drift and diffusion. Energy-dependence of both the diffusion constant and the density of states causes an asymmetric distribution, and additional drift. This causes the centroid $\langle E\rangle$ (shown by black dot) to be become higher than the adiabatic value $E(x(t))$ (shown by dashed line). This drift never becomes larger than the width due to spreading.

As before, squaring this and taking the microcanonical average gives

$$\begin{aligned}
\langle[\mathcal{H}(t) - E(x(t))]^2\rangle_{\rm E} &= (A\omega)^2 \int_0^t\int_0^t dt_1 dt_2 \;\cos(\omega t_1)\cos(\omega t_2)\langle\mathcal{F}(t_1)\mathcal{F}(t_2)\rangle_{\rm E} \\
&\approx \frac{(A\omega)^2}{2}\int_0^t dt' \int_{-\infty}^{\infty} d\tau[\cos(2\omega t') + \cos(\omega\tau)]\, C_{\rm E}(\tau) \\
&= \frac{(A\omega)^2}{2}\left(\frac{\sin(2\omega t)}{2\omega}\nu_{\rm E} \;+\; \tilde{C}_{\rm E}(\omega)\cdot t\right). \qquad (2.17)
\end{aligned}$$

On the second line above, the same transformation of time variables was used as before (Fig. 2.2a, except now multiplied by the 'checkerboard' of the cos functions), and the τ range expanded as before to $[-\infty,\infty]$, appropriate when $t \gg \tau_{\rm cl}$. The final line shows an oscillating term and a diffusive term (linear in t), and is valid for any ω. The diffusion becomes dominant in the limit $\omega t \gg 1$ corresponding to observation over a large number of parameter oscillation periods (the fractional error due to the oscillating term dying like $1/\omega t$), giving linear growth of the second moment. This long-time diffusion rate is

$$D_{\rm E} \;=\; \tfrac{1}{2}\tilde{C}_{\rm E}(\omega)V^2 \qquad (2.18)$$

where V^2 is now defined as the *mean square* parameter velocity $\langle\dot{x}^2\rangle = \frac{1}{2}(\omega A)^2$.

So it is apparent that the diffusion is determined by the spectral density of $\mathcal{F}(t)$ at the driving frequency. As this frequency goes to zero, the linear driving result (2.13) is

recovered. This can be verified by taking the limit $\omega \to 0$ while $\omega A = \dot{x}$ is held constant, using $\tilde{C}_{\rm E}(\omega \to 0) = \nu_{\rm E}$, and noticing that the two terms in (2.17) now contribute equally.

2.1.4 Fokker-Planck equation for energy evolution

Stochastic evolution is generally governed by a *Fokker-Planck equation*, describing drift and diffusion of a density function (see Appendix A and [79]). The derivation of the correct Fokker-Planck equation for the time-evolution of the energy distribution $\rho(E,t)$ is not trivial. Now that we have the energy diffusion rate $D_{\rm E}$ at a given energy E, it is tempting to use it directly as a standard (energy-dependent) diffusion term $\frac{\partial}{\partial E}(D_{\rm E}\frac{\partial \rho}{\partial E})$. This simple answer is wrong, and would give the wrong heating rate. As discussed in Appendix A, care is needed in writing the diffusion term. The equation was first written down correctly (without derivation) by Wilkinson [201]. Jarzynski [106] derives it in a clear fashion (which I will follow), and later he [108] and Berry and Robbins [27] rederive it using expansions of the full phase-space density in powers of a 'slowness parameter' ($\propto V$).

It must be realised that, although the spreading rate was derived in energy, it is Ω that is the 'natural' variable in which to discuss spreading. Not only is the evolution of $\eta(\Omega,t)$ simple, but the consideration of Liouville's theorem becomes obvious. In the adiabatic limit $V \to 0$ we have seen that $\eta(\Omega,t)$ is constant (2.5). At finite V the evolution generally becomes

$$\dot{\eta} = -\frac{\partial}{\partial \Omega}(u\eta) + \frac{\partial}{\partial \Omega}\left(D_\Omega \frac{\partial \eta}{\partial \Omega}\right). \tag{2.19}$$

Liouville's equation of motion [126] for $f(\mathbf{q},\mathbf{p},t)$, the full phase-space probability density, is $\dot{f} = -\{f,\mathcal{H}\}$. Regardless of the Hamiltonian and its time-dependence, this means that a constant density must remain so for all times. Recalling that η is just the average of f over an energy-shell, this means $\eta(\Omega,t) = $ const must be a solution of (2.19). It must also be a solution with zero flux j_η everywhere in Ω (if this we not so, a discontinuity would arise at $\Omega = 0$; see also Appendix B of [106]). This implies that in (2.19) the drift term u is zero, and the diffusion is *unbiased* in Ω-space. By 'unbiased' I mean that at every value of Ω, zero density gradient results in zero flux. Note that the D_Ω may still depend on Ω.

Now this "Liouville trick" has been performed, we can transform the evolution back to energy-space; this is done using Appendix A, and gives a Fokker-Planck equation with non-zero drift,

$$v \;=\; \dot{x}\left.\frac{\partial E}{\partial x}\right|_\Omega + \frac{1}{g}\left.\frac{\partial g}{\partial E}\right|_t D_{\rm E}. \tag{2.20}$$

The first term gives the usual adiabatic energy change, while the second is a drift term entirely due to diffusion. This gives the final energy evolution, written with the $D_{\rm E}$ terms grouped together to give a modified diffusion term,

$$\dot{\rho} \;=\; -\dot{x}\frac{\partial}{\partial E}\left(\left.\frac{\partial E}{\partial x}\right|_\Omega \rho\right) \;+\; \frac{\partial}{\partial E}\left[D_{\rm E}\, g \frac{\partial}{\partial E}\left(\frac{\rho}{g}\right)\right]. \tag{2.21}$$

Remember that (2.21) is an effective equation for timescales longer than the ergodic time. In the context of billiards, energy evolution can be thought of as the biased 'random walk' in energy space due to the sum of 'kicks' due to the particle hitting the moving walls [107].

2.1.5 Irreversible growth of average energy

How does the energy spreading cause growth of the average energy? It is a result of both the diffusion term D_{E} being *energy-dependent*, and of the extra drift (bias) due to transformation from Ω to E (see Fig. 2.2b). In the billiard context, the former effect is easy to visualize: we can imagine that particles which happen to have a slightly increased kinetic energy will hit the walls more often, thus diffusing faster in energy. We will see in Chapter 4 that the dependence is $D_{\mathrm{E}} \propto E^{3/2}$, independent of d. The latter effect is less intuitive. Considering that $g(E) \propto E^{d/2-1}$ for billiards it is not hard to show that the ratio of the latter to the former effect is $(d-2)/3$, independent of energy. Therefore which effect dominates depends on d the system dimensionality.

We compute the rate of change of average energy, substituting (2.21) and integrating by parts,

$$\begin{aligned} \langle \dot{E}(t) \rangle &\equiv \int dE\, E\, \dot{\rho}(E,t) \\ &= \dot{x} \int dE\, \frac{\partial E}{\partial x}\bigg|_{\Omega} \rho - \int dE\, D_{\mathrm{E}}\, g \frac{\partial}{\partial E}\left(\frac{\rho}{g}\right). \end{aligned} \tag{2.22}$$

The first term is the reversible heating (easily verified using the microcanonical case $\rho(E) = \delta(E - E_0)$). The second is irreversible since it relies on diffusion; it can be written (by substituting for D_{E}) as

$$\dot{Q}_{\mathrm{irrev}}(t) \;=\; \mu V^2. \tag{2.23}$$

$\mu(\omega)$ is a *friction coefficient*, named so because he agency moving x has to do work against a nonzero average force. We have *ohmic* dissipation, corresponding to an average force $\langle \mathcal{F} \rangle = -\mu \dot{x}$ which is proportional to velocity. The value of $\mu(\omega)$ depends on the energy distribution ρ at time t, and has the general expression

$$\mu(\omega) \;=\; -\frac{1}{2}\int_0^\infty dE\, \tilde{C}_{\mathrm{E}}(\omega)\, g \frac{\partial}{\partial E}\left(\frac{\rho}{g}\right) \;=\; \frac{1}{2}\int_0^\infty dE\, \frac{\rho}{g}\frac{\partial}{\partial E}(g\tilde{C}_{\mathrm{E}}(\omega)) \qquad \text{general } \rho. \tag{2.24}$$

It is interesting that Koonin and Randrup derived an equivalent expression (Eq. (2.23) of [120]), without explicitly considering energy spreading. If we start in a particular choice of ensemble (initial ρ_0), and diffusion has not yet caused ρ to differ much from ρ_0, then $\mu(\omega)$ therefore takes special forms:

$$\mu_{\mathrm{E}}(\omega) \;=\; \frac{1}{2}\frac{1}{g(E)}\frac{\partial}{\partial E}(g(E)\tilde{C}_{\mathrm{E}}(\omega)) \qquad \text{microcanonical,} \tag{2.25}$$

$$\mu(\omega) \;=\; \frac{g(E_{\mathrm{F}})}{2\Omega_{\mathrm{F}}}\tilde{C}_{\mathrm{F}}(\omega) \qquad \text{Fermi distribution,} \tag{2.26}$$

where the subscript $_{\mathrm{F}}$ means evaluation at the Fermi energy E_{F}. The Fermi distribution has a constant phase-space density of Ω_{F}^{-1}, normalised to represent a single particle. For an application to many non-interacting fermions, see Section 2.2.4. The case of the canonical distribution, relevant to a thermal gas of non-interacting classical particles, is covered by Cohen (Section 4 of [46]) and I do not discuss it here. After long times, if the system

remains isolated, it can be that ρ is very different from ρ_0, and the special forms (2.25) or (2.26) will have to be replaced by (2.24).

What I have presented is only the dissipation appearing to lowest order in V (Jarzynski [107] discusses this by way of an expansion in powers of a 'slowness parameter' $\propto V$). There will be higher-order terms in the friction which become more relevant as V approaches for example the limit (2.12). There are other ways that the classical spreading picture can break down. For instance for t less than the correlation time $\tau_{\rm cl}$ the energy spreading should be much less than E itself. Also for this same t the average energy increase should remain small compared to the spreading. These set additional upper limits on V [46].

2.1.6 Microcanonical averages and trajectory averages

In the preceding sections $F(x)$ and $\tilde{C}_{\rm E}(\omega)$ were defined in terms of the microcanonical average over initial conditions,

$$\langle\cdots\rangle_{\rm E} \equiv \int d\mathbf{q}d\mathbf{p}\; f(\mathbf{q},\mathbf{p},0)(\cdots), \tag{2.27}$$

$$f(\mathbf{q},\mathbf{p},0) = \frac{1}{g(E)}\delta(\mathcal{H}(\mathbf{q},\mathbf{p},x)-E) \tag{2.28}$$

The assumption was made that the change in x over the correlation time $\tau_{\rm cl}$ is insignificant. This means that the 'frozen' Hamiltonian (at fixed x) can be used, so the distribution is unchanging in time $f(\mathbf{q},\mathbf{p},t)=f(\mathbf{q},\mathbf{p},0)$. All resulting averages are constant in time (or, if they involve multiple times, they are functions of time differences only), and the choice of initial time $t=0$ is arbitrary.

However, these averages need not be taken using an average over phase space. By ergodicity [79], they are equal to time averages over a single trajectory. Namely, the conservative force (2.3) can be written

$$F(x) = -\lim_{T\to\infty}\frac{1}{T}\int_0^T dt\;\frac{\partial\mathcal{H}}{\partial x}(\mathbf{q}(t),\mathbf{p}(t),x), \tag{2.29}$$

seen to be the mean force due to motion of a single particle in the system. Similarly, the auto-correlation is written

$$C_{\rm E}(\tau) = \lim_{T\to\infty}\frac{1}{T}\int_0^T dt\;\mathcal{F}(t)\mathcal{F}(t+\tau). \tag{2.30}$$

There is now an issue of convergence: the number of independent samples of any quantity along a trajectory is $N \approx T/\tau_{\rm cl}$, and the fractional error of the estimate of the above quantities converges slowly as $\sim N^{-1/2}$. Therefore very long trajectories are required to get good estimates. However, this is often easier than performing the $2d-1$ dimensional integral over the energy shell which would be required for the explicit evaluation of the phase-space average, especially since the integrand in (2.30) already involves propagation forward in time. This technique of evaluation of a multi-dimensional integral using a random sample of points taken from the distribution function is called *Monte Carlo* integration [161]. However, in practice, rather than compute $C_{\rm E}(\tau)$ and take the Fourier transform, $\tilde{C}_{\rm E}(\omega)$

is most efficiently estimated directly from the Fourier transform of $\mathcal{F}(t)$. This approach is discussed in detail in Appendix B.

In an identical fashion to that shown in Fig. 2.2a, Eq.(2.30) describes $C_{\rm E}(\tau)$ as the projection of the function $\mathcal{F}(t_1)\mathcal{F}(t_2)$ onto the $\tau = t_2 - t_1$ axis. However only a single trajectory is involved, so $\mathcal{F}(t)$ is noisy, and an average over the time axis $t' = \frac{1}{2}(t_2 + t_1)$ is required. In this figure, one can imagine the 'box' of allowed t_1,t_2 values now bounded by $[0, T]$. The average converges in the limit $T \to \infty$.

An instructive convergence issue arises if we extend this single-trajectory estimate to the noise intensity $\nu_{\rm E}$. Naive use of (2.14) and (2.30) would give

$$\nu_{\rm E} \;=\; \lim_{T\to\infty} \frac{1}{T} \int_0^T dt \int_\infty^\infty d\tau \; \mathcal{F}(t)\mathcal{F}(t+\tau) \qquad \text{non-convergent.} \tag{2.31}$$

This does not converge because the τ integral involves an infinite number of similarly-sized contributions (the integrand is similar in size for all t, τ). Averaging over any finite T cannot remove this divergence. Another option is to look back to Eq.(2.8) which is responsible for the appearance of $\nu_{\rm E}$ in the spreading rate. So I could define

$$\nu_{\rm E} = \lim_{T\to\infty} \frac{1}{T} \left[\int_0^T dt \; \mathcal{F}(t) \right]^2 = \lim_{T\to\infty} \frac{1}{T} \int_0^T dt_1 \int_0^T dt_2 \; \mathcal{F}(t_1)\mathcal{F}(t_2) \qquad \text{nearly convergent,} \tag{2.32}$$

which is the same integral (2.31) with different limits. This is simply proportional to the energy variance δE^2 after time T for the single trajectory involved, divided by T. The term inside the square brackets is simply a random walk (on timescales $\gg \tau_{\rm cl}$), giving a Gaussian distribution whose variance grows linearly. Therefore this estimate of $\nu_{\rm E}$ will not converge, rather it will wander for all $T \to \infty$, taking values with a χ^2 distribution whose mean is the correct $\nu_{\rm E}$. In effect this reproduces exactly the stochastic energy spreading whose variance is desired! However, it is 'more convergent' than (2.31). A convergent estimate can only be created by limiting the τ integration further, to give

$$\nu_{\rm E} \;=\; \lim_{t_0\to\infty} \lim_{T\to\infty} \frac{1}{T} \int_0^T dt \int_{-t_0}^{t_0} d\tau \; \mathcal{F}(t)\mathcal{F}(t+\tau) \qquad \text{convergent,} \tag{2.33}$$

which will converge once $T \gg t_0 \gg \tau_{\rm cl}$. In essence the convergence arises because the limits do not allow the number of independent $\tau_{\rm cl}$-squared-sized 'patches' of the integrand to grow as fast as N^2, where N is defined as above. The nearly-convergent case (2.32) corresponds to exactly N^2 growth.

The above considerations will not relate directly to the numerical method of finding $\nu_{\rm E}$ (which will be via $\tilde{C}_{\rm E}(\omega \to 0)$). However they serve to warn and provide intuition about convergence when a single trajectory is used for estimation.

2.2 Review of the linear response theory of quantum dissipation

Here I remind the reader of the quantum-mechanical linear response theory (LRT), as applied to a chaotic system. LRT in electron transport is considered to originate with

Kubo [122] and Greenwood [84], and is often called the 'Kubo formalism' [99]. It has more recently been applied to dissipation in closed chaotic systems [118, 200, 47, 46]. As a primer, I will start with a conventional presentation of perturbation theory in a fixed basis. Then, once stochastic energy-spreading is established, I link up with the classical picture of Section 2.1. In particular, my aim is to show the connection between heating rate and a certain correlation spectrum $\tilde{C}_{\rm E}^{\rm qm}(\omega)$. Finally, because the energy-spreading picture differs markedly from the language of LRT in the condensed matter literature, I will make contact with the latter. I hope this will provide a useful bridge to the unified picture of Cohen[46].

2.2.1 Basis choice, perturbation theory and the FGR

Dissipation in a quantum system will be the result of 'driving' a parameter x on which the Hamiltonian depends. The parameter is externally controllable with any (real) function of time $x(t)$; it does not have any dynamics of its own (it is not a degree of freedom). We start by linearizing about $x = 0$,

$$\mathcal{H}(x) \;\approx\; \mathcal{H}(0) + \frac{\partial \mathcal{H}}{\partial x} x. \tag{2.34}$$

The hermitian operator $-\partial\mathcal{H}/\partial x \equiv \mathcal{F}$ is the 'generalized force' conjugate to the parameter x: it is an actual force if x is a displacement, it is a dipole moment if x is an applied electric field magnetic field, it is a current around a ring (magnetization) if x is magnetic flux enclosed by the ring, and so forth. I shall now use time-dependent perturbation theory, choosing to perform this in a *fixed basis* of the eigenstates of $\mathcal{H}(0)$, which we call the 'unperturbed' Hamiltonian. This choice of basis is conventional in textbooks [128, 99, 118] so is ideal for introductory purposes. However this choice does not allow an understanding of the limitations of conventional LRT; to understand breakdown and possible extensions to LRT it is necessary to consider transitions in the so-called 'adiabatic' basis of the local eigenstates of $\mathcal{H}(x)$ ([200, 46, 48]. The adiabatic basis is much more natural for unifying with the classical picture [46]. It also removes certain difficulties, for instance that of the infinite matrix elements $\frac{\partial \mathcal{H}}{\partial x}$ which would otherwise occur for any deformation $x \neq 0$ of a hard-walled billiard system. As another example, the case of constant velocity driving $x = Vt$ cannot be shown to give diffusive E spreading in a fixed basis—the adiabatic basis is required (Appendix B of [46]).

The Schrodinger-picture wavefunction is represented as

$$|\psi(t)\rangle \;=\; \sum_n a_n(t)\,|n(t)\rangle\,, \qquad\qquad |n(t)\rangle \;=\; e^{\frac{1}{i\hbar}E_n t}\,|n(0)\rangle\,, \tag{2.35}$$

where the time-independent states $|n\rangle \equiv |n(0)\rangle$ are eigenstates of the unperturbed Hamiltonian: $\mathcal{H}(0)\,|n\rangle = E_n\,|n\rangle$. The basis vectors are fixed but include phase evolution under $\mathcal{H}(0)$ (following [128, 160]), thus the coefficients $a_n(t)$ are equivalent to an 'interaction picture' representation. The full time-evolution

$$i\hbar \frac{d}{dt}\,|\psi(t)\rangle \;=\; \mathcal{H}(x(t))\,|\psi(t)\rangle \tag{2.36}$$

gives a first-order differential equation for the coefficients

$$\frac{da_n}{dt} \;=\; \frac{1}{i\hbar}x(t)\sum_m \mathcal{F}_{nm}(t)a_m \tag{2.37}$$

where $\mathcal{F}_{nm}(t) \equiv (\partial\mathcal{H}/\partial x)_{nm}e^{-i\omega_{nm}t}$ and $\hbar\omega_{nm} \equiv E_n - E_m$. The formal solution (using notation $\mathbf{a} \equiv \{a_n\}$, and the matrix $\mathcal{F}$) is a time-ordered exponential (Dyson series [174])

$$\mathbf{a}(t) \;=\; \exp\left(\frac{1}{i\hbar}\int_0^t dt' x(t')\mathcal{F}(t')\right)\mathbf{a}(0). \tag{2.38}$$

I take the initial condition as

$$a_n(0) \;=\; \delta_{nm} \qquad \text{microcanonical distribution,} \tag{2.39}$$

that is, a single initial pure eigenstate $|m\rangle$, well above the ground state. The generalization to any incoherent ensemble of initial eigenstates (*e.g.* canonical ensemble) can be achieved by an occupation-weighted average of the results over m. Formulations of LRT ([122, 84, 99, 118, 200], and Ingold in [65]) using the density matrix are equivalent to this, although the notation may superficially look different. Using (2.39) and keeping only the first-order term in x in (2.38) gives

$$a_n(t) \;=\; \frac{1}{i\hbar}\left(\frac{\partial\mathcal{H}}{\partial x}\right)_{nm}\int_0^t dt'\, x(t')e^{-i\omega_{nm}t'}, \qquad n \neq m, \tag{2.40}$$

which is called first-order perturbation theory (FOPT). Clearly this is a valid approximation only when $a_m(t)$ remains close to unity and $a_n(t) \ll 1$ for all $n \neq m$. If this is true, the response of any expectation value is linear in x: this is what is meant by the word 'linear' in LRT (rather that the linearization (2.34)).

I will now specialize to the case of periodic driving at amplitude A,

$$x(t) \;=\; A\sin(\omega t) \;=\; \tfrac{1}{2}A(e^{i\omega t} - e^{-i\omega t}). \tag{2.41}$$

Substituting into FOPT and squaring the absolute value gives

$$|a_n(t)|^2 \;=\; \frac{1}{\hbar^2}\frac{A^2}{4}\left|\left(\frac{\partial\mathcal{H}}{\partial x}\right)_{nm}\right|^2 \cdot t\left[\mathrm{sinc}^2\tfrac{1}{2}(\omega_{nm}-\omega)t + \mathrm{sinc}^2\tfrac{1}{2}(\omega_{nm}+\omega)t\right] + \text{interference terms,} \tag{2.42}$$

where $\mathrm{sinc}(x)$ as usual means $\sin(x)/x$. After many periods ($t \gg 1/\omega$) the sinc^2 functions become localized delta-like functions of width $\varepsilon \sim 1/t$,

$$\lim_{t\to\infty} t\,\mathrm{sinc}^2(\alpha t) \;=\; \pi\delta_\varepsilon(\alpha), \tag{2.43}$$

and the interference terms become irrelevant as these functions become separated. In order that these delta-like functions do not become narrow enough to resolve discrete energy levels, we are limited by the Heisenberg time, $t \ll t_{\mathrm{H}}$. Substituting the delta functions

(from which I now drop the width subscript ε), the probabilities in states $n \neq m$ can be seen to grow linearly, with transition rates

$$\Gamma_{nm}(\omega) = \frac{2\pi}{\hbar^2}\frac{A^2}{4}\left|\left(\frac{\partial \mathcal{H}}{\partial x}\right)_{nm}\right|^2 [\delta(\omega_{nm}-\omega)+\delta(\omega_{nm}-\omega)] \quad (2.44)$$

$$= \frac{2\pi}{\hbar}\frac{A^2}{4}\left|\left(\frac{\partial \mathcal{H}}{\partial x}\right)_{nm}\right|^2 [\delta(E_n-E_m-\hbar\omega)+\delta(E_n-E_m+\hbar\omega)] \quad (2.45)$$

recognizable as the usual Fermi Golden Rule (FGR) in the presence of both positive and negative frequency driving.

In textbook treatments of the FGR [128], the squared matrix element $|(\partial\mathcal{H}/\partial x)_{nm}|^2$ is assumed to be a smooth function of E_n, enabling it to be taken as constant over the width of the delta-function— this is often true in integrable systems. In contrast, in a chaotic system it takes essentially random, uncorrelated values at each n and m (this assumption underlies any Random Matrix Theory[35, 146] description of chaos). The *average* transition rate is given by the local mean value near the location nm in the matrix (see Fig. 2.3). This will next be formalised using the 'band profile' of the matrix.

2.2.2 Generalized force-force correlation and band profile

Here I show a way to express the average transition rates in terms of the time correlation function

$$C_{\rm E}^{\rm qm}(\tau) \equiv \frac{1}{M}\sum_m \langle m\,|\mathcal{F}(0)\mathcal{F}(\tau)|\,m\rangle = \langle \mathcal{F}(0)\mathcal{F}(\tau)\rangle_{\rm E}, \quad (2.46)$$

where the (generalized force) operator $\mathcal{F}$ is an abbreviation for $-\partial\mathcal{H}/\partial x$. Note that $\langle m\,|\mathcal{F}(0)\mathcal{F}(\tau)|\,m\rangle$ corresponding to the initial state (2.39) can be explicitly written $\sum_n |\mathcal{F}_{nm}|^2 e^{-i\omega_{nm}\tau}$. Averaging over $M \gg 1$ adjacent initial states $|m\rangle$ (which nevertheless span a classically-small energy range) has been performed in order to extract the local average of the random values of $|\mathcal{F}_{nm}|^2$. This smearing will from now be implied by the microcanonical ensemble average at energy E, indicated by the notation $\langle\cdots\rangle_{\rm E}$. The Fourier transform gives the *auto-correlation spectrum*

$$\tilde{C}_{\rm E}^{\rm qm}(\omega) = \int_{-\infty}^{\infty} d\tau e^{i\omega\tau} C_{\rm E}^{\rm qm}(\tau) = \sum_n |\mathcal{F}_{nm}|^2 2\pi\delta(\omega_{nm}-\omega), \quad (2.47)$$

which will play a central role in this and the following chapters.

Here, as before, the delta function must be taken to have a width $\varepsilon > \Delta/\hbar$, where Δ is the mean level spacing in energy. Thus there are *two* components to the smoothing ('smearing') procedure: a smearing by width $M\Delta$ along the diagonal, and by width $\hbar\varepsilon$ off the diagonal. The smearing is chosen to be sufficient to allow $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ to become a well-defined function of ω, and smooth in E. By using the continuum limit substitution $\sum_n \to (1/\Delta)\int dE = (\hbar/\Delta)\int d\omega$, the auto-correlation spectrum can now be interpreted as giving a formula for the mean value of the squared matrix element:

$$\left\langle |\mathcal{F}_{nm}|^2 \right\rangle = \frac{\Delta}{2\pi\hbar}\tilde{C}_{\rm E}^{\rm qm}(\omega=\omega_{nm}). \quad (2.48)$$

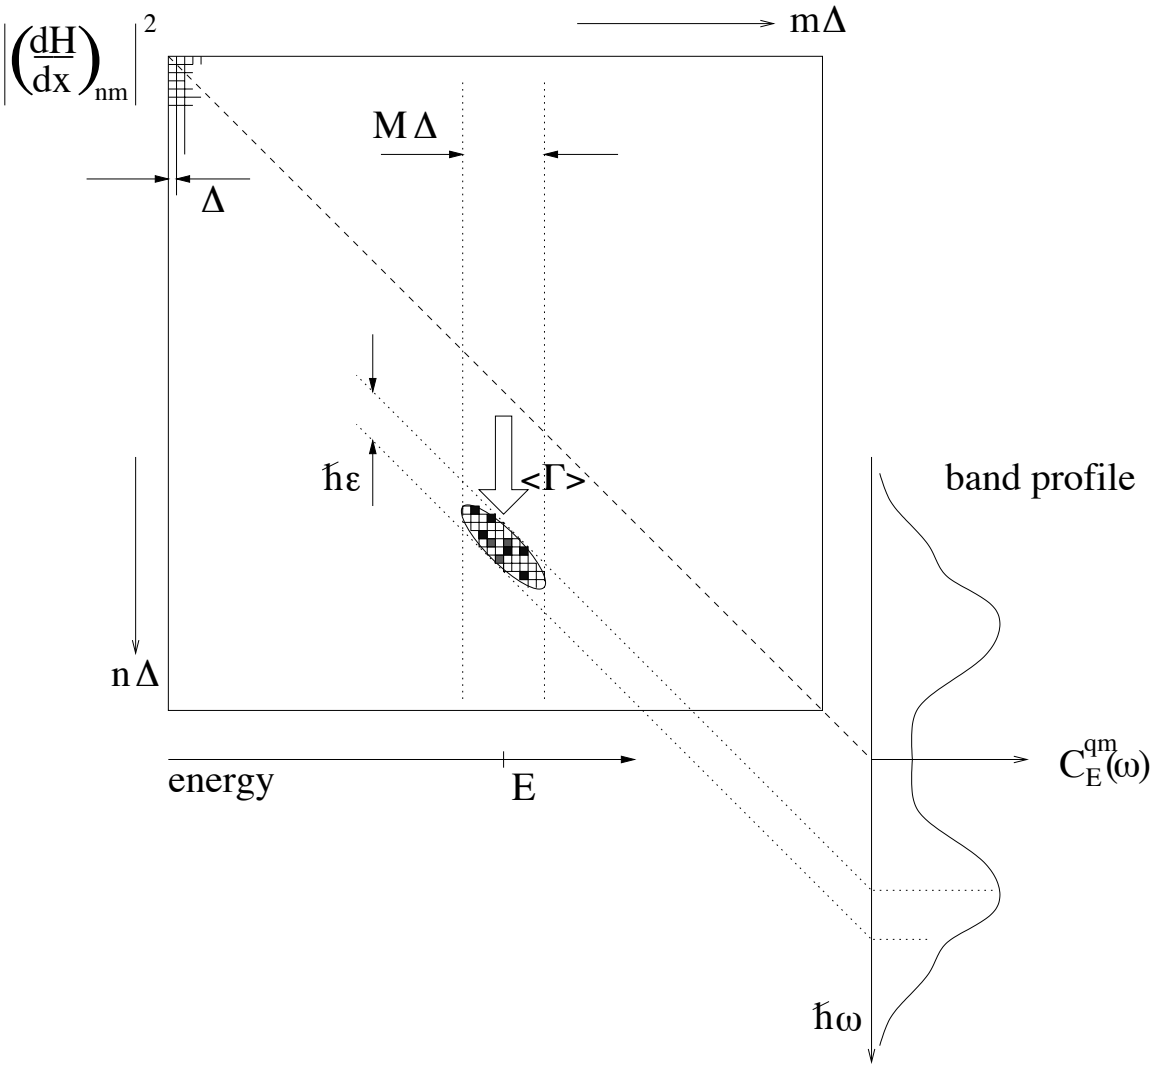

Figure 2.3: Relation of the band profile $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ to the squared matrix elements $|(\partial\mathcal{H}/\partial x)_{nm}|^2$ around $E_n \approx E_m \approx E$. The diagram of the matrix has energy units in both axes. The local average of the transition rate Γ is taken over an ellipse whose horizontal width is an energy range $M\Delta$, and whose width in the off-diagonal direction is and energy range $\hbar\varepsilon$. Both widths are $\gg \Delta$, the mean level spacing. The width $M\Delta \ll E$ is classically small.

The right-hand side, viewed as a continuous function of ω, is called the *band profile*[1] of the matrix $\mathcal{F}_{nm}$ because it measures its off-diagonal ('band') structure. This construction is illustrated by Fig. 2.3. The matrix, and hence $\tilde{C}_{\rm E}^{\rm qm}(\omega)$, has structure on ω-scales similar to inverse correlation times of the chaotic motion (for instance the bouncing rate in a billiard system). However the dependence on E is very much weaker (if $\hbar/M\Delta$ is chosen smaller than the shortest periodic orbit [200], as discussed in Section 2.3.3), only changing over classically-large energy scales, so can be taken as constant in any classically-small range. An example matrix from a real system is shown in Fig. 2.7).

We can now replace the FGR transition rate by its well-defined average over the width of the delta-like functions,

$$\langle \Gamma_{nm}(\omega) \rangle \;=\; \frac{A^2}{4}\frac{\Delta}{\hbar^3}\tilde{C}_{\rm E}^{\rm qm}(\omega_{nm})[\delta(\omega_{nm}-\omega)+\delta(\omega_{nm}+\omega)] \tag{2.49}$$

which is given by the band profile.

2.2.3 Stochastic energy spreading

Returning now to the FGR expression (2.45), one might wonder how it is possible to get dissipation (irreversible growth in the mean energy) at all, given that the up and down transition rates are always equal! The answer will be that heating is possible because the *density of states* $g(E)$ may differ slightly at the two energies $E = E_m \pm \hbar\omega$ being 'pumped' into, and therefore more states may fall under the upper delta function than the lower, giving an increase in mean energy. However to understand this in more detail, the up and down transitions (2.45) must be interpreted as giving *diffusion* in energy-space. A diffusion equation for time-evolution of a continuous energy distribution function $\rho(E)$ can then be written—from this follows the rate of increase of mean energy.

The linear growth of $p_n(t) \equiv |a_n(t)|^2$ resulting from the FGR means a linear growth in the second moment (variance) of the distribution about the initial energy E_m,

$$\delta E^2 \;\equiv\; \langle (\mathcal{H}(t)-E_m)^2 \rangle_{{\rm E}_m} \;=\; \sum_n (E_n - E_m)^2 p_n(t). \tag{2.50}$$

This can be identified with diffusive growth $\delta E^2 = 2D_{\rm E}\cdot t$ with diffusion constant given by

$$\begin{aligned} D_{\rm E}(\omega) &= \frac{1}{2}\sum_n (E_n-E_m)^2\,\Gamma_{nm}(\omega) \\ &= \frac{1}{2}\cdot\tfrac{1}{2}(\omega A)^2\cdot\tfrac{1}{2}[\tilde{C}_{\rm E}^{\rm qm}(\omega)+\tilde{C}_{\rm E}^{\rm qm}(-\omega)] \\ &= \frac{1}{2}\cdot\tfrac{1}{2}(\omega A)^2\cdot\tilde{C}_{\rm E}^{\rm qm}(\omega) \end{aligned} \tag{2.51}$$

Use was made of (2.49), which is appropriate since we are considering a continuum limit. Also, as in the previous section, an averaging over a classically-small range of adjacent states was performed. Thus we see that the band profile at the driving frequency ω determines

[1] Note that I will also refer to $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ as the band profile, since it differs only by a constant factor.

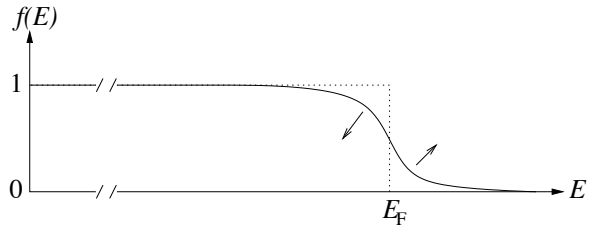

Figure 2.4: Fermi occupation function $f(E)$ evolving under energy diffusion, causing growth of mean energy. The initial ensemble occupation function is shown by the dotted line.

the resulting diffusion rate. The term $\frac{1}{2}(\omega A)^2$ is simply the mean square parametric driving velocity $\langle \dot{x}^2 \rangle$.

So far the evolution of the coefficients $a_n(t)$ has been entirely coherent (this is unaffected by the final averaging over initial state m). Therefore an important issue arises: does the energy diffusion eventually stop at $t \sim t_{\rm H}$, when quantum recurrences due to the discreteness of the spectrum will surely occur? (For instance, in the kicked rotator system [75], periodic driving results in saturation of energy spreading, giving localization). The answer (discussed in detail in Section 8 and Appendix B of [46]) is that preservation of coherent effects is actually very fragile. Either $x(t)$ can have some slight stochasticity (frequency jitter), or there is some slight dephasing process (interaction with the environment, universal in real systems). Correlations due to coherence are then lost on timescales longer than that associated with jitter or dephasing. So if diffusive spreading and loss of coherence can be established before $t_{\rm prt}$, the 'breaktime' (*i.e.* time limit of applicability) of the FGR, then stochastic spreading will continue for ever. In other words the 'rate equation' (2.49) will remain valid, and we will not be surprised by recurrence (interference) effects at long times. The FGR breaktime will be discussed more in Section 2.2.6.

2.2.4 Heating rate example—non-interacting fermions

Once the stochasticity of energy spreading is established, the quantum and classical pictures of dissipation coincide. Therefore the reader is referred to the general diffusion equation and expressions for heating rate in Sections 2.1.4 and 2.1.5 of the classical review. In these expressions, the classical $\tilde{C}_{\rm E}(\omega)$ should be replaced by $\tilde{C}_{\rm E}^{\rm qm}(\omega)$, and the the classical $g(E)$ divided by $(2\pi\hbar)^d$.

However I will for introductory purposes present here the simple example of non-interacting fermions in a driven chaotic system. This is relevant to nucleons in a deforming chaotic nucleus [118, 120, 29], and to electrons in small irregular metal particles[201], or in quantum dots[15], subjected to external driving fields. The simplicity arises because all the heating occurs at the Fermi edge. The Fermi 'occupation function' $f(E)$ is the appropriate initial distribution of single-particle states: the many-particle wavefunction is a Slater determinant of these states. At low temperatures $f(E)$ is close to a unit step-function, and subsequent diffusion smoothes out the step as shown in Fig. 2.4. One might ask how the exclusion principle affects the evolution of the many-particle system: the answer is that at all future times the many-particle wavefunction is correctly given by a Slater determinant

of the single-particle states (this follows from unitarity of time-evolution, also see [201]), as long as any external interaction can be ignored. Thus we can get the heating rate by evolving the single-particle distribution function f. What follows is therefore identical to the classical case. The normalization is such that $N = \int g(E)f(E)$ gives the number of particles in the system.

The diffusion equation is $\dot{f} = (D_{\rm E} f')'$ where a prime means E-derivative. The rate of increase in ensemble-averaged energy is

$$\begin{aligned} \langle \dot{E} \rangle &= \int dE\, E \dot{f}(E) g(E) \;=\; -\int dE\, (Eg(E))' D_{\rm E} f' \\ &= g(E_{\rm F}) D_{\rm F}, \end{aligned} \tag{2.52}$$

where the boundary term in the integration by parts vanishes, f' is replaced by a negative delta-function, and a constant density of states $g(E_{\rm F})$ can be assumed near the Fermi energy. The diffusion rate at the Fermi energy I call $D_{\rm F}$. Heating of the Fermi gas can be interpreted as friction on the parameter x. The friction coefficient μ is defined by $\langle \dot{E} \rangle = \mu \langle \dot{x}^2 \rangle$. Comparing (2.52) and (2.51) gives

$$\mu(\omega) \;=\; \frac{g(E_{\rm F})}{2} \tilde{C}^{\rm qm}_{\rm E}(\omega). \tag{2.53}$$

In the zero-frequency limit, $\tilde{C}^{\rm qm}_{\rm E}(\omega) \to \tilde{C}^{\rm qm}_{\rm E}(0) \equiv \nu_{\rm E}$ which is the (dc) fluctuations intensity of the observable $\mathcal{F}$. This resulting relation $\mu = \frac{1}{2} g(E_{\rm F}) \nu_{\rm F}$ is an example of a fluctuation-dissipation relation. Different relations arise in different ensembles—for instance the canonical ensemble produces the traditional fluctuation-dissipation relation at fixed temperature [127, 65, 46].

2.2.5 Relation to susceptibility

Here, for the sake of intuition, I briefly make contact with the traditional LRT notation arising in condensed matter physics [127, 122, 160, 99, 65]. I consider first the 'response' (time-dependent expectation value) of a general measurable whose operator is $\mathcal{G}$. Looking at Eq.(2.40) we can write

$$\langle \mathcal{G}(t) \rangle \;=\; \mathcal{G}_{mm}(t) + \sum_{n \neq m} [a^*_n(t) \mathcal{G}_{nm}(t) + \text{c.c.}] \tag{2.54}$$

where the first term is the constant equilibrium value $\langle \mathcal{G}(0) \rangle$ and the other terms have been kept only up to first order in the driving x. The response is linear but is not generally local in time, *i.e.* it has memory. However, because of translational invariance in time, the response must be local in *frequency*. So we can relate the Fourier transform of response to that of driving by $\langle \tilde{\mathcal{G}}(\omega) \rangle = \chi_{\rm GF}(\omega) \tilde{x}(\omega)$, where $\chi_{\rm GF}(\omega)$ is a generalized susceptibility. Numerous symmetry relations of $\chi_{\rm GF}(\omega)$ follow from causality[127]; in particular $\chi_{\rm GF}(-\omega) = \chi^*_{\rm GF}(\omega)$. Substitution of (2.40) gives

$$\chi_{\rm GF}(\omega) = \frac{i}{\hbar} \int_0^\infty dt\, e^{i\omega t} \left([\mathcal{G}(t), \mathcal{F}(0)] \right)_{mm} \to \frac{i}{\hbar} \int_0^\infty dt\, e^{i\omega t} \left\langle [\mathcal{G}(t), \mathcal{F}(0)] \right\rangle, \tag{2.55}$$

where the latter form is averaged over m with the appropriate initial distribution p_m, and is a standard result of LRT [65].

We now choose $\mathcal{G} = \mathcal{F}$ whereupon the subscript GF is no longer needed. The susceptibility is split into real and imaginary parts, $\chi = \chi' + i\chi''$. The dissipation (energy absorption) is determined by χ'' [127], since, using the driving (2.41) we have

$$\begin{aligned} \dot{Q}_{irrev} &= \langle \dot{\mathcal{H}}(t) \rangle = \dot{x} \langle \mathcal{F}(t) \rangle = \frac{A^2}{4} i\omega (e^{i\omega t} + e^{-i\omega t}) \left[\chi(\omega) e^{i\omega t} - \chi(-\omega) e^{-i\omega t} \right] \\ &= \tfrac{1}{2}(wA)^2 \cdot \frac{\chi''(\omega)}{\omega}, \end{aligned} \tag{2.56}$$

where at the last stage the oscillatory terms average to zero. This heating rate expression results entirely from the definition of χ. However now our friction coefficient from the previous section can be seen to be simply $\mu(\omega) = \chi''(\omega)/\omega$.

We should also note that in the condensed matter physics context $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ is proportional to the *dynamic form factor* $S_{\mathbf{q}}(\omega)$ if we associate $\partial\mathcal{H}/\partial x$ with the density operator of a mode labelled $\mathbf{q}$ [160]. In the case where x is flux enclosed by a mesoscopic sample (quantum dot), $\partial\mathcal{H}/\partial x$ is a current operator, thus μ determines the conductivity of the system. As one would expect, the conductivity is then given by the time-integral of the current-current correlation $e^2\langle v(0)v(t)\rangle$ [122, 99, 8]. We have performed such a calculation for a chaotic mesoscopic dot [15].

2.2.6 Regime of applicability of linear response

In the adiabatic limit $V \to 0$ an initial state $|m(0)\rangle$ will remain in the local eigenstate $|m(x)\rangle$ for all x. Hence for extremely small V, LRT will not be valid because Landau-Zener transitions (avoided crossings) become the dominant heating mechanism, as explained by Wilkinson [200]. This is called the *quantum adiabatic* regime. For a given system and given parameter x, there will be a typical change in x which results in each energy level encountering about one avoided crossing. More strictly this is expressed by the criterion that eigenstate overlaps (see Section 6.4) should mix about one nearest neighbour. This parameter change is called $\delta x_{\rm c}^{\rm qm}$. If this change takes longer than $t_{\rm H}$ to occur, then the system remains 'localized' on the initial state, and we have adiabaticity. The above also requires that no dephasing processes occur before $t_{\rm H}$.

LRT also has an upper velocity limit, beyond which non-perturbative effects dominate. The key criterion is whether LRT is valid for a single correlation time $\tau_{\rm cl}$ (the timescale required to establish diffusive spreading [46]). Once this is true, stochastic energy spreading is expected to continue forever. FOPT will break down when there is non-perturbative mixing between levels. If the band profile is flat near $\omega = 0$, this happens first between neighbouring levels, when x reaches $\delta x_{\rm c}^{\rm qm}$. Therefore the limit of applicability of FOPT is $V \ll \delta x_{\rm c}^{\rm qm}/\tau_{\rm cl}$. This was once thought [202] to be the point at which the heating rate departs from LRT. Beyond this, a 'core' region will be created which is non-perturbatively mixed, before diffusive growth is established. However, Cohen[46] has realised that the existence of this core need not invalidate the LRT result, because it is the 'tail' region (everything outside the core) which dominates the energy spreading. His modification extends

the applicability of LRT up to $V \sim \delta x_{\rm c}^{\rm qm}/t_{\rm prt}$, where $t_{\rm prt}$ is the timescale for much of the probability to have left the initial level (beyond which perturbation theory is past rescue).

Beyond this, the LRT picture breaks down. However, at the highest V, it is expected that there is semiclassical correspondence, so that the *classical* dissipation rate applies. The above sequence of spreading profiles, FOPT $\rightarrow$ core-tail $\rightarrow$ semiclassical, is demonstrated in our work [48] (the calculations for which appear in Section 6.4 of this thesis).

2.3 Quantum-classical correspondence

The classical calculation of Section 2.1 and the quantum LRT calculation of Section 2.2 gave the same answers for heating rate. This is remarkable because the energy evolution (spreading profile) was different: the classical spreading is gradual (for $t > \tau_{\rm cl}$), whereas the quantum spreading involves jumps in units of $\hbar\omega$ (for driving frequency ω). It is also remarkable because the quantum result is therefore independent of $\hbar$. However the heating rates agree because the expressions for spreading *rate* (second moment of spreading profile) were the same. The only difference is the replacement of the classical definition of $\tilde{C}_{\rm E}(\omega)$ by the corresponding quantum-mechanical definition $\tilde{C}_{\rm E}^{\rm qm}(\omega)$. If these two agree, then there is *quantum-classical correspondence* (QCC) as far as dissipation (friction μ) is concerned [47, 46]. Cohen calls this 'restricted correspondence' because of the different spreading profiles. Here I present numerical evidence supporting this claim of QCC in a real system (a 2D cavity). Hence in the following chapters I will move freely between the quantum and classical pictures. In particular, the term 'band profile' will then refer to both quantum and classical $\tilde{C}_{\rm E}(\omega)$ auto-correlation spectra.

First I briefly mention QCC outside the LRT regime (whose limits were outlined in the previous section). For extremely small V in the quantum-adiabatic regime, there is no QCC expected because the level spacing distribution $P(s)$ for small $s \ll \Delta$ dominates the heating rate [200]. This is a purely quantum effect (determined by certain quantum symmetries [35, 146]), and has no reason to agree with any classical quantity. At the other extreme, as $\hbar \rightarrow 0$ the maximum V where LRT is valid also vanishes. Therefore if one is to have QCC in the semiclassical limit (a fundamental requirement of quantum mechanics being that it reduces to classical mechanics in this limit), some new mechanism is required. As explained by Cohen[47, 46], in this limit QCC is achieved through semiclassical agreement of the spreading profiles ('detailed correspondence').

2.3.1 Semiclassical connection between quantum and classical band profiles

Semiclassical convergence between the classical and quantum $\tilde{C}_{\rm E}(\omega)$ is a particular case of convergence for a general operator, found originally by Feingold and Peres [73], and most clearly expressed (as a 'sum rule') by Wilkinson [199] (see also [72]). The theoretical statement was that the average squared matrix element of any operator $\hat{A}$, a 'distance' $\hbar\omega$ from the diagonal when written in the energy basis, is proportional to the ω-component of the Fourier transform of the corresponding classical auto-correlation $\langle A(0)A(t)\rangle_{\rm E}$. The

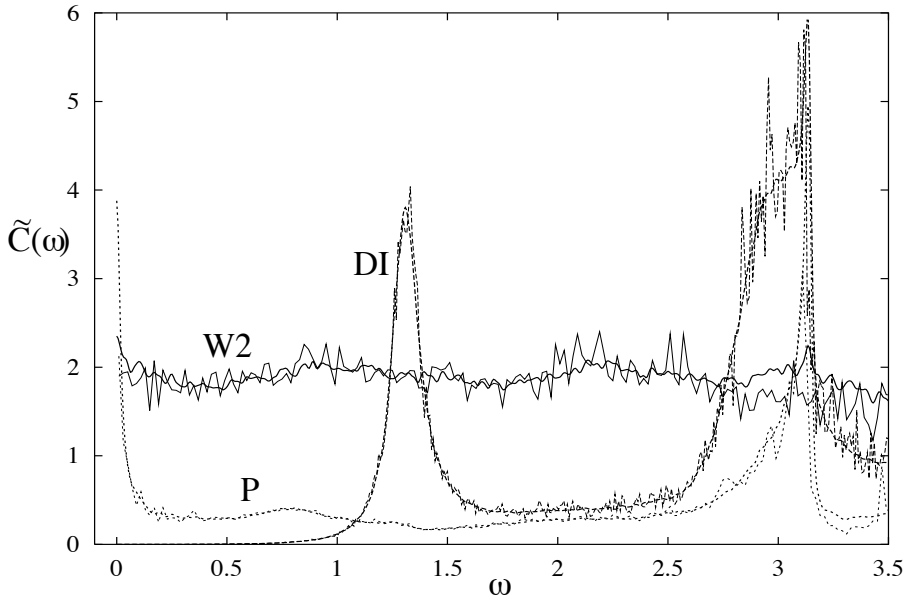

Figure 2.5: Agreement between quantum and classical band profile in the billiard of Fig. 2.6, for three example deformations: DI (dilation), W2 (periodic oscillation around the perimeter), P (wide 'piston' existing only on the top edge). In each case classical is shown as a thick line (RMS estimation error of 3%), and quantum a thin line (RMS error of 10%, increasing at higher ω); the agreement is excellent. Note that the y-axis has been displaced to clearly show the differing $\omega \to 0$ behavior. The singular peak at $\omega = \pi$ is due to the 'bouncing ball' orbit. The units are such that $m = v = 1$. The quantum calculation was performed at wavenumber $k \approx 400$ using 451 adjacent eigenstates.

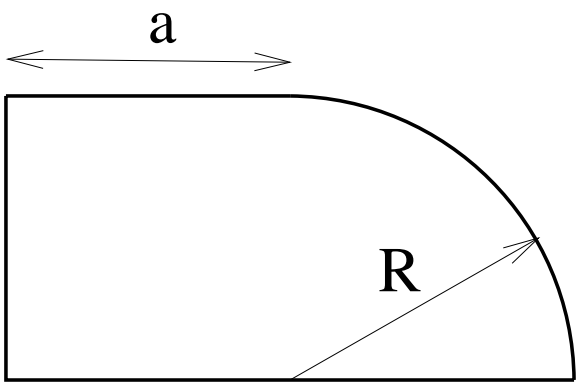

Figure 2.6: The example (undeformed) 2D billiard used for comparison of quantum and classical band profiles: the Bunimovich quarter-stadium. I typically choose $a/R = 1$ since this maximizes the average Lyapunov exponent (a measure of chaoticity). The overall size is set by $R = 1$.

typical energy is E. Convergence is achieved when averaged over a small region, as shown in Fig. 2.3, and the energy smearing ($M\Delta$) is larger than $\hbar/\tau_1$. (τ_1 is the period of the shortest period orbit). Because of this smearing over a scale $O(\hbar)$, all periodic orbit fluctuations (discussed in the following section) are averaged away.

Looking at (2.48), and choosing $\hat{A} = \mathcal{F}$, this corresponds in our notation to

$$\tilde{C}_{\rm E}^{\rm qm}(\omega) \;\approx\; \tilde{C}_{\rm E}(\omega). \tag{2.57}$$

This connection is just another statement of the correspondence principle of Shnirelman [177] (also see [91]) that in the semiclassical limit the expectation value of an operator (averaged over many adjacent states) is given by the microcanonical average of the corresponding classical function. In this case, the operator is $\mathcal{F}(0)\mathcal{F}(\tau)$, and its quantum and classical expectations are (2.46) and (2.10) respectively. The operator involves propagation forward in time. The only significant contribution comes from times $\tau \sim t_{\rm erg}$ or less, therefore *short-time correspondence* of the evolution is an additional requirement for (2.57) to hold.

If the initial energy distribution $\rho(E)$ (or p_m) is a smooth function on this $O(\hbar)$ smearing scale, then the smoothed $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ would give the resulting quantum dissipation rate, and QCC would follow. This will be the usual assumption about $\rho(E)$. However, if the system were prepared in a quantum-mechanically narrow distribution (for example a single pure state (2.39) or a Fermi distribution at temperature $k_{\rm B}T \sim \Delta$), consideration of the unsmeared structure of $|\partial\mathcal{H}/\partial x_{nm}|^2$ is required.

2.3.2 Numerical comparison of band profiles

The classical band profile (found using methods in Appendix B) and quantum band profile (methods in Appendix C) are compared in Fig. 2.5, for the two-dimensional billiard system shown in Fig. 2.6. This system was chosen because efficient diagonalization methods (see Chapter 6) exist for billiards, and a good basis set (Appendix J) is already known for this shape. The agreement is excellent, well within the expected RMS estimation errors for all ω considered. Note that there are no fitted parameters in this comparison. The three different choices of the effect of parameter x are different 'deformations' of the billiard (the subject of the following two chapters). The range of ω studied goes from zero to beyond the frequency of the shortest period orbit (the 'bouncing ball' orbit family). Note that this agreement has also been tested for other example systems [49, 48]. The conclusions can be drawn that classical correlation functions can give a good semiclassical estimate of averaged quantum matrix elements, and QCC in the LRT regime has been tested for (at least some) real chaotic systems.

2.3.3 Quantum structure beyond the band profile

For the purposes of establishing QCC for the heating rate, demonstration of band profile agreement is sufficient. In this section I discuss further structure in the matrix $(\partial\mathcal{H}/\partial x)_{nm}$, which is irrelevant for heating rate, but is interesting in itself.

The classical band profile $\tilde{C}_{\rm E}(\omega)$ is a smooth function of E, changing only on classically-significant scales ($\hbar^0$). To claim this, I must assume that there are no disconti-

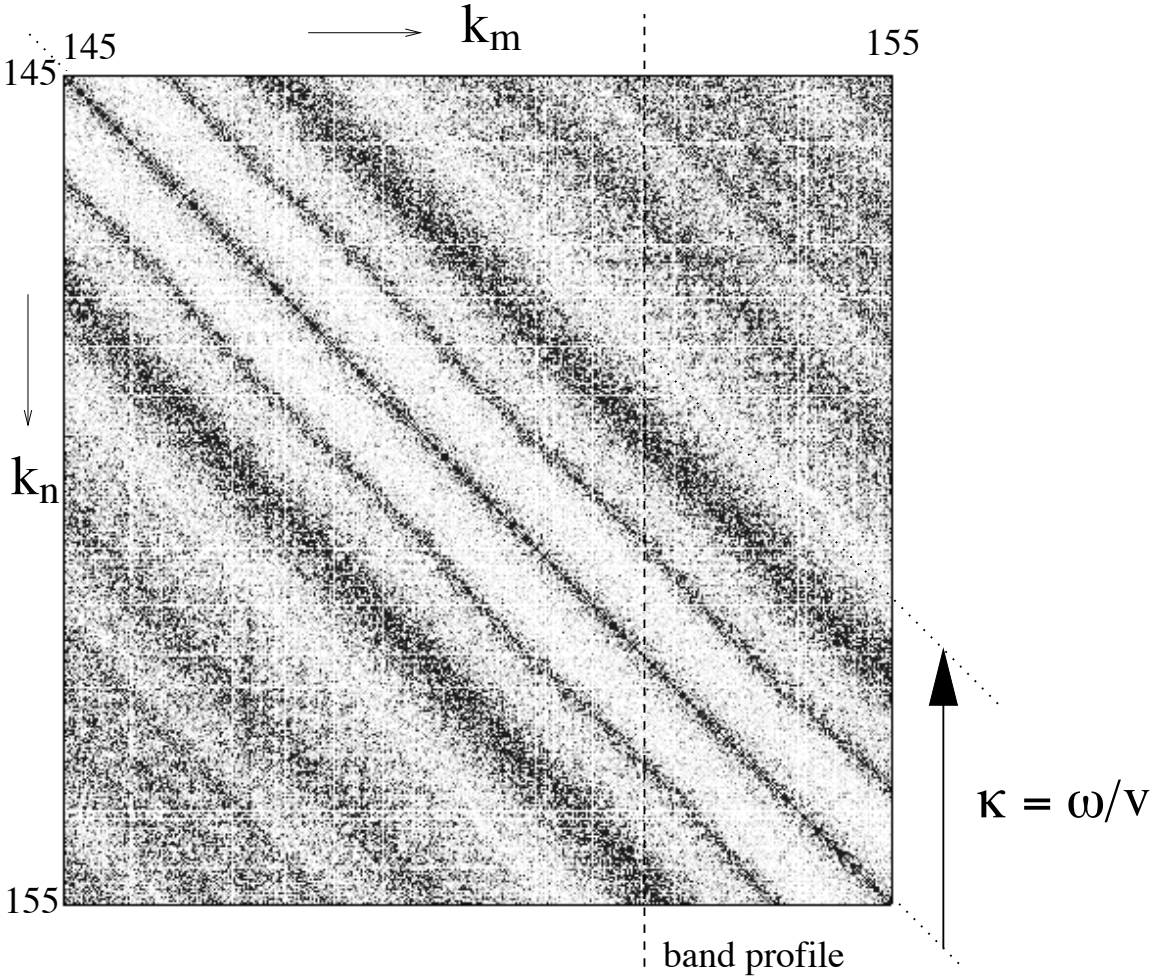

Figure 2.7: Image of the matrix $\partial\mathcal{H}/\partial x$ shown as a density plot of $|\partial\mathcal{H}/\partial x_{nm}|^2$, for the case of the 'bow' deformation of a quarter stadium billiard. This deformation is not sensitive to 'bouncing ball' orbits. Dark pixels correspond to large values. The matrix involves the 422 eigenstates lying in wavenumber range $145 < k < 155$. The band profile gives the average 'slice' as shown (also see Fig. 2.3). For $v = 1$, the units of ω and κ are the same.

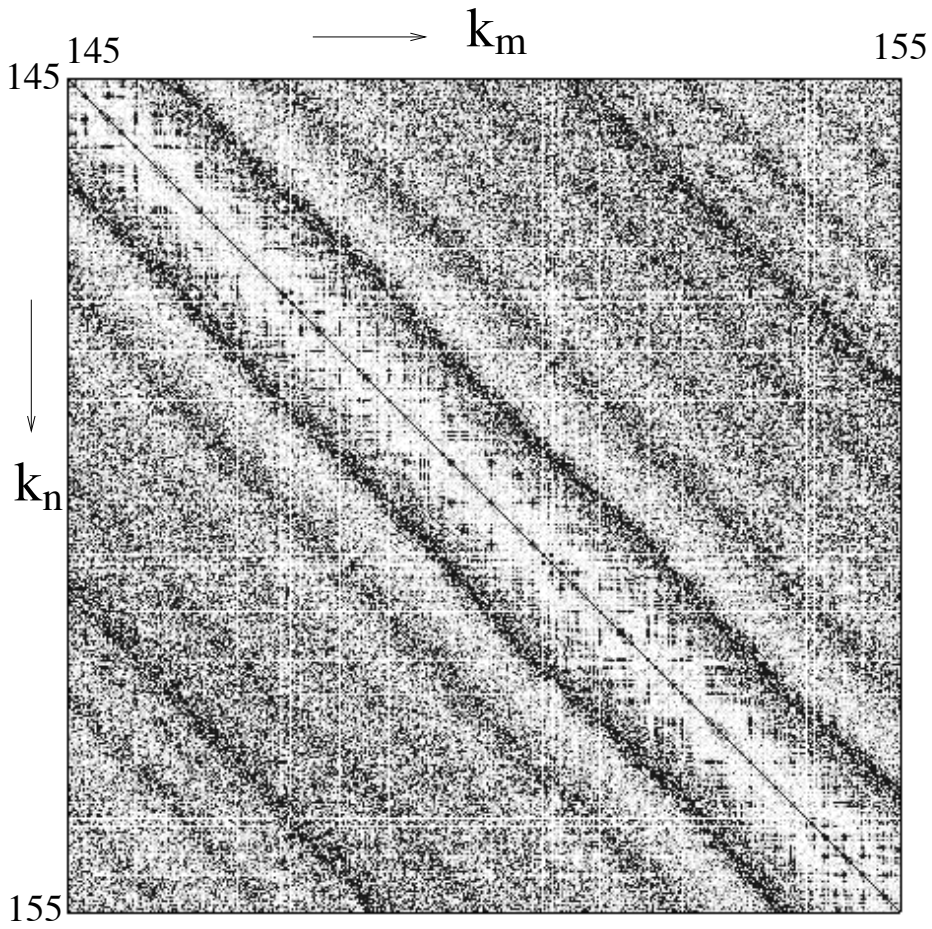

Figure 2.8: The same as Fig. 2.7, except for the case of the 'bend' deformation, which *is* sensitive to 'bouncing ball' orbits. Notice that there is much more sparse structure. In this and the previous figure, periodic dependence in E can best be viewed by holding up the page and looking along the diagonal.

nuities in the phase-space energy-shell as a function of E; such pathologies would be a new area of study.

This constrasts the quantum case, where there is much E-dependent structure before smearing is performed: Figs. 2.7 and 2.8 show images of $|(\partial\mathcal{H}/\partial x)_{nm}|^2$ for parametric changes (deformations) of the billiard of Fig. 2.6. Features to notice are,

1) The overall 'banded' shape—this is what the smeared $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ measures.

2) Generally matrix elements are random and uncorrelated. It is the local mean that follows the banded shape.

3) There is a slight wobble (fluctuation in band profile) *along* the diagonal direction, best viewed by looking down the diagonal.

4) Some matrix elements (or blocks of elements) are anomalously large (especially see Fig. 2.8), and these large elements form periodic structures along the diagonal direction. When the corresponding eigenfunctions are examined, they are often found to be *scarred* states [90].

5) Whenever there is such a large element, the remaining row and column are very small, giving a vertical and horizontal streaking effect. This is an example of *anti-scarring* [111].

I now discuss these effects further. Wilkinson [199] has extended the sum rule for the band profile to include periodic orbit effects which manifest themselves once the smoothing width $M\Delta$ is less than $\hbar/\tau_1$, where τ_1 is the period of the shortest periodic orbit. The contribution from a single orbit of period τ is structure in the matrix at a 'distance' from the diagonal (in energy units) of $2\pi\hbar/\tau$, and its integer multiples. This structure is periodic in E (*i.e.* along the diagonal direction), with period given by the action integral of the orbit. Therefore E-dependent structure up to energy scales $O(\hbar)$ exists, on top of the $O(\hbar^0)$ smooth classical changes. This oscillating correction to the off-diagonal mean value is analogous to Berry's periodic orbit correction to the mean level density [25]. The E-oscillatory peaks (effects 3 and 4 above) can be seen as an example of these oscillations. The clear period of π in the figures corresponds to the shortest periodic orbit (the 'bouncing ball' orbit of length $\tau_1 = 2$).

Scarring [90, 91, 111], a result of the classical dynamics, occurs with much more strength in the stadium than in a generic 'hard chaos' system, because of the marginally-stable bouncing ball orbit family (such non-generic effects are discussed in Section 3.4). However, it is always present to some degree even in hard chaos systems. Its effects (4 and 5 above), that of introducing matrix elements which are anomalously large compared to a Random Matrix Theory prediction, can be viewed in another way. When an integrable system is perturbed (to give a KAM system), the corresponding off-diagonal matrix elements are *sparse* [164]. Scarred states are due to the closest-to-stable classical structures (in this stadium example the main scars are due to marginally stable orbits [91]), so one might expect more sparse rows and columns at these states. Fig. 2.7 is for a deformation which is zero where the bouncing ball orbits hit the wall, and it shows little sparsity. In contrast, Fig. 2.8 deforms this part of the wall, and has much more sparsity and structure.

Why is there 'anti-scarring' (effect 5 above)? To answer this I use a simple formula

for the sum of any row or column of the squared matrix elements,

$$\sum_m \left| \left(\frac{\partial \mathcal{H}}{\partial x} \right)_{nm} \right|^2 \;=\; \langle n| \left(\frac{\partial \mathcal{H}}{\partial x} \right)^2 |n\rangle \;=\; \int d^d\mathbf{r} \, \left[\frac{\partial \mathcal{H}}{\partial x}(\mathbf{r}) \right]^2 \psi_n^2(\mathbf{r}). \tag{2.58}$$

I will be interested in how this quantity varies with n. Clearly if $\frac{\partial \mathcal{H}}{\partial x}(\mathbf{r})^2 = 1$ over all space, then the required sum is unity, independent of n. However, if the positive weighting function $\frac{\partial \mathcal{H}}{\partial x}(\mathbf{r})^2$ covers *any* classically-large region of position space then the required sum will also vary relatively little with n, because of ergodicity of almost all quantum states [91]. The result is a sum rule (note that it does include the diagonal). Unfortunately, the sum diverges in the case of hard-walled systems, because the matrix bandwidth is then infinite [46], but the idea is clear: if one matrix element has a anomalously large value then its row and column must be correspondingly small. The states that are likely to disobey this sum rule are the strongly scarred ones themselves, in the case that the choice of $\frac{\partial \mathcal{H}}{\partial x}(\mathbf{r})$ is sensitive to certain orbits and not others. Further study is needed in this area, in particular quantitative comparison to Wilkinson's predictions.

2.3.4 Effects below the quantum level spacing

Firstly, re-examining Fig. 2.3, it should be noted that if the smoothing energy range $M\Delta$ is sufficiently large, then a smooth $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ results even if the off-diagonal smoothing range $\hbar\varepsilon$ is smaller than the level spacing Δ. In this case, there is meaning to structure in $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ below the level spacing. However, the level spacing can no longer be treated as uniform, as was done in taking the continuum limit of the sum of δ-functions to reach (2.48). In other words, E_m can no longer be taken to be $m\Delta$: the exact energy levels and their spacings are now needed. In Fig. 2.3 this can be visualized by replacing the regular energy grid $(n\Delta, m\Delta)$ with the irregular grid (E_n, E_m).

This point manifests itself most clearly at distances Δ or less from the diagonal. The quantum band profile $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ is defined by a δ-function 'comb' $\sum_n |\mathcal{F}_{nm}|^2 2\pi\delta(\omega_{nm}-\omega)$ averaged over m (see Section 2.2.2). Imagine the case where $|\mathcal{F}_{nm}|^2$ has constant average. The band profile will drop to zero as $\omega \to 0$ because there is level repulsion: the likelihood of a spacing $s \equiv \hbar\omega_{nm}$ drops to zero like a power law s^β for $s \ll \Delta$. The power $\beta = 1, 2, 4$ is determined by the symmetry class of the chaotic system [35]. More precisely, one can say that $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ would be given by the 2-level correlation function $R_2(E)$ [35] at $E = \hbar\omega$. If the mean square matrix element is now given by a continuous function $\sigma^2(\omega)$ which varies with distance $\hbar\omega$ from the diagonal, one expects the band profile to be similarly modified,

$$\tilde{C}_{\rm E}^{\rm qm}(\omega) \;=\; \frac{2\pi\hbar}{\Delta} R_2(\hbar\omega)\, \sigma^2(\omega). \tag{2.59}$$

An important point is that it is $\sigma^2(\omega)$ (rather than $\tilde{C}_{\rm E}^{\rm qm}(\omega)$) which has correspondence with the classical $\tilde{C}_{\rm E}(\omega)$. This is verified numerically in Section 3.3.2. So in effect, the quantum band profile is reached by 'punching a hole' of width $\sim \Delta$ at the origin in the classical band profile. Further than Δ from the diagonal, $R_2 \to 1$ so the two become equal. Level-spacing effects on the quantum band profile are also discussed in [11].

For most of my work, the smoothing width $\hbar\varepsilon$ is larger than Δ, so the above will not have an effect.

Chapter 3

Dissipation rate in chaotic billiards, and 'special' deformations

In this chapter I shall be calculating the dissipation rate resulting from driving a chaotic billiard (also known as a cavity) system containing a single particle or gas of non-interacting particles. The billiard system is in d-dimensional space, and is entirely defined by the location of its closed, hard wall— 'driving' will means the parameter x moves, or *deforms*, this wall, according to a 'deformation function' $D(\mathbf{s})$.

What is the rate in which the 'gas' inside the cavity is heated up? The answer depends on the shape of the cavity, the deformation $D(\mathbf{s})$ involved, as well as on the amplitude A and the driving frequency ω. Also the number of particles N and their energy distribution $\rho(E)$ should be specified. To reach an answer I shall be using the theory outlined in the previous chapter, where it was explained that the dissipation rate due to driving at frequency ω is proportional to a correlation power spectrum $\tilde{C}_{\mathrm{E}}(\omega)$, in both classical and quantum linear response. Hence $\tilde{C}_{\mathrm{E}}(\omega)$, also known as the 'band profile', will now be the main object of study. Of particular interest is its zero-frequency limit ν_{E}. $\tilde{C}_{\mathrm{E}}(\omega)$ will take different forms for the case of different deformations and for different cavity shapes— I will be interested in general deformations which need not preserve the cavity shape nor its volume. I also assume shapes such that the motion of the particle inside the cavity is globally chaotic, meaning no mixed phase space [154]. The criteria for having such a cavity are discussed in [40, 204]. For validity of linear response, the slowness condition of (2.12) is assumed, which in the billiard case becomes $V \ll v_{\mathrm{E}}$, that is the speed of wall movement should be much less than the particle speed at the energy E.

I will introduce the white noise approximation (WNA), which uses a strong chaos assumption to give an estimate for $\tilde{C}_{\mathrm{E}}(\omega)$. In the nuclear application ($d = 3$) this leads to the so-called 'wall formula'. I will then compare computed $\tilde{C}_{\mathrm{E}}(\omega)$ curves to the WNA prediction, and find many deformations for which the WNA fails. In particular, the main result will be the discovery of a class of deformations which have $\tilde{C}_{\mathrm{E}}(\omega)$ vanishing as various powers of ω in the zero-frequency limit, which I name 'special' deformations. This class is the set of deformations that are shape-preserving: they involve only translations, rotations and dilations of the cavity. Note that translations and rotations are also volume-preserving, in which case the associated time-dependent deformations can be described as 'shaking' the

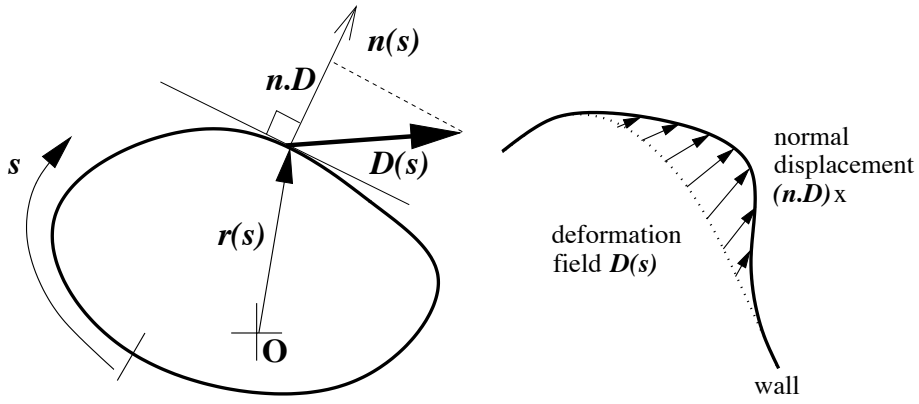

Figure 3.1: Left: Coordinates and local outward normal direction. The boundary deformation function is $D(\mathbf{s}) \equiv \mathbf{n}\cdot\mathbf{D}$. Right: Action of deformation field on wall for finite parameter x (the undeformed wall for $x = 0$ is shown by a dotted line).

cavity. The special class is important for three reasons:

1. It will provide the basis for the improved WNA estimate (improved 'wall formula') for $\nu_{\rm E}$, which takes into account correlations beyond the strong chaos assumption. This is the subject of Chapter 4.

2. The 'special' nature of the deformation involving dilation about an arbitrary origin corresponds to quasi-orthogonality of the cavity quantum eigenstates on the boundary. This in turn has a very direct connection to the success of the numerical method of Vergini and Saraceno [195], the subject of Chapter 6.

3. They have other applications: the band profile of certain special deformations is simply related to the band profile for driving the (charged) particle by uniform external fields. For instance, the mesoscopic version of Drude formula for the *conductance* of a quantum dot in a uniform time-dependent magnetic field reduces to the the calculation of $\tilde{C}_{\rm E}(\omega)$ for the rotation deformation. For more details, I refer the reader to our published work [15].

The Appendices B and C detail the numerical methods used for classical and quantum band profile calculations in this and the following two chapters.

3.1 The cavity system

Consider a single particle whose canonical coordinates are $(\mathbf{r}, \mathbf{p})$ moving inside a cavity of d . The Hamiltonian is

$$\mathcal{H}(\mathbf{r}, \mathbf{p}; x) = \mathbf{p}^2/2m + U(\mathbf{r} - x\mathbf{D}(\mathbf{r})) \tag{3.1}$$

where $U(\mathbf{r})$ is the confining potential. I have introduced a (unitless) deformation 'field' $\mathbf{D}(\mathbf{r})$, so the effect of changing the parameter x is to distort the potential in space. We

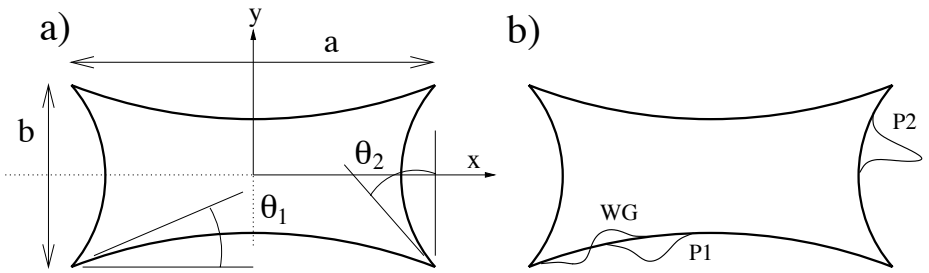

Figure 3.2: **a**) The example (undeformed) cavity used for numerical studies (unless otherwise stated): a generalized two-dimensional Sinai billiard formed from concave arcs of circles with two different radii. Typical parameters used are a=2, b=1, θ_1=0.2, θ_2=0.5, for which the average collision rate with the wall is $(1/\tau_{\rm bl}) \approx 0.63$. **b**) Sketches of the effect of three of the deformation types on the perimeter (here we have chosen three localized deformations; see Tables 3.1 and 3.2 for functional forms of all deformations used). The deformations are shown exaggerated in strength.

assume that $U(\mathbf{r}) = 0$ inside the cavity. Outside the cavity the potential $U(\mathbf{r})$ becomes very large. To be specific, one may assume that the walls exert a normal force f, and we take the hard wall limit $f \to \infty$. With the above assumptions about $U(\mathbf{r})$ it is clear that the deformation is completely specified by the scalar boundary deformation function $D(\mathbf{s}) \equiv \hat{\mathbf{n}}(\mathbf{s}) \cdot \mathbf{D}(\mathbf{s})$, where $\hat{\mathbf{n}}(\mathbf{s})$ is an outwards unit normal vector at the boundary point $\mathbf{s}$. This is shown in Fig. 3.1. The surface area of the cavity is A. Its volume is V, which is related to its typical length L by $\mathsf{V} = \mathsf{L}^d$. Quantum-mechanically, a second length scale $\lambda_{\rm B} \equiv 2\pi/k$ appears, where k is the wavenumber. The other parameters v (particle velocity) and m (particle mass), as well as $\hbar$, could be scaled away with the appropriate choice of units (this will only be done in numerical tests where we set $m = v = \hbar = 1$). The energy is $E = \frac{1}{2}mv^2$. Sometimes I will use $v_{\rm E}$ to denote v corresponding to energy E. Upon quantization the eigenenergies, which in general are x-dependent, are $E_n = (\hbar k_n)^2/2m$.

My numerical calculations of the band profile in this and the following chapter, unless otherwise stated, will refer to the two-dimensional (2D) cavity illustrated in Fig. 3.2, which we call the generalized Sinai billiard. The shape has been chosen because it has 'hard chaos': no mixed phase space, and absence of any marginally-stable orbits (see Section 3.4). In Fig. 3.2b we show three example deformations of this billiard.

The band profile calculations will generally be done classically (Appendix B), since this is easier by far than the quantum calculation (Appendix C). The two have been demonstrated equivalent in Section 2.3. In Section 3.3.2 I will verify this equivalence for special deformations.

3.1.1 Form of correlation spectrum and timescales

As presented in Section 2.1, the generalized time-dependent 'force' associated with the parameter x is $\mathcal{F}(t) = -\partial\mathcal{H}(\mathbf{r},\mathbf{p};x)/\partial x - F(x)$, where $F(x)$ is a constant such that the time-average is zero. It has the t-independent auto-correlation function

$$C_{\rm E}(\tau) \;\equiv\; \langle \mathcal{F}(t)\mathcal{F}(t+\tau)\rangle_{\rm E} \tag{3.2}$$

The subscript E, whenever used, suggests that the average over initial conditions is of microcanonical type, with energy E. For the Hamiltonian (3.1) we can write

$$-\frac{\partial \mathcal{H}}{\partial x}(t) \;=\; \mathbf{D}(\mathbf{r})\cdot\nabla U(\mathbf{r}) = -\mathbf{D}(\mathbf{r})\cdot\dot{\mathbf{p}}. \tag{3.3}$$

Recognizing $\dot{\mathbf{p}}$ as the force on the gas particle, we see that $\mathcal{F}(t)$ is a train of spikes superimposed on a constant average value (see Fig. 3.3a):

$$\mathcal{F}(t) \;=\; \sum_i 2mv\;\cos(\theta_i)\;D_i\;\delta(t-t_i)\;-\;F(x) \tag{3.4}$$

where i labels collisions: t_i is the time of a collision, D_i stands for $D(\mathbf{s}_i)$ at the location $\mathbf{s}_i$ of a collision, and $v\cos(\theta_i)$ is the normal component of the particle's collision velocity. If the deformation is volume-preserving then $\langle\mathcal{F}(t)\rangle = 0$, otherwise it is convenient to subtract the (constant) average value $F(x)$.

The auto-correlation function $C_{\rm E}(\tau)$ can be handled as a time-average of a single trajectory rather than an ensemble-average over trajectories (by ergodicity, as discussed in Section 2.1.6); the resulting construction is illustrated in Fig. 3.3b. The forms of $C_{\rm E}(\tau)$ and its Fourier transform $\tilde{C}_{\rm E}(\omega) \equiv \int C_{\rm E}(\tau)\exp(i\omega\tau)d\tau$ are illustrated schematically in Figs. 3.3c and 3.3d. The removal of the average value of $\mathcal{F}(t)$ ensures that there is no δ-function spike at $\omega = 0$ in $\tilde{C}_{\rm E}(\omega)$. This is required so that $\tilde{C}_{\rm E}(\omega)$ has a well-defined $\omega \to 0$ limit, namely the fluctuations intensity $\nu_{\rm E}$.

The auto-correlation function $C_{\rm E}(\tau)$ consists of a $\tau = 0$ ('self') peak due to the self-correlation of the spikes, and of an additional smooth ('non-self') component due to correlations between successive bounces. (Note that in the hard walled limit the constant value $F(x)$ contributes nothing to the 'self' $\delta(\tau)$ peak, so its only effect is on the 'non-self' component). Thus two time scales are involved: the short time scale is $\tau_0 = 2mv/f \to 0$ in the hard wall limit, and the other time scale is $\tau_{\rm col}$, which involves the time for succesive collisions with the deforming part of the boundary. Note that $\tau_{\rm col}$ can be much larger than the ballistic time $\tau_{\rm bl}$ (the average time to cross the billiard) in the case that only a small piece of the boundary is being deformed. The quantitative definition of the collision rate $1/\tau_{\rm col}$ is postponed to Section 3.2.

I shall be most interested in the noise intensity $\nu_{\rm E}$ defined by (2.14). Observing that $\mathcal{F}(t)$ is linear in $D(\mathbf{s})$, it follows that the noise intensity must (exactly) be a general quadratic functional

$$\nu_{\rm E} \;=\; \oint\!\!\oint d\mathbf{s}_1 d\mathbf{s}_2\, D(\mathbf{s}_1)\gamma_{\rm E}(\mathbf{s}_1,\mathbf{s}_2)D(\mathbf{s}_2), \tag{3.5}$$

where the kernel $\gamma_{\rm E}$ depends on both the cavity shape and the particle energy E [120]. Furthermore, billiards are *scaling systems* in the sense that a change in E leaves the trajectories unchanged. From this and (3.4) we have the scaling relation $\gamma_{\rm E}(\mathbf{s}_1,\mathbf{s}_2) = m^2 v_{\rm E}^3\cdot\hat{\gamma}(\mathbf{s}_1,\mathbf{s}_2)$, where the scaled kernel depends entirely on the geometrical shape of the cavity. However, the reason for being interested in *approximations* for $\nu_{\rm E}$ is that the exact result for the kernel $\hat{\gamma}$ is not analytic: it is very complicated to evaluate, and involves a sum over all classical paths from $\mathbf{s}_1$ to $\mathbf{s}_2$ (see [120]).

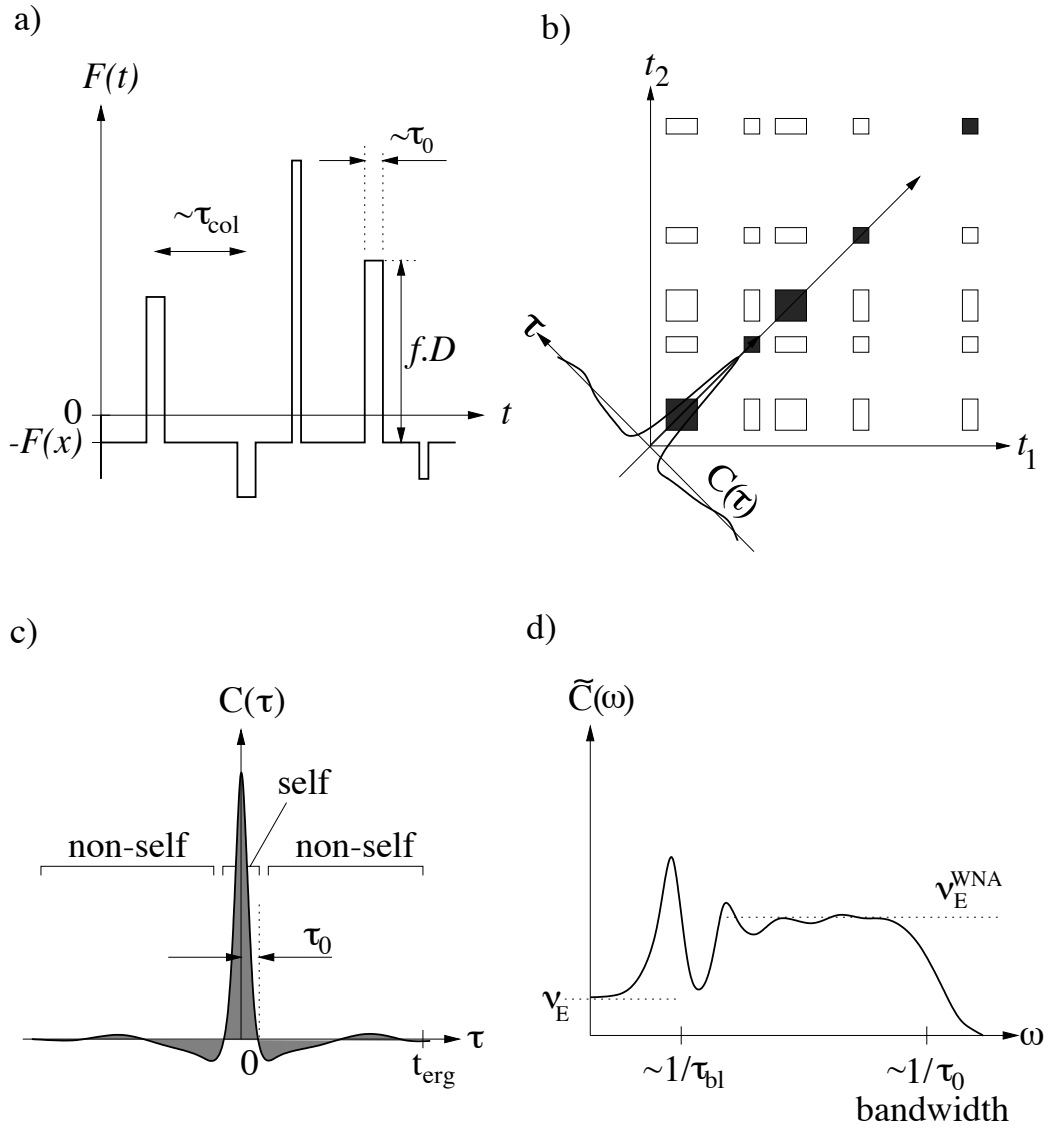

Figure 3.3: **a**) The fluctuating force $\mathcal{F}(t)$ is a series of impulses of maximum duration τ_0. In the hard-walled limit $\tau_0 \to 0$. The average value $F(x)$ has been subtracted. **b**) The autocorrelation function $C(\tau)$ can be viewed as the average projection of the function $\mathcal{F}(t_1)\mathcal{F}(t_2)$ onto the $\tau \equiv t_2 - t_1$ axis (the 'self' components are shaded). **c**) The autocorrelation function has a 'self' peak of half-width τ_0, and becomes negligible beyond the ergodic time $t_{\rm erg}$. **d**) The noise power spectrum $\tilde{C}(\omega)$. Its $\omega \to 0$ limit is equal to the area under $C(\tau)$. The power spectrum typically shows billiard-specific features around the frequency $\omega \sim 1/\tau_{\rm bl}$, where $\tau_{\rm bl}$ is the ballistic time. Note that in some examples $\tau_{\rm col}$ is much longer than $\tau_{\rm bl}$.

3.1.2 Conversion of time averages to averages over collision parameters

This section shows how ergodicity can be used to go from a time average to a collision parameter average, for a quantity which takes values at each of the collisions. This will prove useful for approximating $\nu_{\rm E}$ in the next section, but it can be skipped on first reading. (It is presented in a more terse form in [46]).

I treat time-averages of quantities which are non-zero only during the collisions, and whose value during a collision is a function of the parameters of the collision (angle from the normal, and surface location) only. They can be converted to collision-parameter averages, eliminating the time variable, as follows. For a hard-walled billiard, imagine a function of phase-space

$$B(\mathbf{r},\mathbf{p}) \;\equiv\; \begin{cases} |\cos\theta|\; b(\mathbf{s},\theta) & \text{if } z > -\epsilon \\ 0 & \text{otherwise,} \end{cases} \tag{3.6}$$

which exists only in the thin shell of thickness ϵ near the wall. The position vector $\mathbf{r}$ is represented by $(\mathbf{s}, z)$ where $\mathbf{s}$ is a surface coordinate and z an outward 'radial' coordinate normal to the surface. θ is the angle of the $\mathbf{p}$ vector from the local normal. This leaves B independent of the other $d-1$ momentum degrees of freedom. The function $b(\mathbf{s},\theta)$ is assumed constant upon reflection, *i.e.* it is even about $\theta = -\pi/2$. The time-dependence of B along a given trajectory will be a series of impulses due to passing through the shell; the duration of impulse i will be $2\epsilon/(v_{\rm E}\cos\theta_i)$, and during this time its value is $|\cos\theta_i|\, b(\mathbf{s}_i,\theta_i)$. Therefore the cos factors cancel, and the time-average is

$$\langle B\rangle_t \;=\; \frac{2\epsilon}{v_{\rm E}} \left\langle \sum_i b(\mathbf{s}_i,\theta_i)\; \delta(t-t_i) \right\rangle_t . \tag{3.7}$$

By ergodicity one can equate this to the phase-space average of (3.6) at the energy E,

$$\langle B\rangle_{\rm E} \;=\; \frac{\epsilon \mathsf{A}}{\mathsf{V}} \langle |\cos\theta|\; b(\mathbf{s},\theta)\rangle_{\mathbf{s},\Omega} . \tag{3.8}$$

Here I used the fact that in the $\epsilon \to 0$ limit the fraction of position space occupied by the shell is $\epsilon\mathsf{A}/\mathsf{V}$. This collapses the average to be over $\mathbf{s}$ and over the $d-1$ dimensional solid angle Ω. Choosing the function $b = 1$ gives a useful expression for the mean collision rate (the inverse of the ballistic time),

$$\tau_{\rm bl}^{-1} \;\equiv\; \left\langle \sum_i \delta(t-t_i) \right\rangle_t \;=\; \frac{v_{\rm E}\mathsf{A}}{2\mathsf{V}}\; \langle |\cos\theta|\rangle, \tag{3.9}$$

which is a form of Sabine's formula[1]. The geometrical factor is the solid angle average $\langle |\cos\theta|\rangle = 2/\pi, 1/2, \cdots$ for dimension $d = 2, 3, \cdots$ (see Appendix D of [46]). Or, choosing $b = \cos^2\theta\; D(\mathbf{s})^2$ gives

$$\left\langle \sum_i \cos^2\theta_i\; D(\mathbf{s}_i)^2\; \delta(t-t_i) \right\rangle_t \;=\; \frac{v_{\rm E}\mathsf{A}}{2\mathsf{V}}\; \langle |\cos\theta|^3\rangle\, \langle D(\mathbf{s})^2\rangle_{\mathbf{s}} \tag{3.10}$$

[1] This formula was first found (in $d = 2, 3$) by Wallace Clement Sabine in 1898, in his studies of Harvard's own Jefferson Laboratory main auditorium, room 250. Sabine used this expression for the mean free path to give a formula for the reverberation time of a hall [21, 173].

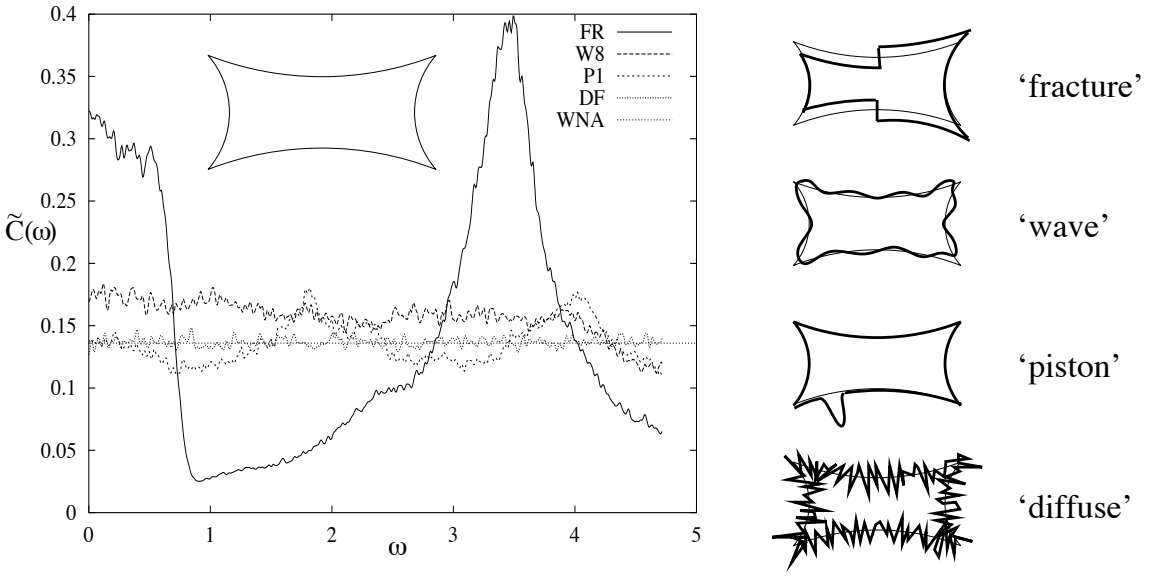

Figure 3.4: The white-noise approximation estimate (WNA is the horizontal dotted line) compared to actual $\tilde{C}_{\rm E}(\omega)$ power spectra for some example deformations of the 2D generalized Sinai billiard, with $m = v = 1$. (The RMS estimation error of 3% can be seen as multiplicative noise with short correlation length in ω). The deformations (sketched on the right), from worst to best agreement of the WNA are: FR (for which sensitivity to the least-unstable vertical periodic orbit causes large correlation effects and large deviations from WNA), W8 (oscillatory deformation changes sign enough to be close to WNA), P1 (localized 'piston' type deformation for which WNA is good), and DF (random function of zero correlation-length along the perimeter, showing complete WNA agreement). Deformation functions are defined in Table 3.1.

where we used the fact that for this choice the averages over surface location and solid angle separate. Again, formulae exist for the dimension-dependent geometric factor $\langle|\cos\theta|^3\rangle = 4/(3\pi), 1/4, \cdots$ for $d = 2, 3, \cdots$ [46].

3.2 The white noise approximation (WNA)

The most naive estimate of the fluctuations intensity is based on the 'white noise approximation' (WNA). Namely, one assumes that the correlation between bounces can be neglected. This corresponds [120] to the *local* part of the kernel (3.5). In such a case only the self-correlation of the spikes $C_{\rm E}^{\rm self}(\tau)$ is taken into consideration. Using (3.2) and (3.4) one obtains

$$\nu_{\rm E} \;\approx\; \nu_{\rm E}^{\rm WNA} \;=\; \int_{-\infty}^{\infty} d\tau\, C_{\rm E}^{\rm self}(\tau) \;=\; (2mv_{\rm E})^2 \left\langle \sum_i \cos^2(\theta_i)\; D_i^2\; \delta(t-t_i) \right\rangle_{\rm E}, \qquad (3.11)$$

where the microcanonical average can then be replaced by the time average. (The same expression also results from considering the time-averaged contribution of the shaded parts

key	*description*	*surface deformation function* $D(s)$
CO	constant	1
Wn	n periods	$\cos(2\pi n s/L)$
DF	diffuse	random[-1,1] (equivalent to W∞)
FR	fracture	$\mathrm{sgn}(x(s))$ if on top or bottom, else 0
SX	shift-x	$\mathrm{sgn}(x(s))$ if on left or right, else 0
P1	piston 1	$10\exp(-\frac{1}{2}\alpha^2)$, $\alpha=(s/L-0.3)/0.01$
P2	piston 2	$10\exp(-\frac{1}{2}\alpha^2)$, $\alpha=(s/L-0.6)/0.005$
WG	wiggle	$5\alpha\exp(-\frac{1}{2}\alpha^2)$, $\alpha=(s/L-0.25)/0.02$

Table 3.1: Key to deformation types used for numerical 2D billiard experiments in this and the following chapter. L is the billiard perimeter. The deformation is described by a function $D(s)$, where s is measured counter-clockwise along the perimeter with $s = 0$ at the upper left corner. In the 'fracture' and 'shift-x' cases we use the horizontal Cartesian coordinate $x(s)$.

in Fig. 3.3b). From here the result (3.10) can be used, giving the WNA result

$$\nu_{\mathrm{E}}^{\mathrm{WNA}} \;=\; 2m^2 v_{\mathrm{E}}^3 \,\langle|\cos\theta|^3\rangle \frac{1}{\mathsf{V}} \oint [D(\mathbf{s})]^2 d\mathbf{s}, \tag{3.12}$$

where the geometric factor $\langle|\cos\theta|^3\rangle$ is given above. If we can use the convention $|D(\mathbf{s})| \sim 1$ over the deformed region (and zero otherwise), then we can write the WNA as $\nu_{\mathrm{E}}^{\mathrm{WNA}} = (2mv_{\mathrm{E}})^2 \times (1/\tau_{\mathrm{col}})$ which defines $(1/\tau_{\mathrm{col}})$ as the effective collision rate (for more discussion of this convention see Appendix F of [46], and [48]). Again, τ_{col} can be much larger than the ballistic time τ_{bl} in the case that only a small piece of the boundary is being deformed.

For hard walls, the band profile $\tilde{C}_{\mathrm{E}}(\omega)$ is non-zero for all omega, corresponding to an infinite bandwidth of the matrix $(\partial\mathcal{H}/\partial x)_{nm}$. The band profile at $\omega \gg \tau_{\mathrm{bl}}^{-1}$ is entirely due to the 'self' component, so is constant and given by the WNA (see Fig. 3.3d). This convergence at large ω has been tested numerically: the decay to the constant WNA value is found to be quite oscillatory.

The use of the WNA can be justified whenever successive collisions are effectively uncorrelated. The applicability of such an assumption depends on the shape of the cavity (which will determine the decay of correlations via the typical Lyapunov exponent) as well as on the type of deformation involved. If we have the cavity of Fig. 3.2a, and the deformation involves only a small piece of the boundary (*e.g.* see Fig. 3.2b), then successive collisions with the *deformed part* of the boundary are effectively uncorrelated. This is so because there are many collisions with static pieces of the boundary before the next effective collision (with non zero D_i) takes place. If the deformation involves a large piece (or all) of the boundary, one can still argue that successive collisions are effectively uncorrelated provided $D(\mathbf{s})$ is 'oscillatory' enough (*i.e.* changes sign many times along the boundary).

These expectations are qualitatively confirmed by the numerical results of Fig. 3.4, where I show a sequence of deformations for which the WNA performs increasingly well. The numerical method and estimation error for finding the classical $\tilde{C}_{\mathrm{E}}(\omega)$ are presented in Appendix B.

key	*description*	*deformation field*
DI	dilation about origin	$\mathbf{D}(\mathbf{r}) = \mathbf{r}$
TX	x-translation	$\mathbf{D}(\mathbf{r}) = \mathbf{e}_x$
TY	y-translation	$\mathbf{D}(\mathbf{r}) = \mathbf{e}_y$
RO	rotation about origin	$\mathbf{D}(\mathbf{r}) = \mathbf{e}_z \times \mathbf{r}$

Table 3.2: Key to the four 'special' deformations in 2D. The unit vectors $\mathbf{e}_x$ and $\mathbf{e}_y$ are in the plane (see Fig. 1), and $\mathbf{e}_z$ is in the perpendicular direction. In the case of dilation and rotation $\mathbf{D}$ could be made unitless by dividing by a constant length.

3.2.1 Relation to random wave approximation

The uncorrelated impulse picture of the WNA is an intuitive estimate for the classical band profile $\tilde{C}_{\mathrm{E}}(\omega)$. However, it gives little understanding in the quantum case. The quantum band profile expresses off-diagonal structure in the matrix $(\partial\mathcal{H}/\partial x)_{nm}$. The matrix elements are given in a hard-walled system by Eq.(C.2), the weighted overlap of eigenstates n and m on the boundary. Therefore a first approximation to the quantum band profile might be reached by assuming that the eigenstates are random sums of plane waves (Berry's postulate). Furthermore one could assume no correlations between states $n \neq m$. The number of independent λ_{B}-sized patches on the billiard surface is $N \sim k^{1-d}\mathsf{A}$, and these will add in a $N^{1/2}$ fashion because they are uncorrelated. Over each patch the typical squared eigenfunction normal derivative is k^2/V. Combining with (2.48) and the fact that in d dimensions the mean level spacing is $\Delta \sim \hbar k^{2-d}\mathsf{V}/m$ gives

$$\tilde{C}_{\mathrm{E}}^{\mathrm{qm}}(\omega) \;\sim\; m^2 v_{\mathrm{E}}^3 \frac{1}{\mathsf{V}}\,\mathsf{A}D^2, \tag{3.13}$$

where D is the typical value of $D(\mathbf{s})$. Note that the predicted band profile is flat (independent of ω) because the overlap of random waves on the boundary does not depend strongly on their wavenumber difference. Also note that $\hbar$ does not appear in this quantum estimate. The similarity to (3.12) is clear. Performing the above calculation more carefully with the correct prefactors (this lengthy result is derived in [46]) gives exactly (3.12). So, remarkably, a random wave estimate in quantum mechanics is equivalent to the WNA in classical mechanics, as far as the naive band profile prediction is concerned.

3.3 'Special' deformations

The WNA dramatically fails (see Fig. 3.5) for dilation, translations and rotations (see Table 3.2 for their definitions in 2D). It is not surprising that the WNA is 'bad' for these deformations because their $D(\mathbf{s})$ are slowly-changing delocalized functions of s. However, what is remarkable is that $\tilde{C}_{\mathrm{E}}(\omega)$ for this type of deformations *vanishes* in the limit $\omega \to 0$. Such deformations we would like to call 'special' [14]. Generally, we would like to define a a deformation as 'special' if the associated fluctuation intensity is $\nu_{\mathrm{E}} = 0$.

A special result that follows from the considerations of Appendix E is that a linear combination of special deformation is also special. Therefore the special deformations constitute a linear space of functions. We believe that this linear space is spanned by the

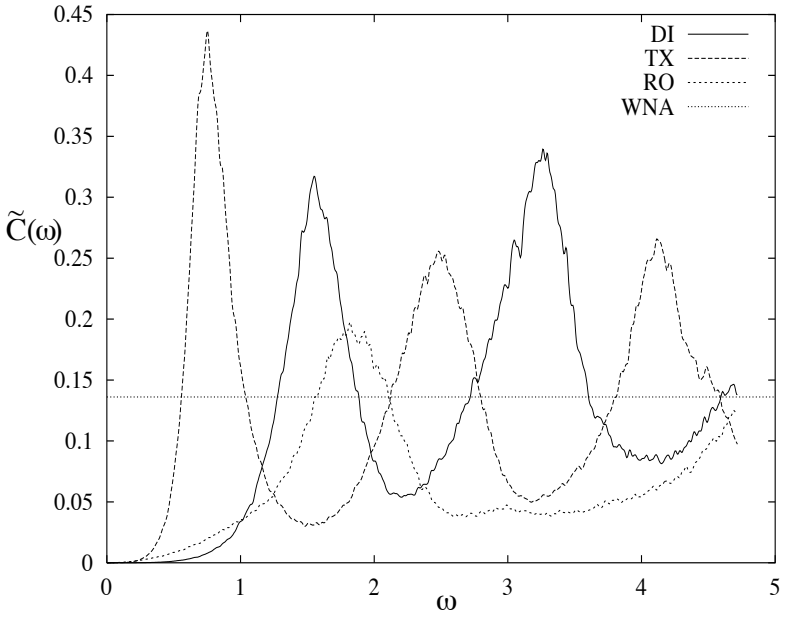

Figure 3.5: The WNA estimate compared to actual $\tilde{C}_{\rm E}(\omega)$ power spectra for example 'special' deformation types: DI (dilation), TX (translation) and RO (rotation). See Table 3.2 for definitions. The WNA fails to predict the vanishing in the small ω limit.

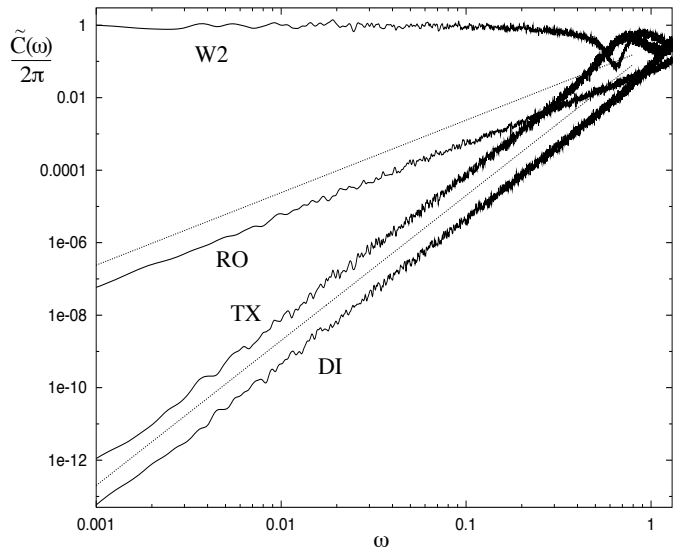

Figure 3.6: Similar to Fig. 3.5 except a log-log plot. This demonstrates the special deformation power laws. The two dotted lines show ω^2 and ω^4 frequency dependence, for purposes of comparison. A non-special deformation (W2) is also shown to contrast its small-ω dependence. Estimation error here is 13% for W2 and RO, 20% for DI and TX.

following basis functions: one dilation, d translations, and $d(d-1)/2$ rotations. However we are not able to give a rigorous mathematical argument that excludes the possibility of having a larger linear space. This is discussed in Appendix D (and the generalization to arbitrary potential $U(\mathbf{r})$ is presented). In other words, we believe that any special deformation can be written as a linear combination of dilation, translations and rotations.

3.3.1 Band profile power laws

Now I explain the observed $\nu_{\rm E} = 0$ for special deformations, using classical considerations. The interested reader should also consult Appendix D. I start with the case of translations and dilations. For translations we have $\mathbf{D} = \mathbf{e}$, where $\mathbf{e}$ is a constant vector that defines a direction in space. We can write $\mathcal{F}(t) = (d/dt)^2\mathcal{G}(t)$ where $\mathcal{G}(t) = -m\mathbf{e}\cdot\mathbf{r}$. A similar relation holds for dilation $\mathbf{D} = \mathbf{r}$ with $\mathcal{G}(t) = -\frac{1}{2}m\mathbf{r}^2$. It follows that $\tilde{C}(\omega) = \omega^4\tilde{C}_G(\omega)$, where $\tilde{C}_G(\omega)$ is the power spectrum of $\mathcal{G}(t)$. If $\tilde{C}_G(\omega)$ is a bounded function (as it must be when correlations are short-range), it immediately follows that $\tilde{C}(0) = 0$. Moreover since $\mathcal{G}(t)$ is a simple function of the particle position, we can assume it is a fluctuating quantity that looks like white noise on timescales $> t_{\rm erg}$. It follows that $\tilde{C}(\omega)$ is generically characterized by ω^4 behavior for either translations or dilations.

I now consider the case of rotations. For rotations we have $\mathbf{D} = \mathbf{e}\times\mathbf{r}$, and we can write $\mathcal{F}(t) = (d/dt)\mathcal{G}(t)$, where $\mathcal{G}(t) = -\mathbf{e}\cdot(\mathbf{r}\times\mathbf{p})$, is a projection of the particle's angular momentum vector [2]. Consequently $\tilde{C}(\omega) = \omega^2\tilde{C}_G(\omega)$. Assuming the angular momentum is a fluctuating quantity that looks like white noise on timescales $> t_{\rm erg}$, we expect that $\tilde{C}(0) = 0$ and that $\tilde{C}(\omega)$ is generically characterized by ω^2 behavior.

Thus we have predictions for the power-laws in the regime $\omega < 1/t_{\rm erg}$ for special deformations (assuming hard chaos). This contrasts the generic case of tending to a constant, that is, ω^0 behavior. These power laws are demonstrated in Fig. 3.6, and have been numerically verified over more than 4 decades in ω. For an estimate of the prefactor for the dilation case, see Section 6.1.2.

For special deformations we have $\tilde{C}(\omega) = 0$ in the limit $\omega = 0$, and consequently the dissipation coefficient vanishes ($\mu = 0$). It should be noted that for the case of a general combination of translations and rotations this result follows from a simpler argument (one which does not rely on the LRT assumption considered in [120, 118]). Taking $\Omega \to 0$ while keeping $A\Omega$ constant corresponds to constant deformation velocity ($\dot{x}$ =const). Transforming the time-dependent Hamiltonian into the reference frame of the cavity (which is uniformly translating and rotating with constant velocity) gives a *time-independent* Hamiltonian. In the new reference frame the energy is a constant of the motion, which implies that the system cannot absorb energy (no dissipation effect), and hence we must indeed have $\mu = 0$.

3.3.2 The quantum band profile case

The understanding of the special deformation class came from classical arguments. We might wonder, is this special property is preserved for the quantum band profile? The

[2] The cross-product form used here for $\mathbf{D}$ and $\mathcal{G}(t)$ is strictly valid in 2D and 3D only. For $d > 3$ the higher-dimensional generalization of a general rotation should be used.

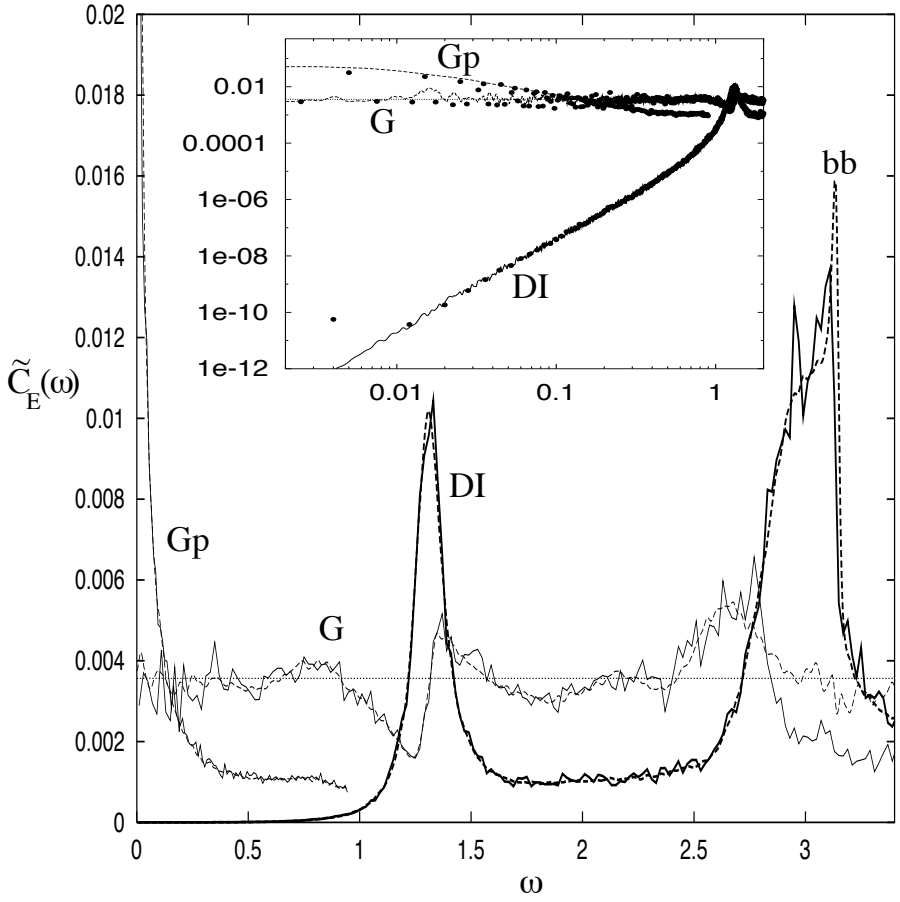

Figure 3.7: The quantum band profile in the 2D quarter stadium at $k \approx 400$ for three choices of deformation: dilation (DI), a generic deformation (G), and a generic deformation restricted to parallel displacement of the stadium upper edge (Gp). In each case, the solid line is the quantum band profile (estimation error 10%); the classical is also shown with a dashed line for comparison. I normalized G and Gp so that they share the same $\nu_{\rm E}^{\rm WNA}$ estimate as DI; this is shown as a horizontal dotted line. The system-specific peak due to 'bouncing-ball' orbits is labelled (bb). The inset is a log-log plot with quantum band profile shown as points.

answer is yes: for instance the graphs labelled DI in Fig. 3.7 show that the special nature of dilation is preserved to the accuracy of the quantum calculation. (The corresponding $(\partial\mathcal{H}/\partial x)_{nm}$ matrix is shown in Fig. 6.2). Here I used the quarter stadium billiard (Fig. 2.6) for the simple reason that efficient quantization methods exist for this shape (Chapter 6), but, as yet, do not exist for the generalized Sinai billiard (Fig. 3.2). Agreement with the classical result, and hence the power law, is maintained down to the point at which errors in the quantum calculation become dominant—this is visible as bottoming-out in the leftmost point of the log-log inset plot, at $\sim 10^{-10}$. Therefore the QCC demonstrated in Section 2.3 also holds for special deformations. This will have profound consequences for the numerical method presented in Chapter 6.

Is there a simpler direct route to the special power-law band profile behavior involving quantum-mechanical (wave) considerations alone? Koonin *et al.* [118] have derived the vanishing of the quantum $\nu_{\rm E}$ for translations and rotations. Berry and Wilkinson [28] have shown that off-diagonal matrix elements vanish for translations, rotations and dilations, in the case of exact degeneracy (*i.e.* $\omega_{nm} = 0$). However, neither of these results addresses the finite ω dependence. It is clear that direct application of a random wave assumption, which leads to (3.12), completely fails to predict the band profile for special deformations (agreement is only reached when $\omega \gg \tau_{\rm bl}^{-1}$). More generally, it can be said that the random wave approximation fails whenever $D(\mathbf{s})$ is significant on a large fraction of the boundary.

It might be that there exists some transformation from the boundary overlap form Eq.(C.2) to another overlap integral which can be estimated well by a random wave approximation. For instance, (H.23) gives $(\partial\mathcal{H}/\partial x)_{nm}$ for translation in terms of the dipole matrix element (a weighted overlap of eigenstates in the billiard interior). A random wave argument for this overlap would predict the correct power law $\gamma = 4$ for the band profile. A similar effort has been made in the dilation case [195], however, as discussed in Section 6.1.2 this leads to the wrong power law. Generally the use of random waves for such overlap estimates is dangerous. No relation has been found relating to rotations. The prediction of special deformation band profiles using wave manipulations alone (*e.g.* of the type in Appendix H) is an area for research.

3.4 The WNA revisited—cavity shape effects

In Section 3.2 we have assumed that generic fluctuating quantities such as $\mathbf{r}^2$ and $\mathbf{e}\cdot\mathbf{r}$ and $\mathbf{e}\cdot(\mathbf{r}\times\mathbf{p})$ have a white noise power spectrum for $\omega \ll 1/\tau_{\rm bl}$. In section 4.1 we are going to suggest that this white noise assumption is approximately true for any fluctuating quantity $\mathcal{F}(t)$ that comes from a normal deformation (the term 'normal' will be defined there).

Obviously, the goodness of the 'white noise assumption' in the two cases mentioned is related to the chaoticity of the system, and should be tested for particular examples. This has been done so for the cavity of Fig. 3.2 (see [14], and Figs. 3.4 and 4.2). This cavity is an example of a 'scattering billiard' and so exhibits strong chaos (no marginally-stable orbits) [40]. If the motion is *not* strongly chaotic we may get a $C(\tau)$ that decays like a power law (say $1/\tau^{1-\gamma}$ with $0 < \gamma \leq 1$) rather than an exponential [40, 78, 54, 37] (time

of crossover to algebraic decay is discussed in [53]). In such case the universal behavior is modified: we get $\omega^{-\gamma}$ behavior for $\tilde{C}_{\rm E}(\omega)$ at small frequencies ($\nu_{\rm E}$ diverges), signifying faster-than-diffusive energy spreading in Eq.(2.8) [37]. The stadium is an example where such a complication may arise: an ergodic trajectory can remain in the marginally-stable 'bouncing ball' orbit family (between the top and bottom edges) for long times, with a probability scaling as t^{-1} [78, 54, 53]. Depending on the choice of $D(\mathbf{s})$ this *may* manifest itself in $C(\tau)$. For example, in Figs. 2.5 and 3.7 the deformations P and Gp respectively both involve a distortion confined to the upper edge, and the resulting sensitivity to the bouncing ball orbit leads to large enhancement of the fluctuations intensity $\tilde{C}(\omega{=}0)$, and is suggestive of singular behavior for small ω. However in the same system the deformations W2 and G, which are zero on the upper and lower edges, show no such enhancement—the band profile is flat as $\omega \to 0$ and the deviation from strong chaos is masked.

If the billiard has a mixed phase space (which is the generic case), then the integrable component does not contribute to diffusive energy spreading. Proposals have been made to account for this via a phase-space volume factor [158, 149].

Chapter 4

Improving upon the white noise approximation: a new 'wall formula'

For calculating the rate of energy absorption due to time-dependent deformation of the confining potential, in this chapter I introduce an improved version of the wall formula. The formulation takes into account the 'special' class of deformations that cause no heating in the zero-frequency limit, which was identified in the previous chapter. Recall that since calculation of the exact kernel (3.5) is very complicated, we are interested in an approximate prediction for $\tilde{C}_{\rm E}(\omega)$, and for the noise intensity $\nu_{\rm E}$ in particular. From this follows the friction coefficient μ, according to the recipe in Section 2.1.5.

The simplest estimate for $\nu_{\rm E}$ is the white noise approximation (WNA) introduced in the previous chapter, and it leads (for a 3D cavity) to the well known 'wall formula' [29]

$$\mu_{\rm E} = \frac{N}{\mathsf{V}} m v_{\rm E} \oint D(\mathbf{s})^2 d\mathbf{s} \tag{4.1}$$

where the subscript $_{\rm E}$ implies that we are considering a microcanonical ensemble $\rho(E)$, the number of particles is N, and the volume of the cavity is V. The deformation of the cavity is described by $D(\mathbf{s})$. A general $\rho(E)$ can be handled simply by replacing $v_{\rm E}$ by the enesemble average particle speed $\bar{v}$. The above version of the wall formula was originally derived for the purpose of calculating the so-called one-body dissipation rate in nuclei. The original derivation of this formula is based on a simplified kinetic (gas particle) picture [29]. For an alternate derivation using linear response see [118, 120]. Cohen[46] provides the generalization to any dimension d.

The main purpose is to introduce an improved version of the wall formula, in the form of an improved estimate for $\nu_{\rm E}$, which I will call the 'IFIF' (Section 4.3). This improvement involves projecting out the special components of a general deformation, and only then to estimate $\nu_{\rm E}$ using the WNA. This will give an estimate which handles many forms of $D(\mathbf{s})$ better than the plain WNA, as I demonstrate numerically.

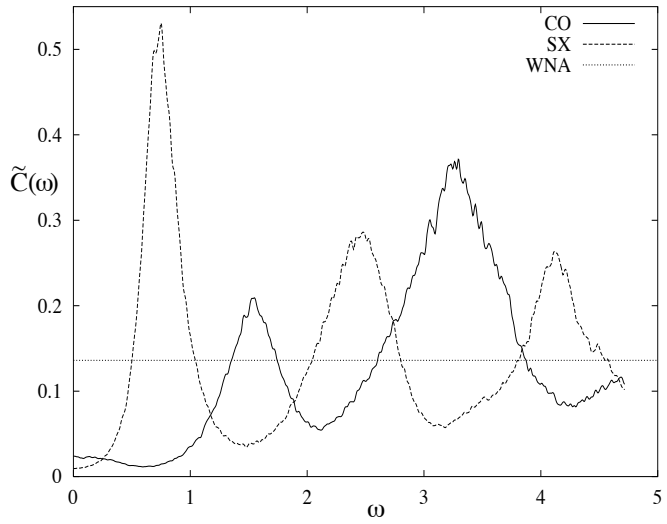

Figure 4.1: The failure of the WNA estimate for $\tilde{C}(\omega)$ for deformation types CO (similar to DI) and SX (similar to TX). The WNA is clearly a vast overestimate of the small-ω limit. See Tables 3.1 and 3.2 for explanation of deformation types.

4.1 Decomposition of general deformations

The failure of the WNA for 'special' deformations also extends to the much wider class of deformations which are *similar* to special. This is demonstrated in Fig. 4.1. It should be emphasized that this failure happens even if the cavity is strongly chaotic.

We seek an analytical estimate for $\tilde{C}(\omega)$, and in particular for its zero-frequency limit ν. This estimate should apply to any (general) deformation, including the case of 'close-to-special' deformations. It would be useful to regard any general deformation as a combination of 'special' component and 'normal' component. The formulation of this idea is the theme of the present section. Supporting numerical evidence is gathered in the next section.

The special deformations (for which we have $\nu = 0$) constitute a linear space, meaning that any sum of special deformations is also a special one. Now we would like to conjecture that there is also a linear space of 'normal' deformations. By definition, for 'normal' deformation $\mathcal{F}(t)$ looks like an uncorrelated random sequence of impulses, and consequently the WNA is a reasonable approximation. The notion of randomness can be better formulated as in Appendix F leading to Eq.(F.4). However in practice (F.4) is not useful, because it cannot be applied as an actual classification tool. (Eq.(F.4) is never satisfied exactly). Still we are going to demonstrate that there is a *unique* way to identify the subspace of normal deformations, if we insist on a maximal (*i.e.* the most inclusive) definition of this subspace.

It is important to clarify the heuristic reasoning for having a linear space of normal deformations. The $\mathcal{F}(t)$ that corresponds to some normal deformation $D(\mathbf{s})$ looks like white noise. It means that only self-correlations of its spikes are statistically significant. If we

have two such generic quantities, say $\mathcal{F}_1(t)$ and $\mathcal{F}_2(t)$, then we expect $\mathcal{F}_1(t)+\mathcal{F}_2(t)$ to share the same property.

The correlation function of $\mathcal{F}(t) = \mathcal{F}_1(t) + \mathcal{F}_2(t)$ can be written formally as

$$C_{1+2}(\tau) = C_1(\tau) + C_2(\tau) + 2C_{1,2}(\tau) \tag{4.2}$$

where $C_{1,2}(\tau)$ is the *cross-correlation function.* In Appendix E we argue the following

$$\int_{-\infty}^{\infty} C_{1,2}(\tau)d\tau \;\; = \;\; 0 \qquad \text{if 1=general, 2=special} \tag{4.3}$$

This can also be proved easily using the fact that $\nu_{\rm E}$ is an exact quadratic form (3.5) in the function space of $D(\mathbf{s})$. (Consider that the special deformations are eigenvectors of this quadratic form with zero eigenvalue, *i.e.* they lie in the null-space [188]). The result is exact, and does not involve any approximation. In Appendix F we argue the following

$$C_{1,2}(\tau) \;\approx\; c \times \left[\oint D_1(\mathbf{s})D_2(\mathbf{s})d\mathbf{s}\right] \delta(\tau) \qquad \text{if 1=normal, 2=general} \tag{4.4}$$

where $c = 2m^2 v_{\rm E}^3 \langle|\cos\theta|^3\rangle/\mathsf{V}$. This result is an approximation, which is expected to be as good as our assumption regarding the 'normality' of the deformation $D_1(\mathbf{s})$. Consider now the case where $D_1(\mathbf{s})$ is normal and $D_2(\mathbf{s})$ is special. Both Eq.(4.3) and Eq.(4.4) should apply. But these equations are consistent if and only if $D_1(\mathbf{s})$ is orthogonal to $D_2(\mathbf{s})$. We say that $D_1(\mathbf{s})$ and $D_2(\mathbf{s})$ are orthogonal ($1 \perp 2$) using the following definition:

$$\text{orthogonality} \qquad \Leftrightarrow \qquad \oint D_1(\mathbf{s})D_2(\mathbf{s})d\mathbf{s} = 0 \tag{4.5}$$

Thus we have proved that normal deformations must be orthogonal (in the sense of (4.5)) to special deformations. Obviously we have proved here a necessary rather than a sufficient condition for 'normality'. However, if we insist on a maximal definition for the subspace of normal deformations, then we get a unique identification. Namely, a deformation is classified as 'normal' if it is orthogonal to the subspace of special deformations.

The practical consequences of Eq.(4.3) and Eq.(4.4) are as follows:

$$\nu_{1+2} = \nu_1 \qquad \text{if 1=general, 2=special} \tag{4.6}$$

and

$$\nu_{1+2} \approx \nu_1 + \nu_2 + 2c \oint D_1(\mathbf{s})D_2(\mathbf{s})d\mathbf{s} \qquad \text{if 1=normal, 2=general} \tag{4.7}$$

These results are tested in the next section.

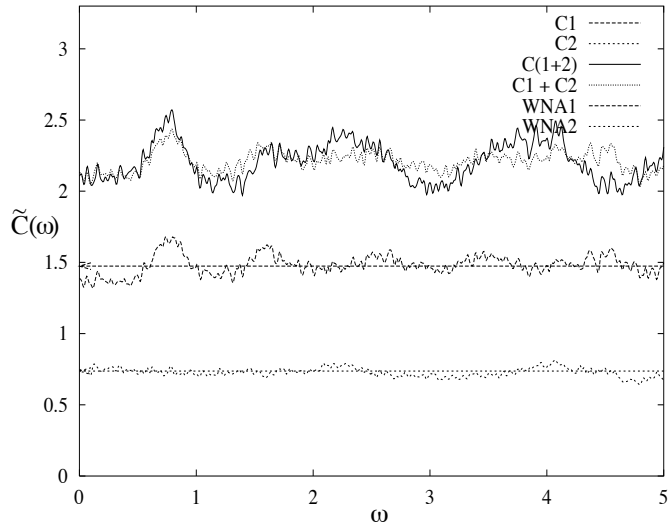

Figure 4.2: Addition of two 'good' normal deformations (1=P2, 2=WG). The two are orthogonal in the sense of (4.5). That they are 'good' can be seen by their good agreement with their WNA results (horizontal arrows). The power spectrum of the sum agrees well with the sum of the power spectra.

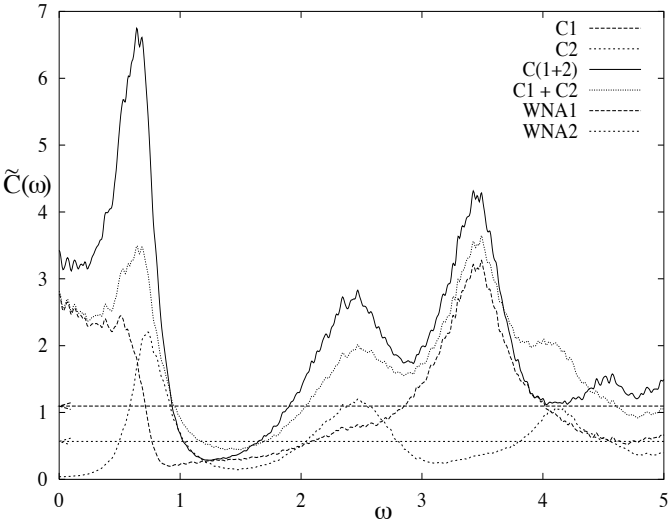

Figure 4.3: Addition of two 'bad' normal deformations (1=FR, 2=SX). The two are orthogonal in the sense of (4.5). That they are 'bad' is shown by a lack of agreement with their WNAs. The power spectrum of the sum is badly approximated by the sum of the power spectra (non-linear addition).

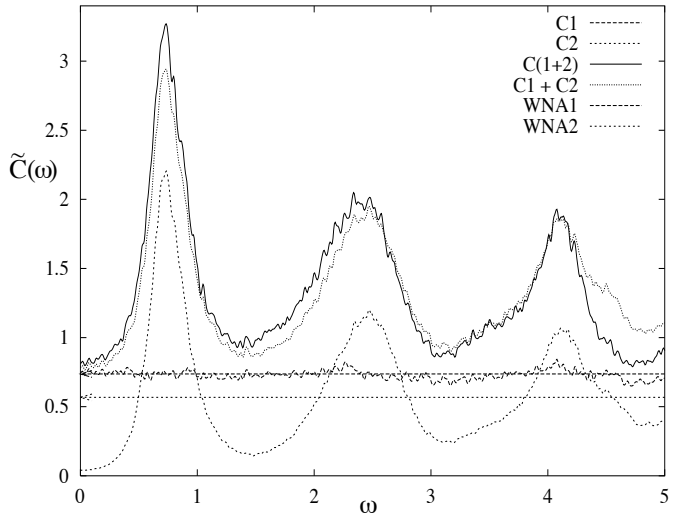

Figure 4.4: Addition of a 'good' normal deformation (1=WG) to a general deformation (2=SX). The two are orthogonal in the sense of (4.5). The power spectrum of the sum agrees well with the sum of the power spectra.

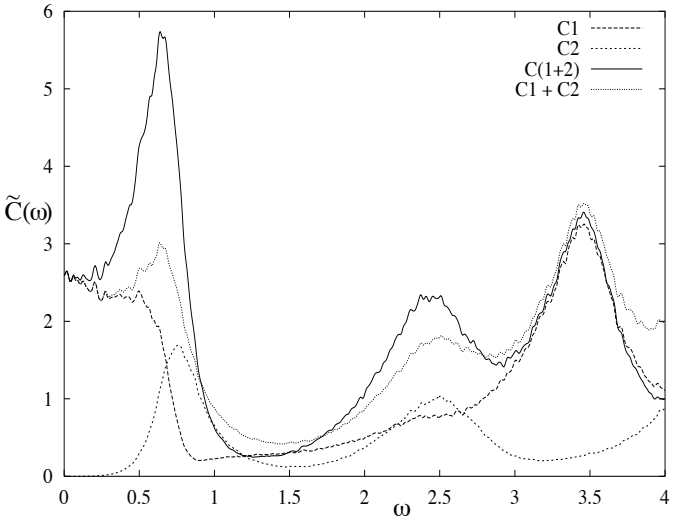

Figure 4.5: Addition of a general deformation (1=FR) to a 'special' deformation (2=TX). The power spectrum of the sum coincides with the sum of the power spectra in the limit $\omega \to 0$, as implied by Eq.(4.6).

4.2 Addition of deformations: numerical tests

On the basis of the discussion in the previous section we *define* normal deformation as those that are orthogonal to all special deformations, in the sense of Eq. (4.5). Obviously there are 'good' normal deformations for which the WNA is an excellent approximation (P1 and W8 in Fig. 3.4, for example), and there are 'bad' normal deformations for which the WNA is not a very good approximation (FR in Fig. 3.4, and the normal component in Fig. 4.8b). In this section we present numerical evidence that verifies the theoretical results of the previous section, and investigate how 'bad' a normal deformation has to be for them to break down.

From what we have claimed it follows that if $D_1(\mathbf{s})$ and $D_2(\mathbf{s})$ are orthogonal normal deformations, then $\nu_{1+2} = \nu_1 + \nu_2$. We could as well write

$$\begin{aligned} \tilde{C}_{1+2}(\omega) \;&\approx\; \tilde{C}_1(\omega) + \tilde{C}_2(\omega) \\ &\text{if 1=normal, 2=normal, and } 1 \perp 2 \end{aligned} \tag{4.8}$$

because by assumption the three correlation functions are approximately flat. We demonstrate this addition rule in the case of two 'good' deformations which are orthogonal in Fig. 4.2. We found that small 'pistons' (P2 is significant on only $\sim 1/50$ of the perimeter) were needed to achieve addition of the accuracy (a few %) shown. However, the restriction on the 'wiggle' type of deformation was somewhat more lenient (WG is ~ 5 times wider than P2 yet obeys the WNA better than P2 does).

In general we observe that the quality of the addition rule is limited by the deviation from the WNA of the *better* of the two deformations. In Fig. 4.3 we see that if both $D_1(\mathbf{s})$ and $D_2(\mathbf{s})$ are bad, then also the addition rule (4.8) becomes quite bad. Fig. 4.4 shows that the addition rule (4.8) is reasonably well satisfied also if *either* $D_1(\mathbf{s})$ or $D_2(\mathbf{s})$ is a 'good' normal deformation. We have chosen $D_1(\mathbf{s})$ as WG (good), and $D_2(\mathbf{s})$ as SX which is almost completely dominated by the special x-translation deformation. The addition rule (4.8) is obeyed at all ω. This proves that our assertions Eq.(4.4) about the vanishing of $C_{1,2}^{\text{non-self}}(\tau)$ is indeed correct. It holds here as a non-trivial statement ($D_2(\mathbf{s})$ is general and 'bad').

Finally, we consider the case where $D_1(\mathbf{s})$ is general and $D_2(\mathbf{s})$ is special. This is illustrated in Fig. 4.5. The addition rule (4.8) becomes exact in the limit of small frequency corresponding to the vanishing of $\tilde{C}_{1,2}(\omega \to 0)$ as implied by Eq.(4.3). In particular this implies that $\nu_{1+2} = \nu_1$. Note that there is *no* condition on the orthogonality of $D_1(\mathbf{s})$ and $D_2(\mathbf{s})$). This will be the key to for improving over the WNA, which we are going to discuss in the next section.

4.2.1 Symmetry effects

In drawing the above conclusions it is important to note that symmetry effects can play a deceptive role if the cavity shape has symmetry. Our example Fig. 3.2a is in the C_{2v} symmetry group, and the symmetry class of the deformations used are given in Table 4.1. In Fig. 4.6 we demonstrate that the addition rule (4.8) is very accurately satisfied at *all* ω if $D_1(\mathbf{s})$ and $D_2(\mathbf{s})$ belong to different symmetry classes of the cavity. Orthogonality of $D_1(\mathbf{s})$

parity		*deformation functions in class*
inversion +	no reflections	Wn (n even)
	x+, y+	DI, CO
	x−, y−	RO
inversion −	no reflections	Wn (n odd)
	x+, y−	TY
	x−, y+	TX, FR, SX
no symmetries		DF, P1, P2, WG

Table 4.1: Categorization of deformations into symmetry classes shared by the (C_{2V}-symmetric) example billiard of Fig. 3.2a. See Tables 3.1 and 3.2 for explanation of deformation types.

and $D_2(\mathbf{s})$ is *not* sufficient to explain this perfect linearity of addition of $\tilde{C}_{\rm E}(\omega)$. Rather, it follows from the symmetry of the kernel $\gamma_{\rm E}(\mathbf{s}_1, \mathbf{s}_2)$ of Eq.(3.5). The cross-terms in (3.5) rigorously vanish when such deformations are added, because the kernel must possess the same symmetries are the cavity itself. The consequence is that in order to demonstrate the assertions of this and of the previous section, we had to choose deformations of the *same* symmetry class, or which break all symmetries of the cavity.

4.3 Beyond the WNA

It is possible now to consider the case of general deformation, and to go beyond the WNA. Given a general deformation $D(\mathbf{s})$ we should project out (subtract) all the special components, leaving the normal component, and only then apply the WNA. In Fig. 4.8 we demonstrate this decomposition for the deformation (CO + W16) and the deformation SX. Here and elsewhere, all boundary deformation function integrals (of the form (4.5)) are evaluated using the technique of Appendix G.

The special deformations constitute a linear space which is spanned by the basis functions: one dilation, d translations, and $d(d-1)/2$ rotations. (For d=2 they are listed in Table 3.2). For a general cavity shape these basis functions are not orthogonal. However, because they are linearly independent, we can use standard linear algebra to build an orthonormal basis $\{D_i(\mathbf{s})\}$ of special deformations. The special ($\|$) and the normal ($\perp$) components of any given deformation $D(\mathbf{s})$ are therefore

$$\begin{aligned} D_{\|}(\mathbf{s}) &= \textstyle\sum_i \alpha_i D_i(\mathbf{s}) \\ D_{\perp}(\mathbf{s}) &= D(\mathbf{s}) - D_{\|}(\mathbf{s}) \end{aligned} \tag{4.9}$$

where the coefficients are

$$\alpha_i = \oint D(\mathbf{s}) D_i(\mathbf{s})\, d\mathbf{s}. \tag{4.10}$$

The improved approximation for ν applies the WNA only to the normal component, giving

$$\nu_{\rm E} \approx 2m^2 v_{\rm E}^3 \langle |\cos\theta|^3 \rangle \frac{1}{\mathsf{V}} \oint d\mathbf{s} [D_{\perp}(\mathbf{s})]^2 \tag{4.11}$$

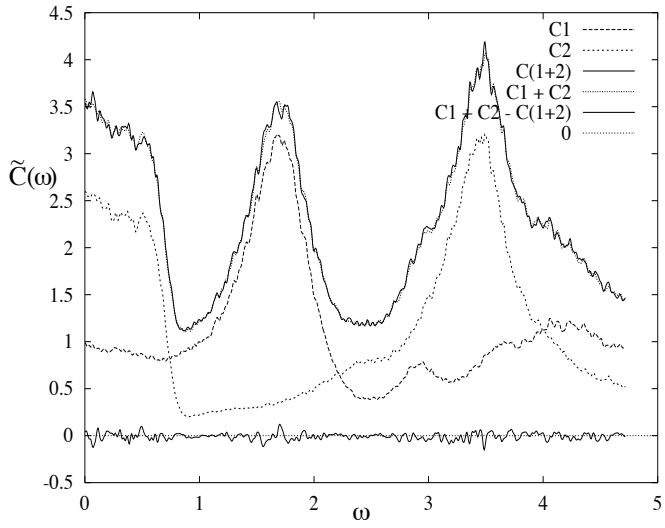

Figure 4.6: Addition of two 'bad' general deformations which come from different symmetry classes of the cavity (1=W2, 2=FR). The two must also be orthogonal, by symmetry. The deviation from linear addition (solid line varying about zero) vanishes at all ω.

which we name the IFIF (Improved Fluctuations Intensity Formula). In the particular case of d=3, substitution of this result into the microcanonical FD relation gives an 'improved wall formula' consisting of the replacement of $D(\mathbf{s})$ by $D_\perp(\mathbf{s})$ in Eq. (4.1).

4.3.1 Numercial tests of the improved formula

In Fig. 4.8 we use the IFIF to estimate ν for two examples. The first is a deformation (CO + W16) whose normal component is 'good', due its oscillatory nature. The deviation from a flat white power spectrum is $\sim 20\%$ for the normal component. The IFIF result Eq.(4.11) is accurate to a few percent. It is a much better estimate of the actual ν compared with the naive WNA Eq.(3.12) which overestimates the correct value by a factor of 2.2. In the second example the deformation is SX. The resulting normal component is 'bad'. Its power spectrum fluctuates by a factor of about 10 in the ω range shown. Consequently the IFIF is limited in its accuracy, and the correct value for ν is underestimated by a factor of 2.5. However, it is still a great improvement over the naive result Eq.(3.12). In this second example we can extract another prediction about $\tilde{C}_{\rm E}(\omega)$. The special component is a factor ~ 10 larger than the normal component. Therefore the ω^2 behavior at small ω is almost entirely due to the 'rotation' component. The prefactor of the ω^2 behavior need only be found once for each billiard shape. This saves computation and gives an extra information about the dissipation rate at finite driving frequency.

4.3.2 Historical remarks and conclusion

A few concluding remarks regarding the history of the wall formula are in order. It has been known since its inception that the naive wall formula gives unphysical answers in

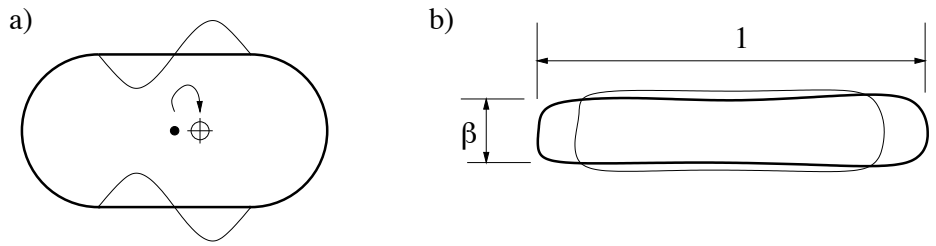

Figure 4.7: **a**) A deformation of the stadium which moves the 'center of mass' (centroid) of the cavity to the right (from the dot to the crosshairs symbol). This deformation is orthogonal (in the sense of (4.5)) to all special deformations, in particular, all translations. **b**) An example volume-preserving deformation of an elongated approximately-rectangular cavity ($\beta \ll 1$) which nevertheless has a large overlap with dilation. It can be shown that this results in an IFIF estimate of $\approx 4\beta$ times that of the naive WNA. In both diagrams the undeformed shape is shown as a heavy line, the deformed one as a thin line.

the case of *constant-velocity* translations and rotations. This was first regarded as a kinetic gas 'drift' effect [29]. It should be noted that the recipe presented in [29], namely to subtract this drift component, is equivalent in practice to the recipe (4.9) that we have presented here, *provided we ignore dilations*. It is also important to realize that the argumentation in [29] for this subtraction appears to be ad hoc, being based on a 'least-structured drift pattern' reasoning. A stated condition on this subtraction was that the resulting deformation preserve the location of the 'center of mass' (centroid) of the cavity, for reasons particular to the nuclear application [29]. This condition seems to have become standard practice in numerical tests of the wall formula [158, 149, 30, 31, 32]. However, as Fig. 4.7a shows, this condition is generally *not* equivalent to the above subtraction of translation and rotation components [1]. This seems to invalidate the theorem presented in Section 7.1 of [29]. Where the flaw in their reasoning lies we are not sure.

The consideration of the special nature of dilations is absent from the literature. Even if we restrict ourselves to volume-preserving deformations (the nuclear dissipation case), then deformations of certain cavities can be found for which the dilation correction is significant. This correction can only be large if the cavity has a large variation in radius (*i.e.* is highly non-spherical). We illustrate this in Fig. 4.7b. We suggest this as a possible reason why major discrepancies due to dilation have not emerged in the numerical tests of the wall formula until now. Such tests have generally been of shapes close to a 3D sphere [29, 158, 149, 30, 31, 32].

Hence we believe that the new recipe presented, along with the associated theory and in conjunction with the particular power-law dependences from the previous chapter, is a significant step in the treatment of one-body dissipation and of dissipation in d-dimensional cavities in general.

[1] The condition that a deformation $D(\mathbf{s})$ not move the 'center of mass' (centroid of the cavity volume) is $\oint D(\mathbf{s})\mathbf{r}(\mathbf{s})\, d\mathbf{s} = \mathbf{0}$. This is in general different from the condition for having zero overlap with translations, namely $\oint D(\mathbf{s})\hat{\mathbf{n}}(\mathbf{s})\, d\mathbf{s} = \mathbf{0}$.

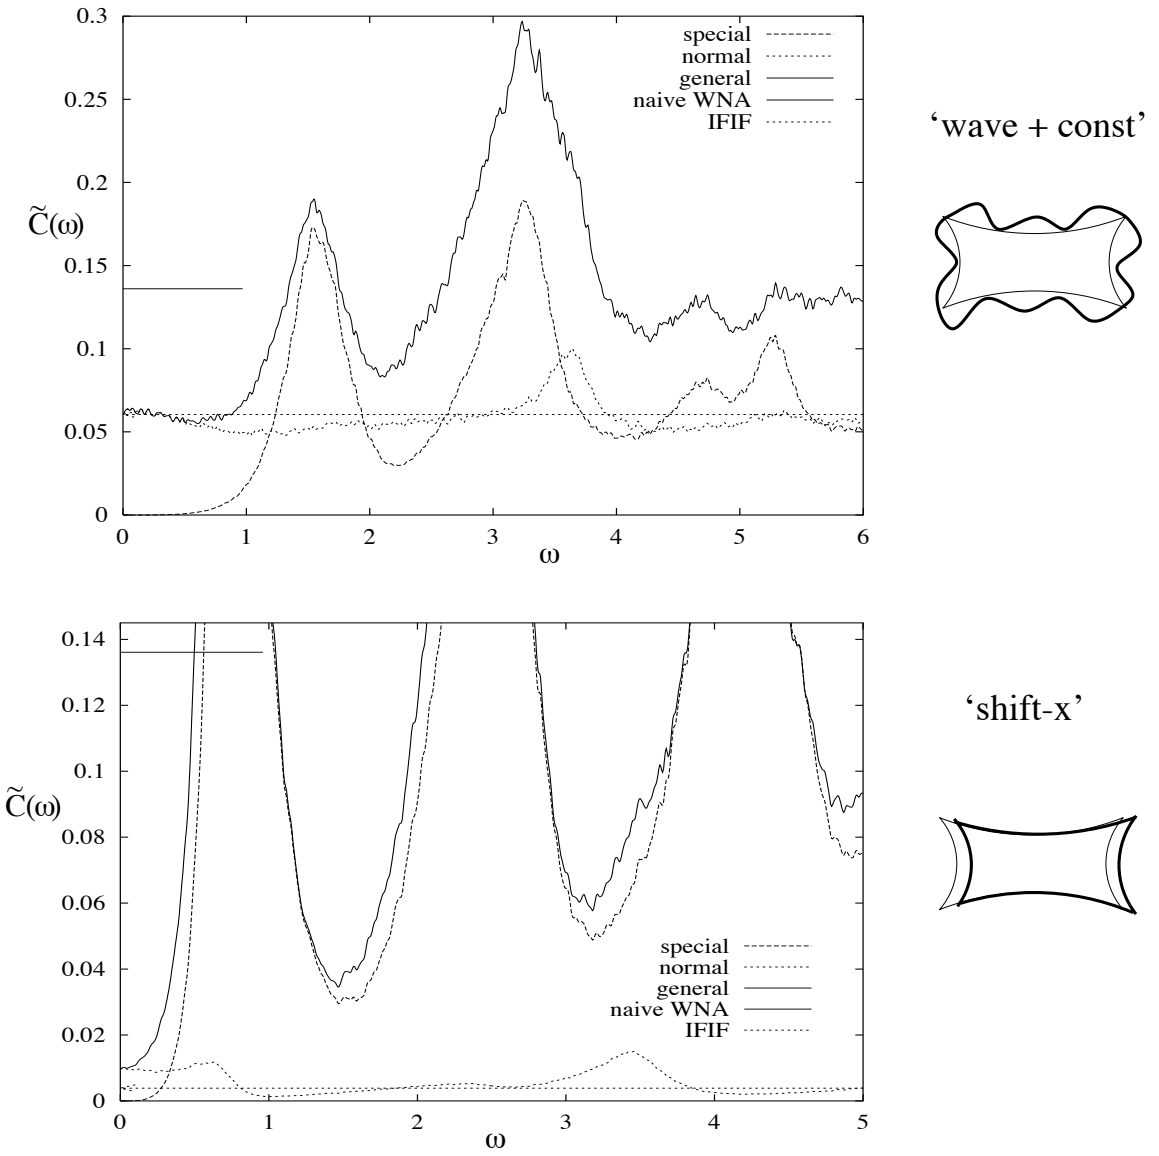

Figure 4.8: Decomposition of general deformations $D(s)$ into orthogonal 'normal' and 'special' components. The general deformation is CO + W16 in subfigure (a), and SX in subfigure (b). The naive WNA Eq.(3.12) is indicated by short solid line. The improved (IFIF) result Eq.(4.11) is indicated by long dashed arrow. In (a) the normal component is quite 'good', giving an accurate IFIF result, but in (b) the normal component of SX is 'bad', limiting the accuracy of the IFIF to 40% of the actual ν.

Chapter 5

Improved sweep methods for billiard quantization

In this chapter I describe a numerical method for finding billiard eigenstates which involves 'sweeping' the energy to locate the states. It is a generalization of and improvement upon Heller's Plane Wave Decomposition Method (PWDM) [90, 91]. The PWDM was invented over 15 years ago and provided a more powerful and rapid way to find eigenstates in certain billiard shapes than had been available beforehand. In this chapter several problems of the PWDM have been solved (missing states, sensitivity to choice of basis set size and number of matching points), it has been simplified, the accuracy is improved, and the speed has been increased considerably. This chapter is also a useful introduction to the following one, where a much more efficient method (which bypasses the 'sweep' altogether) is analysed.

5.1 Introduction and history of quantum eigenproblems

Here I outline the general problem of solving for quantum eigenstates in a finite domain, and categorize some of the various strategies which are used. The case of billiard systems at high quantum number is a relatively recent example. I give a historical overview of which methods have been used for this system.

5.1.1 General definitions

When the Hamiltonian of a quantum mechanical problem is independent of time, the time evolution of the quantum state (wavefunction) can always be written as the sum of components with harmonic time-dependence. Each component, or *stationary state*, ψ_μ, is a solution of the time-independent Schrodinger equation (TISE)

$$\hat{H}\psi \;=\; E\psi. \tag{5.1}$$

The (hermitian) operator $\hat{H}$ is linear, so we have a linear eigenproblem, which only has solutions for (real) energies E_μ. If we ignore all discrete quantum numbers (spin) then ψ is a function of continuous variables only which can be written generally as a vector $\mathbf{r}$. The

dynamical variable $\mathbf{r}$ may represent the location of a particle, or many particles, or more general coordinates such as relative distances, angles, etc. One can define position states $|\mathbf{r}\rangle$ which are the eigenstates of the position operator $\hat{\mathbf{r}}$, and which form a complete basis for the quantum problem. $\psi(\mathbf{r})$ is the position representation of the wavefunction, and $H(\mathbf{r},\mathbf{r}')$ is the position representation of $\hat{H}$. If $\hat{H}$ is *local*, as is the case in a huge variety of physical situations, then $H(\mathbf{r},\mathbf{r}') = \delta(\mathbf{r}-\mathbf{r}')[h_0(\mathbf{r}') + h_1(\nabla_{\mathbf{r}'}) + \cdots]$ is a function only of the local value and spatial derivatives of the wavefunction. The function h_0 is commonly known as the *potential*. For instance a very common form of Hamiltonian arising in the case of isotropic mass m and quadratic dispersion relation is

$$H = -\frac{\hbar^2}{2m}\nabla^2 + U(\mathbf{r}), \tag{5.2}$$

written as an operator in the position representation. The potential is $U(\mathbf{r})$. However, when $\mathbf{r}$ does not simply indicate location of a single particle in a conservative potential, then (5.2) is not sufficiently general.

If the Hamiltonian is infinite (or much larger than the energies of interest) outside a finite region, then $\mathbf{r}$ is effectively confined to this region and the eigenenergies E_μ are *discrete*. This is also true if $\mathbf{r}$ is not strictly confined, but if the energy is below the continuum transition[128], giving bound states. In this latter case ψ can take non-zero values over an infinite region of $\mathbf{r}$; however outside a certain (classically-allowed) region the decay is exponential so the effective region of confinement is finite. I will be interested in solving for these resulting discrete states. The problem is to find the eigenvalues E_μ, or equivalently the eigenfunctions $\psi_\mu(\mathbf{r})$, for which (5.1) is obeyed. If an E_μ is known then it is easy to get the corresponding $\psi_\mu(\mathbf{r})$, and vice versa. If there is no analytical solution, then a numerical approach is required—the subject of this and following chapters.

The issue of *boundary conditions* (BCs) now enters; notice that it is not actually present in the eigenproblem (5.1). BCs arise because of the computational need to consider only a finite region. The finite region is chosen such that $\psi(\mathbf{r})$ is negligible outside the region. I call the region $\mathcal{D}$ and its boundary Γ (see Fig. 5.1). If $\psi(\mathbf{r})$ dies smoothly to zero (as is true for a smooth-potential or soft-walled problem), then the choice of $\mathcal{D}$ is determined by the required accuracy. The smallest region $\mathcal{D}$ is found such that $\psi(\mathbf{r} \in \Gamma) \approx 0$ to a sufficient approximation. In this case, the exact nature of BCs imposed at Γ is irrelevant for the problem—the only important BC is then the asymptotic restriction $\psi(\mathbf{r}) \to 0$ with increasing distance outside the finite region. Another class of problems results if the potential suddenly changes (for instance if it jumps to infinity corresponding to a hard wall) at the edge of the problem region. It is then natural to locate Γ at this edge (often this is essential since it avoids the difficulty of representing singularities which occur at the edge). Therefore $\mathcal{D}$, and the BCs at Γ, are exactly specified. To summarize, the BCs result from truncation of the region of space to $\mathbf{r} \in \mathcal{D}$.

BCs can generally be written

$$(\mathcal{L}\psi)(\mathbf{s}) = 0, \qquad \text{for all } \mathbf{s} \in \Gamma, \tag{5.3}$$

where $\mathcal{L}$ is a linear operator which acts on the function space of ψ, and is sensitive only to its values, gradients, and higher derivatives, at the boundary Γ. The result of the

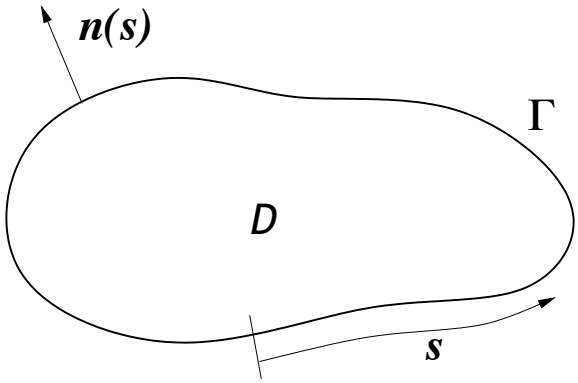

Figure 5.1: d dimensional domain $\mathcal{D}$ and $d-1$ dimensional boundary Γ of a quantum solution region. The local normal $\mathbf{n}(\mathbf{s})$ is also shown.

operator is the smaller function space of functions of the boundary position coordinate $\mathbf{s}$. Generally $\mathcal{L}$ is local on Γ, meaning it measures only properties of $\psi(\mathbf{r})$ local to a point $\mathbf{s}(\mathbf{r})$. For instance, most BCs are described by special cases of the operator $\mathcal{L}$ which returns a (position-dependent) linear combination of the local value and normal gradient (this is known as 'mixed BCs').

BCs may be specified for other reasons. They may enter in the context of solving for resonance states in the continuum, where they can model the effect of an infinite region of space whose propagator is known analytically. These "radiative BCs" may be non-local on Γ, and the resulting E_μ will be complex. The calculation of eigenstates in the continuum (which are usually highly degenerate) is a *scattering* problem, and I shall not consider it further. Also the eigenstates in a region with certain BCs can be needed for use in other methods, for instance R-matrix theory[131] for scattering problems.

5.1.2 Categories of existing numerical solution techniques

Almost all interesting quantum problems involve $d \geq 2$ degrees of freedom. Almost invariably in this case the wavefunction ψ is represented by a linear sum of *basis functions*,

$$\psi(\mathbf{r}) \;=\; \sum_{n=1}^{N} x_n \phi_n(\mathbf{r}). \tag{5.4}$$

This is computationally necessary since it reduces the function space of ψ to a finite number N of degrees of freedom, namely the coefficient vector $\mathbf{x}$. For a *complete* basis, any $\psi(\mathbf{r})$ can be exactly represented by (5.4) in the limit $N \to \infty$. Clearly a computer can only handle finite N, so the basis should be well chosen so that sufficient convergence is achieved for the $\psi(\mathbf{r})$ of interest with a manageable N. The choice of basis falls into two general classes: (A) The basis does not depend on E, or (B) the basis does depend on E. This results in two general strategies, summarized in Table 5.1 and now discussed in detail below. The nomenclature "Class A" etc which I introduce here is new to the literature.

	(A) 'basis diagonalization'	*(B) 'Green's function matching'*
basis funcs.	Independent of E. Obey BCs. Don't obey TISE.	Depend on E. Don't obey BCs. Are solutions of TISE.
basis size	Scales like volume of $\mathcal{D}$	Scales like surface area of Γ
technique	Write TISE in terms of basis coefficients.	Write BC matching equations in terms of basis coefficients and boundary locations.
resulting equation	$H\mathbf{x} = E\mathbf{x}$ Diagonalization of H matrix finds many $(E_\mu, \mathbf{x}_\mu)$ at once.	$A(E)\mathbf{x} = 0$ Hunt in E for $\det[A(E)] = 0$, needs many SVDs per $(E_\mu, \mathbf{x}_\mu)$ found.
examples	Finite Element Method (FEM)[19]. Discrete-Variable Representation (DVR)[17]. Discretization on a lattice[161]. Robnik conformal mapping[171].	Boundary Integral (or Element) Method (BIM or BEM)[121]. Plane Wave Decomposition Method (PWDM) of Heller [91]. Bogomolny's method[34]. Lupu-Sax t-matrix inversion[140]. Methods of this chapter.

Table 5.1: Two general classes of quantum eigenproblem solution strategies using basis sets. The desired solutions are the eigenenergies E_μ and the states (written in terms of basis coefficients) $\mathbf{x}_\mu$. For other physical wave eigenproblems, the TISE (5.1) should be replaced by the appropriate time-independent wave equation obeyed in the domain (*e.g.* Helmholtz equation). In the methods of this chapter a diagonalization will replace the SVD (singular value decomposition). The scaling method of Vergini and Saraceno (not listed) combines the best elements of both classes.

A) 'Basis diagonalization'

Class A has the advantage that the E-independent basis preserves the linearity of the eigenproblem, giving a matrix eigenequation which returns many solutions $(E_\mu, \mathbf{x}_\mu)$ at once. This is the conventional 'basis diagonalization' approach. Generally *all* the eigenstates from 1 up to a maximum useable state are returned. Clearly N cannot be less than this number of states. For instance in the case of a d dimensional billiard, this means N must scale like k^d if states at wavenumber k are sought. This is clearly a huge limitation.

The basis can be chosen as the analytically-known eigenstates of a simpler Hamiltonian $\hat{H}_0$; this I call Class A1. The resulting basis is therefore orthonormal and complete (in the $N \to \infty$ limit) and if $\hat{H}_0$ is 'close' to $\hat{H}$ then N will not need to be much higher than n_{E}, where n_{E} is the typical quantum number of the desired states at energy E.

Alternatively the basis is chosen to be convenient in position space (or momentum space, or a mixture of both). The basis is effectively complete (up to energies of interest) because it entirely covers the domain $\mathcal{D}$. This I call Class A2. The advantage of these localized basis functions is that the resulting H matrix is *sparse*, allowing much faster extraction of eigensolutions. Various smoothnesses of such basis functions are possible, ranging from lattices (corresponding to piece-wise linear 'pyramidal' functions) through other Finite Element [87, 19, 182, 56] basis functions and higher-order spline functions [70], to gaussian packets (coherent states, or the Distributed Gaussian Basis [DGB][57]). The smoother a basis is, the faster the convergence with N can be for a given energy of interest. However, smooth basis functions are more complicated to construct (especially if definite BCs are required), to evaluate, and result in less sparse matrices. Lattice methods (often known as 'finite differencing')[161, 18] generate sparse matrices which can be diagonalized much faster than dense ones, but errors converge only like a power law $\sim N^{-1/d}$. One smooth basis with useful sparsity properties is the Discrete-Variable Representation (DVR) [17, 86, 180]. Most methods involve a compromise. The optimal basis set choice for smooth-potential problems appears to be a covering of phase space by gaussian packets, in which case N need only be a couple of times larger than n_{E} (for small dimensions d) [57, 138]. Because the Wigner function for such problems dies exponentially outside the classically-allowed region of phase space, these phase-space covering methods achieve exponential convergence with N, once N is larger than the semiclassical basis size.

Note that if the basis is not orthogonal, as is frequently the case for Class A2 basis sets, the eigenequation $H\mathbf{x} = E\mathbf{x}$ becomes a generalized eigenequation $H\mathbf{x} = EB\mathbf{x}$ where the matrix B gives the basis function overlaps.

B) 'Green's function matching' (sweep in energy)

Class B can be applied when the Green's function (*i.e.* the eigenstates of $\hat{H}$ the in $\mathcal{D}$ with some different BCs) is known analytically. It has the advantage that this Green's function can now be used to propagate between points on Γ without explicit solution over the volume of $\mathcal{D}$. This reduces the dimensionality of the problem from d to $d-1$. Why is this an enormous reduction in computational effort? Generally the number of bases N is reduced from $O(\mathsf{V}/\lambda_{\mathrm{B}}^d)$ to $O(\mathsf{A}/\lambda_{\mathrm{B}}^{d-1})$, where V is the volume of $\mathcal{D}$ and A is the area of Γ. The typical quantum wavelength at the energy of interest is λ_{B}. This is a typical reduction

in N of $O(\mathsf{L}/\lambda_{\mathrm{B}})$ where L is the typical length scale of the region. Given that the required dense matrix techniques take time $O(N^3)$, this is a huge reduction in effort when there are many wavelengths across the system (semiclassical limit). This corresponds to very high quantum number of interest n_{E}. The problem has now become a one of matching on the boundary. If the length in coefficient space $|\mathbf{x}|^2$ roughly corresponds to a state's norm in $\mathcal{D}$, then the BC matching is equivalent to a zero-determinant condition (see Table 5.1). The E values for which the BCs can be obeyed must be found. However the linearity in E of the matrix equation has been sacrificed. There is usually no way to tell at what E the determinant will vanish, so a 'sweep' or 'hunt' (search routine) in E-space is required, taking many iterations per solution found.

The Boundary Integral Method (BIM or BEM), probably the most common example of Class B methods, has been used extensively in engineering and physics [36, 121, 28, 33, 34]. It emerged in the engineering community in the 1950s (see Chapter 1 of [36]). Also in Class B are similar methods which replace matching on Γ by 'internal boundary conditions' where two (or more) different Hamiltonians may be joined (matched) along an internal surface, which is simply a quantum version of a Poincaré surface of section. Examples are Bogomolny's method [34] (where the propagation in each half-region is found semiclassically), the variation of Prosen [163] (where the Hamiltonians are known analytically), the variation of Smilansky [172], and the Korringa-Kohn-Rostoker (KKR) method (as first used in the quantum chaos community by Berry [23], also see [162] and references within). Equally important for quantum billiard problems is Heller's PWDM, upon which this chapter is based, discussed below. The above methods differ mainly in the choice of basis set, and the way in which the matching is expressed. However they all suffer from the need to 'sweep' along the energy axis.

For specialized problems other Class B methods have been invented, for instance the t-matrix inversion technique of Lupu-Sax [140] for eigenstates of domains containing point-scatterers. Note that the standard 'shooting method' [119] for problems in $d = 1$ (where it is possible to integrate directly the ordinary differential equation resulting from the position-space form of (5.1)) is essentially a Class B method because an energy sweep is required.

C) Other methods

Other methods which do not fall easily into the above two classes include most notably the extremely powerful scaling method of Vergini and Saraceno, which combines the best elements from both classes (see history below), and is the subject of the next chapter. Also deserving mention are 1) time-propagation (by *e.g.* split-operator FFT method[18]) followed by extraction of Fourier coefficients, 2) physical simulation by microwave[185, 123, 93] or acoustic [69] cavities, or even by soap bubble films [6], 3) an interestingly bizarre but impractical recent suggestion by Bulgac [39].

5.1.3 Brief history of the quantum billiard eigenproblem

A billiard is a closed region of space $\mathcal{D}$ in which the wavenumber k is independent of the position $\mathbf{r} \in \mathcal{D}$. Boundary conditions are defined on Γ. The number of degrees of

freedom (dimensionality of the problem) is d. The case of $d = 2$ (or occasionally 3) has become a paradigm system for the study of quantum chaos. For this reason, the solution of billiard eigenstates at high energies (quantum number $n \gg 1$) has been of recurring interest in the last 20 years. I will confine myself to study of the case without magnetic field. (I refer the interested reader to the magnetic field billiard method of Smilansky[98]).

Since the billiard interior is a region of free space whose Green's function is known analytically, Class B is the most common approach. Since Class B can reach much higher quantum numbers (being a surface rather than volume method), it is the only viable way to reach the semiclassical regime (λ_B of order $L/100$ or less). Two methods have dominated the field. The first is the Boundary Integral Method (BIM), where wavefunctions are represented by a dipole source distribution on the boundary. This was used in the pioneering study by MacDonald and Kaufman[145], using the BIM as presented by Riddell[166]. A similar method was then presented by Berry and Wilkinson [28], which has since been used in an essentially unaltered form [33, 137, 121]. The BIM still provides the most reliable (it handles a large variety of billiard shapes), if not the most efficient method for billiard eigenproblems. The second is Heller's Plane Wave Decomposition Method (PWDM), which has the edge on efficiency and simplicity over the BIM, mainly due to the simpler basis functions. It has been used apparently in an unaltered form (for instance [136, 137, 134]) since its invention in 1984 [90, 91]. The only significant improvement to this method has been a slightly better search in E-space [135].

In 1995 a *scaling method* which found a large cluster of states in a single diagonalization was invented by Vergini and Saraceno [195]. This is a Class B method which nevertheless returns many states per diagonalization in the manner of Class A methods. It is without doubt the most significant tool in billiard quantization to have emerged in the last 15 years. However its workings are very poorly and briefly explained [195] (much to the lament of many [39]), and there are in fact errors in this explanation. The aim of Chapter 6 is to remedy this situation.

It should be noted that all the above Class B methods have difficulties with certain billiard shapes. The PWDM (and the scaling method) are most sensitive in this respect, because only certain shapes are known for which a good plane-wave basis set exists. Much more study is needed in this area. The BIM suffers from problems with non-convex shapes [137], an active area of research. Currently I believe there is no method reliable for arbitrary shapes, even in $d = 2$.

Other, semiclassically-motivated methods have been employed for billiards, for instance KKR [23, 162] and the surface of section approaches of Bogomolny [34] and Smilansky [66]. The KKR is applicable only to Sinai-type billiards (cells of periodic lattices). Prosen's matching method [163] (a variant of that of Bogomolny) require the shape to be composed of the union of two *integrable* regions. As an example of the latter case, the stadium billiard so common in quantum chaos (and in this thesis, see Fig. 2.6) is composed of a circular end-cap joined to a rectangular box, and could be thus solved to arbitrary accuracy without any basis-set choice problems. However, we will see in the next chapter than the efficiency of the scaling method will far exceed any other Class B approach.

Finally, the brute-force Class A approach of Robnik [171, 165] should not be overlooked, if only on the grounds of its impressive computational requirements. A banded

matrix of order $N = 35000$ was diagonalized to give highly accurate low-energy billiard eigenstates (these energies have been a useful standard for error analysis of other methods [137]).

5.2 Definition of the billiard problem

Given the isotropic quadratic dispersion relation corresponding to (5.2), we can choose energy units such that $E = (\hbar k)^2/2m$. If the dispersion relation is not isotropic, it can be made so by a re-parametrization of $\mathbf{r}$. However, note that E will not play any further role. The numerical methods described in this thesis are really about finding the *eigenwavenumbers* k_μ. Therefore the methods are entirely applicable to any other Helmholtz eigenproblem regardless of the dispersion relation, or indeed the existence of an 'energy' (for instance E is physically irrelevant in acoustic problems). The only requirement is that the wavenumber be constant (and isotropic) in the interior.

The billiard has d-dimensional 'volume' $\mathcal{D}$ and $d-1$ dimensional 'surface area' A, giving a typical length scale $\mathsf{L} = \mathsf{V}/\mathsf{A}$. Our eigenproblem can be written

$$(\nabla^2 + k^2)\psi(\mathbf{r}) = 0 \qquad \text{for all } r \in \mathcal{D} \qquad \text{(Helmholtz equation)} \tag{5.5a}$$

$$f[\psi] = 0 \qquad \text{(boundary conditions)} \tag{5.5b}$$

$$g[\psi] = 1 \qquad \text{(normalization condition)} \tag{5.5c}$$

where the eigenvalue is k^2. The functionals f and g return a scalar, and are *quadratic* in the wavefunction ψ.

The BCs have been incorporated as (5.5a) rather than the linear condition (5.3) because satisfaction of the BCs by a wavefunction ψ is then revealed by a single number $f[\psi]$. This number measures the 2-norm of some error function, and is therefore a non-negative quantity. The error function (*e.g.* (5.3)) gives the amount by which the desired BCs fail to be obeyed. Heller[91] named f 'tension', and I shall follow suit. The definition of f is

$$f[\psi] \equiv \oint_\Gamma w(\mathbf{s}) d\mathbf{s} \, |(\mathcal{L}\psi)(\mathbf{s})|^2, \tag{5.6}$$

where $w(\mathbf{s}) \geq 0$ is a weighting function over the boundary. Arbitrary weighting schemes are possible, but I will stick to $w = 1$ which treats all locations on Γ equally. The case of Dirichlet BCs is given by $(\mathcal{L}\psi)(\mathbf{s}) \equiv \psi(\mathbf{r}(\mathbf{s}))$, put more simply, f measures the boundary integral of the squared value of the wavefunction. Generally my work will be restricted to this case from now on—however the same techniques would apply to other BCs.

Without a further condition, (5.5) and (5.5a) admit the useless solution, $\psi(\mathbf{r}) = 0$ for all $\mathbf{r}$. Therefore the quadratic functional

$$g[\psi] \equiv \langle \psi \, | \psi \rangle_{\mathcal{D}} \equiv \int_{\mathcal{D}} d\mathbf{r} |\psi(\mathbf{r})|^2 \tag{5.7}$$

which measures the normalization (norm) in the domain is required. Unit norm is fixed by (5.5b).

The solution $\psi(\mathbf{r})$ is now completely determined, when k reaches one of the eigenwavenumbers k_μ. For other k, no solution exists. Therefore in order to be able to define

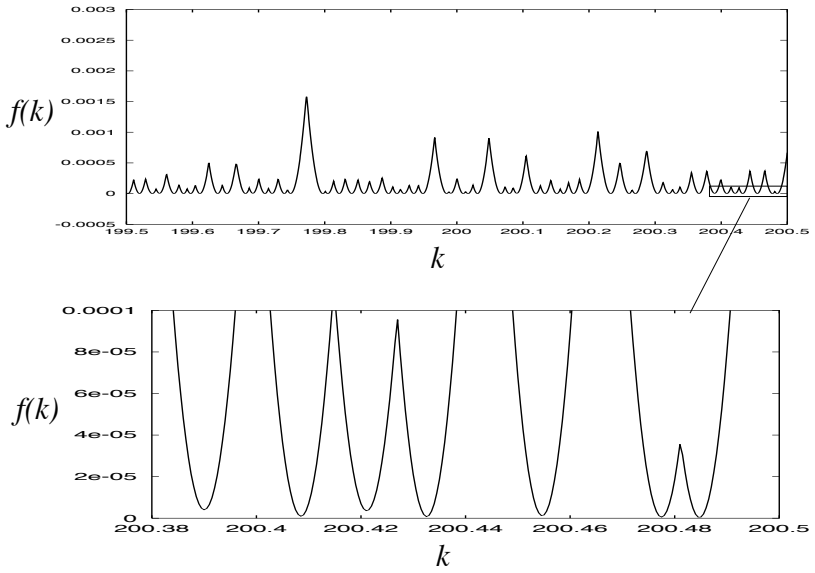

Figure 5.2: Tension (equal to the inverse largest eigenvalue of (5.12)), as a function of k. Minima correspond to billiard eigenstates. The finite value ϵ_μ of these minima is visible (for some cases) in the magnified plot.

a 'best' solution for *any* given guess at k, one of the conditions needs to be relaxed. The condition (5.5a) will be replaced by the minimization

$$f[\psi] \;=\; \min. \tag{5.8}$$

A sweep in k will now show the k-dependent tension $f(k)$, indicating how well the best ψ at each k can match the BCs. When k approaches an eigenwavenumber k_μ, the tension will drop to zero, indicating all the conditions (5.5), (5.5a) and (5.5b) are satisfied simultaneously. Such a sweep is shown in Fig. 5.2.

5.3 Representation by a Helmholtz basis

I choose a basis set composed entirely of solutions to (5.5) in $\mathcal{D}$. Therefore they are k-dependent (equivalently, energy-dependent), and this means we are joining the Class B methods. Since the basis functions do not obey the BCs we can choose the BCs for which the analytic forms of the functions are simplest, namely, free space. The eigenfunctions can be chosen real-valued everywhere, so we can choose entirely real basis functions. Typical basis functions at wavenumber k include: Real Plane Waves (RPWs), Evanescent Plane Waves (EPWs), and angular momentum states (*e.g.* the singular $Y_0(kr)$) based at various origins (which if they are singular, must be outside $\mathcal{D}$, or placed on Γ as in the BIM basis set). Each of these forms is oscillatory with wavenumber k. Following Heller's PWDM, the 'default' basis will be RPWs uniformly spaced in angle, containing both sin and cos terms, and suitably symmetrized for the billiard shape. The basis size N will be discussed below.

Helmholtz basis sets will be further discussed in Sections 5.3.1 and 6.1.3, and examples described in Appendix J.

Substitution of the k-dependent basis representation

$$\psi(\mathbf{r}) \;=\; \sum_{n=1}^{N} x_n \phi_n(k;\mathbf{r}). \tag{5.9}$$

into (5.8) and (5.5b) gives

$$\begin{cases} \mathbf{x}^{\mathrm{T}} F(k) \mathbf{x} = \min \\ \mathbf{x}^{\mathrm{T}} G(k) \mathbf{x} = 1, \end{cases} \tag{5.10}$$

where the k-dependent positive definite matrices are given by

$$F_{ij}(k) \equiv \oint_{\Gamma} w(\mathbf{s}) d\mathbf{s}\, \phi_i(k;\mathbf{s}) \phi_j(k;\mathbf{s}), \qquad G_{ij}(k) \equiv \int_{\mathcal{D}} d\mathbf{r}\, \phi_i(k;\mathbf{r}) \phi_j(k;\mathbf{r}). \tag{5.11}$$

The F matrix can be evaluated using methods in Appendix G, and the evaluation of G is discussed in Section 5.4.

The condition (5.10) is equivalent [7] to finding the maximum-eigenvalue solution to the generalized eigenproblem

$$G(k)\mathbf{x} \;=\; \lambda F(k)\mathbf{x}. \tag{5.12}$$

The maximum eigenvalue is λ_1, and λ_1^{-1} is then simply the tension $f(k)$ corresponding to the 'best' solution. The basis coefficients of this best solution are returned in the eigenvector $\mathbf{x}^{(1)}$.

Essentially, the method proposed in this chapter is then searching for local minima in the inverse largest eigenvalue of (5.12), as a function of k. An example sweep in k is plotted in Fig. 5.2. However there are implementation issues without which this specification is all but useless. The rest of the chapter is devoted to such issues.

5.3.1 Numerical rank of the basis set

Examination of the matrices F and G shows that they both become *singular* [188] at almost exactly the point when the basis set size (matrix order) N becomes sufficiently large to represent the solutions at the desired k to sufficient accuracy. This makes (5.12) a singular generalized eigenproblem. Such singularity causes failure of standard generalized matrix eigensolvers (for instance `dsygv` in LAPACK [5]). This might be seen as a huge setback—however I will show that singularity seems to be an inevitable result of using a basis which obeys the Helmholtz equation (however, the issue of this singularity in the special case of the BIM basis set is unknown). Also I will discuss correct handling of the singular problem.

Why do the matrices become singular? The answer is that physically, wavefunctions $\psi(\mathbf{r})$ which are arbitrary solutions of the Helmholtz equation at a single wavenumber k are effectively limited to a finite number of degrees of freedom inside a given billiard. This can be seen intuitively by realising that two real plane waves (RPWs) much closer in angular separation than $(k\mathsf{L})^{-1}$ (where L is the typical system length) cannot be distinguished

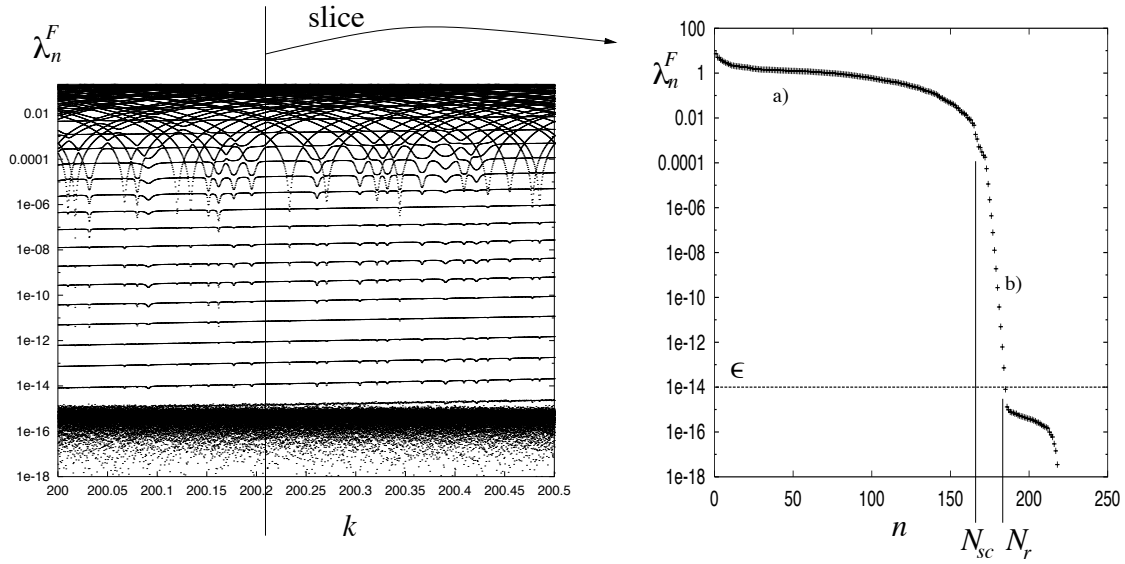

Figure 5.3: Eigenvalues of the 'tension matrix' F. Left side: sweep over k. Right side: slice at a single k, with λ_n sorted into decreasing order. The billiard is the quarter stadium of Fig. 2.6. The basis set used was a set of 250 real plane waves, symmetrized so as to fall into the odd-odd symmetry class. The semiclassical basis size at this k is $N_{sc} \approx 164$.

when only their values inside the billiard are considered. The overlap of the two RPWs in $\mathcal{D}$ becomes very similar to their individual norms in $\mathcal{D}$, *i.e.* they tend to the same function in $\mathcal{D}$. (In contrast, when considered over *all* space, the RPWs are always orthogonal). This would indicate that in $d = 2$ a basis size of $N \sim k\mathsf{L}$ is sufficient to represent all solutions at k. This has been formalised in a more rigorous fashion [64] for convex billiards in $d = 2$, as a *semiclassical basis size* $N_{sc} = kL/\pi$, where L is the perimeter (length of Γ). This is simply the number of half-wavelengths along the perimeter. (I believe the generalization to higher d would be that N_{sc} is given by the Weyl law for a $d-1$ dimensional 'volume' A).

Not only is this basis size sufficient, but attempts to improve accuracy by increasing N beyond this rapidly fail. Above N_{sc} it becomes exponentially hard to introduce new (non-evanescent) functions into the basis. This can be illustrated by considering an angular-momentum (l) basis around some origin inside $\mathcal{D}$. The inside wavefunction is strongly affected by states with $|l| < kR$ (where R is the maximum radius from the origin). For higher $|l|$ the billiard falls into the 'classically-forbidden' region (the circle of radius $\approx |l|/k$), so these l-states have an effect on ψ which dies exponentially with $|l|$. For these states, the only effect on ψ is an exponentially-small evanescent wave.

Fig. 5.3 shows the eigenvalues λ_n^F of the F (tension) matrix. These λ_n^F are simply the tensions produced by eigenvectors $\mathbf{x}^{(N)}$ of unit length in $\mathbf{x}$-space (coefficient space). The figure shows that up to N_{sc} the eigenvalues are $O(1)$, but beyond this they plummet exponentially, finally settling at $\sim 10^{-16}$ which is the rounding error of the double-precision arithmetic used. By analogy with singular values this defines a numerical *rank* [188, 80]: the number N_r of eigenvalues above a certain cut-off ϵ, which is chosen as roughly the working

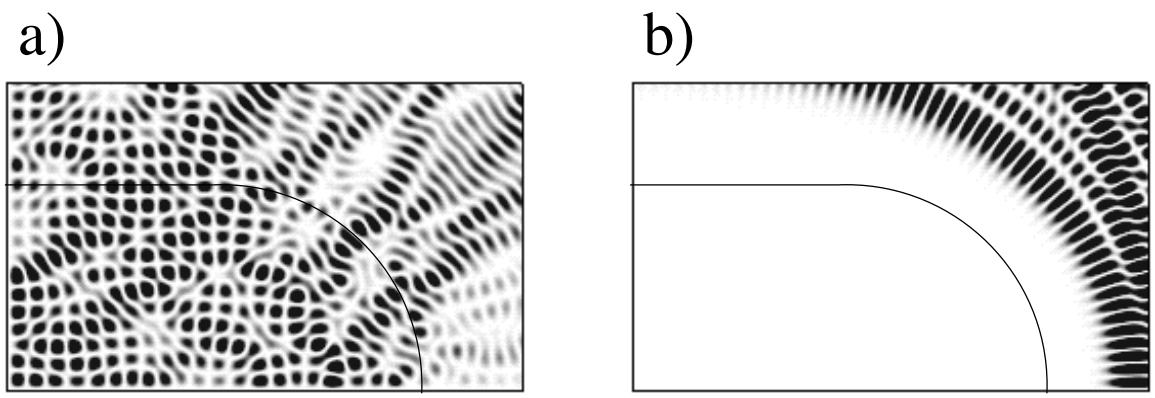

Figure 5.4: Eigenvectors $\mathbf{x}^{(n)}$ of the 'tension matrix' F at a single k: **a)** $\lambda_n^F \approx 1.3$, in the 'semiclassical region', **b)** $\lambda_n^F \approx 7 \times 10^{-10}$, in the rapid decay towards the null-space region. The corresponding regions are shown in the previous figure as a) and b). The billiard (shown in outline) and basis are the same as the previous figure, except at $k = 50$.

precision. Examining the figure, the F matrix for $N = 250$ has a rank of about 182, not much larger than $N_{sc} = 164$. The eigenvectors with $\lambda_n^F < \epsilon$ fall into the numerical *null-space*, and have a numerically irrelevant (tiny) effect on the value of $\mathbf{x}^{\mathrm{T}} F \mathbf{x}$, and therefore on the value of ψ anywhere within $\mathcal{D}$. An eigenvector heading towards the null space is shown in Fig. 5.4b; it has negligible value inside $\mathcal{D}$ and resembles a high angular-momentum state outside. What is surprising is that increasing N further has little or no effect on the rank N_r. As Vergini has realised[194], this behavior has profound consequences for the representation of billiard eigenstates by RPWs. Namely, once the rank of the problem is reached (at a particular k) it is very hard to improved accuracy simply by 'throwing' more real plane wave basis states at the problem (a method usually successful in other basis solution techniques). Most of the basis functions thus 'thrown' will simply fall into the numerical null space without contributing to the accuracy of the ψ found.

This self-truncation of the basis set size would appear quite bizarre to a numerical analyst trained on the likes of lattice and FEM approaches. In these latter approaches the basis set size N can be increased without limit, and increased N always results in higher accuracy of solution. However, because the basis sets of this chapter are already solutions of the wave equation, self-truncation occurs when the basis becomes (semiclassically) overcomplete. The exception to this is if suitable evanescent waves are added close to the boundary (see Appendix J). Therefore if highly accuracy is needed (or if the semiclassical limit is to be reached), EPWs must be added, as demonstrated by Vergini [194]. Discussion of the Helmholtz solution function space is continued in Section 6.1.3.

5.3.2 Truncating the singular generalized eigenproblem

Following [195, 194], the singular eigenproblem (5.12) is handled by restriction of the solution vector $\mathbf{x}$ to the non-singular subspace of F. This means that the numerical null-space of F is ignored. We have seen above (see Fig. 5.4) that this null-space corresponds

to directions in $\mathbf{x}$ which do not contribute to ψ inside the billiard to more than an accuracy of ϵ. There is an exception: if k is very close to one of the k_μ, then restriction of tension to ϵ does not guarantee small values inside, however the only non-small ψ which can possibly exist is the eigenfunction ψ_μ itself which is already represented. In conclusion it is fine to ignore this subspace.

Diagonalizing F gives $F = V\Lambda V^{\mathrm{T}}$ where V is orthogonal and $\Lambda = \mathrm{diag}(\lambda_1^F \cdots \lambda_N^F)$. This representation of F is now *truncated* by removing the eigenvalues smaller than ϵ thus: $\Lambda_r = \mathrm{diag}(\lambda_1^F \cdots \lambda_{N_r}^F)$, assuming descending eigenvalue order. V_r is set to the corresponding first N_r columns of V. The transformation

$$G\mathbf{x} = \lambda F\mathbf{x} \qquad \longrightarrow \qquad G'_r\mathbf{x}'_r = \lambda\mathbf{x}'_r \tag{5.13}$$

then follows, with the truncated F-representation of G defined by

$$G'_r = \Lambda_r^{-1/2} V_r^{\mathrm{T}} G V_r \Lambda_r^{-1/2}. \tag{5.14}$$

The eigenvectors are transformed forwards and backwards by

$$\mathbf{x}'_r = \Lambda_r^{1/2} V_r^{\mathrm{T}}\mathbf{x} \qquad \leftrightarrow \qquad \mathbf{x} = V_r\Lambda_r^{-1/2}\mathbf{x}'_r. \tag{5.15}$$

There is ambiguity in the latter relation since any vector in the numerical null-space could be added to $\mathbf{x}$, but by design this is not numerically significant when it comes to the resulting wavefunction $\psi(\mathbf{r})$ in $\mathcal{D}$.

Therefore the recipe is to diagonalize F, choose ϵ and construct Λ_r and V_r. Then to construct G'_r, which is diagonalized to give eigenvalues λ_n and unit-norm eigenvectors $\mathbf{x}'^{(n)}_r$, for $n = 1 \cdots N_r$. Finally the eigenvectors are rotated back using the second relation of (5.15), to give $\mathbf{x}^{(n)}$. In order to give the correct normalization $\mathbf{x}^{\mathrm{T}}G\mathbf{x} = 1$, corresponding to unit wavefunction normalization in the billiard, the eigenvectors $\mathbf{x}^{(n)}/\lambda_n^{1/2}$ should be used. In particular, the 'best' match to the BCs is given by the largest eigenvalue $n = 1$. The corresponding tension of this best normalized wavefunction is then λ_1^{-1}. This recipe was used to generate Fig. 5.2 by sweeping over k.

Note that it would also be possible to select the null-space using a different criterion. For instance since (5.12) is really symmetric between F and G, the numerical null-space of G could be used instead. Since both are sensitive to ψ inside the billiard only, this would be equally valid. However, it will be convenient to remain with the above choice for reasons apparent in Section 5.4.

5.3.3 Choice of ϵ for truncation

The criterion for choosing ϵ is that it be close to, but larger than, the machine precision[161] $\epsilon_{mach} \approx 10^{-16}$. I have found that pushing ϵ too close to ϵ_{mach} (or smaller) *does* affect the quality of the returned eigenvector $\mathbf{x}^{(1)}$ from the recipe in the previous section. One simple measure of this is the difference $\delta f \equiv [\mathbf{x}^{(1)}]^{\mathrm{T}} F\mathbf{x}^{(1)} - \lambda_1^{-1}$, which should be zero. For $\epsilon = 10^{-16}$ I found $|\delta f| < 10^{-6}$ while for $\epsilon = 10^{-14}$ it drops to $|\delta f| < 10^{-9}$. It is unknown how the contamination by other (error) eigenvectors varies with ϵ. However I believe that since the errors mainly occur in the evaluation of small eigenvalue/eigenvector pairs of F, any contamination is mainly by harmless (null-space) vectors.

On the other hand, ϵ also provides an approximate lower bound for the tension that can be reported. Therefore, in order to preserve the full depth of the tension minima, as small an ϵ as possible should be used. In practice, tension minima rarely are smaller than 10^{-8} using RPWs, but with better basis sets they can reach typical values of 10^{-11} in systems under investigation.

A choice of $\epsilon = 10^{-14}$ has proved optimal for calculations in this and the next chapter.

5.4 The choice of norm matrix G

Accurate evaluation of F is a rapid procedure because it only involves integration over the boundary Γ. It requires $O(NM)$ basis function evaluations, where the number of discretization points on the boundary is $M \sim k^{d-1}\mathsf{A}$ (see Appendix G). However the evaluation of the norm matrix G at first sight seems like a tremendous bottleneck. Eq.(5.11) for G involves computing N^2 overlaps $\langle \phi_i \, | \phi_j \rangle_{\mathcal{D}}$ each of which is a d dimensional integral over the domain. Since the $\phi(\mathbf{r})$ are oscillatory at scale $1/k$, each integral would require $O(k^d)$ basis function evaluations, and the construction of G would scale as k^{d+2}. If this were performed literally, then much of the advantage of a surface method would be lost. In fact, in any implementation of the PWDM as originally described, almost all the time is spent performing normalizations over the domain, which is wasteful.

The aim of this section is to describe more rapid evaluations of G, both approximate and exact, which scale like the boundary. (In Section 5.5 we will see that the accuracy of G needed is actually quite low).

5.4.1 Estimation by interior points—relation to Heller's PWDM

The simplest approximation to G involves approximating the domain integral by a sum over M_I interior points:

$$G_{ij}(k) \;=\; \frac{\mathsf{V}}{M_I} \sum_{m=1}^{M_I} \phi_i(k;\mathbf{r}_m)\phi_j(k;\mathbf{r}_m), \tag{5.16}$$

where $\mathbf{r}_m$ are points inside $\mathcal{D}$. In the case of a uniform grid spacing a for these points, convergence would be reached for $ka \ll 1$ (and resulting error converges like a, in numerical experiments), requiring a very large number $M_I \gg \mathsf{V}k^d$. In this case all the advantages of a boundary method are lost. However, a Monte Carlo estimate can be reached for a much smaller M_I: in general the estimation error will be Gaussian with a width $M_I^{-1/2}$ as a fraction of the value. This requires that the points be chosen in 'statistically independent' regions of the wavefunction, a criterion which for ergodic systems (excluding heavily-scarred states) it is easy to fulfill. The slow convergence means that accuracy cannot be high. Generally an optimal choice for M_I is such that the time spent evaluating G is no more than for F, preserving the 'boundary' scaling of the method. This method has the advantage of simplicity: only basis function values are required (the exact form below requires values and derivatives).

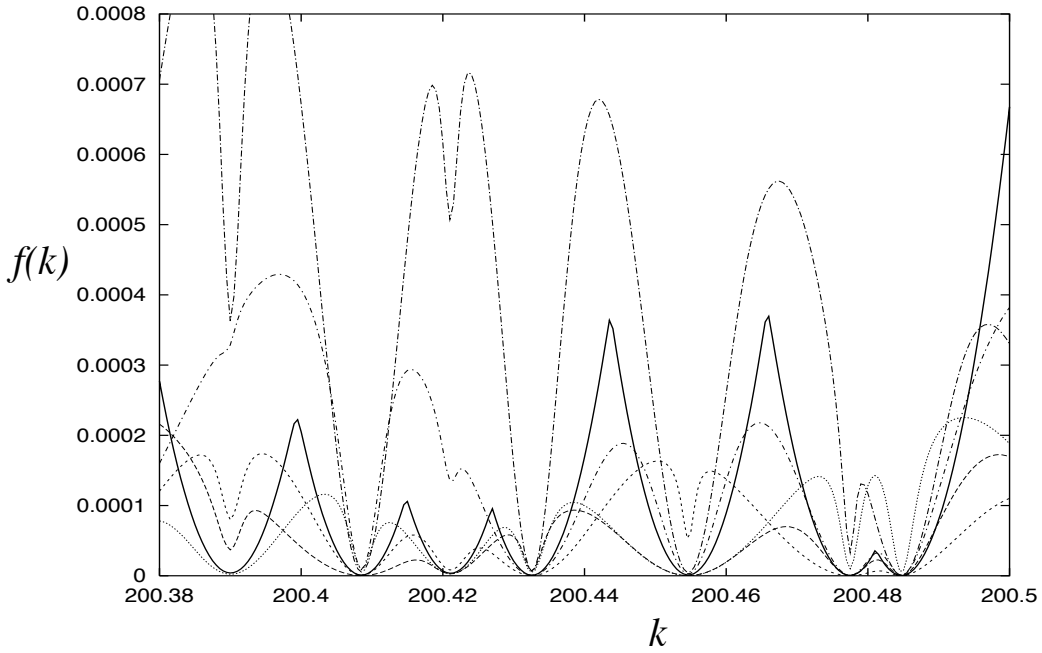

Figure 5.5: Tension minima using single-interior-point method for calculation of G (many dotted lines, each for different interior point choice), versus exact G (solid line). Note that the single-interior-point minima are often indistinct, and one is missed ($k = 200.390$).

The case $M_I = 1$ is interesting, since it corresponds to the original recipe of Heller [91]. This recipe involves finding the closest solution to

$$\begin{cases} A(k)\mathbf{x} &= \mathbf{0} \\ \mathbf{y}^{\mathrm{T}}\mathbf{x} &= 1 \end{cases} \tag{5.17}$$

where $A_{mn}(k) = \phi_n(k; \mathbf{s}_m)$, and $y_n = \phi_n(k; \mathbf{r}_1)$ at the randomly-chosen interior point $\mathbf{r}_1$. Heller's choice of the N, the number of columns of A, was about 1.4 times N_{sc}. The density of boundary matching points $\mathbf{s}_m$ was about 2.6 per wavelength. It is easy to see that taking the 2-norm of the above equations gives (5.10) for the version of G involving a single interior point. The result $F = A^{\mathrm{T}}A$ is needed (shown in Appendix G). The singular value decomposition of A is therefore equivalent to diagonalization of F [81]. Looking at Fig. 5.5 one can see that the effect of some of the eigenstate dips is almost gone, depending on the choice of $\mathbf{r}_1$. Also notice that the transition from one state to another is not very 'clean'. Even though Heller's prescription was then followed by another normalization step and a hunt for the resulting *normalized* tension minima, if an eigenstate never appears in the solution of the above for any k, then this normalization step cannot recover it—it becomes a 'missing state'.

What is the expected fraction of such missing states in Heller's method? An exact answer is hard, but empirically seems to be of the order of 1% for the $d = 2$ stadium billiard at $k < 400$. (Some users have reported as high as 8% in other systems [136]). If a true normalized billiard eigenstate has a very small value of $\psi(\mathbf{r}_1)$, being ε times the

RMS value, then this state's parabolic tension-minimum curve will be multiplied by $1/\varepsilon$ (this multiplication is visible in Fig. 5.5). If the 'dynamic range' r_d (see Section 5.5.3) is less than $1/\varepsilon$, then the parabola will typically be 'blown' to large enough values that it is obscured by those of neighboring states, and will never be noticed. Since the chances of $|\varepsilon| < r_d$ occurring are proportional to $1/r_d$ (the Porter-Thomas probability distribution on $\psi(\mathbf{r})$ is flat around $\psi = 0$), one would expect roughly $1/r_d$ states to be missed. The dynamic range decreases with increasing k so it is expected that this will become more severe a problem at high k.

5.4.2 Exact form on the boundary and Dirichlet approximation

It is possible to calculate G *exactly* using $O(NM)$ basis function (or derivative) evaluations. The equivalent expressions (H.14) and (H.7) both give the domain overlap of two Helmholtz solutions (with no particular BCs) at the same wavenumber, purely in terms of boundary integrals. Taking for example the former expression, it can be approximated (using the discretization of Appendix G) as

$$G \;=\; \frac{\mathsf{A}}{M}\,\frac{1}{2k^2}\left[C^{\mathrm{T}}C + D^{\mathrm{T}}D - k^2 A^{\mathrm{T}}A - (A^{\mathrm{T}}B + B^{\mathrm{T}}A)\right], \tag{5.18}$$

where the four rectangular matrices are defined by $A_{in} = p_i\phi_n(\mathbf{r}_i)$, $B_{in} = p_i^{-1}\mathbf{n}_i\cdot(\nabla\nabla\phi_n)(\mathbf{r}_i)\cdot\mathbf{r}_i$, $C_{in} = p_i\partial_x\phi_n(\mathbf{r}_i)$, and $D_{in} = p_i\partial_y\phi_n(\mathbf{r}_i)$. At each boundary point $\mathbf{r}_i$ for $i = 1\cdots M$, the weight is $p_i \equiv (\mathbf{r}_i\cdot\mathbf{n}_i)^{1/2}$. The outward normal at this point is $\mathbf{n}_i$. Notice that the (Hessian) matrix of second derivatives is required for B.

This formidable expression has been coded and tested. It has been found that for the Dirichlet eigenproblem, it can be replaced by a much simpler form which is valid only when the wavefunctions in question vanish on the boundary, namely (H.9). When written in a basis this gives

$$G_{ij} \;=\; \frac{1}{2k^2}\oint_\Gamma d\mathbf{s}\; r_n\,(\partial_n\phi_i)\partial_n\phi_j, \tag{5.19}$$

where ∂_n means derivative in the normal direction. The discretized version is therefore

$$G \;=\; \frac{\mathsf{A}}{M}\,\frac{1}{2k^2}C^{\mathrm{T}}C, \tag{5.20}$$

where $C_{in} = p_i\partial_n\phi_n(\mathbf{r}_i)$. The largest few resulting generalized eigenvalues λ (and eigenfunctions) of (5.12) are indistiguishable from those resulting from (5.18). That this works is not entirely obvious, because the basis functions ϕ_n do not vanish on the boundary. However all the relevant vectors $\mathbf{x}$ contain combinations of basis functions which *are* very small on the boundary. For the tension plots in this chapter, (5.20) has been used.

As a bonus, the eigenstates with Neumann BCs are to be found at the minima of the *smallest* λ. This is true because exchanging Dirichlet for Neumann in the tension function definition is equivalent to swapping F and G in (5.12).

5.5 The hunt for local tension minima as a function of k

For the reason that the basis is k-dependent (true for all Class B methods), a 'hunt' in k-space is required to locate the tension $f(k)$ minima.

5.5.1 Form of tension minima and resulting accuracy in k

As discussed in Section 5.3.1 (also see Appendix J), the values ϵ_μ of these minima at each k_μ is determined by the quality of the Helmholtz basis set. Around each minimum the tension is quadratic,

$$f(k) \;=\; \epsilon_\mu + \tfrac{1}{2}c_\mu(k-k_\mu)^2, \tag{5.21}$$

where the curvatures c_μ are very similar for all states μ (see Fig. 5.2). This happens because the lowest tension state of constant norm (that is, the state $\mathbf{x}_1$ returned by (5.12)) is simply the *scaling eigenfunction* μ with lowest tension (see Section 6.1.1). The tension expansion of this function about $k = k_\mu$ is quadratic with $c_\mu = 4$ exactly, independent of μ, for a modified weight function $w(\mathbf{s})$ (see Appendix I). Therefore with our choice $w(\mathbf{s}) = 1$, a curvature estimate is $c_\mu \approx 4\langle r_n^2\rangle/\langle r_n\rangle$ where the averages are taken over $\mathbf{s} \in \Gamma$, and $r_n \equiv \mathbf{r}\cdot\mathbf{n}$. If either there is little scarring or the billiard is close to spherical, then we expect this estimate to be good.

Eq.(5.21) defines a natural 'error scale' for the k_μ found: a change k that is much less than $(\epsilon_\mu/\mathsf{L})^{1/2}$ from a tension minimum causes only a small fractional change in tension. Therefore I will call $\Delta k(\epsilon_\mu) = (\epsilon_\mu/\mathsf{L})^{1/2}$ the 'tension rounding error' of the state μ, because it defines an uncertainty in k_μ due to the basis-dependent finite tension minima. In practice, the errors in k_μ found appear to be an order of magnitude better than $\Delta k(\epsilon_\mu)$, as shwon by the lines in Fig. 6.6.

5.5.2 Finding the nearest minimum

The abrupt transition from one parabola to the next therefore happens very close to the midpoints between the k_μ. Therefore it is possible to use the quadratic form to make an efficient search routine for the nearest state to a given k: $f(k_1)$ and $f(k_2)$ can be used to solve for the current k_μ and ϵ_μ, given that k_1 and k_2 are nearest to the same k_μ. This can be iterated until k_μ settles to some accuracy. I will not present detailed results, however the search seems to require a few (less than 10) tension evaluations per state found. Essentially, this is the Newton-Rapheson [161] scheme applied to the derivative $f'(k)$. This method is similar to the 'improved PWDM' of Li [135], who claims 2-3 evaluations per state found (if high accuracy is needed). In the Boundary Integral Method (BIM or BEM) literature, a recent innovation named the 'Multiple Reciprocity BEM' has arisen [101, 100], which seems very similar (it seems to be application of 'power iterations' [81]).

The accuracy of the norm matrix G need not be very high for location of tension minima (a fractional error of $O(1)$ is even acceptable). Therefore estimation by a few interior points is fine (see Section 5.4). But the following question arises: Would an inaccurate norm create false fluctuations in the tension which would be mistaken for minima? The answer turns out to be no. This is because the wavefunction associated with the maximum (generalized) eigenvalue λ_1, being a scaling eigenfunction, is actually very stable (this was noted in [91]) as k is changed through k_μ, until the abrupt transition to the next state happens. Therefore the inaccuracy makes a jump only at these transitions, not around k_μ, so no spurious minima are created.

In the next chapter we will see that special properties of a certain weighting function $w(\mathbf{s})$ allow a vastly superior method, which bypasses the need to 'hunt' for states at

all.

5.5.3 Dynamic range and sweep method breakdown

I define the 'dynamic range' r_d as the ratio of the tension fluctuation to the typical tension minima found over a sweep in k. The tension fluctuation is given roughly by the value reached inbetween states separated by Δ_k, the mean level spacing in wavenumber. Using the above curvature estimate $c_\mu \sim \mathsf{L}$, and the Weyl law $\Delta_k \sim k^{1-d}/\mathsf{V}$, the fluctuation is $\sim k^{2-2d}\mathsf{L}^{1-2d}$. An estimate for dynamic range is therefore $r_d \sim (\epsilon_{\rm typ} k^{2d-2}\mathsf{L}^{2d-1})^{-1}$ where $\epsilon_{\rm typ}$ is a typical value of ϵ_μ. The scaling of $\epsilon_{\rm typ}$ with k is currently not known, but with appropriate basis sets in the 2D stadium it appears that arbitrarily small $\epsilon_{\rm typ}$ can be created, even in the semiclassical limit. (For a discussion of the diffractive contribution to $\epsilon_{\rm typ}$, see [194]). However, in billiards where such good basis sets are not known (and therefore RPWs are used), it seems empirically that $\epsilon_{\rm typ}$ is independent of k. Therefore in $d = 2$, we expect r_d to decrease like k^{-2}.

The method of this chapter, and sweep methods in general which rely on hunting for minima in a single quantity, will certainly break down when r_d approaches 1 or less. By 'break down' I mean that since most minima are no longer distiguishable (or they will be eclipsed by the parabolas from the states with smallest ϵ_μ), and the fraction of missing or erroneous states approaches 1. Encouragingly, we will find that the method of the next chapter can distinguish close eigenstates which are indistinguishable using a sweep method (with the same basis set).

5.5.4 A more intelligent hunt?

What are the prospects for a more intelligent 'hunt' procedure? It is interesting that the individual matrix elements of $F(k)$ and $G(k)$ are linear up to a k-scale of $\sim \mathsf{L}^{-1}$, the inverse system length (this is because the basis functions generally oscillate at a wavenumber k). This scale is $O(k^{d-1}\mathsf{A})$ times larger than the average level-spacing in k. This implies that information about *all* the tension minima in a k range $O(\mathsf{L}^{-1})$ might be contained in F and G and their k-derivatives at a single k value.

To take a 'toy problem' example, imagine the values of k are desired such that a parameter-dependent order-N symmetric matrix $T(k)$ has a zero eigenvalue. We assume linear dependence $T(k) = T + kS$. The zero eigenvalue condition is written $(T + kS)\mathbf{x} = \mathbf{0}$. This is simply the generalized eigenvalue equation between T and S, whose once-off diagonalization predicts all N solutions of k. This sounds promising, however the types of zero-crossings produced in the eigenvalues of $T(k)$ are linear (passing through zero with finite slope, changing sign in the process). Unfortunately the methods of this chapter require detection of eigenvalues of positive definite matrices which reach (close to) zero in a quadratic fashion (they cannot change sign). Therefore the above trick is no help.

In fact, the above linearization of $F(k)$ and $G(k)$ is deceptive, since their positive-definiteness cannot be maintained without considering higher-order powers of k. The positive-definite nature of F and G arises because they are the square of other matrices (*e.g.* $F = A^{\rm T}A$, see Appendix G). It is these other matrices whose entries can unproblematically be linearized in k. If we imagine again the toy problem now with $T(k) = A^{\rm T}(k)A(k)$

and $A(k) = A + kB$. Eigenvalues of T are given by squares of singular values of A [81]. So now the problem is to predict the singular value zero-intersections of $A(k)$ (of course, this is the problem common to all Class b) methods, and is therefore of huge interest). Unfortunately the *generalized singular value decomposition* [81] of A and B does not predict these k values. Even though all the information about the k values is contained in A and B, I am unaware of a suitable decomposition which returns these values. It is an area for future research, and would have a huge impact on the large physics and engineering community currently using Class b) methods.

One untested idea on this front is the presentation of the toy problem as $(T + kD + k^2E)\mathbf{x} = \mathbf{0}$, where $D = B^{\mathrm{T}}A + A^{\mathrm{T}}B$ and $E = B^{\mathrm{T}}B$. It is possible to convert this nonlinear eigenvalue problem into a linear one of order $2N$ (see *e.g.* [161]),

$$\begin{pmatrix} 0 & 1 \\ -E^{-1}T & -E^{-1}D \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = k \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix}. \tag{5.22}$$

The factor 8 reduction in diagonalization time would only be bearable if more than a few k-values were returned each time. Tests of this are needed.

5.6 Conclusion and discussion

In this chapter I have presented the main components of an improved 'sweep' method (so named because a sweep in k is required to locate eigenstates) for billiard diagonalization. It falls squarely in Class b) of Table 5.1, although the zero-determinant condition has been replaced by a maximum-eigenvalue one. The resulting eigenstates once found can be normalized accurately using the boundary formula for $\langle \psi \, | \psi \rangle_{calD}$ presented in Appendix H. The method is a simplified but more rigorous version of Heller's original PWDM, and reduces to the key component of Heller's method in the case of a particular choice of G.

Several problems in Heller's method have been solved. The choice of basis size N in Heller's method requires careful adjustment ($\pm 10\%$ or so), however the generalized eigenproblem approach of this chapter truncates N automatically. The method has been reduced to a simpler algorithm which is entirely expressible on the boundary, and therefore scales with the area of the boundary. By reducing to an order-N symmetric eigenproblem, the choice $M \gg N$ can be made without the disadvantage of slowing down any determinant (SVD) evaluations. This enables very accurate boundary sampling without slowing down the $O(N^3)$ matrix operations. N and M are no longer parameters whose choice has a sensitive effect on the success of the method; they simply need to be sufficiently large. The accurate norm matrix approaches of Section 5.4 gives clean 'switching' between states corresponding to tension minima—this is an improvement over Heller's formulation, and would allow such minima to be distiguished up to a higher k.

The scaling of the numerical effort required per found eigenstate is not hard to compute. The basis size N is of order the semiclassical basis size ($N_{sc} = kL/\pi$ in $d = 2$, $N_{sc} \sim k^{d-1}\mathsf{A}$ in higher d). The matrices F and G are dense, requiring $O(N^3)$ effort to diagonalize (5.12) at a single k. A few of these are required per state found. Therefore the effort per state found scales like $k^{3(d-1)}$. The accuracy in k found is determined entirely by

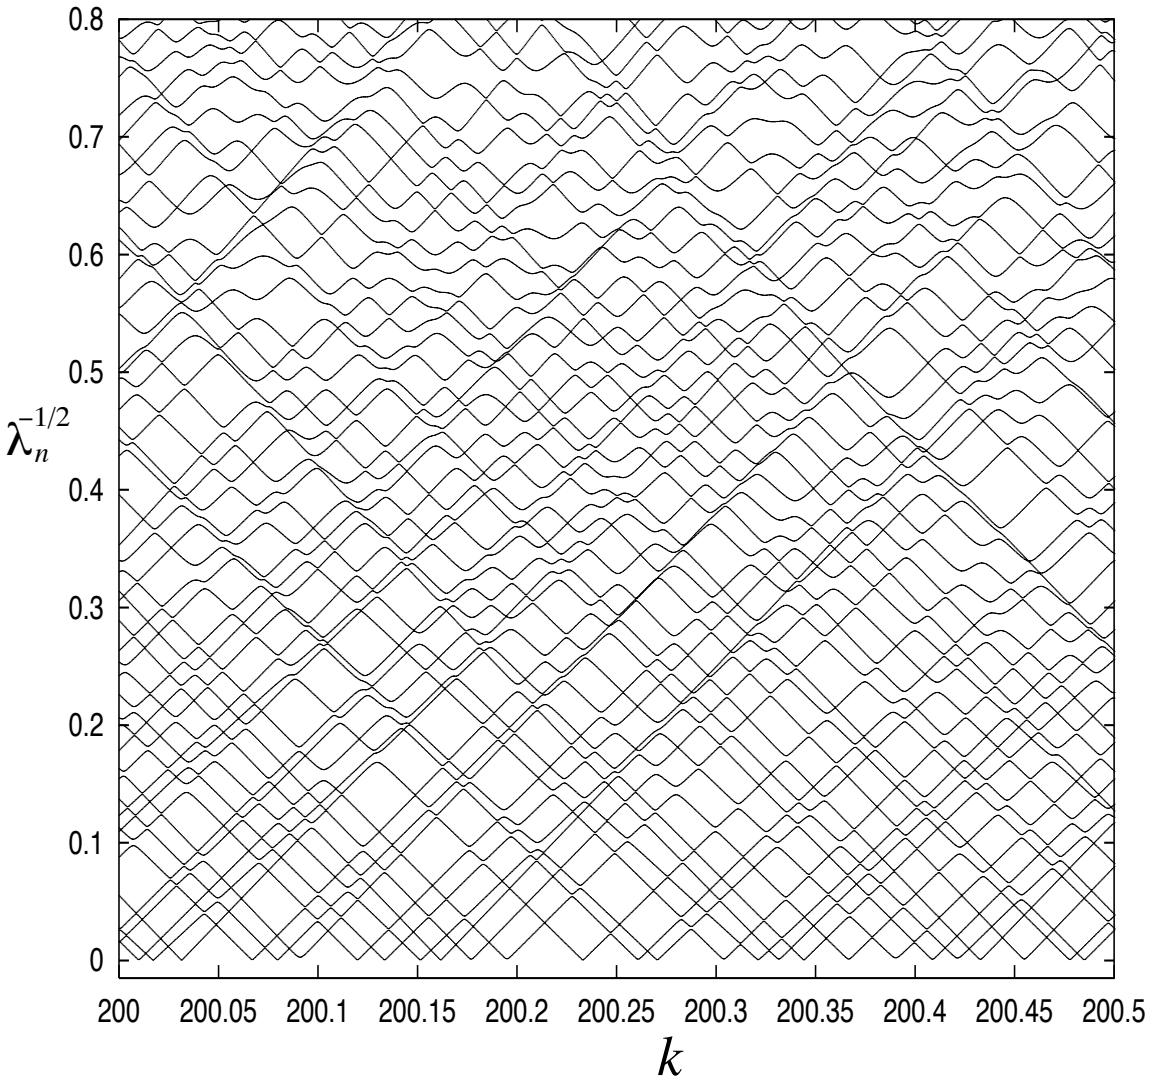

Figure 5.6: Inverse square-root of higher generalized eigenvalues of (5.12), as a function of k. As before, the minima in the largest eigenvalue correspond to billiard eigenstates. The long-lived nature of the parabolic form of λ_n^{-1} is visible as the survival of the straight lines over a range of $O(\mathsf{L}^1)$ around these minima. Note that avoided crossings are small, tending to zero at the bottom of the plot.

the basis set limitations: it appears to be an order of magnitude better than the 'tension rounding error' $\Delta k(\epsilon_\mu)$.

The method of this chapter breaks down when the maxima of tension reached between k_μ values becomes about equal to the typical tension minima ϵ_μ, in which case the fraction of 'missing states' becomes large. Even long before this is reached, close eigenvalues become indistiguishable (see Fig. 6.5). Improved basis sets which can reduce ϵ_μ at a given k require the use of evanescent waves: the reduction can then be by orders of magnitude in certain billiard shapes (see Fig. 6.6).

Finally, it is clear that there is *much* information that is being wasted by the whole 'tension minima hunt' approach described here and very common elsewhere (including in the vast array of engineering applications of the BIM). Fig. 5.6 shows the tension eigenvalues (inverse generalized λ_n of (5.12)). The square root is plotted so that parabolas become straight lines, for clarity. (The finite ϵ_μ actually round the bottoms to give hyperbolae). It is clear that the eigenstates associated with each minimum are in fact long-lived as k changes over many level-spacings, and therefore prediction of a large number of eigenstates from a single diagonalization is possible. This is the essence of the method of the next chapter.

Chapter 6

The scaling method of Vergini and Saraceno

Here I explain and analyse the powerful scaling method recently invented by Vergini and Saraceno (VS) [195, 194]. It allows a much more rapid solution of the Helmholtz eigenvalue problem (5.2) in a hard-walled billiard (with Dirichlet boundary conditions) than the sweep methods of the previous chapter, or indeed than any other known method. The success of the method is essentially due to a single key relation, namely that eigenstates are *quasi-orthogonal on the boundary* (or equivalently, a certain tension matrix is quasi-diagonal near the diagonal in a basis of rescaled eigenstates). However, an adequate explanation for the quasi-orthogonality has been sorely lacking. The knowledge gained about the band profile of special deformations in Chapter 3 will provide, for the first time, an understanding of quasi-orthogonality. This improves upon an (incorrect) guess made by VS, and allows an accurate semiclassical estimate to be made of the deviation from exact orthogonality.

I give a complete, presentation of the scaling method (Section 6.1), discuss the utility of higher-order corrections to the basic method (Section 6.2), and analyse the various sources of error (Section 6.3). Conclusions and open issues are brought up in Section 6.5. Appendix I shows how to evaluate higher-order terms in the tension matrix, correcting an oversight (to do with boundary curvature) in the original formulation of VS.

Generally for illustration I will use results for the 2D stadium billiard, the same shape studied by VS. This is simply for the reason that excellent evanescent basis sets are known for this shape, allowing the errors of the method to be distinguished from those of the basis set. In the last section I will demonstrate an application to computing the eigenstates of a stadium deformed by various *finite* amounts. The overlaps of these states with the undeformed states will be studied. These calculations were published in our work [48]. Note that it is probably useful to have read the previous chapter to fully appreciate this one.

6.1 The basic scaling method

We are interested in finding eigenstates of a d-dimensional billiard whose domain $\mathcal{D}$ has 'volume' V and whose boundary Γ has 'surface area' A. Therefore the typical system

size is $\mathsf{L} \sim \mathsf{V}/\mathsf{A}$. The scaling method will find all the states in a range $O(\mathsf{L}^{-1})$ around a given wavenumber k, using a *single* diagonalization of the same numerical effort required for *each* k evaluation in the previous chapter. Up to $N_{sc}/10$ states of useful accuracy are returned per diagonalization, making the method many orders of magnitude faster than any other method known at this time.

6.1.1 Tension matrix in a scaling eigenfunction basis

It is important to define a spatial rescaling of a wavefunction, corresponding to dilation (or contraction) about the (arbitrary) origin. If the wavefunction $\psi(\mathbf{r})$ is a Helmholtz solution at wavenumber k_0, then the rescaled version $\psi(k, \mathbf{r})$ of the same wavefunction obeys the Helmholtz equation at a new wavenumber k:

$$\psi(k, \mathbf{r}) \;\equiv\; \psi(\frac{k}{k_0}\mathbf{r}). \tag{6.1}$$

The rescaled function value at a given point is found by looking at the original function at a point a fractional distance $k/k_0 - 1$ further from the origin. The derivative with respect to k of the wavefunction value at a point $\mathbf{r}$ is

$$\left.\frac{\partial}{\partial k}\psi(k, \mathbf{r})\right|_{k=k_0} \;=\; \frac{1}{k_0}\left.\frac{\partial}{\partial \alpha}\psi(\alpha\mathbf{r})\right|_{\alpha=1} \;=\; \frac{1}{k_0}\mathbf{r}\cdot\nabla\psi. \tag{6.2}$$

This means that for $k = k_0 + \delta$ with δ small ($\ll 1/\mathsf{L}$), this can be expanded in powers of δ as

$$\psi(k_0 + \delta, \mathbf{r}) \;=\; \left[1 + \frac{\delta}{k_0}\mathbf{r}\cdot\nabla + O(\delta^2)\cdots\right]\,\psi(\mathbf{r}). \tag{6.3}$$

We will not need any higher terms for the basic method (they are given in Appendix I).

We now choose a basis $\psi_\mu(k, \mathbf{r})$ of the exact billiard eigenfunctions scaled as above to have a single wavenumber k; these are called *scaling eigenfunctions* [195]. The unscaled eigenfunctions are taken to have unit normalization in the billiard. The index $\mu = 1\cdots\infty$ labels the eigenfunctions in ascending order of their unscaled wavenumber k_μ. The scaling eigenfunctions no longer go to zero on the original billiard boundary Γ, however they do go to zero on rescaled (dilated or contracted) versions of the boundary Γ_μ (see Fig. 6.1). Imagine that a state μ has a wavenumber shift $\delta_\mu \equiv k - k_\mu$ which is very small ($\ll 1/\mathsf{L}$, the inverse system size). Then Γ will be very close to Γ_μ, and the value of $\psi_\mu(k, \mathbf{r})$ for $\mathbf{r} = \mathbf{r}(\mathbf{s}) \in \Gamma$ will be close to zero. The boundary coordinate $\mathbf{s}$ measures location on Γ. Applying (6.3), the difference from zero is given by

$$\psi_\mu(k, \mathbf{s}) \;=\; \frac{\delta_\mu}{k_\mu} r_n \partial_n \psi_\mu + O(\delta_\mu^2)\cdots \tag{6.4}$$

where $\partial_n\psi_\mu$ is the (unscaled) eigenstate normal derivative at location $\mathbf{s}$. Notice that the zeroth-order term vanished because of the BCs. (This expansion is continued in Appendix I). A general sum of scaling eigenfunctions

$$\psi(k, \mathbf{r}) \;=\; \sum_\mu \tilde{x}_\mu \psi_\mu(k, \mathbf{r}) \tag{6.5}$$

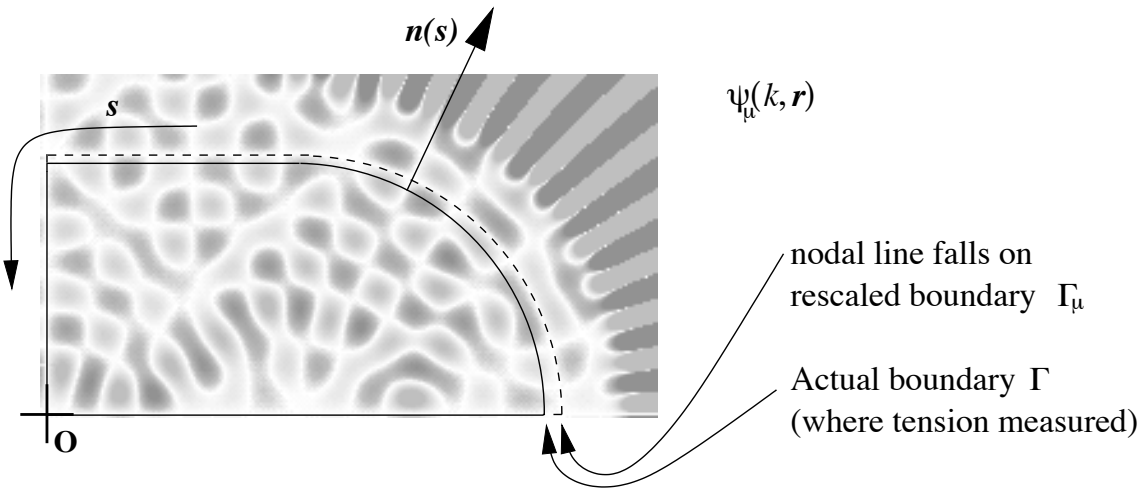

Figure 6.1: Scaling eigenfunction $\psi_\mu(k,\mathbf{r})$ vanishes on a rescaled boundary Γ_μ. It is clear that when Γ_μ is much closer than a wavelength from Γ, the wavefunction can be linearized, giving a value on Γ proportional to the eigenfunction normal derivative. Also the boundary coordinate $\mathbf{s}$ and local surface normal $\mathbf{n}(\mathbf{s})$ are shown.

is also a scaling wavefunction.

The tension (as defined in the previous chapter for Dirichlet BCs) of a general scaling wavefunction is

$$f(k) \;=\; \oint_\Gamma d\mathbf{s}\, w(\mathbf{s})\, |\psi(k,\mathbf{s})|^2, \tag{6.6}$$

where $w(\mathbf{s})$ is a boundary weighting function. If ψ is now written in the form (6.5) then the tension is $\tilde{\mathbf{x}}^{\mathrm{T}}\tilde{F}\tilde{\mathbf{x}}$. The matrix $\tilde{F}$ is the scaling eigenfunction representation of the tension quadratic form, and we will see that it has wonderful properties. For $\delta_\mu, \delta_\nu \ll 1$ we can substitute (6.4) as follows,

$$\begin{aligned}
\tilde{F}_{\mu\nu}(k) &= \oint_\Gamma d\mathbf{s}\, w(\mathbf{s})\, \psi_\mu(k,\mathbf{s})\psi_\nu(k,\mathbf{s}) \\
&= \frac{\delta_\mu\delta_\nu}{k_\mu k_\nu} \oint_\Gamma d\mathbf{s}\, w(\mathbf{s}) r_n^2 (\partial_n\psi_\mu)\partial_n\psi_\nu \;+\; O(\delta^3)\cdots \\
&= 2\delta_\mu\delta_\nu M_{\mu\nu} \;+\; O(\delta^3)\cdots,
\end{aligned} \tag{6.7}$$

where r_n is an abbreviation for $\mathbf{r}\cdot\mathbf{n}$ at the location $\mathbf{s}$. The final form has been written in terms of the boundary inner-product matrix of the eigenstate normal derivatives,

$$M_{\mu\nu} \;=\; \frac{1}{2k^2} \oint_\Gamma d\mathbf{s}\, D(\mathbf{s})\, (\partial_n\psi_\mu)\partial_n\psi_\nu, \tag{6.8}$$

where $D(\mathbf{s}) = r_n^2 w(\mathbf{s})$ is the weight defining the inner-product. The approximation $k \approx k_\mu \approx k_\nu$ has been used, because M has only very slow dependence on k. Note that all the fast dependence of F on k appears in the wavenumber shifts δ_μ and δ_ν.

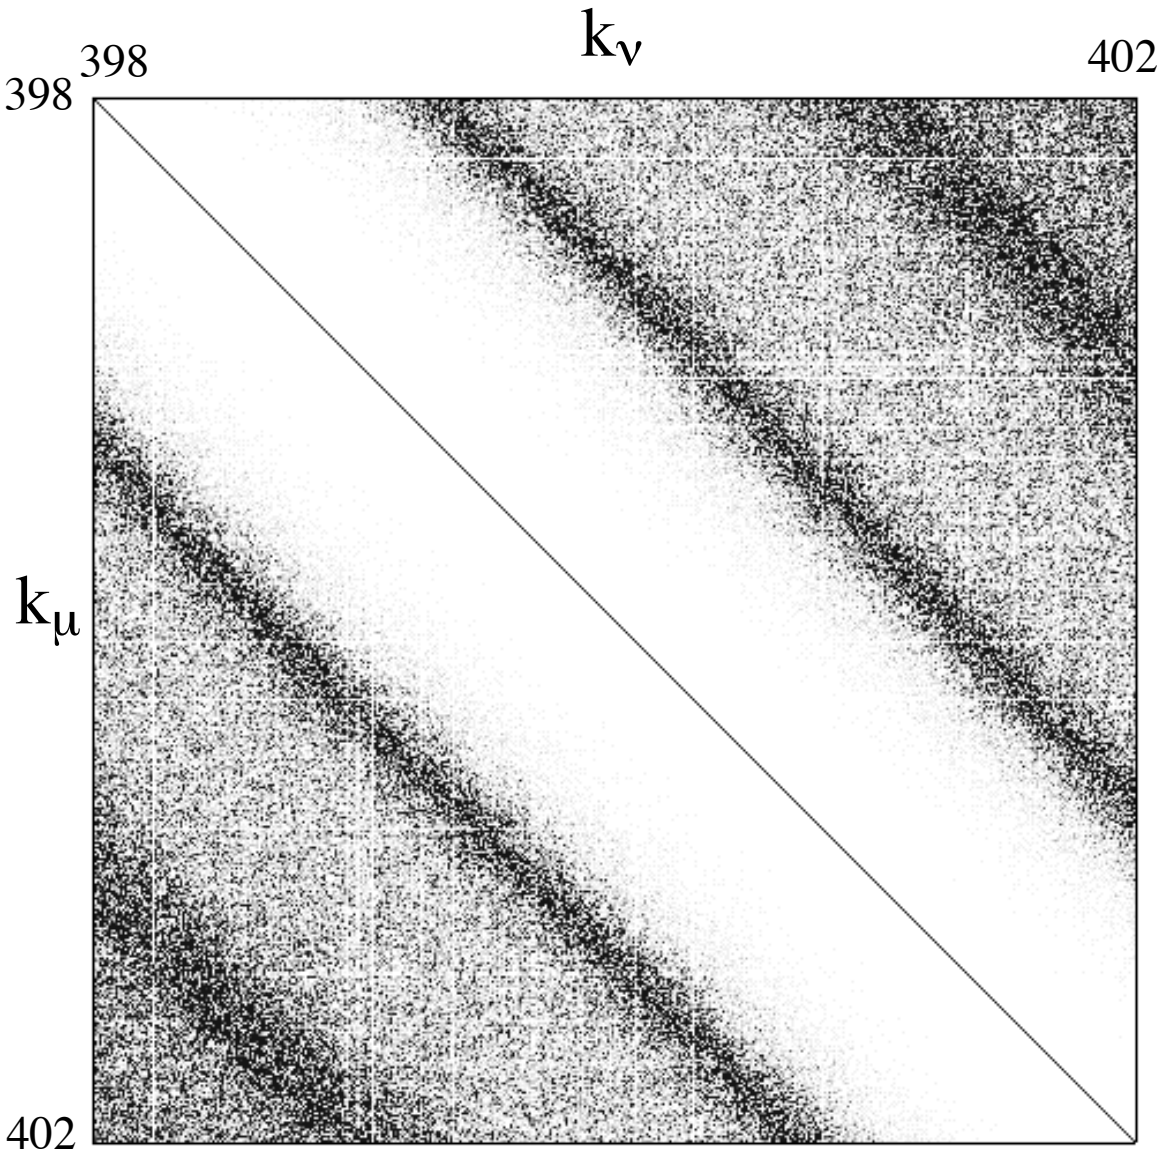

Figure 6.2: Image of the matrix $M \propto \partial\mathcal{H}/\partial x$ shown as a density plot of $|\partial\mathcal{H}/\partial x_{\mu\nu}|^2$, for the case of dilation (a special deformation) of the 2D quarter stadium billiard. Darker pixels correspond to larger values. The wide central white region (of width $\sim 1/\mathsf{L}$) corresponds to quasi-orthogonality: the quantum band profile vanishes like κ^4 as $\kappa \to 0$, where $\kappa \equiv k_\mu - k_\nu$. The diagonal is unity, and is not included in the band profile. The matrix involves all 451 eigenstates falling in the wavenumber range $398 < k < 402$.

6.1.2 Quasi-orthogonality on the boundary

The aim of this section is to show that $M_{\mu\nu}$ can be made very close to the identity matrix $\delta_{\mu\nu}$ via a certain choice of weighting function $w(\mathbf{s})$. The matrix M can be related to the rate of change of the Hamiltonian under deformation of Γ, written in the energy basis. $D(\mathbf{s}) = \mathbf{n}(\mathbf{s}) \cdot \mathbf{D}(\mathbf{s})$ is then interpreted as the deformation function associated with a deformation 'field' $\mathbf{D}(\mathbf{r})$ (see Section 3.1). Looking at Appendix C we have

$$M_{\mu\nu} \;=\; -\frac{m}{(\hbar k)^2}\left(\frac{\partial \mathcal{H}}{\partial x}\right)_{\mu\nu}. \tag{6.9}$$

Thus all the knowledge gained in Chapters 2 and 3 about the structure of $\partial\mathcal{H}/\partial x$ can be applied. In particular, $\partial\mathcal{H}/\partial x$ is a zero-mean random quantity whose variance an energy $\hbar\omega$ from the diagonal is given by the 'band profile' $\tilde{C}_{\rm E}(\omega)$ (see Eq.(2.48)). The band profile is to be evaluated at overall energy $E = (\hbar k)^2/2m$, where k is the typical wavenumber of interest.

In order to allow the scaling method to extract the eigenstates ψ_μ, we seek a choice of $D(\mathbf{s})$ (and therefore $w(\mathbf{s})$) which makes $M_{\mu\nu}$ as close as possible to the identity matrix $\delta_{\mu\nu}$. This is the issue of orthogonality on the boundary. Note that in the domain $\mathcal{D}$ there is an exact orthogonality relation for the eigenstates: $\langle\psi_\mu|\psi_\nu\rangle_{\mathcal{D}} = \delta_{\mu\nu}$. There is no such exact relation on the boundary. One cannot define a Hilbert space which consists of the boundary functions (eigenstate normal derivatives) at a single wavenumber k. Still, the boundary functions 'live' in an effective Hilbert space of dimension $\sim (k\mathsf{L})^{d-1}$. This dimension is just the semiclassical basis size N_{sc} of Section 5.3.1. On the boundary the best that can be done is a *quasi-orthogonality* relation—approximate orthogonality of N_{sc} adjacent eigenstate boundary-derivative functions, which span a range $\sim 1/\mathsf{L}$ in wavenumber.

The smallest off-diagonal elements of M close to the diagonal are found when $D(\mathbf{s})$ corresponds to a 'special deformation'. As discussed in Section 3.3, the band profile $\tilde{C}_{\rm E}(\omega)$ then vanishes as $\omega \to 0$ like a power law ω^γ. On the other hand, the diagonal elements of M are given by the average generalized force $F(x)$ on the deformation (see Section 3.1). So far the only special deformation known to the author (see Appendix D) which has a non-zero $F(x)$ is *dilation* $\mathbf{D}(\mathbf{s}) = \mathbf{r}$. (Translations and rotations do not change the billiard volume so have $F(x) = 0$). For dilation, the diagonal elements of M are all unity (proofs of this for $d = 2$ are known [28, 33]; I have proved it for arbitrary d with Eq.(H.9) of Appendix H—also see our work[14]). The power law is $\gamma = 4$, a result which has a classical origin (Section 3.3.1). The quasi-orthogonality of M for dilation is illustrated by Fig. 6.2.

The weighting function which gives $D(\mathbf{s}) = r_n$ is the one found by VS [195], namely $w(\mathbf{s}) = 1/r_n$. Therefore from now on I will assume this weighting, unless stated otherwise. It seems that in order to prevent w from diverging, this introduces the restriction that there must exist an origin from which all parts of the billiard wall are directly 'visible' (connected by a chord which always remains inside $\mathcal{D}$), so that r_n never passes through zero. Such billiards have been called 'star-shaped'[195]. However (despite protest to the contrary[39]), it seems that a slight generalization can be made to include non-simply-connected billiards composed of a star-shaped exterior surface and a star-shaped (concentric) interior surface, without breaking the above restriction on r_n.

Semiclassical estimate of off-diagonal strength of M

This and the next section can be skipped on first reading. We would like an estimate of the strength of the off-diagonal part of M, which we call V. Writing the strength of V in terms of the band profile, we have

$$M_{\mu\nu} = \delta_{\mu\nu} + V_{\mu\nu}, \tag{6.10}$$

$$\langle |V_{\mu\nu}|^2 \rangle \approx \frac{(2\pi)^d}{\Omega_d} \frac{1}{\mathsf{V}} \frac{m}{2\pi\hbar^3 k^{2+d}} \tilde{C}_{\rm E}(\omega) \tag{6.11}$$

$$\approx \frac{1}{\Omega_d} \frac{1}{\mathsf{V}} \left(\frac{2\pi}{k}\right)^{d-1} \tilde{C}(\kappa), \tag{6.12}$$

which followed from (6.9), (2.48) and the Weyl form[35, 12] for the mean level spacing

$$\Delta = \frac{2}{\Omega_d \mathsf{V}} \left(\frac{2\pi^2\hbar^2}{m}\right)^{d/2} E^{1-d/2}, \tag{6.13}$$

where $\Omega_d \equiv 2\pi^{d/2}/\Gamma(d/2) = 2\pi, 4\pi \cdots$ is the surface area of the unit sphere in $d = 2, 3 \cdots$ dimensions. The 'distance' from the diagonal is expressed either as a frequency $\omega = (E_\mu - E_\nu)/\hbar$ or as the wavenumber difference $\kappa \equiv k_\mu - k_\nu$. The two are related by $\omega = v_{\rm E}\kappa$, where the particle speed is $v_{\rm E}$ at energy E. The band profile has been written in a scaled form $\tilde{C}_{\rm E}(\omega) = m^2 v_{\rm E}^3 \tilde{C}(\kappa)$. This scaled band profile $\tilde{C}(\kappa)$ depends only on the billiard geometry (it is independent of k), because the system has hard walls. It has universal forms at both large and small 'distances' from the diagonal. The crossover happens at typical correlation scales of the system, given by the system size, namely at $\kappa \sim 1/\mathsf{L}$. In this intermediate region there is non-universal system-specific structure—*e.g.* see the band profile in Fig. 3.7 (graph DI). The equivalence of classical and quantum band profiles was demonstrated in Section 2.3.

For $\kappa \gg 1/\mathsf{L}$ we have $\tilde{C}_{\rm E}(\omega) \approx \nu_{\rm E}^{\rm WNA}$ (see Section 3.2), which upon substitution into (6.11) gives

$$\langle |V_{\mu\nu}|^2 \rangle \approx \frac{2\langle |\cos\theta|^3 \rangle}{\Omega_d} \frac{1}{\mathsf{V}^2} \left(\frac{2\pi}{k}\right)^{d-1} \oint_\Gamma ds\, r_n^2 \qquad \text{for } \kappa \gg 1/\mathsf{L}. \tag{6.14}$$

This corresponds to a random-wave assumption for the eigenstates (Section 3.2.1), giving no dependence on κ (flat band profile). Since we have $r_n \sim \mathsf{L}$, this gives roughly $\langle |V_{\mu\nu}|^2 \rangle \sim (\lambda_{\rm B}/\mathsf{L})^{d-1}$, the inverse number of wavelength-sized patches on the boundary.

Quasi-orthogonality (and the resulting success of the VS method) is determined by very small elements close to the diagonal. Here we have

$$\langle |V_{\mu\nu}|^2 \rangle \approx \frac{a}{k^{d-1}} \kappa^\gamma \qquad \text{for } \kappa \ll 1/\mathsf{L}, \tag{6.15}$$

where we seek the prefactor a, which will depend only on the billiard shape. The prefactor is proportional to the fluctuations intensity $\nu_{\rm E}^{(\mathcal{G})} = \tilde{C}_{\mathcal{G}}(\omega \to 0)$ for the signal $\mathcal{G}(t) = -\frac{1}{2}m\mathbf{r}^2(t)$ (see Section 3.3.1). This signal has a variance of $\sigma_{\mathcal{G}}^2 \equiv \frac{1}{4}m^2(\langle r^4 \rangle - \langle r^2 \rangle^2)$ about its mean. A rough estimate is $\sigma_{\mathcal{G}}^2 \sim m^2\mathsf{L}^4$. The intensity is the product of the variance and the correlation

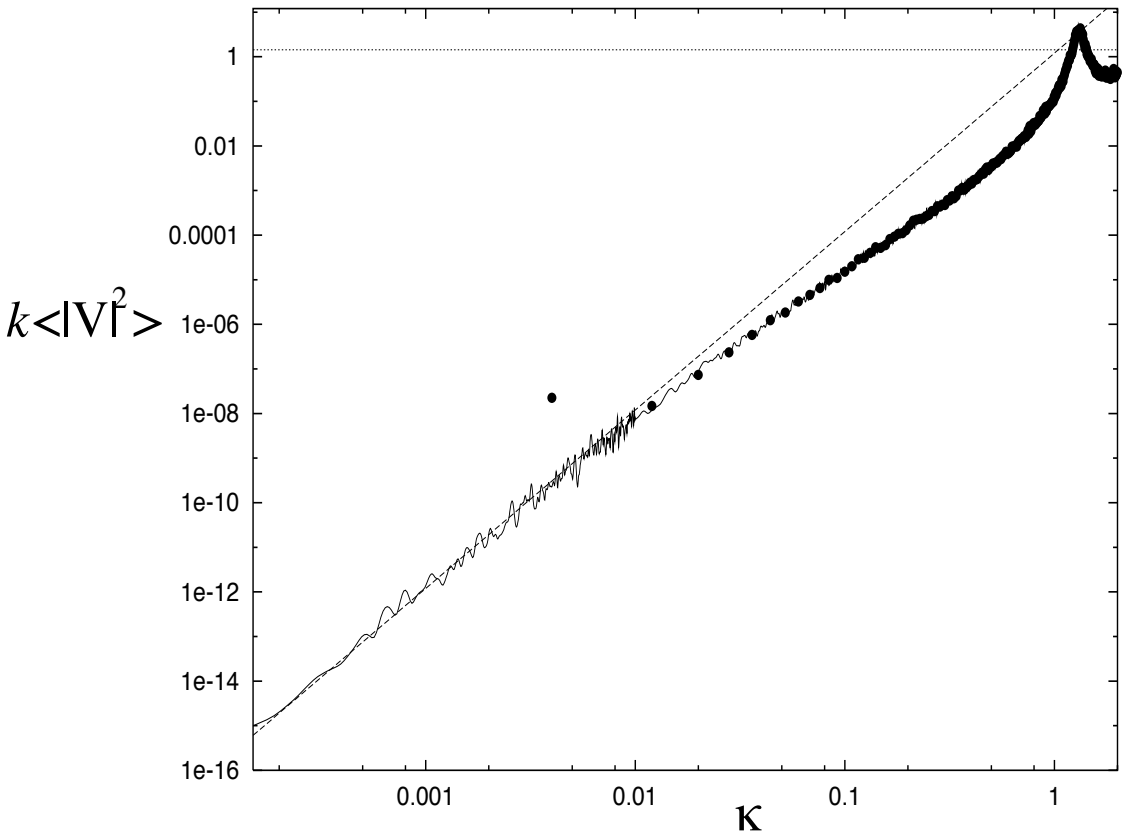

Figure 6.3: Strength of off-diagonal elements V of the quasi-orthogonal matrix M, for the $d = 2$ quarter stadium billiard. The vertical axis $k^{d-1}\langle|V_{\mu\nu}|^2\rangle$ is a quantity dependent only on the billiard system (not on k). The horizontal axis is the wavenumber difference $\kappa \equiv k_\mu - k_\nu$. The solid line is the semiclassical estimate (taken from the classical band profile). The estimation error changes at $\kappa = 0.01$, allowing a dynamic range of 16 orders of magnitude. The quantum band profile (taken from the matrix of Fig. 6.2) is shown as dots, and it becomes unreliable at the level 10^{-8}, for this choice of k. The dashed line is the fit to (6.15); the dotted is the random wave result (6.14). The two lines meet at a κ around the inverse system size $1/\mathsf{L} \sim 1$.

time: $\nu_{\rm E}^{(\mathcal{G})} \sim m^2 \mathsf{L}^4 \tau_{\rm bl}$. Therefore $\tilde{C}_{\rm E}(\omega) = \omega^4 \tilde{C}_{\mathcal{G}}(\omega) \sim (\kappa \mathsf{L})^4 \nu_{\rm E}^{\rm WNA}$. Thus the prefactor is $a \sim \mathsf{L}^{5-d}$. The accuracy is limited because we lack knowledge of the correlation time of the signal $\mathcal{G}(t)$. In the 2D quarter stadium billiard (of size $\mathsf{L} \sim 1$, see Fig. 2.6) the prefactor is measured to be $a \approx 1.2$, remarkably close to unity (see Fig. 6.3). It seems that convergence to the fourth-power law is very slow in this system, not becoming accurate until $\kappa < 0.01$. This is probably due to long tails in the time correlations (Section 3.4) which makes $\tilde{C}_{\mathcal{G}}(\omega)$ close to singular at $\omega = 0$. More study is needed of this convergence to the asymptotic form. In systems without these non-generic effects, we expect more rapid convergence to the fourth-power law (for instance Fig. 3.6), and therefore better quasi-orthogonality.

Both the above estimates are for $|V_{\mu\nu}|^2$ averaged over an energy range as in Section 2.2.2, with κ held fixed. The Random Matrix Theory assumption is that the elements $V_{\mu\nu}$ are independent normally-distributed random quantities. However, deviations from strong chaos (due to periodic orbits and scars; see Section 2.3.3) will affect the distribution and correlation of the elements.

Thus we have estimated a wave overlap quantity semiclassically. In both the above expressions all the dependence on 'particle quantities' ($\hbar$, m, $v_{\rm E}$) has cancelled, which is necessary since the particle picture is merely a construction allowing semiclassical estimates. (For instance in an acoustic Helmholtz application the particle model is completely fictitious). Rather, everything is in terms of lengths, as expected since the matrix elements of V arise purely from wave considerations.

Relation to strength estimate of Vergini and Saraceno

I have not yet been able to find a direct wave-mechanical explanation (which would bypass the above semiclassical estimate) for the quasi-orthogonality of dilation. However a clue is given by Berry and Wilkinson's proof [28] that degenerate ($k_\mu = k_\nu$) off-diagonal elements are exactly zero. An explanation for $k_\mu \neq k_\nu$ was first attempted by VS [195, 194]. They tried to establish quasi-orthogonality using the identity $V_{\mu\nu} = [(k_\nu^2 - k_\mu^2)/2k^2] B_{\mu\nu}$, with $B_{\mu\nu} \equiv \langle \psi_\mu | \, \mathbf{r} \cdot \nabla \, | \psi_\nu \rangle_{\mathcal{D}}$. This identity can be proved in a simple fashion (Eq.(H.25)). However, they then made the assumption $|B_{\mu\nu}| \sim O(1)$ (for $d = 2$) by claiming that eigenstates are uncorrelated across the volume of $\mathcal{D}$ [194]. (In $d = 2$ this would give $\sim (\mathsf{L}/\lambda_{\rm B})^2$ variables whose variances add linearly). This would give a power law $\gamma = 2$ for the off-diagonal growth of $|V_{\mu\nu}|^2$.

This is in error for two reasons. Firstly, we know that the number of degrees of freedom in a constant-wavenumber function is actually $N_{sc} \sim (\mathsf{L}/\lambda_{\rm B})^{d-1}$ scaling like the boundary (see Section 5.3.1), which would imply the naive random-wave estimate $|B_{\mu\nu}| \sim O(\mathsf{L}/\lambda_{\rm B})^{(3-d)/2}$. Secondly, a random wave estimate of an overlap over a large fraction of the region is generally very bad (Section 3.3.2). In fact, comparison to our band profile results (analytical and numerical) shows that $|B_{\mu\nu}| \sim |k_\mu - k_\nu|$, a result which cannot be guessed by random wave assumptions. This result has also since been verified by Vergini [1].

The conclusion is that the original authors' quasi-orthogonality estimate (which they have used in [195, 194, 175, 196, 203]) was in fact pessimistic: the true power law $\gamma = 4$ actually gives *much* smaller elements close to the diagonal.

[1] Personal communication.

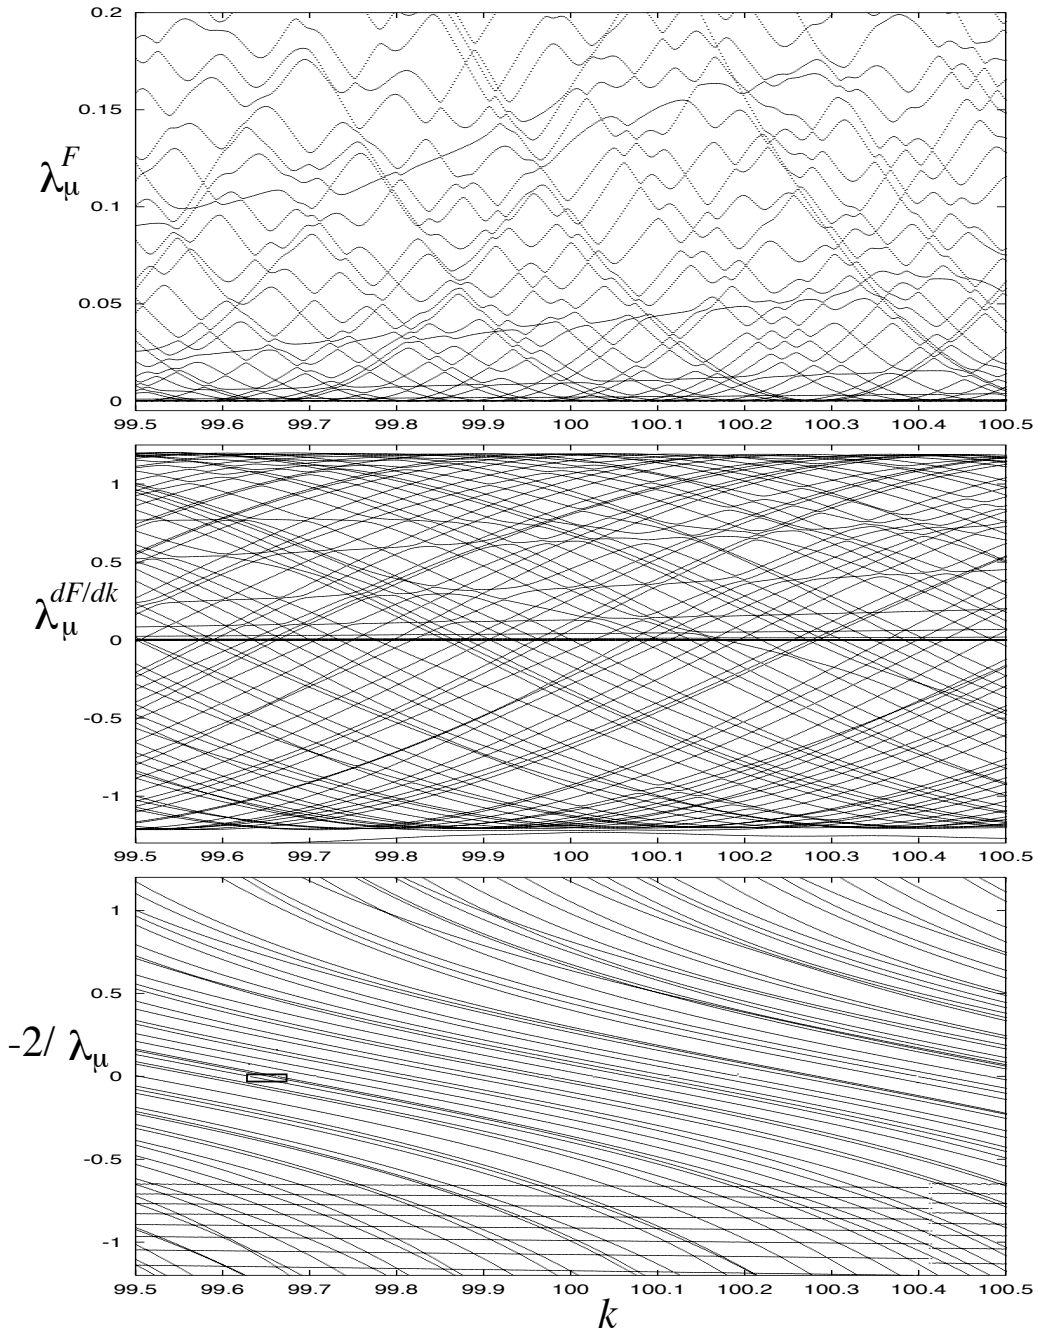

Figure 6.4: Eigenvalues as a function of k. **Top:** tension matrix F. **Middle:** derivative of tension matrix dF/dk. **Bottom:** the generalized eigenvalues of (6.25), plotted as the lowest-order predicted wavenumber shifts $-\delta_\mu = -2/\lambda_\mu$. The small rectangle indicates the enlargement (see Fig. 6.12). In the top two plots, the null-space of the matrix is visible as undesirable interfering eigenvalues climbing slowly upwards from zero. The bottom plot shows the remarkable result that (6.25) is free from such null-space solutions over a wide range. The quarter stadium with a basis of 120 symmetrized real plane waves was used.

6.1.3 Representation in a Helmholtz basis

Linear function space at fixed wavenumber

There is a linear space of functions which are Helmholtz solutions at wavenumber k, which have arbitrary boundary conditions on Γ, and which are normalizable in $\mathcal{D}$. As we saw in Section 5.3.1, the semiclassical size of this space is $N_{sc} \sim (k\mathsf{L})^{d-1}$, corresponding to waves which propagate across the entire billiard, for instance real plane waves (RPWs)—plane waves with real wavevector components. One possible represention of this space is by an angular distribution function $f(\Omega)$, namely

$$\psi(k;\mathbf{r}) \;=\; \int d\Omega\, \psi(\Omega) e^{ik\mathbf{n}_\Omega\cdot\mathbf{r}}, \tag{6.16}$$

where Ω is the angle coordinate in d dimensions and $\mathbf{n}_\Omega$ the corresponding unit normal vector. Propagating waves correspond to functions $\psi(\Omega)$ which oscillate on an Ω-scale $\sim (k\mathsf{L})^{-1}$ but no faster. Another possible representation is using surface 'charge' (Green's function monopole or directed-dipole sources) distributed across Γ—this is the basis set of the Boundary Integral Method (BIM) [36, 121]. Again, the spatial frequency of this surface charge is limited (to k) for propagating waves.

However there are also infinitely many evanescent waves which cling to the surface without penetrating far into the volume [2] [194]. The higher the imaginary part of the wavenumber (the more evanescent), the quicker is the decay into the interior. Evanescent plane waves (EPWs) fall into this category. EPWs *can* be represented by distributions of RPWs in the finite region $\mathcal{D}$, however the required coefficients grow exponentially with the evanescence parameter, and are of rapidly alternating sign [26]. This means that evanescent (surface) waves can be represented by angular functions $\psi(\Omega)$ which oscillate more rapidly than $\sim (k\mathsf{L})^{-1}$ and are exponentially large. (The corresponding surface charge representation oscillates more rapidly than k, however it has the advantage that the charge is *not* exponentially large).

Therefore if all such evanescent waves are included, the dimensionality of the linear space seems infinite. However this issue is a tricky one because of the exponential divergences, and has only been explored quite recently (see also [64, 63]). Practically, the inclusion of some EPWs into a representation of the function space enables the coefficients to avoid exponential blow-up in their closest representation of eigenfunctions (see Appendix J). Therefore in practice, the dimensionality required is not much larger (a factor 1.1 to 1.2 larger in the 2D stadium) than N_{sc}.

The use of EPWs is very beneficial in many systems, particularly the stadium, where the tension minima ϵ_μ can be lowered from the range 10^{-5} to 10^{-7} to the range 10^{-10} to 10^{-12} (see Fig. 6.6), a dramatic discovery made by Vergini[194]. Therefore for all high-accuracy computations for the stadium a basis composed of RPWs plus a few EPWs is used. The particular EPWs used [194] are described in Appendix J. It seems that most eigenfunctions of nonintegrable shapes and even some integrable ones have evanescent character [63, 194]. The problem of finding suitable evanescent waves to add to the basis for a general billiard shape is unsolved. There are many shapes for which no good basis is known (Sinai billiard, generalised Sinai billiard of Fig. 3.2a, etc.).

[2] For an example of appearance of surface waves in an example problem in linear spaces see [77].

Use of a scaling basis

Because the scaling eigenfunctions $\psi_\mu(k,\mathbf{r})$ defined above all have the same k, and are normalizable, they fall into the above linear space. They can be represented by coefficient vectors $\mathbf{x}^{(\mu)}$ of a Helmholtz basis set with a *single* wavenumber k,

$$\psi_\mu(k,\mathbf{r}) \;=\; \sum_{i=1}^{N} x_i^{(\mu)}\phi_i(k;\mathbf{r}) + \varepsilon_\mu(k;\mathbf{r}), \tag{6.17}$$

where ε_μ is some small error function for each state. (The error function is a Helmholtz solution and is orthogonal to the basis set space). The functions ϕ_i will not be orthogonal over $\mathcal{D}$, and there will be some (numerical) null-space of vectors $\mathbf{x}$ that has negligible effect on wavefunctions inside $\mathcal{D}$ (see Section 5.3.1).

The boundary tension $f(k)$ is a quadratic form in the linear space. Applying this form to the error $\varepsilon_\mu(k;\mathbf{k})$ gives the achievable tension minimum ϵ_μ for each state. In the N-dimensional above basis, the quadratic form is written $\mathbf{x}^{\mathrm{T}}F(k)\mathbf{x}$ where F is the tension matrix.

In the methods of the previous chapter the only requirement on the k-dependence of this basis set was that at a given k all the functions are Helmholtz solutions at that k. In this chapter we must specialize to a *scaling basis*,

$$\phi_i(k;\mathbf{r}) \;=\; \phi_i(k,\mathbf{r}) \;\equiv\; \bar{\phi}_i(k\mathbf{r}) \qquad \forall i. \tag{6.18}$$

Note that $(k,\mathbf{r})$ implies a scaling function whereas $(k;\mathbf{r})$ implies some more general k-dependence. The expansion of any scaling function, as in (6.17), now has a coefficient vector $\mathbf{x}$ which is constant as k changes. This will be necessary for our particular case of the scaling eigenfunctions. Note that certain basis sets, most notably that of the BIM, *cannot* be used because their basis functions (Green's functions) have origins at *different* points 'pinned' to the boundary Γ.

6.1.4 Solving for the scaling eigenfunctions

We have a very-nearly-diagonal form for the tension matrix for wavenumbers which are close to k, in a scaling eigenfunction basis (Section 6.1.2). We consider this matrix and its derivative with respect to k,

$$\tilde{F}_{\mu\nu}(k) \;=\; 2\delta_\mu\delta_\nu M_{\mu\nu} \;+\; O(\delta^3)\cdots \tag{6.19}$$

$$\frac{d\tilde{F}_{\mu\nu}}{dk} \;=\; 2(\delta_\mu+\delta_\nu)M_{\mu\nu} \;+\; O(\delta^2)\cdots, \tag{6.20}$$

where the tilde indicates the scaling eigenfunction representation. Use has been made of $d\delta_\mu/dk = 1$. (I will not write the k-dependence of the δ's explicitly). The quasi-diagonality of both these matrices results from that of M, whereas the k-dependence results from the δ's.

The transformation between matrix representations is

$$\tilde{F}(k) \;=\; X^{\mathrm{T}}F(k)X \tag{6.21}$$

$$\frac{d\tilde{F}}{dk} \;=\; X^{\mathrm{T}}\frac{dF}{dk}X, \tag{6.22}$$

where $X_{i\mu} \equiv x_i^{(\mu)}$ are the desired coefficient vectors of the scaling eigenfunctions (the error functions in (6.17) will be dropped). X is rectangular with some number of columns less than N; the small-δ approximation breaks down before this number is reached anyway. The scaling basis representations

$$F_{ij}(k) \;=\; \oint_\Gamma d\mathbf{s}\, \frac{1}{r_n}\, \phi_i(k,\mathbf{r})\phi_j(k,\mathbf{r}) \tag{6.23}$$

$$\frac{dF_{ij}}{dk} \;=\; \frac{1}{k}\oint_\Gamma d\mathbf{s}\, \frac{1}{r_n}\, \phi_i(k,\mathbf{r})\, \mathbf{r}\cdot\nabla\phi_j(k,\mathbf{r}) + \text{transpose} \tag{6.24}$$

can easily be evaluated using the method of Appendix G.

Then diagonalization of the matrix $F(k)$ can give eigenvectors which are very close to the desired rows of X. In particular, an eigenvector very close to the scaling eigenfunction $\mathbf{x}_\mu$ may be returned (if $\delta_\mu \ll 1/\mathsf{L}$). However the eigenvalue will equal the tension of the state resulting from a unit norm $|\mathbf{x}| = 1$ in coefficient space, which has no physical significance in the basis sets used (RPWs + EPWs). Also the null-space vectors in $\mathbf{x}$ produce small-eigenvalue solutions (exponentially spread down to machine precision, as in Fig. 5.3) which interfere (mix) with the desired vectors. However the parabolic tension minima are still visible. The same is true if the matrix dF/dk is diagonalized, only now small eigenvalues correspond to both Dirichlet (upwards-travelling) and Neumann (downwards-travelling) eigenstates. Therefore direct diagonalization is not a good way to extract the desired states. Fig. 6.4 shows the k-dependence of these two diagonalizations.

Generalized eigenproblem

The problem is solved by a forming a generalized eigenequation,

$$\left(\frac{dF}{dk} - \lambda_\mu F\right)\mathbf{x}^{(\mu)} \;=\; \mathbf{0}, \tag{6.25}$$

which finds the eigenstates of the k-derivative of tension, treating the tension quadratic form F as a norm which is held constant. The presence of a norm based on the wavefunction at Γ means that null-space vectors can be excluded by the method of Section 5.3.2 with truncation ϵ set at about 10^{-14}. It is possible to replace this tension norm F by other norms, for instance the exact or Dirichlet G from the previous chapter, however they will not be quasi-diagonal in the scaling eigenfunction basis. This destroys much of the benefit of the quasi-diagonality of dF/dk in this basis. The power of (6.25), as realised by VS, is that *both* matrices are quasi-diagonal in this basis. Therefore, simultaneous diagonalization (that is, solving the generalized eigenproblem) of dF/dk and F in the computational basis ϕ_i returns a very good approximation to the transformation X into the desired eigenfunction basis.

From (6.20) to lowest order the diagonal elements of $\tilde{F}(k)$ are $2\delta_\mu^2$ and those of $d\tilde{F}/dk$ are $4\delta_\mu$. The ratio gives the generalized eigenvalue λ_μ. The prediction for the eigenwavenumber is $k_\mu = k - \delta_\mu$, giving to lowest order (see Section 6.2 for higher orders),

$$k_\mu \;=\; k - 2\lambda_\mu^{-1} + O(\lambda_\mu^{-2})\cdots. \tag{6.26}$$

The method will compute *all* the scaling eigenfunctions within a wavenumber range of up to about 1, for a system size $\mathsf{L} \approx 1$. This corresponds to whose rescaled boundaries lie within

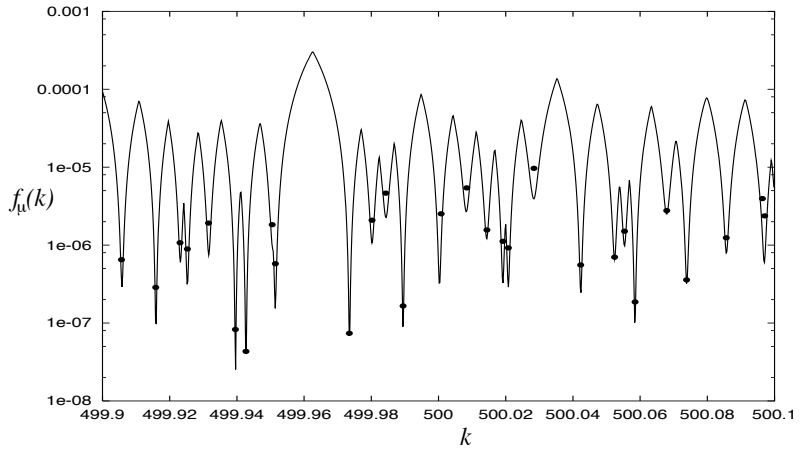

Figure 6.5: Comparison of eigenvalue solutions returned by a single diagonalization of (6.25) (shown by points, at their resulting k_μ and tensions ϵ_μ) against those obtained by the sweep method of the previous chapter (tension shown by line). The agreement is excellent (well within the errors of the sweep minima), apart from the state at $k = 500.0003$ which suffers from error discussed in Section 6.3.2. Close eigenvalues not distiguished by the sweep method *are* found by the scaling method, even though the same basis set (500 symmetrized RPWs in the stadium) was used. The tensions t_μ from the scaling method are a factor of 2–3 larger than the sweep method minima ϵ_μ.

about 1/12 wavelength of the original boundary. The predicted $-\delta_\mu$ as a function of k are shown in Fig. 6.4 (bottom).

The scaling method in action

In Figs. 6.5 and 6.6 I show that the k_μ (and tensions t_μ) found coincide with tension minima found by the sweep method of the previous chapter. The k agreement is well within the sweep method error due to the finite rounding of the minima. Because the scaling method does not just compute the tension of a single state, it is able to distinguish close eigenvalues that cannot be distinguished using the sweep method (using the same basis set). How far this can be pushed has not yet been investigated. The finite tension minima are highly correlated between the methods, and are therefore a function of the limitation of the RPW basis set chosen.

The remarkable improvement in efficiency over the sweep methods of the previous chapter (or any other known method!) comes from the fact that $\sim (k\mathsf{L})^{d-1}$ billiard eigenstates are computed using a single diagonalization. The effort per state therefore scales like $\sim k^{2(d-1)}$ compared to $\sim k^{3(d-1)}$ for sweep methods. In $d = 2$ this gives an improvement in efficiency of $\sim 10^3$ at $k \sim 10^3$ if only moderate accuracy is required. More specifically, the basis set size is $N \approx N_{sc}$, and about $N/10$ usable states are returned per diagonalization (moderate accuracy, $|t_\mu| < 10^{-5}$), or about $N/50$ (very high accuracy, $|t_\mu| < 10^{-10}$). This assumes that the basis set is adequate for the particular billiard shape. Errors will be further analysed in Section 6.3.

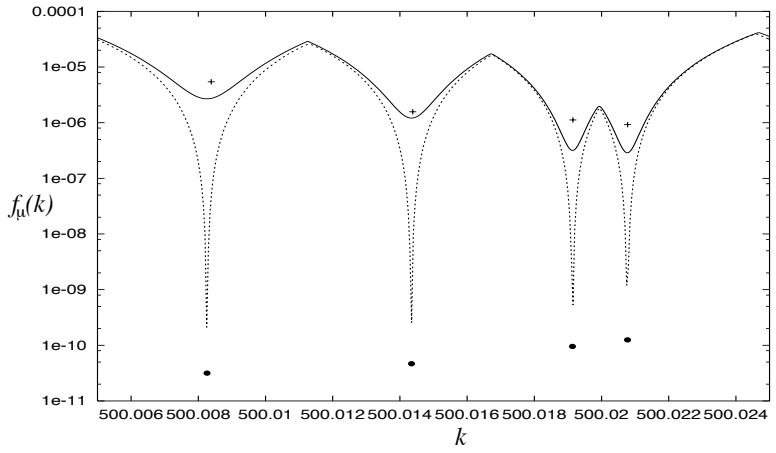

Figure 6.6: Zoom in on Fig. 6.5, with solid line and crosses showing the sweep and scaling methods respectively (basis of 500 RPWs). The dotted line and points show the same with an improved basis (500 RPWs and 30 EPWs). The scaling method returns k_μ that are well within the 'tension rounding error' width $\Delta k(\epsilon_\mu)$ from the sweep method. The evanescent wave improvement is dramatic, and allows the scaling k_μ to reach accuracy $\sim 10^{-6}$ for the $N/50$ most accurate states.

Now I give a practical example. A very high energy state of the stadium billiard (odd-odd symmetry class) is shown in Fig. 6.7, with $k = 1000.00275996$, and tension 8×10^{-12}. The error in k is about ± 1 in the last digit. At this k, corresponding to about the 142000'th odd-odd eigenstate, the basis size is $N = 950$. It takes about 10 seconds to fill F and dF/dk, another 50 seconds for diagonalization, returning about 16–100 usable states (depending on required tension accuracy), and about 2 seconds per state to evaluate the wavefunction and its normal gradient on the perimeter (2000 sample points around the quarter-perimeter of the stadium). It takes a further 300 seconds to generate spatial plots of the type of Fig. 6.7 with resolution 1000*500 samples. Because the plotting algorithm generates plots of the scaling eigenfunctions simultaneously, it does not take much longer to evaluate more than one such plot (in each case $\sim 10^9$ sin or cos evaluations are required). This was performed on a Compaq XP1000 workstation, with alpha processor 21264, 667MHz, using C++ and LAPACK [5]. It is clear that obtaining spatial plots is the limiting factor above (due to the number of basis function evaluations required), rather than solving for the eigenstate k_μ and coefficients themselves!

6.2 Higher-order correction and normalization

In the brief original paper of VS [195], the hint that the authors used higher-order expansions of the tension $f(k)$ as a function of k is given. Here I present what I have found to be the most practical implementation of such corrections. Some of this has been developed through conversations with Vergini. 'Automatic' normalization was already presented in Vergini's thesis [194]; the only new result I have added is the observation of the power-law

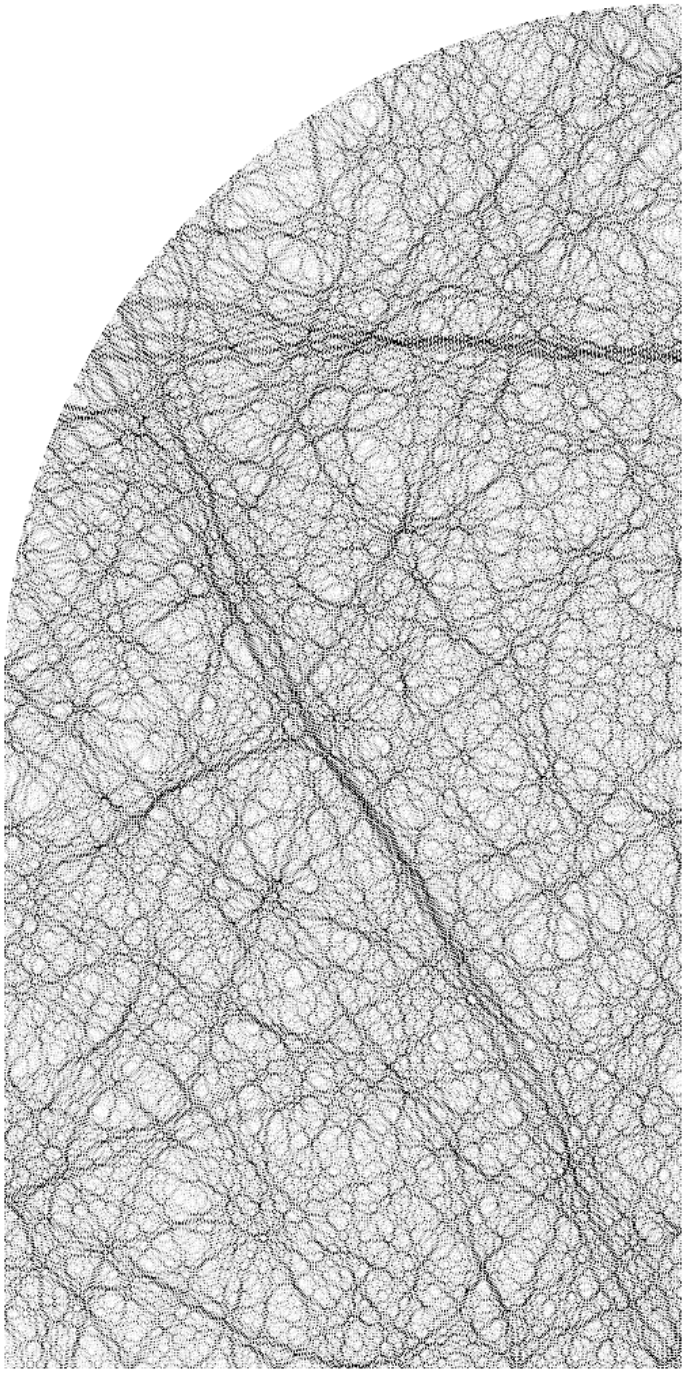

Figure 6.7: Odd-odd symmetry Dirichlet eigenstate of the 2D stadium billiard at $k = 1000.00275996$, shown as a probability density plot. Only the quarter-stadium is shown (spanning about 320 wavelengths). Scarring by the 'bow-tie' orbit is visible.

norm error growth.

6.2.1 Correction of the eigenwavenumbers

In Fig. 6.4 (bottom) the consistent curvature in the predicted shifts δ implies that higher-order effects could be important in predicting the k_μ values. There are many levels of higher-order correction possible once the basic diagonalization of (6.25) has been performed: 1) identification of the higher-order diagonal terms in $\tilde{F}$ and $d\tilde{F}/dk$ for an average state, then correcting all the k_μ according to the same resulting formula, 2) using $\partial_n\psi_\mu$ for each state on the boundary to compute a correction for each corresponding k_μ, 3) using $\partial_n\psi_\mu$ for each state on the boundary and knowledge of off-diagonal terms in $\tilde{F}$ and $d\tilde{F}/dk$ to compute corrections to the eigenvectors $\mathbf{x}_\mu$ as well as the k_μ. I will only report on the first of these ideas, since it is easiest to routinely implement. The other two levels of correction are areas for investigation, however it is not known whether they will prove even worthwhile (there is a trade-off betweeen correcting the results of one diagonalization and simply performing more diagonalizations). My inkling is that in practice the corrections presented here will be the only worthwhile ones.

Appendix I shows that the expansion of the diagonal of the tension matrix is

$$\tilde{F}_{\mu\mu} \;=\; 2\delta_\mu^2 - 2\frac{\delta_\mu^3}{k_\mu} + C_4\delta_\mu^4 + O(\frac{\delta_\mu^5 \mathsf{L}^2}{k_\mu}) + O(\delta^6 \mathsf{L}^4)\cdots \tag{6.27}$$

where the coefficient C_4 is the sum of many complicated boundary integrals of $\partial_n\psi_\mu$, $\partial_{nt}\psi_\mu$, etc, and the local curvature α and its derivative. It is possible in theory to perform a random-wave estimate for these integrals to reach an expression in terms of the billiard shape. However in practice it is much easier to fit for C_4 by measuring the average value for a few states. For instance, in the quarter stadium of length 2 by height 1, $C_4 \approx -1.3$. However, it varies from state to state [3], taking smaller values for states with higher average transverse momentum on the boundary, or for states which are scarred so as to avoid regions of high r_n. Note that C_4 has units of length squared, so that if the billiard changes size, then C_4 will need to be scaled accordingly.

The generalized eigenvalue is $\lambda_\mu = (d\tilde{F}_{\mu\mu}/dk)/\tilde{F}_{\mu\mu}$ which can be found from (6.27). Inversion of the series expansion is needed to get the series for our estimate for δ_μ, giving

$$k_\mu \;\equiv\; k - \delta_\mu \;=\; k - \frac{2}{\lambda_\mu} + \frac{2}{k\lambda_\mu^2} - \frac{4C_4}{\lambda_\mu^3} + O(\frac{1}{\lambda_\mu^4})\cdots, \tag{6.28}$$

where corrections of $O(\delta_\mu/k)$ have been ignored. This is our improved version of (6.26). In tests, use of this formula had little effect on the $N/100$ most accurate states, but the resulting wavenumber shift reduced the tension of the remainder of the useful states by a factor of ≈ 8 (see Fig. 6.9). Also, the $O(\lambda_\mu^{-2})$ term had very little effect (much less than the $O(\lambda_\mu^{-3})$) because of the $1/k$ factor. The utility of higher-order expansion is minimal (the next significant, the sixth-order, has little effect).

[3] as also found by E. Vergini, personal communication.

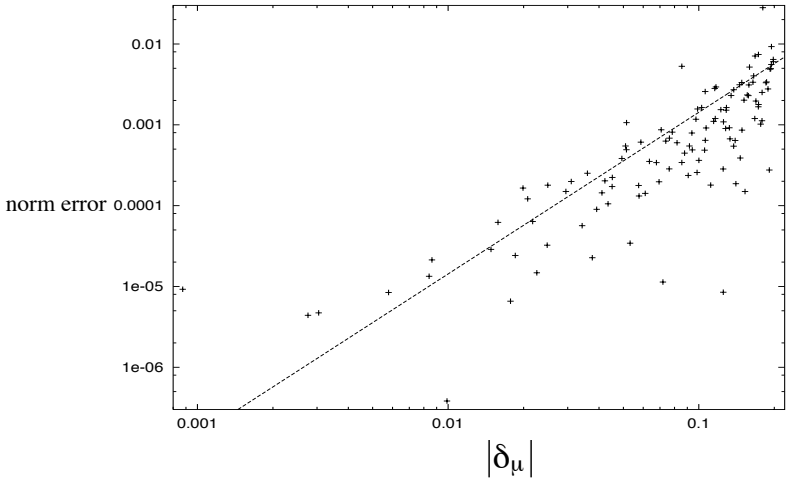

Figure 6.8: Automatic normalization of states returned by the scaling method at $k \approx 10^3$, using (6.29). The error from 1 is shown. The growth of random norm errors roughly follows a second power law (shown by the staight line).

6.2.2 'Automatic' normalisation of states

A feature of the scaling method, as VS realised, is that the eigenstates found are already normalized in a known fashion. (This contrasts the BIM, Heller's PWDM for which explicit normalization is required). The numerical diagonalization of (6.25) usually returns eigenvectors normalized so $\mathbf{x}^{\mathrm{T}} F \mathbf{x} = 1$, corresponding to scaling eigenfunctions with unit tension. However, we already know the tension of such functions when they are normalized to unity in the domain $\mathcal{D}$: this is given by (6.27). Therefore by taking the square root we obtain the amplitude correction factor,

$$\mathbf{x}_\mu \longleftarrow \left(\tilde{F}_{\mu\mu}\right)^{1/2} \mathbf{x}_\mu. \tag{6.29}$$

Of course, with (H.9) and the method of Appendix G we are armed with a rapid tool for checking the normalization of states. This is done in Fig. 6.8, showing that for high k, the normalization is correct to 1% for the $N/10$ useful states, and correct to 0.01% for the $N/50$ highest accuracy ones. The growth of norm errors is random from state to state, and follows a second power law. In practice, since the boundary derivatives of the eigenstates are needed anyway, a final normalization using (H.9) was performed. If the norm deviates much from 1, it is a very useful indicator that something is wrong (*e.g.* a spurious state has been found). This is probably the single most important use of automatic normalisation. Such errors are now analysed in the next section.

6.3 Sources of error in the method

There are four distinct causes of errors in the scaling method: 1) those due to the basis set limitations, 2) those due to deterioration of the eigenstates as $|\delta_\mu|$ grows, 3)

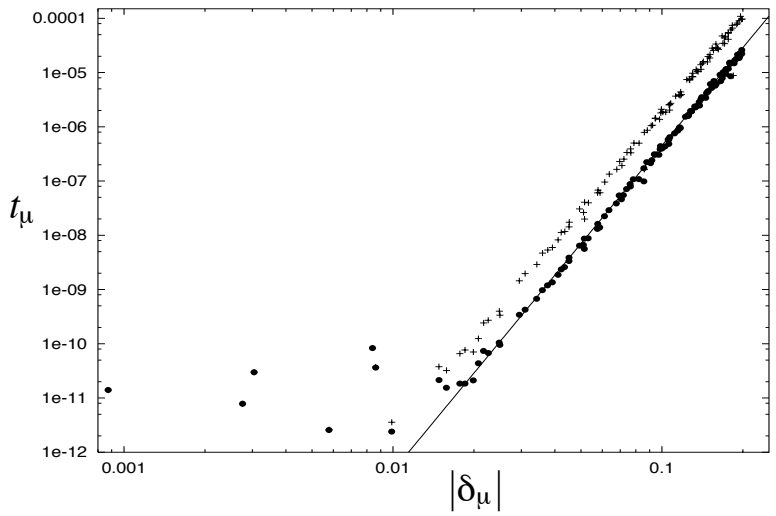

Figure 6.9: Growth of tension (2-norm of error from obeying the boundary conditions) with δ_μ of eigenstates returned from a single scaling diagonalization at $k \approx 10^3$. The basic method using (6.26) (crosses) and corrected wavenumbers (6.28) (dots) are shown. Both show a sixth power law (straight line), with small deviations from state to state. Truncation at $t \sim 10^{-11}$ is due to limitations of the basis set.

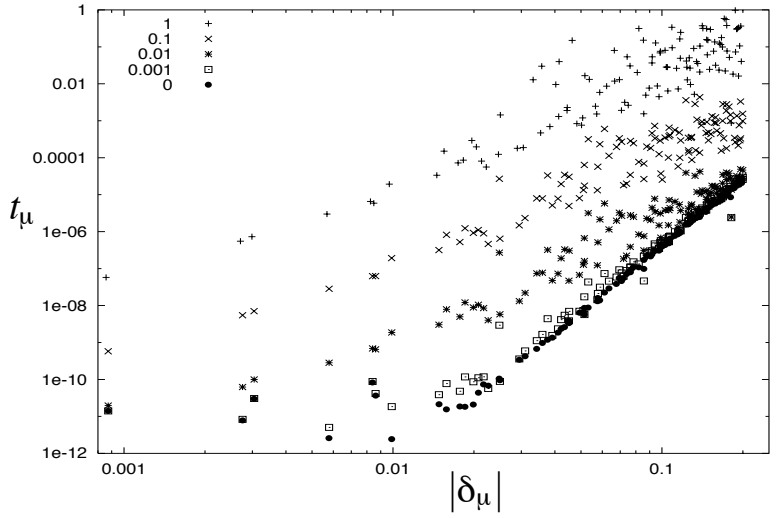

Figure 6.10: Exploration of the effect of weakening the quasi-diagonality of $M_{\mu\nu}$. The dots show the same as Fig. 6.9. The other sets of points show the tension errors for $w(\mathbf{s}) = (1-\beta)r_n^{-1} + \beta r_n^{-2}$, with the different choices of β labelled in the upper left corner. This corresponds to $D(\mathbf{s})$ being the dilation deformation with some 'constant' (CO) deformation mixed in (see Table 3.1). As β increases a new type of error emerges with a different power-law and much greater random state-to-state fluctuations. Examining the ψ boundary errors shows that they are dominated by the effect of the one or two states with smallest $|\delta_\mu|$.

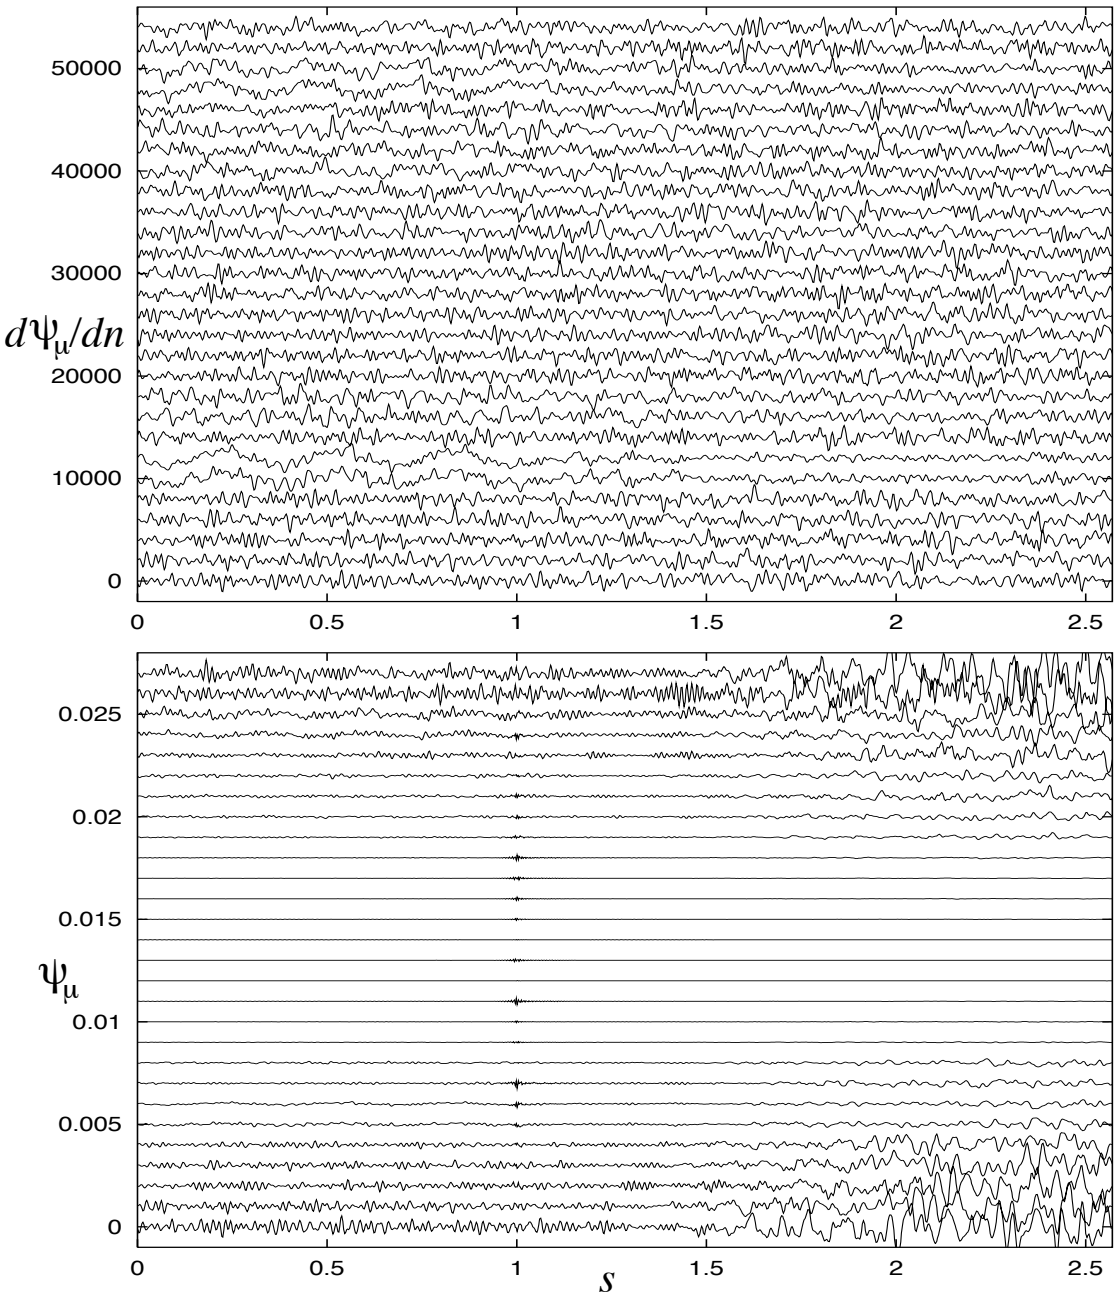

Figure 6.11: Normal gradients $\partial_n \psi_\mu$ (upper) and values ψ_μ (lower) as a function of the boundary coordinate s. The plots have been displaced vertically by state (with increasing k). The states correspond to those shown in Fig. 6.5, except that a better RPW and EPW basis set has been used to reduce basis-related errors. The billiard is the 2D quarter stadium at $k = 500$. The coordinate s is measured from the upper-left to the lower-right 'corner' of the quarter stadium, on the outer edge (which is also present in the full stadium). The small 'blip' at $s = 1$ corresponds to the 'kink' in the stadium boundary.

those due to divergence at very small $|\delta_\mu|$, and 4) those due to the appearance of spurious surface-wave solutions in the 'useful' δ window.

The first of these is already familiar from the previous chapter, and Section 6.3.3. It places a minimum possible tension ϵ_μ on each state—the basis cannot produce a smaller tension with any choice of coefficient vector $\mathbf{x}$. This tension is therefore that of the function $\varepsilon_\mu(\mathbf{s})$ alone (see (6.17)). This is visible as the bottoming-out at $t \sim 10^{-11}$ in Fig. 6.9.

Here and in the future we use tension as the main measure of the error of a found eigenstate, since is gives the 2-norm of the amount by which the boundary conditions fail to be obeyed. I will now discuss the three remaining types of error.

6.3.1 Deterioration of the eigenstates with $|\delta_\mu|$

For a given state μ returned by the generalized eigenproblem, there will be an optimal adjustment to its k_μ that can be made to reduce the tension to a minimum. (Generally this will be close to that returned by the state-independent higher-order correction formula (6.28)). This adjustment has the effect of rescaling the eigenfunction so that its node best approximates the Dirichlet boundary Γ. However the tension minimum t_μ thus achievable is much higher than the intrinsic basis limitation ϵ_μ, for most of the $N/10$ useful states, so is the dominant error in practical applications. Therefore it would be good to understand its origin.

In Fig. 6.9 I show that the growth of tension error is $t_\mu \approx A\delta_\mu^6$ with high accuracy. There is relatively little state-to-state fluctuation, that is, t is not simply a random quantity with growing mean (as would be expected if it were the square of a zero-mean gaussian variable whose variance is given). In the same system at lower k I show the normal gradient, and wavefunction error on the boundary for a set of states, in Fig. 6.11. It is interesting to note that the error arises roughly in proportion to r_n (or some power) on the boundary. Also, the spatial frequency of the error at surprisingly low at large r_n. The constant A seems relatively independent of k (it certainly does not change with any higher power than $k^{1/2}$).

This sixth power law is not affected by changing the basis, or by whether k is near or far from the closest k_μ. The cause, as one would expect, seems to be the non-zero off-diagonal elements of $\tilde{F}$ and $d\tilde{F}/dk$, which causes the eigenvectors $\tilde{\mathbf{x}}_\mu$ (in the scaling eigenfunction representation) to be rotated slightly away from the unit vectors. At first sight it might be thought that the quasi-diagonality of M is to blame, because this appears with the lowest power of δ, namely δ^2. However, Appendix I shows that in fact the 3rd and 4th order terms in δ contain much more off-diagonal strength for $\delta \ll \mathsf{L}^{-1}$. This is evidenced by Fig. 6.10, which shows that by increasing the off-diagonal strength of M by controlled amounts, a *new* type of error results. This shows that the off-diagonal strength of M in the standard case of $w(\mathbf{s}) = 1/r_n$ does *not* seem to cause the sixth power law error growth.

I have attempted to model the effect of the 3rd- and 4th-order in δ off-diagonal terms by applying them in first-order perturbation theory to the generalized eigenproblem (6.25) (in the scaling eigenfunction representation). However, it seems that the perturbing effect of scaling eigenfunctions with distant wavenumbers is divergent, and grows like δ_μ^2 not the observed δ_μ^6. The role of these distant states is dubious because they fall outside

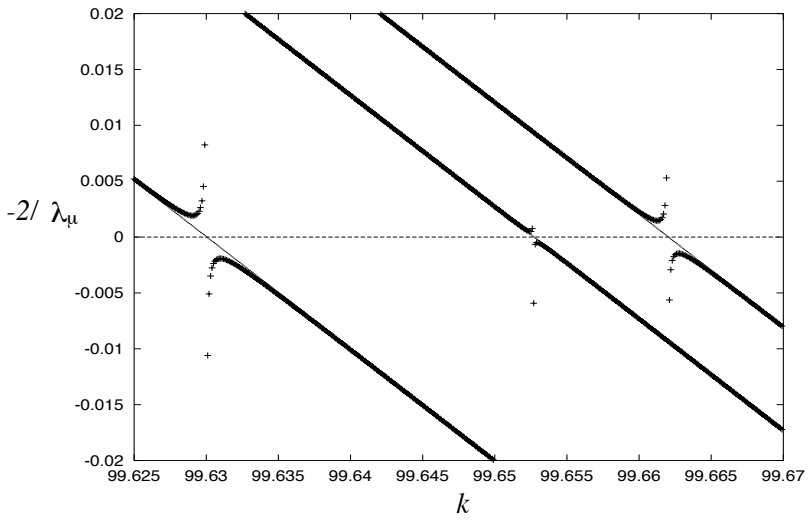

Figure 6.12: Breakdown of the prediction of wavenumber shift $\delta_\mu = -2/\lambda_\mu$ within $\Delta k(\epsilon_\mu)$ of a shift of zero. The crosses are with the generalized problem of (6.25), and the smaller dots using the modified version with F replaced by G as described in the text.

the wavenumber range $O(\mathsf{L}^{-1})$ encompassing N_{sc} adjacent states. It is unknown whether near or distant states are the cause of this error. Clearly this is an area for more study.

6.3.2 Breakdown for a state as $|\delta_\mu|$ vanishes

As observed by Vergini [194], if k falls close to a true eigenstate k_μ, then the scaling method fails to predict this state accurately. Fig. 6.12 shows this happening: the predicted shift heading towards zero suddenly diverges in a first-order pole, giving an inaccurate k_μ prediction. The range of k affected corresponds to the tension rounding error $\Delta k(\epsilon_\mu)$ defined in Section 5.5.1. Therefore the problem does not occur very often, and can be almost (but never quite) eliminated by a good choice of basis. When it does occur it only affects the one state involved. It is very easy to detect because the norm of the state produced automatically (Section 6.2.2) grows to be much larger than 1.

The reason for this type of error is that the tension of the scaling eigenfunction μ has reached the tension minimum allowed by the basis. The diagonal of $\tilde{F}$ behaves like $2\delta_\mu^2$ for a true scaling eigenfunction. However the closest that can be achieved by the basis representation is $2\delta_\mu^2 + \epsilon_\mu$ (see Section 5.5.1). Substitution of this diagonal term into (6.25) and then (6.26) or (6.28) gives the observed pole in the predicted shift.

It is possible to 'cure' the problem completely by replacing F by the true area norm matrix G in (6.25), which has no singularity as δ_μ passes through zero. Then the predicted wavenumber shift is $\delta_\mu = -4\lambda_\mu$. This apparent cure is shown in Fig. 6.12. However, G has no strong quasi-diagonality property when expressed in the scaling eigenfunction basis; the result is a severe loss in the quality of most of the $N/10$ usable states (corresponding to a choice $\beta \approx 0.03$ in Fig. 6.10).

6.3.3 Spurious surface-wave solutions

Imagine the generalized eigenproblem (6.25) represented in the full linear space of Helmholtz solutions normalizable in $\mathcal{D}$ (as described in Section 6.1.3). This eigenproblem finds the Helmholtz solutions with extremal df/dk under the condition that the tension 'norm' f is held constant. The eigenvalue is then $\lambda = (df/dk)/f$. One set of eigensolutions with large λ will be the scaling eigenfunctions with $\delta = -2/\lambda$, as demonstrated in Section 6.1.4. However, highly-evanescent surface waves (Section 6.1.3) form another set of eigensolutions, which we call 'spurious solutions' (they are not relevant to the Dirichlet billiard problem). We cannot find the exact form of these surface waves, but it is easy to show that they exist and to estimate their λ, which I now do[4].

Take an EPW in this function space, with oscillatory wavenumber k_{osc} and evanescent decay constant κ. We have $\kappa^2 = k_{osc}^2 - k^2$. If this function is to have a norm of $O(1)$ in $\mathcal{D}$ then it must be oscillating along the boundary, and decaying along the inward normal. (If this were not true, it would have an exponentially-large norm in $\mathcal{D}$. We ignore the complications that may arise in non-convex shapes). It is therefore a surface wave. The contribution to df/dk at a given point on Γ is $2\psi\, d\psi/dk$, where ψ is the wavefunction value at that point. The contribution to f is just ψ^2. Using $d\psi/dk = (1/k)\mathbf{r}\cdot\nabla\psi$ with $\nabla\psi \approx \mathbf{n}\kappa\psi$ for the EPW, gives $\lambda_{\text{spurious}} \approx 2r_n\kappa/k$ for this wave. Therefore arbitrarily-high λ solutions exist with arbitrarily short decay lengths κ^{-1}. We do not know the exact distribution of ψ on the boundary that forms a solution which is extremal, but it is plausible that they exist. The corresponding wavenumber shift is $\delta_{\text{spurious}} \approx -(1/r_n)(k/\kappa)$. For $k_{osc} > 2k$ we can substitute $\kappa \approx k_{osc}$.

The absence of spurious eigenvalues in a δ range $O(1)$ about zero is clear in the bottom plot of Fig. 6.4. However, spurious solutions (horizontal lines) do exist for $\delta < -0.6$ in this example. Why do the spurious solutions not persist all the way to $\delta = 0$ as the above argument would indicate? The answer is that the *basis set* (in this case RPWs) cannot represent arbitrarily-high k_{osc} EPWs (the required coefficients diverge exponentially [26, 63] so the function rapidly falls into the numerical null-space of F and is truncated away). Therefore a limitation on the maximum k_{osc} representable by the basis (using coefficients $< (\epsilon_{mach})^{-1/2} \sim 10^8$) is what allows a finite window of true scaling eigenfunctions to exist either side of $\delta = 0$. If $\delta_{\text{spurious}} \approx -0.6$ with $r_n \approx 1$ this corresponds to $k_{osc} \approx 2k$.

The 'usable' $N/10$ states returned from the diagonalization fall within a window of about $|\delta| < 0.2$ for the quarter stadium. To keep spurious solutions below this window would therefore suggest limiting $k_{osc} < 5k$. This is in fact very close to the maximum reliable EPW basis state oscillatory wavenumber which I have found to be about $4k$, independent of k.

To conclude, spurious solutions are not a limitation in the basis set. Quite the opposite: they are present in the full normalizable Helmholtz function space, but need to be *excluded* by an appropriate limitation on the k_{osc} representable by the basis set. Even when they do arise, these spurious solutions are easy to detect and ignore, because their 'automatic' normalisation (Section 6.2.2) is much less than one and their tension errors are very large.

[4] This is an elaboration of a suggestion of Vergini (personal communication).

6.4 Application: local density of states at finite deformations

Now I present a calculation of eigenstates in a billiard *deformed* by various finite amounts. Note that the considerations of Chapter 3 only dealt with infinitesimal deformations (rates of change under deformation). We use the same notation as that chapter, namely that the parameter x causes a normal displacement of the billiard wall by $xD(\mathbf{s})$. The matrix of overlaps of eigenstates at deformation x with the undeformed eigenstates will be computed. From this follows an estimate of the *profile* of the matrix, taken by averaging the off-diagonal strength. This is otherwise known as the local density of states (or 'line shape') of the deformation. The physical significance and analysis of this profile is given in [48].

The billiard used was the 2D quarter stadium (Fig. 2.6), chosen because of the excellent RPW and EPW Helmholtz basis set known for this shape (Section 6.3.3, the previous chapter, and Appendix J).

6.4.1 Deformations and eigenstate computation

Examples of a 'special' deformation and a 'generic' deformation were used (see Chapters 3 and 4). The special deformation chosen was rotation about the origin (lower-left corner). This corresponds to $D(s) = \mathbf{n}(s) \cdot \mathbf{e}_z \times \mathbf{r}(s)$ (see Table 3.2). Because the special deformations preserve shape, the eigenstates of a thus-deformed billiard are unchanged (space is invariant under translation and rotation). Therefore only one set of eigenstates is required for all x values.

The generic deformation used was $D(s) = \sin^2(\theta)$, where θ is the angle on the circular part of the boundary measured from the vertical. $D(s) = 0$ on the top, left, and bottom straight sections of the boundary. This has the effect of pushing out the curved part into an ellipse (and a corresponding lengthening of the bottom straight section). This deformation was chosen because it was possible to create large deformations with minimum sacrifice to the quality of the basis-dependent tension minima ϵ_μ.

Generally, it has been found that almost any non-special deformation of the stadium creates a shape for which the RPW basis becomes much worse, and that improvement by addition of EPWs is also limited. For deformations that are localized on the perimeter, this renders the problem unsolvable if x approaches λ_{B} (de Broglie wavelength) or larger. In the above choice of deformation at $k \approx 400$ it was possible to deform from $x = 0$ to $x = 10\lambda_{\mathrm{B}}$ with a corresponding change in typical tension from 3×10^{-11} to $\sim 10^{-7}$. A few states of the deformed system had higher tensions of $\sim 10^{-5}$. This was acceptable for the calculation.

At each x, all the states in the wavenumber range $398 < k < 402$ were gathered using multiple applications of the scaling method at choices of k equally spaced by $2\delta_{\max} = 0.016$ in wavenumber. Thus only states falling within $|\delta_\mu| < \delta_{\max}$ were kept from each scaling diagonalization. This rather small value of $\delta_{\max}$ means that only about 2 states were found per diagonalization. However, because of the limitation of the basis set in such deformed shapes, this gave the highest accuracy. Because the scaling method is so fast, it did not slow down the overall computation by much (the evaluation of eigenstate values and gradients on the perimeter was the main bottleneck).

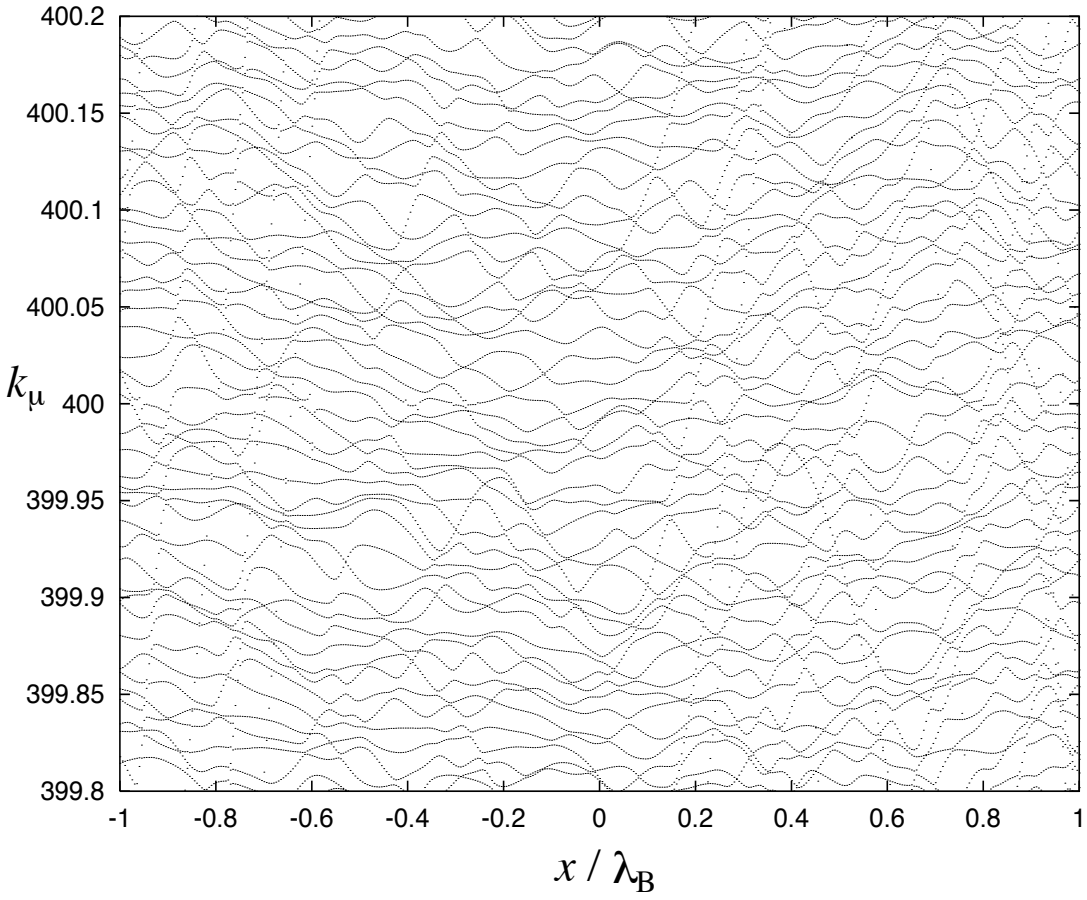

Figure 6.13: Parameter-dependent eigenvalues in the stadium billiard at $k \approx 400$. Avoided crossings are clearly visible. The parameter x controlled a deformation which bent the upper straight section outwards and the curved part inwards. As a result a small integrable pocket of phase-space is created; the resulting fast-moving eigenstates become visible at the extremes of the x range.

To illustrate the effect of deformation on the eigenvalues, Fig. 6.13 shows their parameter-dependence under a (different) generic deformation. This was generated (to lower accuracy) by a single scaling diagonalization at each x value.

6.4.2 Computation of overlap matrix $P(n|m)$ and its profile

The quantity of interest is the matrix of overlaps $P(n|m) \equiv |\langle\psi_n(x)|\psi_m(0)\rangle_{\mathcal{D}'}|^2$. The eigenstates $\psi_n(x)$ are dependent on the deformation parameter x. The overlap is taken in position-space $\mathbf{r}$ over the domain $\mathcal{D}'$, defined as the intersection of the undeformed domain $\mathcal{D}(0)$ and the deformed domain $\mathcal{D}(x)$. We seek a way to perform these d-dimensional domain integrals rapidly (in our case $d = 2$). Appendix H shows that because the eigenvalues $k_n(x)$ and $k_m(0)$ are (in general) different, there exists a very simple formula (H.5) for computing this integral using only the boundary. This will be a huge saving in effort.

Because our eigenstates have Dirichlet boundary conditions, a simplification can be made. Define Γ_1 (Γ_2) as the part of $\Gamma(0)$ ($\Gamma(x)$) inside $\mathcal{D}(x)$ ($\mathcal{D}(0)$). Then $\Gamma_1 + \Gamma_2 = \Gamma'$, where Γ' is the boundary of the intersection region $\mathcal{D}'$. Then the overlap integral over $\mathcal{D}'$ can be written

$$\langle\psi_n(x)\,|\psi_m(0)\,\rangle_{\mathcal{D}'} = \frac{1}{k_n(x)^2 - k_m(0)^2}\left[\oint_{\Gamma_1} d\mathbf{s}\,\psi_n^*(x)\partial_n\psi_m(0) - \oint_{\Gamma_2} d\mathbf{s}\,\psi_m^*(0)\partial_n\psi_n(x)\right]. \quad (6.30)$$

In practice this formula is evaluated using the method of Appendix H, using the wavefunctions evaluated at the set of discretization points from $\Gamma(x)$ and $\Gamma(0)$ which lie in Γ'. The formula becomes inaccurate (depending on the accuracy of the eigenstates) when $|k_n(x) - k_m(0)| \to 0$. In practice this was only a problem at wavenumber differences well below the mean level spacing. The above overlap calculation technique is so efficient that it only took a few seconds to fill each $P(n|m)$ matrix of size about 450 by 450, at $k \approx 400$.

Given the overlap matrix $P(n|m)$ at each x, its average profile was found by the smoothing method of Section C.2. The smoothing width ω_s was chosen to be 0.02 in wavenumber units, which was a couple of times the mean level spacing $\Delta_k \approx 0.0088$.

The resulting sequences of profiles are shown in Fig. 6.14 and 6.15. Comparison with the first-order perturbation theory (FOPT) result,

$$P(n|m) \;=\; \frac{x^2}{(E_n - E_m)^2} \cdot \left|\left(\frac{\partial\mathcal{H}}{\partial x}\right)_{nm}\right|^2 \qquad \text{FOPT}, \quad (6.31)$$

is also made in these plots. $(\partial\mathcal{H}/\partial x)_{nm}$ is explicitly given by (C.2), and its band profile is described in Chapters 2 and 3. The FOPT result simply corresponds to a linearization of the wavefunction $\psi = z \cdot \partial_n\psi + O(z^2)\cdots$ using the notation of Appendix I. The profile shows complete breakdown of FOPT when x reaches $\sim \lambda_{\mathrm{B}}$, and in the case of a generic deformation the formation of a 'core' region before this. For more analysis see our work [48].

6.5 Discussion

I have explained and demonstrated the workings of what is probably the most powerful tool for billiard eigenproblems in existence today: the scaling method of Vergini

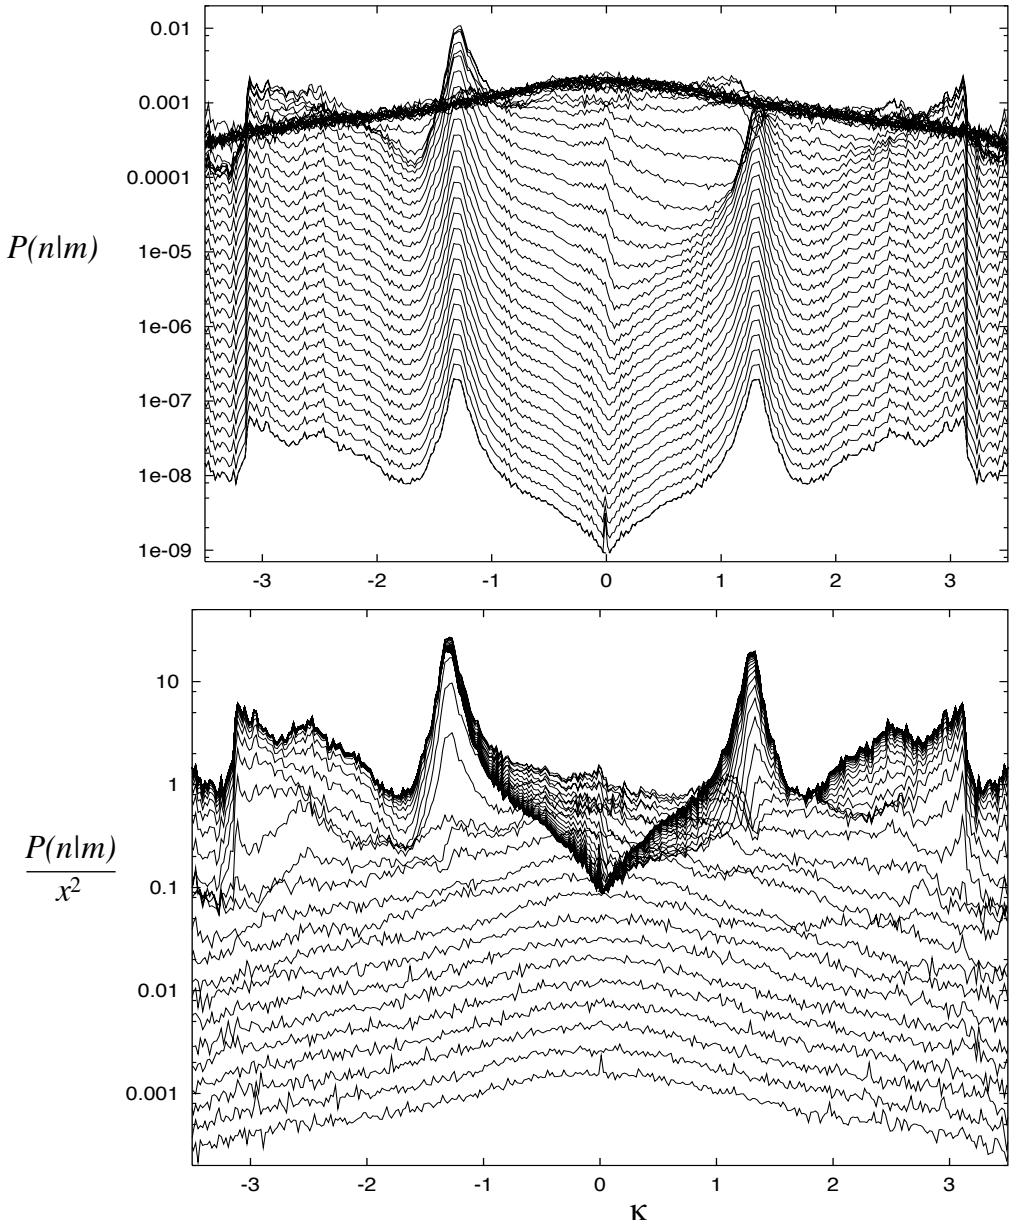

Figure 6.14: Local density of states for finite rotations of the 2D quarter stadium at $k \approx 400$. Profiles of the matrix $P(n|m)$ are shown for 41 choices of x logarithmically spaced from $x/\lambda_{\rm B} = 10^{-3}$ to $x/\lambda_{\rm B} = 10$ (giving the rotation angle in radians). The upper plot shows the profiles, with the FOPT result for the smallest x shown as the bold dashed line. The lower plot shows the profiles divided by x^2, so that the breakdown of FOPT is clear.

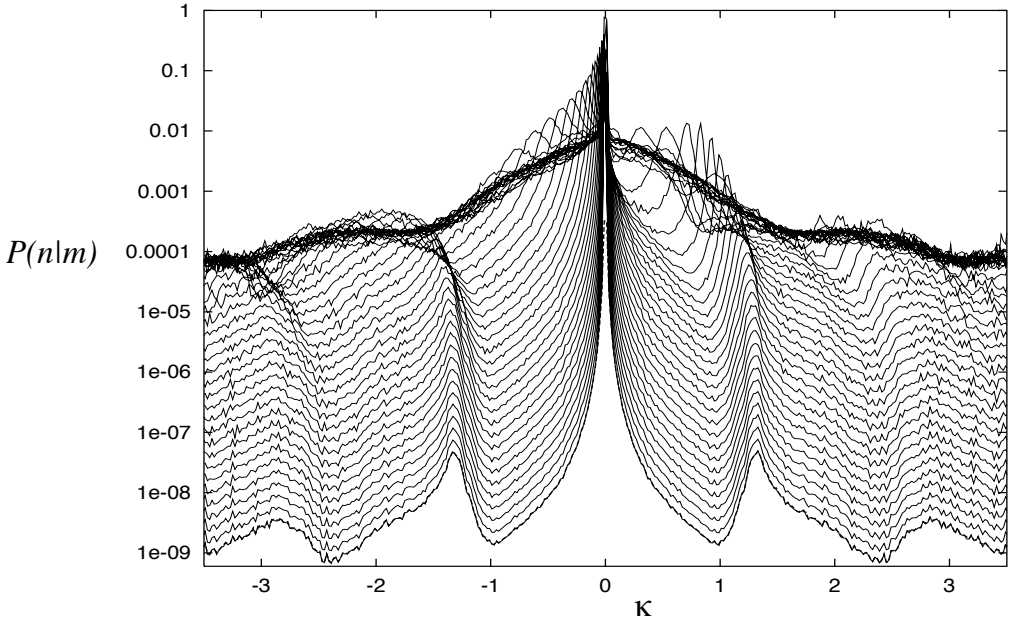

Figure 6.15: Same as the upper plot of Fig. 6.14, except for the generic deformation explained in the text. Drift to the left is due to the lack of volume-preservation of this deformation.

and Saraceno. Certain errors in the original authors' work have been corrected, most notably the understanding of quasi-orthogonality (which has been deepened through a semiclassical estimate), and the derivation of higher-order terms (Appendix I). A comparison between the results of the scaling method and of the sweep method of the previous chapter has also been made. An analysis of the different types of error in the scaling method has been given. Most surprising is the fact that the deviation of M (equivalently the tension matrix) from diagonality seems *not* to be responsible for the dominant error, namely the sixth-power-law growth of tension error with the wavenumber shift δ. The understanding of this type of error is a most pressing area of research in the scaling method.

However there are many more directions to investigate. Work on the deformed stadium shows that much more research is sorely needed in the area of basis set choice, in general billiard shapes. It suggests that Vergini's EPW basis for the stadium is in fact very special. Vergini has noted that for the Sinai billiard no such high-quality basis has yet been found[194]. I have found similarly-disappointing results in the generalised Sinai billiard of Fig. 3.2a (the tension minima barely reach 10^{-3}). Currently investigations are underway of the use of a BIM-type basis set of singular angular-momentum states (the Green's function, or $Y_0(kr)$ Neumann functions or their derivatives).

How could the scaling method be applied elsewhere? The BIM for Helholtz eigenvalue problems is very widely used in the physics, engineering, and acoustics communities [36, 121] (in contrast the PWDM has not been used outside the quantum chaos sub-field to my knowledge). It has an basis set which is 'pinned' to the boundary (the effective

finite basis is due to discretization of the integral equation). Therefore this basis, which has been immensely successful and seems not to suffer from the sensitivity to evanescent waves that a plane-wave basis set does, cannot be used with the scaling method. Either a quasi-orthogonality relation could be found for such a basis, or more likely, a BIM-like basis which *is* a scaling basis could be created for use with the scaling method. Work is in progress on the latter idea. This way, the very expensive 'hunt' for zeros of the BIM determinant could be replaced by a diagonalization returning a large number of states at once. The resulting impact on any research field where Helmholtz eigenstates are needed could be huge.

Other generalizations of the scaling method may also exist. One possibility is that a quasi-orthogonality relation exists for Neumann (or even mixed) boundary conditions. This would allow solution of for example acoustic problems where Neumann BCs are appropriate. Another is the generalization to other *scaling systems*, for instance those given by homogeneous potentials of the form (D.11).

Chapter 7

Conductance of quantum point contacts and half-plane scattering

The quantum point contact (QPC)[20, 65] has played a central role in the understanding of mesoscopic conductance. It is the simplest example of a 2DEG system where the quantum coherent nature of the electron controls the bulk transport properties. The Landauer-Büttiker (LB) formalism [129, 130, 41, 55, 65] reduces the calculation of quantum conductance in the linear response regime to the evaluation of single-particle wavefunction transmission amplitudes. Traditionally, these amplitudes are measured between travelling wave basis states in the 'leads'. Far from the scattering system the leads have constant profiles of finite-width, and support a finite number of transverse modes (channels). Eventually the leads are impedance-matched (that is, without reflection) into 'reservoirs' which act as thermalized sources of electrons at their respective potentials; these potentials are taken to reflect the measured bias voltage (hence the resulting conductance is between points deep within the reservoirs). Such theoretical constructs have been remarkably successful at describing transport phenomena, for instance conductance quantization [193, 20, 65], because the scattering systems involved have generally had good lead-to-reservoir matching.

We consider 'open' 2-terminal mesoscopic systems, namely those where a QPC is *non-adiabatic* (possessing rapid longitudinal variation in transverse profile[1] [20]) and has short or nonexistent leads (for instance if it suddenly abutts onto the 'reservoir' regions), or those where there can be scattering off nearby objects in the 'reservoir' region. We call such systems 'open' because the fully two-dimensional (2D) nature of the 'reservoirs' (*i.e.* the surrounding semi-infinite regions of free space) is important, and therefore they cannot be reduced using the quasi-1D approach described above. This includes a variety of recent mesoscopic experiments, for example the combination of QPCs with nearby resonator structures [112] or with a nearby depletion region underneath an AFM tip [192]. It also includes any QPC system where elastic backscattering from disorder in the reservoirs is significant, or generally where the lead-reservoir matching is bad. In such systems, the

[1] By non-adiabatic, we mean that even at a QPC's narrowest region the transverse profile is changing rapidly. Clearly every QPC becomes 'non-adiabatic' at the coupling to infinite-width reservoirs: this type of non-adiabaticity we do not include because it does not cause significant impedance mismatch, as explained in Ref. [206].

conventional quasi-1D picture does not apply: the scattering system is not coupled to leads in the usual sense, indeed the distinction between leads and reservoirs is no longer clear. (The quasi-1D approach can be retained by attaching very wide leads to such systems, following [189], although this has both numerical and conceptual limitations). The aim of the present work is to introduce a 2D scattering approach which can handle such systems, and to consider some of the consequences for mesoscopic quantum conductance. We will stay within the non-interacting quasiparticle picture, consider zero applied magnetic field, and assume spin degeneracy of 2 throughout.

We imagine a geometry where a 2DEG exists in two semi-infinite half-plane regions, separated by an impenetrable potential barrier which we align with the y-axis (see Fig. 7.1a). Our general 'QPC scattering system' is any gap in this barrier which allows coupling of the wavefunction on the left and right sides. This gap can be defined by an arbitrary form of the elastic potential, and may include other nearby scattering objects, disorder, etc. (which would all be placed within the box shown in Fig. 7.1a). The only important limitation is that this coupling region (the 'system') be of finite y extent, so that electrons which leave the system do so via a well-defined terminal: either the left ($x < 0$) or the right ($x > 0$). We also assume that the system size L is much smaller than both the dephasing length l_ϕ and the momentum relaxation (elastic scattering) length l_e. The former requirement allows treatment using a coherent wavefunction across the system; the latter allows free-space elastic scattering concepts to be applied.

The conventional distinction between 'reservoir' and 'lead' is no longer applicable, however at short distances outside the system ($r > L$ but $r \ll l_\phi$ and $r \ll l_e$) the two semi-infinite free space regions behave like leads, since they support scattering-free 'channels' (see Section 7.2). At large distances the *same* regions behave as reservoirs: for $r \gg l_e$ ergodicity ensures that the momentum distribution is uniform in angle, and for $r \gg l_\phi$ the energy is redistributed to ensure equilibrium at the relevant (experimentally-measured) chemical potential of each terminal. In the intermediate region, there is a broad cross-over from lead to reservoir.

In this work we first derive a general relation between transmission cross section (a concept we define using scattering in the half-plane) and conductance for this open geometry, in Section 7.1. In Section 7.2 we show that partial-wave type states, defined in the half-plane regions, can take the place of transverse lead modes in the Landauer formula. In Section 7.3 we examine an idealized, highly non-adiabatic QPC (a hole in a thin hard wall), for which we discuss the small hole (tunneling), large hole (semiclassical) and intermediate regimes. In Section 7.4 we discuss the maximum conductance through this QPC, which is reached when a resonator is placed on one side of the QPC. We find a universal result *regardless* how small the hole is. This illuminates the findings of a recent experiment[112] in such an open geometry. In Section 7.5 we discuss attempts to exceed the universal quantum of conductance in a single channel. A reciprocity relation for cross section is derived in Section 7.6, and the possibility of breaking this reciprocity, due to a non-thermal reservoir occupation, is described. We discuss an application to matter-wave 'conductance' through a 3D QPC. Finally we conclude in Section 7.7.

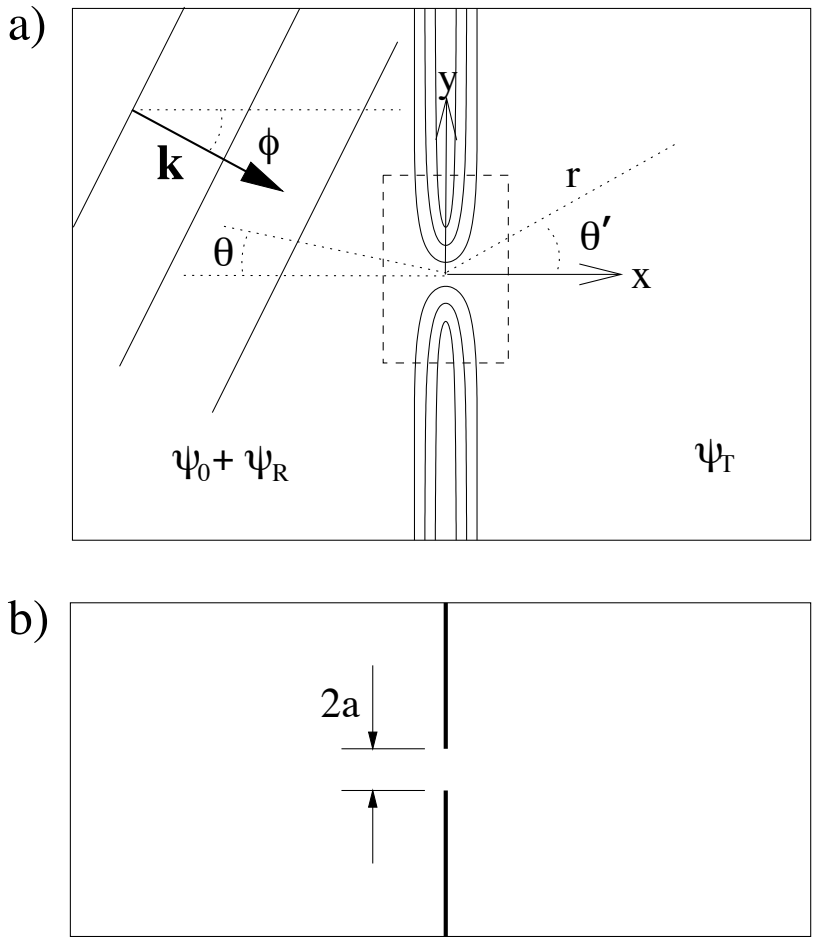

Figure 7.1: Schematic QPC geometry in 2D: **a)** general point contact scatterer coupling two semi-infinite regions of free space. The solid curves are contours of an elastic scattering potential $V(\mathbf{r})$. The 'system' size L (dashed box) we take to be the region where $V(\mathbf{r})$ has not yet reached its asympotic form (which is zero apart from a y-invariant profile around the y-axis). Also shown are an incoming plane wave, and the coordinate system. **b)** The idealized 'slit' aperture in a thin, hard wall considered in Section 7.3.

7.1 Conductance in terms of cross section

We consider scattering of a single-quasiparticle wavefunction from the general 2-terminal system described in the Introduction (see Fig. 7.1a). The Hamiltonian is $\mathcal{H} = -(\hbar^2/2m)\nabla^2 + V(\mathbf{r})$, for a quasiparticle mass m. The elastic scattering potential $V(\mathbf{r})$ completely defines the system. We imagine a monochromatic unit plane wave $\psi_{\rm I} = e^{i\mathbf{k}\cdot\mathbf{r}}$ incident from the free-space left-hand region[2]. The wavevector is $\mathbf{k} \equiv (k, \phi)$ in polar coordinates, ϕ being the angle of incidence. The free-space wavevector magnitude is taken as $k = k_{\rm F}$ (corresponding to a total energy $E = \hbar^2 k^2/2m$ equal to the Fermi energy), unless stated otherwise.

We are at liberty to choose our definition of the 'unscattered' wave ψ_0. We take it to be the wavefunction which would result from reflection of the incident wave off a wall uniform in the y direction. We can imagine creating such a wall by replacing the 'system box' shown in Fig. 7.1a by the surrounding y-invariant wall profile. Note that ψ_0 exists only on the left side. In the left free-space region it is

$$\psi_0 \;=\; e^{i(k_x x + k_y y)} - e^{i(-k_x x + k_y y + \gamma_{\mathbf{k}})} \tag{7.1}$$

where the first term is $\psi_{\rm I}$, and the angle-dependent reflection phase $\gamma_{\mathbf{k}}$ of the second term depends on both (k, ϕ) and the wall profile[3]. Upon introduction of our true system potential, the full wavefunction becomes

$$\psi \;\equiv\; \psi_0 + \psi_{\rm R} + \psi_{\rm T}, \tag{7.2}$$

where the change in reflected wave $\psi_{\rm R}$ exists only on the left side, and the new transmitted wave $\psi_{\rm T}$ exists only on the right. These scattered waves have the asymptotic ($r > L$ and $kr \gg 1$) forms of 2D scattering theory (see [128], and Chapter 7 of [174]) [4],

$$\psi_{\rm R} = f_{\rm R}(\theta)\frac{e^{ikr}}{\sqrt{r}}, \qquad\qquad \psi_{\rm T} = f_{\rm T}(\theta')\frac{e^{ikr}}{\sqrt{r}}. \tag{7.3}$$

See Fig. 7.1a for definitions of θ and θ'.

The transmission cross section $\sigma_{\rm T}(k, \phi)$ is the ratio of $\Gamma_{\rm T}$, the transmitted particle flux (number per unit time), to $j_{\rm I}$, the incident particle flux per unit length normal to the incident beam:

$$\sigma_{\rm T}(k, \phi) \;\equiv\; \frac{\Gamma_{\rm T}}{j_{\rm I}}. \tag{7.4}$$

Physically, $\sigma_{\rm T}(k, \phi)$ is the length required of an aperture oriented normal to the incident beam in order to transmit an equivalent flux of classical particles. (Note that $\sigma_{\rm T}(k, \phi)$ is

[2] Of course, throughout this chapter we could imagine the incident wave on the right-hand side, and the same conductance would result (since we are in linear response); see Section 7.6.

[3] We could equally well imagine that the QPC can be 'closed off' (no transmission) by varying a parameter (this is often true experimentally), and define ψ_0 as the full wavefunction in this closed-off state. Thus ψ_0 would be the sum of an incident plane wave and a more complicated outgoing wave. This alternative definition may be better in systems where the wall has disorder, or where there is more complicated structure on the left-hand side than shown in Fig. 7.1a. The two definitions are equivalent as far as Sections 7.3 and 7.4 are concerned.

[4] Presentations of two-dimensional scattering theory are rare; see [132, 3, 43].

proportional to the *injection distribution* [20] which can be measured in mesoscopic systems [181]). It depends on the incident angle because $V(\mathbf{r})$ has no radial symmetry. $j_{\rm I}$ is the magnitude of the incoming probability flux density vector $\mathbf{j} \equiv (\hbar/m)_{\rm Im}[\psi_{\rm I}^* \boldsymbol{\nabla} \psi_{\rm I}]$, which for a unit wave gives $j_{\rm I} = v$, the particle speed. The transmitted flux is defined as

$$\Gamma_{\rm T} \;\equiv\; \int dl\; \hat{\mathbf{n}} \cdot \mathbf{j} \;=\; \frac{\hbar}{m} \int dl\; \hat{\mathbf{n}} \cdot {}_{\rm Im}[\psi_{\rm T}^* \boldsymbol{\nabla} \psi_{\rm T}], \tag{7.5}$$

where the line integral encloses the entire transmitted wave, and the (rightwards-pointing) surface normal is $\hat{\mathbf{n}}$. Applying this and (7.4) to the asymptotic form gives

$$\sigma_{\rm T}(k,\phi) \;=\; \int_{-\pi/2}^{\pi/2} d\theta' \, |f_{\rm T}(\theta')|^2, \tag{7.6}$$

familiar from scattering theory apart from the restriction to the right half-plane. There is a corresponding form

$$\sigma_{\rm R}(k,\phi) \;=\; \int_{-\pi/2}^{\pi/2} d\theta \, |f_{\rm R}(\theta)|^2, \tag{7.7}$$

for the reflective cross section (removal from the unscattered wave without being transmitted).

We will calculate the conductance by assuming the chemical potential is slightly higher on the left side than the right, and as is usual[20, 55] consider only the left-to-right transport of the states in this narrow energy range. We take the left region to be a large ($\gg l_\phi$) closed region of area A containing single-particle states, and find their decay rate through the QPC into the right side. Semiclassically each single-particle state occupies a phase-space volume h^d, where we have $d = 2$. Therefore the phase-space density in the 2DEG Fermi sea is $2/h^2$ where the factor of 2 comes from the spin degeneracy. We can project this density onto momentum space in order to find the effective number of plane-wave states impinging on the wall[5]: this corresponds to a uniform density of states in $\mathbf{k}$-space given by

$$\rho(k,\phi)\, k dk\, d\phi \;=\; \frac{A}{2\pi^2}\, k dk\, d\phi. \tag{7.8}$$

Each state has an amplitude $A^{-1/2}$ due to the requirement of unity area normalisation in the left region, so has incoming flux density $j_{\rm I} = v/A$. Substituting this into (7.4) gives the decay rate of a state i as

$$\Gamma_{\rm T}^{(i)} \;=\; \frac{v}{A} \sigma_{\rm T}(k_i, \phi_i). \tag{7.9}$$

We can now sum the decay rates of all the left-hand states in a given wavevector range $k_{\rm F}$ to $k_{\rm F} + \delta k$, to get the current

$$\begin{aligned} \delta I = e \sum_i \Gamma_{\rm T}^{(i)} \;&=\; \frac{ev}{A} \int_{-\pi/2}^{\pi/2} d\phi \int_{k_{\rm F}}^{k_{\rm F}+\delta k} k dk\, \rho(k,\phi)\, \sigma_{\rm T}(k,\phi) \\ &=\; \frac{ev\, k_{\rm F} \delta k}{2\pi^2} \int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\rm T}(k_{\rm F},\phi), \end{aligned} \tag{7.10}$$

[5] This argument can also be verified in the more specific case of the left region being a rectangular Dirichlet box, in which case the exact eigenfunctions are known and can be written explicitly in terms of a sum of ψ_0 for incidences ϕ and $-\phi$. However the phase-space presentation is more general, and applies to the real situation where the left region is chaotic (diffusive elastic scattering).

where the last step incorporated the linear-response assumption that $\sigma_{\rm T}$ is constant over the range δk.

When a potential difference δV is applied across the QPC, the energy range carrying current is $\delta E = e\,\delta V$, which we can equate with $\hbar v\,\delta k$ using the dispersion relation. This can be used with (7.10) to write the conductance

$$G \equiv \frac{\delta I}{\delta V} = \frac{2e^2}{h}\cdot\frac{1}{\lambda_{\rm F}}\int_{-\pi/2}^{\pi/2} d\phi\,\sigma_{\rm T}(k_{\rm F},\phi) \tag{7.11a}$$

$$= \frac{2e^2}{h}\cdot\frac{k_{\rm F}}{2}\langle\sigma_{\rm T}\rangle_\phi, \tag{7.11b}$$

where the particle wavelength is $\lambda_{\rm F} \equiv 2\pi/k_{\rm F}$. The latter form is written in terms of the angle-averaged cross section at the Fermi energy. The weighting of this average is uniform because of the ergodic assumption that incoming states are uniformly distributed in angle.

Eq.(7.1) is a key result of this chapter. Like the Landauer formula, it directly connects conductance and scattering. In a scattering measurement from the left side, $\sigma_{\rm T}$ appears to be the QPC's inelastic cross section (since the transmitted waves never return to this side). In a current measurement the corresponding conductance is given by (7.1). An independent verification is provided by the result (L.4) of Appendix L, when combined with the Landauer formula Eq.(7.13). Our derivation was for temperature $T = 0$, but it applies at a finite T as long as $\sigma_{\rm T}$ does not change significantly over the energy range $k_{\rm B}T$. This can be seen by generalizing the above to include integration over the Fermi distribution.

In the limit where a QPC is adiabatic, its conductance is known to be quantized [193, 20, 65]: $G = (2e^2/h)N$ where N is the integer number of open channels at the Fermi energy. Looking at (7.11), this corresponds to quantization of the angular integral of the cross section in units of $\lambda_{\rm F}$.

7.2 Partial-wave channel modes for a 2-terminal system

In traditional scattering theory, partial waves form a basis in which to decompose the asymptotic ($r \to \infty$) form of the full wavefunction ψ into incoming and outgoing states of definite angular momentum l. The basis functions are the cylindrical (in 2D) solutions to the free-space wave equation; the S-matrix which takes incoming to outgoing waves can then be written in this basis [128]. Because there is only a single set of incoming channels and a single set of outgoing channels, this is equivalent to a scattering system (a stub) connected to a single 'lead', with an infinite number of open channel modes. This contrasts the open two-terminal geometry we study, where we need to account for two new facts: 1) in the $r \to \infty$ limit the potential V no longer becomes angular-momentum-preserving, and 2) there are now distinct ways the particle can enter and exit the system, via different leads.

We define a 'half-plane partial-wave basis' as the subset of the cylindrical free-space solutions which go to zero on the entire y-axis. This gives independent basis functions existing on either the left or right side of the y-axis. The basis is expressed in terms of Hankel functions[7] on either side:

$$\phi_l^{-L}(kr) \equiv H_l^{(2)}(kr)\sin[l(\tfrac{\pi}{2}-\theta)]$$

$$\begin{aligned}\phi_l^{+L}(kr) &\equiv H_l^{(1)}(kr)\sin[l(\tfrac{\pi}{2}-\theta)]\\ \phi_l^{-R}(kr) &\equiv H_l^{(2)}(kr)\sin[l(\tfrac{\pi}{2}-\theta')]\\ \phi_l^{+R}(kr) &\equiv H_l^{(1)}(kr)\sin[l(\tfrac{\pi}{2}-\theta')]\end{aligned} \tag{7.12}$$

where on the left (L) side θ is the angle from the negative x-axis and on the right (R) side θ' is the angle from the positive x-axis (see Fig. 7.1a). The channel index is $l = 1, 2, \cdots \infty$, and $+(-)$ refers to outgoing (incoming) travelling waves. We note that the s-wave $l = 0$ is excluded because of the y-axis barrier, leaving the first channel as the p-wave $H_1(kr)\cos(\theta)$. Assuming the width of the barrier is finite and constant as $|y| \to \infty$ (see Fig. 7.1a), then any wavefunction in the $r \to \infty$ limit can be written as a sum of the above basis functions (this is discussed further in Appendix L). The separability of this basis in (r, θ) is directly analogous to the separability of conventional (constant-width) lead basis states [129, 130, 41] into a product of transverse modes and longitudinal travelling waves.

The corresponding outgoing (incoming) amplitude coefficients are p_l^+ (p_l^-) on the left and q_l^+ (q_l^-) on the right. Applying (7.5) to the aymptotic forms of the Hankel functions $H_l^{(1,2)}(x) \approx (2/\pi x)^{1/2} e^{i(x+\text{phase})}$ shows that a unit amplitude carries a flux of $\hbar/m$, independent of channel index l. This results from the fact that all channels have the same asymptotic particle velocity, equal to that in free space. This contrasts the usual constant-width lead case where the non-zero transverse energy results in an l-dependent asymptotic velocity, and thus care is needed to include a velocity factor in the normalisation [74, 187, 55].

Now we give a brief reminder of the recipe for the conductance of a general completely-open single channel [20, 65, 55]. The channel is a 1D Fermi gas corresponding to the longitudinal degree of freedom. This can correspond to a conventional (constant-width) transverse mode, to a Hankel (cylindrical) mode defined above, or indeed to any situation where the Hamiltonian is separable into (orthogonal) transverse modes and longitudinal motion in the asymptotic regime [6]. For a 1D Fermi gas the rightward-moving density of states per unit length (including spin) is $\rho(E) = (\pi\hbar v)^{-1}$, where v is the (group) velocity in the channel. The current carried is $\delta I = ev\,\rho(E)\,\delta E$ in the relevant energy range $\delta E = e\,\delta V$, giving a single 'quantum' of conductance $G \equiv \delta I/\delta V = 2e^2/h$. The important point is that the velocity factors cancel, giving the *equipartition rule* that an equal current is carried by all open channels regardless of their channel velocities.

If the basis is chosen such that unit amplitude coefficients carry equal fluxes in all incoming and outgoing channels, then flux conservation implies the unitarity of the S-matrix when written in this basis. Then the results of the previous paragraph generalise to the familiar Landauer formula

$$G = \frac{2e^2}{h}\mathrm{Tr}(t^\dagger t), \tag{7.13}$$

in the case of general transmission matrix t [187, 55]. Our partial-wave basis (7.12) has this required flux normalisation. The transmission matrix (from incoming left to outgoing right channels) is defined by $q_l^+ = \sum_m t_{lm} p_m^-$. Thus we have shown that the partial-wave basis, appropriate for open 2-terminal systems, is equivalent to the conventional transverse

[6] The only requirement for this 1D argument to hold is that the effective longitudinal potential tend to a constant asymptotically. This allows a 1D density of states to be defined. It precludes for instance the Mathieu basis mentioned in Section 7.3, unless a reparametrization of the 'radial' coordinate is performed.

lead mode basis as far as the Landauer formula is concerned. Note that is possible to 'mix and match' different basis set types (for instance define a transmission matrix between transverse lead modes on the left side and partial-wave modes on the right), as long as equal-flux normalisation and orthogonality are preserved.

7.3 Idealized 'slit' aperture point contact

We consider an idealized QPC model: the potential V is zero everywhere except along a hard, thin wall where it is taken as infinite. (This corresponds to the reflection phase shift $\gamma_{\mathbf{k}} = 0$). The 'scattering system' is an aperture (hole) in this wall, of width $2a$ centered at the origin (see Fig. 7.1b). The desired solutions to the Schrodinger equation in this system are solutions of the Helmholtz equation corresponding to Dirichlet boundary conditions on the wall. This is a paradigm system for the study of classical particle conductance[179], and because of the sharp edges it forms a very non-adiabatic point contact. However, in mesoscopic 2DEG systems where depletion regions are defined electrostatically by gates the potential cannot possess structures sharper than about $1/k_F$ due to screening[20, 55]. Therefore as a model of real 2DEG systems, it is approximate.

In the limit of a small aperture ($ka \ll 1$) the transmission cross-section scales as $\sigma_{\mathrm{T}} \sim a^4$ for constant k, and is dominated by p-wave (dipole) scattering. This is derived in Appendix K, including the correct prefactor. The result is

$$\sigma_{\mathrm{T}}(k,\phi) \;=\; \frac{\pi^2}{16} k^3 a^4 \cos^2\phi. \qquad\qquad ka \ll 1. \tag{7.14}$$

Substituting this into the conductance formula (7.11a) yields

$$G \;=\; \frac{2e^2}{h} \cdot \frac{\pi^2}{64} (ka)^4, \qquad\qquad ka \ll 1, \tag{7.15}$$

which is much less than $2e^2/h$ so corresponds to a tunneling junction, and vanishes as the aperture closes. There is another way to compute the conductance, using the transmission matrix t of Section 7.2. Only the dipole-dipole element t_{11} will be significant (argued in Appendix K). Because transmission is very small, a unit incoming channel amplitude ϕ_1^{-L} will reflect with nearly unit amplitude into ϕ_1^{+L}. The x-gradient of this combined wavefunction a small distance to the left of the QPC will therefore [using (7.12)] be $c_1 = k$. Substituting this into Eq.(K.4) gives the transmitted amplitude q_1^+, corresponding to a matrix element

$$t_{11} \;=\; \frac{i\pi}{8} (ka)^2 \qquad\qquad ka \ll 1. \tag{7.16}$$

Reassuringly, insertion of a transmission matrix with this single non-zero element into the Landauer formula Eq.(7.13) gives the same conductance as the cross section based derivation, Eq.(7.15).

The limit of a large aperture ($ka \gg 1$) is the semiclassical limit, and the conductance is that of a classical gas (first studied in 3D by Sharvin [179].) The cross section is a purely classical quantity (independent of k), being simply the projection of the aperture length $a_{\mathrm{eff}} = 2a$ onto the incident beam direction:

$$\sigma_{\mathrm{T}}(k,\phi) \;=\; a_{\mathrm{eff}} \cos(\phi) \qquad\qquad \text{classical.} \tag{7.17}$$

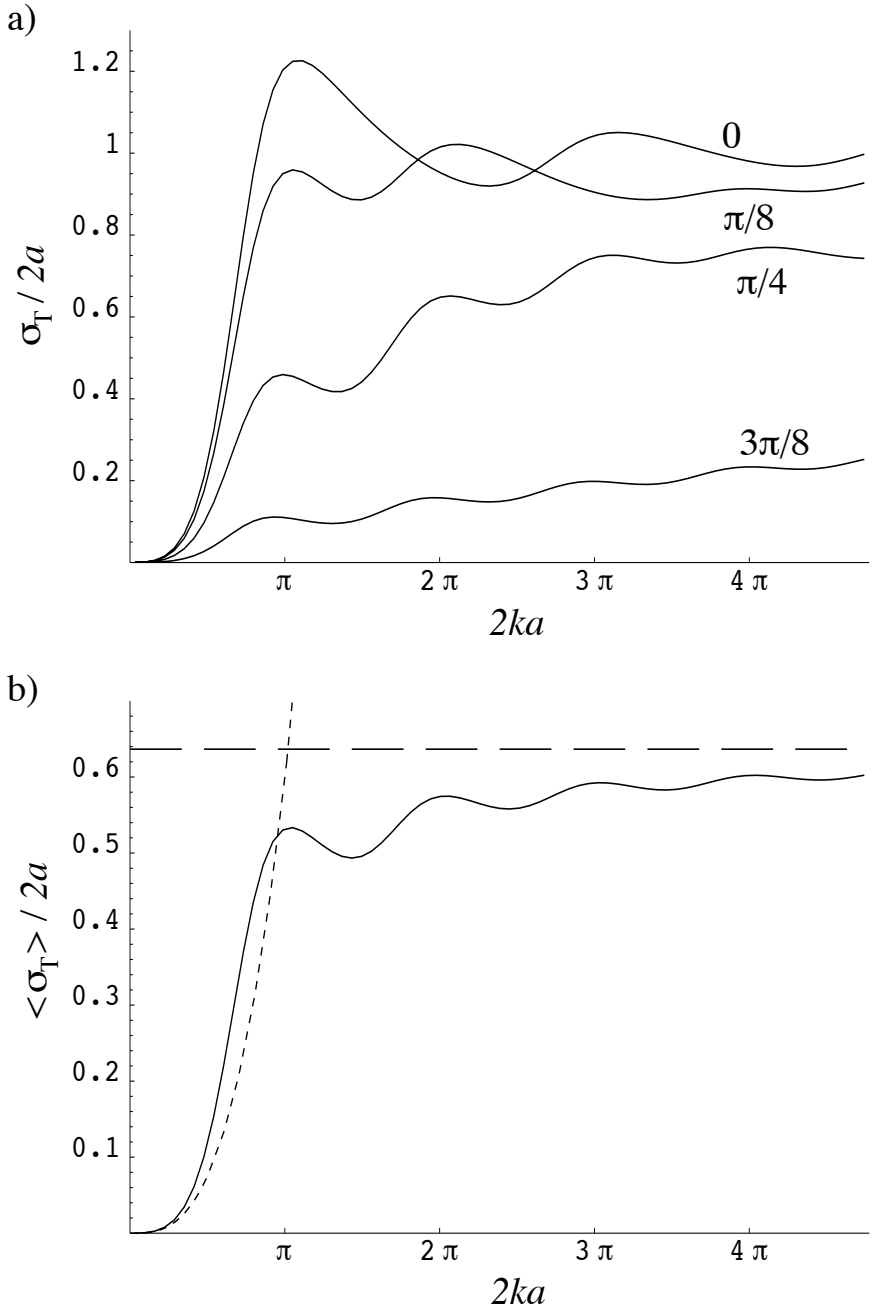

Figure 7.2: **a)** Transmission cross section $\sigma_{\rm T}(k,\phi)$ as a fraction of the classical normal-incidence cross section $2a$, vs dimensionless slit size $2ka$, for our idealised slit QPC. Four different incident angles ϕ are shown (each curve is labelled with its corresponding ϕ). **b)** The incident-angle-averaged transmission cross section as a fraction of $2a$ (solid line), compared to its small ka (dotted line) and large ka (dashed line) asymptotic forms. The conductance is proportional to $\langle\sigma_{\rm T}(k,\phi)\rangle$, and would show very weak quantization steps.

Integrating over incident direction gives

$$\int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\rm T}(k,\phi) \;=\; 2a_{\rm eff}, \tag{7.18}$$

which with (7.11) gives the conductance $G = (2e^2/h)N$ where $N = 2a_{\rm eff}/\lambda_{\rm F}$ is the number of half-wavelengths spanning the aperture. Note that this result is identical to that of an arbitrarily-long uniform scattering-free channel of the same width. When we are no longer in the semiclassical limit, or when we consider a general open 2-terminal system, (7.18) serves as a definition of the *effective* classical area $a_{\rm eff}$. Namely, $a_{\rm eff}$ is the slit aperture width whose classical conductance is equal to that of a given quantum system.

The exact result at arbitrary ka is also known. We refer the reader to Chapter 11.2 of [148] and [144] for the detailed form of the wavefunctions. They can be expressed as Mathieu functions [1] in elliptical coordinates (μ, α) defined by $x = a\sinh\mu\cos\alpha$ and $y = a\cosh\mu\sin\alpha$. The Mathieu function expansion of the incoming wave is also needed. Calculation of $\sigma_{\rm T}(k,\phi)$ requires summing the transmitted flux of the lowest few transverse Mathieu functions [144].

We apply this scheme to generate Fig. 7.2a, which shows the cross section $\sigma_{\rm T}(k,\phi)$ for arbitrary width of the aperture and for various incident angles ϕ. Fig. 7.2b shows the angular average $\langle\sigma_{\rm T}(k,\phi)\rangle_\phi$ (our plot differs from that of [144] only in that we show cross section as a fraction of the normal-incidence classical cross section $2a$). This plot also shows asymptotic convergence to the small- and large-ka results presented above, and gives an idea of when breakdown occurs.

The oscillations in $\langle\sigma_{\rm T}\rangle_\phi$ have the same period as the quantization steps in an adiabatic QPC of the same minimum width, but are much weaker, even though we are at zero temperature: this is because our QPC is very nonadiabatic [189, 206].

7.4 Point contact coupled to a resonator

Fig. 7.3 illustrates a QPC-plus-reflector system whose conductance has been experimentally measured [112]. The circular arc reflector and the vertical wall together form a cavity which can support long-lived resonances; the energy of these resonances can be swept by sweeping the reflector gate voltage. The classical condition for stability of the cavity modes is that the arc center must lie at, or to the left of, the wall (x=0) [112]. The cavity modes are coupled to the left terminal via the QPC, and to the right terminal via leakage of the modes out through the cavity top and bottom. The system is interesting because it is 'open' in the sense that it has no Coulomb blockade [20], but 'closed' in the sense that the dwell time is much greater than the ballistic time (the resonances are long-lived). It has also been studied recently in our laboratory using microwave measurements [94, 93, 95].

The actual potential in a mesoscopic experiment differs from the illustration: it has soft walls (on the scale $1/k_{\rm F}$), it may have deviations from the circle due to lithographic error, and it has modulations of the background potential due to elastic disorder [112]. However, we will not be interested in details of the resonator on the right-hand side. Rather, we will adopt the view of a 2D scattering-theorist 'looking' from the left-hand side. In this section

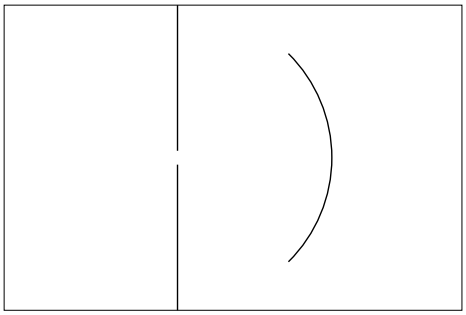

Figure 7.3: A tunneling-regime QPC combined with a nearby circular reflector, forming a stable resonant cavity open at the sides.

we discuss the maximum conductance of this system, when the 'bare' QPC (*i.e.* without the reflector) is in the tunneling regime (conductance $\ll 2e^2/h$).

We use the slit model from the previous section to model the QPC. This simplifies the treatment of the left-hand side scattering problem, and we do not believe it alters the basic conclusion. As before, we consider the incident plus reflected wave Eq.(7.1) when the QPC is closed to be the 'unscattered' wave. This we expand in Bessel functions,

$$\begin{aligned} \psi_0(\mathbf{r}) &= e^{i(k_x x + k_y y)} - e^{i(-k_x x + k_y y)} \\ &= -4iJ_1(kr)\cos(\theta)\cos(\phi) + \text{higher order terms.} \end{aligned} \tag{7.19}$$

The first term in the expansion is the incoming plus outgoing p-wave, which in the tunneling limit, will dominate in our consideration of the absorption (as shown in Appendix K).

Now we open the slit, and replace $2J_1(kr)$ in the above by $H_1^{(2)}(kr) + e^{2i\delta}\, H_1^{(1)}(kr)$, where δ follows the usual definition of partial-wave phase shift [128]. The closed slit corresponds to $\delta = 0$. An open slit leading into a closed resonator (imagine extending the arc in Fig. 7.3 to seal off the entire right side), in the case of infinite dephasing length, corresponds to $\delta =$ real, and would appear from the left side as an elastic dipole scatterer. An open slit with an open resonator corresponds to complex δ with positive imaginary part, and would appear as a general inelastic dipole scatterer. Therefore transmission though the QPC appears, to an observer on the left side, to be *absorption* of incident waves. $\sigma_{\rm T}$ is interpreted as an 'inelastic' cross section (since exiting the right-hand terminal is equivalent to leaving in a new channel), and $\sigma_{\rm R}$ as an 'elastic' one. $\sigma_{\rm T}(k,\phi)$ can be found from integrating the net incoming flux [as in Eq.(7.5)] of the total wavefunction on the *left* side. Substitution into (7.4) then gives $\sigma_{\rm T}(k,\phi) = \frac{4}{k}(1 - |e^{2i\delta}|^2)\cos^2(\phi)$. For $\delta \to i\infty$ the maximal cross section is reached,

$$\sigma_{\rm T,max}(k,\phi) \;=\; \frac{4}{k}\cos^2(\phi), \tag{7.20}$$

which gives effective classical area $a_{\rm eff} = \lambda_{\rm F}/2$. This is analogous to the fact that in 3D the effective area of an arbitrarily-small electromagnetic dipole aerial can be of order λ^2 [105]. To an observer on the left side who was able to 'see' the electron waves living in the energy range $e\,\delta V$ responsible for conductance, the QPC would stand out as a 'black dot' of size

$\sim \lambda_F$ against the surrounding uniform 'grey' thermal luminosity reflected in the vertical wall mirror.

The associated maximum conductance is found easily using (7.20) and (7.1) to be

$$G_{\max} = \frac{2e^2}{\hbar}, \tag{7.21}$$

the universal quantum of conductance (for 2 spin channels), independent of the size of the QPC hole, even for an *arbitrarily small* hole. This universal resonant-tunelling maximum conductance was first found numerically [205, 110] (also see [20]); however our system differs from those of Refs. [205, 110] because the resonance does not involve transmission though an *isolated* (zero-dimensional) quantum dot. The dramatic increase over the bare QPC conductance Eq.(7.15) runs counter to the naive classical expectation, namely that the reflector would *decrease* the left-to-right flow of electrons because it sends back into the QPC particles which would otherwise exit to the right.

How do we know that it is possible to build a resonant geometry which corresponds to $\delta \to i\infty$? The reflector can be described by r, the amplitude with which it returns an outgoing p-wave back to the QPC as an incoming p-wave. If $|r|^2 = 1 - |t_{11}|^2$, where the p-wave transmission of the QPC is t_{11} [*e.g.* see Eq.(7.16)], then the p-wave channel becomes a 1D Fabry-Perot resonator with mirrors of matched reflectivity. Sweeping the round-trip phase then produces peaks of complete transmission (corresponding to complete p-wave absorption on the left side). The ratio of peak separation to peak width is the quality factor $Q \sim 1/|t_{11}|^2$. Such peaks, with heights much greater than the bare tunneling QPC conductance, were observed in the experiments of Katine *et al.*[112]. However, Eq.(7.21) has not yet been tested quantitatively because of the difficulty of matching the Fabry-Perot reflectivities in a real 2DEG experiment. Note that the maximum conductance (7.21) also follows immediately from the Landauer formula, when we realize that there can be complete transmission of the incoming l=1 channel state (in Section 7.2).

An interesting possibility arises when we realize (Appendix K) that higher l channels are still *slightly* transmitted by the bare QPC, when $ka \ll 1$, even though they are increasingly evanescent. If the resonator has a high enough reflectivity for these modes, then additional Fabry-Perot conductance peaks will be produced [205, 38]. The peaks may be extremely narrow, but can carry a full quantum of conductance because they can transmit another incoming l channel. By careful arrangement of the cavity, one or more of these peaks could be brought into conjunction with an already-existing l=1 peak at the Fermi energy. (For instance, the l=1 and l=2 resonances are in different symmetry classes in Fig. 7.3 so there can be an exact level crossing). Therefore, we have the surprising result that, in theory, a conductance of $(2e^2/h)n$ can pass through an arbitrarily small QPC hole if n resonances (from n different channels) coincide at the Fermi energy. However, due to their extremely small width, such large conductance peaks are unlikely to be observable in a real mesoscopic tunneling QPC due to finite dephasing length and finite-temperature smearing [20].

Finally, we should not overlook the fact that our expressions for partial cross sections are a factor of 4 greater than those conventionally arising in 2D scattering theory from a radial potential [128, 132, 3, 43], because we are measuring cross section on the reflective boundary of a semi-infinite half plane. For instance, the maximum inelastic partial

cross section for a single channel in free space is $\sigma_r = 1/k$ [128, 132, 3, 43], compared to our maximum 'inelastic' cross section per channel Eq.(7.20). Similarly, the maximum elastic result in free space is $\sigma_e = 4/k$, compared to our maximum (normal-incidence) 'elastic' cross section per channel $\sigma_{\mathrm{R,max}} = 16/k$. This latter case occurs when $\delta = (\text{integer} + \frac{1}{2})\pi$.

7.5 What is the maximum conductance of a single quantum channel?

The surprising theoretical results of the previous section lead one to question the conductance limit $2e^2/h$ for a single quantum channel of the general 1D-Fermi-gas type described in Section 7.2. For this gedanken-experiment we will consider conventional electron waveguides which are single-mode and long enough that evanescent waves are negligible, but $\ll l_\phi$. We try to encourage more current to pass down a single-mode channel (E) by connecting it to a reservoir via multiple routes (A,B,C,D), as shown in Fig 7.4a, where two routes are used on each side. It is possible to match the junctions so that a wave entering down A,B,C, or D has no reflection back along the same lead. In this case we might guess that the hypothetical left-side observer (from the previous section) would see the single-mode entrances to guides A and B as two 'black dots', giving twice the effective absorption cross section, and therefore infer a conductance of twice $2e^2/h$. We might also justify this by saying that waves travelling down A and B will meet and continue down E, and since they have no particular phase relation, their currents will add to give a doubled current through E, as would be necessary.

However there is a fundamental flaw in the above reasoning. The ABE junction can be designed so that if waves come down A and B in phase, they will be adiabatically transformed into the lowest transverse mode of E, so will propagate through to the right side without reflection, carrying a current of twice that of a usual single-mode guide. However, if A and B are π out of phase, the same adiabatic transformation must result in the second transverse mode, which is evanescent. So this latter wave will reflect perfectly back out of the left side, and carry no current. Plane waves are impinging from the left reservoir uniformly over all angles, and because of the $> \lambda$ separation of the entrances, an average over angles gives an average over relative phase in A and B. Thus we are left with no increase above the single channel conductance. This property of the ABE junction is not merely practical; rather, it is easy to show that its 3×3 S-matrix cannot be unitary if a junction is to couple both A$\rightarrow$E and B$\rightarrow$E with unity transmissions. Such an appealing junction is therefore ruled out on the grounds of flux conservation. A consequence is that the entrances to A and B can at most appear 'half black' to the observer, due to waves which enter A then exit B and vice versa.

This suggests another way to try and defeat the conductance limit: direct the incoming plane waves in a narrow enough angular distribution so that waves *always* come down A and B in phase, and this will double the conductance. (This is similar to experiments [198] where the series resistance of two QPCs was found to be less than the sum of the individual QPC resistances, because collimation at the exit of the first QPC illuminated the second with a narrow beam, increasing its conductance). However, this beam is no longer a *thermal* occupation of incoming states. This illustrates the inextricable link between

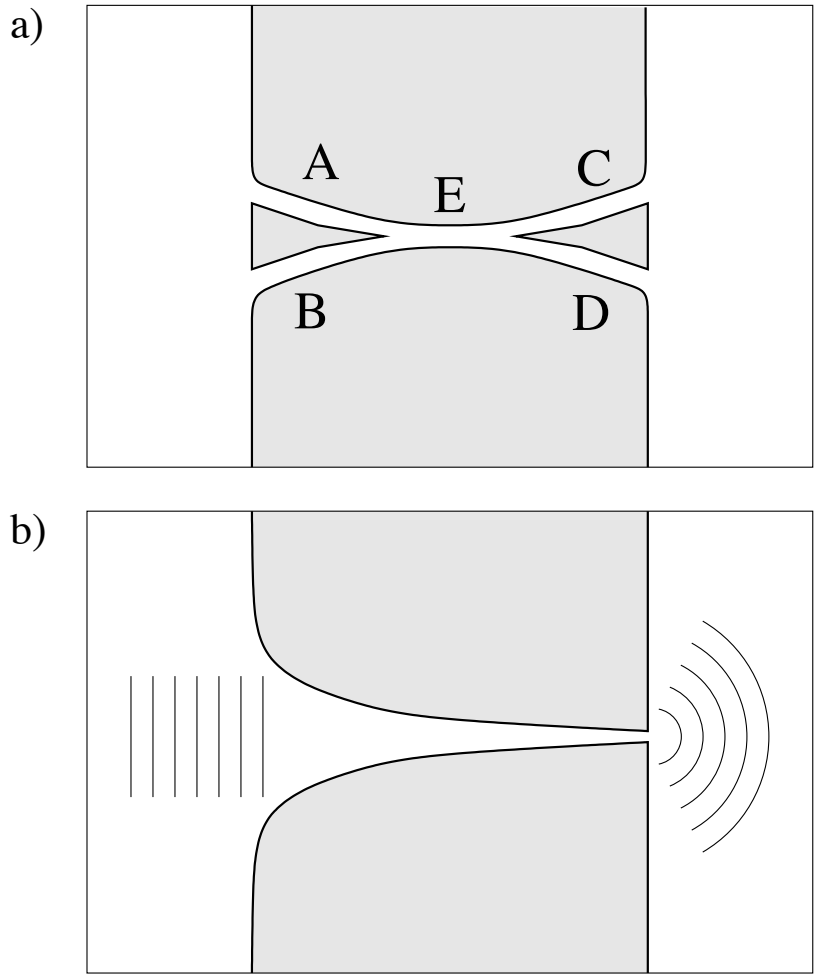

Figure 7.4: **a)** An attempt to increase conductance through a single channel by multiple connections feeding from the reservoirs. All channels are single-mode and sufficiently long that the evanescent tunneling of higher modes is negligible. **b)** An illustrative hard-walled exponential horn system which has differing acceptance angles on each side: very narrow on the left, and very wide on the right. Such a mesoscopic 2DEG system would exhibit symmetric conductance, however, in an atom beam context the conductance can become unsymmetric.

thermal Fermi occupation of reservoir states and the universal quantum of conductance. At T=0, thermal occupation at a given chemical potential difference implies that *all* quantum states lying in the appropriate energy range are filled in the left reservoir and empty in the right. Semiclassically, this corresponds to a uniform distribution in phase space, or when projected into momentum states, uniform in angle, as exemplified by Eq.(7.1). The semiclassical viewpoint allows one to see that since transformations in phase space cannot change the phase space density (Liouville's theorem), neither can the universal conductance per quantum channel be changed. This reminds us that unitarity in quantum mechanics is analogous to Liouville's theorem in classical mechanics.

7.6 Reciprocity and 'conductance' of atom waves

We can ask if the conductance (7.1) computed using transmission of left-side reservoir plane wave states through the QPC is equal to that using right-side reservoir states. Since the two directions correspond to opposite signs of δV, then in order to have linear response (well-defined constant G around $\delta V = 0$) we would hope that they are equal. That the angular average of transmission cross section is equal from the left and right sides is not immediately apparent in a general asymmetric system. For instance, consider Fig. 7.4b which has a small acceptance angle from the left but a large from the right, therefore very different forms of the transmission cross sections $\sigma_{\rm T}^{\rm L\to R}(k,\phi)$ and $\sigma_{\rm T}^{\rm R\to L}(k,\phi)$.

If we assume classical motion then we can imagine a map from a Poincaré Section (PS) (y,p_y) at a vertical slice at $x=-x_0$ to another PS (y',p_y') at $x=+x_0$. At each PS we consider only rightwards-moving ($p_x>0$) particles, and take $x_0>L$. A certain area of phase space (y,p_y) is transmitted and is mapped to an *equal area* in phase space (y',p_y') [88]. Time-reversal invariance holds since we consider magnetic field B=0, so we can negate the momenta (now considering $p_x<0$) and find that the *same* phase space area is transmitted right-to-left. When it is realised that the angle-averaged cross section is proportional to the transmitted phase space area on a PS, then the symmetry of the angle-averaged classical cross sections follows.

This is not obvious either for quantum cross sections, but it is also true. Comparing (7.11) with (7.13) gives

$$\int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\rm T}^{\rm L\to R}(k_{\rm F},\phi) \;=\; \lambda_{\rm F}{\rm Tr}(t^\dagger t), \tag{7.22}$$

where t is measured from left to right states. It is instructive to derive this directly; this is done in Appendix L, giving an independent verification of (7.1). This relation ties together the cross section and Landauer views of conductance. Time-reversal invariance and flux conservation together imply [7] that ${\rm Tr}(t^\dagger t)$ is unchanged by swapping the labelling of the leads [20, 187], thus we immediately have from (7.22) the reciprocity of angle-integrated

[7] With $B\neq 0$, the conductance is *still* symmetric under swapping the leads. This results from the 2-terminal special case of unitarity sum rules[55], namely that the rows and columns of the matrix of absolute-value-squared S-matrix elements must all sum to 1. Thus the reciprocity derived here is preserved for $B\neq 0$. How the classical argument from the previous paragraph generalizes for $B\neq 0$ is not known by the authors.

quantum cross section

$$\int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\rm T}^{\rm L\to R}(k,\phi) \;=\; \int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\rm T}^{\rm R\to L}(k,\phi). \tag{7.23}$$

So in Fig. 7.4b is it now clear that the ratio of acceptance angles is balanced by the ratio of effective areas.

We now discuss a case in which non-thermal occupation of incoming states is possible: the rapidly developing field of coherent matter-wave optics, in which potentials are defined by microfabricated structures ([76, 191, 16] and references therein). There is a recent proposal [191] for observation of quantization of atomic flux through a micron-sized 3D QPC defined by the Zeeman effect potential of a magnetic field. The device is illuminated by a beam of atoms passing through a vacuum, whose angular distribution is an experimental parameter (for instance, a collimated oven source or a dropped cloud of cold atoms [116]). The atomic flux transmitted (per unit k, at wavevector k) will be $F(k) = G_{\rm atom}(k)\, J_0(k)$ where $J_0(k)$ is the flux incident per unit wall area, and we define the atomic 'conductance' by

$$G_{\rm atom}(k) \;\equiv\; \int d\Omega\, w(k,\Omega)\, \sigma_{\rm T}(k,\Omega). \tag{7.24}$$

As before, the quantum transmission cross section is $\sigma_{\rm T}(k,\Omega)$, but now there is a *weighting function* $w(k,\Omega)$ which defines the angular distribution of the incident beam [8]. The weight has the normalization $\int d\Omega\, w(k,\Omega)\, \cos(\theta) = 1$. (All integrals over solid angle $\Omega \equiv (\theta,\phi)$ are over a range of 2π appropriate for the half-sphere). Following the analogy of Ref. [191], $F(k)$ plays the role of current, $J_0(k)$ that of bias voltage. However, the name 'conductance' does not imply any definite chemical potential difference as in the 2DEG case. For classical particles, the 'conductance' of an aperture of area $A_{\rm eff}$ in a thin wall is simply $G_{\rm atom}(k) = A_{\rm eff}$, regardless of the incident angular distribution. Thus $G_{\rm atom}(k)$ gives the effective area $A_{\rm eff}$ of a QPC, in an analogous fashion to $a_{\rm eff}$ in 2D [see Eq.(7.18)].

For an integer number of quantum channels, the 2D quantization of $a_{\rm eff}$ in units of $\lambda/2$ [see Eq.(7.18] becomes in 3D the quantization of $A_{\rm eff}$ in units of λ^2/π [191, 20, 179]. As stated in Ref. [191], this accurate flux quantization requires the incident beam width to be much larger than the QPC acceptance angle.

Eq.(7.24) is the matter-wave equivalent of Eq.(7.11), with the important difference that it has a general weight function. Non-uniformity of this weight function leads to a key result: that *asymmetry* of the conductance is possible given identical illumination on either side, even though the (center of mass) motion is time-reversal invariant. For example, if the incident flux used to illuminate the horn QPC of Fig. 7.4b is narrow in angular spread, then the left-to-right conductance will be much larger than the right-to-left conductance. This contrasts with the 2DEG case where the conductance is always symmetric.

[8] Because we wish to consider general illumination and general $st(k,\Omega)$, our definition of 'conductance' coincides with that of Ref. [191] only in the case of isotropic illumination $w(k,\Omega) = 1/\pi$. The beam *brightness* per unit k range, that is, its phase-space density, is assumed uniform in position space, and is proportional to $J_0(k)\, w(k,\Omega)$. This is also proportional to $a(\mathbf{k})$ defined in Ref. [191].

Finally, it is interesting to note that for the non-thermal incident (reservoir) distributions discussed above, the Landauer formula has the modified form

$$G \propto \mathrm{Tr}(t^\dagger t \rho) \tag{7.25}$$

where ρ is the density matrix of the incident beam.

7.7 Conclusions

Quantum scattering theory in the 2D half-plane can provide an alternative description of the mesoscopic conductance of non-interacting particles. It is especially useful in 'open' systems (*e.g.* those with nearby scatterers in the reservoir regions) where the usual transverse-channel approach is inappropriate. We have considered elastic potentials in zero magnetic field, in linear response in the low temperature limit. Conductance is proportional to the transmission cross section integrated over all incident angles, Eq.(7.1). We also define a half-plane partial-wave basis applicable with the usual Landauer formula, and relate this to our transmission cross section result (Appendix L). A difference between this and previous work is the ability to treat a direct 'leadless' connection to the reservoir.

We have studied the case of a single slit in a thin hard wall in detail. Its conductance (7.15) in the small-aperture limit arises entirely from its p-wave transmission coefficient, which can be found by matching to solutions of Laplace's equation (Appendix K). Using the example of a slit QPC combined with an open cavity structure, we show that an arbitrarily small QPC can carry up to a single quantum of conductance via resonant tunnelling (equal to the limit in the closed-dot resonant tunnelling case). This requires a resonance at the Fermi energy. If n coincident resonances occur for different incoming channels, then n conductance quanta can in theory be achieved through this same tunneling QPC, a result which we believe has not been noted previously.

We emphasize that conductance is proportional to phase-space density of the reservoir states. Therefore the universal quantum of conductance e^2/h per spin in Fermi gas systems is a direct result of the uniform phase-space density (angular distribution) in a thermal occupation of the Fermi sea. This insight is supported by discussion of attempts to exceed this universal value. When the reservoir occupation differs from thermal, the conductance formula requires generalization: an angle-dependent weight is included in the cross section integral (7.24), equivalently the Landauer formula requires inclusion of the incoming ensemble (7.25). This result, and our approach in general, is relevant to the emerging field of matter-wave conductance by microfabricated structures (for instance, a quantum point contact in 3D), under general illumination by atom waves. We hope this work provides new tools for the study of coherent electron and matter-wave systems.

Chapter 8

Substrate-based atom waveguide using two-color evanescent light fields

8.1 Introduction

There has been much recent progress in the trapping and cooling of neutral atoms, opening up new areas of ultra-low energy and matter-wave physics [45, 51, 159]. Waveguides for such atoms are of great interest for atom optics, atom interferometery, and atom lithography. Multimode atom waveguides act as incoherent atom pipes that could trap atoms, transport them along complicated paths or between different environments, or deliver highly localized atom beams to a surface. Single-mode waveguides (or multimode guides populated only by atoms in the transverse ground-state) could be used for coherent atom optics and interferometry [151, 2], as well as a tool for one-dimensional physics such as boson-fermion duality [190, 152, 44] and low-dimensional Bose-Einstein condensation effects [115].

The optical dipole-force has long been used to trap and manipulate atoms [45, 51, 159] as well as dielectric particles [9, 113] (for a review see [10]). The available intensity of lasers has allowed a multitude of such atom traps in the far-detuned regime, giving very low decoherence and heating rates, and storage times on the order of seconds (for a review of dipole-force atom trapping, see [85]).

Evanescent light waves have been popular in many atom mirrors, traps and guides (see the original theoretical work [52, 13] and review [67]) since they can provide potentials with high spatial gradients (decay lengths $\sim \lambda/2\pi$ where λ is the optical wavelength), and use rigid dielectric structures (prisms, fibers) to define the potential shape. For example, there has been a series of repulsive (blue-detuned) evanescent-wave (EW) traps which rely on gravity to provide the counteracting force [4, 155] and recent experiments have shown that hollow optical fibers can guide atoms confined within the hollow core using a repulsive evanescent field guided by the fiber [170, 102, 103].

The idea of using an EW to provide both attractive and repulsive forces is due to Ovchinnikov *et al.* [156], who proposed the use of two colors (*i.e.* red and blue detunings) and differing evanescent decay lengths to achieve a trap with the potential minimum a

distance $\sim \lambda$ from a prism surface. Until now, this design has been restricted to planar traps (weak confinement in the other two dimensions).

In this chapter we discuss a two-color trap based on the EW fields above a single-mode, submicron optical 'channel' waveguide. The trap provides tight confinement in *two* dimensions and allows free de Broglie wave propagation in the third, forming an atomic waveguide that could transport atoms a between $\lambda/4$ and $\lambda/2$ above the optical guide surface. Our proposal is to utilize the differing vertical evanescent decay lengths of the two *polarizations* carried in the single-mode optical guide (see Figure 8.1). The physical origin of this decay length difference is the fact that the TM mode is closer to optical *cut-off* than the TE mode at the same frequency.

Our proposal is reminiscent of some existing resonant enhancement schemes for EW mirrors (demonstrated with surface plasmons [71] and dielectric waveguides [109, 124]) but with a radical change from a planar geometry to a linear, the mechanism for exciting the guide, and the simultaneous guiding of a second frequency of opposite detuning. It also shares the feature of two guided colors with an atom trap proposal using microsphere whispering-gallery modes[141].

Our design has many desirable experimental features: 1) very little optical power is required to obtain large trapping intensities since the optical bound mode has very small cross sectional area ($\sim 0.3\,\mu m^2$), 2) the optical field is *non-divergent*, so can be maintained over distances orders of magnitude further than diffraction-limited propagation in free space allows, 3) the trapping potential is well-known, mechanically stable, and insensitive to experimental parameters other than the optical powers, since it is defined by single-mode intensity distributions fixed relative to a substrate, 4) fabrication of arrays of closely spaced atom waveguides is possible[176], for parallel lithography or measurement, creating "on-chip" integrated atom-optical elements, 5) the atoms are exposed providing additional optical and physical access (a feature not shared with hollow-fiber designs), and 6) the velocity of the atoms along the direction of the waveguide could be controlled by standing waves in the light carried by the optical guide [114].

Compared to a Zeeman-effect magnetic trap for neutral atoms, far-detuned optical dipole-force traps can have comparable trapping times, but typically an order of magnitude less depth and transverse mode spacings than recent magnetic traps [197, 97, 104, 59, 190] (for an introduction see [22]). However, in microfabricated applications the stray magnetic fields decay as a power law with distance, whereas evanescent light fields decay exponentially (ignoring for now any scattering into free space caused by optical defects). We believe this could give guided optical traps a distinct advantage in terms of achievable density of independent atom-optical elements on a single substrate.

Also, optical traps have the advantage that there is no significant loss mechanism which can remove atoms from the trap (assuming the thermal energy is much less than the trap depth): spontaneous events cause a small heating rate, and non-adiabatic changes in m_F can change the optical potential but not the fact that the atom remains trapped. This is to be contrasted with a non-adiabatic spin-flip event in a magnetic trap, which results in loss of the atom. This makes optical waveguides particularly attractive for incoherent transport, when the loss of coherence due to the spontaneous events is unimportant. Finally, optical manipulation has the advantage over magnetic manipulation in terms of high possible

switching speeds.

This chapter is organized as follows. In Section 8.2 we describe the dipole potential, the exponential approximation for the EW fields, and the mechanism for the difference in decay length. We show how we optimized the optical guide dimensions, in the case of a rectangular guide on a substrate of unity refractive index (for $m_F = 0$), and discuss some design objectives and implementation issues. In Section 8.3 we give simulated results for cesium atoms: trap depth, coherence time, transverse mode spacing and Q factor, and spontaneous heating rate. We also show how depth and coherence time are generally limited by only two parameters (the detuning and the normalized decay length difference). We study both the case of a substrate refractive index of unity, and in Section 8.3.3 the more realistic index of 1.32. We describe the numerical electromagnetic finite element technique in Section 8.4, including the accuracy achieved. Section 8.5 is an investigation of two potential causes of loss or decoherence of atoms, namely interactions with the dielectric surface and bending of the waveguide. Finally in Section 8.6 we conclude and give some future prospects for this proposal.

8.2 Trap Concept

8.2.1 Theory of the light potential

An atom in a near-resonant light field of frequency ω experiences both a conservative force (due to stimulated photon exchange) and a dissipative force (due to spontaneous photon emission)[82, 50]. The conservative force is the gradient of a spatially-dependent potential $U_{\rm dip}(\mathbf{r})$ which can be viewed as the time-averaged induced dipole interaction energy (proportional to the real part of the classical polarizability) in the electric field, or equally well as the 'light shift' (that is, energy level shift due to the ac Stark effect) of the atomic ground state [2, 85].

We assume that we apply a monochromatic light field of detuning $\Delta \equiv \omega - \omega_0$ to an alkali atom (with the $n\,s \rightarrow n\,p$ transition resonant at ω_0), in the *far-detuned* regime (Δ greater than the excited state hyperfine splitting, but much less than ω_0) and the *low saturation* regime ($\Omega_{\rm rabi} \ll \Delta$, where the Rabi flopping rate $\Omega_{\rm rabi}$ is defined [1] by $\hbar\Omega_{\rm rabi} \equiv \mu E_0$, the dipole matrix element multiplied by the electric field amplitude). The dipole potential has both a scalar and a magnetic part:

$$\begin{aligned} U_{\rm dip}(\mathbf{r}) &= \beta_s \frac{\hbar\Gamma}{8}\frac{\Gamma}{\Delta}\, s(\mathbf{r}) \\ &+ \mu_{\rm Bohr}\, g_{\rm nuc}(L,S,J,i,F) m_F |\mathbf{H}(\mathbf{r})|\,, \end{aligned} \tag{8.1}$$

where Γ is the spontaneous decay rate, $s(\mathbf{r})$ is the spatially-dependent saturation parameter, and the potential is taken to be much less than the ground state hyperfine splitting. Only the magnetic part is affected by m_F, which is defined as the projection of the total angular momentum F on the direction of the local effective magnetic field $\mathbf{H}(\mathbf{r})$. The constant $g_{\rm nuc}$ is the nuclear Landé g-factor appropriate for the F of the ground state. The scalar potential

[1] Our definitions of $\Omega_{\rm rabi}$ and $I_{\rm sat}$ correspond to those of C. S. Adams *et al.* [2]

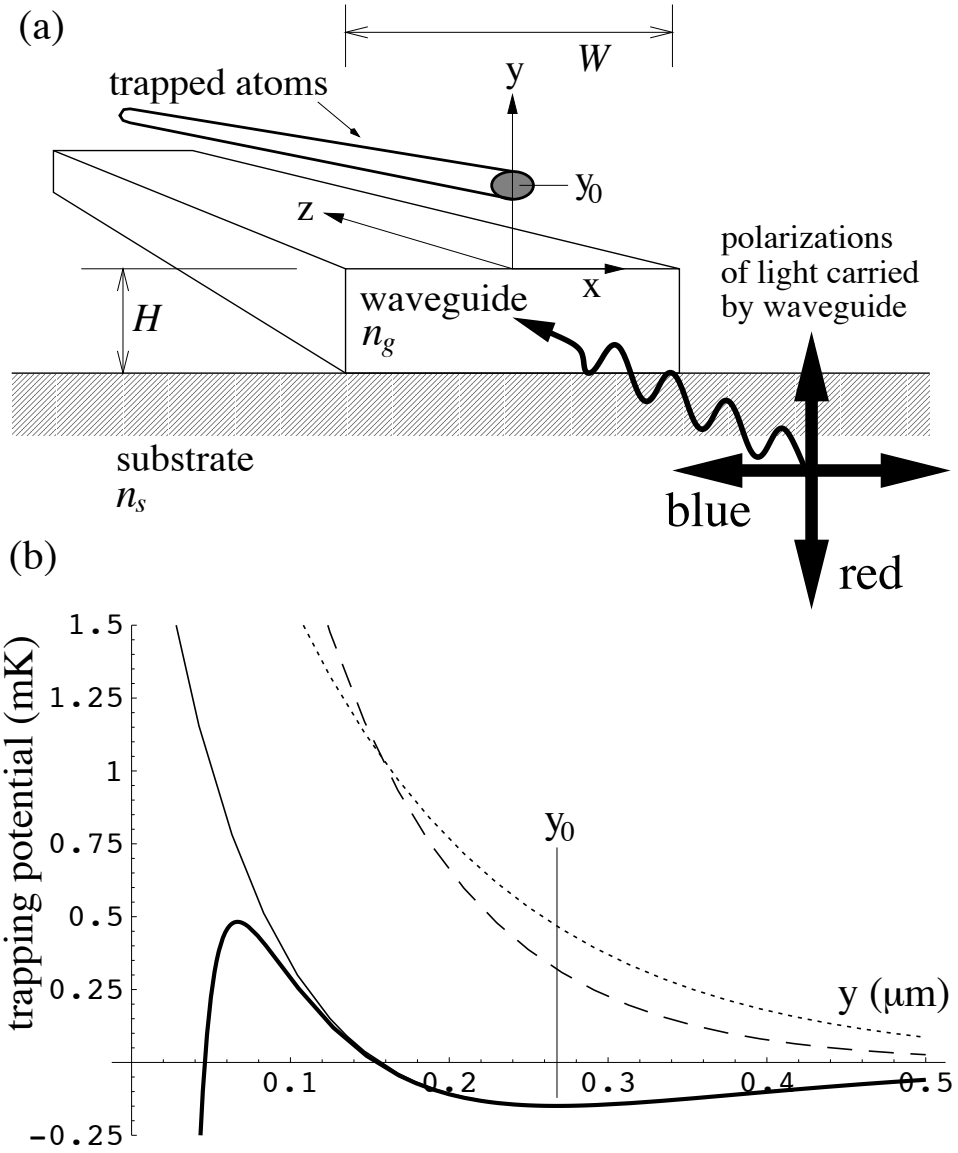

Figure 8.1: (a) Shows trap geometry, dielectric guide dimensions, incoming laser polarizations, and the cartesian axes; (b) shows the trapping potential above the dielectric along a vertical slice at $x = 0$. The component due to red-detuned light (absolute value shown as dotted line) subtracts from that of blue-detuned light (dashed line) to give the total dipole potential $U_{\rm dip}$ (thin solid line). This is modified by the Casimir surface interaction (Section 8.5.1), giving the final potential (thick line). Here the trap depth of 150 μK and coherence time of 12 ms is generated in our design by 2 mW total guided laser power detuned by $\pm$15 nm from the cesium D2 line.

is identical to the case of a two-level atom, apart from the strength factor β_s which is $\frac{1}{3}$ for detuning from the D1 line, $\frac{2}{3}$ for the D2 [85].

It is important to be precise with the definition of the saturation parameter. We write

$$s(\mathbf{r}) \equiv \frac{E_0^2(\mathbf{r})}{E_{0,\mathrm{sat}}^2}, \tag{8.2}$$

where $E_{0,\mathrm{sat}}^2 \equiv 2I_{\mathrm{sat}}/\epsilon_0 c$ (in the MKSA system) is the squared electric field amplitude in a plane wave of intensity I_{sat} (following definition in [85]). This avoids the ambiguities present with the usual definition $s(\mathbf{r}) \equiv I(\mathbf{r})/I_{\mathrm{sat}}$ in the case of a general monochromatic light field (composed of an arbitrary coherent sum of travelling and evanescent waves), and emphasizes that it is the local electric field alone that causes the dipole potential.

The effective magnetic field has a strength and direction given by the circularly polarized component of the electric field[85], which can be written thus:

$$\mu_{\mathrm{Bohr}}\mathbf{H}(\mathbf{r}) \;=\; \beta_m \frac{\hbar\Gamma}{8}\,\frac{\Gamma}{\Delta}\,\frac{\Re[\mathbf{E}_0^*(\mathbf{r})] \times \Im[\mathbf{E}_0(\mathbf{r})]}{E_{0,\mathrm{sat}}^2}, \tag{8.3}$$

where the strength factor β_m is $-\frac{2}{3}$ for D1, $\frac{2}{3}$ for D2, and we take the physical electric field (with amplitude $E_0(\mathbf{r}) \equiv |\mathbf{E}_0(\mathbf{r})|$) to be the real part of a complex field $\mathbf{E}(\mathbf{r},t) \equiv \mathbf{E}_0(\mathbf{r})\exp(-i\omega t)$. The reason for the 'dummy' constant μ_{Bohr} is aesthetic, so that (8.1) can be expressed in a standard magnetic form. Note that for nonzero m_F the magnetic and scalar contributions to the potential are of similar order, if the fields have significant circular polarization (this will be true for our trapping fields, for the reason that the optical guide is close to cut-off).

The fact that $U_{\mathrm{dip}}(\mathbf{r})$ has its sign controlled by the sign of the detuning allows both attractive (red-detuned) and repulsive (blue-detuned) potentials to be created. The potential scales as I/Δ but the spontaneous emission rate scales as I/Δ^2; from this follows the well-known result that, if coherence time is an important factor, it is best to be far off-resonance and use high intensities in order to achieve the desired trap depth [85, 2].

For simplicity, in this chapter we will restrict our further analysis and simulations to $m_F = 0$, although our initial calculations suggest that the effect of the magnetic part of our potential when trapping in other m_F states will not pose major problems (assuming the spin axis adiabatically follows the $\mathbf{H}(\mathbf{r})$ field direction), and can even be used to our advantage by increasing the depth and the transverse oscillation frequency ω_x in the case $m_F > 0$. Also, we will consider the effect of only a single resonance (choosing D2 because it has a larger β_s than D1), which is a valid approximation when the detunings from this resonance are much less than the alkali atom fine structure splitting Δ_{fs}. Even when it becomes advantageous to use a large detuning of the order of Δ_{fs}, it is possible to cancel the effect of the other resonance by a simple shift in the two detunings (as we will see at the end of Section 8.3.1).

If we now have two light fields of differing frequency, the atomic potentials add[82, 156], as long as we assume that the timescale of atomic motion is much slower than the beating period (that is, the inverse of the frequency difference). In our case, atomic motion occurs at $\sim 10^5$ Hz and our light field frequency difference is $\sim 10^{15}$ Hz, so this assumption is

valid. Choosing equal but opposite detunings $\pm\Delta$ about the D2 line, the trapping potential for $m_F = 0$ is

$$U_{\rm dip}(\mathbf{r}) = \frac{2}{3}\frac{\hbar\Gamma}{8}\frac{\Gamma}{\Delta}\left[s_{\rm blue}(\mathbf{r}) - s_{\rm red}(\mathbf{r})\right], \tag{8.4}$$

written in terms of the saturation parameters for the two colors. The spontaneous scattering rate[2] is a factor Δ/Γ smaller than $U_{\rm dip}(\mathbf{r})/\hbar$ (in fact this relation applies for any m_F state), but it depends on the *sum* of the saturation parameters rather than the *difference*. It is also spatially dependent, and has the form

$$\Gamma_{\rm scatt}(\mathbf{r}) = \frac{2}{3}\frac{\Gamma}{8}\left(\frac{\Gamma}{\Delta}\right)^2\left[s_{\rm blue}(\mathbf{r}) + s_{\rm red}(\mathbf{r})\right]. \tag{8.5}$$

8.2.2 Design of the light fields

Our basic task is to create intense evanescent light fields with a potential minimum sufficiently far from a dielectric surface to make the surface interaction potential and heating mechanisms negligible (discussed in Section 8.5.1). The main difficulty arises because the evanescent fields have a typical exponential decay length $\sim \lambda/2\pi$, so if we are to have a trap of useful depth, we are restricted to keep it within roughly λ of the surface (less than a micron).

A potential minimum in one dimension can be obtained using a blue (repulsive) light field of higher intensity at the dielectric surface than the red (attractive) light field, and ensuring the decay lengths obey $L_{\rm red} > L_{\rm blue}$, giving a potential of the form

$$U_{\rm dip}(y) = A_{\rm blue}e^{-y/L_{\rm blue}} - A_{\rm red}e^{-y/L_{\rm red}}. \tag{8.6}$$

This gives a repulsive force at short range, which becomes attractive at long range (see Figure 8.1b), and is the scheme for the planar trap of Ovchinnikov *et al.* [156]. A large amount of insight into our proposed trap can be gleaned from this simple one-dimensional model (which we call the *exponential approximation*), because the squared electric fields above the guide will turn out to approximate exponential forms in the vertical direction quite closely.

If we define a normalized decay length difference $\alpha_L \equiv (L_{\rm red} - L_{\rm blue})/L_{\rm blue}$, then we can give two reasons why increasing α_L is a vital design objective. Firstly, it is easy to show that for small α_L the deepest available trap depth (found by optimizing the ratio of surface intensities $A_{\rm red}/A_{\rm blue}$) scales as α_L. Secondly, a larger α_L is beneficial for trap coherence, (giving a smaller spontaneous decay rate at a given trap depth and detuning), because the sum of the intensities can be kept lower (see equation (8.5)) for a given intensity difference (equation (8.4)). We will quantify this latter connection in Section 8.3.1.

Our two key differences from the proposal of Ovchinnikov *et al.* are as follows. Firstly, we create a non-zero α_L by using two orthogonally-polarized bound modes of a dielectric slab guide, which have different evanescent decay lengths at the same frequency[2].

[2]Even though the frequencies will actually differ by a couple of percent because of their detunings, and cause a difference in decay lengths of the correct sign, this effect alone is not large enough to create large α_L values. Future designs with larger detunings, possibly much larger than the fine splitting, may be able to use this effect alone

This contrasts with Ovchinnikov *et al.* who suggest varying the decay lengths by varying the reflection angles from the inside surface of a glass prism. Secondly, horizontal confinement is achieved by limiting the width of the slab guide to approximately λ (forming what is called a *channel guide* [117]), which automatically creates a maximum in each light intensity field in the horizontal direction. This results in a tight horizontal confinement in the atomic potential, of similar size to the vertical confinement, and is something very hard to achieve in a prism geometry.

A schematic of our design is shown in Figure 8.1a. The optical guide height H and width W are kept small enough to guarantee that there are exactly two bound modes, differing in polarization but not in nodal structure (in optical terminology this is called *single-mode*): E_{11}^x has an electric field predominantly in the x direction[3], and is to be excited by blue-detuned laser light, and E_{11}^y has electric field predominantly in the y direction and is to be excited by red-detuned laser light. We can see why their vertical decay lengths differ by considering the case of the slab (*i.e.* taking the width $W \to \infty$), where these modes are simply the slab TE and TM modes respectively. For both these slab modes the purely transverse field obeys the differential equation

$$\frac{\partial^2 \phi}{\partial y^2} = k_0^2 \left[(n_{\text{eff}}^{(i)})^2 - n(y)^2 \right] \phi \,, \tag{8.7}$$

where $\phi = E_x$ (H_x) for the TE (TM) mode, the eigenvalue $n_{\text{eff}}^{(i)} \equiv k_z^{(i)}/k_0$ is the effective refractive index for the i^{th} mode (k_z being the wavenumber in the propagation direction and k_0 the free space wavenumber), and $n(y)$ is the spatially-dependent refractive index [117]. (This is equivalent to a one-dimensional quantum problem in the direction normal to the slab, in a potential $-k_0^2 n(y)^2$ with $\hbar^2/2m = 1$). However the boundary conditions on the slab surfaces differ for the two mode types: ϕ is always continuous, but $\partial\phi/\partial y$ is continuous for TE as opposed to $n^{-2}\partial\phi/\partial y$ continuous for TM. This asymmetry exists because the permittivity $\epsilon = n^2$ varies in space but the permeability μ is assumed to be constant. This discontinuity in the gradient for the TM mode forces it to have a lower n_{eff} than that of the TE mode, which means it is less tightly bound so has a longer evanescent decay length. This effect becomes more pronounced as the slab index increases or as optical cut-off is approached (which happens when n_{eff} decreases until it reaches n_s and the mode becomes unbound). This tendency is preserved even as the width is decreased to only a few times the height, as in our scheme.

No analytic solution exists for the general rectangular guide, so we used the finite element method discussed in Section 8.4 to solve for the bound mode n_{eff} values and fields as a function of guide dimensions. Figure 8.2 shows the resulting 'cut-off curves', that is, contours of constant n_{eff} in the parameter space (W, H). In this example we chose a guide index $n_g = 1.56$ (typical for a polymer dielectric) and, as a preliminary case, a substrate index $n_s = 1$.

The single-mode region, in which we wish to remain, is bounded below by the E_{11}^y and E_{11}^x curves and above by the E_{21}^y curve. Note that, as in any dielectric guiding

[3]Following the notation of Fernandez & Lu[139] for rectangular guides, mode E_{nm}^k has $n(m)$ antinodes in the $x(y)$ direction, and overall polarization axis $k = x, y$. However as the symmetry line $W = H$ is approached (see Figure 8.2), mixing of polarizations occurs for the case $n \neq m$ and the notation becomes more arbitrary.

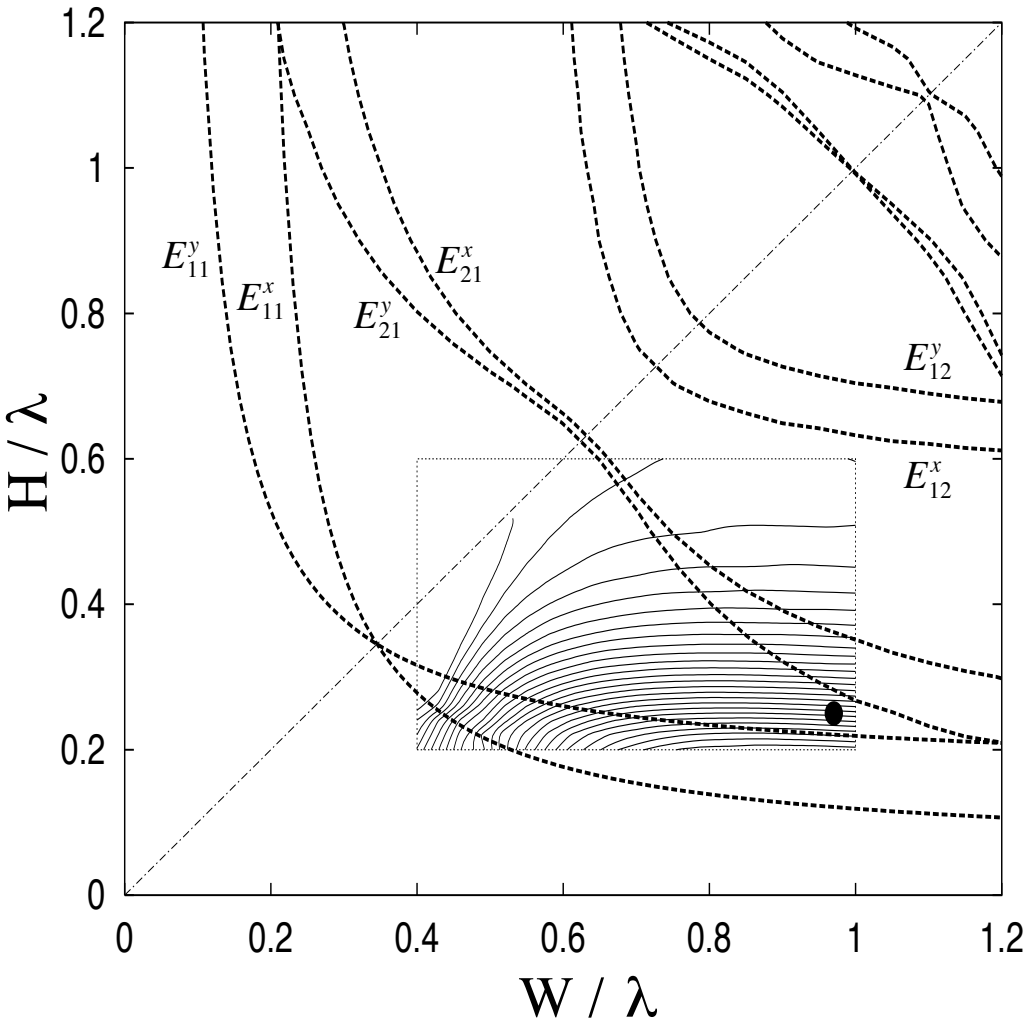

Figure 8.2: Numerically solved cut-off curves for a dielectric optical waveguide of $n_g = 1.56$ (with $n_s = 1$) as a function of its width and height (dashed curves) and, on the same axes, contours of the maximum trapping potential depth achievable at fixed total laser power (thin solid curves in rectangular overlayed box region). Also shown is the symmetry line $W = H$ (thin dash-dotted line). Note that the contours show depth *increasing* as H decreases, almost independent of W, near the suggested operating dimensions (shown as a solid ellipse). Cut-off is defined as reaching an effective refractive index $n_{\text{eff}} \equiv k_z/k_o = 1.05$, except for E_{21}^x and E_{21}^y (thick dashed lines) which we show cut-off at $n_{\text{eff}} = 1.02$.

structure uniform in the z axis, the lowest two modes (E_{11}^y and E_{11}^x in our case) *never* truly reach cut-off, rather, they approach it exponentially as the guide cross-section is shrunk to zero. For this reason, we chose the practical definition of cut-off for these modes to be $n_{\text{eff}} = 1.05$, which corresponds to only about 20% of the power being carried inside the guide. In contrast, higher modes do have true cut-offs[183, 117] (this distinction is illustrated by the dispersion curves of Figure 8.4), and for the E_{21}^y mode our (numerically limited) contour choice of $n_{\text{eff}} = 1.02$ falls very close to the true cut-off curve.

Using the numerically calculated electric field strengths of the E_{11}^y and E_{11}^x modes, we found the red and blue guided laser powers which gave the deepest trap, subject to the constraint of fixed total power (keeping the detuning constant [4]). We also imposed the restriction that the zero of trapping potential come no closer than 100 nm along the vertical line $x = 0$, which kept the trap minimum a reasonable distance from the surface (see Section 8.5.1). Performing this optimization over a region of the parameter space covering the single-mode region gave a contour plot of maximum achievable depth for a given total power, shown within the rectangle overlayed on Figure 8.2. This depth increases from negligible values in the top left to the largest depths in the lower right, indicating that choosing W and H to be in this latter corner of the single-mode region is best for depth. The depth shows very little variation with W in this corner, rather it is clear that varying H to stay within our definition of the single-mode region has become the limiting factor on achievable depth. We indicate a practical choice of $W = 0.97\,\lambda$ and $H = 0.25\,\lambda$ as a small marker on Figure 8.2. Example trapping potentials shapes possible with these parameters are shown in Figure 8.3; we discuss their properties in Section 8.3.

In Figure 8.1a the direct excitation of the optical guide by the two laser beams is shown only schematically. In a realistic experimental setup this coupling into the guide would happen on the order of a centimeter from the atom guiding region, and could involve tapered or Bragg couplers[117] from beams or from other fibers. At this distance we estimate that isotropic stray light due to an insertion loss of 0.5 would have 8 orders of magnitude less intensity than the EW fields in the guiding region. Assuming the light is coherent, this limits the fractional modulation of the guiding potential to 10^{-4}. More improvements are possible, including the use of absorbing shields, bending the guide through large angles away from the original coupling direction, and reducing the coherence length.

8.2.3 Discussion of optical cut-off and substrate choice

We have shown in Figure 8.2 that parameters optimized for trap depth are near optical cut-off. It is worth gathering together the physical reasons for this. Firstly, on general geometric grounds, the typical available intensities in a guide scale inversely with the effective cross-sectional area of the bound mode, which, far from cut-off (that is, when $n_g - n_{\text{eff}} \ll n_g - n_s$) follows very closely the cross-sectional area of the guide. This makes it favorable to shrink guiding structures to areas less than a square wavelength, where they generally become single-mode. Secondly, once we're in the single-mode regime, as we approach cut-off the mode power is carried increasingly outside the guide, increasing the

[4] We made the assumption that the positive and negative detunings caused negligible shifts in the wavelength used for the bound mode calculation. For the largest detunings used in this chapter the fractional shift is 1.8%.

ratio of surface intensity to guide center intensity. Thirdly, the evanescent decay length in the vacuum is longer as we approach cut-off (recall that in the case of the slab, this is exactly expressed by $L^{-1} = 2k_0\sqrt{n_{\text{eff}}^2 - 1}$ [117], where the factor of 2 arises because we are considering intensity decay length rather than amplitude). In the special case of $n_s = 1$, the decay length diverges to infinity as we approach cut-off. Finally, the ratio α_L becomes larger as we approach cut-off, with corresponding beneficial effects on depth and coherence (due to increasing the 'goodness factor' we will introduce in Section 8.3.1).

Unfortunately, these purely theoretical reasons for approaching cut-off are in opposition to more practical ones. The closer to cut-off a guide is, the more sensitive it is to manufacturing variations in cross-section: in our case this will be predominantly a sensitivity to H. The result is that small variations in H cause large variations in mode size, or, at worst, complete cut-off. If the mode size change is rapid (nonadiabatic) along the z axis, (for instance if this change is due to surface roughness or refractive index inhomogeneities) the resulting mismatches will be a source of scattering of the guided power. Any coherent scattering back down the guide will set up periodic modulations of the light field over long distances. (One way to reduce the distance over which coherent addition is possible is to use very broad line-width light sources, which would dramatically reduce modulations due to both guided and stray scattered light). The consequence for the atoms would be a z-dependent trap depth and shape, and this could lead to partial reflection or even localization of the matter waves. In general we expect scattering to limit how close to cut-off we can reliably operate.

With regard to the substrate, further practical issues arise. In the above cut-off calculation we chose the simplest case of $n_s = 1$, corresponding to a guide surrounded by vacuum. A real substrate with $n_s > 1$ has the unfortunate effect of limiting the propagation constant k_z of strictly bound modes to be larger than the freely propagating wavevector in the substrate; in other words, $n_{\text{eff}} > n_s$ must hold or the light field will rapidly tunnel into the 'attractive potential' of the substrate. This in turn limits the decay lengths and α_L that can be achieved. The trap properties quoted in the abstract and in Sections 8.3.1 and 8.3.2 rely on very low n_{eff} values (1.07 for the E_{11}^y mode, 1.18 for E_{11}^x) for the reason that a low n_{eff} is the only way to create long evanescent decay lengths in the vacuum. (This is equivalent to Ovchinnikov *et al.* choosing reflection angles very close to critical[156].) However, in Section 8.3.3 we present preliminary results for a substrate of sodium fluoride (the lowest-index common optical mineral, at $n = 1.32$), and do not believe the substrate alters the basic feasibility of our waveguide.

For completeness, here we list some other possible approaches to the substrate issue. 1) Use an aerogel substrate, which can have exceptionally low refractive indices and low loss (films of several μm thickness with indices of about 1.1 can be produced[147, 150]). 2) Use a dielectric multilayer substrate with an effective index of unity or less (very low loss multilayer mirrors[169] with effective indices less than unity can be created). 3) Investigate if there exist guide shapes which have sufficiently small tunnelling rate into a conventional substrate that the fact that the modes are not strictly bound becomes irrelevant (for instance, a wedge shape with the smallest face in contact with the substrate). 4) Unsupported guiding structures could be produced over short distances[176]. Finally, it is important to note that the idea of replacing the substrate by a metallic reflective layer is

not practical because they are too lossy.

Ultimately, the best values of W and H, the best guide cross-sectional shape, and the substrate choice will depend on many of the above factors and is an area for further research.

8.3 Trap properties

In the bulk of this Section we will examine the atomic waveguide properties for light nearly resonant with the D2 line of cesium, using an optical guide of index 1.56 of the dimensions $W = 0.97\,\lambda$ and $H = 0.25\,\lambda$ from Section 8.2.2, and a substrate of unity index. The saturation intensity for cesium is $11.2\,\mathrm{W/m^2}$ [2], and its resonant wavelength of 852 nm requires that the physical guide size is $0.83\,\mu$m by $0.21\,\mu$m. (At the end of the Section we present preliminary calculations for a $n_s = 1.32$ substrate and a different guide, and discuss how the atom waveguide properties are changed).

Given the guide, we are free to choose three experimental parameters, namely the optical powers carried in the two modes, and the detuning Δ (assumed to be symmetric, that is, to be of equal magnitude for red and blue beams, because little advantage can be gained with an unsymmetric detuning). The first two of these can usefully be reexpressed as total power $P_{\mathrm{tot}} \equiv P_{\mathrm{red}} + P_{\mathrm{blue}}$, and the power ratio $p \equiv P_{\mathrm{red}}/P_{\mathrm{blue}}$. The trap shape will be affected by p alone: we show the trapping potential shapes achievable at the two practical extremes of $p = 0.4$ and $p = 0.2$ in Figure 8.3, where we have chosen P_{tot} and Δ to give identical trap depths and coherence times. Smaller p values cause the trap minimum to move further from the surface (a distinct advantage), to be less "bean" shaped (*i.e.* to have smaller cubic deviations from a 2D harmonic oscillator), and to cause a slight increase in collection area. It is possible to achieve a trap minimum as distant as $y_0 = 0.52\,\lambda$ from the surface when $p = 0.2$. The only disadvantage to implementing these smaller p values is that a higher P_{tot} is required to achieve the same trap depth and coherence time (for instance a factor of 7.5 increase is required as we take p from 0.4 to 0.2). This can be quantified within the exponential approximation, and it can be found that the total power required to maintain a given depth and coherence time with a fixed trap geometry scales as $P_{\mathrm{tot}} \sim (1+p)/p^{1+1/\alpha_L}$.

If we were purely interested in maximizing trap depth at a given P_{tot} and detuning, it would be best to make p as large as possible, however if we take p much larger than 0.4 the trap is brought so close that the corners of the "bean" shape touch the dielectric surface (see Figure 8.3, upper plot) and we will lose effective collection area due to sticking of atoms onto this surface.

8.3.1 Depth, coherence time, and Q factor

We may ask what trade-offs are necessary between trap depth and coherence time. It turns out that, within the exponential approximation (8.6), this is elegantly quantifyable. We can define a 'goodness factor'

$$G \equiv \frac{s_{\mathrm{blue}}(\mathbf{r_0}) - s_{\mathrm{red}}(\mathbf{r_0})}{s_{\mathrm{blue}}(\mathbf{r_0}) + s_{\mathrm{red}}(\mathbf{r_0})} = \frac{\Gamma}{\hbar|\Delta|}\, U_{\max}\tau_{\mathrm{coh}}, \tag{8.8}$$

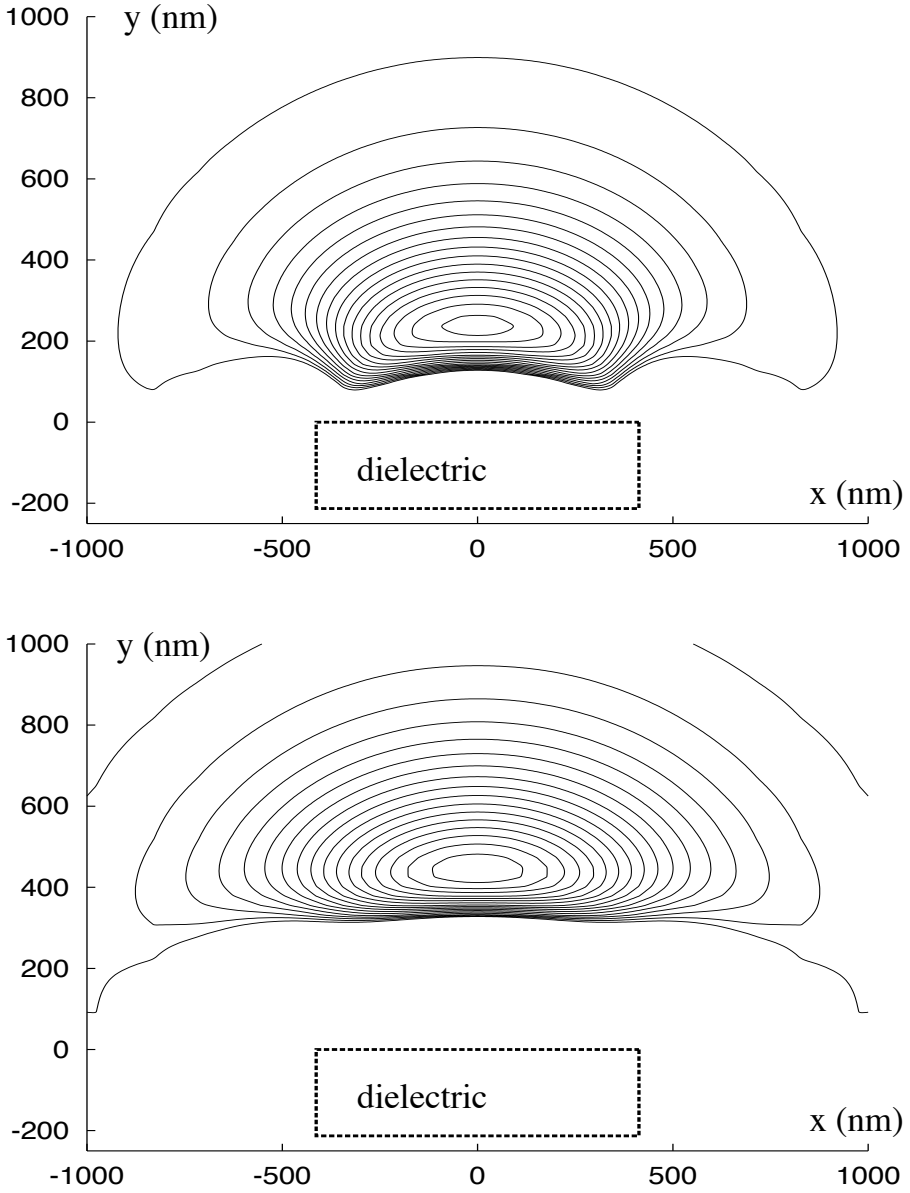

Figure 8.3: Trapping potential contours in the xy plane, with guide dimensions of 0.83 μm by 0.21 μm and an index of 1.56. In each case the maximum depth is 100 μK and the coherence time of atoms in the ground state is 19 ms, achieved with detuning $\pm$15 nm from the cesium D2 line. The outer contour shows a depth of 3 μK, the cesium MOT temperature. Subsequent contours are spaced by 6 μK. The plots illustrate the range of trapping distances acheivable: the upper trap ($p = 0.4$, using a total guided power of 1 mW) has a minimum 0.24 μm from the surface; the lower trap ($p = 0.2$, total guided power of 7.5 mW) has a minimum 0.44 μm from the surface.

where the trap minimum position $\mathbf{r_0}$ is at $(x = 0, y = y_0)$, and the second equality is verified by substitution of (8.4) and (8.5), and defining $U_{\max} \equiv |U_{\rm dip}(\mathbf{r_0})|$ and $\tau_{\rm coh} \equiv \Gamma_{\rm scatt}^{-1}(\mathbf{r_0})$. We use this latter definition because we are interested in the coherence time of atoms spending time close to the trap minimum (which will certainly be true for the transverse ground state.) Using (8.6) to solve for y_0 and evaluate the 'goodness factor', it turns out that the factor is independent of either laser power (*i.e.* of either $A_{\rm red}$ or $A_{\rm blue}$), giving

$$G = \frac{L_{\rm red} - L_{\rm blue}}{L_{\rm red} + L_{\rm blue}} = \frac{\alpha_L}{2 + \alpha_L}\,. \tag{8.9}$$

Combining (8.8) and (8.9) gives

$$U_{\max}\tau_{\rm coh} = \frac{\alpha_L}{2 + \alpha_L}\,\frac{\hbar|\Delta|}{\Gamma}\,, \tag{8.10}$$

fixing the product of achievable depth and coherence time as a constant multiple of the detuning. This is a remarkable result since it shows that increasing α_L is really *the* only objective in the field design of two-color EW traps. We can write this in units more convenient for cesium trap design, thus,

$$\begin{aligned} \frac{U_{\max}}{\mu\rm K}\cdot\frac{\tau_{\rm coh}}{\rm ms} &= (644.2)\frac{\alpha_L}{2+\alpha_L}\cdot\frac{|\Delta|}{\rm nm} \\ &= (122)\cdot\frac{|\Delta|}{\rm nm}\,, \end{aligned} \tag{8.11}$$

where the value $\alpha_L = 0.47 \pm 0.02$ (taken from best-fit exponentials to the numerically-found squared electric fields for the guide dimensions of Section 8.2.2 with $n_s = 1$) has been substituted to give the the final form. This design expression does not give the $P_{\rm tot}$ required to reach a desired balance between $U_{\max}$ and $\tau_{\rm coh}$, however, the total laser power is usually in the mW range, several orders of magnitude less than in most free-space trap designs. For instance, with $P_{\rm tot} = 20\,\rm mW$, $p = 0.4$ and $\Delta = \pm 15\,\rm nm$ we could generate a trap of 2 mK depth with the relatively short coherence time of 0.9 ms. The transverse oscillation frequencies in this trap would be $\omega_x/2\pi = 116\,\rm kHz$ and $\omega_y/2\pi = 490\,\rm kHz$ (the field shapes fix this ratio at about 1:4), giving an atomic mode spacing due to the x motion of 5.6 μK, roughly twice the cesium MOT temperature, and a characteristic ground-state size of 26 nm by 12 nm.

For coherent guiding, we can define a more physically meaningful figure of merit, $Q \equiv \omega_\perp \tau_{\rm coh}$, which tells us the typical number of coherent transverse oscillations we can expect multiplied by 2π (*i.e.* it is the Q-factor of the transverse oscillations). We should choose $\omega_\perp = \omega_x$ since this is the smaller of the transverse frequencies in our case. For $Q \gg 1$ the transverse atomic modes will be well resolved, and our guide can be a useful interferometric device. Using (8.10), in conjunction with the fact that when the trapping potential shape is fixed then ω_x is proportional to the square root of the depth, tells us that for a given trap and detuning, $Q \propto 1/\omega_x$. For a higher Q we should choose smaller transverse oscillation frequencies, that is, shallower traps. For example, the 2 mK trap discussed above has $Q \approx 650$, but if we reduce it to a 20 μK trap of the same Δ (by changing p or the laser powers), the Q is 10 times larger. Increasing Δ would allow even higher Q to be realized.

The dependence on detuning in (8.11) is another way of expressing the advantages already known about using far off-resonant beams[2, 85]. However, our single-resonance approximation will break down if the detuning is too large: we have (somewhat arbitrarily) chosen a detuning limit of 15 nm, as compared to $\Delta_{\rm fs} = 43\,{\rm nm}$ for cesium. At this limit, the additional dipole potential created due to the detunings from the D1 line is very significant. However, by removing the detuning symmetry (changing Δ_+ from $+15\,{\rm nm}$ to $+12.07\,{\rm nm}$ and Δ_- from -15 nm to -17.14 nm), the desired D2 single-resonance approximation potential is recovered in the true physical situation of both resonances present. (These required shifts, which are of order $\Delta^2/\Delta_{\rm fs}$, can easily be found using the expression for the sum of dipole potentials from the two lines.) An additional necessity for our limit is the fact that any larger detunings start to demand separate bound-mode calculations for the two colors, a treatment we reserve for the future. This detuning limit in turn limits the depths, coherence times and Q-factors we quote here, but we anticipate similar future EW atom waveguide designs which explore the region $|\Delta| > \Delta_{\rm fs}$ (or even $|\Delta| \sim \omega_0$), and achieve much better coherence.

8.3.2 Other properties

We estimate the collection area of the trap as the cross-sectional region within which the potential is deeper than the typical cesium MOT energy $k_B T_{\rm MOT}$. For our example $100\,\mu$K traps of Figure 8.3 this area is about $1\,\mu{\rm m}^2$. However, it is not possible to do much better than this with our design: if one tries to increase the area by increasing the overall trap depth, the $k_B T_{\rm MOT}$ contour touches the substrate, indicating that atoms at this energy can reach the substrate surface, where they will stick, limiting the effective collection area.

To investigate the lifetime of atoms transported incoherently (the multi-mode regime), we can calculate the heating rate along similar lines as Grimm and Weidemüller[85]. We start with their equation (23) which gives the rate of change of the average of the total energy of atomic motion $E = E_{\rm kin} + E_{\rm pot}$ as

$$\dot{\overline{E}} = k_B T_R \, \overline{\Gamma_{\rm scatt}}\,, \tag{8.12}$$

T_R being the recoil temperature, and use the assumption that in an equilibrated 3D trap $\overline{E_{\rm kin}} = \frac{3}{2} k_B T$. Since there is harmonic motion in two directions but free motion in the third, the virial theorem gives us $\overline{E_{\rm pot}} = \frac{2}{3}\overline{E_{\rm kin}}$. Combining this with $\overline{\Gamma_{\rm scatt}} = U_{\rm max}\Gamma/G\hbar\Delta$ from (8.8) gives the heating rate

$$\dot{T} = \frac{2}{5G}\frac{\Gamma}{\Delta} T_R \, \frac{U_{\rm max}}{\hbar}\,, \tag{8.13}$$

which is of the order of one recoil temperature per coherence time. For our $100\,\mu$K depth trap at $\Delta = 15\,{\rm nm}$ the rate is $4.4\,\mu{\rm K\,s^{-1}}$, implying that storage and transport for many seconds is possible. For simplicity, we have ignored the fact that there may be distinct longitudinal and transverse temperatures which do not equilibrate over the trapping timescales.

8.3.3 Effect of a realistic substrate

In this section we present calculations, performed using the method of Section 8.4, for a practical substrate choice of sodium fluoride (the lowest refractive index common mineral, with $n_s = 1.32$ at a wavelength of 852 nm), and investigate how this changes the atom waveguide properties from those presented above. We increased n_g to 1.7 (dense flint glass, e.g. BaSF type) in order to provide sufficient index difference from the substrate.

Fixing the width at $W = 1.00\,\lambda$, we found that a height $H = 0.34\,\lambda$ gave the largest α_L of 0.20 ± 0.01, and allowed both modes to be sufficiently far from cut-off (greater than half the power being carried inside the guide for both modes). The result is a goodness factor G which is approximately half that of the $n_s = 1$ case, with a corresponding halving of the achievable product of depth and coherence time according to (8.11), and doubling of the heating rate at a given $U_{\max}$ and Δ according to (8.13). The shorter decay lengths of 56 nm and 68 nm (compared to 93 nm and 137 nm for $n_s = 1$) cause the typical trapping distance y_0 to be reduced by a factor of roughly 1.8.

We found that in order to reproduce the depth of $100\,\mu$K and $y_0 = 0.24\,\mu$m of the upper trap of Figure 8.3 (with Δ unchanged) we needed $P_{\rm tot} = 22$ mW, giving $\tau_{\rm coh} = 9$ ms. The large power increase over the 1 mW required for $n_s = 1$ is explained by the fact that this y_0 is now towards the upper limit practically achievable rather than the lower. (If y_0 is instead scaled in proportion to the new decay lengths, the required increase in $P_{\rm tot}$ is only a factor 1.7). In this example, we find the transverse oscillation frequencies have increased to $\omega_x/2\pi = 81$ kHz and $\omega_y/2\pi = 202$ kHz, compared to the original $\omega_x/2\pi = 26$ kHz and $\omega_y/2\pi = 109$ kHz. The increase in ω_y is explained entirely by the shorter decay lengths, and the increase in ω_x (by a factor of over 3) is attributed to tighter optical mode shapes. It is clear that this latter effect outweighs the decrease in $\tau_{\rm coh}$, implying that the inclusion of the substrate has actually *increased* Q by 50%.

In summary, the effects of including a realistic substrate limit the maximum trapping distance y_0 that can be achieved (because of the reduction in decay lengths), lower the goodness factor, increase the heating rate and the required optical power, but also increase the oscillation frequencies. For our substrate choice, each of these changes was approximately a factor of 2, and we believe that they do not alter the basic practicality of implementing our proposed waveguide.

8.4 Numerical solution of the light fields

The detailed electric field distribution is necessary for calculating the trapping potential above the waveguide. An approximation for the y-dependence of the fields is given by the analytically-known solution for the slab waveguide; this gives exponential decay, as already discussed. To get knowledge of the actual potential shape in the x-y plane, we used a full-vector finite element calculation.

8.4.1 The vector Finite Element Method

The technique represents the electric and magnetic fields as simple piecewise functions over many "elements" (regions) subdividing a slice through the guide and surrounding

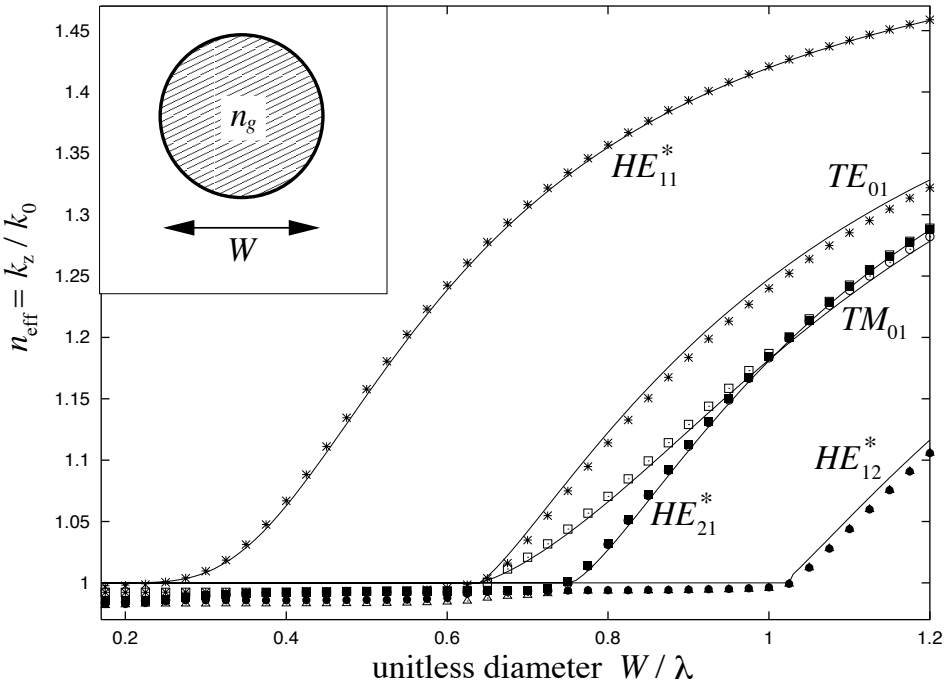

Figure 8.4: Comparison of our optical guide bound-mode numerical implementation against known analytic solutions, in the case of a free-standing dielectric cylinder of $n_g = 1.56$. We used discretization and box-size identical to the rectangular guide case, and observe typical errors of $\pm 0.5\%$ in propagation constant for the first two modes. The mode naming convention and analytic calculation follow Snitzer [183]; an asterisk indicates a doubly-degenerate mode.

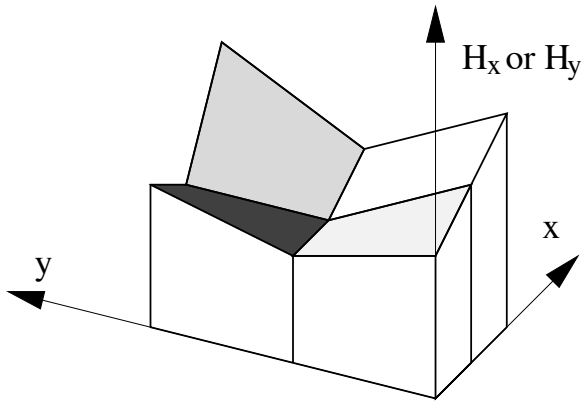

Figure 8.5: A transverse magnetic field component is represented by a patchwork of bilinear (*i.e.* lowest-order) basis functions, each living a rectangle, covering the $x-y$ plane.

media in the xy plane, therefore by a finite number of degrees of freedom (see Fig. 8.4.1). Each element has a dielectric constant associated with it, allowing arbitrary stepwise refractive index distributions in the xy plane to be modelled. Maxwell's equations for propagating solutions of the form $\exp(ik_z z - i\omega t)$ are reduced to a generalized sparse eigenvalue equation with k_z^2 as the eigenvalue and the bound mode field distributions as the eigenvectors. For introductory reviews of electromagnetic waveguide bound-mode techniques see Silvester and Ferari [182] and Davies [56].

Specifically, we used the technique of Fernandez and Lu [139], with H_x and H_y as the field degrees of freedom, for simplicity using first-order (bilinear) functions to represent these fields over each element in a two-dimensional grid. To reach a compromise between flexibility and ease of coding, the grid of elements was chosen to be *non-uniform* (in order to allow high densities in the regions where fields changed rapidly, and low densities in others), but rectangular and separable into a product of elements in the x and y directions. See Fig. 8.4.1 which shows an example grid used for calculation.

This required a generalization of the Fernandez and Lu implementation, and careful consideration of their line-integral terms (which are non-standard for a finite element formulation). Of the many available finite element approaches to dielectric waveguide mode solution, this frequency-domain method was chosen for its absence of 'spurious modes' (unphysical numerical solutions), its ability to handle index step discontinuities, its small number of required degrees of freedom, and its matrix sparsity. It does not perform as well as some more complicated methods which represent more field degrees of freedom (for a comparison of some current methods, see [178]).

Rather than emulating a radiative boundary condition (a notoriously hard task usually requiring an additional outer iterative loop, due to the k_z dependence of the boundary condition), we enclosed the problem in a large, perfectly-conducting box of sufficient size that the bound mode evanescent fields were negligible on its walls, rendering the exact nature of the boundary condition irrelevant. However, the average level spacing of the unbound modes (the 'continuum') decreases with increasing box size, and especially near cut-off we found that this increases the number of iterations required to solve the eigenvalue problem to a given accuracy. Thus we have a trade-off, and found that a box size of 6λ to 7λ gave the best compromise between accuracy and speed.

We used the well-known ARPACK nonsymmeteric sparse eigenvalue solver [133] (essentially a block Lanczos [81, 125] iterative solver) to find the 11 lowest eigenmodes of the sparse matrix. Of these, only the lowest 4 were needed to identify the single-mode region in Fig. 8.2, and only the lowest 2 eigenvectors were needed for their electric field information. Our non-uniform elements allowed us to have a high element density across the waveguide and in the adjacent trapping region, but a low density over the much larger box area, keeping the total number of degrees of freedom manageable.

8.4.2 Tests of accuracy and convergence

For the rectangular guide, the fractional error ϵ in the propagation constant k_z was less than 1%, and the accuracy of the electric field strengths in the trapping region $\approx 3\%$, when we used $N \sim 5000$ degrees of freedom. This was sufficiently accurate for the present work.

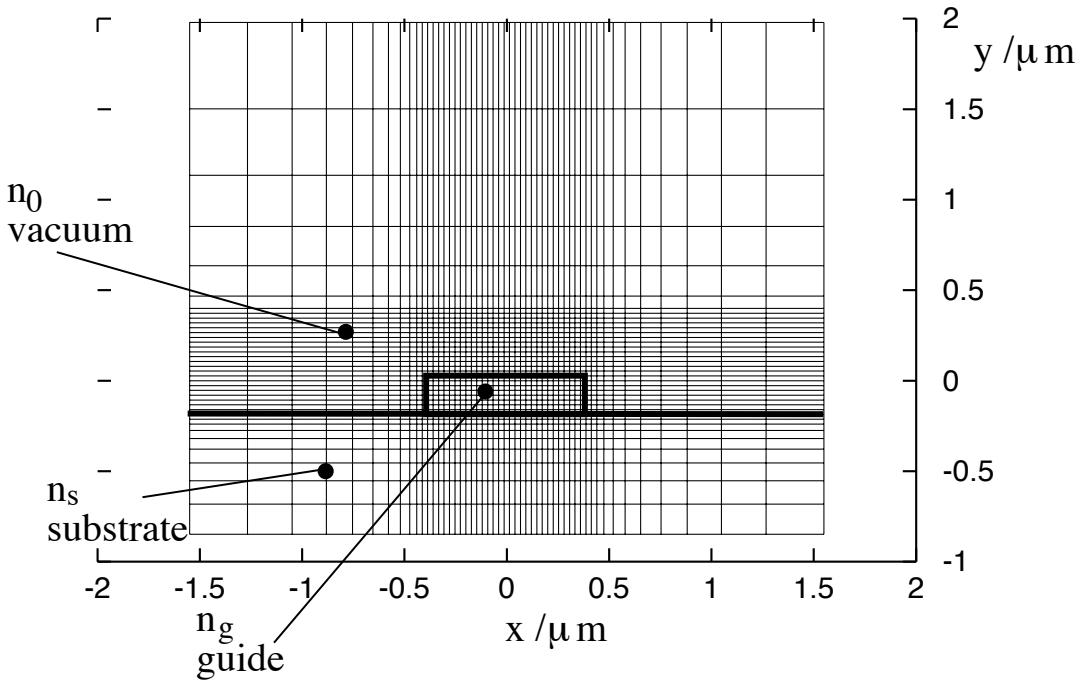

Figure 8.6: Finite elements in the $x-y$ plane: an example rectangular, non-uniform but separable grid used for calculation. The grid fills a box, which is chosen to be large enough that the fields on its boundary are negligible.

We estimated the accuracy of the method by solving a cylindrical guide in an identical fashion with the same N and a very similar non-uniform grid, for which there are known field solutions [183]. This system is quite similar to the rectangular guide in question, although it does not possess sharp corners. However, representing the circular cross-section by a pixellated approximation on a square grid (which cannot represent surfaces other than horizontal and vertical) was taken to be a stringent test of the method's ability to handle corners sensibly.

We performed a sweep in radius of the guide from below cut-off to where there are several bound modes: Fig. 8.4 shows the propagation constant agrees with the analytics to within 1%, even close to cut-off. The convergence with N was measured for the rectangular guide case, and found to be $\epsilon \sim N^{-\gamma}$ with $0.55 < \gamma < 0.7$. This is less than optimal for first-order elements (which have a maximum possible convergence of $\gamma = 1$), and is believed to be due to an inability of the bilinear functions to represent physical in-plane E and H components at dielectric steps, or the weak field singularities which can physically occur at any exterior dielectric sharp edges (regardless of whether acute or obtuse).

8.4.3 Results

Resulting plots of the balanced intensity differences of the two lowest modes were shown in Fig. 8.3. Finding the bound modes for each waveguide parameter choice, with $N \approx 5000$, typically took between 3 and 20 minutes of computing time on a desktop workstation[5]. This was dominated by the sparse eigenvalue diagonalization time. The convergence time of the eigenvalues for the same N could vary by an order of magnitude, and was largely dependent on how close to cut-off the guide was: closer to cut-off took longer. Little attempt was made to optimize this.

Future improvements to the method, which would increase the accuracy or the convergence rate γ, include using higher order elements (if done carefully, this could correctly represent physical E and H components at dielectric steps), and explicit modelling of the field singularities at guide corners using specialized elements.

8.5 Further decoherence and loss mechanisms

8.5.1 Effects of surface interactions

The EW trap has the benefit of creating high field gradients near a surface, but along with this comes the disadvantage that interactions with that surface that can alter the trap potential and even cause heating and loss of trapped atoms.

An atom's change in potential near a surface is known as the van der Waals interaction ($l \ll \lambda$) or the Casimir interaction ($l \gg \lambda$), depending on the distance l from the surface compared to λ, the dominant wavelength responsible for the polarizability of the atom (in our case of Cs this is the D line resonance, the same as our trapping resonance). There is a smooth cross-over from van der Waals ($U \sim l^{-3}$, which can be viewed as the atom's electrostatic interaction with the image of its own fluctuating dipole) to Casimir

[5]Silicon Graphics 75MHz R8000 processor commercial compilers, with a SPEC FP95 rating [186] of between 10 and 12.

($U \sim l^{-4}$, which can be viewed as a retarded van der Waals attraction or equally well as an atomic level shift due to a cavity QED effect) at $l \approx \lambda/10$[96]. In the case of a perfect mirror surface, the full form is known for any l, but for a dielectric surface, the expression becomes much more complicated to evaluate[184].

Since our trapping distances are larger than this cross-over point, we will use the Casimir form, which is correct for asymptotically large l, and is always an overestimate of the true potential[96]. The dependence of the coefficient with dielectric constant is complicated [68, 207], but we will use the simpler approximate form given by Spruch and Tikochinsky[184], to give

$$U_{\mathrm{Cas}}(l) = -\frac{3}{8\pi}\frac{\hbar c\,\alpha(0)}{4\pi\epsilon_0\, l^4}\,\frac{\epsilon - 1}{\epsilon + (30/23)\epsilon^{1/2} + 7/23} \tag{8.14}$$

(in the MKSA system). This approximate form is known to be within 6% of the exact expression for any dielectric constant ϵ [207]. Substituting the recently calculated[60] static polarizability of cesium, $\alpha(0) = 399.9$ a.u., gives a Casimir interaction coefficient of $4.9\,\mathrm{nK}\,\mu\mathrm{m}^4$ for $n_g = 1.56$.

Figure 8.1b shows the effect of this potential on a typical trap of depth $150\,\mu$K and distance $y_0 = 270\,$nm. It is clear that the change is negligible further than 100 nm from the surface, and a WKB tunneling calculation along this straight-line path (at $x = 0$) shows that even if all atoms that reach the surface stick, the loss rate from the first few transverse modes is entirely negligible. However, care should be taken with the multi-mode regime, or in the case of high-p traps, since the tunneling via the corners of the "bean" shape may dominate for $p > 0.4$ (Figure 8.3).

The issue of energy transfer to trapped atoms due to a finite (and possibly room) temperature nearby surface is far less well understood, and may be a problem with many surface-based particle traps, as discussed by Henkel and Wilkens [92]. However, since we are trapping neutral particles and the conductivity of our surface is low, we expect a decoherence rate negligible compared to that already present from spontaneous absorption and emission cycles.

8.5.2 Bending the waveguide

It would be very useful to be able to carry atom beams along curved paths, by bending our atom waveguide in the plane of the substrate, without significant atom loss. Here we briefly estimate three limitations on the waveguide minimum bending radius (in decreasing order of leniency): 1) the limit imposed by optical radiation leakage, 2) the limit needed for incoherent atom transport, and 3) the limit needed for coherent atom transport in the transverse groundstate. This will give us an idea of the practicality of curved atomic guides.

Firstly, whenever a dielectric optical guide has curvature, there is a loss rate (exponentially small in the curvature radius R), which can be viewed as tunneling out of the guide's 'potential well' induced by the addition of an effective centrifugal potential. In the limit $R \gg W$, the effective potential is linear with x (the radial coordinate), and the frac-

tional loss per radian of curvature can be estimated[6], for instance using the one-dimensional WKB formula, to be

$$\alpha = C\,\frac{R}{\lambda}\exp\left(-\frac{1}{6\pi^2}\frac{\lambda^2 R}{L_x^3}\right), \tag{8.15}$$

where L_x is the evanescent decay length in the radial direction, and C is a constant of order unity. Therefore for negligible light loss at a $\pi/2$ bend we need $R > 60\pi^2 L_x^3/\lambda^2$, typically a couple of tens of microns. This is so small chiefly because we are using an optical guide with a large refractive index step[117].

Secondly, we consider atom loss from an incoherent beam with a transverse temperature $k_B T_\perp$ (assumed small compared to the trap depth magnitude $U_{\max}$), and a longitudinal kinetic energy $E_\parallel$. We call the approximate spatial extent of the trap potential in the x direction 2ξ, and restrict ourselves to one-dimensional classical motion in this direction. When in a region of radius of curvature $R \gg \xi$, an effective centrifugal term adds to the trapping potential giving $U(x) = U_{\rm dip}(x) - 2E_\parallel(x/R)$. This causes the atoms to 'slosh' towards positive x, only ever returning if there exists a point where $U(x) > -U_{\max}$ for $x > 0$. We can estimate that this will happen if $U_{\max} > 2E_\parallel(\xi/R)$, giving our lower limit on R as $2\xi E_\parallel/U_{\max}$. In our design $\xi \approx 0.5\,\mu$m, so if we choose $R = 1\,$mm we can expect loss-free transport of a beam at a longitudinal kinetic energy up to 10^3 times the trap depth.

Thirdly, to model coherent matter-wave propagation along a curved guide, we consider the amplitude for remaining in the transverse ground-state, having passed into a curved section and back into a straight section. If again we assume one-dimensional x motion, and assume a harmonic potential $U(x) = \frac{1}{2}M\omega_x^2(x - x_0)^2$ around the trap minimum, then the effect of curvature is to shift the minimum position from $x_0 = 0$ to $x_0 = 2E_\parallel/M\omega_x^2 R$. If this shift is much less than the characteristic ground-state size $(\hbar/M\omega_x)^{1/2}$ then the projection at each transition will be high, resulting in high flux transmission coefficient. This gives $R \gg 2E_\parallel/(\hbar M\omega_x^3)^{1/2}$ as our condition, which for $E_\parallel/k_B = 10\,$mK (that is, $v_\parallel = 1.1\,\mathrm{ms}^{-1}$) and $\omega_x/2\pi = 40\,$kHz corresponds to $R \gg 0.45\,$mm. This limit is very conservative since we have not yet made use of the adiabatic condition $\Omega \ll \omega_x$ (where $\Omega \equiv v_\parallel/R$ is the rate of change of direction of the guided atom), to design a waveguide path without discontinuities in the curvature.

In conclusion, we have shown that it is possible to bend atoms both incoherently and coherently through large angles on a compact substrate structure of a few millimeters in size.

8.6 Conclusion

We have presented a novel substrate-based neutral atom waveguide, based on the optical dipole force, which combines the features of a planar far-detuned two-color evanescent trap[156] with the ability to confine strongly along two axes. We utilized the

[6]For some of the earliest bending loss estimation in a rectangular fiber see [142] and for the case of a cylindrical fiber see [143]. Most of the work done since then has been on the more complicated situation of fibers with cladding.

differing vertical decay lengths of the two bound-mode *polarizations* of a submicron-sized optical waveguide near cut-off. We have shown that only a few milliwatts of guided laser power can create atomic potential depths $\sim 100\,\mu$K with transverse oscillation frequencies $\sim 100\,$kHz, a coherence time $\sim 10\,$ms and a trap minimum 200-400 nm above the optical guide surface, for Cs atoms in the $m_F = 0$ state. Laser powers greater than ten milliwatts can give a transverse mode spacing greater than the temperature of a Cs MOT, opening up the single-mode waveguide regime. The advantages of guiding optical trapping fields on a substrate include mechanical stability and reliability, mass production and the potential for transport along complicated paths.

We have given some design criteria for guided-lightwave two-color atom waveguides (chiefly the maximization of the evanescent decay lengths, and of their normalized difference α_L), and shown that a substrate of low refractive index can be very beneficial. We modelled in detail the trapping potentials for a general rectangular guide of index 1.56 above a unity-index substrate, and have shown that a realistic substrate choice of index 1.32 poses few problems to the viability of the device. We predict that the effect of the surface interaction is generally small, and that coherent guiding is possible around corners of radii $\sim 1\,$mm for a longitudinal velocity $\sim 1\,\mathrm{ms}^{-1}$. We also believe that the magnetic part of the potential felt by nonzero m_F atoms could be used to increase the depth and oscillation frequencies further.

This preliminary work (specifically equation (8.10)) indicates that utilizing detunings much larger than the 15 nm we limit ouselves to here will be very advantageous for coherent guiding. We have only scratched the surface of the design variations possible; for instance, equalizing the horizontal and vertical oscillation frequencies is yet to be attempted. The use of two polarizations is our solution to the problem of maximizing α_L when the detuning is very small compared to the wavelength, but we suspect that there will exist other fruitful schemes where these are comparable ($\Delta \sim \omega_0$, very far-detuning) and where a large α_L is caused simply by the different optical cut-off conditions at the two wavelengths. We have reserved investigation of cooling schemes for future work (although this has already been demonstrated in an EW mirror[61] and proposed in EW traps[157, 62]). We believe that the potential shapes capable of being produced by guided waves on a substrate also include the possibility of funnel-type loading schemes and coherent atom couplers, allowing for a complete "integrated" atom-optical experiment on a substrate.

Bibliography

[1] M. Abramowitz and I.S. Stegun, editors. *Handbook of mathematical functions*. Dover, NY, 1965. see Chapter 20.

[2] C. S. Adams, M. Sigel, and J. Mlynek. *Phys. Rep.*, **240**:143, 1994.

[3] S. K. Adhikari. *Am. J. Phys.*, **54**:362, 1986.

[4] C. G. Aminoff, A. M. Stearne, P. Bouyer, P. Desbiolles, J. Dalibard, and C. Cohen-Tannoudji. *Phys. Rev. Lett.*, **71**:3083, 1993.

[5] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. P.O. Box 7260, Philadelphia, PA 19104, third edition, August 1999. `http://netlib2.cs.utk.edu/lapack/lug/lapack_lug.html`.

[6] E. Arcos, G. Baez, P. A. Cuatlayol, M. L. H. Prian, R. A. Mendez-Sanchez, and H. Hernandez-Saldana. Vibrating soap films: An analog for quantum chaos on billiards. *Am. J. Phys.*, **66**:601, 1998.

[7] G. Arfken. *Mathematical methods for physicists*. Academic Press, 3rd edition, 1985.

[8] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. W. B. Saunders, Philadelphia, 1976.

[9] A. Ashkin. *Phys. Rev. Lett.*, **24**:156, 1970.

[10] A. Ashkin. *Proc Natl. Acad. Sci. USA*, **94**:4853, 1997.

[11] O. M. Auslaender and S. Fishman. *Phys. Rev. Lett.*, **84**:1886–9, 2000.

[12] H. P. Baltes and E. R. Hilf. *Spectra of finite systems*. Bibliographisches Institut, Mannheim, 1976.

[13] V. I. Balykin, V. S. Letokhov, Yu. B. Ovchinnikov, and A. I. Sidorov. *Phys. Rev. Lett.*, **60**:2137, 1988.

[14] A. H. Barnett, D. Cohen, and E. J. Heller. Deformations and dilations of chaotic billiards: dissipation rate, and quasi-orthogonality of the boundary wavefunctions. *Phys. Rev. Lett.*, **85**:1412, 2000. preprint `nlin.CD/0003018`.

[15] A. H. Barnett, D. Cohen, and E. J. Heller. Rate of energy absorption for a driven chaotic cavity. *submitted to J. Phys. A*, June 2000. preprint `nlin.CD/0006041`.

[16] A. H. Barnett, S. P. Smith, M. Olshanii, K. S. Johnson, A. W. Adams, and M. Prentiss. Substrate-based atom waveguide using guided two-color evanescent light fields. *Phys. Rev. A*, **61**:023608, 2000. preprint `physics/9907014`.

[17] D. Baye and P. Heenan. *J. Phys. A*, **19**:2041, 1985.

[18] J. E. Bayfield. *Quantum Evolution: an introduction to time-dependent quantum mechanics*. Wiley, 1999.

[19] E. B. Becker, G. F. Carey, and J. T. Oden. *Finite Elements*. Prentice-Hall, Englewood Cliffs, N.J., 1981.

[20] C. W. J. Beenakker and H. van Houten. Quantum transport in semiconductor nanostructures. *Solid State Physics*, **44**:1–228, 1991.

[21] L. L. Beranek. Wallace Clement Sabine and acoustics. *Physics Today*, page 44, February 1985.

[22] T. Bergeman, G. Erez, and H. J. Metcalf. *Phys. Rev. A*, **35**:1535, 1987.

[23] M. V. Berry. Quantizing a classically ergodic system: Sinai billiard and the KKR method. *Ann. Phys.*, **131**:163–216, 1981.

[24] M. V. Berry. Semiclassical theory of spectral rigidity. *Proc. Roy. Soc. London A*, **400**:229, 1985.

[25] M. V. Berry. Some quantum-to-classical asymptotics. In M. J. Giannoni, A. Voros, and J. Zinn-Justin, editors, *Proceedings of the 1989 Les Houches Summer School on "Chaos and Quantum Physics"*, pages 251–303, North-Holland, 1991. Elsevier Science Publishers B.V.

[26] M. V. Berry. *J. Phys. A*, **27**:L391–8, 1994.

[27] M. V. Berry and J. M. Robbins. Chaotic classical and half-classical adiabatic reactions: geometric magnetism and deterministic friction. *Proc. R. Soc. Lond.*, **442**:659–672, 1993.

[28] M. V. Berry and M. Wilkinson. Diabolical points in the spectra of triangles. *Proc. Roy. Soc. London A*, **392**:15, 1984.

[29] J. Blocki, Y. Boneh, J. R. Nix, J. Randrup, M. Robei, A.J. Sierk, and W. J. Swiatecki. One-body dissipation and the super-viscidity of nuclei. *Ann. Phys.*, 113:330–386, 1977.

[30] J. Blocki, F. Brut, and W. J. Swiatecki. *Nucl. Phys. A*, **554**:107, 1993.

[31] J. Blocki, Y.-J. Shi, and W. J. Swiatecki. *Nucl. Phys. A*, **554**:387, 1993.

[32] J. Blocki, J. Skalski, and W. J. Swiatecki. *Nucl. Phys. A*, **594**:137, 1995.

[33] P. A. Boasman. *Nonlinearity*, **7**:485, 1994.

[34] E. B. Bogomolny. *Nonlinearity*, **5**:805–866, 1992.

[35] O. Bohigas. Random matrix theories and chaotic dynamics. In M. J. Giannoni, A. Voros, and J. Zinn-Justin, editors, *Proceedings of the 1989 Les Houches Summer School on "Chaos and Quantum Physics"*, pages 87–199, North-Holland, 1991. Elsevier Science Publishers B.V.

[36] C. A. Brebbia, editor. *Progress in Boundary Element Methods*, volume 1. Halstead Press, 1981.

[37] R. Brown, E. Ott, and C. Grebogi. *J. Stat. Phys.*, **49**:511, 1987.

[38] G. W. Bryant. *Phys. Rev. B*, **39**:3145, 1989.

[39] A. Bulgac and P. Magierski. Eigenstates for billiards of arbitrary shapes. preprint `physics/9902057`, 1999.

[40] L. A. Bunimovich. *Sov. Phys. JETP*, **62**:842, 1985.

[41] M. Büttiker. *Phys. Rev. Lett.*, **57**:1761, 1986.

[42] by the staff of Engineering Research Associates, Inc. *High-speed computing devices.* McGraw-Hill, 1950.

[43] Y. S. Chan. *Theories and Applications of Multiple Scattering of S-wave Scatterers.* PhD thesis, Harvard University, Cambridge, MA, September 1997.

[44] T. Cheon and T. Shigehara. *Phys. Rev. Lett.*, **82**:2536, 1999.

[45] S. Chu. *Rev. Mod. Phys.*, **70**:685, 1998.

[46] D. Cohen. Chaos and energy spreading for time-dependent hamiltonians, and the various regimes in the theory of quantum dissipation. *Ann. Phys.*, 1999.

[47] D. Cohen. Quantum dissipation due to the interaction with chaotic degrees of freedom and the correspondence principle. *Phys. Rev. Lett.*, **82**:4951–5, 1999.

[48] D. Cohen, A. H. Barnett, W. Bies, and E. J. Heller. Parametric evolution for a deformed cavity. *submitted to Phys. Rev. E*, August 2000. preprint `nlin.CD/0008040`.

[49] D. Cohen and T. Kottos. Parametric dependent hamiltonians, wavefunctions, random-matrix-theory, and quantal-classical correspondence. preprint `nlin.CD/0001026`, 2000.

[50] C. Cohen-Tannoudji. Fundamental systems in quantum optics. In J. Dalibard, J.-M. Raimond, and J. Zinn-Justin, editors, *Proceedings of the 1990 Les Houches Summer School*. North-Holland, Amsterdam, 1992.

[51] C. Cohen-Tannoudji. *Rev. Mod. Phys.*, **70**:707, 1998.

[52] R. J. Cook and R. K. Hill. *Opt. Comm.*, **43**:258, 1982.

[53] P. Dahlqvist. *Phys. Rev. E*, **60**:6639, 1999.

[54] P. Dahlqvist and R. Artuso. *Phys. Lett.*, **219A**:212, 1996.

[55] S. Datta. *Electronic Transport in Mesoscopic Systems*. Cambridge University Press, NY, 1995.

[56] J. B. Davies. Finite element analysis of waveguides and cavities – a review. *IEEE Transactions on Magnetics*, **29**(2):1578, 1993.

[57] M. J. Davis and E. J. Heller. *J. Chem. Phys.*, **71**:4759, 1979.

[58] P. Davis and P. Rabinowitz. *Methods of Numerical Integration*. Academic Press, Orlando, FL, 2nd edition, 1984.

[59] J. Denschlag, D. Cassettari, and J. Schmiedmayer. *Phys. Rev. Lett.*, **82**:2014, 1999.

[60] A. Derevianko, W. R. Johnson, M. S. Safronova, and J. F. Babb. *Phys. Rev. Lett.*, **82**:3589, 1999.

[61] P. Desbiolles, M. Arndt, P. Szriftgiser, and J. Dalibard. *Phys. Rev. A*, **54**:4292, 1996.

[62] P. Desbiolles and J. Dalibard. *Opt. Comm.*, **132**:540, 1996.

[63] B. Dietz, J.-P. Eckmann, C.-A. Pillet, U. Smilansky, and I. Ussishkin. *Phys. Rev. E*, **51**:4222, 1995.

[64] B. Dietz and U. Smilansky. *Chaos*, **3**:581, 1993.

[65] T. Dittrich, P. Hänggi, G.-L. Ingold, B. Kramer, G. Schön, and W. Zwerger. *Quantum Transport and Dissipation*. Wiley-VCH, Weinheim, 1998.

[66] E. Doron and U. Smilansky. Semiclassical quantization of chaotic billiards — a scattering theory approach. *Nonlinearity*, **5**:1055–1084, 1992.

[67] J. P. Dowling and J. Gea-Banacloche. *Adv. At. Mol. Opt. Phys.*, **37**:1, 1996.

[68] I. E. Dzyaloshinskii, E. M. Lifshitz, and L. P. Pitaevskii. *Adv. Phys.*, **10**:165, 1961.

[69] C. Ellegaard, T. Guhr, K. Lindemann, H. Q. Lorensen, J. Nygard, and M. Oxborrow. *Phys. Rev. Lett.*, **75**:1546–9, 1995.

[70] B. D. Esry, C. D. Lin, and C. H. Greene. *Phys. Rev. A*, **54**:394, 1996.

[71] T. Esslingera, M. Weidemüller, A. Hemmerich, and T. W. Hänsch. *Opt. Lett.*, **18**:450, 1993.

[72] M. Feingold, D. Leitner, and M. Wilkinson. *Phys. Rev. Lett.*, **66**:986, 1991.

[73] M. Feingold and A. Peres. *Phys. Rev. A*, **34**:591, 1986.

[74] D. S. Fisher and P. A. Lee. *Phys. Rev. B*, **23**:6851, 1981.

[75] S. Fishman. In G. Casati, I. Guarneri, and U. Smilansky, editors, *Proceedings of the International School of Physics "Enrico Fermi", Course CXIX*. North-Holland, 1993.

[76] R. Folman, P. Krüger, D. Cassettaria, B. Hessmo, T. Maier, and J. Schmiedmayer. *Phys. Rev. Lett.*, **84**:4749, 2000.

[77] B. Friedman. *Principles and techniques of applied mathematics*. John Wiley & Sons, Inc., N.Y., 1956.

[78] B. Friedman and Jr. R. F. Martin. *Phys. Lett.*, **105A**:23, 1984.

[79] C. V. Gardiner. *Handbook of Stochastic Methods*. Springer-Verlag, 1983.

[80] P. E. Gill, W. Murray, and M. H. Wright. *Numerical linear algebra and optimization*, volume 1. Addison-Wesley, 1991.

[81] G. H. Golub and C. F. Van Loan. *Matrix computations*. Academic Press, 2nd edition, 1989.

[82] J. P. Gordon and A. Ashkin. *Phys. Rev. A*, **21**:1606, 1980.

[83] C. Grebogi, S. M. Hammel, J. A. Yorke, and T. Sauer. *Phys. Rev. Lett.*, **65**:1527–1530, 1990.

[84] D. A. Greenwood. The Boltzmann equation in the theory of electrical conduction in metals. *Proc. Phys. Soc.*, **71**:585–598, 1958.

[85] R. Grimm, M. Weidemüller, and Yu. B. Ovchinnikov. *Adv. At. Mol. Opt. Phys.*, **42**:95, 2000. preprint `physics/9902072`.

[86] G. C. Groenenboom and D. T. Colbert. *J. Chem. Phys.*, **99**:315, 1993.

[87] R. Gruber, editor. *Finite Elements in Physics*, Amsterdam, 1987. North-Holland. Proceedings of the the 1st European Graduate Summer Course on Computational Physics, September 1986.

[88] M. C. Gutzwiller. *Chaos in Classical and Quantum Mechanics*. Springer-Verlag, NY, 1990. see Chapter 7.

[89] S. M. Hammel, J. A. Yorke, and C. Grebogi. *J. Complexity*, **3**:136, 1987.

[90] E. J. Heller. Bound state eigenfunctions of classically chaotic hamiltonian systems: Scars of periodic orbits. *Phys. Rev. Lett.*, **53**:1515, 1984.

[91] E. J. Heller. Semiclassical wave packet dynamics and chaos in quantum mechanics. In M. J. Giannoni, A. Voros, and J. Zinn-Justin, editors, *Proceedings of the 1989 Les Houches Summer School on "Chaos and Quantum Physics"*, pages 547–663, North-Holland, 1991. Elsevier Science Publishers B.V.

[92] C. Henkel and M. Wilkens. *Europhys. Lett.*, **47**:414–420, 1999.

[93] J. S. Hersch. *Scattering in the extreme quantum limit*. PhD thesis, Harvard University, Cambridge, MA, September 1999.

[94] J. S. Hersch, M. R. Haggerty, and E. J. Heller. *Phys. Rev. Lett.*, **83**:5342, 1999.

[95] J. S. Hersch, M. R. Haggerty, and E. J. Heller. *Phys. Rev. E (to appear)*, 2000.

[96] E. A. Hinds and V. Sandoghdar. *Phys. Rev. A*, **43**:398, 1991.

[97] E. A. Hindsa, M. G. Boshier, and I. G. Hughes. *Phys. Rev. Lett.*, **80**:645, 1998.

[98] K. Hornberger and U. Smilansky. *J. Phys. A*, **33**:2829, 2000.

[99] Y. Imry. *Introduction to Mesoscopic Systems*. Oxford University Press, 1997.

[100] M. Itagaki, S. Nishiyama, S. Tomioka, T. Enoto, and N. Sahashi. A new complex-valued formulation and eigenvalue analysis of the Helmholtz equation by boundary element method. *Adv. Eng. Software*, **26**:219–227, 1996.

[101] M. Itagaki, S. Nishiyama, S. Tomioka, T. Enoto, and N. Sahashi. Power iterative multiple reciprocity boundary element method for solving three-dimensional Helmholtz eigenvalue problems. *Eng. Anal. with Boundary Elements*, **20**:113–121, 1997.

[102] H. Ito, T. Nakata, K. Sakaki, and M. Ohtsu. *Phys. Rev. Lett.*, **76**:4500, 1996.

[103] H. Ito, K. Sakaki, M. Ohtsu, and W. Jhe. *App. Phys. Lett.*, **70**:2496, 1997.

[104] J. Fortagh and A. Grossmanna and C. Zimmermann and T. W. Hänsch. *Phys. Rev. Lett.*, **81**:5310, 1998.

[105] J. D. Jackson. *Classical Electrodynamics*. Wiley, NY, 1975.

[106] C. Jarzynski. Diffusion equation for energy in ergodic adiabatic ensembles. *Phys. Rev. A*, **46**:7498–7509, 1992.

[107] C. Jarzynski. Energy diffusion in a chaotic adiabatic billiard gas. *Phys. Rev. E*, **48**:4340–50, 1993.

[108] C. Jarzynski. Multiple-time-scale approach to ergodic adiabatic systems: another look. *Phys. Rev. Lett.*, **71**:839–842, 1993.

[109] R. Kaiser, Y. Levy, N. Vansteenkiste, A. Aspect, W. Seifert, D. Leipold, and J. Mlynek. *Opt. Commun.*, **104**:234, 1994.

[110] V. Kalmeyer and R. B. Laughlin. *Phys. Rev. B*, **35**:9805, 1987.

[111] L. Kaplan and E. J. Heller. *Phys. Rev. E*, **59**:6609–28, 1999.

[112] J. A. Katine, M. A. Eriksson, A. S. Adourian, R. M. Westervelt, J. D. Edwards, A. Lupu-Sax, E. J. Heller, K. L. Campman, and A. C. Gossard. *Phys. Rev. Lett.*, **79**:4806, 1997.

[113] S. Kawata and T. Sugiura. *Opt. Lett.*, **17**:772, 1992.

[114] S. Kawata and T. Tani. *Opt. Lett.*, **21**:1768, 1996.

[115] W. Ketterle and N.J. van Druten. *Phys. Rev. A*, **54**:656, 1996.

[116] M. Key, I. G. Hughes, W. Rooijakkers, B. E. Sauer, E. A. Hinds, D. J. Richardson, and P. G. Kazansky. *Phys. Rev. Lett.*, **84**:1371, 2000.

[117] H. Kogelnik. Theory of optical waveguides. In T. Tamir, editor, *Guided-Wave Optoelectronics*, chapter 2, pages 7–88. Springer-Verlag, 2nd edition, 1990.

[118] S. E. Koonin, R. L. Hatch, and J. Randrup. One-body dissipation in a linear response approach. *Nucl. Phys. A*, **283**:87–107, 1977.

[119] S. E. Koonin and D. C. Meredith. *Computational Physics*. Addison-Wesley, 1990.

[120] S. E. Koonin and J. Randrup. Classical theory for one-body nuclear dynamics. *Nucl. Phys. A*, **289**:475–510, 1977.

[121] I. Kosztin and K. Schulten. *Intl. J. Mod. Phys.*, **8**:293–325, 1997. preprint `physics/9702022`.

[122] R. Kubo. A general expression for the conductivity tensor. *Canadian J. Phys.*, **34**:1274–7, 1956.

[123] A. Kudrolli, V. Kidambi, and S. Sridhar. *Phys. Rev. Lett.*, **75**:822–5, 1995.

[124] G. Labeyrie, A. Landgragin, J. Von Zanthiert, R. Kaiser, N. Vansteenkiste, C. Westbrook, and A. Aspect. *Quant. Semiclass. Opt.*, **8**:603, 1996.

[125] C. Lanczos. *J. Res. Nat. Bur. Standards*, **45**:255–281, 1950.

[126] L. D. Landau and E. M. Lifshitz. *Mechanics*. Pergamon, 1967.

[127] L. D. Landau and E. M. Lifshitz. *Statistical Physics*. Pergamon, 3rd edition, 1994.

[128] L. D. Landau and E. M. Lifshitz. *Quantum Mechanics*. Butterworth-Heinemann, 3rd edition, 1998.

[129] R. Landauer. *IBM J. Res. Dev.*, **1**:233, 1957.

[130] R. Landauer. *Z. Phys. B*, **68**:217, 1987.

[131] A. M. Lane and R. G. Thomas. *Rev. Mod. Phys.*, **30**:257, 1958.

[132] I. R. Lapidus. *Am. J. Phys.*, **50**:45, 1982.

[133] R. Lehoucq, K. Maschhoff, D. Sorensen, and C. Yang. *ARPACK: Arnoldi Package*, 1997. `http://www.caam.rice.edu/software/ARPACK/`.

[134] B. Li. *Phys. Rev. E*, **55**:5376–79, 1997.

[135] B. Li and B. Hu. *J. Phys. A*, **31**:483–504, 1998.

[136] B. Li and M. Robnik. *J. Phys. A*, **28**:2799, 1995.

[137] B. Li, M. Robnik, and B. Hu. *Phys. Rev. E*, **57**:4095–4105, 1998.

[138] R. G. Littlejohn. Optimizing basis sets for quantum mechanical calculations. Talk presented to Heller Group, November 1999.

[139] Y. Lu and F. A. Fernandez. An efficient finite element solution of inhomogeneous anisotropic and lossy dielectric waveguides. *IEEE Transactions on Microwave Theory and Techniques*, **41**(6/7):1215, 1993.

[140] A. Lupu-Sax. *Quantum scattering theory and applications.* PhD thesis, Harvard University, Cambridge, MA, September 1998.

[141] H. Mabuchi and J. H. Kimble. *Opt. Lett.*, **19**:749, 1994.

[142] E. A. J. Marcatili. *Bell Syst. Tech. J.*, **48**:2103, 1969.

[143] D. Marcuse. *J. Opt. Soc. Am.*, **66**:216, 1976.

[144] A. Matulis and D. Šegžda. *J. Phys. Cond. Mat.*, **1**:2289, 1989.

[145] S. W. McDonald and A. N. Kaufman. *Phys. Rev. Lett.*, **42**:1189, 1979.

[146] M. L. Mehta. *Random matrices and the statistical theory of energy levels.* Academic Press, NY, 2nd edition, 1991.

[147] P. Mezza, J. Phalippou, and P. Sempere. *J. Non-Cryst. Solids*, **243**:75, 1999.

[148] P. M. Morse and H. Feshbach. *Methods of Theoretical Physics*, volume 2. McGraw-Hill, NY, 1953.

[149] T. Mukhopadhyay and S. Pal. *Phys. Rev. C*, **56**:296–301, 1997.

[150] S. V. Nitta, V. Pisupatti, A. Jain, P. C. Wayner Jr, W. N. Gill, and J. L. Plawsky. *J. Vac. Sci. Tech. B*, **17**:205, 1999.

[151] M. A. Olshanii. *Laser Physics*, **4**:995, 1994.

[152] M. A. Olshanii. *Phys. Rev. Lett.*, **81**:938, 1998.

[153] E. Ott. Goodness of ergodic adiabatic invariants. *Phys. Rev. Lett.*, **42**:1628–31, 1979.

[154] E. Ott. *Chaos in dynamical systems.* Cambridge University Press, NY, 1993.

[155] Yu. B. Ovchinnikov, I. Manek, and R. Grimm. *Phys. Rev. Lett.*, **79**:2225, 1997.

[156] Yu. B. Ovchinnikov, S. V. Shul'ga, and V. I. Balykin. *J. Phys. B*, **24**:3173, 1991.

[157] Yu. B. Ovchinnikov, J. Söding, and R. Grimm. *JETP Lett.*, **61**:21, 1995.

[158] S. Pal and T. Mukhopadhyay. *Phys. Rev. C*, **54**:1333–1340, 1996.

[159] W. D. Phillips. *Rev. Mod. Phys.*, **70**:721, 1998.

[160] D. Pines and P. Nozières. *The Theory of Quantum Liquids.* W. A. Benjamin, Menlo Park, CA, 1966.

[161] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C.* Cambridge University Press, second edition, 1992. Note that I recommend this only as a textbook as an introduction to concepts in numerical methods; I do not endorse the code examples.

[162] H. Primack. *Quantal and Semiclassical analysis of the three-dimensional Sinai billiard.* PhD thesis, Weizmann Institute, Rehovot, Israel, September 1997.

[163] T. Prosen. *Physica D*, **91**:244, 1996.

[164] T. Prosen and M. Robnik. *J. Phys. A*, **26**:1105–14, 1993.

[165] T. Prosen and M. Robnik. *J. Phys. A*, **27**:8059, 1994.

[166] Jr. R. J. Riddell. *J. Comp. Phys.*, **31**:21–41; 42–59, 1979.

[167] Lord Rayleigh. *Phil. Mag.*, **41**:107, 1871.

[168] Lord Rayleigh. *Phil. Mag.*, **43**:259, 1897.

[169] G. Rempe, R. J. Thompson, H. J. Kimble, and R. Lalezari. *Opt. Lett.*, **17**:363, 1992.

[170] M. J. Renn, D. Montgomery, O. Vdovin, D. Z. Anderson, C. E. Wieman, and E. A. Cornell. *Phys. Rev. Lett.*, **75**:3253, 1995.

[171] M. Robnik. *J. Phys. A*, **17**:1049, 1984.

[172] C. Rouvinez and U. Smilansky. *J. Phys. A*, **28**:77–104, 1995.

[173] W. C. Sabine. *Collected Papers on Acoustics.* Peninsula Publishing, Los Altos, 1993.

[174] J. J. Sakurai. *Modern Quantum Mechanics, Revised Edition.* Addison-Wesley, 1994.

[175] M. J. Sanchez, E. Vergini, and D. A. Wisniacki. *Phys. Rev. E*, **54**:4812–18, 1996.

[176] T. A. Savasa, S. N. Shaha, M. L. Schattenburg, J. M. Carter, and H. I. Smith. *J. Vac. Sci. Technol. B*, **13**:1995, 1995.

[177] A. I. Schnirelman. *Usp. Mat. Nauk.*, **29**:181, 1974.

[178] S. Selleri and M. Zoboli. *J. Opt. Soc. Am. A*, **14**:1460, 1997.

[179] Yu. V. Sharvin. *J. Exptl. Theoret. Phys. (U.S.S.R.)*, **48**:984–985, 1965. [trans. in Sov. Phys. JETP **21**, 655 (1965).].

[180] S. E. J. Shaw. Discrete-Variable Representation. Notes for the Heller Group, 2000.

[181] K. L. Shepard, M. L. Roukes, and B. P. Van der Gaag. *Phys. Rev. Lett.*, **68**:2660, 1992.

[182] P. P. Silvester and R. L. Ferari. *Finite Elements for Electrical Engineers*. Cambridge University Press, 2nd edition, 1991.

[183] E. Snitzer. *J. Opt. Soc. Am.*, **51**:491, 1961.

[184] L. Spruch and Y. Tikochinsky. *Phys. Rev. A*, **48**:4213, 1993.

[185] S. Sridhar and E. J. Heller. *Phys. Rev. A*, **46**:1728, 1992.

[186] Standard Performance Evaluation Corporation, 6585 Merchant Place, Suite 100, Warrenton, VA 20187, USA. *CPU95 Benchmark Results*, 2000. `http://www.spec.org/`.

[187] A. D. Stone and A. Szafer. *IBM J. Res. Develop.*, **32**:317, 1988.

[188] G. Strang. *Linear algebra and its applications*. Academic Press, 3rd edition, 1988.

[189] A. Szafer and A. D. Stone. *Phys. Rev. Lett.*, **62**:300, 1989.

[190] J. H. Thywissen, M. Olshanii, G. Zabow, M. Drndić, K. S. Johnson, R. M. Westervelt, and M. Prentiss. *Euro. Phys. J. D*, **7**:261, 1999.

[191] J. H. Thywissen, R. M. Westervelt, and M. Prentiss. *Phys. Rev. Lett.*, **83**:3762, 1999.

[192] M. A. Topinka, B. J. LeRoy, R. M. Westervelt, S. E. J. Shaw, E. J. Heller, K. M. Maranowski, and A. C. Gossard. *Science (to be published)*, September 2000.

[193] B.J. van Wees, H. van Houten, C. W. J. Beenakker, J. G. Williamson, L. P. Kouwenhoven, D. van der Marel, and C. T. Foxon. *Phys. Rev. Lett.*, **60**:848, 1988.

[194] E. Vergini. *Estudio Cuántico y Semiclásico de Billares Clásicamente Caóticos*. PhD thesis, Universidad de Buenos Aires, 1995.

[195] E. Vergini and M. Saraceno. Calculation by scaling of highly excited states of billiards. *Phys. Rev. E*, **52**:2204–2207, 1995.

[196] E. Vergini and D. A. Wisniacki. *Phys. Rev. E*, **58**:R5225, 1998.

[197] J. D. Weinstein and K. G. Libbrecht. *Phys. Rev. A*, **52**:4004, 1995.

[198] D. A. Wharam, M. Pepper, H. Ahmed, J. E. F. Frost, D. G. Hasko, D. C. Peacock, D. A. Richie, and G. A. C. Jones. *J. Phys. C*, **21**:L887, 1998. also see [20] and references therein.

[199] M. Wilkinson. A semiclassical sum rule for matrix elements of classically chaotic systems. *J. Phys. A*, **20**:2415–23, 1987.

[200] M. Wilkinson. Statistical aspects of dissipation by Landau-Zener transitions. *J. Phys. A*, **21**:4021–37, 1988.

[201] M. Wilkinson. Dissipation by identical oscillators. *J. Phys. A*, **23**:3603–11, 1990.

[202] M. Wilkinson and E. J. Austin. *J. Phys. A*, **28**:2277–96, 1995.

[203] D. A. Wisniacki and E. Vergini. *Phys. Rev. E*, **59**:6579–84, 1999.

[204] M. Wojtkowski. Principles for the design of billiards with nonvanishing Lyapunov exponents. *Commun. Math. Phys.*, **105**:391–414, 1986.

[205] W. Xue and P. A. Lee. *Phys. Rev. B*, **38**:3913, 1988.

[206] A. Yacoby and Y. Imry. *Phys. Rev. B*, **41**:5341, 1990.

[207] Z-C. Yan, A. Dalgarno, and J. F. Babb. *Phys. Rev. A*, **55**:2882, 1997.

Appendix A

General transformation of the 1D Fokker-Planck equation

In this Appendix, I show how the drift and diffusion terms in a general Fokker-Planck equation [79] for evolution of a *probability density function* (PDF) transform when the variable (*i.e.* the spatial ordinate of the PDF) is transformed. It is an essential step to understanding the relation between ergodic energy spreading in the energy variable E (in which the diffusion constant is most easily found) and in the enclosed-phase-space-volume variable Ω (the natural variable for discussing spreading because of its conservation under Liouville's theorem). I present known results; in particular, the results derived here are assumed by Jarzynski [106] and Cohen [46]. However it is complicated enough, especially in the case of a time-dependent density of states, to merit a clear explanation.

We start with a known evolution of the PDF $\eta(\Omega, t)$ in Ω-space, and want to find the equivalent evolution of $\rho(E, t)$, the PDF in E-space which describes the same situation. We assume a one-to-one time-dependent mapping $\Omega = \Omega(E, t)$, which can be represented as a (fixed) surface in three-dimensional (E, t, Ω) space (Fig. A.1). Therefore $\eta(\Omega, t)$ and $\rho(E, t)$ are projections of the *same function* on the 3D surface, down the E and Ω directions respectively. The PDFs are related as follows,

$$\eta = \frac{\rho}{g}, \qquad \text{Jacobean (density of states)} \quad g \equiv \left.\frac{\partial \Omega}{\partial E}\right|_t . \tag{A.1}$$

From now on I will use prime to denote spatial derivative and dot time derivative, in the natural variables of each PDF. Namely,

$$\eta' \equiv \left.\frac{\partial \eta}{\partial \Omega}\right|_t , \quad \dot{\eta} \equiv \left.\frac{\partial \eta}{\partial t}\right|_\Omega , \quad \rho' \equiv \left.\frac{\partial \rho}{\partial E}\right|_t , \quad \dot{\rho} \equiv \left.\frac{\partial \rho}{\partial t}\right|_E . \tag{A.2}$$

Time-dependence will not be written explicitly. Partial derivatives will be taken with t held constant unless otherwise specified.

The Fokker-Planck equations in the two variables arise as follows. The probability fluxes (*i.e.* probability per unit time flowing past a fixed point) are assumed linear in the local density and its local gradient:

$$j_\eta(\Omega, t) = u\eta - C\eta', \tag{A.3}$$

$$j_\rho(E, t) = v\rho - D\rho'. \tag{A.4}$$

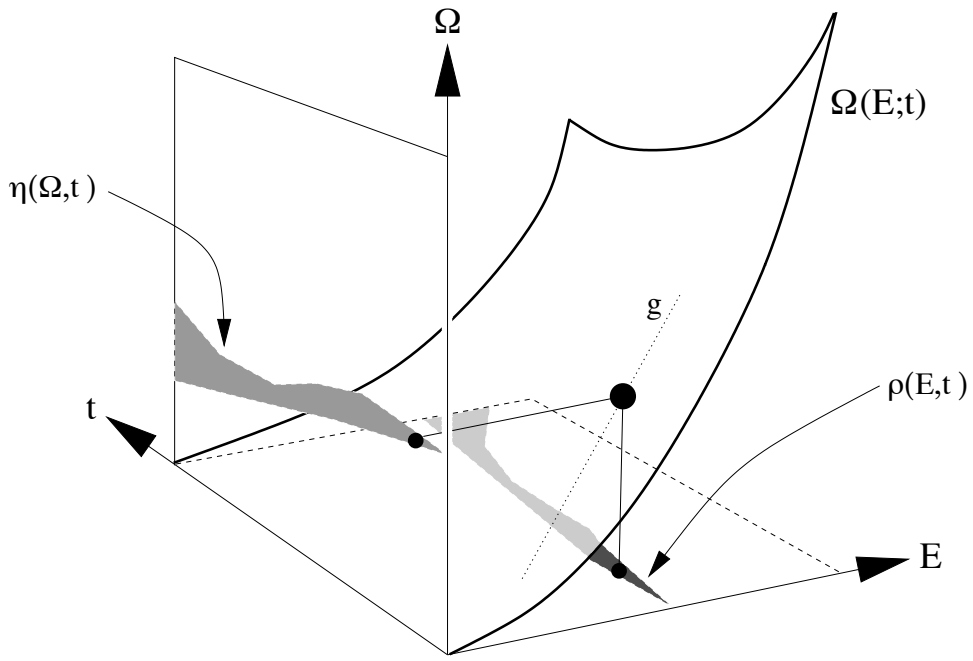

Figure A.1: Projecting a probability density on the surface $\Omega(E,t)$ down to $\rho(E,t)$ in the $E-t$ plane, or across to $\eta(\Omega,t)$ in the $\Omega-t$ plane. The gradient g is also shown.

For the PDF η, the coefficient $u(\Omega,t)$ is a 'drift velocity' field and $C(\Omega,t)$ is a diffusion rate (both are generally space- and time-dependent). For the PDF ρ, similarly, $v(E,t)$ is the drift velocity and $D(E,t)$ the diffusion rate. We also have the probability conservation laws,

$$\dot{\eta} + \frac{\partial j_\eta}{\partial \Omega} = 0, \tag{A.5}$$

$$\dot{\rho} + \frac{\partial j_\rho}{\partial E} = 0. \tag{A.6}$$

Combining (A.3) and (A.5) gives the Fokker-Planck equation obeyed by η,

$$\dot{\eta} = -(u\eta)' + (C\eta')'. \tag{A.7}$$

Similarly the Fokker-Planck equation obeyed by ρ is,

$$\dot{\rho} = -(v\rho)' + (D\rho')'. \tag{A.8}$$

Given functions $u(\Omega,t)$, $C(\Omega,t)$ and (A.7), our task is to find $v(E,t)$ and $D(E,t)$ such that (A.8) and the Jacobean transformation (A.1) are obeyed for all time. The Jacobean results in the following rule for transforming derivatives:

$$\frac{\partial}{\partial \Omega} = \frac{1}{g}\frac{\partial}{\partial E}. \tag{A.9}$$

Use of this and of (A.1) allows (A.7) to be written as

$$\dot{\eta} \;=\; -\frac{1}{g}\frac{\partial}{\partial E}\left(u\frac{\rho}{g}\right) + \frac{1}{g}\frac{\partial}{\partial E}\left(Dg\frac{\partial}{\partial E}\left(\frac{\rho}{g}\right)\right), \tag{A.10}$$

where we have identified the transformed diffusion rate $D = C/g^2$. Note that the diffusion term has a modified form, which we examine later. The LHS now poses a problem. It is possible to transform the LHS, written as $(\partial/\partial t)_\Omega(\rho/g)$, using the Legendre rule for a change of variable:

$$\left.\frac{\partial}{\partial t}\right|_\Omega \;=\; \left.\frac{\partial}{\partial t}\right|_E + \left.\frac{\partial E}{\partial t}\right|_\Omega \cdot \frac{\partial}{\partial E}. \tag{A.11}$$

However the resulting sequence of algebra required to reach the desired form is very complicated[1] and I will not give it here.

Instead I give a geometric approach which bypasses the algebra. We ask how the probability fluxes are related. If $\Omega(E)$ were independent of time, j would transform without a Jacobean, that is $j_\rho = j_\eta$. Now with $\Omega(E)$ time-dependent, ρ acquires an extra drift velocity equal to the gradient of the constant-Ω contour in the $E-t$ plane:

$$j_\rho \;=\; j_\eta + \left.\frac{\partial E}{\partial t}\right|_\Omega \rho. \tag{A.12}$$

This can easily be verified by considering Fig. A.1. The apparent lack of $E-\Omega$ symmetry in this relation is deceptive since the second (flux difference) term can also be written $-(\partial\Omega/\partial t)_E\,\eta$ using the relation

$$\left.\frac{\partial E}{\partial t}\right|_\Omega \left.\frac{\partial t}{\partial \Omega}\right|_E \left.\frac{\partial \Omega}{\partial E}\right|_t \;=\; -1, \tag{A.13}$$

which is a geometric property of any surface in 3D.

Combining (A.12) with (A.5), (A.6) and (A.9) immediately gives

$$\dot{\eta} \;=\; \frac{1}{g}\dot{\rho} + \frac{1}{g}\frac{\partial}{\partial E}\left(\left.\frac{\partial E}{\partial t}\right|_\Omega \rho\right), \tag{A.14}$$

which can now be substituted into the LHS of (A.10) to give the desired Fokker-Planck evolution (A.8). The resulting transformed coefficients are

$$D \;=\; \frac{C}{g^2} \qquad \text{transformed diffusion rate} \tag{A.15}$$

$$v \;=\; \frac{u}{g} + \left.\frac{\partial E}{\partial t}\right|_\Omega + \frac{g'}{g^3}C \qquad \text{transformed drift rate.} \tag{A.16}$$

Looking at the transformed drift, the first term is the (scaled) original drift rate, the second is 'dragging' due to time-dependence of $\Omega(E)$, and the third, which can also be written as $(g'/g)D$, is a drift term entirely due to the modification of the diffusion term. The second term, being (minus) the parameter velocity $\dot{x}$ times the conservative force $F(x) \equiv$

[1] Doron Cohen, personal communication.

$-(\partial E/\partial x)_\Omega$, will only cause reversible energy changes. Note that the third term arises from the the transformation of the diffusion term

$$(C\eta')' \quad \longrightarrow \quad \frac{1}{g}\left(Dg\left(\frac{\rho}{g}\right)'\right)', \tag{A.17}$$

In general, we see that a transformed diffusion term takes the form of the right-hand side of (A.17), which is biased. By 'biased' I mean that it has nonzero local drift when there is zero local density gradient. The bias can be visualized if we imagine that diffusion in Ω-space comes from a the continuum limit of an unbiased random walk on a uniformly-spaced lattice of sites. The same sites appear non-uniformly spaced in E-space, so the same random walk gives additional drift for the PDF in E-space. Only when a diffusion term is written in terms of its 'natural' variable (here Ω) does it take the unbiased form of the left-hand side of (A.17). This illustrates that great care should be taken in writing down an energy diffusion term in the theory of energy spreading.

Appendix B

Numerical evaluation of the classical band profile in billiards

Here I collect the tools needed to calculate $\tilde{C}_{\mathrm{E}}(\omega)$, the classical auto-correlation spectra ('band profiles'), in real systems, in particular billiard systems. Section 2.1.6 is necessary preliminary reading material.

B.1 Trajectories—the issue of exponential divergence

In the hard-chaotic systems we consider, there is exponential divergence of initially-close trajectories [154]. We need to simulate the Hamiltonian dynamics using a computer with necessarily finite-precision arithmetic. This means that numerical round-off error will, after a few times the Lyapunov time t_λ, grow to dominate any trajectory, thereby destroying any relevance of the initial conditions. Might this be a problem for estimating correlation spectra of functions along such a trajectory? Especially one might worry about frequencies $\omega \ll t_\lambda^{-1}$, which are fundamental to the demonstration of the various $\omega \to 0$ behaviours in Chapters 3 and 4.

It turns out that there is no cause for concern. One is saved by the very same exponential divergence, which forces all correlation functions to die exponentially like e^{-t/t_λ}. Therefore times separated by $\gg t_\lambda$ are effectively independent samples and do not contribute to small-frequency components in a spectrum. All the $\omega \to 0$ behaviours in hard-chaos systems result from short time correlations less than or order t_λ, approximately the ergodic time. (This fact alone is sufficient to deduce the possible power laws—done in Appendix D).

This issue is slightly modified for chaotic systems with power-law correlation function decay, for instance the stadium (see Section 3.4). If correlations persist for long times, the Lyapunov time of contributing orbits must be equally long. Therefore these orbits will not become dominated by amplified round-off error noise, rather they will stay numerically accurate. In general, any trajectory can be computed accurately enough that round-off error does not invalidate its contribution to a correlation function.

It turns out that in hyperbolic chaotic systems, a numerically-generated 'pseudo-trajectory' *shadows* (is very close to) a true trajectory with slightly different initial conditions [89, 83]. However, due to the above arguments, this shadowing theorem is not

necessary to give us confidence in correlation function estimates.

B.2 General system—windowed estimation of correlation spectrum

I assume that $\mathcal{F}(t)$ (which I call the 'signal') for a single, ergodic trajectory has been generated and stored at a fine enough time-resolution t_r to capture all the desired features. This can be done by evaluating $-\partial\mathcal{H}/\partial x$ along the trajectory, and subtracting the average value (estimated, or often analytically known). We need $t_r < \omega_0$, where ω_0 is the upper frequency limit desired for the band profile.

The auto-correlation spectrum of this single signal is $\tilde{C}_{\mathcal{F}}(\omega)$, and if the signal exists over all times $-\infty < t < \infty$ then $\tilde{C}_{\mathcal{F}}(\omega)$ is a stochastic quantity uncorrelated in ω. The microcanonical average $\tilde{C}(\omega) = \langle \tilde{C}_{\mathcal{F}}(\omega) \rangle_{\mathrm{E}}$ (I drop the energy subscript) is reached by averaging over initial conditions of trajectories. It is $\tilde{C}(\omega)$ that is the desired 'band profile'. However, $\tilde{C}(\omega)$ is also the *local mean* of $\tilde{C}_{\mathcal{F}}(\omega)$, so can be more easily estimated by smoothing $\tilde{C}_{\mathcal{F}}(\omega)$. This smoothing (convolution with Gaussian of width ω_s) is performed using the Fast Fourier Transform (FFT) [161]. Since computationally the signal must be of necessarily finite length $\sim t_0$, correlations in ω-space arise in $\tilde{C}_{\mathcal{F}}(\omega)$ on the scale t_0^{-1}. Therefore the smoothing operation corresponds to averaging $N \approx \omega_s t_0$ independent samples, and a fractional error of $N^{-1/2}$ results for the band profile estimate. This estimation error appears as multiplicative noise with frequency correlation scale ω_s. Typical parameters were $\omega_0 = 4, \omega_s = 10^{-2}$, $t_0 = 10^5$, giving about 3% RMS estimation error.

There is another 'failure mode' of this estimation procedure: There are only $M =\sim t_0/\tau_{\mathrm{cl}}$ independent samples in the finite sample length, therefore only M independent values can appear in cfw. If $M \ll \omega_0 t_0$ then the errors may exceed that quoted above, because samples are no longer independent. This failure mode only arises when $\tilde{C}(\omega)$ is desired at large ω, which forces t_r to shrink, and therefore t_0 to also shrink if the memory requirements are not to grow. In this case, true averaging over trajectories is the only remedy.

Now I describe how to find $\tilde{C}_{\mathcal{F}}(\omega)$ given $\mathcal{F}(t)$. I will always use the Fourier Transform (FT) convention

$$\begin{aligned} \tilde{A}(\omega) &= \int_{-\infty}^{\infty} dt\, e^{i\omega t} A(t) && \text{forward,} \\ A(t) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} d\omega\, e^{-i\omega t}\tilde{A}(\omega) && \text{backward.} \end{aligned} \tag{B.1}$$

For this general real signal $A(t)$, the Wiener-Khinchin Theorem (easy to prove [161]) equates the power spectrum to the FT of the auto-correlation,

$$|\tilde{A}(\omega)|^2 \;=\; \tilde{C}_A(\omega), \qquad \text{with} \quad C_A(\tau) \;\equiv\; \int_{-\infty}^{\infty} dt\, A(t)A(t+\tau). \tag{B.2}$$

I choose $A(t)$ to be a windowed version of the signal,

$$A(t) \;=\; W(t)\,\mathcal{F}(t), \qquad W(t) = \text{windowing function}, \tag{B.3}$$

where $W(\pm\infty) \to 0$ and has finite characteristic width t_0. This windowing function defines the maximum trajectory length needed. Therefore the auto-correlation is related to the desired $C_{\mathcal{F}}(\tau)$ by

$$C_A(\tau) \;=\; \int_{-\infty}^{\infty} dt\, W(t)W(t+\tau)\mathcal{F}(t)\mathcal{F}(t+\tau) \;=\; C_W(\tau)C_{\mathcal{F}}(\tau), \tag{B.4}$$

where the last step involved identifying $W(t)W(t+\tau)$ as a τ-dependent weighting function whose t-integral is $C_W(\tau)$, giving a generalization of the 'top-hat' weighting function of *e.g.* Eq.(2.30). Applying the Convolution Theorem[161] and dividing by Z,

$$\frac{|\tilde{A}(\omega)|^2}{Z} \;=\; \frac{1}{Z}(\tilde{C}_{\mathcal{F}} * \tilde{C}_W)(\omega). \tag{B.5}$$

$\tilde{C}_W(\omega)$ is a narrowly peaked function $\approx Z\delta(\omega)$ where the normalization is $Z = 2\pi \int_{-\infty}^{\infty} dt\, W(t)^2$. Therefore we have the LHS as our estimate for $\tilde{C}_{\mathcal{F}}(\omega)$. Structure below frequency scale t_0^{-1} (the δ-function width) is lost. I typically used the Gaussian window $W(t) = e^{-t^2/(2t_0^2)}$, for which $Z = 2\pi^{3/2}t_0$. The maximum signal length T_0 needed to evaluate the windowed signal $A(t)$ was typically chosen as $10t_0$, beyond which $W(t)$ is negligible. Because $W(t)$ dies to zero very smoothly, the windowed signal $A(t)$ of length T_0 can now be used as a periodic discretely-sampled array, and the FFT used to find $\tilde{A}(\omega)$.

All FFTs were implemented on a Compaq XP1000 (21264 667MHz Alpha processor) running C++ and the Compaq Extended Math Library (CXML). Using this setup, an 8 million (2^{23}) point double-precision FFT takes about 5 seconds.

B.3 Considerations in a billiard system

In billiards the signal $\mathcal{F}(t)$ consists of an irregular series of spikes (corresponding to the collisions with the walls—they are infinitely narrow in the hard-wall limit), with the average value subtracted (see Section 3.1.1).

In order to use the FFT as in the previous section, $\mathcal{F}(t)$ is needed on a *regular* time grid. The problem arises: how to represent irregular δ-functions on a regular grid? Is the limit of grid spacing $\Delta t \to 0$ needed? The answer is no. The spikes can be replaced by (*i.e.* convolved with) Gaussians of width t_r, if $\tilde{C}(\omega)$ is only desired up to frequency ω_0. This modified signal is then windowed to give $A(t)$ and resampled on regular grid of spacing $\Delta t < t_r$. For efficiency of the FFT, Δt is chosen such that the number of samples, $T_0\Delta t$, is an integer power of 2. In frequency-space, the convolution has the effect of multiplying $\tilde{A}(\omega)$ by a high-frequency 'roll-off' function $R(\omega) = e^{-t_r^2\omega^2/2}$. Since this roll-off is known, it can be exactly corrected for by division of $\tilde{A}(\omega)$ by $R(\omega)$. This correction allowed accurate $\tilde{C}(\omega)$ estimation up to $\omega \approx 2/t_r$.

Note that the δ-function nature of the signal might suggest that its FT could best be performed by direct summation of the terms $e^{i\omega t_i}$ arising from each collision i. However this turns out to scale like N^2 where N is the number of collisions needed ($\approx T_0/\tau_{\rm bl}$), *i.e.* it is a 'slow' FT. Resampling onto an even grid and using the FFT, which scales like $N \log_2 N$, is vastly superior.

Appendix C

Numerical evaluation of quantum band profile in billiards

Here I gather the tools used to estimate the quantum 'band profile' $\tilde{C}_{\rm E}^{\rm qm}(\omega)$ given the eigenstates of a billiard. Sections 2.2.2 and 3.1 are required preliminary reading. I start with the very important derivation of the matrix elements.

C.1 Deformation matrix elements

When a billiard is deformed, the potential $U(\mathbf{r})$ changes by an infinite amount (in the hard-wall limit) in narrow regions of space on the boundary. These regions are simply those excluded from or newly-included into the cavity interior by the deformation. The operator $\partial\mathcal{H}/\partial x$ is therefore either infinite or zero in a position representation. Therefore it might be a surprise that the matrix elements $(\partial\mathcal{H}/\partial x)_{nm}$ in the energy representation are in fact finite. The following derivation has been known for some time [28, 118, 46].

The position of a particle in the vicinity of a wall element is conveniently described by $\mathbf{r} = (\mathbf{s}, z)$, where $\mathbf{s}$ is a $d-1$ dimensional surface coordinate and z is a perpendicular 'radial' coordinate. I set $U(\mathbf{r}) = U_0$ outside the undeformed billiard; this will relax the usual Dirichlet condition $\psi = 0$ on the boundary (at $z = 0$). Later I will take the limit $U_0 \to \infty$. I have

$$\frac{\partial\mathcal{H}}{\partial x} \;=\; -[\mathbf{n}(\mathbf{s})\cdot\hat{\mathbf{D}}(\mathbf{s})]\, U_0\delta(z). \tag{C.1}$$

The *normal derivative* (derivative with respect to z) of an eigenfunction $\psi(\mathbf{r})$ measured on the boundary is $\varphi(\mathbf{s}) \equiv \mathbf{n}\cdot\nabla\psi$. The 'logarithmic derivative' is defined as the ratio of gradient to value, $\varphi(\mathbf{s})/\psi(\mathbf{s})$. For $z > 0$ the wavefunction $\psi(\mathbf{r})$ is a decaying exponential in z, using simple one-dimensional considerations which hold true if the decay length is much smaller than any radius of curvature. Hence the logarithmic derivative of the wavefunction on the boundary should be equal to $-\sqrt{2mU_0}/\hbar$. Consequently one obtains

$$\left(\frac{\partial\mathcal{H}}{\partial x}\right)_{nm} \;=\; -\frac{\hbar^2}{2m}\oint \varphi_n(\mathbf{s})\varphi_m(\mathbf{s})\,(\mathbf{n}\cdot\hat{\mathbf{D}})d\mathbf{s}. \tag{C.2}$$

The weighting $(\mathbf{n}\cdot\hat{\mathbf{D}})$ is simply the deformation function $D(\mathbf{s})$. Note that for any *finite* deformation x, the matrix elements of $\mathcal{H}(x)$ in the original energy basis become infinite—this problem is avoided by always remaining in the local adiabatic energy basis [46, 48].

Therefore, given a subset of N adjacent eigenstates, an on-diagonal block of the matrix $(\partial\mathcal{H}/\partial x)_{nm}$ can be found. A resulting example is shown in Fig. 2.7. The integrals are evaluated using the techniques of Appendix G.

C.2 Smoothing

As is clear from Fig. 2.7, individual matrix elements are random quantities. Recalling the definition (2.48), the average square matrix element a certain distance ω from the diagonal is proportional to the desired band profile $\tilde{C}_{\mathrm{E}}^{\mathrm{qm}}(\omega)$. I estimate $\tilde{C}_{\mathrm{E}}^{\mathrm{qm}}(\omega)$ only at equally-spaced frequency points $\omega_j = (j + \frac{1}{2}\omega_s)$, where integer $j \geq 0$. The value at each point j is found by averaging the squared elements whose frequency difference $|\omega_n - \omega_m|$ falls within the 'bin' from $j\omega_s$ to $(j+1)\omega_s$. The constant bin width ω_s determines the resulting frequency resolution. This is seen to be simply smoothing via 'histogram binning'. All such bin averages are multiplied by $2\pi\hbar/\Delta$, if needed to give $\tilde{C}_{\mathrm{E}}^{\mathrm{qm}}(\omega)$ in the correct units. Since the distance from the diagonal is classically small ($\hbar|\omega_n - \omega_m| \ll E$), linearization of the dispersion relation is a good approximation giving $k_n - k_m = (\omega_n - \omega_m)/v$. In practice, I bin in terms of the wavenumber difference $\kappa \equiv k_n - k_m$. Note that the phrase 'distance from the diagonal' would strictly imply an *integer* $r \equiv n - m$. The difference between this interpretation (r) and the continuous version based on the corresponding eigenvalues (κ) is small, having jitter on the scale of a single level-spacing. This small size of jitter is due to spectral rigidity[24]. The choice of continuous over integer variable is therefore arbitrary if only features larger than the level spacing are desired. The only time my choice is important is in Section 2.3.4.

The diagonal elements $n = m$ are not treated as part of the band profile, and are removed before binning. Since the band profile is symmetric (the matrix is Hermitian), a 'single-sided' band profile was taken (discarding the sign of κ). However the above can easily be extended to a two-sided version (preserving the sign of κ) if the band profile of a non-Hermitian, possibly rectangular, matrix is desired.

Statistical errors result from any estimation of average value. If the number of matrix elements collected in a bin is $\mathcal{N}_j$, then the fractional error is Gaussian with standard deviation $\mathcal{N}_j^{-1/2}$ (for $\mathcal{N}_j \gg 1$). This assumes that all elements were uncorrelated random variables. This was found to be a good assumption, *except* at certain κ in billiards which have strong scarring. As discussed in Section 2.3.2, at these κ the average is determined by a few very large values of $|(\partial\mathcal{H}/\partial x)_{nm}|^2$, giving poorer estimation errors at certain places in the band profile. For the quarter-stadium example at $k \approx 400$, I used $N = 451$ states, giving an estimation error (varying with distance from the diagonal) at the 10% level.

Appendix D

How Many Special Deformations Are There?

In Chapter 2, the noise power spectrum $\tilde{C}_{\rm E}(\omega)$ of the generalized 'force' on the parameter was identified as proportional to the dissipation rate due to driving that parameter at frequency ω. Restricting to that case of a single particle Hamiltonian in a deformable potential field, it would be interesting to know if there can exist 'special' deformations (those which give vanishing $\tilde{C}_{\rm E}(\omega)$ as the frequency $\omega \to 0$), other than the class of dilations, translations and rotations identified in Chapter 3.

I will restrict myself to systems with 'hard chaos' (see Section 3.4), where all dynamical correlation functions die exponentially at long times $\gg t_{\rm erg}$, in other words they are *short range*.

The moments of $C(\tau)$ (I omit subscripts E) are defined by

$$\mathsf{M}_p \equiv \int_{-\infty}^{\infty} C(\tau)\tau^p d\tau, \qquad \text{integer } p \geq 0, \tag{D.1}$$

and give the Taylor expansion coefficients of $\tilde{C}(\omega)$ about $\omega = 0$, as follows:

$$\left(-i\frac{\partial}{\partial\omega}\right)^p\Big|_0 \tilde{C}(\omega) = \mathsf{M}_p \tag{D.2}$$

which follows from the definition of the fourier transform. Since $C(\tau) = C(-\tau)$, all the odd moments are zero. We will restrict all the even moments to be finite; this is equivalent to saying $\tilde{C}(\omega)$ has a well-defined $\omega = 0$ limit, namely M_0, and a well-defined local expansion in non-negative even powers of ω. This will exclude certain $C(\tau)$ with power-law (long range) tails.

For a generic deformation, the moments M_0, M_2, $\mathsf{M}_4 \cdots$ will numbers with no particular reason to take the value zero. The smallest p with $\mathsf{M}_p \neq 0$ will determine the dominant power-law ω^γ seen as $\omega \to 0$; in this case it will be $\gamma = 0$. Now a 'special' deformation has the property that the (even) moments $p < \gamma$ vanish, giving the dominant power-law ω^γ with γ an even integer greater than zero.

In this way, we see that the special nature of certain deformations in a hard-chaos billiard is *not* due to any long-time conspiracy in $\mathcal{F}(t)$ (all correlations are lost beyond $t_{\rm erg}$), rather to short-time correlations with vanishing lower moments.

The basic issue is whether $\mathcal{F}(t)$ can be written as an *exact time derivative*: if it can, we have a 'special' deformation. We will formalize this more carefully.

D.1 Requirement on time-integral of $\mathcal{F}(t)$

For any $\mathcal{F}(t)$ we can define $\mathcal{G}(t) = \int_0^t \mathcal{F}(t')dt'$, whose autocorrelation function $C_{\mathcal{G}}(\tau)$ is related to that of $\mathcal{F}$ (using the time-average form (2.30)) by

$$C(\tau) \;=\; -\frac{d^2}{d\tau^2}C_{\mathcal{G}}(\tau) + \lim_{T\to\infty}\frac{1}{T}\left[\mathcal{G}(t)\frac{d}{dt}\mathcal{G}(t+\tau)\right]_{t=-T/2}^{t=T/2}. \tag{D.3}$$

$C_{\mathcal{G}}(\tau)$ is well-defined if $\mathcal{G}(t)$ is a *stationary process* [79], *i.e.* its statistical properties, in particular its average, do not change with time. (We have already assumed $\mathbf{r}(t), \mathcal{F}(t)$, etc. are stationary processes). This condition causes the second term to vanish, since then $\mathcal{G}(t)$ remains bounded. Integrating the above over all τ gives an expression for the zeroth moment: if M_0 vanishes this implies $\mathcal{G}(t)$ must be stationary.

D.2 General potential case

We start with the general potential case (with constant isotropic mass). The deformation field $\mathbf{D}(\mathbf{r})$ must be defined everywhere in space (note that it does not change in time). The deformation parameter x controls the hamiltonian

$$\mathcal{H}(\mathbf{r},\mathbf{p};x) = \frac{\mathbf{p}^2}{2m} + U\big(\mathbf{r} - x\mathbf{D}(\mathbf{r})\big) \tag{D.4}$$

by distorting the potential in space. Note that there is already an infinite class of 'trivially special' deformation fields such that $\mathbf{D}\perp\nabla U$ everywhere: these cause *no* deformation. Therefore a scalar field $D(\mathbf{r}) \equiv \nabla U\cdot\mathbf{D}$ is sufficient to uniquely define the deformation. However working with the vector field $\mathbf{D}(\mathbf{r})$ is more convenient because every function $\mathbf{D}(\mathbf{r})$ is a valid deformation. (Not every function $D(\mathbf{r})$ is valid; consider regions of constant potential). At a frozen (fixed) x value, the fluctuating force on the x parameter, defined by (2.6), can be written

$$\mathcal{F}(t) \;=\; -m\ddot{\mathbf{r}}\cdot\mathbf{D}(\mathbf{r}) - F(x), \tag{D.5}$$

where $\mathbf{r} = \mathbf{r}(t)$ is the (time-dependent) particle position, and $m\ddot{\mathbf{r}} = \dot{\mathbf{p}} = -\nabla U(\mathbf{r})$ is the force on the particle. Therefore $\mathcal{F}$ is a function of the particle position alone. The constant $F(x)$ is chosen such that the time-average of $\mathcal{F}$ is zero; this is necessary for the asymptotic value $C(\tau\to\pm\infty)$ to be zero, to allow finite moments of $C(\tau)$. We will not restrict ourselves to phase-space-volume-preserving deformations (*i.e.* ones which have $F(x) = 0$).

The integral of (D.5) gives the general expression

$$\mathcal{G}(t) \;=\; -m\int_0^t dt'\;\ddot{\mathbf{r}}\cdot\Big[\mathbf{a} + B\mathbf{r} + O(r^2)\cdots\Big] \;-\; F(x)t. \tag{D.6}$$

where the square brackets enclose the expansion of $\mathbf{D}(\mathbf{r})$ (each vector component is assumed to be a multi-dimensional Taylor series about $\mathbf{r} = \mathbf{0}$), and $\mathbf{r}$ is the position at time t'. The

coefficients $\mathbf{a}, B, \cdots$ are rank $1, 2, \cdots$ objects. The linear growth rate $F(x)$ is determined uniquely for a given $\mathbf{D}$. The linear growth (final term) is capable of cancelling deterministic linear growth in the time-integral term; however, if the time-integral grows *diffusively*, then no value of $F(x)$ can prevent $\mathcal{G}$ from being non-stationary. Therefore a choice of $\mathbf{D}(\mathbf{r})$ which is 'special' will have no terms in the time-integral which grow diffusively. Note that, since $\mathbf{r}(t)$ decorrelates on a timescale of $t_{\rm erg}$, diffusive growth (random walk) is the generic case.

We treat the time-integral terms in (D.6) in increasing powers of r, (writing with Einstein summation notation):

$$\int_0^t dt' \ddot{r}_i a_i = a_i[\dot{r}_i]_0^t \tag{D.7}$$

$$\int_0^t dt' \ddot{r}_i B_{ij} r_j = -B_{ij} \int_0^t dt' \dot{r}_i \dot{r}_j + B_{ij}[\dot{r}_i r_j]_0^t \tag{D.8}$$

$$\int_0^t dt' \ddot{r}_i C_{ijk} r_j r_k = -C_{ijk} \int_0^t dt' \dot{r}_i (\dot{r}_j r_k + r_j \dot{r}_k) + C_{ijk}[\cdots]_0^t, \tag{D.9}$$

and so on. What constraints arise on the coefficients for there to be no diffusive growth terms? The first result shows that since the integrand is an exact time derivative, there is no diffusive growth for any $\mathbf{a}$. In general, any boundary-type terms (those coming from exact time derivatives, shown with square brackets) cannot cause diffusion since all functions of $\mathbf{r}(t)$ are stationary and bounded (since the energy-surface is bounded in phase-space); so these terms can be ignored. The time-integral term in (D.8) will vanish if B is any antisymmetric tensor. Any symmetric component in B will give a quadratic form in the momentum, and therefore a fluctuating function whose time-integral grows diffusively. The one important exception to this is when this quadratic form is proportional (up to an exact time derivative) to $\mathcal{H}$, the Hamiltonian. $\mathcal{H}$ is a constant of the motion, so gives uniform linear (non-diffusive) growth of the time-integral. This will correspond to the special nature of dilation in the hard-walled case discussed in Sec. D.3.

We can show that the higher terms (D.9), etc cause diffusive growth, unless either all the coefficients C_{ijk}, D_{ijkl}, etc are zero, or unless the coefficients fall into the 'trivially special' subspace corresponding to $\mathbf{D} \cdot \ddot{\mathbf{r}} = 0$ at all accessible $\mathbf{r}$. (This latter case imposes linear conditions on the coefficients; we exclude the resulting deformation-less subspace). For any given k^{th}-order term, for instance (D.9) with $k = 3$, the coefficient tensor must be totally symmetric with respect to interchange of its last $k-1$ indices (nonsymmetric parts cancel in the summation because of this symmetry in the r factors). Note that this symmetry forbids the type of argument that allowed an antisymmetric B to give no diffusive growth for $p = 2$. Manipulation by parts conserves the symmetry of the r terms, the power of r (namely k), and the number of time-derivatives, or 'dots' (namely 2). (It also generates ignorable boundary-type terms). There are only two cases:

1. Both the dots are on the same r, in which case the integrand is a (non-zero) function of position alone (since $\ddot{r}_i$ is a function of position alone).

2. The dots are on different r's, giving a quadratic form in $\mathbf{p}$ multiplied by a (non-zero) function of $\mathbf{r}$.

It is clear that by such manipulation, no exact time-derivative can be produced.

The only other way to prevent diffusive growth is by making the integrand a multiple of $\mathcal{H}$ (plus exact time- derivatives). Case 2) cannot give the required p^2 term so it is no help. (The only p^2 term can come from an isotropic B tensor). It is unknown whether case 1) can give a term proportional to $U(\mathbf{r})$, which could add to an existing p^2 term to give $\mathcal{H}$. It seems this latter occurrence could only happen for particular forms of $U(\mathbf{r})$, for instance the scaling potentials given below [1]. If it happens, it would give a 'dilation-type' special deformation corresponding to conservation of $\mathcal{H}$.

In conclusion this gives the form of a 'special' deformation (in a general potential) as

$$\mathbf{D}(\mathbf{r}) \;=\; \mathbf{a} + B^a\mathbf{r}, \qquad\qquad B^a \text{ an antisymmetric tensor.} \tag{D.10}$$

The first term accounts for all possible translations, and the second all possible rotations, in d dimensions. The existence of a 'dilation-type' special deformation for a general potential is still open.

D.3 Scaling potentials and the billiard case

A scaling system [126] is created when the potential obeys

$$U(\alpha\mathbf{r}) \;=\; \alpha^c U(\mathbf{r}), \qquad\qquad \text{all } \alpha. \tag{D.11}$$

The potential contours all have the same shape, and the radial dependence is a c^{th}-power-law in any direction (c is even for bounded motion). This gives the property $\mathbf{r}\cdot\nabla U = cU$ everywhere in space, so that the constant of the motion $\mathcal{H}$ can be written $\frac{1}{2}m\dot{r}^2 - (m/c)\mathbf{r}\cdot\ddot{\mathbf{r}}$. This last expression is composed of integrand terms in (D.8), in the case of isotropic $B_{ij} \propto \delta_{ij}$, corresponding to dilation. Thus in this case the integrand is constant [and equal to $F(x)$, in order to have zero time-average of $\mathcal{F}(t)$]; there is no diffusive growth and dilation is 'special'. The limit of hard-walled billiards corresponds to $c \to \infty$, in which case the constant of the motion is simply the kinetic term p^2.

To summarize, in d dimensions in a scaling potential (of which the billiard is a special case), counting the 'special' degrees of freedom gives: d for translations (vector $\mathbf{a}$), $\frac{1}{2}d(d-1)$ for rotations (antisymmetric part of B), and 1 for dilation (isotropic part of B). The total is $\frac{1}{2}d(d+1)+1$.

I have strong numerical evidence that dilation is the only new special deformation which always arises when a hard-walled limit is taken of a general potential. Certainly the above arguments are sufficient to exclude simple cases, such as shear-type deformations. However I cannot exclude the possibility that a $U(\mathbf{r})$ which is not differentiable everywhere in space allows new special $\mathbf{D}(\mathbf{r})$ functions to arise, which are not expressible as the Taylor series of (D.6). Also worthy of study is the general Hamiltonian system $\mathcal{H}(\mathbf{r},\mathbf{p};x)$, no longer restricted to a constant mass tensor $\mathbf{p} = M\dot{\mathbf{r}}$. This restriction played a key role in the above arguments.

[1] I have found other examples, for instance the deformation field $\mathbf{D}(\mathbf{r}) = (\frac{1}{2} + \alpha r^2)\mathbf{r}$ is special for the potential $U(\mathbf{r}) = U_0 e^{-1/(2r^2)}$. Such a deformation is dilation-like, but not a linear dilation given by the symmetric part of B.

Appendix E

Cross correlations I: general-special

Consider two noisy signals $\mathcal{F}(t)$ and $\mathcal{G}(t)$. We assume that $\langle \mathcal{F}(t) \rangle = \langle \mathcal{G}(t) \rangle = 0$. The angular brackets stand for an average over realizations. The auto-correlations of $\mathcal{F}(t)$ and $\mathcal{G}(t)$ are described by functions $C_{\rm F}(\tau)$ and $C_{\rm G}(\tau)$ respectively. We assume that both auto-correlation functions are short-range, meaning no power-law tails (this corresponds to the hard chaos assumption of this paper), and that they are negligible beyond a time $\tau_{\rm c}$. We call a signal 'special' if the algebraic area under its auto-correlation is zero. The cross-correlation function is defined as

$$C_{\rm F,G}(\tau) \equiv \langle \mathcal{F}(t')\mathcal{G}(t'') \rangle, \qquad \tau \equiv t' - t''. \tag{E.1}$$

We assume stationary processes so that the cross-correlation function depends only on the time difference τ. We also symmetrize this function if it does not have $\tau \mapsto -\tau$ symmetry. We assume that $C_{\rm F,G}(\tau)$ is short-range, meaning that it becomes negligibly small for $|\tau| > \tau_{\rm c}$. We would like to prove that if either $\mathcal{F}(t)$ or $\mathcal{G}(t)$ is special then the algebraic area under the cross-correlation function equals zero.

Consider the case where $\mathcal{F}(t)$ is general while $\mathcal{G}(t)$ is special. The integral of $C_{\rm F}(\tau)$ will be denoted by ν. Define the processes

$$X(t) = \int_0^t \mathcal{F}(t')dt' \tag{E.2}$$

$$Y(t) = \int_0^t \mathcal{G}(t'')dt''. \tag{E.3}$$

From our assumptions it follows, disregarding a transient, that for $t \gg \tau_{\rm c}$ we have diffusive growth $\langle X(t)^2 \rangle \approx \nu t$. (It may help the reader to review the discussion in Section 2.1.6). However since $Y(t)$ is a stationary process [79], $\langle Y(t)^2 \rangle \approx$ const. Therefore for a typical realization we have $|X(t)| \leq \text{const} \times \sqrt{\nu t}$ and $|Y(t)| \leq$ const. Consequently, without making any claims on the independence of $X(t)$ and $Y(t)$, we get that $\langle X(t)Y(t) \rangle$ cannot grow faster than const $\times \sqrt{\nu t}$. Using the definitions (E.2), (E.3) and (E.1) we can write

$$\int_{-\infty}^{\infty} C_{\rm F,G}(\tau)d\tau = \frac{\langle X(t)Y(t) \rangle}{t} \approx \frac{\text{const}}{\sqrt{t}} \to 0 \tag{E.4}$$

where the limit $t \to \infty$ is taken. Thus we have proved our assertion.

Appendix F

Cross correlations II: normal-general

In this section we further discuss some features of the cross-correlation function. For the purpose of presentation we we would like to view the time as an integer variable $t = 1, 2, 3....$ One may think of each instant of time as corresponding to a bounce.

Let us assume that we have functions $f(s)$ and $g(s)$, and a time-sequence $(s_1, s_2, s_3, ...)$. This gives two stochastic-like processes $(\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, ...)$ and $(\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3, ...)$. The cross correlation of these two processes is defined as follows:

$$C_{\mathrm{F,G}}(i-j) = \langle \mathcal{F}_i \mathcal{G}_j \rangle = \langle f(s_i) g(s_j) \rangle \tag{F.1}$$

It is implicit in this definition that we assume that the processes are stationary, so the result depends only on the difference $\tau = (i - j)$. The angular brackets stand for an average over realizations of s-sequences.

If the sequences are ergodic on the s domain, then it follows that

$$\begin{aligned} \langle \mathcal{F} \rangle &= \int f(s) ds \\ \langle \mathcal{G} \rangle &= \int g(s) ds \\ C_{\mathrm{F,G}}(0) &= \int f(s) g(s) ds \end{aligned} \tag{F.2}$$

The $\tau \neq 0$ cross-correlations requires information beyond mere ergodicity. In case that the s sequence is completely uncorrelated in time we can factorized the averaging and we get $C_{\mathrm{F,G}}(\tau \neq 0) = \langle \mathcal{F} \rangle \times \langle \mathcal{G} \rangle$. If $\langle \mathcal{F} \rangle = 0$ then

$$C_{\mathrm{F,G}}(\tau \neq 0) = 0 \tag{F.3}$$

irrespective of $\langle \mathcal{G} \rangle$.

However, we would like to define circumstances in which Eq.(F.3) is valid, even if the s sequence is *not* uncorrelated. In such case either the $\mathcal{F}$ or the $\mathcal{G}$ may possess time correlations. (Such is the case if $\mathcal{G}$ is 'special'). So let us consider the case where the $\mathcal{F}$

sequence *looks random*, while assuming nothing about the $\mathcal{G}$ sequence. By the phrase 'looks random' we mean that the conditional probability satisfies

$$\text{Prob}(\mathcal{F}_i|s_j) = \text{Prob}(\mathcal{F}_i) \qquad \text{for any } i \neq j \tag{F.4}$$

Eq. (F.3) straightforwardly follows provided $\langle \mathcal{F} \rangle = 0$, irrespective of the $g(s)$ involved. Given $f(s)$, the goodness of assumption (F.4) can be actually tested. However, it is not convenient to consider (F.4) as a practical definition of a 'normal' deformation.

Appendix G

Numerical evaluation of wavefunction boundary integrals

Boundary methods are a central component of this thesis. Closed integrals of a function $f(\mathbf{s})$ over the boundary coordinate $\mathbf{s}$ are ubiquitous. Generally $f(\mathbf{s}) = g(\mathbf{s})h(\mathbf{s})$ and a square matrix F of integrals

$$F_{nm} \;=\; \oint d\mathbf{s}\, g_n(\mathbf{s}) h_m(\mathbf{s}), \qquad\qquad n, m \;=\; 1 \cdots N, \tag{G.1}$$

is required, where the indices label multiple functions. For evaluation of the quantum band profile (Chapters 2 and 3), the local density of states (Chapter 6), the tension and area-norm matrices (Chapter 5), and the Vergini matrix and its derivative (Chapter 6), $g(\mathbf{s})$ and $h(\mathbf{s})$ are basis functions or eigenstates which oscillate about zero on the length scale λ_{B}, the quantum (de Broglie) free-space wavelength. The deformation function boundary integrals (Chapter 4) do not involve any quantum scale, but are also evaluated using the method below. I will present only the $d = 2$ case where boundary integrals over $\mathbf{s}$ become line integrals over s; the generalization to higher d is simple.

My tool for evaluation of an integral on a closed curve is the discretization

$$F \;\equiv\; \oint ds f(s) \;\longrightarrow\; \frac{L}{M} \sum_{i=1}^{M} f(s_i), \tag{G.2}$$

where L is the range of s, that is, the length of the line integral (billiard perimeter). The M points are spread uniformly (equidistant in s) along the closed curve. Because no point is special, no special quadrature [161] weights arise near any endpoints: all weights are equal. More sophisticated and accurate approximations exist for closed line integral evaluation [58], however this is sufficient for my needs and is very simple to code. Its errors will be discussed and tested below.

A single integral (G.2) requires M function evaluations of f. Naively one might guess that filling a matrix F using (G.1) requires $O(N^2 M)$ evaluations. However, the correct way to compute (G.1) requires only $O(NM)$ such evaluations: First fill the rectangular matrices $G_{im} = g_m(s_i)$ and $H_{im} = h_m(s_i)$, from which follows

$$F \;=\; \frac{L}{M} G^{\mathrm{T}} H. \tag{G.3}$$

This matrix multiplication *does* require N^2M operations, but being simple adds and multiplications (and using optimized library code *e.g.* BLAS), it is very fast and does not affect the scaling. If you like, the matrix multiply 'performs' the integration over i. In the case where g and h are the same function, only NM evaluations are required. Note that if a general weighting function $w(\mathbf{s})$ is required in the integrand (G.1), it can easily be incorporated into G or H, or equivalently be included as a diagonal matrix W inserted between G^{T} and H in (G.3).

G.1 Convergence with number of sample points—theory

What is the minimum M required for a given accuracy in (G.2)? Writing $f(s)$ as a complex Fourier series on the periodic interval L,

$$f(s) \;=\; \sum_{n=-\infty}^{\infty} \tilde{f}_n e^{ik_n s}, \qquad\qquad \text{where } k_n = \frac{2\pi n}{L}. \tag{G.4}$$

it is clear that the desired integral F is just $L\tilde{f}_0$. The approximation (G.2) instead gives the sum $L(\tilde{f}_0 + \tilde{f}_{\pm M} + \tilde{f}_{\pm 2M} + \cdots)$. The notation $\tilde{f}_{\pm n}$ implies the sum of $\tilde{f}_n$ and $\tilde{f}_{-n}$ both with arbitrary phase factors. This is an example of the Nyquist sampling theorem[161]: f is completely represented up to spatial frequencies $\pm k_{M/2}$, but beyond this the spectrum is folded back into this range. In particular the frequencies k_{jM} for integer $j \neq 0$ get folded back to zero-frequency, corrupting the estimate of F. Thus one should choose sample spacing $L/M < 2\pi/k_{\max}$, where $k_{\max}$ is the *bandwidth* of $f(s)$. That is, $f(s)$ is assumed to have negligible components beyond wavenumber $k_{\max}$.

The functions $g(s)$ and $h(s)$ are wavefunctions (Helmholtz equation solutions existing in the plane), measured along a closed curve, which I first take as smooth and of large curvature radius. So to a first approximation $g(s)$ and $h(s)$ are *bandwidth-limited* to $k_{\max} = k$, where the free-space wavenumber is $k \equiv 2\pi/\lambda_{\mathrm{B}}$. Therefore $\tilde{f} = \tilde{g} * \tilde{h}$ is bandwidth-limited to $2k$. This would suggest that just over 2 samples per free-space wavelength are needed to achieve high accuracy, and that convergence is exponential beyond this.

Two facts heavily modify this simple semiclassical picture: 1) there exist evanescent wave components in eigenstates or basis functions (for instance, evanescent basis functions can oscillate at up to $4k$, see Section 6.3.3), and 2) the boundary may not be infinitely differentiable (required for exponential convergence), rather corners (and 'kinks') may introduce discontinuities into the function $f(s)$ which give power-law Fourier tails (and power-law convergence).

G.2 Convergence in practice

Following [137], the number of samples per wavelength I will call $b \equiv L/M\lambda_{\mathrm{B}}$. A practical choice of b will depend on what type of wavefunction and billiard we have, as well as the required accuracy of F.

I found $b = 6$ is sufficient for most work involving the stadium billiard and evanescent waves components of $2k$ or less. In this shape, the convergence with b does not follow

any obvious power-law, but for $b = 6$ I found the fraction error on F is $< 10^{-6}$ for wavefunctions composed of random (real) plane waves. If highly evanescent waves are present (up to $3k$), b needs to be slightly higher (about 10). Certainly the b of 30 quoted by [137] is excessive, considering the fact that their PWDM basis set only includes real plane waves. In the scaling method, I find that errors due to a smaller b cause spurious states (Section 6.3.3) to enter earlier than is necessary. However, the accuracy of the valid states found does not seem very dependent on b.

Further study on the accurate evaluation of these integrals is necessary; in particular a way to handle corners and 'kinks' which has high-order (preferably exponential) convergence is sought.

Appendix H

Boundary evaluation of matrix elements of Helmholtz solutions

Here I collect some useful tools for expressing wavefunction domain integrals purely on the boundary. The time-independent (constant-energy) wave equation in uniform space is the Helmholtz equation,

$$(\nabla^2 + k_a^2)a(\mathbf{r}) \;=\; 0, \tag{H.1}$$

where the Helmholtz solution $a(\mathbf{r})$ is in general a complex-valued function of space. At a given k, Green's theorem tell us that a Helmholtz solution in a region $\mathcal{D}$ is determined uniquely by its value or normal derivative on the enclosing surface Γ. Therefore domain integrals of the type

$$\langle a|\,\hat{O}\,|b\rangle_{\mathcal{D}} \;\equiv\; \int_{\mathcal{D}} d\mathbf{r}\; a^*(\mathbf{r})(\hat{O}b)(\mathbf{r}) \tag{H.2}$$

can be reduced to integrals of local quantities over Γ alone, when a and b are Helmholtz solutions. This is true whether the corresponding wavenumbers k_a and k_b are equal or unequal. This has a vast numerical advantage in the semiclassical limit: a reduction of the number of function evaluations by a factor of a few times $k\mathsf{L}$ is expected for simple geometries, for a typical system size L. Once an expression in terms of surface integrals is known, the technique of Appendix G can be used for the numerical evaluation. Note that since a and b are constant-energy solutions, (H.2) can be called a 'matrix element' of the operator $\hat{O}$ in the energy basis.

I collect together useful results for such evaluations, in the cases $\hat{O} = 1$ (corresponding to a simple overlap integral), $\hat{O} = \mathbf{r}$ (dipole operator), $\hat{O} = \nabla$ (momentum operator), and $\hat{O} = \mathbf{r}\cdot\nabla$ (differential dilation operator). The results apply to all spatial dimensions $d \geq 2$. For each operator the cases of equal and unequal wavenumbers need to be handled separately. I have maintained complete generality of the boundary conditions (BCs) on Γ, unless I state otherwise. This allows Γ to be chosen arbitrarily, in particular, independently of the location of wavefunction nodes, billiard walls, or other features. From these the special cases corresponding to Dirichlet, etc., BCs can be derived—I will present some of these simpler forms below. Throughout this appendix the symbol $a^*\leftrightarrow b$ will imply a duplication of the previous term except with a^* and b swapped.

Some of the results have already been known to other authors [28, 33], and some are new, especially the case of general d. For those that are known, my derivations are much simpler than the complicated vector manipulations used by those authors. The most powerful new techniques I have found to be the use of scaled solutions (see Section 6.1.1), and a matrix inversion trick of M. Haggerty. Both are presented below. Remember that the goal in what follows is to 'push' all volume terms onto the boundary.

H.1 Overlaps: $\hat{O} = 1$

H.1.1 Differing wavenumbers $k_a \neq k_b$

The following is an application of Green's theorem. I start with the vector identity

$$\nabla \cdot (a^* \nabla b) \;=\; a^* \nabla^2 b + \nabla a^* \cdot \nabla b. \tag{H.3}$$

Integrating over $\mathcal{D}$, applying Gauss' Theorem to the LHS and Helmholtz to the first term on the RHS gives

$$\oint_\Gamma d\mathbf{s}\, a^* \partial_n b \;=\; \int_\mathcal{D} d\mathbf{r}\, (-k_b^2 a^* b + \nabla a^* \cdot \nabla b). \tag{H.4}$$

Note that ∂_n is an abbreviation for $\mathbf{n} \cdot \nabla$, the local *normal derivative* operator at a location $\mathbf{s}$ on Γ, and should not be confused with the cartesian derivatives ∂_i, ∂_j, etc which appear later. Subtraction of the same equation with the swap $a^* \leftrightarrow b$ gives the useful overlap formula

$$\langle a \,|\, b \rangle_\mathcal{D} \;=\; \frac{1}{k_a^2 - k_b^2} \oint_\Gamma d\mathbf{s}\, (a^* \partial_n b - b \partial_n a^*). \tag{H.5}$$

This formula is undefined for $k_a = k_b$, where the integral vanishes. Note that if a and b are chosen to obey the same (general mixed) BCs, the formula gives zero as one would expect for orthogonal eigenstates.

H.1.2 Equal wavenumber $k_a = k_b = k$

The case $k_a = k_b$ can be reached most easily by setting $k_a = k$ and $k_b = k + \delta$ while taking the limit $\delta \to 0$. The solution b now has to be spatially *scaled* (dilated about the arbitrary origin) appropriately for its new wavenumber. Section 6.1.1 (also Appendix I) gives the expansion of the dilation operator at wavenumber k in powers of δ, thus

$$b(k + \delta; \mathbf{r}) \;=\; \left[1 + \frac{\delta}{k} \mathbf{r} \cdot \nabla + O(\delta^2) \cdots \right] b(k; \mathbf{r}). \tag{H.6}$$

Substitution into (H.5), noting $k_a^2 - k_b^2 = -2\delta k + O(\delta^2)$, and taking the limit (in which only the lowest order $O(\delta^0)$ survives) gives

$$\langle a \,|\, b \rangle_\mathcal{D} \;=\; \oint_\Gamma d\mathbf{s} \, \left[(\mathbf{n} \cdot \nabla a^*)(\mathbf{r} \cdot \nabla b) - a^* \mathbf{n} \cdot \nabla (\mathbf{r} \cdot \nabla b) \right]. \tag{H.7}$$

This can be shown to be a generalization of the overlap formula given by Boasman [33] to general d and differing solutions a and b. Despite the derivation using scaling, (H.7) is in

fact an identity, which can be proved using a messy algebra sequence[1]. Unlike the LHS, the RHS is not manifestly $a^* \leftrightarrow b$ symmetric. However, it can be written in the symmetrized form (H.14) derived in the next section. Notice that the choice $k_a = k_b$ has increased the order of derivative required on the boundary by one.

In the case of Dirichlet BCs the second term in (H.7) vanishes, and the replacement $\mathbf{r} \cdot \nabla b = (\mathbf{r} \cdot \mathbf{n})\mathbf{n} \cdot \nabla b$ (which follows from the fact that ∇b and $\mathbf{n}$ are parallel), gives

$$\langle a \,|\, b \,\rangle_{\mathcal{D}} \;=\; \frac{1}{2k^2} \oint_{\Gamma} d\mathbf{s}\, (\mathbf{r} \cdot \mathbf{n})\, \partial_n a^* \partial_n b, \qquad \text{(Dirichlet BCs).} \tag{H.8}$$

However the knowledge that the LHS is zero for $b \neq a$ (orthogonal eigenstates) shows that (H.8) is an exact boundary orthogonality relation for degenerate Dirichlet eigenstates. Note that the RHS is proportional to the matrix element $(\partial \mathcal{H}/\partial x)_{ab}$ of the billiard dilation deformation between degenerate states (Appendix C). On the other hand, choosing $b = a$ we have

$$\langle a \,|\, a \,\rangle_{\mathcal{D}} \;=\; \frac{1}{2k^2} \oint_{\Gamma} d\mathbf{s}\, (\mathbf{r} \cdot \mathbf{n})\, |\partial_n a|^2, \qquad \text{(Dirichlet BCs).} \tag{H.9}$$

This very useful boundary formula for the norm of Dirichlet eigenstates appears to have been found first for $d = 2$ by Berry and Wilkinson (appendix of [28]). It was since derived in a different way by Boasman[33], and for general d was derived in our work[14]. However, I believe the new derivation above to be the simplest yet.

H.2 Matrix trick for pushing integrals onto the boundary

I now present an algebraic technique for expressing integrals over $\mathcal{D}$ as surface integrals over Γ. This was invented, as far as I am aware, by Mike Haggerty (unpublished), in response to the ugly vector algebra performed by Boasman [33] needed to achieve the result (H.14) in $d = 2$. I have generalized this technique to the case of singular matrices, and to the tensor divergence theorem.

H.2.1 Volume integrals of scalar quantities

By taking the divergence of all the types of terms appearing in relevant vector identities, the following matrix relation is reached,

$$\nabla \cdot \begin{pmatrix} \mathbf{r} a^* b \\ a^* \nabla b \\ b \nabla a^* \\ \mathbf{r} \nabla a^* \nabla b \\ a^* (\mathbf{r} \cdot \nabla) \nabla b \\ b (\mathbf{r} \cdot \nabla) \nabla a^* \end{pmatrix} = \begin{pmatrix} d & 1 & 1 & 0 & 0 & 0 \\ -k_b^2 & 0 & 0 & 1 & 0 & 0 \\ -k_a^2 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & d & 1 & 1 \\ -k_b^2 & -k_b^2 & 0 & 0 & 1 & 0 \\ -k_a^2 & 0 & -k_a^2 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} a^* b \\ a^* \mathbf{r} \cdot \nabla b \\ b \mathbf{r} \cdot \nabla a^* \\ \nabla a^* \cdot \nabla b \\ \nabla a^* \cdot (\mathbf{r} \cdot \nabla) \nabla b \\ \nabla b \cdot (\mathbf{r} \cdot \nabla) \nabla a^* \end{pmatrix}. \tag{H.10}$$

[1] This will not be given, but briefly it involves expanding the second term, use of Gauss' theorem and the identity $\nabla \cdot (\mathbf{r} \cdot \nabla) \nabla b = -k^2 (1 + \mathbf{r} \cdot \nabla) b$. Meanwhile the first term is compared to the volume integral of $\nabla \cdot [\nabla a^* (\mathbf{r} \cdot \nabla b)]$ and its expansion, using the identity $\nabla (\mathbf{r} \cdot \nabla b) = (1 + \mathbf{r} \cdot \nabla) \nabla b$. Both these identities require Einstein summation notation to prove. Bypassing such a tricky maze of formulae is one purpose of this appendix.

The scalar coefficient matrix is formed by finding the linear amounts of each of the 6 scalar functions which result from performing the divergences on each vector term. Deducing a few of the entries requires Einstein summation form, and substitution of the Helmholtz equation, but all are simple. It is convenient that 6 inputs gives 6 outputs since the matrix is then possibly invertible. Its determinant is found to be $(k_a^2 - k_b^2)^2$, implying that it is invertible for $k_a \neq k_b$. The relation (H.10) can be written

$$\nabla \cdot \mathbf{v}_\alpha \;=\; \mathcal{M}_{\alpha\beta}\, s_\beta, \tag{H.11}$$

where the sum over β is implied. Care should be taken not to confuse spatial coordinates (boldface vectors) with coefficient indices $\alpha, \beta = 1 \cdots 6$ (Greek letters). Integrating over the domain $\mathcal{D}$, applying Gauss' theorem and premultiplying by the inverse of $\mathcal{M}$ gives

$$\int_{\mathcal{D}} d\mathbf{r}\, s_\alpha \;=\; (\mathcal{M}^{-1})_{\alpha\beta} \oint_\Gamma d\mathbf{s}\, \mathbf{n} \cdot \mathbf{v}_\beta. \tag{H.12}$$

Thus we have closed-form expressions for the volume integrals of any of the 6 scalar functions, in terms of boundary integrals. That is, we have 'pushed' the problem onto the boundary, as desired.

The symbolic inverse can be found (*e.g.* using the *Mathematica* software package),

$$\mathcal{M}^{-1} \;=\; \frac{1}{\varepsilon} \begin{pmatrix} 0 & 1 & -1 & 0 & 0 & 0 \\ k_a^2 & (k_a^2+k_b^2)/\varepsilon & d-(k_a^2+k_b^2)/\varepsilon & -1 & 1 & 1 \\ -k_b^2 & -d-(k_a^2+k_b^2)/\varepsilon & (k_a^2+k_b^2)/\varepsilon & 1 & -1 & -1 \\ 0 & k_a^2 & -k_b^2 & 0 & 0 & 0 \\ k_a^2 k_b^2 & 2k_a^2 k_b^2/\varepsilon & k_b^2(d-2k_a^2/\varepsilon) & -k_b^2 & k_a^2 & k_b^2 \\ -k_a^2 k_b^2 & k_a^2(-d-2k_b^2/\varepsilon) & 2k_a^2 k_b^2/\varepsilon & k_a^2 & -k_a^2 & -k_b^2 \end{pmatrix}, \tag{H.13}$$

where $\varepsilon \equiv k_a^2 - k_b^2$ is the 'energy difference'. Rows of this matrix give the desired expressions for volume integrals. For instance, (H.5) results from row 1.

When $k_a = k_b = k$ rows 2 and 3 of $\mathcal{M}$ become identical, so the matrix is singular and the above inversion fails. However, it is still possible to extract solutions to this singular problem by finding a vector ξ such that $\mathcal{M}^T \xi = \mathbf{e}$ for a desired unit vector $\mathbf{e}$. This is only possible if the unit vector lies entirely within the *row space* of M [188]. For instance, the first unit vector $e_\alpha = \delta_{\alpha 1}$ gives $\xi = (k^2, d/2, d/2, -1, 1, 1)^T$, corresponding to

$$\begin{aligned} \langle a\,|b\,\rangle_{\mathcal{D}} \;=\; & -\frac{1}{2k^2} \oint_\Gamma d\mathbf{s}\, \Big[\tfrac{1}{2} d(a^* \partial_n b + b \partial_n a^*) + (\mathbf{n}\cdot\mathbf{r})(k^2 a^* b - \nabla a^* \cdot \nabla b) \\ & + a^* \mathbf{n} \cdot (\mathbf{r} \cdot \nabla) \nabla b + a^* \leftrightarrow b \Big]. \end{aligned} \tag{H.14}$$

This is a manifestly symmetric form of (H.7). Its derivation is much more routine than that of Boasman[33] because the matrix method automatically handles the coefficients.

H.2.2 Volume integrals of vector quantities

A tensor divergence analogy of the above trick can easily be found. Using Einstein summation convention, one can write tensors on the LHS whose divergence gives vectors

on the RHS,

$$\partial_i \begin{pmatrix} \delta_{ij} a^* b \\ r_j a^* \partial_i b \\ r_j b \partial_i a^* \\ a^* \partial_i \partial_j b \\ b \partial_i \partial_j a^* \\ \delta_{ij} \partial_m a^* \partial_m b \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & -k_b^2 & 1 & 0 & 0 \\ 0 & 1 & -k_a^2 & 1 & 0 & 0 \\ -k_b^2 & 0 & 0 & 0 & 1 & 0 \\ 0 & -k_a^2 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} a^* \partial_j b \\ b \partial_j a^* \\ r_j a^* b \\ r_j \partial_m a^* \partial_m b \\ (\partial_m a^*) \partial_m \partial_j b \\ (\partial_m b) \partial_m \partial_j a^* \end{pmatrix}. \tag{H.15}$$

Notice that no dependence on d has entered. The relation can be written

$$\partial_i t_{ij}^{(\alpha)} \;=\; \mathcal{T}_{\alpha\beta}\, w_j^{(\beta)}. \tag{H.16}$$

Remember that Roman indices $i, j, m = 1 \cdots d$ are spatial, whereas Greek indices $\alpha, \beta = 1 \cdots 6$ label coefficients. Note that the vector of vectors $\mathbf{w}^{(\beta)}$ does contain some different functions than $\mathbf{v}_\beta$ from Section H.2.1. The matrix determinant is $\det \mathcal{T} = \varepsilon^2$, the same as for $\mathcal{M}$. The tensor Gauss' theorem is

$$\int_{\mathcal{D}} d\mathbf{r}\, \partial_i A_{ij} \;=\; \oint_\Gamma d\mathbf{s}\, n_i A_{ij}. \tag{H.17}$$

Applying this to the domain integral of (H.16) gives

$$\int_{\mathcal{D}} d\mathbf{r}\, w_j^{(\alpha)} \;=\; (\mathcal{T}^{-1})_{\alpha\beta} \oint_\Gamma d\mathbf{s}\, n_i\, t_{ij}^{(\beta)}, \tag{H.18}$$

which are the desired closed-form expressions for volume integrals of (the j component of) the vector functions $\mathbf{w}^{(\alpha)}$. The symbolic inverse is found to be,

$$\mathcal{T}^{-1} \;=\; \frac{1}{\varepsilon} \begin{pmatrix} k_a^2 & 0 & 0 & 1 & 1 & -1 \\ -k_b^2 & 0 & 0 & -1 & -1 & 1 \\ -(k_a^2 + k_b^2)/\varepsilon & 1 & -1 & -2/\varepsilon & -2/\varepsilon & 2/\varepsilon \\ -(k_a^4 + k_b^4)/\varepsilon & k_a^2 & -k_b^2 & -(k_a^2 + k_b^2)/\varepsilon & -(k_a^2 + k_b^2)/\varepsilon & (k_a^2 + k_b^2)/\varepsilon \\ k_a^2 k_b^2 & 0 & 0 & k_a^2 & k_b^2 & -k_b^2 \\ -k_a^2 k_b^2 & 0 & 0 & -k_a^2 & -k_b^2 & k_a^2 \end{pmatrix}. \tag{H.19}$$

Again, singular solutions can still be found when $k_a = k_b$. The full consequences of these types of matrix relations have not been explored.

H.3 Other matrix elements

Results are quoted below for other simple choices of $\hat{O}$. They are believed to be new results. They arise without effort from the matrix trick of the previous section.

H.3.1 Momentum matrix elements: $\hat{O} = \nabla$

For $k_a \neq k_b$ we have for each component $j = 1 \ldots d$ of the grad operator,

$$\langle a |\, \partial_j \,| b \rangle_{\mathcal{D}} \;=\; \frac{1}{k_a^2 - k_b^2} \oint_\Gamma d\mathbf{s} \left\{ a^* n_m \partial_m \partial_j b + a^* \leftrightarrow b + n_j \left[k_a^2 a^* b - (\partial_m a^*) \partial_m b \right] \right\}. \tag{H.20}$$

This follows from row 1 of the matrix $\mathcal{T}^{-1}$. Summation over m is implied.

For $k_a = k_b = k$ it is unambiguous to write all components together as a vector relation thus,

$$\langle a|\, \nabla\, |b\rangle_{\mathcal{D}} \;=\; \frac{1}{2}\oint_\Gamma d\mathbf{s}\; [\mathbf{n} a^* b + \mathbf{r}(a^*\partial_n b - b\partial_n a^*)]\,. \tag{H.21}$$

This follows from the singular solution $\mathcal{T}^T\xi = \mathbf{e}$ for the first unit vector $\mathbf{e}$. It is interesting that the equal wavenumber case involves one order of derivative less than the unequal case. In the Dirichlet BC case, (H.21) vanishes.

H.3.2 Dipole matrix elements: $\hat{O} = \mathbf{r}$

For $k_a \neq k_b$ we have for each component $j = 1 \ldots d$ of the dipole (position) operator,

$$\begin{aligned}\langle a|\, r_j\, |b\rangle_{\mathcal{D}} \;=\; \frac{1}{\varepsilon^2}\oint_\Gamma d\mathbf{s}\; \Big\{ & n_j\left[2(\partial_m a^*)\partial_m b - (k_a^2 + k_b^2)a^* b\right] + \varepsilon r_j (a^* n_m \partial_m b - a^* \leftrightarrow b) \\ & -\; 2(a^* n_m \partial_m \partial_j b + a^* \leftrightarrow b)\Big\}, \end{aligned} \tag{H.22}$$

where as before $\varepsilon \equiv k_a^2 - k_b^2$. This follows from row 3 of the matrix $\mathcal{T}^{-1}$. Summation over m is implied. A mixture of vector and Einstein notations has been avoided. In the Dirichlet BC case this becomes

$$\langle a|\, r_j\, |b\rangle_{\mathcal{D}} \;=\; \frac{2}{\varepsilon^2}\oint_\Gamma d\mathbf{s}\, n_j\, \partial_n a^* \partial_n b, \qquad \text{(Dirichlet BCs).} \tag{H.23}$$

Note that the integral is proportional to the matrix element $(\partial\mathcal{H}/\partial x)_{ab}$ of the billiard deformation corresponding to translation in the j direction.

No form for $k_a = k_b$ has been found. The corresponding unit vector $e_\alpha = \delta_{\alpha 3}$ does not lie in the row space of $\mathcal{T}^{-1}$, indicating that other boundary derivatives are needed as input.

H.3.3 Differential dilation matrix elements: $\hat{O} = \mathbf{r}\cdot\nabla$

For $k_a \neq k_b$ we have for the differential dilation operator,

$$\begin{aligned}\langle a|\, \mathbf{r}\cdot\nabla\, |b\rangle_{\mathcal{D}} \;=\; \frac{1}{\varepsilon}\oint_\Gamma d\mathbf{s}\; \Big\{ & (\mathbf{r}\cdot\mathbf{n})(k_a^2 a^* b - \nabla a^*\cdot\nabla b) + \frac{k_a^2 + k_b^2}{\varepsilon}(a^*\partial_n b - b\partial_n a^*) \\ & + d\, b\partial_n a^* + a^*\mathbf{n}\cdot(\mathbf{r}\cdot\nabla)\nabla b + a^*\leftrightarrow b\Big\}, \end{aligned} \tag{H.24}$$

where again $\varepsilon \equiv k_a^2 - k_b^2$. This follows from row 2 of the matrix $\mathcal{M}^{-1}$. In the Dirichlet BC case this becomes,

$$\langle a|\, \mathbf{r}\cdot\nabla\, |b\rangle_{\mathcal{D}} \;=\; -\frac{1}{\varepsilon}\oint_\Gamma d\mathbf{s}\, (\mathbf{r}\cdot\mathbf{n})\, \partial_n a^* \partial_n b, \qquad \text{(Dirichlet BCs).} \tag{H.25}$$

The integral is proportional to the matrix element $(\partial\mathcal{H}/\partial x)_{ab}$ of the billiard dilation deformation.

No general form for $k_a = k_b$ has been found. The corresponding unit vector $e_\alpha = \delta_{\alpha 2}$ does not lie in the row space of $\mathcal{M}^{-1}$, indicating that other boundary derivatives are needed as input. However, for the particular case of Dirichlet BCs and $b = a$ there exists the simple formula

$$\langle a|\, \mathbf{r}\cdot\nabla\, |a\rangle_{\mathcal{D}} \;=\; -\frac{d}{2}\,\langle a\,|a\,\rangle_{\mathcal{D}}\,, \qquad \text{(Dirichlet BCs)}, \tag{H.26}$$

which can be derived by expanding $\nabla\cdot(\mathbf{r}|a|^2)$ and integrating over the domain.

Appendix I

Scaling expansion of eigenfunctions and tension matrix

Here I derive the expansion of the value of a scaling eigenfunction on the unscaled boundary Γ, in powers of the wavenumber shift δ. Section 6.1.1 is a necessary prerequisite. The expansion will be in terms of normal and tangential derivatives of the unscaled eigenfunction on Γ. Dirichlet boundary conditions will be assumed for this eigenfunction. From this will follow the expansion of the tension on Γ, in the scaling eigenfunction basis. This is a vital part of the understanding of the scaling method of Vergini and Saraceno (VS).

It is my belief that although the lowest-order term found by those authors [195, 194] is correct (giving the basic explanation of the scaling method), the higher-order terms are not, because the curvature (metric) of the boundary was not taken into account. My aim here is to correct this oversight. This will affect the understanding of the growth of errors in the method. I will stick to $d = 2$, although the generalization to higher dimension is believed to be simply a matter of introducing vector notation for the tangential coordinate.

I.1 Expansion of the dilation operator

A *scaling wavefunction* ψ (see Eq.(6.1)) whose unscaled wavenumber is k_μ will be rescaled to a new wavenumber $k = k_\mu + \delta$, where δ is the wavenumber shift. Its value at location $\mathbf{r}$ is

$$\psi(k_\mu + \delta, \mathbf{r}) \;=\; \psi(\frac{k_\mu + \delta}{k_\mu}\mathbf{r}) \;=\; \psi(\mathbf{r} + \mathbf{a}), \tag{I.1}$$

where the translation vector is $\mathbf{a} = (\delta/k_\mu)\mathbf{r}$. Formally this translation can be performed by exponentiation of the generator of translation [174, 7]:

$$\psi(k_\mu + \delta, \mathbf{r}) \;=\; e^{\mathbf{a}\cdot\nabla}\psi \;=\; \left[1 + \mathbf{a}\cdot\nabla + \frac{1}{2!}(\mathbf{a}\cdot\nabla)(\mathbf{a}\cdot\nabla) + \cdots\right]\psi. \tag{I.2}$$

Thus we have written the scaled value in terms of local derivatives of the unscaled wavefunction. This becomes a definition of *dilation* only when the $\mathbf{r}$-dependence of $\mathbf{a}$ is finally substituted. Note that $\mathbf{a}$ is held constant as far as ∇ is concerned. (This differs from other possible definitions of dilation [14], where the commutator $[\partial_i, r_j] = \delta_{ij}$ plays a role; however

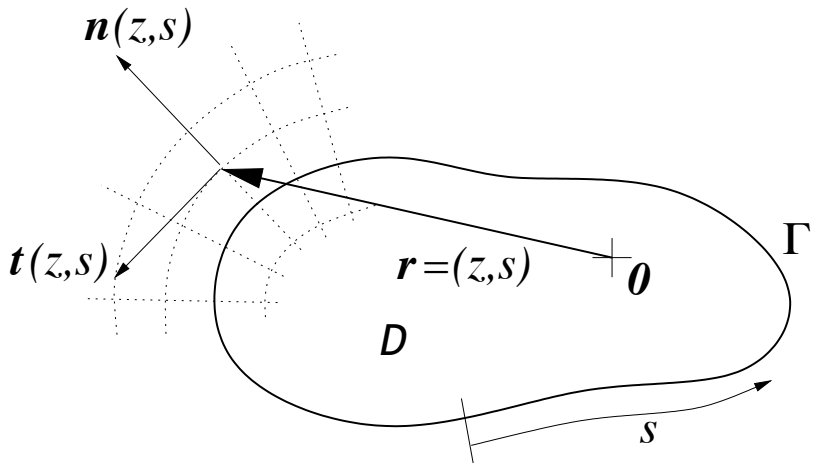

Figure I.1: Curvilinear coordinates (z, s) used to represent points $\mathbf{r}$ near the billiard boundary Γ. The orthogonal unit vectors are the vector fields $\mathbf{n}$ and $\mathbf{t}$.

my form will be simpler to expand to high order). For instance in Cartesian coordinates, $(\mathbf{a}\cdot\nabla)(\mathbf{a}\cdot\nabla)$ is to be interpreted as $a_i a_j \partial_i \partial_j$ (Einstein summation assumed), because in Cartesian coordinates $\partial_i a_j = 0$.

I.2 Curvilinear boundary coordinates

We express $\mathbf{r}$ for $d = 2$ in the form (z, s) where z is the outward normal coordinate, and s is a (periodic) tangential location on the boundary Γ, increasing in an anti-clockwise fashion (see Fig. I.1). I will assume Γ is a curve differentiable to sufficient order. The corresponding locally-orthonormal unit vectors are $\mathbf{n}$ (normal) and $\mathbf{t}$ (tangential). In our curvilinear coordinates they are vector fields defined by $\mathbf{n} \equiv \partial\mathbf{r}/\partial z$ and $\mathbf{t} \equiv \partial\mathbf{r}/\partial s$. For an introduction to such noncartesian systems see [7]. For the local components of the gradient of a scalar field I will use the abbreviations $\partial\psi/\partial z \equiv \partial_n\psi \equiv \psi^n$ in the normal direction, and $\partial\psi/\partial s \equiv \partial_t\psi \equiv \psi^t$ in the tangential direction. A general vector field $\mathbf{b}$ will have the components (b_n, b_t) with $b_n \equiv \mathbf{n}\cdot\mathbf{b}$ and $b_t \equiv \mathbf{t}\cdot\mathbf{b}$.

Our task is to express (I.2) in terms of ψ, ψ^n, ψ^t, and higher derivatives ψ^{nn}, ψ^{nt} etc. We will choose a point $\mathbf{r} \in \Gamma$, so that these derivatives are evaluated at the boundary. The boundary conditions will later be applied. We need rules for handling derivatives. The metric is such that derivatives of the unit vectors obey

$$\partial_n\mathbf{n} = \mathbf{0}, \qquad \partial_n\mathbf{t} = \mathbf{0}, \tag{I.3a}$$

$$\partial_t\mathbf{n} = \alpha\mathbf{t}, \qquad \partial_t\mathbf{t} = -\alpha\mathbf{n}, \tag{I.3b}$$

on the boundary ($z = 0$). The local inverse radius of curvature of the boundary is

$$\alpha(s) \equiv \frac{1}{R(s)} \equiv -\mathbf{n}(s)\cdot\frac{\partial\mathbf{t}}{\partial s}. \tag{I.4}$$

Thus α gives the 'connection' of the metric (or Christoffel symbol [7]). This means that while derivatives of scalars obey the usual Cartesian rules, derivatives of vectors (or their local components) introduce extra terms. This key observation was omitted in the work of VS [195, 194]. For a general vector field $\mathbf{b}$ we have,

$$\begin{aligned}
\partial_n b_n &\equiv \partial_n(\mathbf{n}\cdot\mathbf{b}) = \mathbf{n}\cdot\partial_n\mathbf{b} && \text{(I.5)}\\
\partial_n b_t &\equiv \partial_n(\mathbf{t}\cdot\mathbf{b}) = \mathbf{t}\cdot\partial_n\mathbf{b} && \text{(I.6)}\\
\partial_t b_n &\equiv \partial_t(\mathbf{n}\cdot\mathbf{b}) = \mathbf{n}\cdot\partial_t\mathbf{b} + \alpha b_t && \text{(I.7)}\\
\partial_t b_t &\equiv \partial_t(\mathbf{t}\cdot\mathbf{b}) = \mathbf{t}\cdot\partial_t\mathbf{b} - \alpha b_n. && \text{(I.8)}
\end{aligned}$$

The four derivative terms on the RHS are the components of the *covariant derivative* tensor. If we apply the above to the choice $\mathbf{b} = \mathbf{a}$ which we take to be a constant (translation) field according to the prescription of the previous section, then all the derivatives of $\mathbf{a}$ vanish leaving

$$\begin{aligned}
\partial_n a_n &= \partial_n a_t = 0, && \text{(I.9a)}\\
\partial_t a_n &= \alpha a_t, && \text{(I.9b)}\\
\partial_t a_t &= -\alpha a_n, && \text{(I.9c)}
\end{aligned}$$

which we will use later. We will also need the divergence of a general vector field $\mathbf{b}$, which is *not* simply $\partial_n b_n + \partial_t b_t$, but rather

$$\nabla\cdot\mathbf{b} = \mathbf{n}\cdot\partial_n\mathbf{b} + \mathbf{t}\cdot\partial_t\mathbf{b} = \partial_n b_n + \partial_t b_t + \alpha b_n, \qquad \text{(I.10)}$$

the trace of the covariant derivative tensor. Choosing $\mathbf{b} = \nabla\psi$ finally gives the Laplacian

$$\nabla^2\psi = \nabla\cdot\nabla\psi = \psi^{nn} + \psi^{tt} + \alpha\psi^n, \qquad \text{(I.11)}$$

equivalent to the form in locally-cylindrical coordinates.

I.3 Applying boundary conditions and simplifying

We now proceed term-by-term in (I.2). The zeroth-order is

$$\psi \xrightarrow{\text{BCs}} 0, \qquad \text{(I.12)}$$

where the arrow implies application of the Dirichlet boundary conditions (BCs), namely $\psi(z{=}0, s) = 0$ for all s (remember that we already chose $z = 0$ in order to be able to write the metric). The first order is

$$\mathbf{a}\cdot\nabla\psi = a_n\psi^n + a_t\psi^t \xrightarrow{\text{BCs}} a_n\psi^n, \qquad \text{(I.13)}$$

where the result $\psi^t = 0$ was used. In will use in future the fact that all higher t-derivatives $\psi^{tt\cdots t}$ vanish upon application of the BCs. The second order requires application of the rules

(I.9):

$$\begin{aligned}\frac{1}{2!}(\mathbf{a}\cdot\nabla)(\mathbf{a}\cdot\nabla)\psi &= \frac{1}{2}(a_n\partial_n + a_t\partial_t)(a_n\psi^n + a_t\psi^t) \\ &= \frac{1}{2}\left[a_n^2\psi^{nn} + 2a_n a_t\psi^{nt} + a_t^2\psi^{tt} + \alpha a_t(a_t\psi^n - a_n\psi^t)\right] \\ &\stackrel{\text{BCs}}{\longrightarrow} a_n a_t\psi^{nt} + \frac{\alpha}{2}(a_t^2 - a_n^2)\psi^n. \end{aligned} \tag{I.14}$$

The Helmholtz equation and the BCs together imply $\psi^{nn} = -\alpha\psi^n$, a relation used both above and below. In general our task is to reduce the number of n-derivatives; it is always possible using such manipulations to leave only terms containing a single n-derivative. This will be desirable for manipulation of boundary integrals by parts later.

The third-order result (included below) requires use of the following. $\psi^{nnn} = \partial_n(-\psi^{tt} - \alpha\psi^n - k_\mu^2\psi)$ simplifies to $-\psi^{ntt} + (2\alpha^2 - k_\mu^2)\psi^n$ when the BCs are applied. This required the normal derivative of curvature $\partial_n\alpha = -\alpha^2$ since α now needs to be regarded as a scalar field with z-dependence. The tangential derivative is given the name $\partial_t\alpha = \alpha'$. Also $\psi^{nnt} = \partial_t(-\psi^{tt} - \alpha\psi^n - k_\mu^2\psi)$ simplifies to $-\alpha'\psi^n - \alpha\psi^{nt}$ when BCs are applied.

Combining everything and finally substituting for $\mathbf{a}$ gives the expansion of a Dirichlet scaling eigenfunction at location s on Γ:

$$\begin{aligned}\psi(k_\mu + \delta, s) &= \frac{\delta}{k_\mu} r_n\psi^n \\ &\quad + \frac{\delta^2}{k_\mu^2}\left[\alpha\frac{r_t^2 - r_n^2}{2}\psi^n \;+\; r_n r_t\psi^{nt}\right] \\ &\quad - \frac{\delta^3}{6k_\mu} r_n^3\psi^n \;+\; \frac{\delta^3}{6k_\mu^3}\left\{\left[2\alpha^2 r_n(r_n^2 - 3r_t^2) + \alpha' r_t(r_t^2 - 3r_n^2)\right]\psi^n \right. \\ &\qquad\qquad \left. +\; 3\alpha r_t(r_t^2 - 2r_n^2)\psi^{nt} + r_n(3r_t^2 - r_n^2)\psi^{ntt}\right\} \\ &\quad + O(\delta^4)\cdots \end{aligned} \tag{I.15}$$

The expression is believed to be correct to order δ^3. The complexity increases greatly with each power of δ, and higher derivatives of the curvature α enter. However, in the case of the circle billiard (where α is constant) it is easy to verify the expansion. One general pattern is that the m^{th}-order terms all contain m powers of $r\delta$, and for each term the number of spatial derivatives of ψ must equal the power of $1/k_\mu$ minus the power of α. It is important to notice that the $O(\delta^2)$ term differs by the presence of a factor of α from the $O(\delta^2)$ term given by VS.

I.4 Tension matrix expansion

The tension matrix (using the 'dilation' weighting function $w(s) = 1/r_n$) in the scaling eigenfunction basis has elements

$$\tilde{F}_{\mu\nu}(k) \;=\; \oint \frac{ds}{r_n}\,\psi_\mu(k_\mu + \delta_\mu)\,\psi_\nu(k_\nu + \delta_\nu), \tag{I.16}$$

where the wavenumber shifts for the two states involved are $\delta_\mu \equiv k - k_\mu$ and $\delta_\nu \equiv k - k_\nu$. As usual I define the wavenumber difference $\kappa \equiv k_\mu - k_\nu = \delta_\nu - \delta_\mu$. Substitution of (I.15) gives a power series in the δ's, which we will study and make estimates for unknown boundary integrals to grasp the general structure. We will assume for simplicity that the billiard size $\mathsf{L} \approx 1$, and remain in $d = 2$. We will ignore differences in k_μ and k_ν from k when appropriate, and look only for the dominant off-diagonal terms.

The $O(\delta^2)$ term in $\tilde{F}$ comes from the lowest order in (I.15),

$$\tilde{F}^{(2)}_{\mu\nu}(k) \;=\; \frac{\delta_\mu \delta_\nu}{k_\mu k_\nu} \oint ds\, r_n\, \psi^n_\mu \psi^n_\nu \tag{I.17}$$

which we can write as $2\delta_\mu \delta_\nu M_{\mu\nu}$ (see Section 6.1.1), where M is quasi-diagonal. The diagonal is $M_{\mu\mu} = 1$, and we have from the considerations of Section 6.1.2 small off-diagonal elements of size $M_{\mu\nu} \sim k^{-1/2}\kappa^2 O(1)$, for $d = 2$.

The $O(\delta^3)$ term in $\tilde{F}$ is, after some rearrangement,

$$\tilde{F}^{(3)}_{\mu\nu}(k) \;=\; -\frac{2\delta_\mu^3}{k} M_{\mu\nu} - \frac{\delta_\mu \delta_\nu}{2k^3} \kappa \oint ds\, r_n r_t\, (\psi^n_\nu \psi^{nt}_\mu - \psi^n_\mu \psi^{nt}_\nu), \tag{I.18}$$

where an integral of $r_n r_t \partial_t (\psi^n_\mu \psi^n_\nu)$ was performed by parts. This required use of $\partial_t(r_n r_t) = r_n + \alpha(r_n^2 - r_t^2)$, a different expression than found by VS. It can be proved easily using $\partial_t r_n = \alpha r_t$ and $\partial_t r_t = 1 - \alpha r_n$, which follow from (I.7) and (I.8). Remarkably, the α dependence then cancels out, giving the same $\tilde{F}^{(3)}$ diagonal term as that of VS. We believe that the integral in (I.18) does not have any quasi-orthogonal property, so can be estimated using random waves. The estimate gives $k^{5/2}O(1)$ for this integral, and shows that this term dominates over any off-diagonal contribution from the first term (involving M). The factor of κ in this term means that there is a weak form of quasi-diagonality at this order. Importantly, for $\kappa \ll 1$ the off-diagonal $O(\delta^3)$ error renders the $O(\delta^2)$ error insignificant. Hence we expect the quasi-diagonality property of M to play no role in errors in the scaling method.

The $O(\delta^4)$ and higher terms in F become very messy. I believe that the dominant $O(\delta^4)$ terms, both on and off the diagonal, are

$$\tilde{F}^{(4)}_{\mu\nu}(k) \;=\; -\frac{\delta_\mu \delta_\nu (\delta_\mu^2 + \delta_\nu^2)}{6k^2} \oint ds\, r_n^3\, \psi^n_\mu \psi^n_\nu \;+\; \frac{\delta_\mu^2 \delta_\nu^2}{k^4} \oint ds\, r_n^2 r_t^2\, \psi^{nt}_\mu \psi^{nt}_\nu. \tag{I.19}$$

This can be seen by comparing powers of k and using random-wave estimates. A random-wave estimate of the integrals then gives $\tilde{F}^{(4)} \approx \delta_\mu^4 O(1)$ on the diagonal and $\delta_\mu^4 k^{-1/2} O(1)$ off-diagonal. Note that the off-diagonal has no quasi-diagonality property at this or higher orders. For higher orders $O(\delta^m)$ for m even, we expect $\tilde{F}^{(m)} \approx \delta_\mu^m O(1)$ on the diagonal and $O(\delta^m) k^{-1/2} O(1)$ off-diagonal. For m odd, the leading diagonal terms are down by a factor of k which renders them insignificant.

To summarize, the diagonal of the tension matrix has the form given in Eq.(6.27), and for $\delta_\mu, \delta_\nu \ll 1$ the dominant off-diagonal terms are

$$\tilde{F}_{\mu\nu}(k) \;\approx\; \frac{\delta_\mu \delta_\nu}{k^{1/2}} \Big[\kappa O(1) + (\delta_\mu^2 + \delta_\nu^2) O(1) + \delta_\mu \delta_\nu O(1) \Big]. \tag{I.20}$$

Here contributions from orders 3 and 4 were included because it may be that the 4th order (the lowest order with *no* quasi-diagonality, *i.e.* no powers of κ) contributes most to errors in the scaling method. It is important to note that the 2nd order term (due to off-diagonal strength of M) is negligible. It is clear that more research is needed on the properties of the higher-order terms, especially if an explanation of the tension error δ^6 power-law growth (Section 6.3.1) is sought.

I.5 Useful geometric boundary integrals

Here I give a couple of geometrical results in billiards which found no other place to hide. In arbitrary d we have

$$\oint_\Gamma d\mathbf{s}\, r_n \;=\; 2\mathsf{V}, \tag{I.21}$$

where as usual V is the billiard volume. For $d = 2$ we have

$$\oint ds\, r_t \;=\; 0, \tag{I.22}$$

$$\oint ds\, \alpha \;=\; 2\pi, \tag{I.23}$$

where the local curvature $\alpha(s)$ is defined in Section I.2.

Appendix J

Helmholtz basis functions for two-dimensional billiards

We are interested in a set of basis functions appropriate to represent to sufficient accuracy the eigenstates of a billiard close to a given wavenumber k. The basis functions are solutions to the Helmholtz equation (5.5) at wavenumber k (so that the diagonalization problem can be reduced to the boundary). The basis need not obey the boundary conditions (BCs)—if it did, we would already have the desired eigenstates—so we can choose to ignore the BCs and consider the basis as composed of free-space solutions. The exception will be when the billiard shape has symmetries: then the basis should obey certain BCs on the lines (or planes) of symmetry (see below in Section J.4). Since the stadium is the quantum system that I have investigated most, the basis sets presented here are optimized for that shape. The issue of Helmholtz basis sets for more general shapes is open, and is sorely in need of investigation.

Note that because the scaling method of Chapter 6 requires it, I consider only the special case of *scaling basis sets*,

$$\phi_i(k;\mathbf{r}) \;=\; \bar{\phi}_i(k\mathbf{r}) \qquad\qquad \forall i, \tag{J.1}$$

where the dimensionless position will be written $k\mathbf{r} \equiv \bar{\mathbf{r}} \equiv (\bar{x}, \bar{y})$. This is no great restriction for the plane-wave type basis we present here, however, it disallows basis functions (*e.g.* singular $Y_0(kr)$) 'pinned' to the boundary as is required in the BIM.

J.1 Basis 'badness' γ

Generally the best (in the least squares sense in the linear Helmholtz function space) coefficient vector $\mathbf{x}$ able to represent a true (domain-normalized) eigenstate has exponentially large coefficients. This can be expressed (following Vergini [194]) in terms of a basis 'badness',

$$\gamma \;\equiv\; |\mathbf{x}|^2, \tag{J.2}$$

which diverges exponentially once the basis set size N is increased beyond N_{sc}. Numerically, because of the null-space truncation of Section 5.3.2, γ has an upper limit of $\epsilon^{-1} \sim 10^14$.

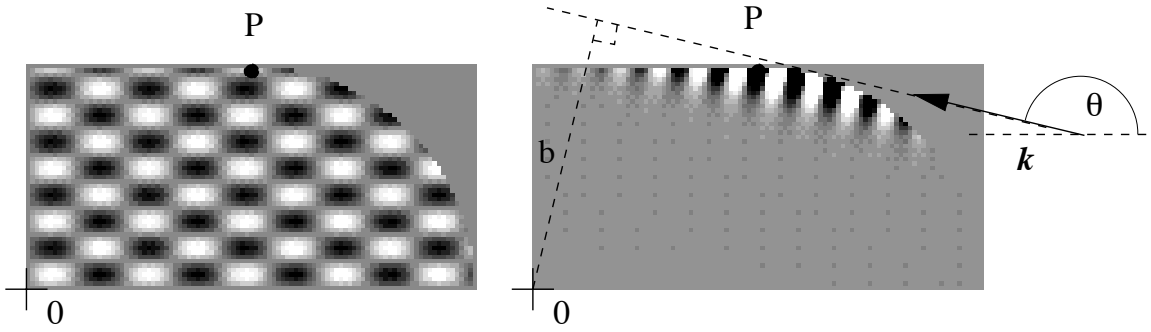

Figure J.1: Basis functions in the stadium billiard. **Left:** Real plane wave, and **Right:** Evanescent plane wave (showing geometry). Both are symmetrized, and are shown at $k = 30$.

This issue is identical to the divergence of basis coefficients found by Dietz *et al.* when attempting to expand a 'cake billiard' eigenfunction in terms of plane waves (or angular-momentum states) [63]. The problem is that of representing evanescent components, which are generally present in billiard eigenfunctions, by distributions of real plane waves. As Berry has shown [26], rapidly-oscillating distributions are required which grow exponentially with the number of wavelengths across the system and with the evanescence parameter. This was well understood by Vergini, who demonstrated that improving the basis by adding a small number of evanescent plane waves (in the case of the stadium only) enabled γ to reach $O(1)$. This improved the basis tension error $\epsilon_{\rm typ}$ from about 10^{-6} to about 10^{-11}.

In general the requirements of a good basis are small $\epsilon_{\rm typ}$ and, if possible, $\gamma \sim O(1)$ for most states. I first give the real plane waves, which work well in many billiard shapes, then give Vergini's evanescent plane waves which currently are well-adapted only to the stadium billiard (and shapes close to it).

J.2 Real plane waves (RPWs)

This is the default basis set:

$$\bar{\phi}_{2i-1}(\bar{\mathbf{r}}) = \cos(\mathbf{n}_i \cdot \bar{\mathbf{r}}) \tag{J.3}$$

$$\bar{\phi}_{2i}(\bar{\mathbf{r}}) = \sin(\mathbf{n}_i \cdot \bar{\mathbf{r}}), \tag{J.4}$$

for $i = 1 \cdots N/2$, where $\mathbf{n}_i$ is the unit direction vector at angle θ_i. The angular distribution can be uniform in the range $[0, \pi]$, giving $\theta_i = 2\pi(i - \frac{1}{2})/N$. A slight improvement results if there is more concentration in the directions perpendicular to the billiard's largest distances from the origin. In the 2-by-4 stadium (following [195]) I used $\theta_i = \frac{1}{4}[5 - 4i/N] \cdot 2\pi(i - \frac{1}{2})/N$.

Note that Li[137] claims he can half the RPW basis set size required, by randomizing phases (removing the need for both cos and sin functions). From our considerations in Section 5.3.1, it seems unlikely that he has somehow defeated the semiclassical basis size requirement by a factor of 2, however this is an area for investigation.

J.3 Adding evanescent plane waves (EPWs)

Evanescent plane waves (EPWs) are given a good introduction in [26]. Vergini [194] developed a set of EPWs for the stadium which involved 'families' with increasing evanescence parameter, which he arranged to become increasingly localized at the stadium curvature-discontinuity point P (see Fig. J.1). I found a single such family to be sufficient. Following Vergini's recipe, the EPWs are

$$\left.\begin{array}{r}\bar{\phi}_{2i-1}(\bar{\mathbf{r}})\\ \bar{\phi}_{2i}(\bar{\mathbf{r}})\end{array}\right\} \;=\; e^{-\sinh\,\alpha_i[(-n_{y,i}\bar{x}+n_{x,i}\bar{y})-b_i]}\cdot\left\{\begin{array}{c}\cos\\ \sin\end{array}\right\}\Big(\cosh\,\alpha_i(n_{x,i}\bar{x}+n_{y,i}\bar{y})\Big)\,, \qquad \text{(J.5)}$$

for $i = 1\cdots N_{\rm EPW}/2$. For each i, the propagation direction $\mathbf{n}_i$ (parallel to the wavevector $\mathbf{k}$) points along θ_i, and the growth and decay directions are perpendicular to this. The evanescence parameters are given by an algebraic progression,

$$\alpha_i \;=\; \frac{3+i}{2k^{1/3}}. \qquad \text{(J.6)}$$

The angles (restricted to the 2nd quadrant $[\pi/2,\pi]$) are chosen so as to 'balance' the EPW to roughly equal weight on the boundary either side of the point P:

$$\sin\theta_i \;=\; \frac{1}{\beta}\left(\frac{2}{k\sinh\alpha_i}\right)^{1/2}, \qquad \text{(J.7)}$$

where the balancing is controlled by the constant β. I found $\beta = 1.5$ optimal (Vergini reports $\beta = 2$ and uses a second family at $\beta = 10$). The impact parameters b_i (see Fig. J.1) are such that the maximum value of ϕ on the boundary is unity. The number $N_{\rm EPW}$ is chosen so that the maximum α_i is about 3 (the limit on α is discussed in Section 6.3.3). At $k = 10^3$, $N_{\rm EPW} \approx 50$ is sufficient. So even though the EPW functions take longer to evaluate than RPWs (especially when symmetrized, see below), there are so few that their computational effort remains small.

J.4 Symmetrization and reduction of effort

If the billiard has symmetry, then all its quantum eigenstates will have definite parities under the various symmetry operations [174]. If a basis set is chosen which does not share the problem symmetry, then all the eigenstates of the various classes will be found. However it is very beneficial to instead symmetrize the basis set, restricting the problem to that of finding only eigenstates in a single symmetry class. This can be repeated for all symmetry classes if desired.

Why is this beneficial? Consider a billiard with a single plane (or line) of reflection, $\mathcal{R}$. States fall into the odd or even classes. An unsymmetrized basis approach requires consideration of the full boundary area A and the full domain volume V, leading to a basis set size N proportional to A. Consideration of only the odd class (by using an odd-symmetrized basis set) requires matching of the BCs on only half of the area (no matching is required on $\mathcal{R}$), and therefore N is halved. If the effort per state found scales like N^3

(N^2) with the sweep (scaling) method, then symmetrization has increased efficiency by a factor 8 (4). Also the level density of states is halved, making the states easier to distinguish (see Section 5.5.3), raising the maximum k reachable.

Generally the symmetrization of the basis functions means that they take n times as long to evaluate (where n is the symmetry factor). This has an effect on the speed of filling matrices, and of evaluating the final eigenfunctions (and gradients) in position-space.

For the stadium example common in this thesis, a fourfold symmetry is present (C_{4v}): reflection about both the x and y axes. Therefore all basis functions will undergo the following symmetrization:

$$\phi_i(x,y) \;\longleftarrow\; \frac{1}{4}\left[\phi_i(x,y) + p_x\phi_i(-x,y) + p_y\phi_i(x,-y) + p_xp_y\phi_i(-x,-y)\right], \tag{J.8}$$

where the parities are $p_x, p_y = \pm 1$. The choice of the odd-odd symmetry class corresponds to $p_x = p_y = -1$. In the case of RPWs, such symmetrization results in only a single basis function per direction vector. For instance, the odd-odd class gives

$$\bar{\phi}_i(\bar{\mathbf{r}}) \;=\; \sin(n_{x,i}\bar{x})\sin(n_{y,i}\bar{y}), \tag{J.9}$$

where the directions θ_i need now only cover the range $[0, \pi/2]$. For an example see Fig. J.1. So it turns out that for RPWs in C_{4v} only twice as much time is spent per basis function evaluation than for the unsymmetrized case (J.3). For EPWs no such reduction in symmetrization time occurs.

Appendix K

Transmission cross section in the narrow slit limit

In this appendix we derive the transmission cross-section of the idealised slit from Section 7.3, in the tunneling limit $ka \ll 1$. We first consider a simple scaling argument which gives the dependence on k and a, and then find the correct prefactor.

The slit system possesses a scaling property (shared by all hard-walled quantum systems): it is invariant under a rescaling of space $V(\mathbf{r}) \mapsto V(\alpha\mathbf{r})$ provided there is a corresponding rescaling of wavevector $k \mapsto \alpha k$. Our slit potential is defined by a single parameter a, the slit half-width. The particle wavelength $\lambda \equiv 2\pi/k$ is the only other length scale in the problem. This means that for a given incident wave, the form of the full scattering wavefunction, and hence the ratio of the quantum to classical transmission cross-sections, can only depend on ka. (In other words, if the wavelength changes in proportion to the slit size, we are viewing the original problem with a new unit of length.) The classical cross-section is proportional to a. Therefore the quantum cross-section must have the exact form

$$\sigma_{\mathrm{T}}(k, \phi) = a \cdot f_\phi(ka), \tag{K.1}$$

where f_ϕ is a function dependent on the incident angle ϕ.

We seek f_ϕ in the $ka \ll 1$ limit. For a hard-walled system solutions of Schrodinger's equation become those of the Helmholtz equation with Dirichlet boundary conditions. Close to the slit (a distance $r \ll \lambda$) they become well approximated by the $k \to 0$ limit, namely solutions of Laplace's equation. Thus we can use 'electrostatics' at these small distances, and then match to the small-ka tails of the Helmholtz solutions to find the transmitted flux. The solutions obtained in this way have been known for over a century, and were first found by Lord Rayleigh ([168] in which the scattering solutions are derived for $ka \ll 1$ Dirichlet and Neumann apertures and their inverses, in both 2D and 3D. See p. 270-1 for our particular case). The transmission is dominated by scattering of incoming dipole (l=1) radiation to outgoing dipole radiation; we will justify this shortly.

The dipole parts of the 2D Laplace solution can be written generally as

$$\psi \;=\; \begin{cases} c_1 r\cos\theta + (c_{-1}/r)\cos\theta & \text{left side} \\ d_1 r\cos\theta' + (d_{-1}/r)\cos\theta' & \text{right side.} \end{cases} \tag{K.2}$$

The solution of Laplace's equation in the slit problem with this asymptotic ($r \gg a$) form is well known (see [168, 148]), and can be found for instance using a conformal map $w(z) = (z - ia)^{1/2}(z + ia)^{1/2}$, where $z \equiv x + iy$. (ψ is then a simple linear function of w.) There results a linear relation between the coefficients

$$c_{-1} \;=\; d_{-1} \;=\; A(c_1 + d_1), \tag{K.3}$$

where A is a constant dependent on the slit size a alone. The term in brackets is the difference in wavefunction x-gradients between the left and right sides. It is determined by the 'unscattered' wave ψ_0, to be precise by the l=1 component of ψ_0. This gradient difference 'drives' the dipole terms c_{-1} and d_{-1}. The transmitted flux is in turn proportional to $|d_{-1}|^2$.

The dimensional scaling of (K.3) implies $A \sim a^2$, so at constant k the transmitted flux, and hence σ_{T}, grows like a^4. Comparison with (K.1) immediately gives the scaling $\sigma_{\mathrm{T}} \sim k^3 a^4$. This λ^{-3} dependence of cross section for a given aperture size is the 2D equivalent of Rayleigh's famous λ^{-4} law for dipole light scattering in 3D explaining why the sky appears blue [167] (see also [105] p.418).

How does the contribution of higher multipole radiation scale? The exact Laplace solution with general asymptotic form will involve generalizing Eq.(K.3) to a linear matrix relation connecting all the $\{c_m\}$ and $\{d_m\}$ with $m = 1 \cdots \infty$ to all those with $m = -1 \cdots -\infty$. However because of the associated r^m and r^{-m} factors, by dimensional arguments all the elements of this matrix (other than the A already considered above) will scale like powers of a higher than 2. It is important to realise that the Laplace coefficients $\{c_m\}$ ($\{d_m\}$) directly connect to the respective coefficients of Bessel functions of J_m type for $m > 0$ and Y_{-m} type for $m < 0$ on the left (right) sides, which in turn determine the incoming and outgoing waves of angular momentum $l = |m|$ on those same sides. Therefore the resulting contribution to σ_{T} due to excitation by the l channel on the left and re-radiation into the l' channel on the right will scale as $\sim a^{l+l'}$. This contribution will therefore be smaller than the dipole-dipole (l=1, l'=1) component by a factor $\sim (ka)^{l+l'-2}$. In fact, because of symmetry about the x-axis, the system only couples even l to even l' and odd l to odd l'; this guarantees that the most significant correction is in fact a factor $(ka)^2$ down. Thus in our $ka \ll 1$ limit we are justified in using only p-wave scattering: dipole 'absorption' from ψ_0, and dipole re-radiation into ψ_{T}. (There will also be equal dipole re-radiation into ψ_{R}, which interferes with ψ_0 but does not affect the conductance).

We now can find the prefactor. The connection between the Laplace dipole solution ($m = -1$) and the outgoing dipole coefficient q_1^+ (see Section 7.2) is found by matching to the small-kr form [7] of the Neumann part $Y_1(kr)$ of the Hankel function. This gives $q_1^+ = (i\pi k/2)d_{-1}$. Combining this with (K.3) and the correct Laplace slit result $A = a^2/4$ [168, 148] gives

$$q_1^+ \;=\; \frac{i\pi k a^2}{8}\, c_1. \tag{K.4}$$

Here we assumed that $d_1 = 0$ which is strictly true only when the slit is closed, since the unscattered wave is zero on the right side. However at $ka \ll 1$ it is a good approximation because the transmitted wave is much weaker than the incident. With a unit incident plane wave at angle ϕ, the unscattered wave is Eq.(7.1) with $\gamma_{\mathbf{k}} = 0$, giving $c_1 = 2ik\cos\phi$. Now taking the large-kr form of the Hankel function and comparing to Eq.(7.3) gives the

scattering amplitude

$$\begin{aligned} f_{\mathrm{T}}(\theta') &= \sqrt{\frac{2}{\pi k}}\, q_1^+ \cos\theta' \\ &= -\sqrt{\frac{\pi k^3}{8}}\, a^2 \cos\phi\, \cos\theta', \end{aligned} \tag{K.5}$$

where in the second equality we used (K.4) with c_1 substituted. Applying Eq.(7.6) gives the desired transmission cross section σ_{T} as stated in Eq.(7.14).

Note that σ_{R} is equal to σ_{T}, because the reflected dipole strength is equal to the transmitted. (The dipole is also 'two-headed', radiating *in phase* on both sides).

It is interesting to note that in *any* hard-walled scattering system, the conductance can only be a function of the product ka. This follows from substitution of Eq.(K.1) into Eq.(7.11a).

Appendix L

Explicit relation of cross section to Landauer formula

In this appendix we derive the relation (7.22) formally in a self-contained fashion. No reference to the semiclassical derivation of Section 7.1 is made. Therefore this provides an independent derivation of our main result (7.1).

We first present the special case $\gamma_{\mathbf{k}} = 0$ corresponding to a QPC system embedded in a hard wall (at $x = 0$). The reflection off the unbroken wall is diagonal in the half-plane partial-wave basis of Section 7.2. We can expand the 'unscattered wave' in Bessel functions [continuing the expansion of Eq.(7.19)], to give

$$\psi_0(\mathbf{r}) \;=\; 4\sum_{l=1}^{\infty}(-i)^l J_l(kr)\sin[l(\tfrac{\pi}{2}-\theta)]\sin[l(\tfrac{\pi}{2}-\phi)]. \tag{L.1}$$

Recognising this as being composed entirely of a sum of the half-plane partial-wave bases (7.12) on the left side, we can read off the coefficients $p_l^- = 2(-i)^l \sin[l(\frac{\pi}{2}-\phi)]$. (Note that the outgoing coefficients p_l^+ are also equal to this). With the QPC system now placed in the wall, the outgoing amplitudes on the right side are $q_n^+ = \sum_l t_{nl} p_l^-$. This gives a flux

$$\begin{aligned}\Gamma_{\mathrm{T}}(\phi) \;&=\; \frac{\hbar}{m}\sum_n |q_n^+|^2 \;=\; \frac{\hbar}{m}\sum_{ll'}(p_l^-)^*(t^\dagger t)_{ll'}p_{l'}^- \\ &=\; 4\frac{\hbar}{m}\sum_{ll'}(i)^{l-l'}\sin[l(\tfrac{\pi}{2}-\phi)]\;(t^\dagger t)_{ll'} \\ &\qquad \sin[l'(\tfrac{\pi}{2}-\phi)].\end{aligned} \tag{L.2}$$

Using the definition of cross section (7.4) we can write

$$\begin{aligned}\int_{-\pi/2}^{\pi/2} d\phi\,\sigma_{\mathrm{T}}(k,\phi) \;&=\; \frac{1}{v}\int_{-\pi/2}^{\pi/2} d\phi\,\Gamma_{\mathrm{T}}(\phi) \\ &=\; \frac{4}{k}\sum_{ll'}(t^\dagger t)_{ll'}\cdot\int_{-\pi/2}^{\pi/2} d\phi \sin[l(\tfrac{\pi}{2}-\phi)] \\ &\qquad \sin[l'(\tfrac{\pi}{2}-\phi)],\end{aligned} \tag{L.3}$$

from where the orthogonality of the radial functions allows us to replace the final integral by $(\pi/2)\delta_{ll'}$, whereupon the sum becomes a trace, giving

$$\int_{-\pi/2}^{\pi/2} d\phi\, \sigma_{\rm T}(k,\phi) \;=\; \lambda {\rm Tr}(t^\dagger t), \tag{L.4}$$

completing the proof.

Now considering a general wall profile with a general $\gamma_{\mathbf{k}}$, the situation becomes more subtle. An attempt to expand ψ_0 in Bessel functions as in (L.1) gives additional terms which are not part of the half-plane basis (7.12), namely an l=0 (s-wave) term and terms with angular dependence $\cos[l(\frac{\pi}{2}-\theta)]$ for all $l \geq 1$. This would suggest that the basis (7.12) is incomplete. However, all these new terms can be expressed as sums of the $\sin[l(\frac{\pi}{2}-\theta)]$ functions already present, if we are careful to consider the large radius limit $kr \to \infty$ before the maximum angular momentum limit $l_{\rm max} \to \infty$. In this limit $l_{\rm max} \ll kr$, a given basis state will have negligible x-wavevector near the wall, compared to the wall profile width, so the 'softness' of the wall profile will not be apparent. Rather, the wall will appear as a hard reflector at $x \approx 0$. We write the entire left side wavefunction

$$\psi(\mathbf{r}) \;=\; f^-(\theta)\frac{e^{-ikr}}{\sqrt{r}} + f^+(\theta)\frac{e^{ikr}}{\sqrt{r}}, \tag{L.5}$$

in a general form where the incoming and outgoing angular distributions are apparent. Therefore, even for general $\gamma_{\mathbf{k}}$, the apparent hardness of the wall forces $f^-(\theta)$ and $f^+(\theta)$ to go smoothly to zero as $|\theta| \to \pi/2$. This ensures that the original basis set is sufficient to represent all asymptotic wavefunctions, since $\sin[l(\frac{\pi}{2}-\theta)]$, $l = 1 \cdots \infty$ form a complete set in the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$. The set $\cos[l(\frac{\pi}{2}-\theta)]$, $l = 0 \cdots \infty$ form *another* complete set in the same interval; both sets are not required and we are able to choose just one [7]. Because of the hard wall boundary condition in θ-space, the former set is more appropriate.

We are still left with the issue of finding the incoming p_l^- coefficients given an unscattered wave ψ_0 for general $\gamma_{\mathbf{k}}$. The problem is subtle, but can be understood when we consider the order of limits above: restricting $l \leq l_{\rm max}$ allows the incoming representation $f^-(\theta)$ of a single plane wave $e^{i\mathbf{k}\cdot\mathbf{r}}$ to become a well-defined, narrowly-peaked delta-like function about the incident angle. This single plane wave will also cause a similar delta peak in the opposite direction, which it turns out is irrelevant because it contributes instead only to *outgoing* $f^+(\theta)$. Thus we have the important result that we can *ignore* the reflected wave (making its phase shift $\gamma_{\mathbf{k}}$ irrelevant) in calculating p_l^-, because this wave can only contribute to $f^+(\theta)$ in the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$. We can use stationary-phase integration (method of steepest descents) [7] applied to a single plane wave to show this, and find the delta strength, as follows. Taking care with the definitions of angle (see Fig. 7.1a) we have

$$e^{i\mathbf{k}\cdot\mathbf{r}} \;=\; e^{-ikr\cos(\theta-\phi)} \;=\; f^-(\theta)\frac{e^{-ikr}}{\sqrt{r}}, \tag{L.6}$$

where $f^-(\theta) = e^{ikr[1-\cos(\theta-\phi)]}$. Expanding the cos as a quadratic about the stationary point $\theta \approx \phi$, and making the usual stationary-phase replacement

$$e^{ig(x)} \;\to\; \sqrt{\frac{2\pi}{-ig''(x_0)}}\, e^{ig(x_0)}\, \delta(x-x_0) \tag{L.7}$$

gives $f^-(\theta) = e^{i\pi/4}(2\pi/k)^{1/2}\,\delta(\theta-\phi)$. The other stationary point $\theta \approx \phi + \pi$ results in the same delta strength but an additional factor e^{2ikr}, which turns the incoming e^{-ikr} to an outgoing e^{ikr}, thus contributing only to $f^+(\theta)$. Generally, we can write the expression for incoming partial-wave expansion

$$p_l^- \;=\; \sqrt{\frac{2k}{\pi}}\,e^{-i\pi/4}\,(i)^l \int_{-\pi/2}^{\pi/2} d\phi\, f^-(\theta)\,\sin[l(\tfrac{\pi}{2}-\theta)], \tag{L.8}$$

which is easily derived using orthogonality of the angular functions and the asymptotic form of $H_l^{(2)}(kr)$. Substituting the narrowly-peaked $f^-(\theta)$ resulting from the single incident plane wave gives

$$p_l^- \;=\; 2(-i)^l \sin[l(\tfrac{\pi}{2}-\phi)]. \tag{L.9}$$

These are the same incoming partial wave coefficients as derived above for the case $\gamma_{\mathbf{k}} = 0$, which is as expected since the reflected wave phase has not entered into our considerations. Thus our proof (L.4) holds for general $\gamma_{\mathbf{k}}$.