Professional statement — Alex H. Barnett, November 2010

As an applied mathematician, I create numerical and analytical tools that enable the solution of challenging problems in diverse areas of science and technology. This requires cross-disciplinary collaboration with scientists as well as other mathematicians, and the pursuit of excellence in developing innovative computational methods. In turn, questions arising in these areas have led me to interesting mathematics. As an educator, I aim to share both the joy that comes when a theorem has been understood deeply, and the satisfaction of using the power of computational mathematics to help reliably solve problems we face in the wider world.

1 Scholarship

My work focuses on solving partial differential equations (PDEs), which are the mathematical objects that describe much of the physical world around us: the flow of heat, the propagation of sound, light and radio signals, and the quantum wavefunctions that lie at the root of all matter. For example, we may represent an electric field throughout a region of space by a function; the PDE is then an equation relating various *rates of change* of this function. An example at the core of my work is the *Helmholtz equation*, which states essentially that if a function is positive its curvature must be negative, and vice versa, leading to oscillations and waves. Solving the PDE then means finding a function which obeys the equation, yet also behaves in prescribed ways at the *boundaries* of geometric objects such as reflectors or other devices with which the waves interact.

It is possible to write down a formula for the exact solution only for boundaries with few simple geometric shapes. Thus finding a *numerical approximation* to the solution when the boundary shape is more complex has been crucial to progress in science and technology. With the invention of electronic computers, the size and complexity of problems we routinely solve has grown exponentially. An interesting consequence is that the choice of a good solution method matters now more than ever before, since it may turn a problem requiring several days of computer time into one requiring mere minutes. I create and apply new algorithms for the efficient and accurate solution of challenging PDE problems—ones that would be prohibitively expensive to solve with current conventional tools. In many cases I also analyze these methods mathematically, so that their understanding is given a more rigorous foundation, and their computational speed and accuracy is guaranteed.

My research achievements and interests are in the following areas:

- Numerical analysis and scientific computing, particularly computational waves at high frequency and/or high accuracy: the Helmholtz equation; eigenvalue, scattering and periodic problems; high-order and spectral methods, corner singularities, boundary integral equations; the heat equation.
- Mathematical physics: numerical studies in quantum chaos, semiclassical (high frequency asymptotic) properties of the spectrum and eigenfunctions of the Laplacian; billiards, and scarring.

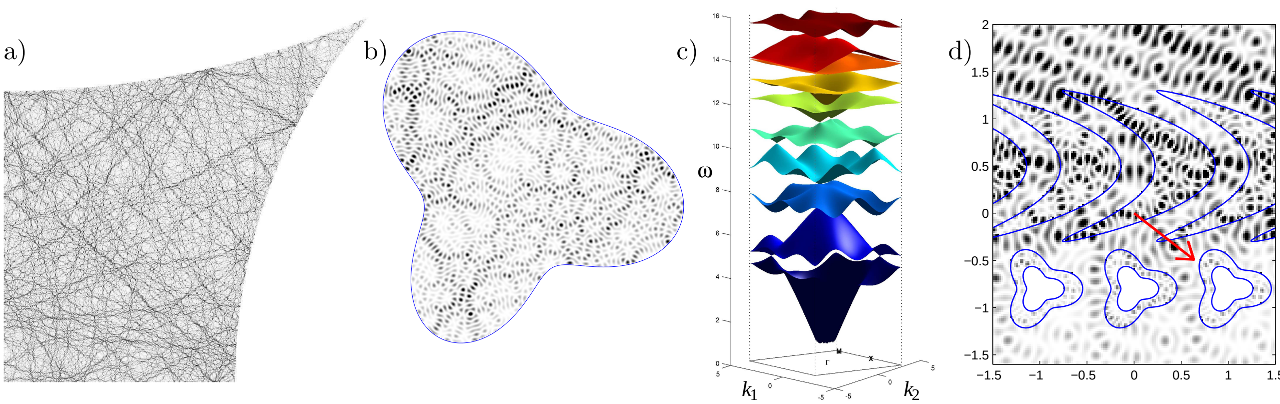

Figure 1: a) 10^5th Dirichlet eigenfunction of a chaotic 2D domain (size 320 wavelengths; see Section 1.2.) b) Eigenfunction of 2D domain, 10 digit accuracy (Section 1.1.2.) c) Band structure of 2D photonic crystal (Section 1.1.1.) d) Scattering from infinite periodic array of dielectric obstacles, 11 digit accuracy (Section 1.1.1.)

- Applied mathematics: mathematical ecology (animal home range modeling), inverse problems in medical imaging (diffuse optical tomography), the mathematics of music.

I have given 46 research talks, all but four of which were by invitation. I have 18 refereed journal publications that have appeared or are in press, and a further preprint in review, all in top-level international journals. I have also authored and co-authored freely-available and documented software, as part of my goal to enable other researchers to take advantage of the tools I have developed. My work has been supported over the last five years by two single-PI 3-year research awards from the Division of Mathematical Sciences of the National Science Foundation, on high-frequency wave eigenvalue and scattering problems, totalling $413,037. My research trajectory reflects my origins in a Ph.D in physics, followed by a postdoc in medical imaging, before finding what I find to be a truly satisfying home in applied mathematics. A significant part of my work remains interdisciplinary, collaborating with biologists, engineers, pure mathematicians, and recently musicologists. Grappling with challenges and modes of thought outside of mathematics is a key part of what excites me.

1.1 Numerical analysis for wave computations

Advances in numerical methods have been of paramount importance for scientific progress and the optimization of physical devices in today's age of computers and communications. Society has been transformed in the last century by technologies that use waves: one need only mention television, cell-phones, radar, ultrasound, and fiber optics. Finite elements and other standard numerical tools have enabled a large variety of devices to be modeled, at least to low accuracy, but there remain many situations which demand high frequency (where the many wavelengths across the system renders it a challenging multiscale problem) and/or high accuracy. These include acoustic and radar scattering problems, resonant devices such as filters, and devices containing periodic arrays or lattices. There is a growing range of modern applications to nano-optics, lasers, optical signal processing, meta-materials, communications, remote sensing, and cost-efficient solar photovoltaics. Efficiency and robustness are paramount because modeling is a bottleneck: predicting real-world device performance often requires thousands of solutions at different parameters.

A key theme of my work has been to develop more efficient representations for waves that allow solutions to the Helmholtz equation to be approximated with high order (often spectral) accuracy. Often these use *particular solutions* (for example, plane waves) as global basis functions, turning the PDE problem into the simpler problem of matching on boundaries via least-squares or other collocation methods. With an efficient basis (often needing only three unknowns per wavelength on boundaries), large-scale problems several hundred wavelengths in size become tractable with, and indeed best handled by, dense linear algebra. For instance, with collaborator Timo Betcke (Reading, UK) I have designed such Fourier-Bessel basis sets to handle the singularities occurring at each corner of a polygonal domain, solving the two-dimensional (2D) scattering problem with proven exponential convergence [15] [1]. By matching on a few artificial boundaries we achieve typical errors of 10^{-10} in only a few seconds of computation, for Dirichlet and Neumann problems that are challenging for both standard discretization and integral equation methods.

Another efficient global basis is a set of fundamental solutions ('point sources') lying on a curve outside the physical domain. This idea is popular with engineers, yet had not been analyzed for interior Helmholtz boundary value problems before our work [13]. We provided a complete convergence analysis in the disc, and demonstrated spectral accuracy in a variety of analytic planar domains. Crucially, it is not the condition number (which is very high), but rather the norm of the coefficient vector that limits the accuracy in finite-precision arithmetic. Keeping this norm small in turn requires understanding the singularities in the analytic continuation of the solution field via Vekua's analytic theory of elliptic PDEs, a realization important for future progress.

I have released a user-friendly, object-oriented and documented MATLAB toolbox (`MPSpack`, co-authored with Betcke) which implements the above methods, as well as the periodic methods described below, allowing the efficient solution of 2D wave scattering and eigenvalue problems on piecewise-homogeneous domains. Codes are human-readable and have a simple interface. For example, `s = segment(20, [0 1])` creates a line segment from 0 to 1 with 20 Clenshaw-Curtis quadrature points; from this, domains, basis sets, and problems may be set up and solved. The 70+ pages of tutorial and user manual include many code examples and figures. Publications [15] and [18] include `MPSpack` codes of a few lines which solve complete example problems with the

[1] Numbered citations are to the works enumerated in the publications section of my CV

new algorithms, allowing other researchers to validate and make use of my work.

One future goal is to exploit analytic and contour integral formulae for the fundamental solution in graded index media in 2D and 3D, continuing work arising with former Ph.D. student Matt Mahoney. This will enable integral equation methods to be applied to continuously-varying material problems for the first time.

1.1.1 Periodic problems: diffraction gratings and photonic crystal band structure

Many devices designed to guide and control waves rely on periodic structures on the wavelength scale: these include antennae, diffraction gratings (used to squeeze multiple signals onto a single optical fiber, and in our highest-powered lasers), photonic crystals (the most promising route to energy-efficient ultra-fast optical computation on a chip), meta-materials (allowing the control of waves in ways impossible in naturally-occurring media), sound absorbers, and solar cells. In recent work with Leslie Greengard (NYU) I have devised new efficient approaches to solving such periodic frequency-domain problems using integral equations, which, unlike existing methods, are robust for all problem parameters.

Integral equations (as with the global bases discussed above) reduce the dimensionality of the problem: a 2D wave problem becomes a 1D integral equation involving the Green's function kernel on the material boundaries. To solve the problem of an infinite array of objects, one must apply quasi-periodic boundary conditions to the solver for a single object. The conventional route is via the *quasi-periodic Green's function*. This turns out to have major drawbacks: i) it breaks down near certain parameters (the so-called Wood's anomalies) even though the boundary value problem remains well-posed, and ii) in large-scale problems its computation via lattice sums is cumbersome, and inefficient in geometries with high aspect ratio. We bypass both problems by using the free space Green's function alone, adding auxiliary layer potentials on the unit cell boundary, while imposing quasi-periodicity as a new linear condition.

With this new approach, in [16] we demonstrated a new spectrally-accurate solver for 2D photonic crystal band structure (see Fig. 1c), that is, the eigenvalue problem for the frequencies of traveling waves with given quasi-periodicity. We periodized using only around 100 extra degrees of freedom, achieving errors approaching machine precision using only small systems that are set up and solved in less than 1 second of computer time. The scheme is shown to be robust at parameters where the conventional quasi-periodic Green's function blows up (the resonances of the empty unit cell). Along the way we provided the first proof (via Calderón projectors) that the conventional scheme is equivalent to the band-structure problem.

We proceeded to the problem of scattering from a grating (infinite 1D array) of smooth 2D dielectric obstacles in [18]: a complication arises in evaluating layer potentials on unit cell walls since they now are infinite in extent. We tackle this using spectral quadrature in the Fourier domain, explicitly handling poles in the complex plane near Wood's anomalies. Thus, as above, our method is robust at parameters where the standard approach fails. Because of careful quadrature in the Fourier variable, we achieve 10^{-14} errors with a few hundred unknowns (see Fig. 1d), whilst maintaining the 2nd kind nature of the system (essential for a low condition number in iterative solvers).

We designed the above schemes to scale well to large-scale problems (10^4 or more unknowns), and allow fast multipole acceleration to be included without fuss. My future plans are to adapt the method to multi-layer dielectrics with embedded inclusions, to solve the scalar (and eventually, full Maxwell) band structure for 3D crystals, and to the technologically-important case of doubly-periodic layered scattering problems in 3D.

1.1.2 Laplacian eigenvalue problems

Eigenvalue problems for the Laplace operator ('modes of the drum') have a wealth of applications in engineering and physics, such as acoustic and electromagnetic resonators and waveguides, vibrations of membranes, quantum mechanics, and, more recently, the design of higher-power micro-cavity lasers for communications applications. This classical problem also has growing mathematical interest in the fields of spectral geometry and quantum chaos (see Section 1.2). Many of these areas demand efficient numerical solution at high frequencies, that is eigenvalue numbers from 10^2 up to 10^6, where conventional discretization methods fail. I have developed, and analyzed rigorously, new particular solution based methods—for instance the scaling method—which solve this problem in regimes where all other known methods are impractical.

In [14], motivated by my numerical observations, I proved that for large Dirichlet eigenvalues, that is, high frequency applications, the method of particular solutions (MPS) in fact has error bounds a factor of the square-root of the eigenvalue *better* than was previously known (in exact arithmetic). In some applications this guarantees 3 extra digits of accuracy with no extra effort. The tools needed were Kato-Rellich perturbation theory and Cauchy interlacing, applied to a certain boundary operator that depends analytically on an 'energy' (or trial eigenvalue) parameter E. This result applies for E in neighborhoods of unknown size around each eigenvalue.

Then, in collaboration with Andrew Hassell (ANU), in [19] we improved this result to apply for all parameters E, and also for the L^2 error bounds of the corresponding eigenfunctions. We gave a numerical example at around the 2500th eigenvalue where the relative error in eigenvalue is of order 10^{-14}, and eigenfunction error 10^{-10}, for a smooth planar domain (Fig. 1b). A central tool in our analysis was the proof of a new form of boundary *quasi-orthogonality* of a domain's Dirichlet eigenfunctions (more specifically, their normal derivative functions), a PDE result with potential impact beyond numerical analysis. Our work in progress extends the above sharp bounds to large Neumann eigenvalues (with applications to acoustic resonators); it turns out that, for this boundary condition, microlocal analysis becomes essential if similar sharp bounds are to be achieved.

In recent work that we are currently writing up, we give the first rigorous analysis of the *scaling method* of Vergini–Saraceno. The latter is little-known accelerated variant of the MPS which is $O(\sqrt{E})$ times more efficient; in high-frequency applications the speed-up can be a factor of 10^3. Its mechanism deserves a brief explanation. To find Dirichlet eigenvalues with the MPS one must search along the E parameter axis to locate minima of a function (a Rayleigh quotient); the search is laborious and unreliable since the function 'zigzags' in an unpredictable fashion. Using only the effort required above for a single function evaluation, the scaling method instead approximates directly the set of *all* eigenpairs in a spectral window of size $\epsilon = O(\sqrt{E})$. We study this using the linearization of the E-dependence of the eigenvalues of a weighted Neumann-to-Dirichlet map for the interior Helmholtz boundary value problem. We prove that the errors in the computed eigenvalues are $O(\epsilon^2/E + \epsilon^3/E)$, its first known error bound. This is important for the reliable use of the method.

It is worth noting that such a collaboration between pure and applied mathematicians is quite rare, and in this case is opening avenues that neither researcher could tackle alone. My study of the scaling method originated in my physics thesis work [2], however at a less rigorous level. I believe that the method has the potential to be useful in a broad range of engineering and other eigenmode applications, and I plan to continue to advance that goal. For instance, I aim to publish (this time in collaboration with Betcke) side-by-side comparisons of this against other numerical eigenmode tools (including integral equation and finite elements), for a range of test cases relevant to applications in 2D and 3D.

1.2 Numerical study of quantum chaos

Quantum chaos is the study of how the allowed vibrational patterns and vibrational frequencies of a region of space (imagine for instance a drum-head) depend on the shape of its boundary. Mathematically, these patterns are given by eigenfunctions of the Laplace operator, and the frequencies by the corresponding eigenvalues. It turns out that the statistical behavior of the patterns as one tends to very high frequencies (hence rapid spatial oscillations) depends on the type of motion that a "billiard ball" (or ray of light reflecting off the walls) would undergo in that same domain. If the billiard ball motion is regular (such as in a rectangular billiard table), one gets a distribution of eigenvalues and eigenfunctions that is very different from the case of *chaotic* billiard motion (for instance a table with curved walls). This idea has deep connections to quantum physics, geometry, and number theory, and involves many open questions at the forefront of contemporary mathematics.

More formally, the asymptotics as eigenvalue tends to infinity depends on the integrability of the geodesic flow on the manifold. For ergodic flow, we have a partial understanding in the form of rigorous results such as the Quantum Ergodicity Theorem, which states that *almost all* eigenfunctions become asymptotically equidistributed (spatially, and also in phase space). The Quantum Unique Ergodicity (QUE) conjecture of Rudnick–Sarnak is that, for uniformly-hyperbolic flow, this holds for *every* eigenfunction. It is believed that we are far from a proof of this conjecture (apart from arithmetic cases, for which it has only recently been proven).

This motivated my large-scale numerical study [9] in a Sinai-type billiard (flat manifold with boundary) with uniformly-hyperbolic flow; see Fig. 1a. I analyzed 30000 eigenfunctions up to eigenvalue number 7×10^5, and found strong evidence for QUE, as well as for conjectures arising in the physics literature on the power-law decay of the fluctuations in the spatial equidistribution of eigenfunctions. In a subsequent study with collaborator Timo Betcke [11], we computed 16000 eigenfunctions of the 'mushroom' billiard, a manifold whose flow is *mixed*: it has two invariant regions in phase space, one integrable and one ergodic. We verified (to an accuracy of 1%) Percival's conjecture on the localization of eigenfunctions, and proposed and tested a model for so-called dynamical tunneling from the integrable into the ergodic region.

In both of the above works, developing specialized numerical tools enabled me to reach eigenvalues 10–100 times higher than previously attained, in very reasonable computation times. In the case of the Sinai-type domain, a fundamentiel solutions global basis was developed. For the mushroom, a Fourier-Bessel basis set adapted to the corner singularity was required. Only by adapting the scaling method (see Section 1.1.2) to these new basis sets could high accuracy and high efficiency be achieved. The mathematical significance of the work [9] is highlighted by its appearance as images on the cover of *Notices of the American Mathematical Society* (January 2008). Related eigenmode images of mine have appeared in four recent review articles, including one by Peter Sarnak (arguably the world leader in this field).

I am excited to continue the study of chaotic eigenfunctions, since it is an area of 'experimental mathematics' where advanced numerical methods can make an impact. In particular, I am keen to study the so-called *bouncing-ball modes* that occur between two parallel flat walls (in collaboration with Andrew Hassell), and to measure the statistics of the numbers of *nodal domains* (regions of positive or negative sign) in eigenfunctions to test the percolation conjectures of Bogomolny–Schmit. The latter project will involve undergraduates, including the proposed senior thesis of Kyle Konrad '12.

1.3 Other applied mathematics research avenues

1.3.1 Ecology and home range modeling

The way animals use their habitat is controlled by factors such as food, water, terrain, predators, their 'home ranges', and social interactions; revealing these effects is key for conservation efforts. In a collaboration with Paul Moorcroft (Harvard), we have developed a new animal movement Markov model [10], [12]. This unifies two standard approaches: resource selection analysis (a spatially implicit approach), and mechanistic home range analysis (an explicit probabilistic kinetic model). Our model's kernel may be factorized to derive an analytic steady state solution, yet has the flexibility to model variation in habitat via a general spatial *preference function* w. As a result, such models could be fit to data much more rapidly than before, and the time-evolution (probability density functions) of moving animals can be simulated quickly using fast Fourier transform methods. The analytic solution allowed us to predict that steady state animal density scales like w on small scales but like w^2 on large scales, with consequences for current practices of real-world telemetry data analysis and model-fitting.

1.3.2 Diffuse optical tomography

Near-infrared light is a promising tool for non-invasive medical imaging, yet requires the solution of a complex inverse problem to reconstruct the images. Following on earlier work on Bayesian statistical solution of the inverse problem [6], and the development of a rapid solver for the PDE forward problem (heat equation) in multi-layer geometries [7], I have continued to contribute in this area with work validating the diffusion approximation against the more accurate, but much more computationally intensive, radiative transport model [8].

1.3.3 Mathematics of music

In a new collaboration [17] with Larry Polanksy (Dartmouth) and Michael Winter (UCSB), we have analyzed a framework for the random generation of musical melodies with varying amounts of temporal correlation. This simplifies and generalizes methods of James Tenney, one of the first composers to apply computer code to the creation of music in the 1960s. Our cross-disciplinary approach combines stylistic analysis, numerical experimentation, and rigorous mathematical analysis.

2 Teaching

In the classroom I make learning an interactive discovery process: most of my class periods involve a 10-15 minute session where students work in groups of two or three to solve problems on worksheets that I have devised. I am then able to move around the classroom, asking probing questions or giving hints, and learning how each individual thinks. This builds their communication skills (the strong students often teach the weaker ones better than I could), while giving me instant feedback on their areas of difficulty, and helping me steer the remainder of lecture. The classroom buzzes with discussion, and the energy level is high. Students take an active rather than passive role: by the end of lecture they have solved one or more core problems themselves, rather than merely copied down solutions from the board. I scan and post all worksheets and solutions online for later study (see any of my course webpages). Student feedback on the use of worksheets is consistently enthusiastic. In developing this idea I was influenced by my time as a teaching assistant with Eric Mazur (Harvard), a pioneer in such peer instruction techniques. Peer instruction has been proven to be successful with students of different levels and personalities, and to increase female enrollment and retention in math and the sciences.

Another feature of many of my lectures is a 'live demo' involving some real world object that illustrates or brings a concept to life. Examples include: waving around huge colored 'vectors' to explain cross product, using masses on springs or torsional wave machines for differential equations, casting shadows of solids with a flashlight (or penetrating them with knitting needles) for triple integrals, sounding coupled tuning forks to illustrate resonance, hitting drums (or whirling flexible plastic tubing above my head!) while measuring their frequency spectrum to illustrate eigenvalue problems, and taking students outside to measure the speed of sound themselves using clapping and echoes. I also use computer demos, for instance showing students applets that let them see vector fields using fluid flow. I find that visualization is key to student understanding and engagement.

Since computation plays an important and growing role in applied mathematics and the sciences, I have woven a computer component (mostly MATLAB/Octave, but also online applets and audio analysis software) into almost every course I teach. Thus students coming out of my courses have (to varying degrees) experience with computational tools that modern scientists use—a research and career advantage. At the major and graduate level I emphasize programming skills more strongly, so that students learn how to debug their own mistakes, and build more complex routines reliably.

In my five years at Dartmouth so far I have developed from scratch four new courses, each of which attracts healthy enrollments (I have been fortunate to be able to teach each two or three times). These courses greatly expand the applied offerings at Dartmouth, and the first two now form a core part of our applied math major:

- Math 46: *Introduction to Applied Mathematics*, a math- and science-major course covering core analytical tools such as dimensional analysis, perturbation series, integral equations, Green's functions, Fourier transforms, and analysis of PDEs.

- Math 53: *Chaos!*, a math- and science-major course introducing phenomena in dynamical systems (maps and flows), their rigorous analysis, numerical simulation, and applications.

- Math 5: *The Mathematics of Music and Sound*, a non-science major course that fulfills the QDS (quantitative distribution) requirement, covering musical acoustics, functions and signals, Fourier analysis, waves, musical instruments, human hearing, and digital sound technology (all without calculus!)

- Math 116: *Numerical Methods for PDEs and Waves*, a graduate-level course on numerical analysis, scientific computing, and integral equations, incorporating material from my research.

In Math 5, 53, and 116, I make independent project work a key part of the assessment (in Math 53 this replaces the final exam); students respond very enthusiastically, since now they feel ownership of the work they are doing, and experience the excitement of research. In fact, I have seen Dartmouth students 'pull out all the stops' when finishing projects like at no other time, and experience great satisfaction when completed. Undergraduate research projects have included computing the Lyapunov exponents in chemical chaos, testing new numerical methods for solving PDEs in the plane, and measuring the temporal sensitivity of the human ear. In Math 5 and Math 50 (*Mathematical Statistics*) I also set up an online wiki/comments page where students can post examples of each week's material from everyday life and culture, connecting the content to applications, and providing discussion material relevant to their lives. In all of the above classes I grade students on the quality of their communication skills, and work with them to create engaging and clear class presentations and write-ups. I often have a final presentation night with pizza—a mini math conference—where students learn about each other's projects, and learn how to communicate in a microcosm of a professional science community.

My future goals include developing the numerical analysis and computational sides of the undergraduate and graduate curricula, and continuing to create and refine interactive teaching tools in all of my classes.

3 Service to Department and College

During my time at Dartmouth I have mentored nine undergraduate research projects on topics such as numerical methods, waves, imaging and chaos. Four of these were senior theses (of which three received *high honors*): examples include Yong Su '09, on spectral methods for computing the capacity of the unit square and Laplacian eigenmodes on the sphere (Yong is now in graduate school at Stanford), and Zoe Lawrence '10 on PDE models for the spread of insect-borne disease across farmland (co-advised with Dorothy Wallace). Two more students have been involved in an interdisciplinary project to develop image-processing algorithms to track the genealogy of cell nuclei in 3D microscopy videos, co-advised with Amy Gladfelter (Biology). Preliminary work suggests that we can significantly boost the amount of data the Gladfelter lab can process. Through the *Women in Science Project* I have mentored Vissuta Jiwariyavej '09 on acoustics and number theory, will be mentoring a *WISP* first-year Hahn Nguyen '14 starting Winter 2011, on numerical analysis. Other outreach to undergraduates includes giving four presentations to the Dartmouth Mathematics Society. For the last three years I have served on the Faculty Advisory Board to the *Dartmouth Undergraduate Journal of Science*, and in this time worked with two students to edit their mathematics articles for publication in *DUJS*.

I have served on various department committees including Recruiting, both the Undergraduate and Graduate program committees, and Graduate Admissions (see CV for a full list). I was also heavily involved in revamping the Department's applied math major curriculum during my first couple of years here, as part of an ad-hoc committee with Scott Pauls, Dan Rockmore, and Peter Doyle. In terms of mentoring other faculty, I have given two presentations at the *Dartmouth Center for the Advancement of Learning*, one solo workshop on the use of worksheets to create interaction in science classes, and one group panel on teaching for those new to Dartmouth.

At the graduate level, I advised Matt Mahoney (Ph.D. '09) from 2006–2007 on the method of particular solutions for graded index materials. In conjunction with Scott Pauls and Dan Rockmore, I have developed our Department's new applied math graduate qualifying exam. I then served on the applied qualifying exam committee of Katherine Kinnaird. I have also given three introductory talks to attract prospective graduate students to our Department.

In my first year at Dartmouth I started the Applied and Computational Mathematics Seminar, which has now hosted 34 talks in total (many being invited visitors). This research seminar brings together a community of researchers across campus interested in numerical modeling; its audience includes engineers, physicists and biologists.

Finally, I was one of four organizers of the 8-day-long *International Conference on Spectral Geometry*, which brought around 80 international visitors, including most of the major people in the field, to Dartmouth for a very successful research conference this July. The other organizers were Carolyn Gordon, Peter Perry (Kentucky) and Alejandro Uribe (Michigan). We feel that this has enhanced Dartmouth's already strong reputation in quantum chaos and spectral geometry.