

Nanostructures and Computation

RLE Group

Nanostructures and Computation Group

Academic and Research Staff

Professor Steven G. Johnson

Graduate Students

Ardavan Farjadpour

Undergraduate Students

Alejandro Rodriguez

Having joined both the faculty of Applied Mathematics and RLE in the fall of 2004, our group has focused on two main directions of research. The first is that of understanding problems of wave propagation in nanostructured media, primarily *photonic crystals*—periodic dielectric media that can mold the flow of light to yield unusual optical phenomena from confinement in air cavities to negative refraction. Because these are complicated structures, there are many open problems in their theoretical understanding, from what phenomena are possible to how to predict their behavior without resorting to brute-force computation. In this modeling, we have made strides in developing semi-analytical approaches that have shed new light on the problems of loss from disorder and other mechanisms. However, large-scale computation is of inevitable importance as well, and enabling such computation is the second focus of our group. In particular, we have developed one of the most widely used software libraries for fast Fourier transforms (FFTs) by continuing innovations in *self-optimizing* software, which deals in an automated fashion with the burgeoning complexity of today's computer hardware.

1. Novel Waveguide Phenomena in Periodic Systems

Sponsors

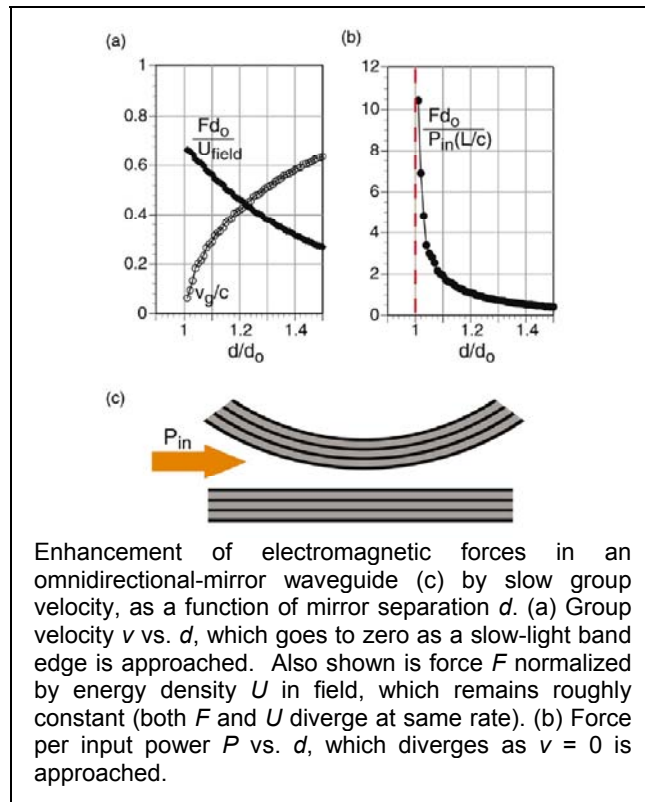
National Science Foundation - DMR-9400334

Project Staff

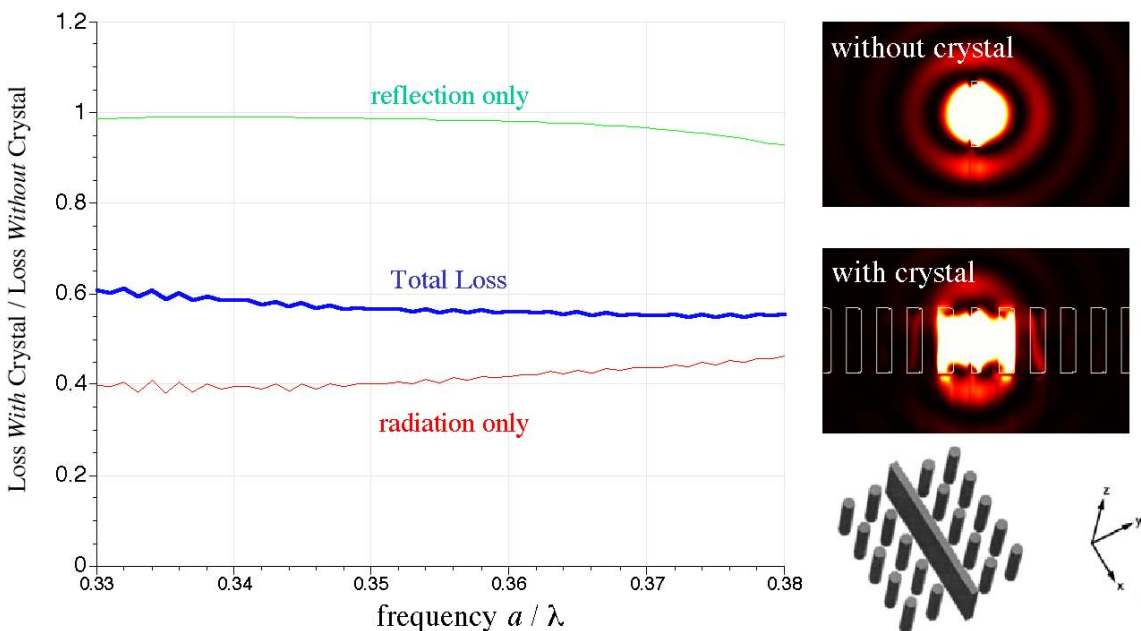
M. L. Povinelli, M. Ibanescu, Prof. M. Soljacic, Prof. S. G. Johnson, and Prof. J. D. Joannopoulos

One ongoing effort has been to explore novel phenomena in waveguides that is enabled by nanostructured media—what can a photonic crystal waveguide *do* that an ordinary waveguide cannot? Here, we will briefly summarize two examples of such phenomena: divergent forces and anomalous roughness losses.

Unlike a conventional dielectric waveguide, a photonic-crystal waveguide can guide light in an *air* core, and can have points in the frequency spectrum where the *group velocity* v (the rate at which energy propagates down the waveguide) goes to zero. We exploit both of these effects in the



structure shown in part (c) of the figure at right, which uses a pair of omnidirectional multilayer mirrors to guide light in the air region between them. As light propagates between the mirrors, there is a repulsive force, due to electromagnetic radiation pressure, that pushes the mirrors apart. Although ordinarily tiny, this force *diverges* as $1/v$ as the velocity goes to zero, simply because the field intensity for a fixed input power P diverges at that rate. Most interestingly, this divergence can be induced *dynamically* by changing the separation of the mirrors, since for a certain separation d_0 the band edge of the mode is shifted to the operating wavelength. So, instead of a fraction of a nano-Newton, the force can suddenly become tens or hundreds of nano-Newtons (assuming mW power levels at infrared lengthscales), well within the range that can be measured in MEMS devices—where sub-nN forces are regularly measured [see e.g. Chan et al., *Science* 291: 1941 (2001)]. This divergence is shown in part (b) of the figure, which plots the force F in dimensionless form as Fcd_0/PL , where L is propagation distance and c is the velocity of light. One therefore achieves an optically actuated mechanical device (instead of the more-typical mechanically actuated optical device).

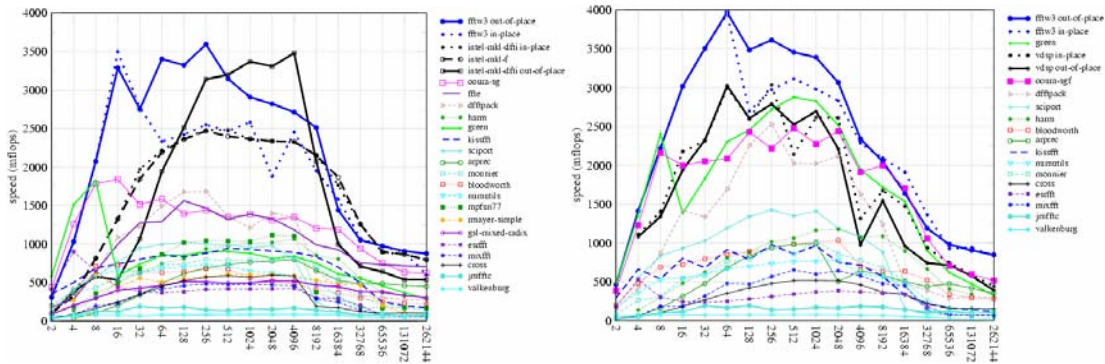


Loss from surface-roughness “bump,” modeled as a point source in the volume-current method, in a silicon strip waveguide surrounded by a 2d-periodic (period a) photonic crystal (lower right). Graph shows loss with/without crystal (1 = equal loss) vs. frequency. The crystal does not affect reflection loss, but greatly reduces radiation loss by blocking in-plane radiation, as seen in $\epsilon|\mathbf{E}|^2$ scattered-field cross-sections at right.

Another ongoing project has involved the study of losses due to disorder, such as surface roughness, in photonic-crystal waveguides. Although of tremendous practical impact, roughness is particularly challenging to study theoretically, for two reasons. First, because realistic levels of roughness require extremely high resolution to model by brute-force simulation. Second, because perturbative techniques that were developed decades ago to handle such small deviations analytically tend to be specialized to low-contrast, uniform waveguides—we have found recently that many classic perturbative methods tend to fail when applied to high-contrast, strongly periodic media such as photonic crystals. One powerful such method is the *volume-current method*, which is essentially a first Born approximation to the Green’s function that models a small perturbation $\Delta\epsilon$ in the dielectric function ϵ as a current $\mathbf{J} \sim \Delta\epsilon \mathbf{E}$. Unfortunately, this method requires correction for high-contrast media because of the discontinuities in the electric field \mathbf{E} , which is the subject of upcoming publications by our group. Fortunately, however, there exist

interesting 3d systems in which these discontinuities are negligible: in particular, for a 2d-periodic lattice of dielectric rods, shown at lower-right of the above figure, the field is mostly “TM” polarized (**E** parallel to the rods), in which case the field is continuous across the vertical interfaces and the classic volume-current method is applicable to the study of sidewall roughness. One simple case is analyzed in the figure above: a point-like bump will scatter light as a point source, whose radiated power we can compute with the finite-difference time-domain (FDTD) method. We then plot the *relative* loss of an isolated strip waveguide compared to the *same* waveguide surrounded by the photonic crystal [a waveguide described in Johnson et al., *Phys. Rev. B* 62:8212 (2000)], in which case the experimental roughness amplitude cancels and we obtain a universal figure of merit. There are two interesting features of this graph. First, it confirms prediction by our paper last year [Povinelli et al., *Appl. Phys. Lett.* 84:3639 (2004)] that a photonic bandgap does *not* affect reflection loss (which would lead to noise), all other things equal—this contradicts an intuitive picture that the crystal will reflect radiative scattering “back” into increased reflection. Second, the radiative loss is substantially reduced (by ~ 60%) by the crystal due to the photonic bandgap blocking radiative power in the plane of the crystal—even though a point scatterer like this radiates in all directions and this 2d crystal only has a bandgap in-plane.

There are other strategies to reduce waveguide losses, of course, such as using a “rib” waveguide that has lower losses than a high-contrast “strip” waveguide because the latter has a field that is spread out far from the interfaces. These conventional strategies, however, involve substantial sacrifices in waveguide confinement strength, requiring large-radius bends and other tradeoffs. Using a photonic crystal does *not* involve a tradeoff in modal volume, since it works merely by *removing* radiative loss channels from the local density of states. Although the reduction in total losses is limited because of the unchanged reflection loss (which, at least, is not worse), a factor of two reduction should still be possible and quite measurable. Moreover, this represents the first theoretical prediction that a photonic crystal can *reduce* disorder-induced losses of on-chip waveguides—most previous studies have focused on photonic-crystal waveguides that differ substantially from strip waveguides, and thus lead to many other effects that disguise the underlying loss reduction. A more extensive study and detailed description of the analytical methods should be available from our group in mid-2005.



Benchmarks of FFT performance, for power-of-two sizes on a 2.8GHz Pentium IV (left) and a 2GHz PowerPC G5 (right). The codes benchmarked are too numerous to describe, except for two: FFTW (solid blue line) and the vendor-tuned code for each CPU (solid black line = Intel Math Kernel Library or Apple vDSP library). Speed is given in “mflops” = $5 N \log_2 N /$ (time for FFT in μ s).

2. The Fastest Fourier Transform in the West

The demands of scientific and engineering computation have continued to outpace the performance of computers, so that extracting every ounce of speed from the hardware continues to be desirable in many applications. However, modern architectures have made it increasingly

difficult to perform this optimization: complicated cache and CPU pipeline behaviors mean that old heuristics, such as minimizing the number of operations performed, no longer even come close to the optimum. Moreover, thanks to the rapid range of change in the hardware arena, performance tuning of code is a never-ending process.

Instead, we believe that a more stable solution may be possible by changing the question: instead of asking what is the best algorithm, one should ask what is the smallest collection of simple algorithmic fragments whose composition spans the optimal algorithm on as many computer architectures as possible. We have pioneered this strategy—combining self-optimization, code generation, and recursive cache-oblivious algorithms—in the context of FFTW, our free and widely used fast Fourier transform (FFT) library (www.fftw.org).

“FFTW” stands for the “Fastest Fourier Transform in the West,” a title that is self-bestowed but is not entirely whimsical. We have benchmarked dozens of FFT implementations on numerous platforms, a small sampling of which is given in the figure above. While there is no single program that is the fastest for *all* FFT transform sizes on *all* computers, FFTW is always comparable to the best programs—including vendor-tuned codes painstakingly optimized by manufacturers such as Intel and Apple for their own machines. Moreover, it should be noted that the different programs benchmarked here differ in their arithmetic operation counts (the classic measure of FFT algorithm efficiency) by no more than 20%—the factor-of-ten performance differences are due, instead, entirely to the degree to which they exploit the memory hierarchies and CPU pipelines.

Although FFTW was first released in 1997 and has since become ubiquitous (for example, it is used for the `fft` function in Matlab), we continue to produce new developments. These graphs show the latest version 3 of FFTW, for which the inner workings were completely written to explore a much wider class of FFT algorithms, and to include such low-level improvements as support for SIMD (single-instruction multiple-data) instructions on modern CPUs.

We found that the key question was not so much what set of algorithmic steps to compose, but rather to identify what computational problem these steps should solve: an FFT works by decomposing a DFT into sub-problems, but in what form should these problems be described? One cannot implement algorithmic steps to solve a computational sub-problem that is not described, and conversely the problem should be simple enough that the corresponding algorithmic steps are simple and can be combined without restriction. By decomposing any DFT recursively in terms of a multi-dimensional “loop” of multi-dimensional transforms and corresponding data strides, we found it possible to define a few algorithmic steps whose compositions span a surprisingly wide variety of FFT algorithms.

Publications

Journal Articles, Published

- M. Frigo and S. G. Johnson, “The design and implementation of FFTW3,” *Proc. IEEE* 93(2): 216–31 (2005). Invited paper, special issue on adaptive software.
- P. T. Rakich, H. Sotobayashi, J. T. Gopinath, S. G. Johnson, J. W. Sickler, C.-W. Wong, J. D. Joannopoulos, and E. P. Ippen, “Nano-scale photonic crystal microcavity characterization with an all-fiber based 1.2–2.0 μm supercontinuum,” *Opt. Express* 13(3): 821–25 (2005).
- S. Assefa, P. T. Rakich, P. Bienstman, S. G. Johnson, G. S. Petrich, L. A. Kolodziejewski, E. P. Ippen, and H. I. Smith, “Guiding 1.5 μm light in photonic crystals based on dielectric rods,” *Appl. Phys. Lett.* 85(25):6110–12 (2004).

- A. Karalis, S. G. Johnson, and J. D. Joannopoulos, "Discrete-mode cancellation mechanism for high-Q integrated optical cavities with small modal volume," *Opt. Lett.* 29(19): 2309–11 (2004).
- M. L. Povinelli, M. Ibanescu, S. G. Johnson, and J. D. Joannopoulos, "Slow-light enhancement of radiation pressure in an omnidirectional-reflector waveguide," *Appl. Phys. Lett.* 85(9): 1466–68 (2004).
- M. Qi, E. Lidorikis, P. T. Rakich, S. G. Johnson, J. D. Joannopoulos, E. P. Ippen, and H. I. Smith, "A three-dimensional optical photonic crystal with designed point defects," *Nature* 429: 538–42 (2004).

Invited Meeting Papers, Presented

- S. G. Johnson *et al.*, "Designing for disorder in photonic crystals," MIT CIPS annual meeting, May 20 (2005).
- S. G. Johnson *et al.*, "Design and disorder in photonic crystals," at Nanophotonics for Information Systems OSA Topical Meeting, San Diego, CA, April 14 (2005).
- S. G. Johnson and M. Frigo, "FFTW: The 'Fastest Fourier Transform in the West,'" at Computational Research in Boston lecture series, MIT, April 1 (2005).
- S. G. Johnson, "Photonic crystals: A crash course in designer electromagnetism," at Central New England Chapter of IEEE LEOS, MIT Lincoln Laboratories, March 23–30 (2005).
- S. G. Johnson *et al.*, "Photonic crystals from order to disorder," at IBM Almaden Science Colloquium, San Jose, CA, Feb. 18 (2005).
- S. G. Johnson and M. Frigo, "FFTW: Towards a minimal compositional framework for high-performance FFTs," at ScalPerf 2005, Bertinoro, Italy, Sep. 19–23 (2004).