© DIGITAL VISION

A Tutorial on Fast Fourier Sampling

[How to apply it to problems]

[Anna C. Gilbert,
Martin J. Strauss, and
Joel A. Tropp]

Suppose that x is a discrete-time signal of length N that can be expressed with only m digital frequencies where $m \ll N$:

$$x[t] = \frac{1}{\sqrt{N}} \sum_{k=1}^{m} a_k e^{2\pi i \omega_k t/N}, \qquad t = 0, 1, 2, \ldots, N - 1.$$

We study the problem of identifying the unknown frequencies $\omega_1, \ldots, \omega_m$ that participate and their coefficients $a_1, \ldots, a_m$. Conceptually, the easiest way is to perform an N-point discrete Fourier transform (DFT):

IEEE SIGNAL PROCESSING MAGAZINE **[57]** MARCH 2008

$$X[\omega] = \frac{1}{\sqrt{N}} \sum_{t=0}^{N-1} x[t] e^{-2\pi i \omega t / N}, \qquad \omega = 0, 1, 2, \ldots, N-1.$$

Having obtained all N Fourier coefficients, it is straightforward to locate the m nonzero frequencies and their coefficients. Although you can compute the DFT efficiently by means of the fast Fourier transform (FFT), the fact remains that you must compute a very large number of zero coefficients when the signal involves few frequencies. This approach seems rather inefficient.

The discrete uncertainty principle [8] suggests that it might be possible to use fewer samples from the signal. Indeed, if the spectrum of a length-N discrete-time signal contains only m nonzero frequencies, then the time domain has at least N/m nonzero positions. As a result, even if we sample the signal at relatively few points in time, the samples should carry significant information about the spectrum of the signal.

This article describes a computational method, called the *Fourier sampling algorithm*, that exploits this insight [10]. The algorithm takes a small number of (correlated) random samples from a signal and processes them efficiently to produce an approximation of the DFT of the signal. The algorithm offers provable guarantees on the number of samples, the running time, and the amount of storage. As we will see, these requirements are exponentially better than the FFT for some cases of interest.

This article describes in detail how to implement a version of Fourier sampling, it presents some evidence of its empirical performance, and it explains the theoretical ideas that underlie the analysis. Our hope is that this tutorial will allow engineers to apply Fourier sampling to their own problems. We also hope that it will stimulate further research on practical implementations and extensions of the algorithm.

THE FOURIER SAMPLING ALGORITHM

A SUMMARY

We begin with a discussion of performance guarantees, so it is clear what the Fourier sampling algorithm can accomplish and what it cannot. The algorithm requires random access to the time domain of a signal x of length N. The input parameter m is the number of frequencies sought. As output, the algorithm produces a signal y that approximates x with only m frequencies:

$$y[t] = \frac{1}{\sqrt{N}} \sum_{k=1}^{m} a_k e^{2\pi i \omega_k t / N}. \qquad (1)$$

This approximation is represented by the set $\{(\omega_k, a_k) : k = 1, 2, \ldots, m\}$ of frequency/coefficient pairs. In a moment, we will see that the approximation error is comparable with the minimal error possible using a signal of the form (1).

The algorithm also involves several design parameters. The number $\varepsilon > 0$ determines the quality of the computed

approximation in comparison with the best approximation. The number $\delta > 0$ is the probability that the algorithm fails with respect to the random choices it makes during its execution. Both these quantities can be controlled by taking additional samples from the signal. The following theorem shows how all the factors interact.

THEOREM 1 [10]
Let x be an arbitrary signal of length N. The Fourier sampling algorithm takes m poly $(\varepsilon^{-1}, \log(\delta^{-1}), \log(N))$ random samples of the signal. (The term poly$(\cdot)$ indicates an unspecified polynomial in its arguments.) With probability at least $1 - \delta$, the algorithm returns an approximation y of the form (1) that satisfies the error bound

$$\|x - y\|_2 \le (1 + \varepsilon)\|x - x_{\text{opt}}\|_2 + \varepsilon$$

where x_{opt} is the best approximation to x of the form (1). It produces this approximation using time and storage m poly $(\varepsilon^{-1}, \log(\delta^{-1}), \log(N))$.

Let us elaborate on the statement of this theorem. When $m \ll N$, the algorithm takes far fewer samples than the total length of the signal. We emphasize that the sample set depends on random choices, but it does not depend on the signal or the progress of the algorithm. Therefore, the sample locations can be established before execution. Moreover, the run time and storage requirements of the algorithm are roughly proportional to the number of frequencies it outputs, rather than the signal length. All the resource requirements are logarithmic in N, the signal length, so Fourier sampling has the potential to be exponentially faster than the FFT.

Second, let us discuss how to interpret the approximation guarantee. When the signal is well approximated by a set of m frequencies, the right-hand side of the error bound is small, so the algorithm produces an approximation that is competitive with the optimal m-frequency approximation. In contrast, when it takes many frequencies to approximate the signal, the algorithm will return a poor result. In this setting, Fourier sampling is not an appropriate tool.

Even if the true signal consists of m frequencies contaminated with heavy noise, the algorithm may not return the m ideal frequencies. Indeed, the theorem only promises that the error in the output approximation is comparable to the amount of noise. Nevertheless, a careful analysis shows that the energy of the noise must be substantial compared with the signal energy before the algorithm delivers frequencies different from the ground truth.

A more familiar way to analyze the quality of the approximation y is to compute its reconstruction signal-to-noise ratio (SNR). Suppose that our signal x consists of m frequencies plus an orthogonal noise vector ν. Then

$$\text{SNR} = 10 \log_{10} \left(\frac{\|x\|_2}{\|x - y\|_2} \right) \ge 10 \log_{10} \left(\frac{\|x\|_2}{(1 + \epsilon)\|\nu\|_2} \right).$$

Consequently, the SNR of the reconstructed signal is smaller than optimal by an additive term. We can reduce this loss by decreasing the design parameter ϵ, although this revision results in additional samples of the signal and increased computation time.

RISING TO THE CHALLENGE

The fundamental challenge for the Fourier sampling algorithm is to divine information about the frequency spectrum of a signal under severe constraints on the number of samples and arithmetic operations. To do so, the algorithm makes random choices to avoid worst-case scenarios. This means that the procedure has access to random bits separate from and in addition to its input. In its execution, the algorithm uses those random bits to guide its behavior. *For each input*, it succeeds with high probability with respect to the source of randomness. This idea is substantially different from the use of statistical signal models, so practitioners of signal processing may be less familiar with it. Here are the key observations:

1) Random time samples of a signal allows us to estimate certain characteristics, such as its zero-frequency coefficient and its energy.

2) Random permutation of the spectrum allows us to separate significant tones into different frequency bands. The tones can then be isolated with bandpass filters.

The algorithm also exploits many standard methods from the DSP toolbox:

1) Sampling in time corresponds to summing modulated Fourier coefficients.

2) Dilation in time corresponds to dilation of the spectrum.

3) Modulation in time corresponds to translation of the spectrum.

4) The FFT can be used to apply a filter bank, which multiplies the spectrum of the signal by a collection of transfer functions.

5) (Nonuniform) FFTs allow fast evaluation of exponential polynomials at multiple points.

The algorithm combines these ideas in a way that is complicated and—perhaps—unusual. Later on in this article, we provide more detailed information about how these methods allow us to approximate the spectrum of an unknown signal.

BACKGROUND AND RELATED WORK

The Fourier sampling algorithm differs from traditional spectral estimation techniques in a variety of ways. First, unlike Prony's method [6] and its more stable variations [17], the algorithm is not predicated upon evenly spaced samples—just the opposite. Second, the reconstruction algorithm uses the samples in a nonlinear fashion, unlike the procedures of [12]. It does not form a linear combination of the sample values. Third, the algorithm and its analysis are inherently discrete. The samples are drawn from a discrete-

time signal (rather than an underlying continuous-time signal) and the output of the algorithm is an approximation to the discrete spectrum.

The Fourier sampling algorithm is related to the compressive sampling paradigm, but the two approaches focus on different resource constraints. Let us consider the case where signals of interest have few significant frequencies in comparison with their length. The primary concern of compressive sampling is to reconstruct the spectrum of the signal from as few samples as possible with extremely strong guarantees on the probability of success. Researchers have established that several different randomized sampling schemes are compatible with this goal [5], [4]. Most of the literature concentrates on reconstruction algorithms such as convex programming, but other methods are available [3], [11]. The Fourier sampling algorithm is closest in spirit to the algorithms in [14] and [15].

EMPIRICAL PERFORMANCE

The Fourier sampling algorithm has been implemented and tested in a variety of settings to assess its empirical performance. We discuss one particular implementation [13], the Ann Arbor fast Fourier transform (AAFFT), and we provide evidence that it is both powerful and resource efficient.

First, we consider a problem inspired by communication devices that use frequency-hopping modulation schemes. Suppose we wish to recover a synthetic signal consisting of two tones that change at regular intervals. These signals are contaminated with white Gaussian noise so the SNR is -17 dB. We apply the AAFFT implementation to identify the location of the two tones. Figure 1 exhibits the output using a *sparsogram*, which is a time-frequency plot that displays only


[FIG1] The sparsogram for a synthetic frequency-hopping signal consisting of two tones in noise, as computed by the AAFFT.

the dominant frequencies in the signal. As a benchmark, we also computed the sparsogram with fastest Fourier transform in the west (FFTW), a highly optimized implementation of the FFT. Both AAFFT and FFTW obtain the correct result in the same amount of time, but AAFFT samples only 3% of the signal—a factor $33\times$ *undersampling*.

This first experiment provides evidence that AAFFT uses far fewer samples than FFTW. The AAFFT implementation is also substantially faster than FFTW for long signals. To prove this point, we constructed (noiseless) signals of different lengths by selecting 60 frequencies at random and assigning them unit coefficients. We compared the running time for AAFFT to identify these 60 frequencies with the running time for FFTW using a log-log scale. The result appears in Figure 2. Notice

[FIG2] The execution time of FFTW and AAFFT for recovering 60 frequencies without noise. The error bars indicate the minimum and maximum run times. The AAFFT runs faster than FFTW when the signal length exceeds 10^6.

[FIG3] Phase transition diagram for recovering one frequency in noise. The plot uses shades of gray to indicate the probability of successful recovery as a function of the SNR and the percentage of the signal that is sampled. The undersampling rate is the reciprocal of the sampling percentage.

that the execution time of FFTW grows dramatically while the speed of the AAFFT remains virtually constant as the signal length varies across several orders of magnitude.

At the beginning of the tutorial, we mentioned that the algorithm may fail completely to approximate the signal. The failure probability can be controlled by increasing the number of samples of the signal that we take. We constructed signals of length $N = 2^{22}$ (about 4 million) containing one tone in additive white Gaussian noise, and we attempted to locate the frequency with AAFFT. For each sampling rate, we performed 1,000 trials and computed the fraction of those trials in which the tone was successfully identified. Figure 3 is a phase transition chart that indicates the probability of recovering a single frequency in heavy noise. We see, for example, that AAFFT can recover the tone 90% of the time at an SNR of -15 dB with $100\times$ undersampling. This rate is fully two orders of magnitude below Nyquist.

We also studied the number of samples necessary to recover a larger number of frequencies. We fixed the signal length at $N = 2^{22}$ and measured the number of samples necessary to recover m tones at least 99% of the time. Figure 4 displays the results. For example, if we sample 10% of the signal, the AAFFT implementation can recover 1,000 tones over 99% of the time. If an application can tolerate a higher failure probability, then AAFFT can recover more tones with the same level of undersampling.

IMPLEMENTATION

This section gives an overview of a simplified version of the AAFFT implementation, including explicit pseudocode (see "Algorithm 1"). This version assumes that N is a power of two, and it removes some failure-control mechanisms. The complete AAFFT algorithm is somewhat more complicated than the code here, but this basic implementation still works quite well.

Let us note that the upcoming description of the algorithm interleaves the sampling of the signal with the other actions of the algorithm. We have elected this description to make it clear precisely where samples are required. Nevertheless, we emphasize that the samples used by the algorithm are totally nonadaptive. In particular, it is possible to select the sample points and draw the samples from the signal prior to runtime.

The algorithm iteratively constructs an approximation y to the input signal. As it runs, the algorithm represents the approximation as a list Λ of at most K frequency/coefficient pairs: $\Lambda = \{(\omega_k, a_k) : k = 1, 2, \ldots, K\}$. The approximation y is implicitly determined via (1). The approximation also induces a *residual signal $r = x - y$*. The most critical parameter in the algorithm is the size K of frequency list. In the pseudocode, we have chosen $K = 8m$. Increasing the factor eight improves accuracy at the cost of additional samples and arithmetic.

The algorithm often needs to determine the value of the residual signal at designated sampling locations. Since the approximation has the form (1), the SAMPLE-RESIDUAL subroutine is able to perform this computation efficiently with a nonuniform FFT. The literature describes several approaches to computing nonuniform FFTs, including [2], [1]. Explicit pseudocode is also available [7,

---

ALGORITHM 1: SIMPLIFIED FOURIER SAMPLING ALGORITHM

Fourier Sampling $(\boldsymbol{x}, m)$
Input: Input signal $\boldsymbol{x}$ of length $N = 2^\alpha$ and number m of frequencies to find
Output: Set $\Lambda = \{(\omega, a_\omega)\}$ containing $O(m)$ frequency/coefficient pairs

$K \leftarrow 8m$ and $\Lambda \leftarrow \emptyset$
for $j = 1$ to 5
    $\Omega \leftarrow$ Identification $(\boldsymbol{x}, \Lambda, K)$ {Identify K frequencies in the residual}
    $\boldsymbol{c} \leftarrow$ Estimation $(\boldsymbol{x}, \Lambda, \Omega)$ {Estimate Fourier coefficients}
    for each frequency $\omega \in \Omega$ and corresponding coefficient c_ω
            if $(\omega, a_\omega) \in \Lambda$ for some a_ω then replace the pair with $(\omega, a_\omega + c_\omega)$
                  else add the new pair (ω, c_ω) to Λ
    Retain K pairs from Λ whose coefficients have greatest magnitude
Retain m pairs from Λ whose coefficients have greatest magnitude

Sample-Residual $(\boldsymbol{x}, \Lambda, t, \sigma, K)$
for $k = 1$ to K
    $u_k \leftarrow x[t + \sigma(k-1) \bmod N]$ {Arithmetic progression from signal}
    $v_k \leftarrow \sum_{(\omega, a_\omega) \in \Lambda} (a_\omega e^{2\pi i \omega t/N}) e^{2\pi i(\omega \sigma/N)(k-1)}$ {In parallel, via nonuniform FFT}
return $(\boldsymbol{u} - \boldsymbol{v})$ {Residual is signal minus approximation}

Identification $(\boldsymbol{x}, \Lambda, K)$
reps $\leftarrow 5$ and $\omega_k \leftarrow 0$ **for** $k = 1, 2, \ldots, K$
Draw $\sigma \sim$ Uniform $\{1, 3, 5, \ldots, N-1\}$ {Random shattering parameter}
for $b = 0$ to $\log_2(N/2)$ {Loop from LSB to MSB}
    $\text{vote}_k \leftarrow 0$ for $k = 1, 2, \ldots, K$
    for $j = 1$ to reps
        Draw $t \sim$ Uniform $\{0, 1, 2, \ldots, N-1\}$ {Random sample point}
        $\boldsymbol{u} \leftarrow$ Sample-Shattering (t) {Samples correlated for testing bth bit}
        $\boldsymbol{v} \leftarrow$ Sample-Shattering $(t + N/2^{b+1})$
        for $k = 1$ to K
            $E_0 \leftarrow u_k + e^{-\pi i \omega_k/2^b} v_k$ {Apply bit-test filters to demodulated signal}
            $E_1 \leftarrow u_k - e^{-\pi i \omega_k/2^b} v_k$
            if $|E_1| \geq |E_0|$ **then** $\text{vote}_k \leftarrow \text{vote}_k + 1$ {Vote when bit is one}
    for $k = 1$ to K
        if $\text{vote}_k > \text{reps}/2$ **then** $\omega_k \leftarrow \omega_k + 2^b$ {Majority vote for bit value}
return Unique$(\omega_1, \omega_2, \ldots, \omega_K)$ {Remove duplicates}

Sample-Shattering (p)
$z \leftarrow$ Sample-Residual $(\boldsymbol{x}, \Lambda, p, \sigma, K)$ {Get arithmetic progression of samples}
$z \leftarrow \text{FFT}(z)$ {Apply sub-band decomposition filter bank}
return z

Estimation $(\boldsymbol{x}, \Lambda, \Omega)$
reps $\leftarrow 5$
for $j = 1$ to reps
    Draw $\sigma \sim$ Uniform $\{1, 3, 5, \ldots, N-1\}$ and $t \sim$ Uniform $\{0, 1, 2, \ldots, N-1\}$
    $\boldsymbol{u} \leftarrow$ Sample-Residual $(\boldsymbol{x}, \Lambda, t, \sigma, K)$
    for $\{\ell = 1$ to $|\Omega|$
        $c_\ell(j) \leftarrow \sum_{k=1}^{K} u_k e^{2\pi i(\omega_\ell \sigma/N)(k-1)}$ {In parallel, via nonuniform FFT}
        $c_\ell(j) \leftarrow (N/K) e^{-2\pi i \omega_\ell t/N} c_\ell(j)$ {Demodulate and scale estimates}
$c_\ell \leftarrow$ Median$\{c_\ell(j) : j = 1, 2, \ldots, \text{reps}\}$ **for** $\ell = 1, 2, \ldots, |\Omega|$ {Do real, imaginary separately}
**return** $c_1, c_2, \ldots, c_{|\Omega|}$

Alg. 1]. Alternatively, the exponential sums can be evaluated directly at somewhat higher cost.

At the highest level, the algorithm proceeds as follows. First, the approximation is set to zero, so the residual signal equals the input signal. The algorithm iteratively refines the approximation, as described in the next paragraph. After the iteration is complete, the algorithm reduces the list Λ by picking m frequencies with the largest coefficients.

The main loop consists of the three steps. First, the identification stage constructs a list Ω containing K frequencies that are likely to carry a significant amount of the energy from the residual. Second, the estimation stage finds estimates of the coefficients of the frequencies in Ω. Third, the algorithm adds the new approximation to the previous approximation to get a total of $2K$ terms (or fewer). Frequencies with small coefficients are likely to be spurious, so the list is reduced to K terms (or fewer) by retaining only the frequencies with the largest coefficients. Our experience suggests that the main loop should be repeated about five times.

The IDENTIFICATION subroutine employs a randomized filtering process to find up to K significant frequencies from the residual signal. Beginning with the least-significant bit, it determines each bit from all K frequencies in parallel. In the inner loop, the subroutine performs several repetitions to drive down the failure probability. Our experience suggests that three to five repetitions are adequate.

The ESTIMATION subroutine uses a related randomized filtering process to estimate simultaneously the coefficients of K given frequencies in the residual signal. This calculation involves the adjoint of the nonuniform FFT. Explicit pseudocode appears in [7, Alg. 2]. The sums can also be evaluated directly at higher cost.

> THE FUNDAMENTAL CHALLENGE FOR THE FOURIER SAMPLING ALGORITHM IS TO DIVINE INFORMATION ABOUT THE FREQUENCY SPECTRUM OF A SIGNAL UNDER SEVERE CONSTRAINTS ON THE NUMBER OF SAMPLES AND ARITHMETIC OPERATIONS.

The subroutine takes the median of several copies of the estimator to improve robustness. In practice, three to five copies suffice.

THE CONCEPTS BEHIND THE CODE
The Fourier sampling algorithm must perform computations on the frequency spectrum of a signal under severe constraints on the number of samples and arithmetic operations. It is possible to achieve an economy of scale by attempting to find all the significant frequencies at once. The central design principle in the algorithm is to exploit this economy whenever possible by means of filter banks, nonuniform FFTs, and random sampling. This section describes the intuitions behind the key steps in the algorithm. In the sequel, we assume that the signal length N is a power of two, that the signal takes complex values, and that all arithmetic on the indices is performed modulo N.

THE ROLE OF RANDOMNESS
In contrast with the field of statistical signal processing, we do not make any assumptions about the input signal. Instead, the algorithm makes random choices during its execution to enable it to succeed with high probability for any given input signal. In this section, we attempt to share the flavor of these techniques.

RANDOM SAMPLING
Random sampling is a very efficient method for estimating some key characteristics of a signal. Let x be a signal of length N, and let T be the random variable that takes each value from $\{0, 1, 2, \ldots, N-1\}$ with equal probability.

First, the squared magnitude of a random time sample gives a good estimate for the signal energy because $\|x\|^2 = N \, \mathbb{E} \, |x[T]|^2$. Owing to Markov's inequality, it is unlikely that a random sample, suitably normalized, has magnitude much greater than the norm of the signal.

Second, consider a signal containing one large frequency plus noise: $x[t] = a e^{2\pi i \omega t/N}/\sqrt{N} + v[t]$. A short argument involving the triangle inequality and Jensen's inequality yields

$$|a| - \|v\|_2 \leq \sqrt{N} \, \mathbb{E}|x[T]|.$$

That is, we can approximate the magnitude of the tone by random sampling. Therefore, we can find the location of a tone that dominates a signal.

Finally, the scaled expectation of a random sample equals the zero-frequency component of a signal. This point follows from the simple fact $X[0] = \sqrt{N} \, \mathbb{E} x[T]$. The algorithm uses

[FIG4] The proportion of the signal that AAFFT samples to recover a fixed number of tones in a signal of length $N = 2^{22}$ at least 99% of the time.

this fact to estimate the coefficient of a specified frequency.

RANDOM SPECTRAL PERMUTATION

A major difficulty is that significant tones in the spectrum of a signal can be clustered together or spread out. One of the central innovations in the algorithm is a randomized technique for isolating significant tones from each other so we can perform spectral analysis using bandpass filters.

To explain, we need some basic number theory. Two numbers are *relatively prime* if they have no common integer factor except ± 1. Since N is a power of two, the numbers relatively prime to N are precisely the odd integers. Given an odd number σ, the Euclidean division algorithm furnishes a number σ^{-1}, called its *multiplicative inverse*, that satisfies $\sigma \cdot \sigma^{-1} \equiv 1 (\mod N)$.

Let σ be odd, and consider the dilation $d_\sigma : t \mapsto \sigma t \mod N$. It is not hard to see that this map is a permutation on the set $\{0, 1, 2, \ldots, N-1\}$ and that its inverse is the map $d_{\sigma^{-1}}$. In discrete Fourier analysis, these observations lead directly to the identity

$$y[t] = x[\sigma t] \quad \text{for all } t \iff Y[\omega] = X[\sigma^{-1}\omega] \quad \text{for all } \omega.$$

Succinctly, time dilations generate frequency permutations. See Figure 5 for an illustration.

The key idea is to choose σ at random from the set $\{1, 3, 5, \ldots, N-1\}$ of invertible numbers. Applying the dilation d_σ to the signal, we produce a random permutation of its spectrum. It is unlikely that a given pair of frequencies is mapped to the same part of the spectrum. Roughly speaking, random permutation of the spectrum isolates significant frequencies from each other.

IDENTIFICATION

The first stage in the Fourier sampling algorithm is to identify a collection of frequencies whose coefficients are large relative to the signal energy. The identification process consists of two conceptual steps, *shattering* and *bit testing*. Shattering generates a collection of signals, many of which have a single dominant frequency. Then bit tests are applied to each signal to find the location of the dominant frequency, one bit at a time.

[FIG5] The impact of dilations on the spectrum of a frequency-sparse signal. The top-right panel is generated by the time dilation $t \mapsto 11t \mod N$ (equivalently, by a frequency dilation $\omega \mapsto 3\omega \mod N$); the bottom-right panel by $t \mapsto 9t \mod N$ (or $\omega \mapsto 9\omega \mod N$). Note that the zero frequency always remains fixed.

SHATTERING

A *shattering* of a signal x is a collection $\{x_0, x_1, \ldots, x_{K-1}\}$ of signals that are formed by a three-step filtering process. First, we randomly permute the spectrum of the signal to isolate significant frequencies from each other. Second, we apply a sub-band decomposition filter bank to create K signals that each carry a chunk of the permuted spectrum. Each significant tone in the original signal is likely to dominate one signal in the shattering. Finally, we invert the dilation to restore the frequencies to their original places in the spectrum. Figure 6 exhibits a block diagram, and Figure 7 illustrates the effects of shattering on a signal.

The design of the subband decomposition filter bank is simple. Let h be a low-pass filter with K taps whose cutoff frequency is about π/K rad. The filter bank consists of K frequency translates of this filter, spaced $2\pi/K$ rad apart. In the time domain, this amounts to convolution with $h_k[t] =$

[FIG6] A conceptual block diagram for the shattering process. A shattering of x contains K elements $x_0, x_1, \ldots, x_{K-1}$.

$e^{-2\pi i k t/K} h[t]$ for each $k = 0, 1, \ldots, K - 1$. The analysis in [10] suggests that the ideal filter has minimal energy among all normalized filters with K taps. This observation recommends the boxcar filter $h[j] = \sqrt{N}/K$ for $j = 0, 1, \ldots, K - 1$. It is possible that more sophisticated low-pass filters or windows will sometimes yield better results [16, Ch. 7].

Let us emphasize that the algorithm never forms a shattering explicitly. Instead, the filter bank is constructed so we can take one time sample from each element of the shattering by processing K samples from the input signal. If σ is the parameter of the random dilation, the kth signal in the shattering satisfies

$$x_k[t] = \sum_{j=0}^{K-1} h[j] \, x[t - \sigma j] \, e^{-2\pi i j k/K}.$$

Given a point t, we can simultaneously calculate $x_0[t]$, $x_1[t]$, $\ldots, x_{K-1}[t]$ by extracting an arithmetic progression from the input signal, multiplying it with the filter, and applying an FFT. The subroutine SAMPLE-SHATTERING performs these actions.

BIT TESTING

Shattering generates a collection of signals, some of which contain a single dominant frequency. The bit-test process is designed to locate the dominant frequency in such a signal. (The bit tests are likely to return spurious frequencies for other elements of the shattering.)

Suppose that x is a length-N signal in which a single frequency carries most of the energy. We find the bits of the frequency sequentially, beginning with the least significant bit $b = 0$. Assuming we alreday know the least-significant $(b - 1)$ bits, we can demodulate the signal so that the binary expansion of the dominant frequency ends in $10 \ldots 0$ or in $00 \ldots 0$. We apply the frequency mask filters

$$g_b^{\text{even}} = \frac{1}{2}(\delta_0 + \delta_{N/2^{b+1}}) \quad \text{and} \quad g_b^{\text{odd}} = \frac{1}{2}(\delta_0 - \delta_{N/2^{b+1}}).$$

The filter g_b^{even} passes the even frequencies and zeros the odd frequencies (mod 2^b). In a similar fashion, the filter g_b^{odd} passes the odd frequencies and zeros the even ones (mod 2^b). If E_0 is a random sample from the output of the even filter and E_1 is a random sample from the odd filter, then the inequality $|E_1| \geq |E_0|$ is evidence that the least-significant bit is one.

Therefore, the bit test compares the magnitude of a random sample from each of the two filtered signals. It repeats the comparison several times, and it takes a majority vote to reduce the failure probability. See Figure 8 for a block diagram of the bth bit test. Note that bit testing is computationally efficient since each filter has only two taps.

IMPLEMENTATION

We separate the *concepts* of shattering and bit testing, but the code must intertwine them for efficiency. Recall that we can simultaneously compute one sample from each of the K elements in the shattering. To exploit this fact, we simultaneously test the bth bit of the dominant frequency in each element of the shattering using two correlated samples, demodulated by the first $(b - 1)$ bits of that frequency. Details appear in the IDENTIFICATION subroutine.
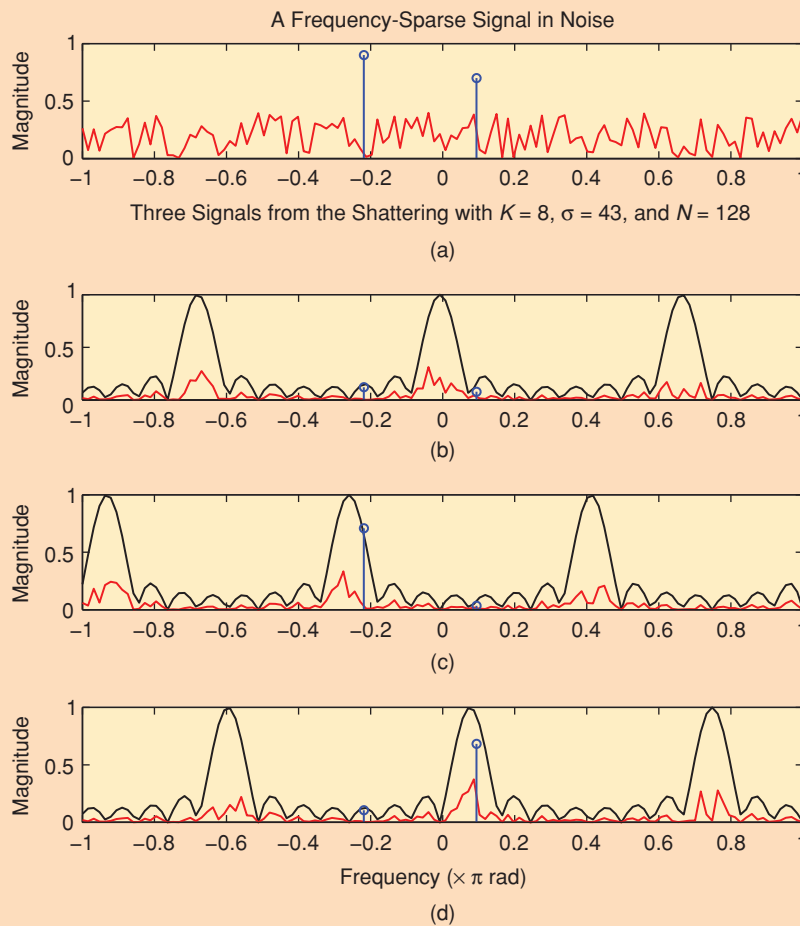
[FIG7] A signal consisting of two tones in noise, along with three elements from a shattering. In (b), both the tones are attenuated so neither is recoverable. In (c) and (d) elements are shown where one tone is preserved and the other is attenuated. In each plot, the transfer function is traced in black.

ESTIMATION OF FOURIER COEFFICIENTS

The identification phase of the algorithm returns a list of K or fewer frequencies, but it does not provide sufficient information about their coefficients. The next stage of the algorithm estimates the coefficients using a randomized filtering technique.

Suppose we want to estimate the coefficient of a significant frequency ω in a signal x. First, we demodulate the signal by ω so we can estimate the zero frequency instead. Next, we randomly dilate the signal. This operation fixes the zero frequency and shuffles the other significant frequencies around (probably away from zero). Third, we apply a filter to pass the zero frequency and attenuate the rest of the spectrum. Then, we invert the dilation. At this point, the zero frequency is likely to dominate the signal. We estimate its coefficient by taking a random sample and scaling appropriately. See Figure 9 for a block diagram.

As in shattering, we cannot afford to use a filter with more than K taps. It turns out that the boxcar filter $h[k] = \sqrt{N}/K$ for $k = 0, 1, 2, \ldots, K-1$ remains a good choice in this setting.

We can write down the cumulative effect of this random filtering process. Let σ be the parameter of the random dilation, and let t be the sample location. Then the coefficient estimate c_ω is

$$c_\omega = \sqrt{N} e^{-2\pi i \omega t/N} \sum_{j=0}^{K-1} h[j]\, s[t - \sigma j]\, e^{-2\pi i (\omega \sigma K/N)\, j/K}.$$

One should view this expression as the DFT of a sequence of length K, evaluated at the nonintegral frequency $2\pi \omega \sigma K/N$ radians, then demodulated and scaled. Therefore, we can simultaneously estimate the coefficients of a collection of K frequencies (or fewer) using the adjoint nonuniform FFT. Afterward, we can demodulate

each coefficient individually. Finally, we make each coefficient estimator more robust by taking the median of several copies. (The medians of the real and imaginary parts are performed separately.)

ITERATION

The *recoverable energy* in a signal is the energy carried by the largest m frequencies. It is impossible to collect more since our approximation contains only m frequencies. By performing identification and estimation once, the algorithm finds a constant proportion of the recoverable energy in the residual. Therefore, after a constant number of iterations, the algorithm can find a fixed proportion of the recoverable energy in the original signal.

EXTENSIONS AND IMPROVEMENTS

Although the Fourier sampling algorithm is designed for discrete-time signals, we can use it in certain analog settings with some modifications. We can acquire a few random structured samples of a wide-band, continuous-time signal that has a few significant tones and recover quickly those tones present. To build a practical system, we must analyze carefully the required minimum sample spacing as it is costly to acquire signal samples close in time. We

[FIG8] A conceptual block diagram for the b th bit test. The test yields a 0/1 bit value ε and a demodulated copy of the input signal for the next bit test.

[FIG9] A conceptual block diagram for the coefficient estimator. The system returns an estimate for the coefficient c_ω of the frequency ω in the signal x. Note that the estimate must also be scaled by $\sqrt{N}$.

must also increase the flexibility of the output representation so as to improve the reconstruction SNR; instead of returning exactly m, we return a (tunable) multiple of m. All of these modifications are possible while still preserving the structure and quality guarantees of the algorithm. We must be realistic, however, in assessing the quality of our output. The algorithm returns a compressed or approximate representation of the discrete spectrum of an inherently analog signal. It only approximates the significant portions of the discrete spectrum, which, itself, is an approximation to the true spectrum.

> **THE FOURIER SAMPLING ALGORITHM IS RELATED TO THE COMPRESSIVE SAMPLING PARADIGM, BUT THE TWO APPROACHES FOCUS ON DIFFERENT RESOURCE CONSTRAINTS.**

AUTHORS
Anna C. Gilbert (annacg@umich.edu) has an S.B. degree from the University of Chicago and a Ph.D. from Princeton University, both in mathematics. In 1997, she was a postdoctoral fellow at Yale University and AT&T Labs-Research. From 1998 to 2004, she was a member of technical staff at AT&T Labs-Research in Florham Park, NJ. Her research interests include analysis, probability, networking, and algorithms. She is especially interested in randomized algorithms with applications to harmonic analysis, signal and image processing, networking, and massive data sets.

Martin J. Strauss (martinjs@umich.edu)is an assistant professor in the Mathematics and EECS Departments at the University of Michigan. He received an A.B. in mathematics from Columbia and a Ph.D. in mathematics from Rutgers. Between finishing his Ph.D. in 1995 and starting at Michigan in 2004, Strauss was a postdoctoral fellow at Iowa State University and moved to AT&T Labs–Research. His research interests include fundamental algorithms, especially randomized and approximation algorithms; algorithms for massive data sets; signal processing and and computational harmonic analysis; computer security and cryptography; and complexity theory.

Joel A. Tropp (jtropp@acm.caltech.edu) recieved the B.A. in liberal arts and the B.S. in mathematics in 1999 and the M.S. and Ph.D. in computational and applied mathematics in 2001 and 2004, respectively, all from the University of Texas, Austin. In 2004, he joined the Mathematics Department of the University of Michigan at Ann Arbor as Research Assistant Professor. In 2005, he won an NSF Mathematical Sciences Postdoctoral Research Fellowship, and he was appointed T.H. Hildebrandt Research Assistant Professor. Since August 2007, he has been an assistant professor of computational and applied mathematics at the California Institute of Technology. He is a Member of the IEEE.

REFERENCES
[1] C. Anderson and M.D. Dahleh, "Rapid computation of the discrete Fourier transform," *SIAM J. Sci. Comput.*, vol. 17, no. 4, pp. 913–919, 1996.

[2] G. Beylkin, "On the Fast Fourier Transform of functions with singularities," *Appl. Comp. Harmonic Anal.*, vol. 2, no. 4, pp. 363–381, 1995.

[3] G. Cormode and S. Muthukrishnan, "Combinatorial algorithms for compressed sensing," in *Proc. 2006 IEEE Int. Conf. Information Sciences Systems*, Princeton, NJ, Apr. 2006, pp. 230-294.

[4] E.J. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[5] D.L. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[6] G.R. de Prony, "Essai éxperimental et analytique: Sur les lois de la dilatabilité de fluides élastique et sur celles de la force expansive de la vapeur de l'alkoo, á différentes températures," *J. de l'École Polytechnique*, vol. 1, no. 22, pp. 24–76, 1795.

[7] A. Dutt and V. Rokhlin, "Fast Fourier transforms for nonequispaced data," *SIAM J. Sci. Comput.*, vol. 14, no. 6, pp. 1368–1393, 1993.

[8] D.L. Donoho and P.B. Stark, "Uncertainty principles and signal recovery," *SIAM J. Appl. Math.*, vol. 49, no. 3, pp. 906–931, 1989.

[9] A.C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M.J. Strauss, "Near-optimal sparse Fourier representations via sampling," in *ACM Symp. Theoretical Comput. Sci.*, pp. 152-161, 2002.

[10] A.C. Gilbert, S. Muthukrishnan, and M.J. Strauss, "Improved time bounds for near-optimal sparse Fourier representations," in *Proc. SPIE Wavelets XI*, M. Papdakis, A.F. Laine, and M.A. Unser, Eds., San Diego, CA, 2005, pp. 59141A.1–15.

[11] A.C. Gilbert, M.J. Strauss, J.A. Tropp, and R. Vershynin, "One sketch for all: Fast algorithms for Compressed Sensing," in *Proc. 39th ACM Symposium on Theory of Computing*, San Diego, June 2007, pp. 237-246.

[12] G. Harikumar and Y. Bresler, "FIR perfect signal reconstruction from multiple convolutions: minimum deconvolver orders," *IEEE Trans. Signal Processing*, vol. 46, no. 1, pp. 215-218, 1998.

[13] M. Iwen, A.C. Gilbert, and M.J. Strauss, "Empirical evaluation of a sub-linear time sparse DFT algorithm," *Commun. Math. Sci.*, vol. 5, no. 4, pp. 981-998, 2007.

[14] E. Kushilevitz and Y. Mansour, "Learning decision trees using the fourier spectrum." in *Proc. ACM Sym. Theory Computing*, 1991, pp. 455–464.

[15] Y. Mansour, "Randomized interpolation and approximation of sparse polynomials," *SIAM J. Sci. Comput.*, vol. 24, no. 2, pp. 357–368, 1995.

[16] A.V. Oppenheim, R.W. Schafer, and J.R. Buck, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1999.

[17] G.K. Smyth and D.M. Hawkins, "Robust frequency estimation using elemental sets," *J. Comput. Graph. Stat.*, vol. 9, no. 1, pp. 196–214, 2000.

SP