FUNDAMENTAL PERFORMANCE BOUNDS FOR A COMPRESSIVE SAMPLING SYSTEM

A. C. Gilbert   M. J. Strauss

Department of Mathematics
University of Michigan
Ann Arbor, MI 48109

ABSTRACT

In this paper we explore several fundamental bounds for a compressive sampling system which uses the Fourier sampling algorithm of Gilbert et al. Beginning with the theoretical bounds on the number of samples necessary to reconstruct high fidelity approximations of input signals, we refine those theoretical bounds with empirical values in several practical input models. We show that the performance is consistent with traditional sampling systems, and, in certain cases, much better.

Index Terms— analog-digital conversion, algorithms

1. INTRODUCTION

Recent results in compressive sampling [1, 2, 3, 4, 5, 6] demonstrate that it is theoretically possible to record a few linear observations of a signal and, from those measurements, nearly reconstruct the signal. The number of observations is considerably smaller than the extrinsic dimension of the signal if the signal is highly compressible. The Fourier sampling algorithm of Gilbert et al. [6] presents itself as an attractive, potentially practical algorithm for carrying out a type of compressive sampling—it yields fast algorithms (exponentially faster than those in [1, 2]) and both the sampling pattern and the reconstruction algorithm are easily implementable in hardware [7].

While the Fourier sampling algorithm is attractive in theory, suggests novel analog-to-digital converter (ADC) designs that defy the state-of-the-art [7], and in implementation exhibits good empirical performance [8, 9], we do not yet understand its fundamental limitations, nor do we have a complete picture of its performance under both signal and hardware nonidealities. We seek to identify these limitations and to characterize the performance in this paper.

To complete the performance profile of the Fourier sampling algorithm, we seek to determine the number of samples, the number of bits per sample, and the total number of bits we need to recover a signal. This is the first goal of our paper. Although theoretical results for both compressive sampling-type algorithms and the Fourier sampling algorithm are expressed in terms of the ℓ_2 error or MSE, $\|x - \widehat{x}\|_2$, between the original signal x and its reconstruction $\widehat{x}$, this measure is not the standard metric for evaluating ADC systems. It is more natural to use signal-to-noise ratio (SNR). It is more useful to express the fundamental bounds as a tradeoff between SNR of recovery and the total bit budget, the sample rate, and the sample precision. This is our second major goal of this paper.

We present a combination of theoretical results along with a suite of empirical results for several realistic signal scenarios, consisting of a signal of interest and various types of interference. We determine several fundamental theoretical bounds for these scenarious. In the remainder of this section, we review the Fourier sampling algorithm and detail the signal scenarios. In the second section, we prove several fundamental theoretical bounds on the Fourier sampling algorithm. We then extend our analysis to the SNR of the signal recovery in Section 3. Finally, we illustrate these relationships empirically in Section 4. We conclude in Section 5 with a discussion and interpretation of the results.

1.1. Review of Fourier sampling algorithm

Let us detail the theoretical guarantees of the algorithm in a discrete setting. Let x be a discrete-time signal of length N. Suppose that x consists of a superposition of m pure tones. Gilbert et al. [6] have developed an algorithm that uses at most $m \operatorname{poly}(1/\epsilon^2 \log N)$ space and time and outputs a representation $\widehat{x}$ which consists of those m pure tones. In the case that x contains some noise, the MSE of the output representation is no more than $(1 + \epsilon)$ times the energy of the noise. The samples are chosen at random (with some structure) and, for each signal, with high probability, the algorithm successfully finds a good compressed representation of the signal. We emphasize that the algorithm returns a short list of significant frequencies, not the entire original signal nor its spectrum. We can view this list as a compressed or denoised version of the signal.

The algorithm iteratively constructs an approximation to the signal by maintaining a list of frequency and coefficient

The authors are supported by NSF DMS 0354600 and DARPA/ONR N66001-06-1-2011. MJS is also supported by NSF DMS 0510203. ACG is an Alfred P. Sloan Fellow. Email:{annacg,martinjs}@umich.edu

pairs, thus representing the approximation implicitly. Each loop of the iteration consists of three steps: (i) identifying candidate dominant frequencies, (ii) estimating the coefficients associated with the candidates, and (iii) updating the previous approximation. The identification step uses a random filterbank to find significant frequencies in the residual signal. It does this via group testing—it finds each bit of dominant frequencies, starting with the least significant bit. Then the estimation step uses a related random filterbank to estimate the coefficients.

1.2. Signal models

We frequently encounter signals in practice which are comprised of a signal of interest (SOI) and some sort of interference. Let us denote our received signal as $x = x_S + x_I$, a superposition of the SOI x_S and the interference x_I. Because compressive sampling is a paradigm in which the measured signals are highly compressive, we assume that the SOI x_S is sparse in the Fourier dictionary Ψ; that is, x_S is well-approximated by a short linear combination of pure tones. With this model, we can classify the interference x_I as one of two signal types:

1. coherent interference: x_I is also sparse in Ψ, or
2. incoherent interference: x_I is not sparse in Ψ.

To simplify our theoretical analysis, we model incoherent interference as Gaussian white noise ν with standard deviation σ that is restricted to the space spanned by the frequency support of x_S. We do not use this restriction for empirical performance evaluation. We assume that we observe the signal x by sampling it and that those samples are quantized (by analog hardware). Once we have quantized those samples, we assume that any postprocessing algorithm has access to full double precision. That is, we do not use internal quantization in postprocessing.

2. THEORETICAL RESULTS

In this section, we prove several fundamental lower bounds on the total number of bits and samples needed to recover a signal.

2.1. Coherent noise

Let us define F, the **Interference Factor**, as the logarithm of the ratio of the interference maximum amplitude to the SOI maximum amplitude.

$$\mathrm{F} = 10 \log_{10} \left(\frac{\|x_I\|_\infty}{\|x_S\|_\infty} \right).$$

While the original Fourier sampling algorithm uses *nonadaptive* samples, we note that there are many situations in which it may be feasible (and desirable) to use *adaptive* samples. The number of rounds of adaptivity is F, the interference factor. This result generalizes to capture the total number of bits needed even in a nonadaptive setting.

Proposition 1. *Let $x = x_S + x_I$ be a superposition of two pure tones, a powerful interferer x_I and a relatively weaker signal of interest x_S. Using* adaptive *samples, the Fourier sampling algorithm can recover x_S with a constant number of bits of precision per sample. The number of rounds of adaptivity is F. More generally, the total number of bits which the Fourier sampling algorithm needs in order to estimate x_S is F, which can be gotten through any tradeoff of number of samples with precision of samples.*

Proof. **Sketch.** Our algorithm proceeds by recovering the interferer, subtracting it off, then recovering the SOI. Our algorithm proceeds in a greedy fashion—it first finds a coarse approximation to the interferer, subtracts off the coarse approximation, then finds a better approximation to the interferer, etc., until the remnants of the interferer is of amplitude comparable to that of the SOI.

We now observe that a constant number of bits (say, 2 bits) is sufficient, if the samples can be made adaptively. Assuming the interferer dominates the SOI, our goal is merely to get a coarse approximation to the interferer that leaves behind an interferer whose power is reduced by, say, the factor $1/4$. It is easy to see that 2 bits suffice for this. We then need to synthesize the current approximation to the interferer with precision depending on the iteration, and growing to F bits. Finally, we subtract the current coarse approximation from the received signal, and iterate. □

We need not change much of the above analysis for multiple tones in either the SOI or the interference (or both). Without loss of generality, we may assume that the SOI contains a single pure tone and the (coherent) interference consists of ℓ tones. The sufficient bit rate is $O(\log \ell - F)$ in this case.

2.2. Incoherent noise

In this subsection, we give theoretical bounds for the number of samples and number of bits of precision required by the Fourier sampling algorithm to recover a single sinusoid amidst Gaussian noise. We note that the analysis is similar for multiple tones and omit it for brevity. Suppose we have a signal $x = x_S + \nu$ consisting of single sinusoid, x_S, at frequency ω and Gaussian noise ν that is white on the space orthogonal to ω. We assume that the sinusoid has amplitude 1; let σ denote the variance of Gaussian noise that inhabits all frequencies *except* ω. The goal is to determine ω and the amplitude.

Proposition 2. *Given a signal of the form $x = x_S + \nu$, the Fourier sampling algorithm can recover ω with $O(\sigma^2)$ samples with* $\log \sigma$ *bits of precision each.*

Proof. **Sketch.** Before we prove the proposition, we find the appropriate lower bound on the number of samples necessary to separate a pure tone from Gaussian noise. Because these types of lower bounds are standard, we simply state the result.

Lemma 3. *We need $\Omega(\sigma^2)$ samples to distinguish a pure sinusoid plus noise from non-constant Gaussian noise alone, even if the frequency of the sinusoid is known.*

We continue with the sketch of the proposition. Our algorithm will make approximately $n = \sigma^2$ measurements, and average them. Our algorithm begins by simulating h, the signal x convolved with an impulse-response that is equal to $1/n$ on its support of size n. Note that, if the samples have infinite precision, then we can arrange that the ω'th coefficient of h has at least 2/3 the energy of h. This is sufficient for the rest of the algorithm to proceed. The question now is how many bits of precision are necessary for this.

We assume that the quantizer rounds up or down, whichever case is *worse*. Regardless of the rounding procedure, we need precision $\pm 1/2$. The average of several such samples is also in the range $\pm 1/2$, which is small compared with the amplitude of the ω'th coefficient. Since sample values can be as large as σ, we need $\log(\sigma)$ bits of precision in our quantization. ☐

3. RECOVERY METRICS

To characterize the fundamental performance of the Fourier sampling algorithm beyond the theoretical lower bounds, we need a metric for the quality of the reconstructed signal. Our strategy for assessing the quality of the output is as follows. If the received signal has coherent interference, we reconstruct both the SOI and the interference; if it has incoherent interference, only the SOI. We eliminate the coefficients for the interference and refer to the remaining coefficients as those of the SOI. We call this oracle reconstruction or reconstruction of the SOI with oracle knowledge of the interferer. Finally, we compare the fidelity of the reconstruction $\widehat{x}_s$ with that of the original SOI x_s. The performance bounds we obtain with this procedure are the same as those one would obtain for an ideal interference rejection system. More formally, the **reconstruction SNR** (SNR_s), assuming oracle reconstruction of the SOI, is

$$\mathrm{SNR}_s = 10\log_{10}\left(\frac{\|x_s\|_2}{\|x_s - \widehat{x}_s\|_2}\right) \quad \text{dB}.$$

This is the metric we report in our empirical evaluation.

Let us express the theoretical error guarantees of the Fourier sampling algorithm in terms of SNR_s. Let x_k denote the top k significant frequencies in the signal x (regardless of the model of the signal). First, we note that if we have any noise μ in the samples themselves (e.g., from quantization), we modify the original result to obtain $\|x - \widehat{x}\|_2 \leq (1+\epsilon)\|x - x_k + \mu\|_2$. Note that $x - x_k$ for incoherent noise is simply ν and for coherent noise $x - x_k = 0$. For coherent noise, the reconstruction SNR is bounded below by

$$\mathrm{SNR}_s \geq 10\log_{10}\left(\frac{\|x_s\|_2}{(1+\epsilon)\|\mu\|_2}\right) = 10\log_{10}\left(\frac{2^B\|x_s\|_2}{(1+\epsilon)\sqrt{M}}\right)$$

where B is the bit rate and M is the total number of samples and where we use a standard approximation for quantization noise. This bound implies that as we increase the bit rate, we increase the SNR but that we will suffer some degradation from the ϵ approximation to the true signal. This is a small additive decrease in the SNR compared to a traditional system. We note that, for a fixed bit budget, we suffer more degradation the more samples M we take.

For incoherent interference ν and oracle reconstruction, the reconstruction SNR is bounded below by

$$\mathrm{SNR}_s \geq 10\log_{10}\left(\frac{\|x_s\|_2}{(1+\epsilon)\|\nu + \mu\|_2}\right).$$

Again, we obtain a small additive decrease in the SNR because of the ϵ approximation. In both cases, we suffer the usual degradation in traditional sampling systems in SNR from measurement noise and from interference.

4. EMPIRICAL ANALYSIS

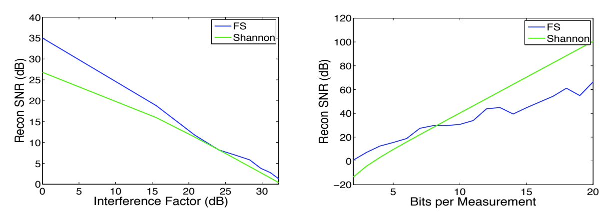

(a) SNR vs. F (B = 6 bits/sample) (b) SNR vs. B (F = 7.8 dB)

Fig. 1. SOI with coherent interference: 3 pure tones each in SOI and interference, signal length 1024.

To verify empirically our theoretical analysis, we build an experimental testbed. We generate two types of signals, a signal of interest with coherent interference and one with incoherent interference. We model the SOI in each case as a superposition of pure tones (with each randomly chosen frequency on the Nyquist grid). We do the same for the coherent interference and scale each randomly chosen tone by the interference factor. For incoherent interference, we use additive white Gaussian noise. To test the fidelity of the reconstruction algorithm, we generate a signal, sample it on the random sample set, quantize each sample, and then use the quantized

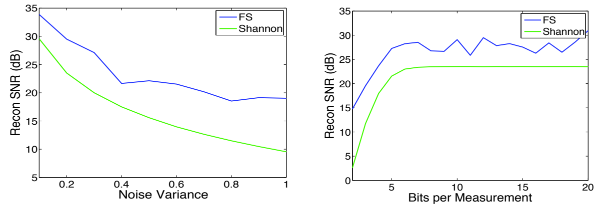

(a) SNR vs. F (B = 12 bits/sample) (b) SNR vs. B (σ = 0.2)

Fig. 2. SOI with incoherent interference: 3 pure tones in SOI, signal length 1024.

samples in the algorithm to reconstruct the SOI. Using oracle knowledge of the interference, we calculate the SNR of the reconstructed signal as a function of the interference factor (for coherent interference) and of the noise variance (for incoherent interference). For both types of interference, we also calculate the reconstruction SNR as a function of the bit rate. We vary the number of tones in the SOI (and in the coherent interferer), the noise variance, the interference factor, and the bit rate for each signal type. We sample the signal at a high enough rate to guarantee successful reconstruction over 95% of the time. In order to validate our approach and to understand its relative performance, we compare the reconstruction SNR to that obtained by a straightforward sampling and reconstruction algorithm we refer to in Figures 1 and 2 as Shannon sampling. In this case, we sample the signal at Nyquist rate, quantize the samples, and compute a DFT. We also use oracle knowledge of the interference to compute the reconstruction SNR.

We observe that the Fourier sampling algorithm's performance is comparable (if not better at times) to the Shannon sampling. It shines when the signal is especially noisy as it returns a denoised version of the SOI, instead of a complete spectrum. Furthermore, because the algorithm uses *fewer* samples than Shannon, it suffers *less* of a degradation in SNR from quantization noise. We note that the performance is similar when the pure tones in the SOI do not fall on a Nyquist grid, but instead, are well-approximated by a few pure tones on the grid. For the sake of space, we do not show these figures. Figures 1 and 2 are example figures for one choice of number of tones or bit rate (depending on the figure) but our complete set of experiments demonstrate comparable performance, as long as we have enough samples. Again, we omit this data for brevity.

5. CONCLUSIONS

We prove several fundamental lower bounds for the Fourier sampling algorithm and demonstrate that its empirical performance is not only consistent with our theoretical analysis but also competitive with traditional sampling systems. These lower bounds not only provide a characterization of the algorithm but also help to optimize the hardware design of novel ADCs and to assess their practical limitations.

Acknowledgments

The authors are grateful to Richard Baraniuk and Dennis Healy for stimulating conversations and research ideas. They also thank Jason Laska for an implementation of the random sampling algorithm and for many discussions about empirical validation.

6. REFERENCES

[1] E. J. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math.*, vol. 59, no. 8, pp. 1208–1223, 2006.

[2] D. L. Donoho, "Compressed Sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

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

[4] A. C. Gilbert, S. Guha, P. Indyk, S. Muthukrishnan, and M. J. Strauss, "Near-optimal sparse Fourier representations via sampling," in *ACM Symposium on Theoretical Computer Science*, 2002.

[5] Martin Vetterli, Pina Marziliano, and Thierry Blu, "Sampling signals with finite rate of innovation," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, June 2002.

[6] A. C. Gilbert, S. Muthukrishnan, and M. J. Strauss, "Improved time bounds for near-optimal sparse Fourier representation via sampling," in *Proc. SPIE Wavelets XI*, San Diego, 2005.

[7] Tamer Ragheb, Sami Kirolos, Jason Laska, Martin Strauss, Anna Gilbert, Richard Baraniuk, and Yehia Massoud, "Implementation models for analog-to-information conversion via random sampling," in *IEEE International Midwest Symposium on Circuits and Systems*, Montreal, Canada, 2007.

[8] Jing Zou, Anna Gilbert, Martin Strauss, and Ingrid Daubechies, "Theoretical and experimental analysis of a randomized algorithm for sparse Fourier transform analysis," *J. Comput. Phys.*, vol. 211, no. 2, pp. 572–595, 2006.

[9] Anna Gilbert, Mark Iwen, and Martin Strauss, "Empirical evaluation of a fast approximate dft algorithm," *J. of Comm. on Math. Sciences*, 2007, to appear.