

Compressive Sensing for Streaming Signals Using the Streaming Greedy Pursuit

Petros Boufounos, M. Salman Asif

TR2010-102 December 2010

Abstract

Compressive Sensing (CS) has recently emerged as significant signal processing framework to acquire and reconstruct sparse signals at rates significantly below the Nyquist rate. However, most of the CS development to-date has focused on finite-length signals and representations. In this paper we present a new CS framework and a greedy reconstruction algorithm, the Streaming Greedy Pursuit (SGP), explicitly designed for streaming applications and signals of unknown length. Our sampling framework is designed to be causal and implementable using existing hardware architectures. Furthermore, our reconstruction algorithm provides explicit computational guarantees, which makes it appropriate for real-time system implementations. Our experimental results on very long signals demonstrate the good performance of the SGP and validate our approach.

MILCOM 2010

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Compressive Sensing for Streaming Signals using the Streaming Greedy Pursuit

Petros T. Boufounos
Mitsubishi Electric Research Laboratories
Cambridge, MA 02139
petrosb@merl.com

M. Salman Asif
School of ECE, Georgia Institute of Technology
Atlanta, GA 30332
sasif@gatech.edu

Abstract—Compressive Sensing (CS) has recently emerged as significant signal processing framework to acquire and reconstruct sparse signals at rates significantly below the Nyquist rate. However, most of the CS development to-date has focused on finite-length signals and representations. In this paper we present a new CS framework and a greedy reconstruction algorithm, the Streaming Greedy Pursuit (SGP), explicitly designed for streaming applications and signals of unknown length. Our sampling framework is designed to be causal and implementable using existing hardware architectures. Furthermore, our reconstruction algorithm provides explicit computational guarantees, which makes it appropriate for real-time system implementations. Our experimental results on very long signals demonstrate the good performance of the SGP and validate our approach.

I. INTRODUCTION

The ever-increasing demand of modern systems for better signal quality and higher data rate has driven the development of new techniques to acquire and process signals of interest. Compressive Sensing (CS) [1]–[3] has emerged as a powerful signal processing framework that enables extremely efficient acquisition by exploiting the sparse structure of most natural and man-made signals. The success of CS has driven both algorithm and hardware development to implement the theory, which places significant emphasis on randomized incoherent measurements at the acquisition stage and increased computation at the reconstruction stage.

Most of the CS results to date focus on finite-dimensional signals. Although the hardware exists which can acquire streaming signals, such as audio, video and radio [4]–[8], but the unstated assumption is that the signal is processed in finite-length blocks. Each block is compressively sampled and reconstructed using one of the known finite-dimensional algorithms. Such an approach can introduce significant blocking artifacts and significant input-output delay. Furthermore, most of the existing finite-dimensional reconstruction algorithms do not provide strict guarantees on the computation time, a critical requirement in streaming real-time systems. Thus, significant buffering of the input or excessive allocation of computation time for reconstruction might be required to ensure that the system satisfies timing requirements for on-line processing.

In this work we examine a CS framework, originally proposed in [9], with features explicitly designed for streaming signals. While [9] focuses on developing the Streaming Greedy Pursuit (SGP) for signals sparse in the frequency domain, and

its application in high-speed video acquisition [10], here we develop the SGP for signals sparse in the time domain.

As with the frequency domain version, the time-domain SGP is a causal algorithm inspired by the Compressive Sampling Matching Pursuit (CoSaMP) [11] and designed to process streaming measurements. The SGP operates at a fixed input rate, computes the reconstruction with a fixed cost per input measurement and outputs the estimated streaming signal at a fixed output rate. The streaming framework we present explicitly avoids processing the signal in blocks and therefore avoids blocking artifacts. Furthermore, in our framework we restrict ourselves to fixed computational cost per measurement which provides explicit trade-offs between input-output delay, reconstruction performance and computation. These features make our framework very well-suited for real-time applications. A significant advantage of our framework is that we can directly accommodate a variety of signal models by modifying CoSaMP iteration, such as [12]–[14].

A formulation for streaming signals poses significant difficulties compared to a fixed-length one. A streaming signal and the corresponding measurements are infinite dimensional vectors. Thus, the usual CS definitions of sparsity and dimensionality reduction need to be reformulated as rates. Streaming time-domain sparsity can potentially be captured in a wide variety of definitions. In this paper we focus on measuring and recovering discrete-time signals. Our choice in defining sparsity is the discrete-time equivalent of signals with finite rate of innovation (FRI) [15]. Although we make connections with the FRI literature, we do not explore them significantly and we do not claim any progress in establishing them.

Our goal in this paper is to describe the measurement framework and present a system architecture using random demodulation [4], [6] or random LTI filtering [5], [16]. We also provide a reconstruction algorithm for signals sparse in the time-domain. While we mention and briefly discuss some of the theoretical difficulties and implications of our endeavor, we defer detailed analysis to future publications. Instead, we present experimental results validating our approach.

To establish the background and notation we briefly review signal acquisition and CS in Sec. II. Section III formulates the streaming models and constructs, possible implementations, and connections with the FRI literature. Section IV describes the SGP algorithm. We present experimental results in Sec. V.

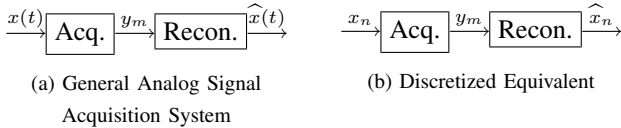


Fig. 1. Basic Streaming Acquisition System

II. BACKGROUND

A. Signal Acquisition

To introduce CS to streaming applications we consider the general streaming system depicted in Fig. 1(a). The signal $x(t)$ is acquired using an acquisition system (“Acq.” in the figure) at an average rate of M measurements per time unit and subsequently reconstructed (“Recon.” in the figure) using those measurements. In most classical signal acquisition systems the acquisition component is an analog-to-digital converter (ADC) which obtains linear measurements using a low-pass anti-aliasing filter followed by uniform time sampling and quantization. The reconstruction component is the linear bandlimited interpolation of the measurements y_m .

An equivalent discrete-time system is shown in Figure 1(b). The discrete-time signal x_n may represent $x(t)$ under a suitable representation or it might be an inherently discrete-time signal. For example, in classical bandlimited sampling and interpolation $x_n = x(nT)$, where T is the signal Nyquist period. In this case the acquisition and reconstruction components are the identity (i.e., $m = n$ and $\hat{x}_n = y_n = x_n$) and the implied reconstructed signal $\hat{x}(t)$ is the bandlimited interpolation of \hat{x}_n . More general cases with less trivial acquisition and reconstruction components are discussed in [17]. In this paper we only consider discrete-time systems. We assume that if the signal of interest is continuous-time, a discretization exists and describes the acquisition system to sufficient accuracy.

Acquiring a general signal x_n requires an acquisition rate of M measurements y_m per time period that is greater than or equal to the input rate of N signal coefficients per time period. Otherwise, the system is not invertible for all input signals and information is lost. However, with additional information on the signal structure it is possible to acquire a signal at a lower measurement rate $M \ll N$ and still reconstruct it.

B. Compressive Sensing

Compressive Sensing [1]–[3] demonstrates that a signal sparse or compressible can be efficiently sampled and reconstructed using very few linear measurements. The signal of interest, $\mathbf{x} \in \mathbb{R}^N$, is measured using the system

$$\mathbf{y} = \mathbf{A}\mathbf{x}, \quad (1)$$

where \mathbf{y} denotes the measurement vector and \mathbf{A} an $M \times N$ measurement matrix with $M \ll N$. The signal \mathbf{x} is assumed K -sparse, i.e., it contains only K non-zero coefficients.

The measurement matrix \mathbf{A} satisfies the *Restricted Isometry Property* (RIP) of order $2K$ if there exists a constant $\delta_{2K} < 1$ such that

$$(1 - \delta_{2K})\|\mathbf{x}\|_2^2 \leq \|\mathbf{A}\mathbf{x}\|_2^2 \leq (1 + \delta_{2K})\|\mathbf{x}\|_2^2, \quad (2)$$

for all $2K$ -sparse \mathbf{x} . If the RIP constant δ_{2K} is sufficiently small, then the signal can be exactly reconstructed using the convex optimization [3]

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{y} = \mathbf{A}\mathbf{x}. \quad (3)$$

or a greedy algorithm such as CoSaMP [11]. Furthermore, a small RIP constant provides robustness guarantees for recovery in the presence of measurement noise and for sampling signals that are not exactly sparse but can be well-approximated by a sparse signal [3], [11].

Significant work has been performed on hardware implementations of CS on streaming signals such as [4], [6], [8]. Such efforts focus on the hardware architectures that enable random projections. However, the unstated assumption on the signal processing side is that the signal is processed in finite-length blocks. Each block is compressively sampled and reconstructed using one of the known finite-dimensional algorithms.

III. STREAMING COMPRESSIVE ACQUISITION

A. Measurement System Model

We measure a streaming, infinite-dimensional signal x_n using a time varying linear system $a_{m,n}$ as

$$y_m = \sum_n x_n a_{m,n} = \langle \mathbf{x}, \mathbf{a}_m \rangle, \quad (4)$$

where y_m is the sequence of measurements, and \mathbf{a}_m is sequence of infinite-dimensional measurement vectors. The system has an input rate of N coefficients per unit time for x_n and an output (measurement) rate of $M = N/R$ measurements per unit time, where R denotes the downsampling rate. For notational simplicity, in the remainder of this paper we assume that R is an integer¹.

We consider causality an essential aspect of a streaming formulation, and, therefore, we assume the measurement system is causal. We further assume it has a finite response of maximum length LR . Finite response length is not essential but a convenience in our development. Infinite response measurement systems can be incorporated to our framework, but it is not an extension we consider in this paper.

Under those assumptions, the system $a_{m,n}$ has support in

$$n = (m - L)R + 1, \dots, mR, \quad (5)$$

$$\Leftrightarrow m = \lceil n/R \rceil, \dots, \lceil n/R \rceil + L - 1, \quad (6)$$

and is zero outside. Figure 2(a) summarizes the relevant parameters of the measurement system $a_{m,n}$. The shaded area demonstrates the support of the system relative to the axes. The system coefficients are zero in the unshaded area. The causality requirement imposes a lower triangular structure to the system. The finite response length assumption further constrains the nonzero coefficients to lie within a band of height L below the main diagonal.

¹In several instances in this paper we assume rates and lengths are integer multiples of each other. Non-integer rates are straightforward to accommodate with appropriate use of rounding operators. However, the notation becomes cumbersome without adding any insight to the fundamental concepts.

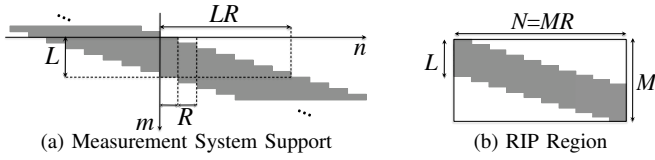


Fig. 2. Overview of the acquisition system and its parameters

B. Sparsity Model and RIP

We assume the signal is sparse in the time-domain and impose a streaming sparsity model. Specifically, we assume there are at most K nonzero coefficients in any sequence of N signal coefficients, where K and N are determined at design time. We call $S = K/N$ the sparsity rate. Using this sparsity model, we formulate the RIP for the measurement system by considering time-windows of length N in the signal domain, as shown in Fig. 2(b). In any such window we define the RIP as in (2).

This model is fairly flexible since it allows the signal to be dense in a small region—as long as the preceding and the subsequent coefficients have sufficient sparsity to meet the overall sparsity requirement for a length- N window—or to be uniformly sparse over time.

In a streaming system establishing the RIP over a fixed window is only part of the system design considerations. In a randomized infinite-dimensional setting, even though the probability of RIP failure is small, the RIP might still fail for some window. Often the streaming signal sparsity might temporarily exceed the original assumptions at the design stage of the system. The reconstruction algorithm we describe in the next section is robust to such failures and sufficiently stable to ensure that local reconstruction errors do not propagate to the subsequent reconstruction. Furthermore, the top and bottom L measurements in each window are also affected by signal components outside that window. A properly designed algorithm should take those measurements correctly into account.

C. Implementation

The measurement model we explore is fairly general. There is no requirement that all the coefficients in the support are non-zero. Therefore, with appropriate choice of parameters the model encompasses several sampling methods proposed for CS such as random sampling, periodic non-uniform sampling, random demodulation, and random filtering, and block-diagonal, among others (for examples, see [4]–[8]).

Of course, not all of those methods are appropriate for signals sparse in the time domain; a simple random demodulator or a random sampler, for example, can certainly fail. The RIP property, as described above, is a fairly good predictor of how appropriate is the system for time-domain sparse signals. Intuitively, we need the columns of the measurement system to measure and significantly “mix” many time domain samples, i.e., the rows of the system to be incoherent with time domain impulses and sufficient in number. This requires that most of the support of the system be populated with non-zero elements;

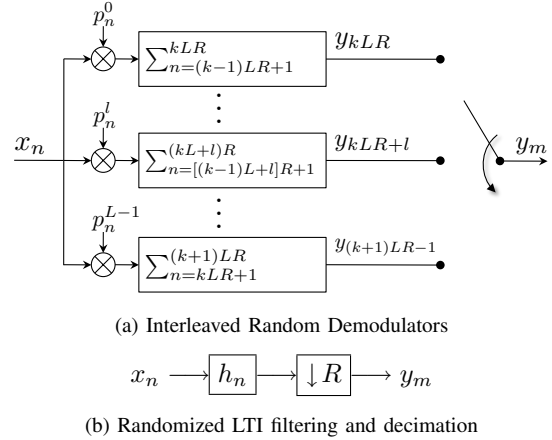


Fig. 3. Possible implementation of the streaming measurement system

we describe two different implementation that achieve that goal.

The first, shown in Fig. 3(a), uses a bank of L interleaved random demodulators operating at rate of one measurement per LR signal samples each, and sampled sequentially in a round-robin fashion at a rate of a measurement per R signal samples. The advantage of this approach is that it correspond to a very efficient analog implementation, as described in [4]. The p_n^l —where $l = 0, \dots, L - 1$ denotes the demodulator index—is referred to as the chipping sequence of the demodulator and is typically chosen as a random sequence taking values of ± 1 . Thus the multiplication is easily implementable using a simple switch reversing the polarity of the input signal. The more expensive component of an analog acquisition system, the quantizer, only needs to be implemented once for all L demodulators, after the round-robin switch.

An alternative approach, shown in Fig. 3(b), is to build a length LR LTI filter with random impulse response h_n , followed by decimation by R [5], [16]. Randomization of the impulse response ensures incoherence with its shifts but makes analog implementation of the filter difficult. Thus it is more appropriate for naturally discrete-time systems. Thanks to the decimation following the LTI filter the system can be implemented efficiently using standard polyphase filterbanks [18].

D. Connections to signals with Finite Rate of Innovation

Our model is the discrete-time equivalent to signals with finite rate of innovation (FRI) (see [15], [19], [20], and references within). The FRI model assumes that the continuous-time signal of interest is composed of a sequence of at most K time-domain pulses within any finite time interval T . The intent is to sample and recover the timing and amplitude of each of the K pulses—equivalent to the support and amplitude in our discrete-time model. The rate of innovation is $2K$ unknowns per time interval T , and the signal can in principle be recovered using $2K$ samples of signal using uniform sampling with an appropriate filtering kernel.

The original work [15] assumes a periodic sequence of Dirac impulses and provides an elegant but numerically

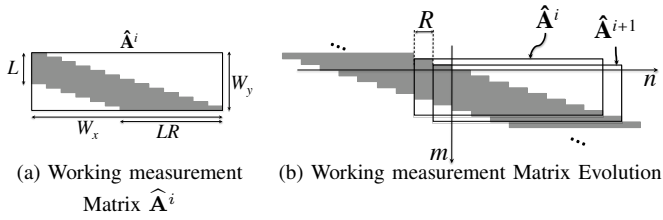


Fig. 4. The working measurement matrix and its correspondence with the measurement system response.

unstable solution. Followup work focused on stability and on generalizing beyond periodic signals. Most infinite-length approaches treat the problem as a sequence of finite-length or periodic problems, either suffering from significant instability or requiring boundary assumptions such as a bursty signal model. To our knowledge the FRI model is not easily amenable to theoretical analysis on the amplitude and the timing errors under noise or quantization, equivalent to the signal and support recovery analysis, respectively, in the CS literature.

While our current work does not directly address most of the issues, it is a first step in achieving streaming reconstruction for FRI signals. Specifically, using a bank of interleaved random demodulators operating at rate N pulses per time interval T , as described in Sec. III-C above, we can sample and recover arbitrary streams of FRI impulses and provide guarantees on the recovery error under noise. Furthermore, any impulse correctly recovered by such a system will be due to an FRI impulse within a timing interval of T/N captured by the demodulators; thus we can leverage standard CS support recovery analysis. The drawback of our approach is that exact timing recovery is not possible even under the noiseless case. Furthermore a slight oversampling by $O(\log N/K)$ over the rate of innovation $2K$ is required to guarantee the RIP of the measurement system. We defer the details of such an approach to a future publication.

IV. STREAMING GREEDY PURSUIT

To simplify the presentation of the Streaming Greedy Pursuit (SGP) algorithm, we first discuss the fundamental algorithm iteration step in Sec. IV-A, followed by more detailed discussion of the signal estimation part in Sec. IV-B.

A. Streaming Iteration

The SGP operates by estimating the signal on a sliding window of W_x signal coefficients using W_y measurements. After each iteration R new signal coefficients (to be estimated) are included in the coefficient window and the oldest R are considered final estimates, removed from the window and committed to the output. Similarly, a new measurement is incorporated at the end of the measurement window and the oldest measurement is removed from the beginning. The measurements are also updated to remove the influence of the committed coefficients such that the updated measurements can be explained only from the uncommitted coefficients. This approach in handling the streaming measurements is the main difference between the iterations of SGP and CoSaMP.

Algorithm 1 Streaming Iteration for SGP

- 1: **loop**
 - 2: *Increase iteration count:* $i \leftarrow i + 1$
 - 3: *Refine working estimate:*
 $\tilde{\mathbf{x}}^i \leftarrow \text{Refine}(\tilde{\mathbf{x}}^{i-1}, \hat{\mathbf{y}}^{i-1}, \hat{\mathbf{A}}^{i-1})$,
 where $\text{Refine}(\cdot)$ is described in Algorithm 2.
 - 4: *Commit coefficients estimate:*
 $\hat{x}_{Ri+j} = \tilde{x}_j^i, j = 1, \dots, R$,
 where \hat{x}_n is the streaming signal estimate at the output.
 - 5: *Slide working coefficients window:*
 $\tilde{\mathbf{x}}^i \leftarrow \left[\frac{\tilde{\mathbf{x}}_{\{(R+1), \dots, W_x\}}^i}{\mathbf{0}_R} \right]$,
 where $\mathbf{0}_R$ is the zero vector.
 - 6: *Remove committed coefficients from working measurements:*
 $\tilde{\mathbf{y}}^i \leftarrow \hat{\mathbf{y}}^{i-1} - \hat{\mathbf{A}}^{i-1} \Big|_{\{1, \dots, R\}} \tilde{\mathbf{x}}_{\{1, \dots, R\}}^i$
 - 7: *Slide working measurement window:*
 $\hat{\mathbf{y}}^i \leftarrow \left[\frac{\tilde{\mathbf{y}}_{\{2, \dots, W_y\}}^i}{y_{i+W_y}} \right]$
 - 8: *Slide working measurement matrix:*
 $\hat{\mathbf{A}}^i \leftarrow \left[\frac{\hat{\mathbf{A}}_{\{2, \dots, W_y\}, \{R+1, \dots, W_x\}}^{i-1} \Big| \mathbf{0}_{(W_x-R)}}{a_{(i+W_y), \{Ri+1, Ri+W_x\}}} \right]$,
 which sets $(\hat{\mathbf{A}}^i)_{k,l} = a_{(k+i), (Ri+l)}$ for $k = 1, \dots, W_y$
 and $l = 1, \dots, W_x$.
 - 9: **end loop**
-

In every i^{th} iteration, the SGP maintains a working signal estimate of length W_x , denoted $\tilde{\mathbf{x}}^i$, a working measurement matrix of dimension $W_y \times W_x$, denoted $\hat{\mathbf{A}}^i$, and a measurement vector of length W_y , denoted $\hat{\mathbf{y}}^i$. The iteration is presented in Algorithm 1, where we use the $(\cdot)_m$ notation to denote the m^{th} element of a vector and $(\cdot)_{\{m_1, \dots, m_2\}}$ to denote a range of elements. Non-boldface characters denote the underlying streaming signals, as described in Sec. II-A and III.

The working matrix $\hat{\mathbf{A}}^i$ and its evolution after each iteration with respect to the system $a_{m,n}$ is shown in Fig. 4(a) and (b). From the figure it becomes obvious that the true measurements y_m do not reflect the working measurements $(\mathbf{y})_{\{1, \dots, L\}}$ corresponding to the first L rows of the working matrix $\hat{\mathbf{A}}^i$. The working measurements, reflecting only the uncommitted part of the signal, are updated accordingly at Step 6 in Algorithm 1.

Any error in the committed coefficients is propagated back to the working measurements and increases the error in subsequent estimates. Such errors can be due to initial conditions, RIP failures of the measurement system, or sparsity model mismatches for the streaming signal. Still, it is possible to demonstrate that the SGP algorithm is stable and quickly recovers from such errors.

B. Signal Estimation and Refinement

The signal estimation and refinement step 3 of the SGP, summarized in Algorithm 2, is similar to one CoSaMP iteration. The algorithm computes the unexplained residual in the current working measurement window and forms a signal

proxy to explain it. In that proxy it identifies T largest coefficients and merges their support set with the support of the current signal estimate. The algorithm then re-estimates the signal by solving a least squares (LS) problem over the merged support. Finally, the SGP truncates the solution to its K largest coefficients.

Although the refinement steps are similar to one CoSaMP iteration, they differ from CoSaMP in two significant ways:

(i) In order to keep the computational cost of every iteration fixed and as small as possible, the number of support elements T selected from proxy at every iteration is chosen to be much smaller than the sparsity K of the signal. This results in a small computational cost for solving LS problem in step 4 at every iteration, where $2T$ rank-one updates can be used to update the LS solution from the previous iteration [21]. Alternatively, instead of using direct method to update LS solution in step 4, one can compute an approximate solution using a fixed number of iterations of some suitable iterative solver such as conjugate gradient (CG) [11], [22].

(ii) The proxy and the LS problem are computed using a weight \mathbf{w} on the measurements, with less weight on the most recent L measurements. This discourages adding the last LR coefficients $(\tilde{\mathbf{x}})_{\{W_x-LR+1, \dots, W_x\}}$ in the support; not all measurements affected by those coefficients are available yet, making the proxy and their estimates less reliable. These correspond to the lower right $L \times LR$ triangular portion of $\hat{\mathbf{A}}^i$, which makes $\hat{\mathbf{A}}^i$ in Fig. 4(a) differ from Fig. 2(b).

The choice of L is an important trade-off in the algorithm. With a fixed working window W_x , as L increases, the probability of satisfying the RIP for denser signals increases and the estimation improves. However, a larger portion of the working matrix $\hat{\mathbf{A}}^i$ is occupied by the lower right triangular part, making the estimation less reliable. This trade-off is further explored in the experimental section. A longer working window W_x reduces this effect but increases the computational cost and the input-output delay of the algorithm.

For the same reason, SGP does not guarantee exact reconstruction. The performance is a function of the computational cost. The longer the window length W_x , the more iterations are used to refine the estimate of each signal coefficient and the better the reconstruction performance. In addition, the guarantees provided by the CoSaMP iteration ensure the stability of the error propagation for the parts of the signal committed to the output. Thus, even if part of the signal is misidentified for any reason and then committed, the feedback in the algorithm does not become unstable and quickly reverts to the correct estimate.

V. EXPERIMENTAL RESULTS

In this section we present some experimental results for the recovery performance of the SGP.

In the first set of experiments we study the recovery performance of SGP for time-sparse signals. We explore different values of signal sparsity rate (S) and downsampling rate (R) using measurements from both the random demodulator and the random filter. We perform the experiments for $R = 2$,

Algorithm 2 Signal Refinement function Refine(\cdot)

1: *Calculate Residual:*

$$\mathbf{r}^i \leftarrow \hat{\mathbf{y}}^{i-1} - \hat{\mathbf{A}}^{i-1} \hat{\mathbf{x}}^{i-1},$$

2: *Compute Signal Proxy:*

$$\mathbf{p}^i \leftarrow (\hat{\mathbf{A}}^{i-1})^T \mathbf{W}^2 \mathbf{r}^i,$$

where $(\cdot)^T$ denotes the matrix transpose, and \mathbf{W} is a diagonal matrix with a weight vector \mathbf{w} in its diagonal.

3: *Identify and Merge Support:*

$$\Omega \leftarrow \text{supp}(\tilde{\mathbf{x}}^{i-1}) \cup \text{supp}(\mathbf{p}^i|_T),$$

where $\text{supp}(\cdot)$ denotes the support index set of a sparse vector, and $\mathbf{p}|_T$ denotes truncation of the vector \mathbf{p} to its T largest in magnitude coefficients

4: *Estimate Signal over Support:*

$$\mathbf{b}^i \leftarrow (\mathbf{W} \hat{\mathbf{A}}^{i-1}|_{\Omega})^{\dagger} \mathbf{W} \tilde{\mathbf{y}}^{i-1}$$

5: *Truncate Estimate:*

$$\tilde{\mathbf{x}}^i \leftarrow \mathbf{b}^i|_K$$

4 and 8, with signal window length $W_x = 400, 800$ and 1200 respectively, and $W_y = W_x/R$. To generate a streaming signal x with average sparsity rate S , we select the nonzero locations independently with probability S by flipping a biased coin at every location. The signal values at nonzero locations are independently chosen to be ± 1 with equal probability. Generating the support using a coin flip does not guarantee a fixed maximum sparsity assumed in Sec. III-B. This allows us to test the stability of the algorithm under model mismatches.

The streaming signal is measured using a random demodulator or random filter, as described in sec. III-C). The elements in the chipping sequences p_n^l for the random demodulator and the impulse response h_n for the random filter are chosen from ± 1 with equal probability. Gaussian noise $N(0, \sigma^2)$ at 35db SNR is added to the measurements. In this experiment we choose $L = W_y/3$; we investigate this choice in the second set of experiments, below. The signal estimate \hat{x} is reconstructed from streaming noisy measurements using the SGP with $K = 1.2SW_x$ and $T = K/2$, by running 5 CG iterations for the LS update in step 4 of Algorithm 2. We measure the recovery performance in terms of the probability of exact recovery (PER) and the signal to error ratio (SER) of the reconstructed signal. The PER is the ratio of number of coefficients which are misclassified as zero or nonzero in \hat{x} , within a small tolerance (e.g., 5σ) to their original value, over the total number of coefficients in x . The SER is equal to

$$\text{SER} = 20 \log_{10} \frac{\|x\|_2}{\|x - \hat{x}\|_2}.$$

Figure 5 plots the results, where each point represents the average performance over 20 trials on 20000 samples long signal. Plots (a) and (c) on the left correspond to the random demodulator experiments and plots (b) and (d) on the right correspond to random filtering experiments. The PER and SER curves for both systems show gradual performance degradation as the S increases for fixed value of R , and recovery is almost exact when the average sparsity rate is $S < 1/10R$. This

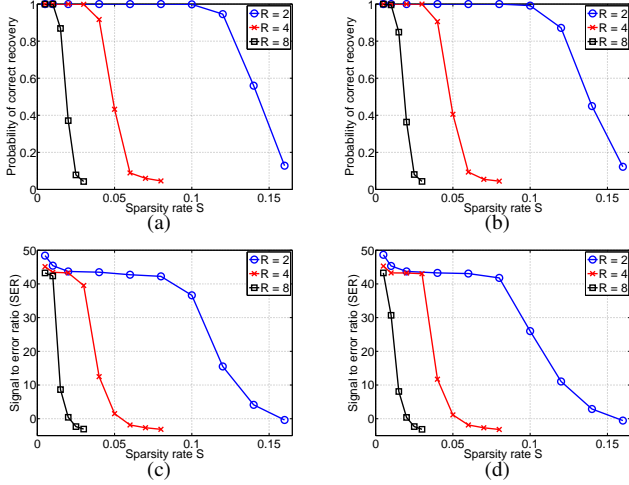


Fig. 5. Performance of the SGP for different downsampling rates (R) against varying sparsity rate (S) of the streaming signal. Left: Random demodulator. Right: Random filter.

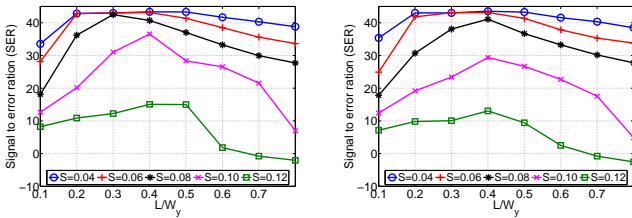


Fig. 6. Performance of the SGP for different sparsity rates (S) and response lengths (L). Left: Random demodulator. Right: Random filter.

behavior is similar to the usual CS rule of thumb, dictating a small constant oversampling factor compared to the sparsity rate.

In the second set of experiments we study the effects of response length L on the system performance, assuming a fixed window length W_x and downsampling rate R . The simulation setup is same as above, where the streaming signal x at sparsity rate S is measured with random demodulator or random filter at downsampling rate R . The SER results for recovery with $W_x = 400$, $R = 2$, and $W_y = W_x/R$ using different values of L and S are presented in Fig. 6, where the left plot corresponds to the random demodulator and the right plot corresponds to the random filter based measurements. The figure demonstrates the trade-off in designing L , where performance degrades significantly for very large and very small values of L . Empirically, the best performance is achieved when $L \approx W_y/3$. At that level of L we obtain the best trade-off between the width of the RIP block and the lower right triangular portion of the working measurement matrix.

While we have only presented the tip of the iceberg in this rich topic, we have demonstrated that CS principles can be effectively used to causally sample and reconstruct streaming time sparse signals. We believe that the SGP is a significant tool for the designer of real-time systems and we have already used it in some of our designs [9], [10].

ACKNOWLEDGMENT

We would like to thank Justin Romberg for fruitful discussion and suggestions in completing this work.

REFERENCES

- [1] E.J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [2] D.L. Donoho, “Compressed sensing,” *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, April 2006.
- [3] E. Candès, “Compressive sampling,” *Proceedings of the International Congress of Mathematicians, Madrid, Spain*, vol. 3, pp. 1433–1452, 2006.
- [4] J.N. Laska, S. Kirolos, M.F. Duarte, T.S. Ragheb, R.G. Baraniuk, and Y. Massoud, “Theory and implementation of an analog-to-information converter using random demodulation,” in *Circuits and Systems, IEEE International Symposium on*, May 2007, pp. 1959–1962.
- [5] J. Romberg, “Compressive sensing by random convolution,” *SIAM J. Imaging Science*, 2008.
- [6] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, “Beyond Nyquist: Efficient sampling of sparse, bandlimited signals,” *IEEE Trans. on IT*, vol. 56, no. 1, pp. 520–544, Jan. 2010.
- [7] M. Mishali, Y.C. Eldar, and A. Elron, “Xampling—Part I: Practice,” *Arxiv preprint arXiv:0911.0519*, 2009.
- [8] A. Veeraraghavan, D. Reddy, and R. Raskar, “Coded strobing photography for high-speed periodic events,” *IEEE Trans. Pattern Anal. Mach. Intell.*, to appear.
- [9] P. T. Boufounos and M. S. Asif, “Compressive sampling for streaming signals with sparse frequency content,” *44th Annual Conference on Information Sciences and Systems*, March 2010.
- [10] M. S. Asif, D. Reddy, P. T. Boufounos, and A. Veeraraghavan, “Streaming compressive sensing for high-speed periodic videos,” (submitted).
- [11] D. Needell and J.A. Tropp, “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples,” *Appl. Comp. Harmonic Anal.*, June 2008.
- [12] R.G. Baraniuk, V. Cevher, M.F. Duarte, and C. Hegde, “Model-based compressive sensing,” *IEEE Trans. Info. Theory*, vol. 56, no. 4, pp. 1982–2001, April 2010.
- [13] P. Boufounos, “Greedy sparse signal reconstruction from sign measurements,” in *Asilomar Conference On Signals Systems and Computers*, Pacific Grove, CA, Nov. 2009.
- [14] C. Hegde, M.F. Duarte, and V. Cevher, “Compressive sensing recovery of spike trains using a structured sparsity model,” *Signal Processing with Adaptive Sparse Structured Representations (SPARS)*, pp. 13–16, 2009.
- [15] M. Vetterli, P. Marziliano, and T. Blu, “Sampling signals with finite rate of innovation,” *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [16] J. A. Tropp, M. B. Wakin, M. F. Duarte, D. Baron, and R. G. Baraniuk, “Random filters for compressive sampling and reconstruction,” *IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 3, 2006.
- [17] M. Unser and A. Aldroubi, “A general sampling theory for nonideal acquisition devices,” *Signal Processing, IEEE Transactions on*, vol. 42, no. 11, pp. 2915–2925, Nov 1994.
- [18] A.V. Oppenheim, R.W. Schaffer, and J.R. Buck, *Discrete-Time Signal Processing*. Prentice Hall Englewood Cliffs, NJ, 2nd edition, 1999.
- [19] T. Blu, P.-L. Dragotti, M. Vetterli, P. Marziliano, and L. Coulot, “Sparse Sampling of Signal Innovations [Theory, algorithms, and performance bounds],” *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 31–40, March 2008.
- [20] R. Tur, Y. C. Eldar, and Z. Friedman, “Low rate sampling of pulse streams with application to ultrasound imaging,” *CoRR*, vol. abs/1003.2822, 2010.
- [21] Å. Björck, *Numerical Methods for Least Squares Problems*, Society for Industrial Mathematics, 1996.
- [22] G.H. Golub and C.F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, 1996.