

A Nonuniform Sampler for Wideband Spectrally-Sparse Environments

Michael Wakin, *Member, IEEE*, Stephen Becker, *Member, IEEE*, Eric Nakamura, Michael Grant, *Member, IEEE*, Emilio Sovero, *Senior Member, IEEE*, Daniel Ching, Juhwan Yoo, *Member, IEEE*, Justin Romberg, *Senior Member, IEEE*, Azita Emami-Neyestanak, *Member, IEEE*, and Emmanuel Candès, *Member, IEEE*

Abstract—We present a wide bandwidth, compressed sensing based nonuniform sampling (NUS) system with a custom sample-and-hold chip designed to take advantage of a low average sampling rate. By sampling signals nonuniformly, the average sample rate can be more than a magnitude lower than the Nyquist rate, provided that these signals have a relatively low information content as measured by the sparsity of their spectrum. The hardware design combines a wideband Indium-Phosphide heterojunction bipolar transistor sample-and-hold with a commercial off-the-shelf analog-to-digital converter to digitize an 800 MHz to 2 GHz band (having 100 MHz of noncontiguous spectral content) at an average sample rate of 236 Ms/s. Signal reconstruction is performed via a nonlinear compressed sensing algorithm, and the challenges of developing an efficient implementation are discussed. The NUS system is a general purpose digital receiver. As an example of its real signal capabilities, measured bit-error-rate data for a GSM channel is presented, and comparisons to a conventional wideband 4.4 Gs/s ADC are made.

Index Terms—Compressed sensing, indium-phosphide heterojunction bipolar transistor (HBT), nonuniform sampler, sample-and-hold, wideband analog-to-digital converter (ADC).

I. INTRODUCTION

IN such far-ranging fields as radio, telephony, radar, image, audio and seismic acquisition, most analysis techniques follow the same pattern: 1) digitize an analog signal, 2) perform

DSP, and, optionally, 3) convert back to the analog domain. The common piece of hardware in this chain is the analog-to-digital converter (ADC). The current trend in systems is wider bandwidths and larger dynamic ranges, and designing a single ADC to meet both of these requirements simultaneously is difficult. To get around this, systems typically use time-interleaved ADCs or channelize the band and digitize each channel separately. However, these approaches do nothing to reduce the output data rate and can require prohibitively high power.

A. A New Paradigm

The *effective instantaneous bandwidth* (EIBW) of an ADC is the total bandwidth of the spectrum that can be unambiguously recovered. Although it is generally desirable to design receivers with high EIBW—say, for applications involving cognitive radio or communications intelligence—it may also often be the case that, at any given time instant, much of the spectrum within this bandwidth is unoccupied. One can define the *information bandwidth* of such signals to be the actual amount of occupied spectrum. In this paper, we present a receiver design intended for signals with high EIBW but low information bandwidth. We do so by adopting concepts from the field of compressed sensing (CS).

The theory of CS [7], [15] suggests that randomized low-rate sampling may provide an efficient alternative to high-rate uniform sampling. For a survey of the modern CS literature, the reader is referred to [8].

To put CS on a concrete footing, we give an explicit (but for the moment, discrete-time) example. Let x be a length- N signal, and suppose the discrete Fourier transform (DFT) of x , denoted X , is K -sparse, meaning that it has only $K \ll N$ nonzero entries.

Now, collect only a *subset* Ω of all the entries of x . Suppose the sample locations Ω are chosen uniformly at random, and let M be the size of Ω . Because $M < N$,¹ it is generally not possible to recover x using a linear method. The remarkable fact of CS is that if M is merely proportional to $K \log N$, then with very high probability (which can be made precise [7]), it is possible to *exactly* recovery x by solving the linear program

$$\min_{x'} \|X'\|_1 \quad \text{subject to} \quad \forall k \in \Omega, x'(k) = x(k).$$

Here, $\|X\|_1 = \sum_{k=1}^N |X(k)|$. There are related approaches, such as greedy methods, that offer similar guarantees; see [27] for a survey.

¹In our implementation, M is approximately $19\times$ smaller than N .

Manuscript received February 28, 2012; revised June 29, 2012; accepted July 30, 2012. Date of publication September 28, 2012; date of current version December 05, 2012. This work was supported by the DARPA/MTO Analog-to-Information Receiver Development Program under Air Force Research Laboratory (AFRL) Contract FA8650-08-C-7853. The views expressed in this paper are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. This is in accordance with DoDI 5230.29, January 8, 2009. Statement “A” (Approved for Public Release, Distribution Unlimited) [DISTAR case #18837]. This paper was recommended by Guest Editor G. Setti.

M. Wakin is with the Department of Electrical Engineering and Computer Science, Colorado School of Mines, Golden, CO 80401 USA (e-mail: mwakin@mines.edu).

S. Becker is with the Applied and Computational Mathematics Option, California Institute of Technology, Pasadena, CA 91125 USA, and also with the Laboratoire Jacques-Louis Lions, Paris 6 University, 75005 Paris, France.

E. Nakamura, E. Sovero, and D. Ching are with the Northrop Grumman Corporation, Redondo Beach, CA 90278 USA (e-mail: eric.nakamura@ngc.com).

M. Grant is with the Applied and Computational Mathematics Option, California Institute of Technology, Pasadena, CA 91125 USA.

J. Romberg is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA.

A. Emami-Neyestanak and J. Yoo are with the Electrical Engineering Option, California Institute of Technology, Pasadena, CA 91125 USA.

E. Candès is with the Departments of Mathematics and Statistics, Stanford University, Stanford, CA 94305 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JETCAS.2012.2214635

This result itself has limited application to signal processing since (1) it is unlikely that a digital signal has an exactly sparse DFT, and (2) the model does not account for noise. Fortunately, there are robust versions of the above statement, which allow signals to be only *approximately sparse*, and which allow noise [6]. In this case, exact recovery is not possible, but the recovered signal agrees with the true signal up to the noise level.

This finite dimensional model does not fully cover the continuous case since an analog signal, unless it is bandlimited *and* periodic, must be treated in an infinite dimensional setting. The infinite dimensional setting may be attacked directly, and there is recent theory [1], [14] that connects the finite and infinite dimensional problems. On a practical side, some nonperiodic (also known as multi-coset) sampling results for multi-band signals [19] have recently been extended to multi-band signals when the band locations are unknown [20]. The approach in [20] and [21] is a hybrid finite-infinite approach that solves a finite dimensional problem to determine the band locations and then processes the samples directly in analog.

The approach taken in this paper deals with the infinite dimensional problem indirectly. Through extensive numerical simulation, and by using standard signal processing techniques such as windowing, it is shown that the error incurred by using a large but finite number of samples is insignificant compared to circuit nonidealities. Numerical simulations are required regardless since CS theories rely on possibly conservative constants and also on signal-dependent parameters, such as the sparsity of the signal.

B. Approach

1) *Overview*: The CS example suggests that signals with high EIBW but low information bandwidth can be efficiently captured using nonuniform samples. Our implementation is such a nonuniform sampling (NUS) approach, which we describe here and treat in more detail in Section II. The key ideas are to leverage the high resolution that can be achieved with lower-rate ADCs and to exploit the fact that the electromagnetic spectrum in our bandwidth is typically not full.

There are two sets of signal restrictions for the NUS. The first is a familiar restriction requiring the EIBW to be less than half the (equivalent) Nyquist sampling rate. The second restriction is an algorithmic one: CS theory dictates that the input signal should have spectral sparsity in order to achieve accurate reconstruction. Roughly speaking, in our implementation the information bandwidth may be up to 10% of the EIBW, although this figure is dependent on the rate at which we acquire NUS samples, and by increasing this rate one could accommodate a larger information bandwidth.

The idea behind the NUS is explained in Fig. 1. For our setup, the maximum EIBW is 2.2 GHz because of an underlying “Nyquist rate” clock with a frequency of $f_s = 4.4$ GHz. For the sake of explanation, assume there is a Nyquist rate ADC which samples the input signal—the actual implementation does not use a Nyquist rate ADC, since the point of the NUS is to avoid a high-rate ADC. A pseudo-random bit sequence (PRBS), generated off-chip, controls which of these samples are collected and which samples are ignored. Of every 8192 Nyquist-rate samples, only 440 are collected. Note that our method of “on grid”




Fig. 1. Conceptually, the NUS takes Nyquist-rate samples of the input signal and then randomly discards most of the samples. The implemented version uses a clock rate of 4.4 GHz and effectively keeps only one of every 19 samples (on average) for a mean output sample rate of 236 MHz.

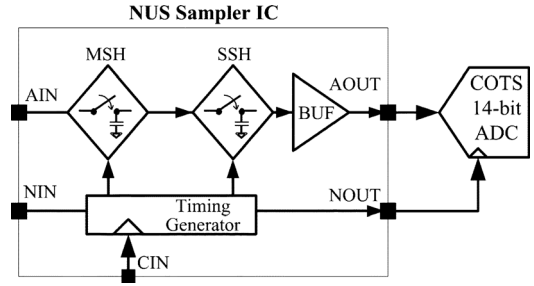


Fig. 2. Simplified block diagram of NUS receiver. The NUS sampler IC (left block) was implemented with the NGAS InP HBT process.

nonuniform sampling is very different from allowing arbitrarily spaced samples that are not integer multiples of the underlying Nyquist rate, since the latter approach would be nearly impossible to calibrate.

The actual implementation, shown in Fig. 2, replaces the theoretical sub-sampled Nyquist-rate ADC with a nonuniformly clocked sample-and-hold (S/H). The sample times of the S/H are controlled by the PRBS sequence, and the same sequence controls a single low-rate ADC which performs the final quantization step. The custom S/H is necessary because the ADC is not designed for 2.2 GHz bandwidth signals.

CS theory confirms that a sufficient number of NUS samples can capture all of the information in the incoming signal—it is possible, after all, to reconstruct all of the Nyquist-rate samples from just the NUS samples. One can think of the NUS as an analog-to-*information* converter,² and given the NUS samples, there are many possible strategies that one could employ for extracting the signal information. Certain very simple questions can be answered with very little processing of compressive samples; examples of such questions include detection or classification of the incoming signal [13] and estimation of signal (e.g., radar pulse) parameters [30]. In other applications, of course, one might wish to actually reconstruct the signal, restoring all of the missing Nyquist-rate samples. This problem is more computationally demanding, if only due to the extremely high number of Nyquist-rate samples that must be computed; any digital processing on billions of samples per second is guaranteed to be expensive. The unique benefit of a NUS receiver over a conventional ADC is that *the signal information is captured directly in compressed form*. Thus, for remote sensing and other applications that do not require information extraction at the sensor, the amount of data that must be transmitted from the sensor can be greatly reduced, and the computational burden can be shifted downstream where power may not be a constraint.

²In fact, this work was supported by DARPA’s Analog-to-Information program.

TABLE I

NUS SPECIFICATIONS AT A GLANCE. POWER INCLUDES THE COMMERCIAL ADC, BUT NOT THE CLOCK OR THE PRBS PATTERN GENERATION. CHIP ALLOWS UP TO 2.2 GHz OF BANDWIDTH THOUGH WE DO NOT HAVE SNDR MEASUREMENTS FOR THIS RANGE. AMOUNT OF OCCUPIED SPECTRUM IS CONSERVATIVELY ESTIMATED, AND FURTHERMORE THE AMOUNT OF OCCUPIED SPECTRUM CAN BE INCREASED WHEN LOWER-FIDELITY RECOVERY IS ACCEPTABLE

	Bandwidth	Occupied spectrum	Power consumption	Resolution of samples
NUS	2.0 GHz	100 MHz	5.8 W	8.8 ENOB, 55 dB SNDR

Due to the unique challenges of the reconstruction problem, we devote part of this paper to proposing an efficient computational framework for reconstruction and to characterizing the computational burden of performing reconstruction. In Section III, we discuss a customized block algorithm for reconstruction that we have implemented on a standard workstation. Each block is composed of $N = 65536$ Nyquist-rate samples, which corresponds to a time interval of length $T = 14.9 \mu s$. When calculating the *occupied* bandwidth, the entire time interval must be considered, so the algorithm-based restriction is that at most 10% of the length- N DFT of the sample block should be nonnegligible. (This 10% figure could be increased if we were to keep more than 440 out of every 8192 Nyquist-rate samples.) Extremely short duration signals are automatically excluded from the signal model since they have a broad frequency spectrum which ruins the sparsity. Our block-by-block reconstruction strategy allows for the capture of frequency hopping signals, although for the specific tests reported in Section IV we require some stationarity of the spectrum (over, say, a period of 2 ms) to allow the spectral support to be identified.

2) *Hardware Specifications*: The NUS integrated circuit (IC) is designed in an Indium-Phosphide (InP) heterojunction bipolar transistor (HBT) technology. A TI ADS5474 14-bit 400 Ms/s ADC (10.9 ENOB at 230 MHz) is used to digitize the samples, and the data is transferred to a computer for processing.

The specifications of the NUS are described in Table I. Power consumption is relatively high because the front-end is designed for wide bandwidth and high dynamic range. However, we stress that lower system power is possible for applications that require data transmission, since the NUS produces $19 \times$ fewer samples.

Comparing the NUS to a Nyquist ADC is difficult since recovery error depends on spectral sparsity. To measure the resolution of the samples, the NUS can be operated in a uniform sampling mode. This measurement shows that the NUS has 8.8 ENOB performance across the frequency band from 800 MHz to 2 GHz. It is important to note that this measurement does not assess the reconstruction accuracy and does not directly relate to ADC ENOB. Instead of a direct comparison, the results in Section IV show promising GSM bit error rate (BER) performance.

C. Related Work in Compressed Sensing

1) *Nonuniform Samplers*: To the best of our knowledge, there have been no IC implementations of the NUS that fully reconstruct the signal. The interesting work [2] on optical sub-

Nyquist sampling is similar in spirit, but it works in the optical domain with commercial off-the-shelf (COTS) components and uses a least-squares fit to reconstruct pure tones rather than a CS-based recovery algorithm to reconstruct an information-carrying modulated signal.

2) *Other CS Devices*: The random modulation pre-integrator (RMPI), a type of random demodulator (RD), is a CS receiver and digitizer which has recently been implemented in [28]–[30]; see also [28]–[30] for references and related approaches. The RMPI uses a far more powerful CS approach that takes nearly random linear combinations of samples, akin to multiplying the digital vector x by a random matrix. The theory predicts that such an approach is optimal for nearly all types of sparse signals, not just signals that are sparse in the Fourier domain. The cost of this generality is that the RMPI is more difficult to implement, the signal processing is less straightforward, there is nontrivial calibration, and recovery is slower than for the NUS. The NUS also allows simple postprocessing windowing techniques.

3) *Xampling*: The modulated wideband converter (MWC) [22], which follows the principles of xampling [21], works directly in the analog domain when possible. The approach requires signals that have a few dense bands of spectrum, such as three or four bands. The digital step is a continuous-to-finite (CTF) block that finds the location of the bands and must be run every time the band structure changes. However, the MWC does not naturally handle our sparse spectral model since the signals may not be easy to group into contiguous blocks. The hardware prototype in [22] has yet to be extended to an IC implementation.

In summary, these other CS and xampling approaches all have their own merits, but for the sparse spectral sensing model defined in the preceding subsection, we believe that the NUS is the best candidate.

D. Outline

In Section II, the implementation of our approach is described. Section III covers the recovery process in detail, describing the general CS recovery method as well as the necessary changes and improvements for our specific architecture. Experimental hardware results are presented in Section IV, and the results of the prototype compare with previous state-of-the-art ADCs. The paper concludes in Section V with some learned wisdom and with a discussion of future challenges.

II. HARDWARE IMPLEMENTATION

A simplified block diagram of the nonuniform sampler (NUS) receiver is shown in Fig. 2. The low-jitter 4.4 GHz clock is used to re-clock the nonuniform pattern to accurately set the sampling instances. For flexibility in testing, the NUS pattern is set by a repeating pseudo-random bit sequence (8192 bits in length) provided by an external pattern generator. A commercial 400 Ms/s ADC (TI ADS5474 [25]) is used to digitize the samples which are captured by a logic analyzer. In order to recover the signal, the samples must be aligned to the NUS pattern; this is accomplished with a synchronization pulse from the pattern generator.

The main building blocks of the NUS receiver are the master and slave sample-and-hold circuits (MSH and SSH), the timing




Fig. 3. Schematic of TG circuit.

generator (TG), and the output buffer. In order to achieve the full bandwidth required, the MSH circuit was designed for 2.4 GHz of bandwidth. The function of the TG is to re-clock the NUS pattern (NIN) with the Nyquist clock (CIN). The re-clocked NIN is then used as the mode control input to drive the MSH/SSH circuits. The output of the SSH is then buffered and amplified so that the signal can drive the external ADC. The output buffer bandwidth determines the settling time of the step-and-settle interface. The chip is designed to perform a sample-and-hold function at a period as long as 6.1 ns and as short as 2.7 ns between consecutive samples, which is under the 400 MHz sampling limit of the ADC.

A. Circuit Description

The main function of the TG, shown in Fig. 3, is to generate timing signals for the master and slave sample-and-holds. This is accomplished by using a low-jitter clock to re-time and delay the NUS pattern input in a chain of flip-flops. Four flip-flop delays are used to delay the ADC clock (NOUT) to provide adequate settling time for AOOUT as well as to account for finite aperture delay of the ADC.

The master and slave sample-and-hold bridge circuits both use diode sampling bridges but have different power consumption based on bandwidth/spur requirements. Fig. 4 shows the basic design of the sampling element. The circuit is controlled by the re-timed NIN coming from the TG to switch the bridge on or off. The schematic shown is only for one of the two pseudo-differential circuits.

The analog input/output interface of the NUS IC was carefully designed to optimize performance. Fig. 5(a) shows the NUS IC input receiver, which acts as a 50- Ω load termination for the analog input signal (AIN) and provides a low impedance drive to the subsequent sampling bridge. The analog output driver [Fig. 5(b)] is a differential quartet design [10]; this was chosen because it offered the required performance while satisfying our finite power consumption goal. The driver was optimized for a large dynamic range and designed to be dc coupled to the TI ADS5474. The gain of the buffer is 2 \times (6 dB), and it was designed to have an SFDR better than 70 dB and 500 MHz bandwidth. To increase simulation accuracy, a detailed interconnect model was used on the step-and-settle ADC interface. This model included bondwires and pallet traces for the NUS packaging, transmission line models for printed circuit board (PCB) traces, a termination network, and the equivalent loading model from the ADC datasheet.

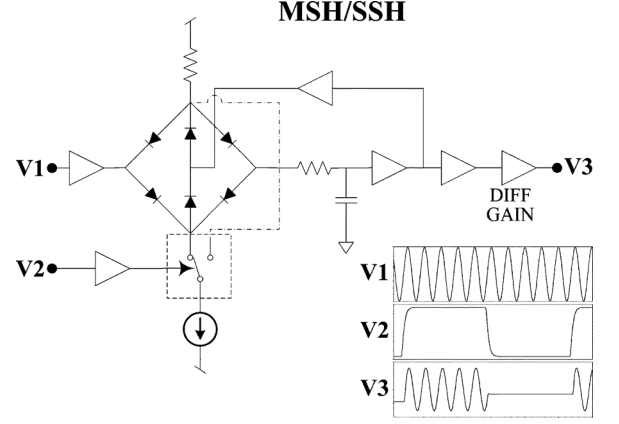


Fig. 4. Sample-and-hold circuit. MSH and SSH are functionally identical. All signal paths and circuits are differential with the exception of the diode bridges. The diode bridges are implemented as two single-ended bridges. The plot shows the output at the transition from the tracking to the holding state (at the MSH output). V1 is the analog input signal, V2 is the MSH clocking signal, and V3 is the MSH output signal. The MSH is tracking V1 when V2 is high. When V2 goes low, the MSH is placed in hold mode.




Fig. 5. NUS IC input/output interface circuits. (a) The analog input receiver is a differential 50 Ω terminated emitter follower. In addition to providing load termination and buffering, the receiver also sets the common-mode level for the master sampling bridge. (b) The output buffer is a differential quartet design.

The overall timing relationship between the NUS circuit and the external ADC is shown in the left panel of Fig. 6. Also displayed in the right panel of Fig. 6 is a simulation showing the NUS operation with a 1.6 GHz input sine wave and the resulting NUS samples.

B. Circuit Fabrication

The NUS IC is fabricated in Northrop Grumman Aerospace Systems' (NGAS) 0.45 μm InP HBT technology featuring f_T and $f_{\text{max}} > 300$ GHz, four-layer metal stack and precision TFR and MIM caps [18]. A die photograph of the 4.0 mm \times 2.6 mm NUS IC is shown in Fig. 7. The NUS die is larger than it needs to be to allow the dicing of other die on the wafer. Pictured in Fig. 8 is the NUS test fixture containing the NUS IC and TI ADS5474 along with various signal and power connectors. The PCB draws a total of 5.8 W: 3.2 W for the NUS IC and 2.6 W for the ADC.

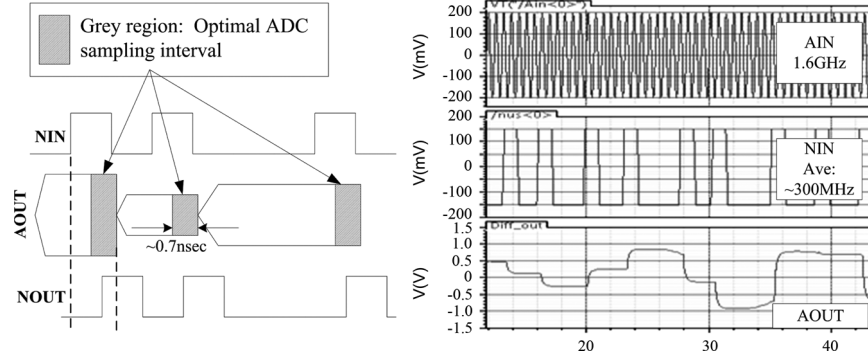


Fig. 6. NUS IC sampling timing and waveforms. Left panel: Interface timing between NUS die and ADC. Right panel: Simulated waveforms before and after being sampled by NIN. Horizontal scale is in nanoseconds.




Fig. 7. NUS IC die photo. Die size is 4.0×2.6 mm.




Fig. 8. NUS test fixture. The NUS IC is mounted on a custom pallet. Also shown is the 14-bit ADC as well as various test equipment connector interfaces.

C. Nonuniform Sample Pattern

The NUS sampling pattern is a pulse train with nonuniform spacing between pulses. It is selected to meet a list of certain criteria. First, the pattern is clocked at 4.4 GHz, and reconstruction produces the equivalent of Nyquist samples taken at this rate. Second, the pulse widths and spacings must satisfy the clocking requirements of the ADS5474. Specifically, the minimum pulse width is 6 clock cycles, the minimum spacing (T_{\min}) between pulses is 12 clock cycles, and the maximum spacing (T_{\max}) between pulses is 27 clock cycles. An example pattern illustrating these specifications is shown in Fig. 9. In effect, as the sample

spacings vary between 12 and 27 clock cycles, the instantaneous sampling rate of the NUS receiver varies between 163 and 367 MHz, which is within the range of the 400 MHz ADC.

We designed the NUS pattern to repeat every 8192 Nyquist samples, during which time there are 440 pulses which set the sampling locations. This corresponds to an average sample rate of 236 MHz. We evaluate the quality of our pattern using a third criterion: the Fourier transform of favorable patterns will tend to have a flat, noise-like spectrum. Fig. 10 compares two NUS patterns with different inter-sample spacings. The pattern shown in the top plots has strong resonances across the Nyquist band. In contrast, the pattern shown in the bottom plots, which has undergone a randomization of its sample locations, has a much whiter spectrum. The flat spectrum is preferred since then all signals have equal gain.

III. DATA PROCESSING

In this section, we describe our computational techniques for efficiently recovering a Nyquist-rate signal from the NUS data by filling in the missing samples. Sections III-A–III-D describe our procedures for windowing the NUS data and recovering the missing samples. Section III-E then briefly discusses additional practical concerns and the computational complexity of implementing this reconstruction algorithm.

A. Windowing

While the NUS produces an arbitrarily long sequence of samples, the recovery algorithm can only deal with a finite number of them at any given time. It is, therefore, necessary to segment the data stream, and we achieve this by windowing the signal. An effective windowing process must guard against edge effects as well as the well-known spectral spreading effect, which would destroy the very Fourier sparsity we seek to exploit. Fortunately, the concept of a *perfect reconstruction filter bank* (PRFB) [24] can be readily adapted to our purposes. A windowing procedure breaks the infinite signal into a series of (possibly overlapping) vectors by using an analysis window. After signal processing, the infinite length signal can be recovered by stitching together the finite series using the analysis window. Using windows from a PRFB ensures that the windowing process itself does not introduce any errors. An example of a PRFB is a rectangular analysis and synthesis window with no overlap, but of course this causes spectral




Fig. 9. NUS pattern example. High pulses have a width of 6 clock cycles. The minimum pulse spacing T_{\min} is 12 clock cycles, and the maximum pulse spacing T_{\max} is 27 clock cycles.

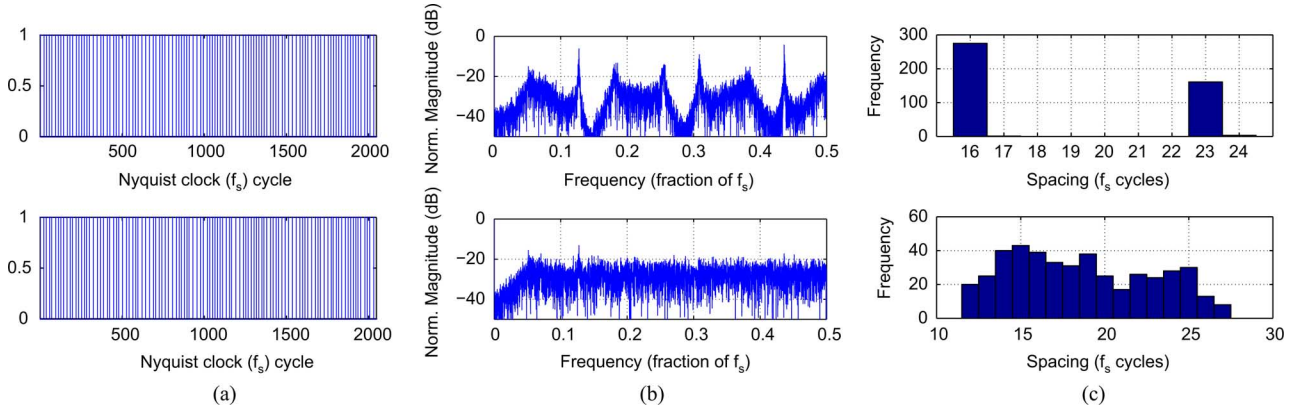


Fig. 10. Comparison of a nonoptimized NUS pattern (top plots) and a properly randomized pattern (bottom plots). (a) Sample patterns in the time domain. (b) Spectral plots. (c) Histograms of inter-sample spacings.

spreading; our implementation uses smooth windowing and 50% overlap.

To describe our processing chain, we let y denote the raw stream of NUS data, i.e., the discrete-time stream of samples coming from the nonuniformly clocked ADC. The processing begins by inserting zeros into y to produce a Nyquist-rate sample stream \underline{y} ; the zeros are inserted at all locations where the NUS did not sample. A delay and downsampling chain then partitions \underline{y} into overlapping windows of length N and multiplies them by an *analysis window function* w_1 . The result is a stream of N -point signals $\underline{y}_j \in \mathbb{R}^N$

$$\underline{y}_j(i) = w_1(i) \underline{y}\left(i + \frac{jN}{2}\right), \quad i = 0, 1, \dots, N-1.$$

Here, $\underline{y}_j(i)$ denotes sample position i in the j th windowed signal. Because these windows are overlapping, each sample $\underline{y}(k)$ maps to two different entries (i, j)

$$k \rightarrow \left(k \bmod \frac{N}{2} + \frac{N}{2}, \left\lceil \frac{2k}{N} \right\rceil - 1\right), \left(k \bmod \frac{N}{2}, \left\lceil \frac{2k}{N} \right\rceil\right).$$

The signals \underline{y}_j are delivered to the sparse recovery engine, which produces a stream of estimates $\hat{x}_j \in \mathbb{R}^N$. The upsampler delay chain stitches these estimates together using an N -point *synthesis window function* w_2 to yield the reconstructed Nyquist-rate sample stream \hat{x} : for each integer j and each $i \in \{0, 1, \dots, N-1\}$

$$\begin{aligned} \hat{x}\left(i + \frac{jN}{2} + d\right) &= w_2(i) \hat{x}_j(i) \\ &+ w_2\left(i + \frac{N}{2}\right) \hat{x}_{j-1}\left(i + \frac{N}{2}\right). \end{aligned}$$

Here, d is the total system delay. The downsample and upsample chains introduce a combined delay of $N-1$ Nyquist samples, so $d = N-1$.

The perfect reconstruction criterion requires that the window functions w_1 and w_2 must satisfy

$$w_1(i)w_2(i) + w_1\left(i + \frac{N}{2}\right)w_2\left(i + \frac{N}{2}\right) = 1$$

for $i = 0, 1, 2, \dots, N/2 - 1$. With this criterion satisfied, we can ensure that the performance of the system is limited by our precise choices of N , w_1 , and w_2 , and by the fidelity of our windowed sparse recovery algorithm. The design of the analysis window w_1 is critical, because it directly affects the spectrum of the estimated signals \hat{x}_j . To minimize the effect of spectral spreading, we must choose an analysis window function w_1 whose spectral sidelobes are well below the system noise floor (high dynamic range), and a main lobe that is as narrow as possible (high sensitivity). We found experimentally that the square of a *Kaiser-Bessel derived* (KBD) window used in audio coding produces excellent results. Furthermore, this choice leads to a rectangular synthesis window (i.e., $w_2(i) \equiv 1$) [24], which is not only convenient but ensures that our reconstruction errors are weighted equally in time. We use $N = 65536$ and KBD parameter $\pi\alpha = 8$, for which the amplitude of the analysis window (in the frequency domain) falls below our system noise floor within six bins. This means that a windowed sinusoid will deliver nontrivial signal energy to no more than 11 DFT bins (of 32768 total bins). It is possible to improve upon this result by designing a window using convex optimization methods, but the KBD window is sufficient for our purposes.




Fig. 11. Block diagram of the practical processing steps.

While the PRFB-inspired architecture has proven useful for verifying the fidelity of our system design, our actual implementation differs in important practical respects. In particular, we choose NUS sampling patterns such that each half-window contains an identical number of NUS samples. This allows us to eliminate the zero-padding step altogether: the NUS samples can be partitioned into overlapping windows of M points each, and these windows can be multiplied by appropriately sampled versions of the analysis window function. The result is a more practical arrangement of processing steps depicted in Fig. 11. The key difference from our earlier description is that the sparse recovery engine no longer receives the zero-padded and windowed N -point signal y_j ; instead its input is y_j , a windowed (but not zero-padded) M -point NUS signal.

With this architecture in place, we can treat each window separately if we choose—although we can (and do) take advantage of the spectral similarity between adjacent windows to improve performance. From this point forward, therefore, we shall focus solely on the generation of estimates $\hat{x}_j \in \mathbb{R}^N$ of signals $x_j \in \mathbb{R}^N$, given nonuniform sample sets $y_j \in \mathbb{R}^M$ and exact knowledge of the sampling pattern. When it is clear from context that we are dealing with a single window, we will drop the j subscript altogether.

B. Frequency Domain Representation

The natural initial choice to compute the frequency domain representation of an N -point signal x is the discrete Fourier transform (DFT). For real signals, the DFT exhibits real symmetry, but the dc and Nyquist component are real-valued; for a more elegant and convenient treatment, we define a slight variant of the DFT that shifts the computed frequency bins by one half

$$X_{\text{shift}}(k) = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} e^{-j2\pi i(k+1/2)/N} x(i)$$

for $k = 0, 1, 2, \dots, N-1$. This modified DFT remains orthonormal and preserves the sparsity behavior of compressible signals, but for real signals it exhibits a simpler symmetry: $X_{\text{shift}}(k) = \overline{X_{\text{shift}}(N-1-k)}$. Thus the frequency domain behavior is captured in $N/2$ complex values. If we preserve only the first $N/2$ frequencies and scale by $\sqrt{2}$ to preserve orthonormality, the result is what we call the *half-bin FFT* (HBFFT)

$$X(k) = \sqrt{\frac{2}{N}} \sum_{i=0}^{N-1} e^{-j2\pi i(k+1/2)/N} x(i)$$

for $k = 0, 1, 2, \dots, N/2-1$. Let $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{C}^{N/2}$ denote the real-to-complex HBFFT operation, so $X = \mathcal{F}(x)$. Because \mathcal{F}

is orthonormal, we have $x = \mathcal{F}^*(X)$, where \mathcal{F}^* denotes the complex-to-real adjoint of \mathcal{F} .

It turns out that the HBFFT \mathcal{F} can be computed as the composition of a single custom butterfly, a standard $N/2$ -point complex DFT, and a simple reshuffling; the inverse HBFFT can also be computed using similar steps. The tight relationship between the HBFFT and the standard complex FFT allows us to achieve high performance with standard FFT libraries.

C. Reprojection on Estimated Support

If the support of the signal is known—that is, if we know which frequency bins contain active signal content (see Section III-D)—and is sufficiently sparse, then we can reduce the reconstruction process to a standard least-squares problem we now introduce. Let $x \in \mathbb{R}^N$ represent the Nyquist-rate signal we wish to estimate, and let $y \in \mathbb{R}^M$ be the NUS samples. Those samples are selected from indexes $\mathcal{I}_T \subseteq \{0, 1, 2, \dots, N-1\}$, $|\mathcal{I}_T| = M$. Thus, x and y satisfy $y = E_T x$, where $E_T \in \mathbb{R}^{M \times N}$ is assembled from rows $i \in \mathcal{I}_T$ of the $N \times N$ identity matrix. Our task is to construct an estimate $\hat{x} \in \mathbb{R}^N$ of x given these samples y . Using the real-to-complex HBFFT operator defined in Section III-B, we let $X = \mathcal{F}(x)$ and $\hat{X} = \mathcal{F}(\hat{x})$ denote the frequency domain representations of x and \hat{x} , respectively.

Let $\mathcal{I}_F \subseteq \{0, 1, 2, \dots, N/2-1\}$, $|\mathcal{I}_F| = P \leq M/2$, denote the support of the signal. We can write our estimate as $\hat{X} = E_F^* z$, where $z \in \mathbb{C}^P$ is a set of nonzero coefficients to be determined below, and $E_F \in \mathbb{R}^{P \times N/2}$ is a frequency sampling matrix assembled from rows $i \in \mathcal{I}_F$ of the $N/2 \times N/2$ identity matrix.

With these definitions in place, the reprojection problem can be cast as

$$\hat{x} = \mathcal{F}^*(E_F^* z), \quad z \triangleq \arg \min_z \|E_T \mathcal{F}^*(E_F^* z) - y\|_2.$$

To help explain the above notation, let us note that $E_T \mathcal{F}^*(E_F^* z)$ is computed from z by constructing a length- $N/2$ vector containing the P entries of z in the positions indexed by \mathcal{I}_F , computing the complex-to-real inverse HBFFT of this vector, and extracting from the result the M values in the positions indexed by \mathcal{I}_T . The minimization on z is a least-squares problem and can be expressed as normal equations

$$E_F \mathcal{F}(E_T^* E_T \mathcal{F}^*(E_F^* z)) = E_F \mathcal{F}(E_T^* y).$$

The linear operation $z \rightarrow E_F \mathcal{F}(E_T^* E_T \mathcal{F}^*(E_F^* z))$ is positive definite, but it cannot be expressed as a complex matrix due to the presence of the complex-to-real operation \mathcal{F}^* . However, we have chosen to solve this form using conjugate gradients, which allows us to utilize the HBFFT and sampling operators directly. Furthermore, CS theory shows that the linear operator is well-conditioned, so conjugate gradients will converge rapidly.




Fig. 12. Three iterations of the greedy selection procedure for preliminary support estimation. In each iteration, blocks of indexes with high aggregate energy estimates $h_{\text{agg}}(i)$ are identified (top row). A reprojection step then removes their influence from the measurements so that additional active blocks can be identified in subsequent iterations. (The aggregate statistics $h_{\text{agg}}(i)$ are plotted before being reset with $h_{\text{old}}(i)$ in step 1c of our algorithm below.) Plots in the bottom row show, for one of J windows used in the estimation, the reprojected spectrum estimate $|\hat{X}_j(i)|$ using the current support estimate.

D. Spectral Occupancy Estimation

The problem of recovering a signal from NUS samples can be partitioned into two subproblems: 1) first estimate the support of the unknown signal in the frequency domain (we refer to this step as *spectral occupancy estimation*), and then 2) reproject to estimate the signal given this estimated support. We have developed a customized algorithm for spectral occupancy estimation that is inspired by existing techniques in CS but adapted to the specific nuances of our problem.

The most unique aspects of our problem that we seek to exploit are as follows. 1) The nonzero HBFFT coefficients for a given window—while few in number—also tend to cluster into an even smaller number of contiguous groups. In the CS literature, this is known as a *structured sparsity model* [3]. Our algorithm is inspired by existing ones in the model-based CS literature designed to exploit block-sparse [17] and clustered-sparse [11] models. 2) Although our data are partitioned into finite windows (as described in Section III-A), the spectral occupancy is often stationary over the duration of multiple windows. In cases like this—where multiple sparse signals share the same support—the signals are said to obey a *joint sparsity model*. Like others in the distributed CS literature [5], our algorithm processes the data from multiple windows jointly in order to better identify the support.

Our spectral occupancy estimation algorithm is greedy: we first run a few iterations of a greedy selection rule that builds an estimate of the support, and we then perform a pruning procedure to remove false positives. The greedy selection procedure (step 1 below) is iterative because it is difficult to identify all of the active frequencies at once; at each iteration, only the largest frequencies can be accurately estimated, since artifacts from these large signals will hide smaller signals. As illustrated in Fig. 12, however, once some blocks of active frequencies have been identified, a reprojection step removes their influence from the measurements, and weaker active frequencies can then be identified. We find the subsequent pruning (step 2 below) to be helpful because setting the thresholds in step 1 low enough to detect weak signals tends to also introduce a number of false positives.

To describe our algorithm, let us set the following notation. Let $x_j \in \mathbb{R}^N$ represent the unknown Nyquist-rate samples from window number j , and let $y_j \in \mathbb{R}^M$ represent the NUS samples. Those samples are selected from indexes $\mathcal{I}_{T,j} \subseteq \{0, 1, 2, \dots, N-1\}$, $|\mathcal{I}_{T,j}| = M$, and so following the notation defined in Section III-C, $y_j = E_{T,j}x_j$, where

$E_{T,j} \in \mathbb{R}^{M \times N}$. From y_j we would like to estimate the positions \mathcal{I}_F of the non-negligible entries of $X_j = \mathcal{F}(x_j)$. We accomplish this by considering an ensemble of windows $j \in \{0, 1, \dots, J-1\}$ simultaneously and exploiting the assumption of stationarity (i.e., we assume that \mathcal{I}_F does not change from window to window). Our algorithm consists of the following steps.

1) Preliminary support estimation

- a) Set the iteration count $\ell = 1$, and for each window j , define a residual vector $r_j = y_j$. Set the initial support estimate to be empty: $\hat{\mathcal{I}}_F := \emptyset$. Set the maximum allowable size for the support estimate $P_{\text{max}} \approx 0.8M/2$. Set $h_{\text{old}}(i) = 0$ for each $i \in \{0, 1, \dots, N/2-1\}$.
- b) For each window j , compute the correlation statistics $h_j = \mathcal{F}(E_{T,j}^* r_j)$. Then, square and sum the correlation statistics over multiple windows: for each $i \in \{0, 1, \dots, N/2-1\}$, compute the aggregate statistic $h_{\text{agg}}(i) = \sum_{j=0}^{J-1} |h_j(i)|^2$.
- c) Identify a set Γ of possible active frequencies. We add an index $i \in \{0, 1, \dots, N/2-1\}$ to Γ if $|h_j(i)|$ is frequently among the largest entries of h_j across multiple windows j , or if $|h_{\text{agg}}(i)|$ is among the largest entries of h_{agg} . We also include all entries of the previous support estimate $\hat{\mathcal{I}}_F$. For the purpose of computations below, we then set $h_{\text{agg}}(i) = h_{\text{old}}(i)$ for each $i \in \hat{\mathcal{I}}_F$ and reset $\hat{\mathcal{I}}_F = \emptyset$.
- d) Pad the set Γ with some number p of indexes on both sides of each index in Γ . For example, if $p = 1$ and $\Gamma = \{2, 3, 4, 6, 18, 19\}$, update Γ to $\{1, 2, 3, 4, 5, 6, 7, 17, 18, 19, 20\}$.
- e) Identify contiguous blocks of indexes in Γ , and compute the estimated energy of each block. For example, for the updated Γ given above, two blocks are identified, and their corresponding energies are $h_{\text{agg}}(1) + \dots + h_{\text{agg}}(7)$ and $h_{\text{agg}}(17) + \dots + h_{\text{agg}}(20)$.
- f) Populate the new support estimate $\hat{\mathcal{I}}_F$ with all of the indexes from the highest energy blocks, such that $|\hat{\mathcal{I}}_F|$ does not exceed $(\ell/10)P_{\text{max}}$. This increasing threshold allows slightly larger support estimates at each iteration.
- g) On each window, use a reprojection step to project the observations y_j orthogonal to the chosen support $\hat{\mathcal{I}}_F$, and let r_j denote the resulting residual.

- h) Store the aggregate energy estimates for use in future iterations, setting $h_{\text{old}}(i) = h_{\text{agg}}(i)$ for all i . Then, increment the iteration counter ℓ . Stop when $\ell = 10$ or the energy in the residual vectors r_j is sufficiently small. Otherwise, repeat steps 1b–1g.
- 2) Final pruning
 - a) Set the iteration count $\ell = 1$.
 - b) Reproject each set of samples y_j onto the estimated support $\widehat{\mathcal{I}}_F$ to obtain an estimate \widehat{X}_j for the HBFFT coefficients. Square and sum these estimates: for each $i \in \{0, 1, \dots, N/2 - 1\}$, compute $\widehat{X}(i) = \sum_{j=0}^{J-1} |\widehat{X}_j(i)|^2$.
 - c) For each contiguous block of indexes in $\widehat{\mathcal{I}}_F$, compute the largest value of \widehat{X} , e.g., if $\widehat{\mathcal{I}}_F = \{2, 3, 4, 6, 18, 19\}$, compute $\max\{\widehat{X}(2), \widehat{X}(3), \widehat{X}(4)\}$, $\widehat{X}(6)$, and $\max\{\widehat{X}(18), \widehat{X}(19)\}$.
 - d) Remove blocks from $\widehat{\mathcal{I}}_F$ whose maximum \widehat{X} value does not exceed some threshold designed to eliminate false positives. Increment the iteration counter ℓ .
 - e) Repeat steps 2b–2d for a small number of iterations.
 - f) Following the same procedure as in step 1d, pad the support estimate $\widehat{\mathcal{I}}_F$ with some number p of indexes on each side of each estimated block. (This procedure operates only so long as $|\widehat{\mathcal{I}}_F| \leq P_{\max}$.)

After running the entire support estimation algorithm on an ensemble of J windows, one can rerun the algorithm on one or more subsequent ensembles of J windows, either for cross-validation purposes or to detect changes in the spectral occupancy.

Overall, our algorithm most closely resembles the well-known OMP and CoSaMP algorithms in CS [23], [26], although we have also experimented with reweighted ℓ_1 minimization [9] and believe that, with appropriate modifications, it could be competitive as well. Our algorithm does enjoy some important advantages over OMP and CoSaMP for this problem. First, OMP estimates a support by adding one index per iteration and thus would require many hundreds of iterations for the scenarios we have tested; our algorithm is guaranteed to terminate in a small number of iterations (at most 10 in step 1 plus a small fixed number in step 2). Second, neither OMP or CoSaMP can exploit the fact that the nonzero HBFFT coefficients should tend to cluster into contiguous groups; our algorithm improves support estimates by exploiting this block structure. Finally, our algorithm processes the data from multiple windows jointly, and this allows supports to be estimated up to a maximum size of $P_{\max} \approx 0.8M/2$. Neither CoSaMP nor OMP (when run on a single window) can tolerate sparsity levels that are such a significant fraction of the measurement rate.

E. Additional Implementation Concerns

1) *Model Violations:* Our system is designed to support signals with up to 100 MHz of information bandwidth. There are a number of strategies that one could use to confirm that the input signal obeys this model assumption. For example, in step 1 of the support estimation algorithm described in Section III-D, the energy of the residual vectors r_j should decrease substantially as the number of iterations increases. If significant energy remains in the residual vectors after the maximum number of iterations, this means that the estimated support is not sufficient to

fully capture the structure in the input signal. A second possible strategy for detecting model violations could be cross validation. For example, 95% of the NUS samples could be used for support estimation and recovery, and the remaining 5% of the NUS samples could be checked against the reconstructed estimates. A close match suggests that the information bandwidth is well captured in the estimated support.

In cases where small model violations are detected, CS theory guarantees that the reconstruction on the estimated support will be relatively accurate, although the small signal components away from this support will of course not be reconstructed. In cases where substantial model violations are detected, reconstruction on the estimated support will not be accurate. There is research into additional analog (preprocessing) and digital (postprocessing) safeguards that could be added to our system to protect against such situations. If the model is violated too frequently, then the NUS is simply not the correct device for the task.

2) *Computation:* Because our system is designed to have a high EIBW, the calculations described in Sections III-C to III-D can be expensive. To quantify these costs, suppose we solve the reprojection problem by applying conjugate gradients (CG) to the normal equations. In stationary signal environments, one could alternatively employ a pseudoinverse to solve the reprojection problem, and in some problems—especially when the spectrum is very sparse—this could be faster than CG. The downside is that computing the pseudoinverse is expensive, and it must be recomputed anytime the spectral occupancy changes.

The cost of solving the reprojection problem by applying CG to the normal equations is dominated by the FFTs, each of which requires $2.5N \log_2 N$ flops. Two FFTs are required per CG iteration, and one each for initialization and finalization. Because the windows overlap by 50%, reconstruction occurs at a rate of $N/2$ samples per window. Therefore, the throughput required to perform real-time reprojection is

$$\begin{aligned} C &= f_r \cdot (2I + 2) \cdot 2.5N \log_2 N / (N/2) \\ &= 10f_r(I + 1) \log_2 N \quad \text{flops/s} \end{aligned}$$

where I is the number of CG iterations and f_r is the reconstructed sample rate. Although we omit the details here, in some problems where the EIBW is less than half of the device Nyquist rate f_s we can envision using digital downconversion to reconstruct at a rate f_r equal to just twice the EIBW. When we are interested in the 800 MHz–2 GHz band, for example, it is possible to reconstruct at a rate of just $f_r = 2.4$ Gs/s.³ Taking this number as an example, with $(N, I, f_r) = (65536, 30, 2.4 \text{ Gs/s})$ we estimate that real-time reconstruction would require $C \approx 12$ Tflops/s. Because C is super-linear in the window size N , the required throughput could be reduced slightly by decreasing N , but this would diminish the accuracy of the reconstructed signal. Regardless, no single processor achieves performance in this range; parallelism must certainly be exploited.

Because reprojection is an elemental computational problem that would appear in many conceivable reconstruction algo-

³This would require using a device Nyquist rate of 4.8 Gs/s, which is part of our future specification for the system. In addition, we note that a system using downconversion would more realistically output complex-valued samples at 1.2 Gs/s rather than real-valued samples at 2.4 Gs/s, although the cost is the same.

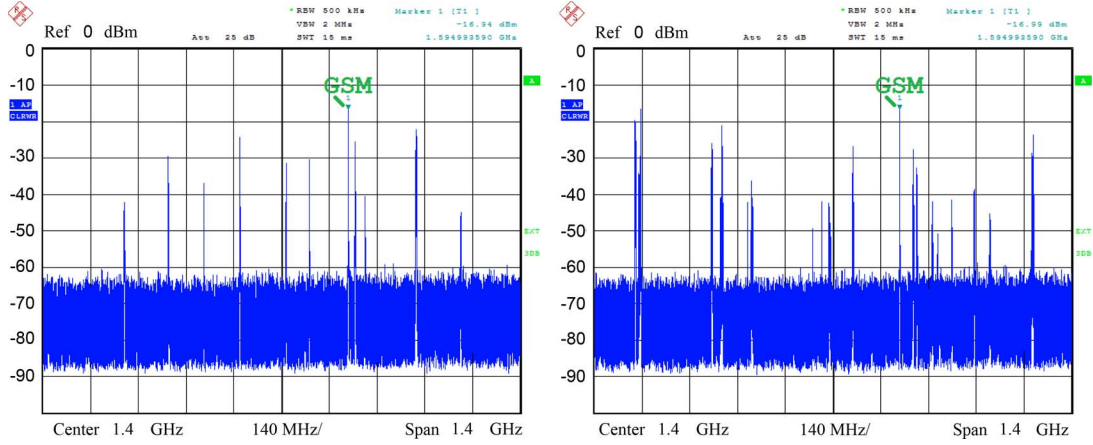


Fig. 13. Example GSM test spectra with 20 MHz clutter (left) and 100 MHz clutter (right). Spectra are plotted from 700 MHz–2.1 GHz. The GSM signal is located at 1.595 GHz and is indicated with a green marker.

rithms, we have focused our discussion more on this topic than that of support identification. However, we can also mention that for our support identification algorithm, the bulk of the effort is consumed by the rejections, and about 14 of these are performed on each window in a typical run. We do not envision this support identification algorithm being applied to every single window but rather being repeated every so often on ensembles of J windows. To choose some example numbers (which match those in Section IV), we estimate the cost of applying this algorithm on an ensemble of $J = 31$ windows of length $N = 65536$ to be approximately 71 Gflops. Once the spectral occupancy has been estimated, of course, the same support can be used for reprojecting multiple windows of NUS samples. The frequency with which the support must be recomputed will depend on the stationarity of the signal environment.

IV. TESTING AND VALIDATION

A. Experimental Setup

In order to test our complete NUS architecture in a representative environment, we conducted a series of experiments using realistic GSM data. To be clear, we do not intend for GSM to be the only possible application of the NUS architecture; our choice to demonstrate our system using GSM signals was motivated by their complexity and by the widespread familiarity with the standard. In the experiments that follow, all signals were generated by an arbitrary waveform generator (AWG) and a vector signal generator (VSG).

For each experiment, we construct a 270.833 kb/s GSM signal that has a bandwidth of ~ 1 MHz (at -50 dBc) and is located at a center frequency of 1.595 GHz. The GSM signal power is scaled between -20 dBFS and -80 dBFS. To the scaled GSM signal, we add various levels of “clutter” consisting of narrowband RF signals having random amplitudes and centered at random frequencies between 800 MHz and 2 GHz. Clutter signals with bandwidths of 20, 50, or 100 MHz are used, and two different cases were generated for each bandwidth. The clutter is used to increase the information bandwidth of the signal, even though the clutter itself is not of interest.

Fig. 13 shows example signals, including the GSM signal and the clutter. It is important to note that our measure of the clutter bandwidth includes spectral “tails” down to -90 dBFS (measured after windowing); as the experiments illustrate, however, it is possible to reconstruct the signal with high accuracy while omitting these tails and thus requiring less overall bandwidth.

After measurement, the NUS data are arranged into blocks of size $M = 3520$ (each corresponding to $N = 65536$ Nyquist rate samples) and run through the spectral occupancy estimation algorithm (Section III-D) using $J = 31$ overlapping windows (these span a total of 1 048 576 Nyquist-rate samples). As a means of cross validation, the support estimation procedure is repeated 10 times (each on a fresh set of $J = 31$ overlapping windows). A frequency bin is included in the final support estimate if it appears in at least 2 out of the 10 preliminary estimates.

With the estimated support, we reproject the NUS data on each window to recover an estimate of the Nyquist-rate signal samples. The windowed samples are then recombined as described in Section III-A. These estimated Nyquist-rate samples are then passed through a GSM decoder (that has *a priori* knowledge of the center frequency of 1.595 GHz) to measure the BER. Note that input powers in the range of -55 to -75 dBFS yield measurable BER. Above -55 dBFS the BER drops to a rate that makes collecting and processing an adequate number of samples difficult. Very low power inputs yield high BERs and make synchronization of the decoder difficult; the GSM signal will only be present in our reprojected samples if the band around 1.595 GHz is correctly identified as part of the support, and this becomes less likely when the input power is very low.

A simplified diagram of the NUS test setup used in the experiment is shown in Fig. 14. Not shown are the controller connections, differential lines, filtering, and power supplies.

For the sake of comparison, we also sampled the test signals using an NGAS developed 5 Gs/s 8-bit Nyquist ADC that uses the same InP technology. A description of an earlier version of this ADC chip can be found in [12]. This ADC uses folding-interpolating architecture, has a greater than 7 ENOB performance, and draws 9.6 W. Because the uniform ADC automatically produces Nyquist samples, its output can be passed directly to the GSM decoder. Testing was done using a 4.4 GHz

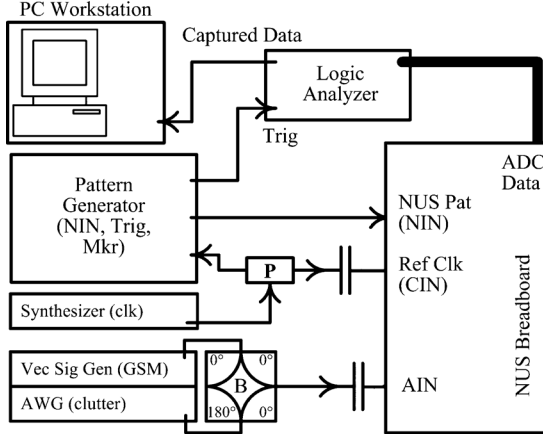


Fig. 14. NUS test setup. The GSM signal is produced by the VSG, clutter is produced by the AWG, and the NUS pattern is produced by the pattern generator. ADC data is captured via the logic analyzer and is then downloaded and processed by the computer. Component “P” is a power splitter and “B” is a balun.

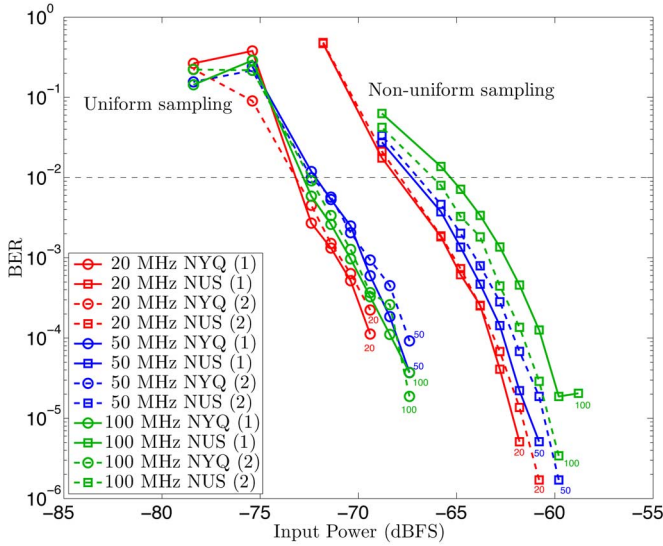


Fig. 15. BER of the decoded GSM signal as a function of input power. Circular markers indicate the performance of the uniform ADC for each of two randomly generated signals (denoted (1) or (2)) at each of three levels of clutter (20, 50 and 100 MHz). Square markers indicate the performance of the NUS on the same signals. The solid and dashed colored lines correspond to separate trials. The dashed black line indicates a BER of 10^{-2} , at which the dynamic range is evaluated for generating Fig. 16.

clock and the output was subsampled by four because of equipment limitations.

B. Experimental Results

Fig. 15 plots the BER as the input power of the GSM signal is reduced. Circular markers indicate the performance of the uniform ADC for each of two randomly generated environments (distinguished by solid and dashed lines) at each of three levels of clutter (distinguished by color). Square markers indicate the performance of the NUS on the same signals. With the uniform ADC at a given BER (e.g., 10^{-2} , as indicated by the dashed horizontal line), we see relatively little variation in dynamic range across the various signals. That is, the performance of the uniform ADC does not depend on the information bandwidth of the

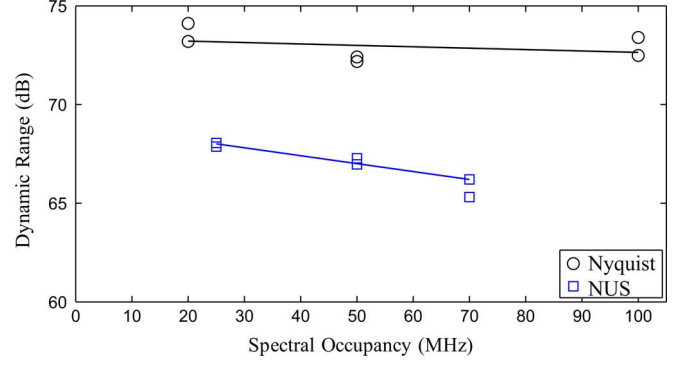


Fig. 16. Dynamic range of the uniform ADC and the NUS as a function of the information bandwidth. The information bandwidth plotted for the NUS curve indicates the total amount of bandwidth identified by the spectral occupancy estimation algorithm. For the signal with 100 MHz of clutter, it is possible to achieve high dynamic range while identifying less than 80 MHz of occupied bandwidth; this is because only the small tails of the clutter signal are omitted.




Fig. 17. Comparison of the SINAD and SFDR versus frequency of the Nyquist ADC and the NUS in uniform sampling mode.

input signal. With the NUS, in contrast, for signals with higher levels of information bandwidth (more clutter), we do see a decrease in the dynamic range. This is to be expected, since the difficulty of accurately estimating the spectral occupancy increases, and the accuracy of the reprojected signal will degrade slightly.

Fig. 16 captures this trend more clearly by plotting the dynamic range of the two systems as a function of the information bandwidth. For this graph, dynamic range is defined as the minimum input power (dBFS) that yields a BER of 10^{-2} ; the displayed values are obtained by taking a “cross section” of the plot in Fig. 15 across the dashed horizontal line at the BER of 10^{-2} . For all levels of the information bandwidth, the uniform ADC has a higher dynamic range, but this is not surprising because the uniform ADC collects more samples in total (approximately $4.7\times$ more than the NUS in this case).

As a separate experiment, we operate our NUS architecture in a 300 MHz uniform sampling mode: samples are not spaced according to the PRBS but rather occur in equispaced intervals. This allows the system to be treated like a low-rate ADC and characterized without the need for reconstruction. Fig. 17 compares the SINAD and SFDR (as a function of input frequency) of the Nyquist ADC and the NUS. The difference in sampling rates between the ADC and the NUS is not a problem since it

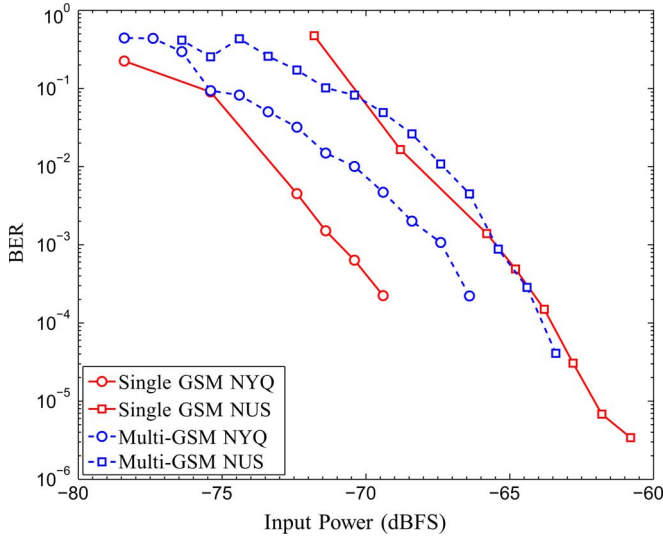


Fig. 18. BER of the decoded GSM signal as a function of input power. Circular markers indicate the performance of the uniform ADC for a single GSM signal (red, solid) and a multi-GSM signal (blue, dashed) with 20 MHz of clutter. Square markers indicate the performance of NUS on the same signals.

does not affect the measurements, only the amount of folding that occurs. The plots show the NUS has a 10 dB advantage in both SINAD and SFDR while consuming about half the power. The NUS SINAD advantage partially offsets the noise penalty of undersampling.

The NUS SFDR advantage is a result of the lower average sampling rate, which allows for more settling time in the sample-and-hold. However, this performance advantage is not evident in the single-carrier GSM testing described above. An example in which the SFDR advantage is important is the case of a signal with multiple active GSM channels that, when passed through the nonlinear samplers, will produce inter-modulation products that fall directly on adjacent channels. A particularly stressing case is when there are three channels, equally spaced in frequency, with two channels having higher power than the third test channel. Thus, we have repeated the GSM BER experiments with multiple GSM channels active simultaneously; the results for the uniform ADC and for the (nonuniform) NUS are shown in Fig. 18 with 20 MHz of clutter. For the multi-GSM signal, we see that the Nyquist ADC dynamic range reduces by approximately 3 dB, while the NUS dynamic range is unchanged.

This phenomenon is also illustrated by the spectra in Fig. 19, in which GSM signals in channels 2 and 3 produce inter-modulation products that appear in channels 1 and 4. The amplitudes of the GSM signals in channels 2 and 3 were skewed to produce a quiet test channel 1 (not active in the plot), so that the measurements would not be limited by the signal generator. The black line shows the signal generator output measured on a spectrum analyzer. The blue line shows the reconstructed spectrum from the NUS using an oracle support estimate spanning the four channels. Notice that the NUS is faithfully reproducing the signal generator output. The red line shows the spectrum from the Nyquist ADC, and here the ADC produces inter-modulation products in channels 1 and 4 that are larger than the signal generator distortion and will affect the dynamic range in




Fig. 19. Reconstructed spectrum of a multi-GSM signal from the NUS (blue, dash-dot) and Nyquist ADC (red, dashed) as compared to the original signal generator output (black, solid).

the test channel. In this case the NUS has a 3–5 dB advantage over the Nyquist ADC. This example was chosen to highlight the SFDR advantage of the NUS and may not be representative of a power controlled GSM spectrum. However, GSM is not the only possible application for the NUS and many scenarios demand a wide dynamic range.

V. CONCLUSION

In summary, we have presented a wide bandwidth receiver system based on nonuniform sampling. A custom InP sampler front-end extends the instantaneous bandwidth of a commercial ADC (14 bits, 400 Ms/s max) to 2.4 GHz. The system performance was verified with a GSM signal embedded in RF environments of varying sparsity levels. As a proof-of-concept the entire RF spectrum was reconstructed prior to using standard digital signal processing techniques to isolate and decode the GSM signal. While the NUS can be similar in function to an ADC, we emphasize that it is not a drop-in replacement for an ADC. Rather, it is a powerful tool in the signal processing toolkit that is useful when the signal is supported on a small (unknown) bandwidth inside a large frequency range. Our experimental results show that the hardware and algorithms are performing as expected and that the custom S/H is not a bottleneck of the system. Thus, by using a faster off-the-shelf ADC, the NUS could be scaled to even higher bandwidths.

We have taken care to outline a computational framework for recovering the Nyquist-rate signal from the NUS data and to characterize the computational complexity of this procedure. In short, even an optimized reconstruction algorithm can present a significant computational burden, if only due to the extremely high number of Nyquist-rate samples that must be computed. Do we believe that real-time reconstruction is possible? We have conducted a variety of tests to demonstrate the feasibility of using graphics processing units (GPUs) to accelerate the key computations; for example, we achieved 93% linearity on a 6-GPU implementation of our reprojection algorithm. While a real-time system would be expensive (and would certainly

require more than 6 GPUs), we are optimistic that it could be built. We reiterate that any digital processing on billions of samples per second is guaranteed to be expensive. Our reprojection estimates in Section III-E2 involve only a few tens of FFTs per data window; the complexity is really dominated by the sheer number of Nyquist-rate samples that must be computed. We also stress that for remote sensing and other applications that do not require information extraction at the sensor, a NUS system allows the amount of data that must be transmitted to be greatly reduced compared to a conventional ADC. In settings such as these, reconstruction can be performed downstream where power and computation may be less of a constraint.

We do believe that further reductions to the computational burden of reconstruction may be possible. For example, in settings where only a portion of the spectrum changes at a time, our support identification algorithm could be modified to use a “warm start,” and this would reduce the required number of iterations. The remaining limitation of our system is the assumption of sparsity. The underlying issue is that the signal information rate cannot exceed the sampling rate, and spectral sparsity is a convenient proxy for the information rate. For specific applications that permit a tighter signal model (e.g., involving a manifold [4] or a union-of-subspaces [3]) it may be possible to devise improved recovery algorithms that have a more relaxed sparsity assumption; we leave the application of these ideas to the NUS for future work.

Finally, although we have focused our study on the reconstruction problem, it is also worth emphasizing that certain other questions can be answered with very little processing of the NUS samples. We give one example that is relevant to the Analog-to-Information program. Suppose we receive a chirp signal

$$x(t) = A \exp \left(j2\pi \left(f_c(t - t_0) + \frac{\alpha}{2}(t - t_0)^2 \right) \right)$$

over some time interval, where f_c denotes the known starting frequency, α denotes the known chirp rate, A denotes the complex amplitude, and t_0 denotes the unknown time-of-arrival. We can “de-chirp” this signal over this interval, computing

$$\tilde{x}(t) = x(t) \exp \left(-j2\pi \left(f_c t + \frac{\alpha}{2} t^2 \right) \right) = \tilde{A} e^{-j2\pi \alpha t_0 \cdot t}$$

where \tilde{A} is a complex amplitude. The signal $\tilde{x}(t)$ is merely a complex sinusoid, and from a small number of NUS samples it is quite simple [16] to determine its frequency t_0 and therefore deduce the time-of-arrival of $x(t)$. (The computational cost to do this is essentially just one FFT.) Since $\tilde{x}(t)$ is defined just via point-wise multiplication of $x(t)$, it is straightforward to transform NUS samples of $x(t)$ into NUS samples of $\tilde{x}(t)$. Therefore, this example provides one alternative application in which processing of NUS samples could more easily be performed at the sensor. In preliminary experiments, we have confirmed that such chirp times-of-arrival can indeed be estimated from real NUS data.

ACKNOWLEDGMENT

The authors are indebted to Dr. Dennis Healy for his foresight, encouragement, and support of this endeavour.

REFERENCES

- [1] B. Adcock and A. C. Hansen, Generalized sampling and infinite dimensional compressed sensing, Tech. Rep. 2011/NA12, DAMTP, 2011.
- [2] M. B. Airoola, S. R. O'Connor, M. L. Dennis, and T. R. Clark, Jr., “Experimental demonstration of a photonic analog-to-digital converter architecture with pseudorandom sampling,” *IEEE Photon. Tech. Lett.*, vol. 20, no. 24, pp. 2171–2173, Dec. 2008.
- [3] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, “Model-based compressive sensing,” *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1982–2001, Apr. 2010.
- [4] R. G. Baraniuk and M. B. Wakin, “Random projections of smooth manifolds,” *Foundat. Computat. Math.*, vol. 9, no. 1, pp. 51–77, 2009.
- [5] D. Baron, M. F. Duarte, M. B. Wakin, S. Sarvotham, and R. G. Baraniuk, “Distributed Compressive Sensing,” 2009, Arxiv Preprint arXiv:0901.3403.
- [6] E. J. Candès and Y. Plan, “A probabilistic and RIPless theory of compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 57, no. 11, pp. 7235–7254, Nov. 2011.
- [7] E. J. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [8] E. J. Candès and M. Wakin, “An introduction to compressive sampling,” *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Feb. 2008.
- [9] E. J. Candès, M. B. Wakin, and S. P. Boyd, “Enhancing sparsity by reweighted ℓ_1 minimization,” *J. Fourier Anal. Appl.*, vol. 14, no. 5–6, pp. 877–905, Dec. 2008.
- [10] R. Capio, “Precision differential voltage-current convertor,” *Electron. Lett.*, vol. 9, no. 6, pp. 147–148, 1973.
- [11] V. Cevher, P. Indyk, C. Hegde, and R. Baraniuk, “Recovery of clustered sparse signals from compressive measurements,” in *Proc. Int. Conf. Sampl. Theory Appl. (SampTA)*, 2009.
- [12] B. Chan, B. Oyama, C. Monier, and A. Gutierrez-Aitken, “An ultra-wideband 7-Bit 5-Gsps ADC implemented in submicron InP HBT technology,” *IEEE J. Solid-State Circuits*, vol. 43, no. 10, pp. 2187–2193, Oct. 2008.
- [13] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk, “Signal processing with compressive measurements,” *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 445–460, Apr. 2010.
- [14] M. A. Davenport and M. B. Wakin, “Compressive sensing of analog signals using discrete prolate spheroidal sequences,” *Appl. Comp. Harm. Anal.*, vol. 33, no. 3, pp. 438–472, Nov. 2012.
- [15] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [16] A. Eftekhar, J. Romberg, and M. B. Wakin, “Matched filtering from limited frequency samples,” 2011, Arxiv Preprint arXiv:1101.2713.
- [17] Y. C. Eldar, P. Kuppinger, and H. Bölcskei, “Block-sparse signals: Uncertainty relations and efficient recovery,” *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 3042–3054, Jun. 2010.
- [18] A. Gutierrez-Aitken *et al.*, “Advanced InP HBT technology at Northrop Grumman Aerospace Systems,” in *Proc. IEEE Compound Semicond. Integrated Circuit Symp.*, Oct. 2009, pp. 1–4.
- [19] Y.-P. Lin and P. P. Vaidyanathan, “Periodically nonuniform sampling of bandpass signals,” *IEEE Trans. Circuits Syst. II*, vol. 45, no. 3, pp. 340–351, Mar. 1998.
- [20] M. Mishali and Y. C. Eldar, “Blind multiband signal reconstruction: Compressed sensing for analog signals,” *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.
- [21] M. Mishali and Y. C. Eldar, “From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals,” *IEEE J. Sel. Top. Signal Process.*, vol. 4, no. 2, pp. 375–391, Apr. 2010.
- [22] M. Mishali, Y. C. Eldar, O. Dounaevsky, and E. Shoshan, “Xampling: Analog to digital at sub-Nyquist rates,” *IET Cir. Dev. Syst.*, vol. 5, no. 1, pp. 8–20, Jan. 2011.
- [23] D. Needell and J. A. Tropp, “Cosamp: Iterative signal recovery from incomplete and inaccurate samples,” *Appl. Comput. Harmon. Anal.*, vol. 26, pp. 301–321, 2009.
- [24] J. O. Smith, Spectral Audio Signal Processing [Online]. Available: <http://ccrma.stanford.edu/~jos/sasp/>
- [25] Texas Instruments, 14-Bit, 400-MSPS analog-to-digital converter Aug. 2008 [Online]. Available: <http://www.ti.com/lit/ds/sym-link/ads5474.pdf>
- [26] J. A. Tropp and A. C. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.
- [27] J. A. Tropp and S. J. Wright, “Computational methods for sparse solution of linear inverse problems,” *Proc. IEEE*, vol. 98, no. 6, pp. 948–958, Jun. 2010.

- [28] J. Yoo, S. Becker, M. Loh, M. Monge, E. Candès, and A. Emami-Neyestanak, "A 100 MHz-2 GHz $12.5\times$ sub-Nyquist rate receiver in 90 nm CMOS," in *Proc. IEEE Radio Freq. Integr. Circ. Conf.*, 2012, pp. 31–34.
- [29] J. Yoo, S. Becker, M. Loh, M. Monge, E. Candès, and A. Emami-Neyestanak, "Design and implementation of a fully integrated compressed-sensing signal acquisition system," in *Proc. Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, 2012, pp. 5325–5328.
- [30] J. Yoo, C. Turnes, E. Nakamura, C. Le, S. Becker, E. Sovero, M. Wakin, M. Grant, J. Romberg, A. Emami-Neyestanak, and E. Candès, "A compressed sensing parameter extraction platform for radar pulse signal acquisition," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 3, 2012.

Michael Wakin (M'06) received the B.A. degree in mathematics, in 2000, and the B.S., M.S., and Ph.D. degrees in electrical engineering, in 2000, 2002, and 2007, all from Rice University, Houston, TX.

He was an NSF Mathematical Sciences Postdoctoral Research Fellow at the California Institute of Technology, from 2006 to 2007 and an Assistant Professor at the University of Michigan from 2007 to 2008. He is now an Assistant Professor in the Department of Electrical Engineering and Computer Science at the Colorado School of Mines.

Dr. Wakin shared the Herschel M. Rich Invention Award from Rice University for the design of a single-pixel camera based on compressive sensing in 2007. In 2008, he received the DARPA Young Faculty Award for his research in compressive multi-signal processing and in 2012, he received the NSF CAREER Award for research into dimensionality reduction techniques for structured data sets.

Stephen Becker (M'10) received the B.A. degree in mathematics and in physics from Wesleyan University, Middletown, CA, in 2005, and the Ph.D. degree in applied and computational mathematics from the California Institute of Technology, Pasadena, in 2011.

He is currently a postdoctoral researcher hosted at the Laboratoire Jacques-Louis Lions in Paris 6, thanks to a fellowship by Foundation Sciences Mathnueur from the Ecole Polytechnique in 1993. His research interests are in computational harmonic analysis, mathematical signal processing, information theory, mathematical optimization, statistical estimation and detection with applications to the imaging sciences, scientific computing, and inverse problems.

Eric Nakamura received the B.S., M.S., and Ph.D. in electrical engineering, all from the University of Hawaii, Manoa.

He joined Northrop Grumman Aerospace Systems (NGAS)—Mixed Signal and Power Center in 2002 and is the current principle investigator for the NGAS A-to-I program.

Michael Grant (M'10) received the B.S. degree in electrical and computer engineering from the University of Texas, Austin, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA.

He is a principal at CVX Research, Inc., and the primary developer of CVX, a popular modeling framework for disciplined convex programming. The research described in this paper was conducted while he served as Staff Scientist in the Department of Applied and Computational Mathematics at the California Institute of Technology. He has served as a Guest Lecturer at the University of Texas at Austin, as a Consulting Assistant Professor in the Information Systems Laboratory at Stanford University, and as a Research Associate in the Department of Energy Resources Engineering at Stanford University.

Emilio Sovero (SM'05) received the B.S. degree in engineering, the M.S. degree in mechanical engineering, and the Ph.D. degree in applied physics and information science, all from the California Institute of Technology, Pasadena.

He is currently VP of Technology at Waveconnex Inc. (October 2011). Prior to his current position, in 2006 he joined Northrop Grumman Aerospace Systems (NGAS)—Mixed Signal and Power Center as Chief Scientist. He was the Program Manager at NGAS of the A-to-I development.

Daniel Ching received the B.S. and M.S. degrees in electrical engineering, both from the University of Hawaii, Manoa.

He is currently Circuit Design Engineer at Northrop Grumman Aerospace Systems (1996 to present). At Northrop, he has worked on high-speed mixed-signal designs which ranged from PCB level down to custom transistor level IC design.

Juhwan Yoo (M'12) received the B.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 2006 and 2012, respectively.

He is currently a postdoctoral researcher in the Department of Electrical and Computer Engineering at Duke University, Durham, NC. During the summer of 2004 he held a summer internship at Finisar Corporation, where he worked on multi-mode optical transceivers. His current research interests include: mixed-signal circuit design, analog signal processing, and the application of dimensionality reduction techniques to hardware signal processing platforms.

Mr. Yoo was the recipient of the Analog Devices Outstanding Student Designer award in 2009.

Justin Romberg (SM'10) received the B.S.E.E., M.S., and Ph.D. degrees from Rice University, Houston, TX, in 1997, 1999, and 2004, respectively.

He is an Associate Professor in the School of Electrical and Computer Engineering at the Georgia Institute of Technology, Atlanta. From Fall 2003 until Fall 2006, he was a Postdoctoral Scholar in Applied and Computational Mathematics at the California Institute of Technology, Pasadena. He spent the Summer of 2000 as a researcher at Xerox PARC, the Fall of 2003 as a visitor at the Laboratoire Jacques-Louis Lions in Paris, and the Fall of 2004 as a Fellow at UCLA's Institute for Pure and Applied Mathematics. In the Fall of 2006, he joined the Electrical and Computer Engineering faculty as a member of the Center for Signal and Image Processing. During 2006 to 2007, he was a consultant for the TV show "Numb3rs".

Dr. Romberg received an ONR Young Investigator Award, in 2008, a PECASE award and a Packard Fellowship, in 2009, and he was named a Rice University Outstanding Young Engineering Alumnus, in 2010. From 2008 to 2011, he was an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY.

Azita Emami-Neyestanak (S'97–M'04) was born in Naein, Iran. She received the B.S. degree with honors from Sharif University of Technology, Tehran, Iran, in 1996, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1999 and 2004, respectively.

She is currently an Assistant Professor of Electrical Engineering at the California Institute of Technology, Pasadena. From July 2006 to August 2007 she was with Columbia University, New York, as an Assistant Professor in the Department of Electrical Engineering. She also worked as a Research Staff Member at IBM T. J. Watson Research Center, Yorktown Heights, NY, from 2004 to 2006. Her current research areas are high-performance mixed-signal integrated circuits and VLSI systems, with the focus on high-speed and low-power optical and electrical interconnects, clocking, biomedical implant, and compressed sensing.

Emmanuel Candès (M'12) received the Diplôme d'Ingénieur from the Ecole Polytechnique, in 1993, and the Ph.D. degree in statistics from Stanford University, Stanford, CA, in 1998.

He is the Simons Chair in Mathematics and Statistics, and Professor of Electrical Engineering (by courtesy) at Stanford University, Stanford, CA. Prior to this, he was the Ronald and Maxine Linde Professor of Applied and Computational Mathematics at the California Institute of Technology, Pasadena. His research interests are in computational harmonic analysis, mathematical signal processing, information theory, mathematical optimization, statistical estimation and detection with applications to the imaging sciences, scientific computing, and inverse problems. Prof. Candès received numerous awards, most notably the 2006 Alan T. Waterman Medal, which is the highest honor bestowed by the National Science Foundation and which recognizes the achievements of scientists who are no older than 35, or not more than seven years beyond their doctorate. Other awards include an Alfred Sloan Fellowship in 2001, the Popov Prize in Approximation Theory in 2001, the 2005 James H. Wilkinson Prize in Numerical Analysis and Scientific Computing awarded by the Society of Industrial and Applied Mathematics (SIAM), the 2008 Information Theory Society Paper Award, the 2010 George Polya Prize awarded by SIAM. In 2011, he received the Collatz Prize awarded every four years by the International Council for Industrial and Applied Mathematics (ICIAM) established to "provide international recognition to individual scientists under 42 years of age for outstanding work on industrial and applied mathematics".