A Mixer Front End for a Four-Channel Modulated Wideband Converter With 62-dB Blocker Rejection

Douglas Adams, Yonina C. Eldar, *Fellow, IEEE*, and Boris Murmann, *Fellow, IEEE*

***Abstract*— The modulated wideband converter receiver architecture leverages compressive sensing techniques to improve flexibility for cognitive radio applications. We present a prototype integrated circuit that adds signal reception to previously demonstrated signal detection. By refactoring the mixing sequence between signal detection and reception, we enable targeted reception and blocker rejection. We algorithmically design a three-level mixing sequence and additionally employ delay-based harmonic cancellation. When applied together in our 65-nm proof-of-concept integrated circuit, we measure 62 dB of in-band blocker rejection, while receiving up to four channels with independently defined locations anywhere up to 900 MHz.**

***Index Terms*— Cognitive radio, harmonic rejection mixer, modulated wideband converter (MWC).**

I. Introduction

FREQUENCY-AGILE radios are an active area of research and receivers based on compressive sensing (CS) show promise for future systems [1]–[4]. The modulated wideband converter (MWC) architecture, introduced in [5] and shown in Fig. 1, downmixes the entire sparse input spectrum to baseband using a spectrally diverse mixing sequence. By making multiple linearly independent measurements with parallel hardware branches, the entire input spectrum can be reconstructed. Among the different CS receiver variants, the MWC is architecturally close to a bank of direct conversion receivers and is, hence, attractive for CMOS realization. In [6] and [7], a variant of the MWC was employed for wideband spectrum sensing and interference detection. The MWC is particularly attractive for cognitive radio applications, because, unlike conventional receivers, it can dynamically adjust the number of bands it receives, the width of those bands, and its performance is independent of their locations [9]. In this paper, we report a front end for the MWC that supports signal reception for up to four channels. A key feature of the presented front end is that, unlike previous realizations, it can attenuate a single in-band blocker by 62.8 dB, or two blockers by 50.2 dB. As described in the following, we achieve this through two complementary techniques that may be used in concert or applied independently—algorithmic three-level digital mixing sequence design and a parallel mixing path that provides harmonic cancellation. The latter is enabled by an online calibrated, digital-to-delay converter (DDC) with a step size of 1.2 ps. This paper extends that presented in [9] with a detailed description of the implemented sequence-design algorithm, analysis of the delay-based rejection scheme, and computation of the system sensitivity, which provides performance limits and motivates the circuit design.

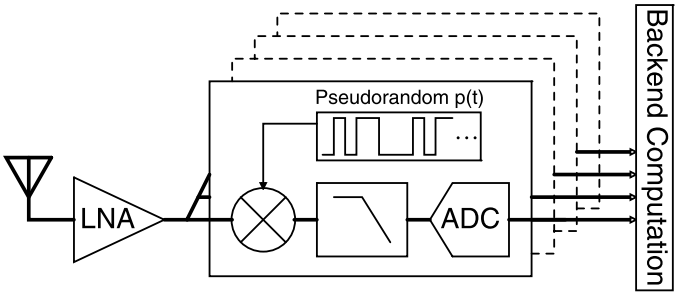

Fig. 1. MWC architecture as described in [5].

In Section II, we introduce the dichotomy between detection and reception in the MWC and leverage it to develop both an algorithm for digital mixing sequence design and a delay-based harmonic cancellation technique that together provide harmonic rejection and sensitivity improvements to our prototype. In Section III, we develop the mathematical framework necessary to estimate the sensitivity of the MWC and describe how we utilize that framework to optimize our design. Section IV describes the details of our circuit implementation, while Section V provides our measurement results. We conclude in Section VI.

II. Harmonic Rejection in the MWC

A. Detection Versus Reception

The MWC architecture, introduced in [5], is shown in Fig. 1. All previous work has used a spectrally rich pseudorandom sequence implemented with a digital mixing sequence for both detection and reception to guarantee that power from each input band is mixed to baseband. However, a pseudorandom sequence provides no blocker rejection (see Fig. 2) and incurs a severe noise penalty, since the effective noise bandwidth spans the entire spectrum. We assume that the input spectral support (both signal and blocker-band locations) is constant on the timescale of interest, so that once the frequency support is known, that information may be used to optimize the mixing sequence and receiver configuration to improve performance. Specifically, by maximizing the mixer gain in the signal bands and nulling the mixer gain in a blocker band,

Manuscript received August 17, 2016; revised October 23, 2016 and December 21, 2016; accepted December 22, 2016. Date of publication April 12, 2017; date of current version April 20, 2017. This paper was approved by Guest Editor Danilo Manstretta.

D. Adams was with Stanford University, Stanford, CA 94305 USA (e-mail: doadams@gmail.com).

Y. C. Eldar is with the Technion–Israel Institute of Technology, Haifa 32000, Israel (e-mail: yonina@ee.technion.ac.il).

B. Murmann is with Stanford University, Stanford, CA 94305 USA (e-mail: murmann@stanford.edu).

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/JSSC.2017.2647941

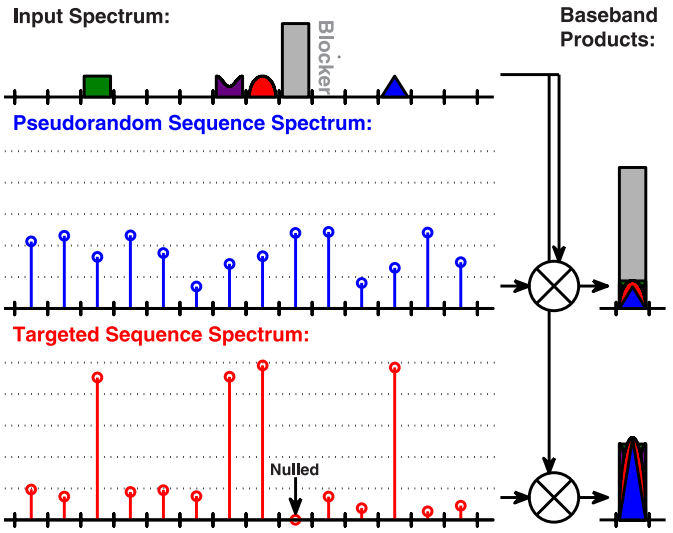

Fig. 2. Pseudorandom versus targeted mixing sequence (which provides blocker rejection and noise attenuation).

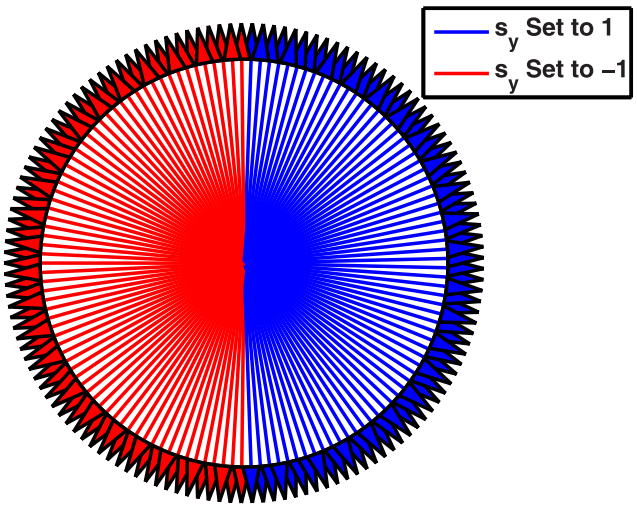

Fig. 3. Row F_x plotted in the complex plane.

we simultaneously provide blocker rejection and improve the signal-to-noise ratio (Fig. 2).

The detection step, which determines both the signal and blocker locations, is demonstrated theoretically in [5] and [8], and practically in [6], [7], and [10]. In addition, [11] develops an algorithm that greatly reduces the number of branches required for detection under the assumption that the signal band locations are stationary (which applies here). We introduce an algorithm that produces targeted mixing sequences that also rejects a single known blocker. In addition, we introduce an independent, delay-based harmonic rejection technique that can be used to either further null the same blocker, or additionally null a separate blocker.

B. Digital Mixing Sequence Design

In each branch of the MWC, we implement a time-shifted copy of the same mixing signal $p(t)$. This simplifies the hardware, and, as shown in [5], does not lead to a loss of generality. We constrain $p(t)$ to a digitally implementable piecewise constant signal, which allows us to fully describe $p(t)$ with both the discrete mixing sequence $\boldsymbol{s} = [s_1, s_2, \ldots, s_L]^T$ and the digital clock frequency. The period T_p of $p(t)$ sets the separation of the mixer harmonics in the frequency domain, which, in turn, specifies the channel bandwidth of our receiver. In addition, as is shown in [11], adjusting the sequence elements only provides independent control of the mixer harmonics up to half of the digital clock frequency, thus requiring our digital clock to operate above twice the maximum input frequency of our system. The sequence length L is, therefore, set by the desired channel bandwidth and the highest frequency of interest. In our design, we target channels that are 1.4 MHz wide at up to 900 MHz, and, therefore, use a 1332-length sequence running on a 1.85-GHz clock. Recovering higher input frequencies is difficult for the MWC, because both clock frequency and sequence length scale with input frequency. However, higher frequency inputs can be received (as in [6] and [7]) by adding an IF mixer, albeit at the cost of requiring effective image rejection.

The Fourier series coefficients $\boldsymbol{c} = [c_1, c_2, \ldots, c_N]^T$ of $p(t)$ in the frequency range of interest are uniquely defined by the mixing sequence elements, and can be efficiently computed with the matrix equation $\boldsymbol{c} = \boldsymbol{F}\boldsymbol{s}$, where the elements $F_{x,y}$ are precomputed via direct application of the Fourier series integral, such that

$$F_{x,y} = \frac{j}{2\pi x}\left[e^{-2\pi j x \frac{y}{T_p f_{clk}}} - e^{-2\pi j x \frac{y-1}{T_p f_{clk}}}\right]. \tag{1}$$

Each $F_{x,y}$ represents the complex contribution of sequence element s_y onto Fourier coefficient c_x. We use the information contained in the Fourier series computation matrix $\boldsymbol{F}$ to drive our sequence selection algorithm. In describing our algorithm, we introduce a simplified version and then elaborate upon it.

To begin, suppose we wish to design a binary mixing sequence that maximizes the mixer harmonic $|c_x|$ corresponding to a single desired frequency band. The value of c_x is computed via the dot product $F_x\boldsymbol{s}$, where F_x corresponds to the xth row of $\boldsymbol{F}$. Maximizing $|c_x|$ is thus equivalent to choosing the elements of $\boldsymbol{s}$, so that each $F_{x,y}s_y$ component adds constructively. Fig. 3 shows the elements of F_x graphically in the complex plane, from which we see that choosing $s_y = \text{sgn}(\text{Re}\{F_{x,y}\})$ guarantees each term in the dot product will have a positive real component and the composite c_x will have some large real component. Since the elements of F_x are uniformly distributed in phase, bisecting along the imaginary axis and choosing s_y according to the real component of $F_{x,y}$ maximize the amplitude of c_x. In Fig. 4, we plot $p(t)$ in the time domain, to illustrate that defining $p(t)$ as such is equivalent to quantizing a sinusoid at the desired frequency to single-bit precision.

We extend our sequence design to target multiple Fourier coefficients by generating a composite vector

$$G = \frac{F_{d_1}}{\text{sinc}\left(\frac{\pi d_1}{L}\right)} + \frac{F_{d_2}}{\text{sinc}\left(\frac{\pi d_2}{L}\right)} + \ldots \tag{2}$$

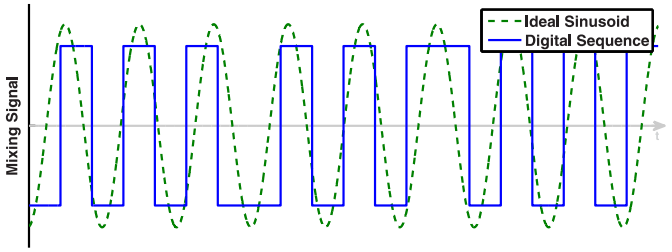

Fig. 4. Single-target mixing sequence $p(t)$ plotted in the time domain.

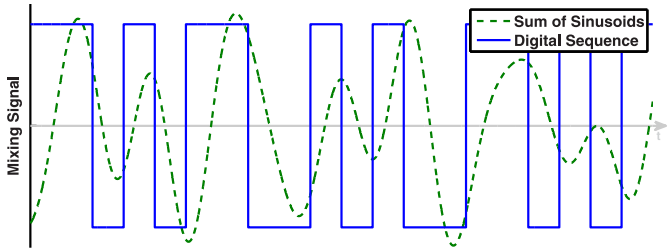

Fig. 5. Multitarget mixing sequence $p(t)$ plotted in the time domain.

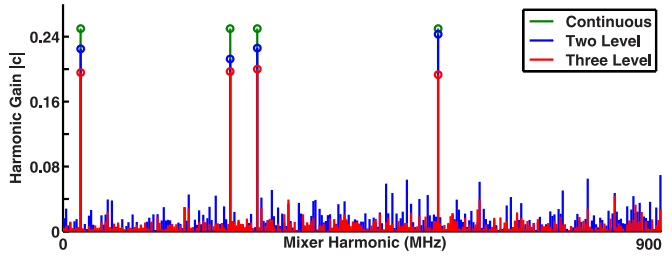

Fig. 6. Harmonic gains $|\boldsymbol{c}|$ for the two-level, three-level, and continuous ideal mixing sequences.

where we normalize the magnitude of the individual rows via a sinc function. This normalization is a result of the clocked, piecewise nature of the mixing signal. Here, d_i indexes the set of the desired Fourier coefficients. We then choose $s_y = \text{sgn}(\text{Re}\{G_y\})$ to produce a sequence that approximately maximizes $|c_{d_1} + c_{d_2} + \ldots|$. Fig. 5 shows the multitarget $p(t)$ in the time domain, and once again, we see that our definition of $p(t)$ is equivalent to quantizing the sum of a set of sinusoids with the desired frequencies. Fig. 6 shows that time discretization and quantization combine to produce lots of spurious harmonics that fold both noise and blockers into our baseband measurement.

To reduce the power in these spurs, we add an additional quantization step to decrease the quantization noise. Thus, we define the mixing sequence according to

$$s_y = \begin{cases} 1 & \text{sgn}(\text{Re}\{G_y\}) > T \\ -1 & \text{sgn}(\text{Re}\{G_y\}) < T \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where the quantization threshold T is chosen to maximize the system sensitivity. Optimizing the value of T is discussed in Section III. Fig. 6 shows the magnitudes $|\boldsymbol{c}|$ for a two-level mixing sequence, three-level mixing sequence, and the continuous ideal mixing sequence. As expected, increasing the number of quantization levels increases the power in the desired harmonics relative to the background spurs.

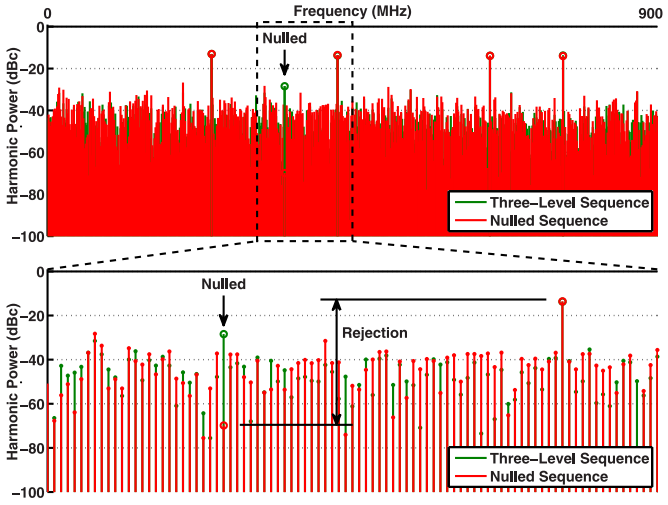

Fig. 7. Log-scale representation of $|c|$ before and after the nulling algorithm is applied.

In addition to increasing sensitivity, adding the third, zero level to the mixing sequence provides an additional degree of freedom that we can use to null any undesired harmonic. In effect, the selection process described in (3) identifies those sequence elements that do not strongly contribute to the desired harmonics and sets them to zero. Therefore, toggling them to either ± 1 will not strongly affect the desired harmonics, but can have a strong impact on the undesired spurs. To null an undesired harmonic, we project $F_{u,y}$ onto c_u (where u indexes the undesired harmonic), sort those projections by magnitude, toggle the zero-level sequence element with the largest projection to either ± 1 so as to decrease $|c_u|$, and iterate.

Fig. 7 shows a log-scale representation of $|\boldsymbol{c}|$ for the three-level mixing sequence before and after the undesired harmonic is nulled. We see that although the average background spur lies only 30 dB below the minimum desired harmonic, the nulling algorithm rejects the unwanted harmonic by more than 50 dB.

The choice of signal and blocker-band locations affects the finite set of available zero-level sequence elements, and thus also affects the blocker rejection that our algorithm can achieve. To demonstrate consistency of performance in spite of this limitation, we run a 5000 iteration Monte Carlo sweep, where the signal and blocker-band locations are randomly selected, and calculate the minimum mixer gain and blocker rejection for each run. Fig. 8 shows that the algorithm is able to achieve greater than 50 dB of blocker rejection for over 90% of the test cases. Note that a nulled mixer harmonic will translate no power from the corresponding signal band to baseband, regardless of signal shape.

C. Delay-Based Harmonic Rejection

In addition to the harmonic rejection achieved through digital sequence design, we introduce a generalization of a power inverter harmonic-rejection technique [12] that uses

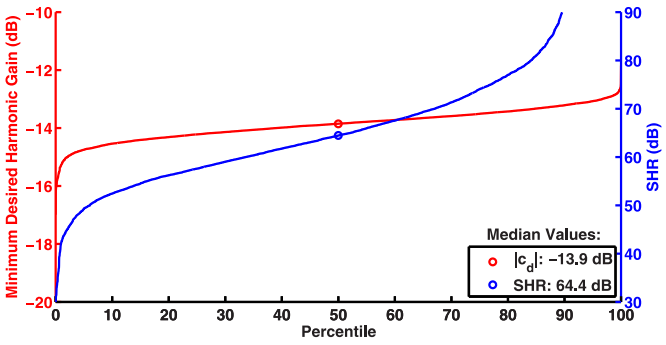

Fig. 8. Performance (harmonic rejection and gain) versus randomized signal-band locations.

the precise delay control achievable in CMOS processes to null any mixer harmonic, independent of the digital mixing sequence.

For any periodic signal with angular frequency ω_p, applying a delay τ will rotate the phase of the nth harmonic by $n\omega_p\tau$ radians. If we choose τ, such that the nth harmonic is rotated by $k\pi$ radians (where k is any odd number), then adding it to the original will result in a signal with no nth harmonic. Thus, we choose τ, such that

$$n\omega_p\tau = k\pi. \tag{4}$$

Adding a delayed copy of any periodic signal to itself modifies the original periodic signal's Fourier coefficients according to

$$c_x' = c_x(1 + e^{-jx\omega_p\tau}) \tag{5}$$

which gives a normalized per-harmonic attenuation of

$$\frac{1}{2}\left|\frac{c_x'}{c_x}\right| = \left|\cos\left(\frac{\pi x k}{2n}\right)\right|. \tag{6}$$

As is apparent from Fig. 9, this harmonic cancellation technique applies a comb filter to all of the harmonics, both desired and undesired, so that it is possible to attenuate the desired signal bands as well as the undesired blocker. However, in choosing τ, both k and n represent degrees of freedom—n sets the harmonic that we wish to cancel, and k determines the bandwidth of the notches and the frequency at which they occur. We, therefore, choose $k \in \{1 \text{ to } n, \text{odd}$ to minimize the attenuation of the desired harmonics (for $k > n$, the shape of the comb filter repeats). Once again, the achievable harmonic rejection depends on the signal and blocker-band locations. Fig. 10 shows a 5000 iteration Monte Carlo sweep, where the signal and blocker bands are chosen at random. For each iteration, n is chosen to null the blocker band, and k is chosen to maximize the achievable harmonic rejection. It is clear that we can choose a comb filter that avoids attenuating the desired harmonics with high probability.

Introducing a delay offset error ϵ into (6), we have

$$\frac{1}{2}\left|\frac{c_x'}{c_x}\right|\Bigg|_{x=n} = \left|\sin\left(\frac{n\omega_p\epsilon}{2}\right)\right| \tag{7}$$

from which we determine that in order to guarantee 50 dB of attenuation, we must specify the delay to within 0.56 ps. This, in turn, dictates the resolution of the DDC. The circuit

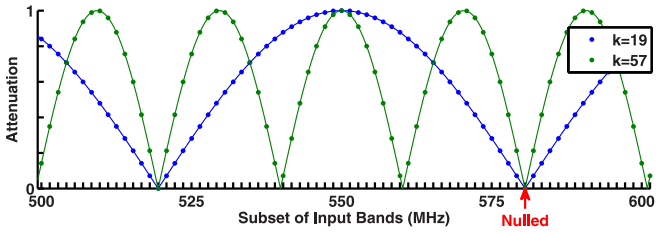

Fig. 9. Harmonic attenuation according to (6) for (a) $k = 19$ and (b) $k = 57$.

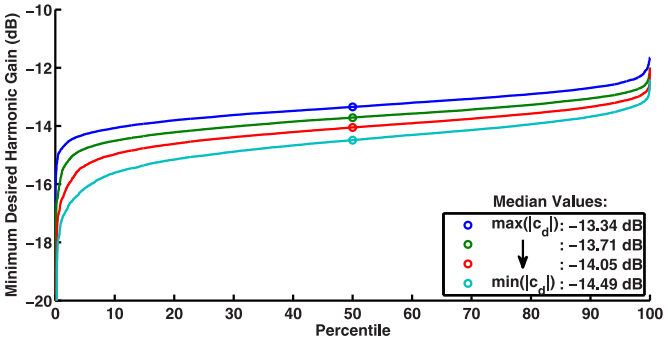

Fig. 10. Harmonic gain versus randomized signal-band locations.

that we use to provide this resolution, as well as the calibration algorithm necessary to guarantee its accuracy, is described in Section IV.

The maximum possible sequence-based harmonic rejection is limited by the finite number of ways in which the zero-level sequence elements can be adjusted to null an undesired harmonic. Similarly, delay-based harmonic rejection is limited by mismatch between the two signal paths. In addition, both techniques are sensitive to jitter, since as the jitter increases, the accuracy to which the harmonics are known decreases. Furthermore, at low jitter, when we apply both techniques to the same blocker, they compound directly, although this combined rejection suffers the same jitter-to-rejection limit as the individual techniques. This behavior is verified by simulation, as shown in Fig. 11, and matches the measured results as well (also shown).

III. System Sensitivity

The sequence generation algorithm and delay-based harmonic rejection scheme described previously are designed to attenuate a single strong blocker, but additionally provide rejection against the wideband noise that our digital mixing sequence folds down to baseband. In this section, we develop the mathematical framework to analyze the noise performance of the multibranch, multiband, MWC architecture and use that understanding to maximize the sensitivity of the MWC. To do so, we begin with the noise folding present in each individual branch and then evaluate the noise averaging achievable in this multibranch architecture.

Fig. 12 shows a simplified signal-path schematic of a single branch of the MWC and highlights the noise sources inherent to each stage. From this noise model, we calculate the total output referred noise power N_T as

$$N_T = \left(\left((N_A + N_L)A_L^2 + N_M\right)G_M^2|c_d|^2k + N_F\right)Z_F^2 + N_{\text{ADC}} \tag{8}$$

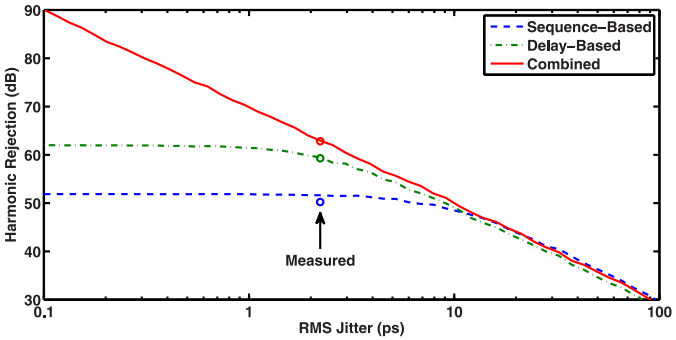

Fig. 11. Simulated HR versus jitter (with 0.3% delay path mismatch).

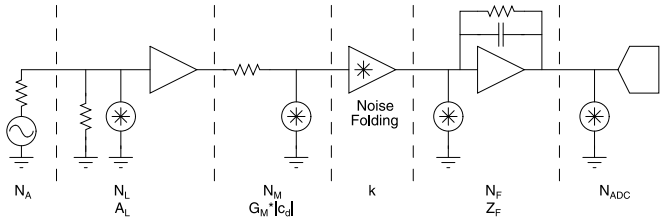

Fig. 12. Simplified schematic of the signal path for a single branch of the MWC highlighting individual noise contributions.

for individual block noises N_A, N_L, N_M, N_F, and N_{ADC}, and individual block gains A_L, $G_M|c_d|$, and Z_F. For our mixer, the signal gain and noise gain are not identical. Specifically, while signal from each desired band is mixed to baseband with gain $|c_d|$, noise is mixed from these similar bands and every other band in the input range as well. To encapsulate the additional noise folded to baseband by the background harmonics, we introduce the noise folding parameter k, which is given by

$$k = \frac{\sum |c|^2}{|c_d|^2} \tag{9}$$

and note that this noise folding multiplies all upstream noise sources.

Parseval's theorem states that $\sum |c|^2 = (1/T)\int_T |p(t)|^2 dt$, which allows us to quickly compute k for any mixing sequence, since $p(t) \in \{-1, 0, 1\}$. Fig. 13 shows the results of a Monte Carlo simulation, where for 5000 sets of possible signal-band locations, we plot the average noise folding parameter against the zero-value threshold. We thus choose the zero-value threshold in our sequence-design algorithm (3) to minimize the noise folding parameter k from these curves according to the number of desired signal bands. Intuitively, a minimum in the noise folding factor must exist, because as we increase the zero-value threshold, the first sequence elements set to zero will have a disproportionate effect on the undesired harmonics, and conversely, the last sequence elements set to zero will have a disproportionate effect on the desired harmonics. In addition, as is evident, increasing the number of target bands makes it more difficult for the algorithm to differentiate between sequence elements that contribute to the desired and the spurious harmonics, so the noise folding factor increases accordingly. For comparison, the noise folding factor of a pseudorandom sequence in this scenario is 70 dB, independent of target bands.

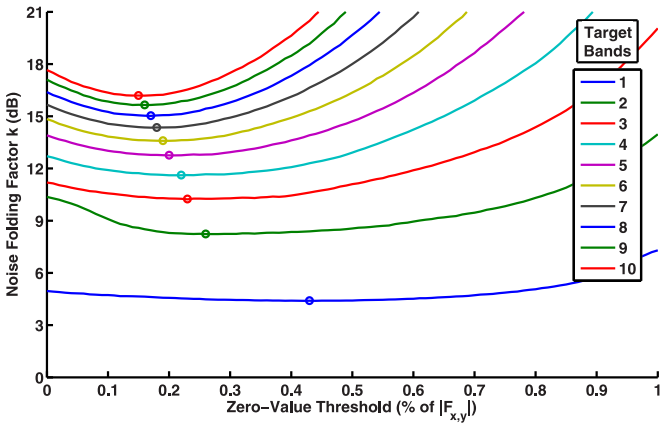

Fig. 13. Noise folding factor k and the number of branches versus the zero-value threshold.

In our implementation of the MWC, each branch utilizes an identical mixing sequence subjected to a branch-dependent relative delay that guarantees the linear independence of the baseband measurements [5]. We can, therefore, express the mixing behavior of the MWC according to $\boldsymbol{Y} = \boldsymbol{CX}$, where we have

$$\boldsymbol{C} = A_L G_m Z_F \begin{bmatrix} c_{d_1} e^{j\omega_1\tau_1} & c_{d_2} e^{j\omega_2\tau_1} & \cdots & c_{d_M} e^{j\omega_M\tau_1} \\ c_{d_1} e^{j\omega_1\tau_2} & c_{d_2} e^{j\omega_2\tau_2} & \cdots & c_{d_M} e^{j\omega_M\tau_2} \\ \vdots & \vdots & \ddots & \vdots \\ c_{d_1} e^{j\omega_1\tau_M} & c_{d_2} e^{j\omega_2\tau_M} & \cdots & c_{d_M} e^{j\omega_M\tau_M} \end{bmatrix} \tag{10}$$

where M represents both the number of signal bands and the number of branches in our single-sideband implementation of the MWC.

In the presence of additive noise, we have $\boldsymbol{Y} = A_L G_m Z_F \boldsymbol{CX} + \boldsymbol{N}$, with the noise power in each term of $\boldsymbol{N}$ given by (8) and the elements of $\boldsymbol{N}$ assumed to be zero mean and mutually independent. The error this noise introduces into our estimate of the original signal is given by

$$\hat{X} - \boldsymbol{X} = \frac{1}{A_L G_m Z_F} \boldsymbol{C}^{-1} \boldsymbol{N} \tag{11}$$

and the estimated error covariance is given by

$$E[(\hat{X} - \boldsymbol{X})(\hat{X} - \boldsymbol{X})^T] = \frac{1}{(A_L G_m Z_F)^2} \boldsymbol{C}^{-1} E[\boldsymbol{N}\boldsymbol{N}^T] \boldsymbol{C}^{-T}. \tag{12}$$

For independent noise sources $E[N_i N_{j\neq i}] = 0$ and $E[\boldsymbol{N}\boldsymbol{N}^T] = N_T \boldsymbol{I}$, so that (12) simplifies to

$$E[(\hat{X} - \boldsymbol{X})(\hat{X} - \boldsymbol{X})^T] = \frac{N_T}{(A_L G_m Z_F)^2} \boldsymbol{C}^{-1} \boldsymbol{C}^{-T}. \tag{13}$$

As in [11], we choose $\tau_{(m)}$ so as to minimize the power of the noise induced error. Under the assumption that each c_d has roughly equal amplitude, this occurs when $(1/(\sqrt{M}|c_d|))\boldsymbol{C}$ approaches a unitary matrix. In this case, we have $\boldsymbol{C}^{-1}\boldsymbol{C}^{-T} \approx (1/(M|c_d|^2))$. Combining (8) and (13) provides

an overall system sensitivity given by

$$\text{Sensitivity} \approx \frac{\left((N_A + N_L)A_L^2 + N_M\right)G_m^2 k|c_d|^2 + N_F\right)Z_F^2 + N_{\text{ADC}}}{M|c_d|^2 A_L^2 G_m^2 Z_F^2} W \tag{14}$$

where W is the channel bandwidth.

IV. Circuit Implementation

We designed a prototype IC that incorporates both sequence and delay-based harmonic rejection techniques and is optimized for its sensitivity in order to demonstrate the ability of the MWC to compete with traditional receiver architectures. To simplify the design, we limited the scope of our chip to the critical core circuitry, which for the MWC is the mixer. Therefore, we implement data storage for the mixing sequence, the DDC, the mixer itself, and its subsequent antialiasing filter on chip, and implement the LNA, input balun, and ADC externally.

A. Mixing Sequence Data Storage

Each mixing sequence is computed off-chip and stored on-chip in two closed-loop shift registers of length 1332, which run at 1.85 GHz. We implement the local storage with shift registers to simplify design and facilitate debugging; however, as each shift register stores the same data, future implementations may share a single SRAM to save power and area.

Our chip is designed to recover one to four signals between 0 and 900 MHz. The channel bandwidth is dictated by the period of the mixing sequence, which for our implementation is achieved by adjusting the effective length of the mixing sequence; 1.4 MHz with a 1332-element mixing sequence, 2.8 MHz with a 666-element mixing sequence, and 4.2 MHz with a 444-element mixing sequence. Of note, as the sequence length decreases so does the efficacy of the sequence-design algorithm, limiting the useful bandwidth modes.

B. DDC

The optimum branch-to-branch delay may lie anywhere within the entire period of the mixing sequence, and in order to achieve 50-dB rejection, must be specified to within 0.56-ps accuracy. To achieve this dynamic range, we implement the DDC with a coarse-fine architecture. Coarse delay control is provided by delaying the sequence the appropriate number of clock cycles within the shift register. Fine control is implemented in a digital delay line with adjustable load capacitors [see Fig. 14(a)]. The DDC capacitors are implemented with deep n-well MOS devices to minimize stray capacitance (and thus jitter) and the total extracted per-stage capacitance ranges between 470 aF and 28.2 fF.

C. Mixer Implementation

Switches used for calibration and the three-level current-mode passive mixer are combined into a single circuit block

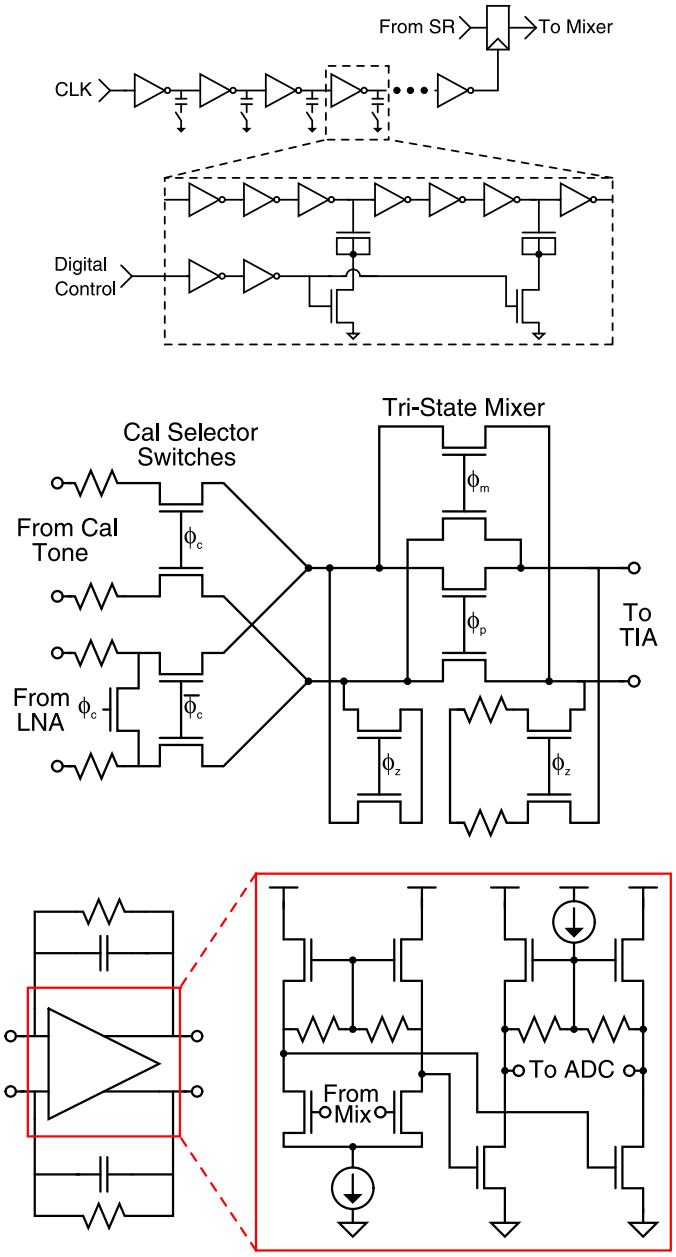

Fig. 14. Simplified schematics of key circuit blocks. (a) DDC. (b) Mixer. (c) Antialias filter.

shown in Fig. 14(b). Together the input resistors provide a 50-Ω impedance match and provide voltage-to-current conversion for the passive mixer. We implement the mixer's zero level by adding switches that tie both the input and output ports of the mixer to common mode, such that there is no transmission across the mixer. Furthermore, we add additional resistors to this mode, so that both mixer ports see a first-order impedance match regardless of mixer state. The input resistors and feedback resistors are sized for an expected 0.3% mismatch.

D. Antialiasing Filter

The filter is implemented as an active RC design using a two-stage amplifier [see Fig. 14(c)]. The op-amp is designed

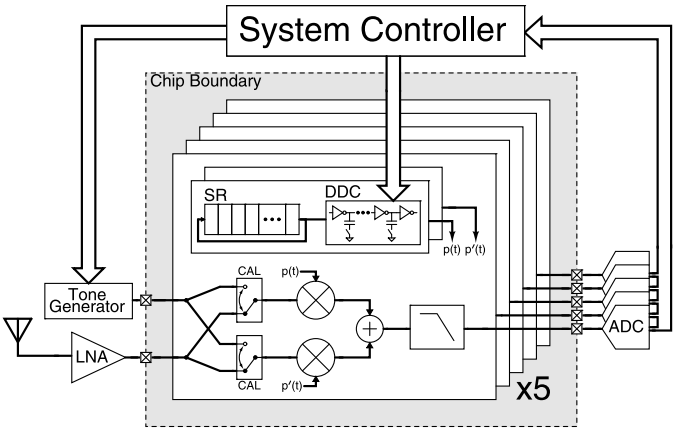

Fig. 15. Test setup and chip architecture.

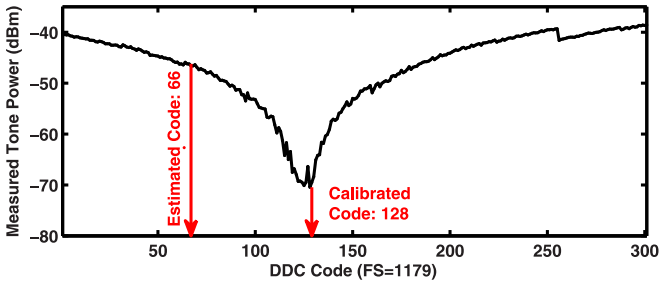

Fig. 16. DDC calibration sweep.

for high gain (500), because the virtual ground is used as the summing node for the two signal paths. The bandwidth is set to 5.6 MHz to allow for the various bandwidth modes of our MWC implementation.

E. Calibration

As shown in Fig. 15, our four-channel design contains five identical branches, so that one branch can be pulled offline for calibration. In order to calibrate a specific active branch, the extra calibration branch matches the delay of its mixing sequence to that of the active branch. After which, calibration switches are toggled so that the calibration branch substitutes for the active branch. The formerly active branch is then fed with a sinusoid generated off-chip that lies in the blocker band. Next, the DDC code is swept and the baseband power of the calibration tone is measured. The DDC is set to the code that minimizes that power and the now-calibrated primary branch may be switched back into the signal path. Calibration then proceeds with the next branch. Since the delay-based cancellation rejects all odd multiples of the target harmonic, future work may utilize an on-chip square-wave generator as the calibration tone.

V. Measurement Results

Fig. 16 shows the calibration tone power near the optimal digital delay. The difference in measured tone power between the expected code and calibrated code demonstrates both the need for calibration and its efficacy. The measured harmonic rejection is shown in Fig. 17, wherein we apply two equal-power test tones, one in a blocker band and one in a signal

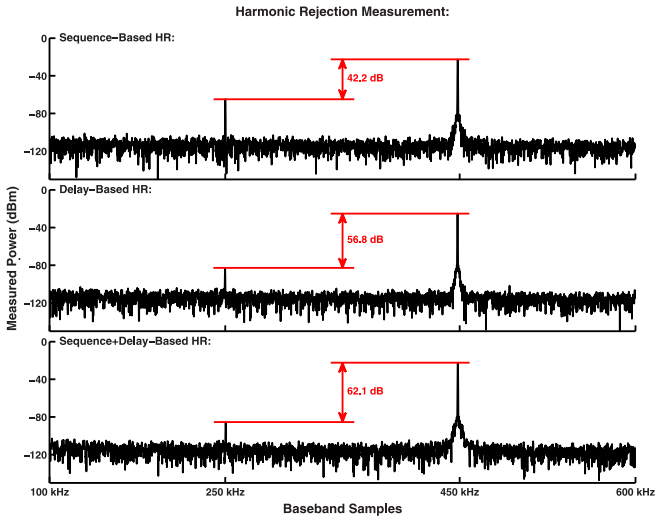

Fig. 17. Harmonic rejection measurement.

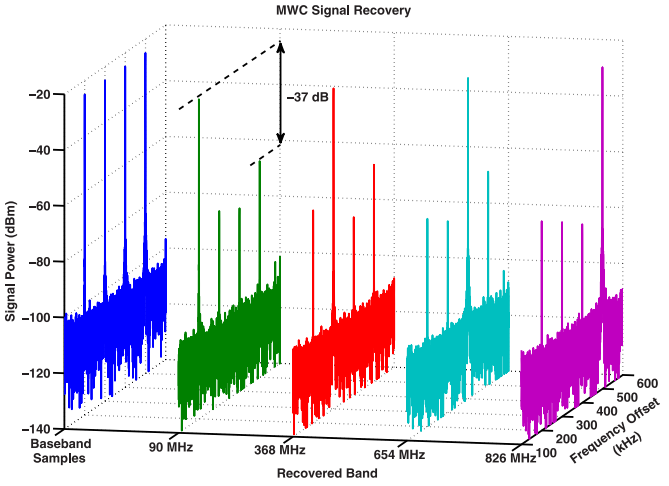

Fig. 18. Measured recovery of four input signals from four independent baseband measurements.

band (in this case 878 and 881 MHz, respectively). At baseband, we measure the downconverted power for both tones—the difference corresponds to the achieved harmonic rejection. We demonstrate 50.2 dB of sequence-based rejection, 59.3 dB of delay-based rejection, and 62.8 dB of rejection with both applied to the same blocker. Furthermore, we note that the measured rejection for each scheme closely matches the respective simulated values 51.4, 59.5, and 62.9 dB shown in Fig. 11 with an estimate of 2.2-ps rms clock jitter. For our implementation, calibration time was limited by the digital I/O speed, and the algorithm was not optimized for speed.

Fig. 18 shows the back-end signal recovery of the MWC in the 1.4-MHz configuration. We simultaneously drive the MWC with four sinusoids in bands spread across the input spectrum and recover those bands (and their respective sinusoids) from the baseband samples concurrently measured in each branch. By appropriately setting the branch-to-branch mixing sequence delay, we ensure linearly independent baseband measurements. We recover the input signal by inverting a matrix of complex mixer harmonics and multiplying it with the frequency-domain

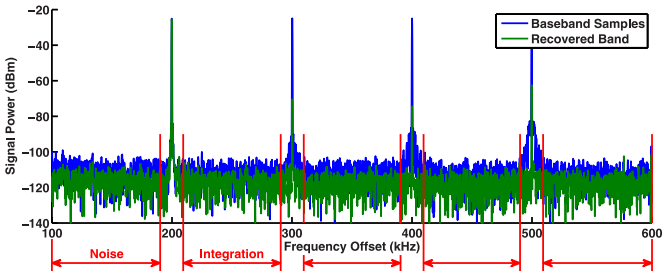

Fig. 19. Demonstrated noise averaging. The integrated noise power from a single ADC's measurements is 6 dB higher than that of a recovered band.

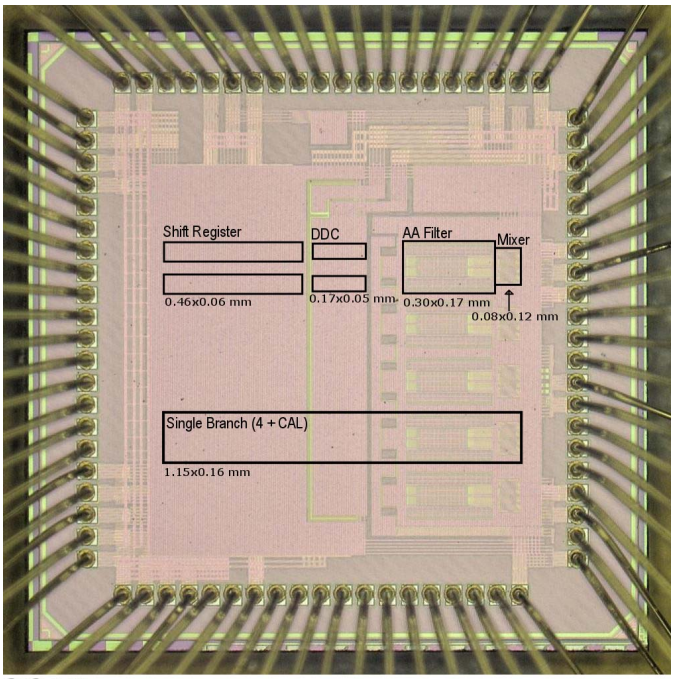

Fig. 20. Die photo.

baseband measurements [5]. Our system demonstrates 37 dB of band-to-band rejection, which is limited by the accuracy with which we estimate the complex mixer harmonics. Due to equipment limitations, signal recovery was limited to continuous wave signals, and BER measurements were infeasible. However, both harmonic rejection techniques attenuate the baseband image of an entire input band, so for spread-spectrum signals, the behavior of the MWC is not expected to change.

Our mixer shows 64.8 dBm of IIP2 and 14.3 dBm of IIP3. We measure the IIP2 by applying a 300-MHz tone to the input and using a mixing sequence designed to target both the 300- and 600-MHz bands, thus directly capturing the mixer's second harmonic. The IIP2 is calculated by extrapolating the input power at which the second harmonic equals the fundamental. Calculating the IIP2 this way has the added benefit of obviating the IIP2 error introduced by any differential imbalance. We measure the IIP3 by applying two-tones separated by 150 kHz in the 500-MHz band. We measure a single-branch noise figure of 37 dB using direct application of the Y-factor method (which includes noise folding) and factoring out the

TABLE I
COMPARISON TO PRIOR WORK

		Rafi [13]	Sundström [14]	Yang [15]	**This Work**
System		HR IF mixer	Dual-carrier aggregation HR IF mixer	Wideband single-channel mixer for SDR	**Multi-channel mixer for spectrum sensing and reception**
Technology	nm	110	65	45	**65**
Area (Per Branch for this Work)	mm^2	0.034	0.48	0.352	**0.085**
Supply Voltage	V	2.7 (mixer) 1.3 (clock)	1.2	1	**1.2**
Min HR	dB	52	68	55	**62.8 (SHR)**
IIP2	dBm	>75	-	-2	**64.8 (300 MHz)**
IIP3	dBm	12	-	-3	**14.3 (500 MHz)**
NF	dB	11 (DSB)	-	35 (SSB)	**31 (37-6) (SSB)**
Signal Path Power (Per Branch for this Work)	mW	59.4	23.6	17	**22.8+3.8**
Frequency	MHz	100-300	390	500-1500	**0-900**

effects of peripheral devices using Frii's formula. However, the MWC makes an independent measurement of the incident signal in each of its branches, so with four active branches, averaging improves the noise figure by 6 dB. This is shown in Fig. 19.

Adding a hypothetical LNA with 2.5-dB NF, 0-dBm IIP3, and 26-dB gain, and applying cascading relationships yield an overall system with 10 dB of SNDR at −25-dBm input power and −96.5 dBm of sensitivity. As discussed in [11], these numbers include the additional wideband noise folding and wideband intermodulation introduced by the LNA. Accounting for the 3-dB SSB noise penalty (by doubling the mixing hardware) and a 0.5-dB increase in signal bandwidth, these results closely parallel the LTE specifications for 5-MHz bandwidth signals up to 900 MHz, and a survey of recent work demonstrates that these specifications are achievable for wideband LNAs [16], [17]. Furthermore, our mixer's performance is competitive with recently published architectures (see Table I) that do not demonstrate the flexibility of our design.

The measured per-branch power draws for the mixer + TIA, DDC, and shift register are 22.8, 3.8, and 149.5 mW, respectively, from a single 1.2 V supply. The entire chip occupies 3.84 mm^2 and is implemented in a 65-nm CMOS process. The per-branch active area is 0.0845 mm^2, split between the shift registers (0.0376 mm^2), DDCs (0.00317 mm^2), TIA (0.0402 mm^2), and mixer (0.00349 mm^2). The rest of the area is used for decoupling capacitors and guard rings. An annotated die photo is provided in Fig. 20.

VI. CONCLUSION

The MWC is a promising circuit architecture for cognitive radio applications, because it can quickly determine the spectrum characteristics [6], [7] and dynamically adjust the number and bandwidth of its reception channels across a wide frequency range [5]. In this paper, we implement both sequence-based and delay-based harmonic rejection to the MWC that provide blocker rejection and sensitivity comparable to traditional receiver architectures while maintaining the flexibility advantage of the system.

ACKNOWLEDGMENT

The authors would like to thank TSMC for chip fabrication through the University Shuttle Program.

REFERENCES

[1] J. Yoo, S. Becker, M. Loh, M. Monge, E. J. Candés, and A. Emami-Neyestanak, "A 100MHz–2GHz 12.5x sub-Nyquist rate receiver in 90nm CMOS," in *Proc. IEEE Radio Freq. Integr. Circuits Symp. (RFIC)*, Jun. 2012, pp. 31–34.

[2] X. Chen *et al.*, "A sub-Nyquist rate compressive sensing data acquisition front-end," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 3, pp. 542–551, Sep. 2012.

[3] D. Adams, C. S. Park, Y. C. Eldar, and B. Murmann, "Towards an integrated circuit design of a compresssed sampling wireless receiver," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2012, pp. 5305–5308.

[4] Y. C. Eldar, *Sampling Theory, Beyond Bandlimited Systems*. Cambridge, U.K.: Cambridge Univ. Press, Apr. 2015.

[5] M. Mishali and Y. C. Eldar, "From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 375–391, Apr. 2010.

[6] R. T. Yazicigil, T. Haque, M. R. Whalen, J. Yuan, J. Wright, and P. R. Kinget, "A 2.7-to-3.7GHz rapid interferer detector exploiting compressed sampling with a quadrature analog-to-information converter," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 1–3.

[7] R. T. Yazicigil, T. Haque, M. Kumar, J. Yuan, J. Wright, and P. R. Kinget, "A compressed-sampling time-segmented quadrature analog-to-information converter for wideband rapid detection of up to 6 interferers with adaptive thresholding," in *Proc. IEEE Radio Freq. Integr. Circuits Symp. (RFIC)*, May 2016, pp. 282–285.

[8] 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.

[9] D. Adams, Y. Eldar, and B. Murmann, "A mixer frontend for a four-channel modulated wideband converter with 62 dB blocker rejection," in *Proc. IEEE Radio Freq. Integr. Circuits Symp. (RFIC)*, May 2016, pp. 286–289.

[10] M. Mishali, Y. C. Eldar, O. Dounaevsky, and E. Shoshan, "Xampling: Analog to digital at sub-Nyquist rates," *IET Circuits, Devices Syst.*, vol. 5, no. 1, pp. 8–20, Jan. 2011.

[11] D. Adams, "A practical implementation of the modulated wideband converter compressive sensing receiver architecture," Ph.D. dissertation, Dept. Elect. Eng., Stanford Univ., Stanford, CA, USA, 2016.

[12] J. G. Kassakian, M. F. Schlecht, and G. C. Verghese, *Principles of Power Electronics*, vol. 2, 1st ed. Reading, MA, USA: Addison-Wesley, 1991, ch. 8.2, pp. 176–179.

[13] A. A. Rafi, A. Piovaccari, P. Vancorenland, and T. Tuttle, "A harmonic rejection mixer robust to RF device mismatches," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2011, pp. 66–68.

[14] L. Sundstrom *et al.*, "Complex IF harmonic rejection mixer for non-contiguous dual carrier reception in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 48, no. 7, pp. 1659–1668, Jul. 2013.

[15] T. Yang, K. Tripurari, H. Krishnaswamy, and P. R. Kinget, "A 0.5GHz–1.5GHz order scalable harmonic rejection mixer," in *Proc. IEEE Radio Freq. Integr. Circuits Symp. (RFIC)*, Jun. 2013, pp. 411–414.

[16] Q. T. Duong and A. Alvandpour, "Low noise linear and wideband transconductance amplifier design for current-mode frontend," in *Proc. 14th Int. Symp. Integr. Circuits (ISIC)*, Dec. 2014, pp. 476–479.

[17] D. Murphy *et al.*, "A blocker-tolerant wideband noise-cancelling receiver with a 2dB noise figure," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2012, pp. 74–76.

Douglas Adams received the B.S. degree in electrical engineering and applied mathematics, and the M.Eng. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2009 and 2010, respectively, and the Ph.D. degree from Stanford University, Stanford, CA, USA, in 2016.

After receiving the Ph.D., he joined FINsix, Menlo Park, CA, USA, where he worked on high-frequency resonant power converters.

Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering from Tel-Aviv University, Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002.

From January 2002 to July 2002 she was a Postdoctoral Fellow at the Digital Signal Processing Group at MIT. She is currently a Professor in the Department of Electrical Engineering at the Technion – Israel Institute of Technology, Haifa, Israel, where she holds the Edwards Chair in Engineering. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT and was a Visiting Professor at Stanford University, Stanford, CA. She is the author of the book *Sampling Theory: Beyond Bandlimited Systems* and co-author of the books *Compressed Sensing* (Cambridge University Press) and *Convex Optimization Methods in Signal Processing and Communications* (Cambridge University Press). Her current research interests include statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar was a member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing Technical Committees, and is a member of the Young Israel Academy of Science and Humanities, the Israel Committee for Higher Education, the IEEE Sensor Array and Multichannel Technical Committee, and serves on several other IEEE committees. She was a Horev Fellow of the Leaders in Science and Technology program at the Technion, and an Alon Fellow. She was a recipient of many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award in 2013, the IEEE/AESS Fred Nathanson Memorial Radar Award in 2014, the IEEE Kiyo Tomiyasu Award in 2016, the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, the Wolf Foundation Krill Prize for Excellence in Scientific Research, the Henry Taub Prize for Excellence in Research (twice), the Hershel Rich Innovation Award (three times), the Award for Women with Distinguished Contributions, the Andre and Bella Meyer Lectureship, the Career Development Chair at the Technion, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technion's Award for Excellence in Teaching (two times). She was also a recipient of several best paper awards and best demo awards together with her research students and colleagues, including the SIAM Outstanding Paper Prize, the UFFC Outstanding Paper Award, the Signal Processing Society Best Paper Award, and the IET Circuits, Devices and Systems Premium Award, and was selected as one of the 50 most influential women in Israel. She is the Editor-in-Chief of Foundations and Trends in Signal Processing. She was a Signal Processing Society Distinguished Lecturer, and served as an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*. She was the Co-Chair and the Technical Co-Chair of several international conferences and workshops.

Boris Murmann (S'99–M'03–SM'09–F'15) received the Dipl.Ing. (FH) degree in communications engineering from Fachhochschule Dieburg, Dieburg, Germany, in 1994, the M.S. degree in electrical engineering from Santa Clara University, Santa Clara, CA, USA, in 1999, and the Ph.D. degree in electrical engineering in 2003, all from the University of California at Berkeley, Berkeley, CA, USA.

From 1994 to 1997, he was with Neutron Mikrolektronik GmbH, Hanau, Germany, where he developed low-power and smart-power ASICs in automotive CMOS technology. Since 2004, he has been with the Department of Electrical Engineering, Stanford University, Stanford, CA, where he currently serves as a Full Professor. His research interests are in the area of mixed-signal integrated-circuit design, with special emphasis on data converters and sensor interfaces.

In 2008, Dr. Murmann was a co-recipient of the Best Student Paper Award at the VLSI Circuits Symposium and the recipient of the Best Invited Paper Award at the IEEE Custom Integrated Circuits Conference (CICC). He received the Agilent Early Career Professor Award in 2009 and the Friedrich Wilhelm Bessel Research Award in 2012. He has served as an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and as the Data Converter Subcommittee Chair of the IEEE International Solid-State Circuits Conference (ISSCC). He is the Program Vice-Chair of ISSCC 2017.