www.ietdl.org

Published in IET Circuits, Devices & Systems
Received on 6th May 2010
Revised on 11th August 2010
doi: 10.1049/iet-cds.2010.0147

ISSN 1751-858X

Xampling: analog to digital at sub-Nyquist rates

M. Mishali Y.C. Eldar O. Dounaevsky E. Shoshan

Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa, Israel
E-mail: moshiko@tx.technion.ac.il

Abstract: The authors present a sub-Nyquist analog-to-digital converter of wideband inputs. The circuit realises the recently proposed modulated wideband converter, which is a flexible platform for sampling signals according to their actual bandwidth occupation. The theoretical work enables, for example, a sub-Nyquist wideband communication receiver, which has no prior information on the transmitter carrier positions. The present design supports input signals with 2 GHz Nyquist rate and 120 MHz spectrum occupancy, with arbitrary transmission frequencies. The sampling rate is as low as 280 MHz. To the best of the authors' knowledge, this is the first reported hardware that performs sub-Nyquist sampling and reconstruction of wideband signals. The authors describe the various circuit design considerations, with an emphasis on the non-ordinary challenges the converter introduces: mixing a signal with a multiple set of sinusoids, rather than a single local oscillator, and generation of highly transient periodic waveforms, with transient intervals on the order of the Nyquist rate. Hardware experiments validate the design and demonstrate sub-Nyquist sampling and signal reconstruction.

1 Introduction

Analog to digital conversion (ADC) is the key enabling many of the advances in signal processing. The Shannon–Nyquist theorem [1, 2] lies at the heart of essentially all ADC devices. In wideband communication, the signal spectrum can reach prohibitively large frequencies, so that direct uniform sampling, at the high rate of twice the maximal frequency component, becomes practically infeasible. Instead, a common practice in engineering is demodulation. The input is multiplied by the carrier frequency of a band of interest, so as to shift the contents of the narrowband transmission from the high frequencies to the origin. Then, commercial ADC devices at low rates are utilised. Demodulation, however, requires knowing the exact carrier frequency. To date, the alternative to demodulation is to sample the entire wideband spectrum, for example, by using a series of low-rate samplers through time interleaving [3] or parallel bandpass oversampling schemes [4]. This requires excessive hardware solutions with extremely high analog bandwidths, and in addition necessitates sophisticated digital algorithms for timing synchronisation and recovery [5, 6].

In this paper, we present a circuit-level hardware prototype of a sub-Nyquist sampling system. Our design realises the modulated wideband converter (MWC) strategy of [7]. The MWC can sample wideband inputs at a low rate, proportional to the actual bandwidth occupation, without knowledge of the carrier positions. Specifically, our hardware can treat 2 GHz Nyquist-rate input signals with spectrum occupancy up to 120 MHz. The sampling rate can be made as low as 280 MHz, which is only 14% of the Nyquist rate, approaching the lowest possible rate for multiband signals with unknown carrier positions [8]. We report on a series of lab experiments which verify the circuit design and prove that the system accommodates a wide dynamic range of input powers, reaching up to 50 dB. The crowning glory of our experiments is a demonstration of sub-Nyquist sampling and reconstruction. We use a mixture of popular communication signals, lying anywhere in the wideband spectrum, and affirm accurate estimation of the unknown carrier frequencies and faithful reconstruction of each of the wideband transmissions.

The MWC strategy, as briefly overviewed in Section 2, involves mixing the signal with highly transient periodic waveforms, which poses two interesting circuit challenges. First, it requires to mix the input simultaneously with the multiple sinusoids comprising the periodic waveforms, in contrast to the standard mixing set-up with a single local oscillator (LO). The second challenge pertains to generating a waveform that duplicates a given highly transient pattern on consecutive time intervals, thereby requires meeting strict timing constraints. Instead of resorting to expensive hardware solutions, we preferred cheap off-the-shelf devices and had to accomplish the desired functionality by wisely operating commercial devices beyond their intended specifications. Specifically, a standard switching-type mixer is used in conjunction with equalisation and power control schemes. By capitalising on the properties of specific shift-register (SR) devices we successfully managed to satisfy the strict timing requirements, with no time-delay elements, so as to avoid synchronisation problems that are popular in time-interleaved ADCs [5, 6].

The main contribution of this paper is in proving feasibility of the MWC strategy using existing analog devices. Several alternative approaches have been proposed in the literature, that can, at least in theory, achieve sub-Nyquist rates. In practice, however, previous methods have shown limited

sub-Nyquist performance. The strategy of [9] relies on dense discretisation of the input; the analog signal is modelled as a finite sum of integral harmonic frequencies. As a result of this discretisation step, the approach is sensitive to slight model mismatches, which are inevitable in practice. Another consequence is severe computational loads, which effectively limit the method to a narrow input range [10]. The implementations of [11, 12] demonstrate the method for signals with Nyquist rate below 1 MHz, while our interest here is on output bandwidths on the order of GHz. An interesting alternative approach, proposed in [13], suggests jittered undersampling for reducing the sampling rate. So far, the method was demonstrated for pure sinusoids sampled at a rate that is substantially higher than the minimum, so that effectively no aliasing occurs in the folded spectrum. Reconstruction of wideband transmissions at sub-Nyquist rate has not yet been reported. In contrast, our hardware treats 2 GHz Nyquist-rate inputs, with lab experiments demonstrating correct reconstruction from minimal rate sampling, even in the worst-case scenario in which multiple input transmissions alias to the same baseband frequencies. Underlying the MWC is the Xampling methodology [10], which establishes the framework for reasoning the above differences. To the best of our knowledge, based on [10] and the discussion on related work herein, our system is the first wideband hardware accomplishing minimal sub-Nyquist rates without knowledge of the carrier positions.

This paper is organised as follows. Section 2 introduces the MWC system, motivates the hardware specifications we chose to realise and elaborates on the circuit challenges. The design work spans two sections: Section 3 provides details on the mixing–filtering stage, whereas Section 4 explains the periodic waveform generation. Results of lab experiments and sub-Nyquist reconstruction demonstration appear in Sections 5 and 6, respectively. Section 7 concludes with a discussion on related work and future directions. Throughout, constants are specified in either linear (e.g. 1 Vptp) or logarithmic (e.g. 4 dBm) formats.

2 Modulated wideband converter

2.1 Theoretical background

The MWC is aimed at sampling wideband sparse signals at sub-Nyquist rates according to the scheme of Fig. 1*a*. To accomplish this goal, a multiband input model is assumed. An analog signal $x(t)$ is termed multiband if its Fourier transform $X(f)$ is concentrated on N frequency intervals, or bands, with individual widths that are no greater than B Hz. An illustration of a multiband spectra appears in Fig. 1*b*. The maximal possible frequency of $x(t)$, denoted by $f_{\max}$, dictates the Nyquist rate

$$f_{\mathrm{NYQ}} = 2f_{\max} \tag{1}$$

The values of N, B, $f_{\max}$ depend on the specifications of the application at hand, such as the expected number of concurrent transmissions and the communication technology which defines the narrowband widths. The multiband model does not assume knowledge of the carrier locations f_i, and these can lie anywhere below $f_{\max}$.

The analog front-end of the MWC preprocesses a multiband signal $x(t)$ using m simultaneous channels. In the ith channel the signal is multiplied by a periodic function $p_i(t)$ with period $T_p = 1/f_p$. The product is lowpass filtered by $h(t)$ with cutoff $f_s/2$, and then sampled uniformly every $T_s = 1/f_s$ seconds.

The mixing operation scrambles the spectrum of $x(t)$, such that a portion of the energy of all bands appears in baseband. More specifically, since each $p_i(t)$ is periodic, it has a Fourier

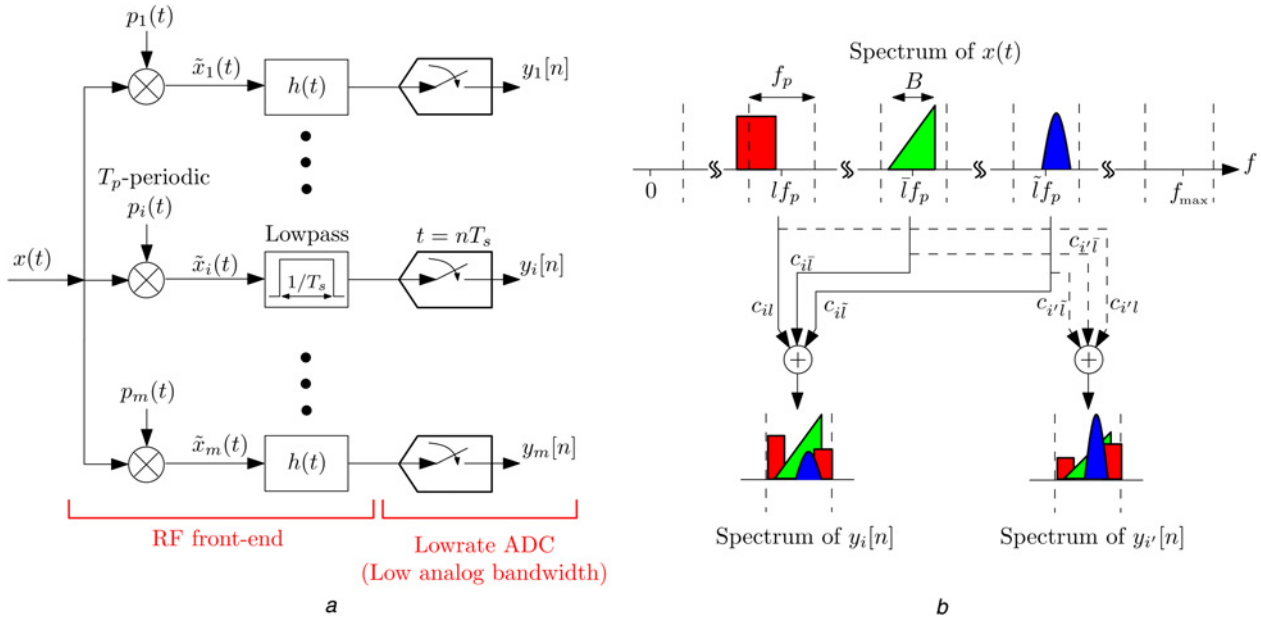

Fig. 1 *Modulated wideband convertor: a practical sub-Nyquist sampling system for multiband signals*

a Sampling stage aliases the wideband energy to low frequencies
b Spectrum of the output sequences $y_i[n]$ consists of overlayed energies, shifted from consecutive f_p-length spectrum slices of $x(t)$
In the example, channels i and i' realise different linear combinations of the spectrum slices centered around lf_p, $\tilde{l}f_p$ and $\tilde{l}f_p$. For simplicity, the aliasing of the negative frequencies is not drawn

expansion

$$p_i(t) = \sum_{l=-\infty}^{\infty} c_{il} e^{j(2\pi/T_p)lt} \quad (2)$$

for some coefficients c_{il}. Therefore the mixing results in a weighted-sum of f_p-shifted copies of $X(f)$, with weights given by the Fourier coefficients c_{il} [7]. The lowpass filter $h(t)$ transfers only the narrow band frequencies up to $f_s/2$ from that mixture to the output sequence $y_i[n]$. The aliased output is illustrated in Fig. 1*b*. While aliasing is often an undesired artifact in sampling, here it is deliberately utilised to generate mixtures at baseband. Intuitively, since there are m branches with different periodic waveforms $p_i(t)$, the system captures different mixtures of the spectrum, which enable recovery of $x(t)$ from the samples $y_i[n]$, $1 \le i \le m$. The exact parameter values and periodic waveforms are specified below.

Low-rate ADCs follow the radio-frequency (RF) front-end. Commercial devices can be used for the task because of the preceding lowpass. The fact that the ADCs see only a lowpass input is an advantage over classic non-uniform methods [3, 14, 15] which acquire pointwise values of $x(t)$, and therefore necessitate an ADC with Nyquist-rate track-and-hold circuitry. In contrast, the MWC shifts the Nyquist burden to RF technology, thereby avoiding interaction of wideband signals with the ADC devices. We also point out that [3, 14, 15] assume the band positions, whereas the MWC does not require this knowledge.

It is important to note that the conceptual division of the spectrum to f_p-width slices does not restrict the multiband spectra. The carrier positions f_i are arbitrary, so that a band is not limited to reside within a single spectrum slice; it can occupy consecutive slices, as illustrated in Fig. 1*b*. Digital post-processing is used to detect the active slices (those intersecting with signal support), reconstruct their content, and stitch together an appropriate data stream, per each input narrowband transmission [10].

2.2 Hardware specifications

The MWC is a flexible system with various parameters, of which we chose to realise the specifications that appear in Table 1. We next motivate these choices.

The multiband model, $N = 6$, $B = 19$ MHz, matches a scenario of three concurrent transmissions in the set-up that was numerically simulated in [7]. The Nyquist rate $f_{NYQ} = 2.075$ GHz stems from the frequency of a voltage-controlled oscillator (VCO) in our system and can be increased for higher VCO frequencies.

Table 1 Prototype specifications

Parameter	Choice
signal model	$N = 6$, $B = 19$ MHz, $f_{NYQ} = 2$ GHz
number of channels m	4
waveform type	periodic sign alternation
alternation rate	2.075 GHz
sign pattern length M	108
period f_p	2.075/108 = 19.212 MHz
filter cut-off	33 MHz
sampling rate/channel f_s	70 MHz
total rate mf_s	280 MHz

The basic parameter choice that is studied in [7] is $m \ge 4N$ and $f_s = f_p \ge B$. This choice results in a sampling rate of $mf_s \simeq 4NB$, which can be significantly smaller than f_{NYQ}. The basic choice implies $m = 4N = 24$ sampling branches. In order to save hardware size and price, we utilised an advanced MWC version, also proposed in [7], in which the number of physical channels is collapsed by a factor of q, at the expense of increasing the sampling rate of each channel by the same factor, so that $f_s = qf_p$ and the total sampling rate mf_s is unchanged. Our design uses a collapsing factor $q = 3$. To further reduce m, an extra factor of 2 is accomplished by combining the SBR2 procedure from [8] into the reconstruction algorithm. With SBR2, the sampling rate can approach $mf_s \simeq 2NB$, which is theoretically optimal [8]. Summarising, $m = 4N/q/2 = 4$ is the actual number of sampling branches chosen, with a total sampling rate of $mf_s = 280$ MHz, which is about 14% of the Nyquist rate and approaches the optimal $2NB = 228$ MHz rate.

We followed a suggestion of [7] and chose cyclic sign-alternating functions for $p_i(t)$. Our study [16] shows that popular binary patterns, for example, the Gold or Kasami sequences that are widely used in communication for their low mutual correlation, fit the MWC system. In practice, we designed a single SR with 2 GHz clock rate, and tapped the register at four locations, so that $p_i(t)$ are delayed versions of a single binary pattern. This choice reduces the hardware size and was also suggested in [7]. Section 4 explains how we initialise the SR.

Our prototype implements the RF front-end of the MWC using two physical boards. An analog board realises the four analog paths from $x(t)$ to the relevant filter outputs, whereas a digital counterpart provides the periodic waveforms. This hardware scope captures the innovative theoretical parts as well as the challenges in the actual design, which are explained in the sequel. The filter cut-off of 33 MHz allows a sampling rate of $f_s = 70$ MHz. In our lab experiments, a four-channel scope is used for digitising the $m = 4$ filter outputs. Future work will embed commercial ADC devices for that task; relevant devices with 16 bits resolution include AD9460 and ADS5562. The 33 MHz cut-off is slightly higher than half of the working band, $qf_p/2 = 28.8$ MHz, to allow flat phase response over the passband. Non-ideal amplitude response is calibrated after manufacturing using the algorithm developed in [17].

2.3 Circuit challenges

In designing an analog circuit to realise the MWC we encountered two main difficulties: (i) analog mixing with spectrally rich waveforms $p_i(t)$, and (ii) constructing the periodic waveforms with the required alternation speed of 2.075 GHz. RF mixers are tailored to multiplication with a single sinusoid, which is the standard procedure for modulating and demodulating an information band onto a given carrier f_c. In contrast, the MWC requires simultaneous multiplication with many sinusoids – those comprising $p_i(t)$. This results in attenuation of the output and substantial non-linear distortion not accounted for in datasheet specifications. The next section describes our circuit solutions for mixing with periodic waveforms.

The second challenge pertains to constructing $p_i(t)$. The waveforms can be generated either by analog or digital means. Analog waveforms, such as sinusoid, square or sawtooth waveforms, are smooth within the period, and therefore do not have enough transients at high frequencies

doi: 10.1049/iet-cds.2010.0147

which is necessary to ensure sufficient aliasing. On the other hand, digital waveforms can be programmed to any desired number of alternations within the period, but require meeting timing constraints on the order of the clock period. In our setting, the clock interval of $1/f_{\mathrm{NYQ}} = 480\,\mathrm{ps}$ leads to severe timing constraints that are difficult to satisfy with existing digital devices. Section 4 elaborates on the solution we found for this challenge.

Our goal in the present paper is to focus on the hardware realisation of the RF front-end of Fig. 1*a*. Reconstruction of the signal is carried out digitally [7]. Section 6 briefly overviews the recovery algorithm and demonstrates signal reconstruction at sub-Nyquist rates.

3 Analog board

The analog board consists of three consecutive stages: splitting the input into four channels, mixing with a given set of periodic waveforms $p_i(t)$ and lowpass filtering. Fig. 2 presents a block diagram of the analog path end-to-end.

3.1 Dynamic range

In realising our sub-Nyquist hardware, we bear in mind the representative application of a wideband receiver, which needs to amplify the input to a level that the ADC can treat. Commercial ADCs at rates around $f_{\mathrm{s}} = 70\,\mathrm{MHz}$ can work with amplitude levels as low as tens to hundreds of millivolts, see for example [18, 19]. We decided to target an output amplitude of at least 250 millivolts peak-to-peak (mVptp). Equivalently, we require an output power satisfying

$$\mathrm{Pout}_{\mathrm{req.}} \geq -8.06\,\mathrm{dBm} \tag{3}$$

In addition, it is important to maintain a sufficiently high signal-to-noise ratio (SNR) so that the consecutive digital algorithms can function properly. Reconstruction in the presence of noise was demonstrated in [7], where it was shown that an SNR of about 15 dB allows correct recovery of the spectrum support, an essential operation preceding any further processing. RF mixers usually produce undesired spurious distorted images of the signal, which are also considered as noise. Therefore

$$\mathrm{SNDR}_{\mathrm{req.}} \geq 15\,\mathrm{dB} \tag{4}$$

stands for the required signal to noise-and-distortion ratio at the filter outputs.

Link budget calculation is a common RF analysis that predicts the amplitude level and SNDR of an RF chain, such as Fig. 2. The analysis takes into account the gain, noise figure (NF) and third intersection-point (IP3) of each of the devices along the path. In what follows, we analyse the chain of devices in our realisation with emphasis on design considerations because of the wideband nature of the front-end.

An immediate difficulty encountered with the non-ordinary mixing is that datasheet specifications of commercial mixer devices do not cover the case of multiple harmonics at the LO port. To estimate the actual mixer's parameters, we had to perform our own lab evaluations for the chosen device, which led to the following adjustments

$$\begin{aligned} &\text{Loss:} && -16\,\mathrm{dB} \rightarrow \ -6\,\mathrm{dB} \\ &\text{IP3:} && 27\,\mathrm{dBm} \rightarrow 30\,\mathrm{dBm} \\ &\text{LO power:} && 17\,\mathrm{dBm} \rightarrow 20\,\mathrm{dBm} \end{aligned} \tag{5}$$

These modifications impact the actual design of the RF path, since without adjusting the datasheet values, conventional budget analysis would falsely indicate that a smaller gain is needed. Taking (5) into account, we are now in position to compute the dynamic range of our system.

The amplitude requirement (3) boils down to minimal input power

$$P_{\mathrm{IN}} \geq \mathrm{Pout}_{\mathrm{req.}} - \sum G_i = -8.06 - 47 = -55\,\mathrm{dBm} \tag{6}$$

where the gains G_i are reported in Fig. 2, and the equality above is achieved when the variable attenuators are set to 0 dB. To determine the input powers that satisfy the required SNDR (4), we consider two extreme scenarios. One of $x(t)$ with a very low power, in which case the SNDR is mainly affected by the additional noise that the system generates. Technically, the NF and the gains of the devices along the path dictate this contribution. The other scenario assumes several narrowband transmissions, whose total power is high and is equally distributed. In this setting, spurious images of the true signal appear because of the non-linearities of the devices. Since the input power is high, these images become dominant and distort the ability of the system to distinguish between the true input and the undesired products. In a wideband system, the effect of non-linearities is prominent since besides the undesired products that appear close to the original frequencies, there are many other spurious images which fall well within the wideband range of the front-end and increase the total distortion level.

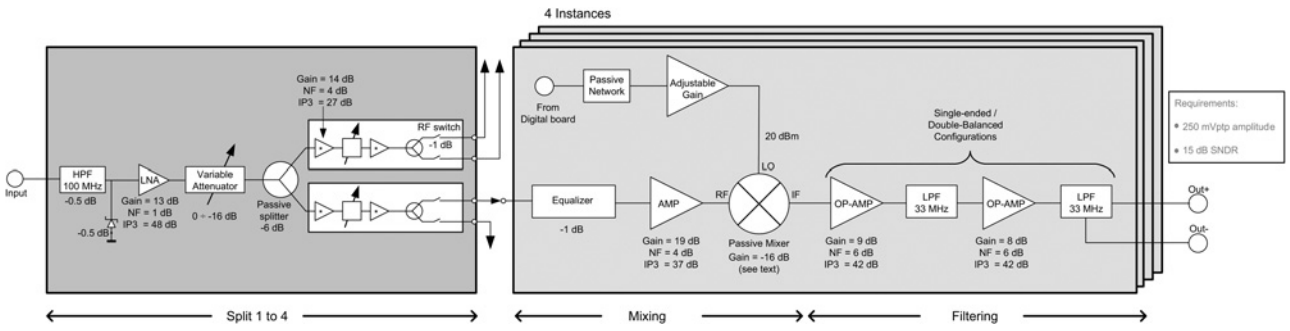

Fig. 2 *Block diagram of the analog board*

Consider the case of $x(t)$ with low power. The thermal noise at the input of the system is

$$\mathrm{n_{thermal}} = \mathrm{KT}W = -114 + 10\log_{10} W = -80.82\,\mathrm{dBm} \tag{7}$$

where $\mathrm{KT} = -114$ dBm is the thermal noise power per MHz, and $W = 2075$ MHz is the intended front-end bandwidth. It follows from Frii's formula [20]

$$F_{\mathrm{system}} = \frac{\mathrm{SNR_{in}}}{\mathrm{SNR_{out}}} = F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 G_2} + \cdots \tag{8}$$

that the equivalent NF of the system is $F_{\mathrm{system}} = 3.13$ dB. Consequently, the signal power at the input must satisfy

$$P_{\mathrm{IN,\,min}} \geq \mathrm{SNDR_{req.}} + \mathrm{n_{thermal}} + F_{\mathrm{system}} = -62.5\,\mathrm{dBm} \tag{9}$$

which is implied by (6).

To analyse the other extreme of high input power, the variable attenuators are set to their maximal level of -15.5 dB in order to deliberately reduce the aggregated gain, and thus decrease the non-linear effects. The contribution of the thermal noise for high input powers can be neglected. To derive the distortion level in this setting, we begin with the output signal power

$$P_{\mathrm{out}} = N_{\max} G_{\mathrm{total}} P_{\mathrm{IN,\,max}} \tag{10}$$

where $N_{\max} = 10$ is taken for the maximal number of concurrent narrowband transmissions that the system is designed for. Note that $N_{\max}$ is greater than the three transmissions taken in our signal model, since the system is modular and we would like the design to support future applications with signal sets beyond the specifications of Table 1. The total gain $G_{\mathrm{total}} = \sum G_i = 15$ dB is calculated from Fig. 2, and $P_{\mathrm{IN,\,max}}$ is the unknown to be computed. We calculate the equivalent IP3 of the entire system, using the formula [21]

$$(\mathrm{IP3_{system}})^{-1} = \sum_{i=1}^{I} \left(\mathrm{IP3_i} \prod_{j=i+1}^{I} \mathrm{G_j} \right)^{-1} = (30.98\,\mathrm{dBm})^{-1} \tag{11}$$

where $I = 15$ is the number of RF devices in the path. Then, the SNDR at the output is given by [22]

$$\mathrm{SNDR} = 2(\mathrm{IP3_{system}} - P_{\mathrm{out}}) - 7.3\,\mathrm{dB} \tag{12}$$

We subtract 7.3 dB in order to account for the fact that IP3 measures the mutual distortion in the presence of two tones only. Compensating the equation for a larger number of input transmissions, $N_{\max} = 10$, requires this term [22]. Taking into account the number of concurrent transmissions is extremely important in a wideband system whose frequency bandwidth covers more than an octave. In narrowband systems, various intermodulations of input tones fall beyond the input bandwidth. The distortions are rejected by the system filters and thus the number of transmissions hardly affects the SNDR. In contrast, intermodulations in a wideband system fall within the front-end bandwidth, modulate with the other tones, and thus raise the total distortion level. We conclude that the signal input power must satisfy

$$P_{\mathrm{IN,\,max}} \leq \mathrm{IP3_{system}} - \frac{1}{2}(\mathrm{SNDR_{req.}} + 7.3) - 10\log_{10} N_{\max} - G_{\mathrm{total}} = -5.17\,\mathrm{dBm} \tag{13}$$

Our system can therefore treat input signals at a dynamic range of 49 dB, as follows from (6), (9) and (13). We note that the calculations of the distortion level for high-power signals assume that all narrowband transmissions have equal power. Refining the theoretical calculations for different power levels is possible. However, since the link budget analysis has spares we did not pursue this direction; Section 5 reports on an actual 50 dB dynamic range of the manufactured system. Fig. 3 depicts the expected output level and the SNDR for varying input power, when optimising for the highest SNDR over all possible

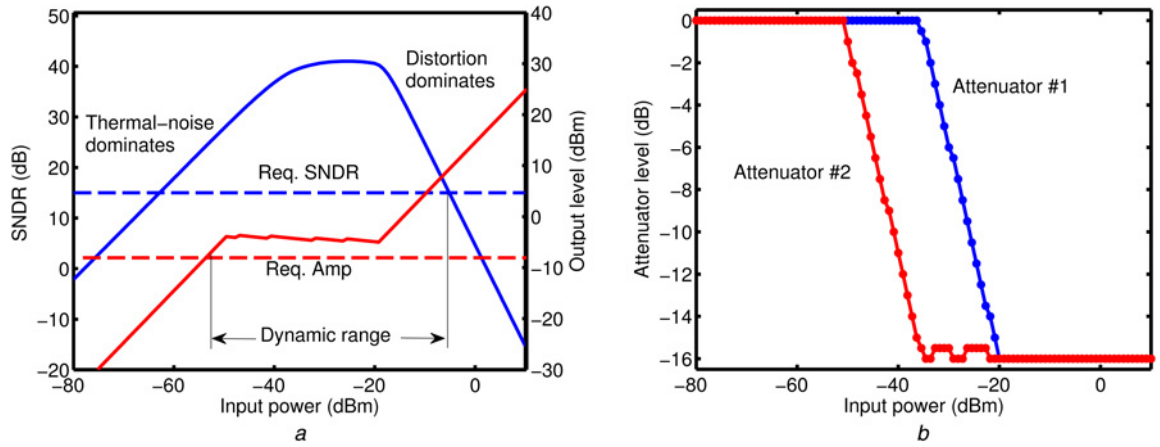

Fig. 3 *Optimising the output SNDR by controlling the analog again*

a Output power level and SNDR for varying input power (solid lines). The former is monotone in the input power. The dynamic range of the system is achieved when both requirements (3) and (4) are satisfied

b The optimal attenuator levels (in the SNDR sense)

doi: 10.1049/iet-cds.2010.0147

attenuator levels using (9) and (13). We point out that the increase in the attenuators level does not breach (3) for all input powers higher than −55 dBm.

3.2 Mixing with multiple sinusoids

The key idea behind the MWC strategy is periodic mixing. Besides the effect on the dynamic range, mixing with periodic waveforms requires supplementary circuitries to support the non-ordinary functionality. The schematic of that block (Fig. 4) highlights three consequences of the non-ordinary mixing. A passive mixer was chosen since active ones typically allow a narrow range of LO frequencies, whereas $p_i(t)$ spans a wide spectrum. The tunable power control on $p_i(t)$ was inserted to allow flexible post-manufacturing refinements of the datasheet adjustments (5). Such power control is not needed in standard mixing for which datasheets guarantee the performance.

The third block in Fig. 4 is a wideband equaliser. The equaliser is located near the LO port, just before the mixer, instead of the common knowhow of equalisation towards the end of the chain, near the ADC terminals or even in the digital domain, where it can flatten the response of the entire analog path. The reason is again the non-ordinary mixing. More specifically, recall that conventional equalisation corrects the effective system response

$$H_{\text{total}}(f) = \prod_i G_i(f - f_c) \prod_i H_i(f) \tag{14}$$

where $G_i(f)$, $H_i(f)$ are the responses of the devices preceding and following the mixer, respectively. The shift in the responses of $G_i(f)$ accounts for the effect of standard demodulation, in which $x(t)$ is multiplied by a single frequency source of f_c Hz. In contrast, our sub-Nyquist system produces the product $x(t)p_i(t)$ which contains a weighed-sum of f_p-wide spectrum slices of $x(t)$; (see Fig. 1*b*). Consequently, there does not appear to be a simple way to translate the frequency response of the devices preceding the mixer, as in (14), to an equivalent response at the output. We therefore chose to place the equaliser circuit before the mixer. Compensating for the rest of the path, namely for the non-ideal response of the filter $h(t)$, is done digitally [17]. A circuit simulation of the equaliser is presented in the top-right corner of Fig. 4, ensuring an equalised flat response over the wideband regime until 2 GHz.

3.3 Lowpass filtering of spectrum mixtures

The filtering is the last processing step in the analog domain. In standard bandlimited sampling, a lowpass filter has the role of rejecting out of band noise. The noise is typically small and consequently does not require a sharp cut-off. In the MWC prototype, the input to the filter contains energy spread all across the spectrum, with non-eligible power beyond the cut-off. For this reason, we realised a sharp cut-off around 33 MHz using two elliptic filters of order seven concatenated and buffered. Further design considerations and schematics are given in the technical report [23]. We note that the modern trend in engineering is to shift filtering tasks to the digital domain. Here, analog processing is inevitable, since the filter is responsible for reducing the signal bandwidth before entering the ADC.

4 Digital board

The choice of the periodic functions $p_i(t)$ dictates the Nyquist rate of the input signals that the MWC can handle. In particular, if $p_i(t)$ has Fourier coefficients c_{il} with non-negligible amplitudes for all $-L \le l \le L$, then the MWC can capture signals with band locations anywhere below Lf_p, corresponding to a Nyquist rate of $2Lf_p$. The precise requirement on the magnitude of c_{il} depends on the

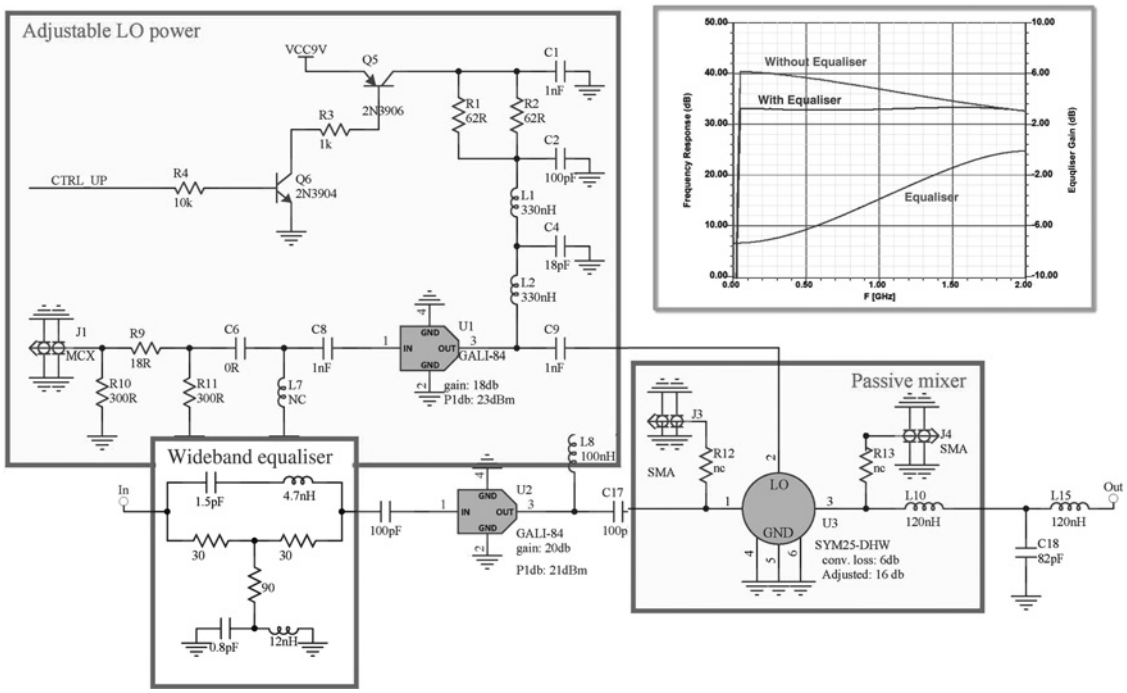

Fig. 4 *Schematics of mixing with multiple sinusoids*

Wideband equaliser flattens the frequency response at the RF port of the mixer. Low frequencies (up to 500 MHz) are attenuated by approximately 8 dB, whereas high frequencies (above 1.5 GHz) are not altered

properties of the digital reconstruction algorithm and is beyond the current scope [7]. Nonetheless, in principle, every periodic function with high-speed transitions within the period T_p can be appropriate. The digital board in our design generates a sign-alternating function, with $M = 108$ sign intervals within the period interval T_p. In the sequel, we elaborate on the circuit design that ensures the periodicity of $p_i(t)$ and explain the choice of the sign pattern.

4.1 Periodicity and high-speed alternations

A straightforward approach to realise a ± 1 logic waveform would be to program the desired pattern into a field programmable gate array (FPGA) device. Unfortunately, this approach is difficult for realising a high-rate SR. Popular FPGA devices do not stand a clock rate of 2.075 GHz and premium devices are power consuming and very expensive. Therefore we departed from conventional logic design and realised the SR using discrete devices.

The difficulty in implementing a discrete SR at 2.075 GHz is in satisfying the set-up/hold timing constraints. For correct functionality, the data must propagate between the flops sufficiently fast and accurate before the next clock edge arrives. Any solution which is based on fine tuning of short time delays (on the order of $1/2.075 \simeq 480$ ps) is prone to errors; the wide literature on time mismatches in interleaved-ADCs is evidence of this difficulty (cf. [6]). To overcome this challenge, our design is based on the MC10EP142MNG device, which implements an internally wired 8 bits SR at emitter-coupled-logic (ECL) technology. We exploited a specific property of that device. The datasheet reports on a maximal shifting frequency of 2.8 GHz (equivalent to a minimal clock period of 357 ps), and a delay from clock to output of about 670 ps. This means that the manufacturer guarantees correct functionality in the internal register, up to a shifting rate of 2.8 GHz; however, once the signal leaves the device the delay is almost doubled.

Capitalising on this difference is the key for satisfying the set-up and hold requirements with no additional hardware. As illustrated in Fig. 5, since the delay of the last bit in each package is larger than the clock period, it reaches the next package one clock later than what would normally be expected in a standard FPGA design. This behaviour seems erroneous and is likely to be reported as a violation of timing constraints by software tools. Fortunately, it follows from Fig. 5 that we can view the large delay as an additional flop in between adjacent packages, with an equivalent clock-to-output delay of $670 - 480 = 190$ ps. In other words, the data arrive about half a clock period before the next edge, in which case both the set-up and time requirements are satisfied. To emphasise, the proposed solution requires no time-delay elements or synchronisation mechanisms. This allows a sign pattern of length $M = 108$ bits to be realised by only 96 physical flops.

4.2 Sign-pattern selection

The SR of the digital board is programmed to the initial value

$$\text{Pattern} = \texttt{43 A7 A5 D7 96 AB 62 B7 2A B3 5C AC} \tag{15}$$

encoded in hexadecimal bytes, which is reasoned as follows. Our study [16] introduced a theoretical measure, termed the expected restricted isometry property (ExRIP), for the quality of a given set of periodic sign-alternating functions $p_i(t)$. Among the conclusions of [16] is that Gold, Kasami and even randomly chosen binary patterns are suitable for the MWC system, whereas orthogonal Hadamard sequences are inadequate. We therefore generated a set of candidate patterns and retained only a subset with high ExRIP values. In order to choose between these patterns, we observed $p_i(t)$ in a spectrum analyser and decided on (15) since the power of its coefficients c_{il} were relatively balanced, as depicted in Fig. 9*a*. Such a balance ensures approximately equal power levels of the multiple sinusoids comprising $p_i(t)$. Experimentally, we noticed that a balanced Dirac set contributes to mitigating non-linear distortions in the non-ordinary mixing.

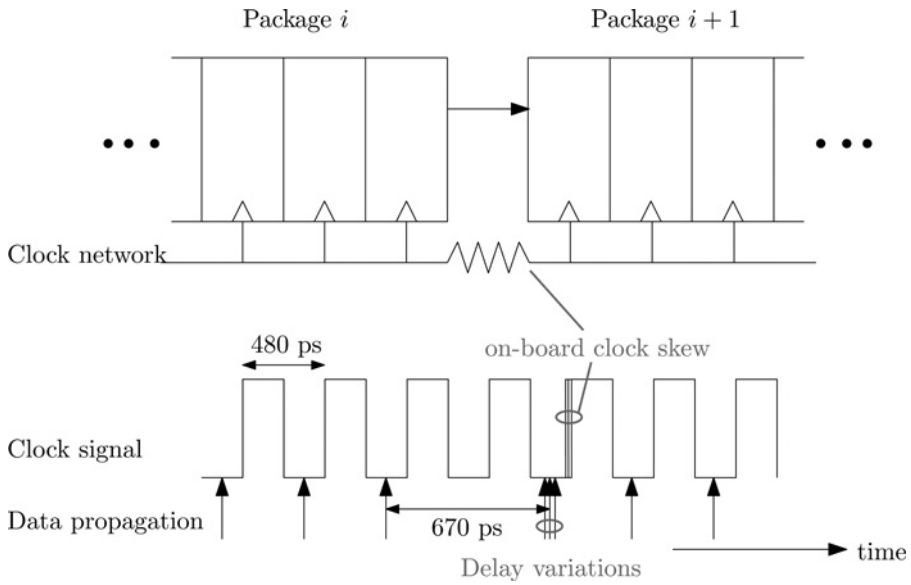

Fig. 5 *Interface of adjacent SR packages involves uncertainty in the clock network because of on-board routing*

Timing diagram illustrates the propagation of data in the SR chain. Set-up and hold constraints are satisfied internally by the manufacturer, and externally by viewing the clock to output delay as an additional flop in between adjacent packages

doi: 10.1049/iet-cds.2010.0147

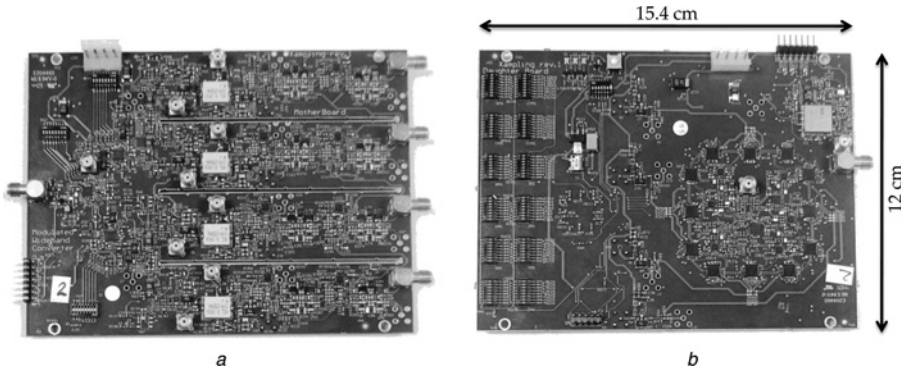

Fig. 6 *Photos of the prototype system*

a Analog board realises $m = 4$ sampling channels, including the amplification path, mixing with $p_i(t)$ (connectors on the rear), and lowpass filtering
b Digital board provides sign-alternating periodic waveforms, which are derived from a single SR chain

Photos of the analog and digital boards are attached in Fig. 6. Table 2 lists the reference names of the devices in our system, allowing to reproduce the design in a standard electronic laboratory.

5 Experiments

The analog and digital designs were verified by the following experiments. The lab equipment that was used throughout includes: HP-8563E spectrum analyser, Agilent Infiniium 54855A four-channel scope, HP-E4432B signal generator, Agilent 8753E network analyser and HP 436A power meter.

5.1 Power level, SNDR and dynamic range

The purpose of this test is to verify that the outputs stand the required power level (3) and SNDR (4). We synthesised a pure sinusoid signal at the system input, with 92.4 MHz frequency and −45 dBm input power. The on-board attenuators were set to 0 dB. Fig. 7 displays the four outputs, as measured by the four-channel scope (with 50 Ω impedance input terminals). The amplitude levels are all around 1 Vptp, or equivalently 4 dBm. A more accurate measurement, using a power meter, revealed that the output power is about 5 dBm. It follows that the minimal input power satisfying (3) is −55 dBm, conforming with (6). As explained earlier, this test assures (3) over all input powers higher than −55 dBm, as long as the attenuators are set according to Fig. 3*b*. Note, that the result also validates our modification (5) of the mixer gain from −6 dB to −16 dB.

To check the SNDR requirement (4), we used an 1:2 RF combiner to sum two signal generators to the board input, with carriers 146 and 145.8 MHz and equal powers. These inputs convolve with the 137 MHz sinusoid of $p_i(t)$. Fig. 8*a* displays the spectrum at the output of the first channel for −55 dBm input powers. The original tones, with the 200 kHz spacing, are noticed, with −14.2 dBm output power per tone. The noise level is −78 dBm at 1 kHz resolution bandwidth, which aggregates to −33 dBm noise power that passes the 33 MHz filter bandwidth. This verifies 18.8 dB SNR. Fig. 8*b* shows the results for input powers of −5 dBm, when the attenuators are set to −15.5 dB. The third order non-linear distortions are 200 kHz aside the desired signals, with 15.83 dB signal to distortion ratio. These two measurements affirm a 50 dB dynamic range of input powers that satisfy (4). The last pane in Fig. 8 suggests a recommended working point with an SNDR of 36.2 dB, obtained for −35 dBm input powers. The attenuator levels were set according to Fig. 3*b*, namely −10.5 dB and −15.5 dB, respectively.

Table 2 Device list

Device	Reference	Manufacturer
Analog board		
LNA	SPF5043	RFMD
Amplifier	Gali-21+	Mini-Circuits
Amplifier	Gali-84	Mini-Circuits
Attenuators	DAT-15R5-SN	Mini-Circuits
RF-switch	HMC284MS8G	Hittite Microwave Corporation
Mixer	SYM25-DHW	Mini-Circuits
Filter amplifier	ADA4817 (single-ended)	Analog Devices
Filter amplifier	ADA4932-1 (balanced)	Analog Devices
Digital board		
Shift-register	MC10EP142MNG	ON Semiconductors
VCO	ROS-2082-119+	Mini-Circuits
Synthesiser	ADF4106	Analog Devices
Crystal	5597ASX3SVT	European Crystal Organisation
ECL clock splitter	ADCLK925	Analog Devices
Amplifier	Gali-21+	Mini-Circuits
CPLD	EPM570T144	Altera

5.2 Periodicity and sign waveforms

To verify the periodicity of the mixing waveforms we used a spectrum analyser. Fig. 9*a* depicts the spectrum of $p_i(t)$, which consists of equalispaced Diracs, namely highly concentrated energy peaks, as expected for periodic waveforms. The spacing between the Diracs was measured as 19.212 MHz, validating the design choice of f_p in Table 1. The Dirac spectral lines appear steady, ensuring the periodicity of $p_i(t)$. The technical report [23] elaborates on additional verifications of the clock network which guarantee periodicity. Similar Dirac patterns were observed for $p_i(t)$, $1 \le i \le 4$. This proves that the internal clock-to-output delays of the chosen taps are deterministic. The exact values of these delays do not play a role in the MWC system, since the coefficients c_{il} are calibrated post-manufacturing.

The time-domain appearance of $p_i(t)$ is plotted in Fig. 9*b*, showing sign alternations that are far from rectangular

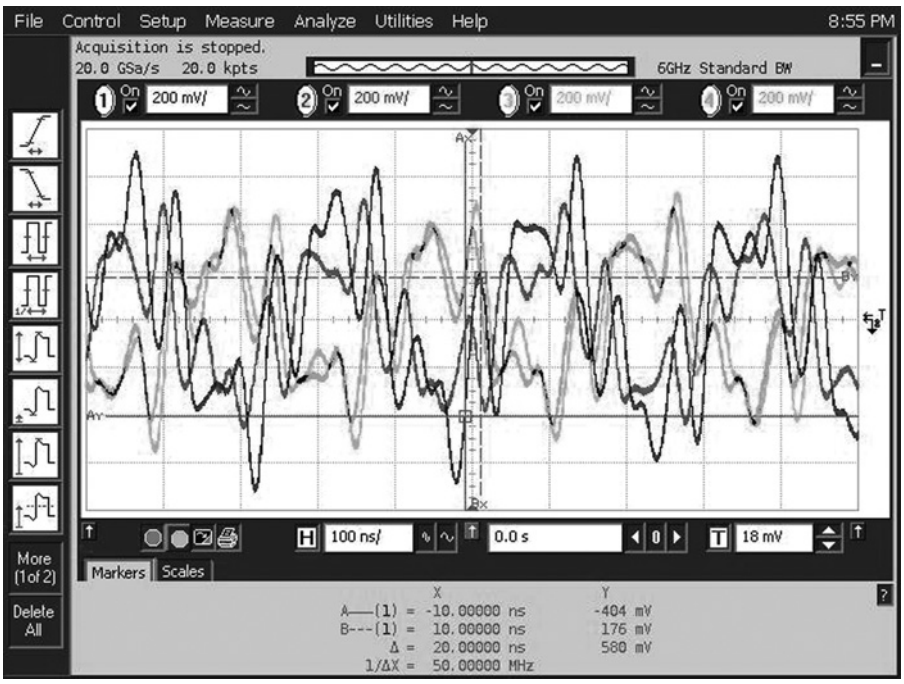

Fig. 7 *System outputs acquired by a four-channel scope*

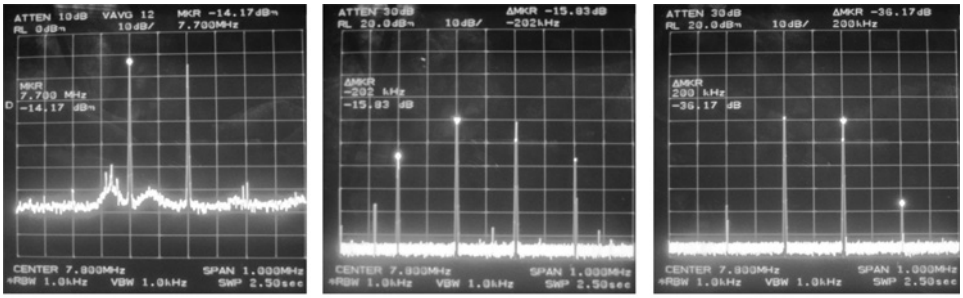

Fig. 8 *Output spectrum for two equal-power input tones*

a Measures the SNR for low input power
b and *c* Verify the non-linear distortion levels for high signal energies Noise floor in these panels is attributed to the spectrum analyser. In *c*, an external 20 dB attenuator is connected in series to the system output

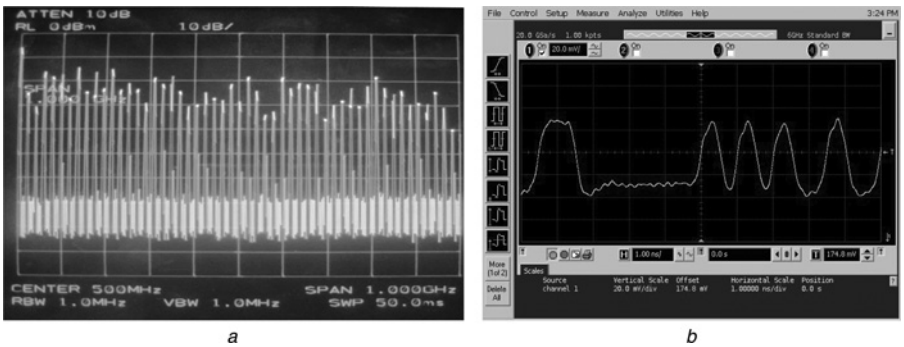

Fig. 9 *Sign-alternating periodic waveform*

a Spectrum
b Time-domain appearance

doi: 10.1049/iet-cds.2010.0147

transitions on the Nyquist grid. Fortunately, since periodicity is the only essential requirement of the MWC, the non-ideal time-domain appearance has no effect in practice. This is in contrast to the approach proposed in [9], which we discuss in more detail in Section 7.1. In this method, accurate time-domain properties of the periodic functions are required, thereby complicating the hardware design.

6 Sub-Nyquist demonstration

In this section, we demonstrate sub-Nyquist sampling and reconstruction using our designed system. For the paper to be self-contained, we first briefly outline the steps of the digital recovery algorithm.

6.1 Background: reconstruction algorithm

The heart of the reconstruction algorithm is a block, named continuous to finite (CTF), which constructs a finite-dimensional optimisation program from the MWC samples over a short time duration [7, 8].

Mathematically, the analog mixture boils down to the linear system [7]

$$\boldsymbol{y}[n] = \boldsymbol{C}\boldsymbol{z}[n] \tag{16}$$

where the vector $\boldsymbol{y}[n] = [y_1[n], \ldots, y_m[n]]^{\mathrm{T}}$ collects the measurements at $t = nT_{\mathrm{s}}$. The matrix $\boldsymbol{C}$ consists of the coefficients c_{il}, $\boldsymbol{z}[n] = [z_{-L}[n], \ldots, z_L[n]]^{\mathrm{T}}$ with $z_l[n]$ denoting an f_{p}-rate samples sequence of the lth spectrum slice of $x(t)$. The CTF builds a frame (or a basis) from the measurements using

$$\boldsymbol{y}[n] \xrightarrow{\text{Frame construct}} \boldsymbol{Q} = \sum_n \boldsymbol{y}[n]\boldsymbol{y}^{\mathrm{H}}[n] \xrightarrow{\text{Decompose}} \boldsymbol{Q} = \boldsymbol{V}\boldsymbol{V}^{\mathrm{H}} \tag{17}$$

where the (optional) decomposition allows the removal of the noise space. The finite dimensional system

$$\boldsymbol{V} = \boldsymbol{C}\boldsymbol{U} \tag{18}$$

is then solved for the sparsest matrix $\boldsymbol{U}$, namely the matrix with the minimal number of non-identically zero rows. This problem has been treated extensively within the compressed sensing community, and is often referred to as a multiple measurement vector (MMV) system. By now, there are plenty of off-the-shelf MMV solvers, any of which can be utilised in the recovery scheme of Fig. 10. Example techniques include mixed ℓ_1/ℓ_2 minimisation and greedy-type algorithms [24–28].

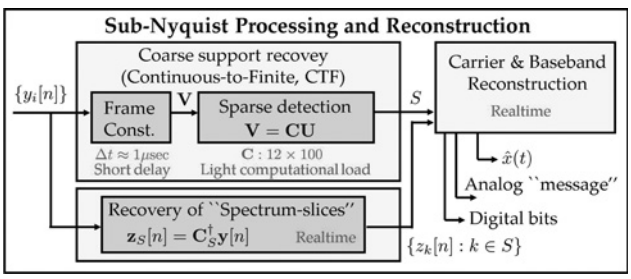

Fig. 10 *Reconstruction flow [7]*

The gist of the recovery flow is that $\boldsymbol{U}$, which solves the finite-dimensional system (18), indicates the support set S of active spectrum slices of the continuous input $x(t)$. The set S consists of the indices of those f_{p}-width spectrum slices that contain signal energy. Once the active slices are detected, their content is reconstructed in real time, including recovering the unknown carrier of each transmission and the information it encodes [10]. Together, the digital processing blocks of Fig. 10 provide a seamless interface to standard DSP packages.

6.2 Signal reconstruction

In order to prove the sub-Nyquist sampling capability of our design, we conducted two experiments. In the first, an HP-E4432B signal generator inputs a sinusoid to the MWC hardware. The four output channels were recorded using an Agilent Infiniium 54855A four-channel scope. All digital computations were carried out in Matlab. We varied the sinusoid frequency from 100 to 1100 MHz in steps of 5 MHz. The CTF outputs the spectral support at resolution $f_{\mathrm{p}} = 20$ MHz. We also executed the additional recovery blocks of Fig. 10 so that the algorithm estimates the input carrier frequency as well. The results, in the left pane of Fig. 11, demonstrate that out of these 200 experiments, there are only two outliers, which means 99% correct support and carrier estimation. This experiment provides initial confirmation that the hardware interacts smoothly with the digital algorithms.

Before proceeding to a more challenging reconstruction scenario, it is important to understand the reason for the outliers in Fig. 11. Theorem 2 in [7] ensures 100% correct recovery, provided that the sparse solution $\boldsymbol{U}$ of (18) is found. The latter, however, is known to be NP-hard. A common practice is to approximate sparse solutions with polynomial-time algorithms. We solve (18) using the multiple orthogonal matching pursuit (M-OMP) algorithm [29] which coincides with the true solution over a wide range of possible inputs; 99% of the cases in our experiments. The detection performance could have been improved for a higher number of sampling channels, say $m = 5$. Our design choice of a four-channel system represents a customary engineering compromise; saving the extra 25% in hardware size and digital computations of the $m = 5$ system, at the expense of not improving the last 1% of system performance. A possible route for improving the detection performance is to utilise other polynomial-time algorithms. We prove the concept with M-OMP due to its simplicity.

Our lab experiments also indicate an average of 10 ms duration for the digital computations, including the CTF support detection and the carrier estimation. The small dimensions of $\boldsymbol{C}$ (12×100) is what makes the MWC practically feasible from a computational perspective.

In the second experiment we verified correct support detection and signal reconstruction in the presence of three narrowband transmissions. Fig. 11 depicts the set up of three signal generators that were combined at the input terminal of the MWC prototype: an amplitude-modulated (AM) signal at 807.8 MHz with 100 kHz envelope, a frequency-modulation (FM) source at 631.2 MHz with 1.5 MHz frequency deviation and 10 kHz modulation rate, and a pure sine waveform at 981.9 MHz. The carrier positions were chosen so that their aliases overlay at baseband, as the photos in Fig. 11 demonstrate. The digital recovery algorithm was executed and detected the correct

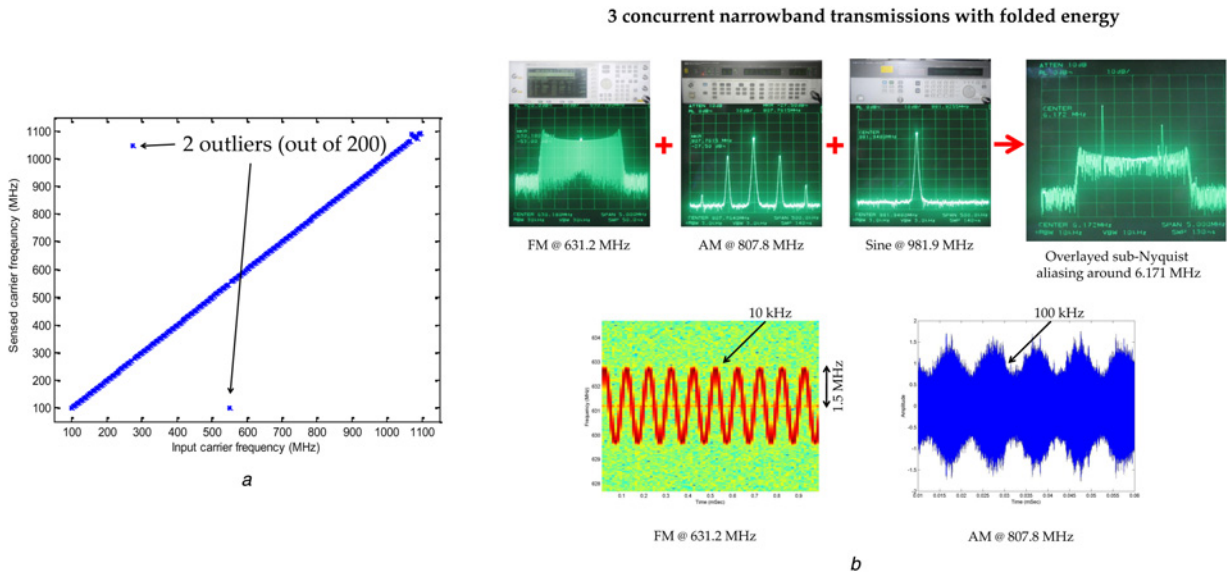

Fig. 11 *Demonstration of sub-Nyquist sampling and reconstruction*

a Recovered carrier frequency for a single sinusoid input
b Three signal generators are combined to the system input terminal
Spectrum of the lowrate samples (first channel) reveals overlapped aliasing at baseband. The recovery algorithm finds the correct carriers and reconstructs the original individual signals

support set S (CTF). The unknown carrier frequencies were estimated up to 10 kHz accuracy. In addition, the figure demonstrates correct reconstruction of the AM and FM signal contents.

The results of Fig. 11 connect between theory and practice. The same digital algorithms that were used in the numerical simulations of [7] are successfully applied herein on real data, acquired by the hardware. Reconstruction proof for a sampling strategy is a certificate that the theoretical principles are sufficiently robust to accommodate circuit non-idealities, that are inevitable in practice. A video recording of these experiments, relevant software packages and additional documentation are available online on the authors' websites:

http://www.technion.ac.il/~moshiko/hardware.html;
http://webee.technion.ac.il/Sites/People/YoninaEldar/Info/hardware.html.

7 Discussion

7.1 Related work

In [10] we present the Xampling framework, which proposes a high-level architecture for low-rate sampling in a union of subspaces, motivated by a series of pragmatic considerations. Sub-Nyquist sampling of multiband signals is a special case of this more general framework. Utilising Xampling, we conducted in [10] a thorough comparison between existing sub-Nyquist methods, including non-uniform sampling [8], random demodulation (RD) [9, 11], Nyquist folding [13] and the MWC [7]. As we discuss in detail in [10], non-uniform methods necessitate ADC devices whose front-end bandwidth exceed the wideband range of the input. In practice, non-uniform sampling can be realised by degenerating a full-rate time-interleaved ADC topology. Therefore while the number of interleaving branches is reduced, the actual savings are not substantial since each branch still requires a Nyquist-rate bandwidth. In addition, the remaining branches necessitate sophisticated digital algorithms for timing synchronisation and recovery [5, 6].

The RD approach [9] relies on dense discretisation of the input. The analog signal is modelled as the sum of Q integral harmonic frequencies

$$f(t) = \sum_{l=1}^{Q} a_l e^{j2\pi lt} \tag{19}$$

where only $K < Q$ harmonics are assumed to have non-zero amplitudes a_l. The RD flips the polarity of $f(t)$ randomly at the Nyquist rate Q. Measurements are generated serially by integrating over a duration of M Nyquist intervals and dumping the result. Owing to the dense discretisation, the reconstruction requires solving a linear system with dimensions $Q/M \times Q$, that is recovery complexity that scales with the Nyquist rate. Reconstructing a 1 MHz Nyquist-rate input boils down to an optimisation program with approximately 1 million unknowns. The approach is sensitive to model mismatch and requires large Q to approximate a truly analog input. In addition, it relies on sharp polarity flipping and accurate integration [10]; refer back to Fig. 9*b*, which demonstrates the difficulty in achieving exact time-domain properties at high rates. The analysis in [10] estimates that computational loads limit the strategy to inputs with Nyquist-rate less than $Q = 10$ MHz. A hardware prototype of the RD was demonstrated in [11] for a 800 kHz Nyquist-rate input and 100 kHz sampling rate. The digital reconstruction was carried out by a 160 MHz processor. A bank of RD samplers is reported in [12] with the same reconstruction techniques and complexities. The demonstration in [12] also uses signals with Nyquist rate below 1 MHz. For comparison, our system treats 2 GHz inputs while the number of unknowns in (18) is as low as 100.

Another interesting approach to sub-Nyquist sampling is the folding system of [13]. It introduces a deliberate jitter in

doi: 10.1049/iet-cds.2010.0147

an undersampling grid, which induces a phase modulation at baseband such that the modulation magnitude depends on the unknown carrier position. The experiments in [13] use a mixture of pure sinusoids, with a sampling rate substantially higher than minimal, thereby resulting in modulated versions that do not overlap at baseband. Most importantly, the method was not proved to reconstruct a mixture of RF transmissions. At present, it is difficult to compare between [13] and the MWC, since details on the reconstruction algorithm of [13] were not released yet for publication, although the approach seems interesting. To the best of our knowledge, this paper is the first to present a fully functioning sub-Nyquist system, with a wideband acquisition hardware, proven reconstruction capabilities for general inputs, and only light computational loads.

7.2 Future directions

We have chosen to implement a specific configuration of the MWC, in which the Nyquist rate of the input is around 2 GHz and the spectrum occupancy reaches 120 MHz. The sampling rate is as low as 280 MHz and slightly above the lowest possible rate for multiband signals with unknown carrier positions. In parallel to proving the feasibility of the method, our experiments revealed several directions for improvements.

Our board-level design uses a 2 GHz VCO which defines the input wideband range. Increasing the front-end bandwidth is possible for higher VCO frequencies. It would be interesting to approach the highest wideband range. The proposed digital board relies on the specific SR device we have chosen, which accepts 2.8 GHz clock rate. High-speed registers, up to 23 GHz and even 80 GHz clock rates, are already reported in the literature [30, 31]. RF technology is also capable of treating these frequencies. Therefore an extension of this work can be a prototype design which covers input wideband range, orders of magnitude higher than the conversion rates offered today by state-of-the-art commercial Nyquist-rate ADC devices.

Another interesting direction is optimising power consumption and area. These factors were not addressed in our proof-of-concept design. The current design consumes about 20 W, most of which is consumed by 50 Ω biasing currents of the ECL registers. An alternative realisation for the periodic waveform can take advantage of the fact that Gold/Kasami patterns are suitable for the MWC [16]. For example, a 511-length Gold sequence requires no more than a single SR device, which would save about 90% of the biasing power consumption. The area would also scale appropriately.

In addition to presenting the sub-Nyquist sampler, we emphasised two circuit challenges, that can be further investigated in a broader context: mixing a signal with multiple sinusoids and generating periodic waveforms with transients at the Nyquist rate. In order to achieve the desired effects, we employed standard devices, modified their specifications because of the non-ordinary usage and added auxiliary circuitry accordingly. Further study of these circuit structures, beyond the current application of sub-Nyquist sampling, may assist in developing alternative solutions for these tasks.

8 Acknowledgments

The authors appreciate the insightful comments and suggestions of the anonymous referees. The authors also thank the technical teams of both the Communication Laboratory (ComLab) and the High-Speed Digital Systems Laboratory (HSDSL) at the technion for investing time and resources in this project, Moshe Namer, Dr. Avraham Saad and Prof. Moshe Nazarathy for fruitful discussions on the preliminary architecture, and Shraga Krauss and Dr. Yossi Hipsh for many insightful comments. The help of Pavel Shpilberg and Mordechay Orbach is also gratefully acknowledged. M. Mishali is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

9 References

1 Nyquist, H.: 'Certain topics in telegraph transmission theory', *Trans. AIEE*, 1928, **47**, (2), pp. 617–644
2 Shannon, C.E.: 'Communication in the presence of noise', *Proc. IRE*, 1949, **37**, pp. 10–21
3 Eldar, Y.C., Oppenheim, A.V.: 'Filter bank reconstruction of bandlimited signals from nonuniform and generalized samples', *IEEE Trans. Signal Process.*, 2000, **48**, (10), pp. 2864–2875
4 Tian, Y., Zeng, D., Zeng, T.: 'Design and implementation of multifrequency front end using bandpass over sampling'. IET Int. Radar Conf., 2009, April 2009, pp. 1–4
5 El-Chammas, M., Murmann, B.: 'General analysis on the impact of phase-skew in time-interleaved adcs', *IEEE Trans. Circuits Syst. I, Reg. Papers*, 2009, **56**, (5), pp. 902–910
6 Nikaeen, P., Murmann, B.: 'Digital compensation of dynamic acquisition errors at the front-end of high-performance A/D converters', *IEEE Trans. Signal Process.*, 2009, **3**, (3), pp. 499–508
7 Mishali, M., Eldar, Y.C.: 'From theory to practice: sub-Nyquist sampling of sparse wideband analog signals', *IEEE J. Sel. Topics Signal Process.*, 2010, **4**, (2), pp. 375–391
8 Mishali, M., Eldar, Y.C.: 'Blind multi-band signal reconstruction: compressed sensing for analog signals', *IEEE Trans. Signal Process.*, 2009, **57**, (3), pp. 993–1009
9 Tropp, J.A., Laska, J.N., Duarte, M.F., Romberg, J.K., Baraniuk, R.G.: 'Beyond Nyquist: efficient sampling of sparse bandlimited signals', *IEEE Trans. Inf. Theory*, 2010, **56**, (1), pp. 520–544
10 Mishali, M., Eldar, Y.C., Elron, A.: 'Xampling: signal acquisition and processing in union of subspaces'. CCIT report no. 747, EE Dept., Technion; [Online] arXiv.org 0911.0519 Report, October 2009
11 Ragheb, T., Laska, J.N., Nejati, H., Kirolos, S., Baraniuk, R.G., Massoud, Y.: 'A prototype hardware for random demodulation based compressive analog-to-digital conversion'. 51st Midwest Symp. on Circuits and Systems, 2008, MWSCAS 2008, 2008, pp. 37–40
12 Yu, Z., Chen, X., Hoyos, S., Sadler, B.M., Gong, J., Qian, C.: 'Mixed-signal parallel compressive spectrum sensing for cognitive radios', *Int. J. Dig. Multimedia Broadcast.*, 2010
13 Fudge, G.L., Bland, R.E., Chivers, M.A., Ravindran, S., Haupt, J., Pace, P.E.: 'A Nyquist folding analog-to-information receiver'. Proc. 42nd Asilomar Conf. on Signals, Systems and Computers, October 2008, pp. 541–545
14 Lin, Y.-P., Vaidyanathan, P.P.: 'Periodically nonuniform sampling of bandpass signals', *IEEE Trans. Circuits Syst. II*, 1998, **45**, (3), pp. 340–351
15 Herley, C., Wong, P.W.: 'Minimum rate sampling and reconstruction of signals with arbitrary frequency support', *IEEE Trans. Inf. Theory*, 1999, **45**, (5), pp. 1555–1564
16 Mishali, M., Eldar, Y.C.: 'Expected-RIP: conditioning of the modulated wideband converter'. IEEE Information Theory Workshop, 2009, ITW 2009, October 2009, pp. 343–347
17 Chen, Y., Mishali, M., Eldar, Y.C., HeroIIIA.O.: 'Modulated wideband converter with non-ideal lowpass filters'. ICASSP 2010, 2010, pp. 3630–3633
18 Analog Devices Corp.: 'A/D converters', http://www.analog.com/en/analog-to-digital-converters/ad-converters/products/index.html, accessed 2009
19 Texas Instruments Corp.: 'Data converters', http://focus.ti.com/analog/docs/dataconvertershome.tsp, accessed 2009
20 Friis, H.T.: 'Noise figures of radio receivers', *Proc. IRE*, 1944, **32**, (7), pp. 419–422
21 RF cafe website: 'Cascaded 2-tone, 3rd-order compression point (IP3)', http://www.rfcafe.com/references/electrical/ip3.htm
22 Carvalho, N.B., Pedro, J.C.: 'Compact formulas to relate ACPR and NPR to two-tone IMR and IP3', *Microw. J.*, 1999, **42**, (12), pp. 70–84

23 Mishali, M., Eldar, Y.C., Dounaevsky, O., Shoshan, E.: 'Xampling: analog to digital at sub-Nyquist rates'. CCIT report no. 751, EE Dept., Technion; [Online] arXiv.org 0912.2495 Report, December 2009

24 Eldar, Y.C., Mishali, M.: 'Robust recovery of signals from a structured union of subspaces', *IEEE Trans. Inf. Theory*, 2009, **55**, (11), pp. 5302–5316

25 Tropp, J.A.: 'Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit', *Signal Process. (Special Issue on Sparse Approximations in Signal and Image Processing)*, 2006, **86**, pp. 572–588

26 Tropp, J.A.: 'Algorithms for simultaneous sparse approximation. Part II: Convex relaxation', *Signal Process. (Special Issue on Sparse Approximations in Signal and Image Processing)*, 2006, **86**, pp. 589–602

27 Davies, M.E., Eldar, Y.C.: 'Rank awareness in joint sparse recovery', *IEEE Trans. Inf. Theory*, 2010 submitted [Online] arXiv:1004.4529

28 Mishali, M., Eldar, Y.C.: 'Reduce and boost: recovering arbitrary sets of jointly sparse vectors', *IEEE Trans. Signal Process.*, 2008, **56**, (10), pp. 4692–4702

29 Cotter, S.F., Rao, B.D., Engan, K., Kreutz-Delgado, K.: 'Sparse solutions to linear inverse problems with multiple measurement vectors', *IEEE Trans. Signal Process.*, 2005, **53**, (7), pp. 2477–2488

30 Laskin, E., Voinigescu, S.P.: 'A 60 mW per Lane, 4×23-Gb/s $2^7 - 1$ PRBS generator', *IEEE J. Solid-State Circuits*, 2006, **41**, (10), pp. 2198–2208

31 Dickson, T.O., Laskin, E., Khalid, I.: 'An 80-Gb/s $2^{31} - 1$ pseudorandom binary sequence generator in SiGe BiCMOS technology', *IEEE J. Solid-State Circuits*, 2005, **40**, (12), pp. 2735–2745

doi: 10.1049/iet-cds.2010.0147