Title
Filtering by Aliasing and its application to Reconfigurable Filtering and Compressive Signal Acquisition

Author
Rachid, Mansour

Publication Date
2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Filtering by Aliasing

and its application to

Reconfigurable Filtering and Compressive Signal Acquisition

A dissertation submitted in partial satisfaction of the

requirements for the degree Doctor of Philosophy

in Electrical Engineering

by

Mansour Rachid

2013

ABSTRACT OF THE DISSERTATION

Filtering by Aliasing and its application to

Reconfigurable Filtering and Compressive Signal Acquisition

by

Mansour Rachid
Doctor of Philosophy in Electrical Engineering
University of California, Los Angeles, 2013
Professor Babak Daneshrad, Chair

The communication systems community has been working towards integrated Software-Defined Radios (SDRs) and Cognitive Radios (CRs) that can reduce cost and enhance connectivity. In light of the technology bottleneck at the analog-to-digital convertor (ADC) and the inapplicability of off-chip filters, the integrated analog front-end is entrusted with the task of sharp, linear, and programmable signal selection required for SDRs and CRs. Traditional analog filtering techniques, however, incur a high penalty in power consumption, area, and linearity to provide the required sharpness and programmability. Similarly, recent efforts that use compressive sensing to acquire wideband spectra have also faced a bottleneck in the complexity and re-configurability of the analog measurement frontend.

Towards enabling SDRs and CRs, this dissertation proposes a new perspective on the design of anti-alias filters that defies the traditional tradeoff between cost, linearity, and programmability. The technique, termed *Filtering by Aliasing* (FA), anticipates the aliasing operation at the sampler instead of avoiding it. The pre-sampling circuitry is modulated, using the high-speed switching techniques popular in state-of-the-art receivers, to provide significantly enhanced filtering responses at the sampling instances. The dissertation describes how the FA technique, by varying the resistor of a single-pole passive RC filter for example, provides programmable anti-alias filtering comparable to a 7^{th}-order Butterworth filter.

On the compressive sensing front, this dissertation proposes a new approach to the acquisition of sparse spectra using *Random Filtering by Aliasing* (RFA). RFA acknowledges the existence of noise in realistic spectra and accordingly simplifies the analog measurement stage, moving most of the complexity to the low-cost, highly reconfigurable digital domain. RFA achieves significantly better resolution, lower cost, and better programmability than existing schemes.

The dissertation of Mansour Rachid is approved.

Milos Ercegovac

Lieven Vandenberghe

Sudhakar Pamarti

Babak Daneshrad, Committee Chair

University of California, Los Angeles

2013

To Mom and Dad

TABLE OF CONTENTS

VITA

2006	**B.E.**, Computer and Communications Engineering, *distinction* American University of Beirut Beirut, Lebanon
2006	Dean's Award for Creative Achievement Computer and Comm. Engineering Department American University of Beirut
2006-2012	Graduate Student Research Assistant Electrical Engineering Department University of California, Los Angeles
2007-2012	Teaching Assistant Electrical Engineering Department University of California, Los Angeles
2008	Hardware Design Intern Silvus Technologies, Inc. Los Angeles, California
2009	**M.S.**, Electrical Engineering University of California, Los Angeles
2009	Departmental Fellowship Award Electrical Engineering Department University of California, Los Angeles
2011	Communication Systems Intern Maxlinear, Inc. Carlsbad, California

I. INTRODUCTION

This dissertation proposes a novel technique, termed Filtering by Aliasing (FA), for the design and implementation of anti-alias filtering. The dissertation also proposes, based on FA, a novel technique to perform compressive signal acquisition termed Random Filtering by Aliasing (RFA). The Filtering by Aliasing technique evolved from a reconsideration of the traditional standalone design of the continuous-time front-end of a communications receiver. This chapter describes said evolution of the FA concept and introduces the material in this dissertation. The applications to which the proposed FA and RFA are well suited will be motivated in the following chapters as appropriate.

In simple words, a modern communications receiver is intended to extract digital discrete-time (DT) data about the signal of interest (SOI) from some analog continuous-time (CT) input waveform, as illustrated in Fig. I-1. The analog front-end is thus entrusted with suppressing unwanted signals (in addition to down-conversion and amplification) such that the sampling rate and resolution of the analog-to-digital converter (ADC) are minimized. We have been traditionally taught that unwanted signals need to be suppressed before sampling to avoid spectral aliasing. Aliasing, however, is a

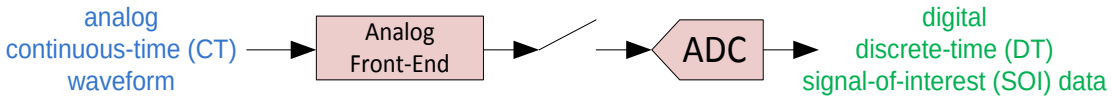

Fig. I-1 An illustration of the role of a communications receiver. Traditionally, the analog front-end is designed to suppress unwanted interferers prior to sampling, the proposed Filtering by Aliasing reconsidered this approach.

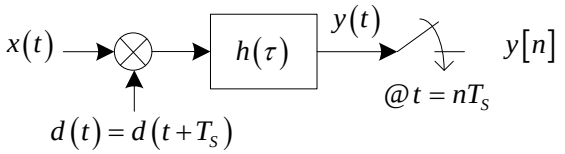

Fig. I-2 The variable-gain FA system: the mixer and sampler operate at the same periodicity rate allowing for aliasing to be explicitly involved in the signal processing.

known deterministic transformation. We need only be concerned with the suppression of unwanted signals at the output of the sampler. *Can we then incorporate the aliasing operation into the anti-aliasing task itself?*

I.1 EXPLICIT FILTERING BY ALIASING

The system in Fig. I-2 explicitly involves the aliasing operation in the anti-aliasing task. The system, which will hereafter be referred to as the *variable-gain FA* system, performs spreading of the input $x(t)$ with a T_S-periodic $d(t)$, filtering with some filter $h(\tau)$, and sampling at rate $F_S = 1/T_S$. The spreading with $d(t)$ creates shifted and weighted images of the input spectrum $X(f)$ as shown in Fig. I-3b for a hypothetical $X(f)$ in Fig. I-3a. The spread spectrum is then filtered with the physical filter $H(f)$ and folded by aliasing as illustrated in Fig. I-3d,e. As shown in Fig. I-3e, we hope to design the FA system such that the unwanted signals appear successfully suppressed after aliasing.

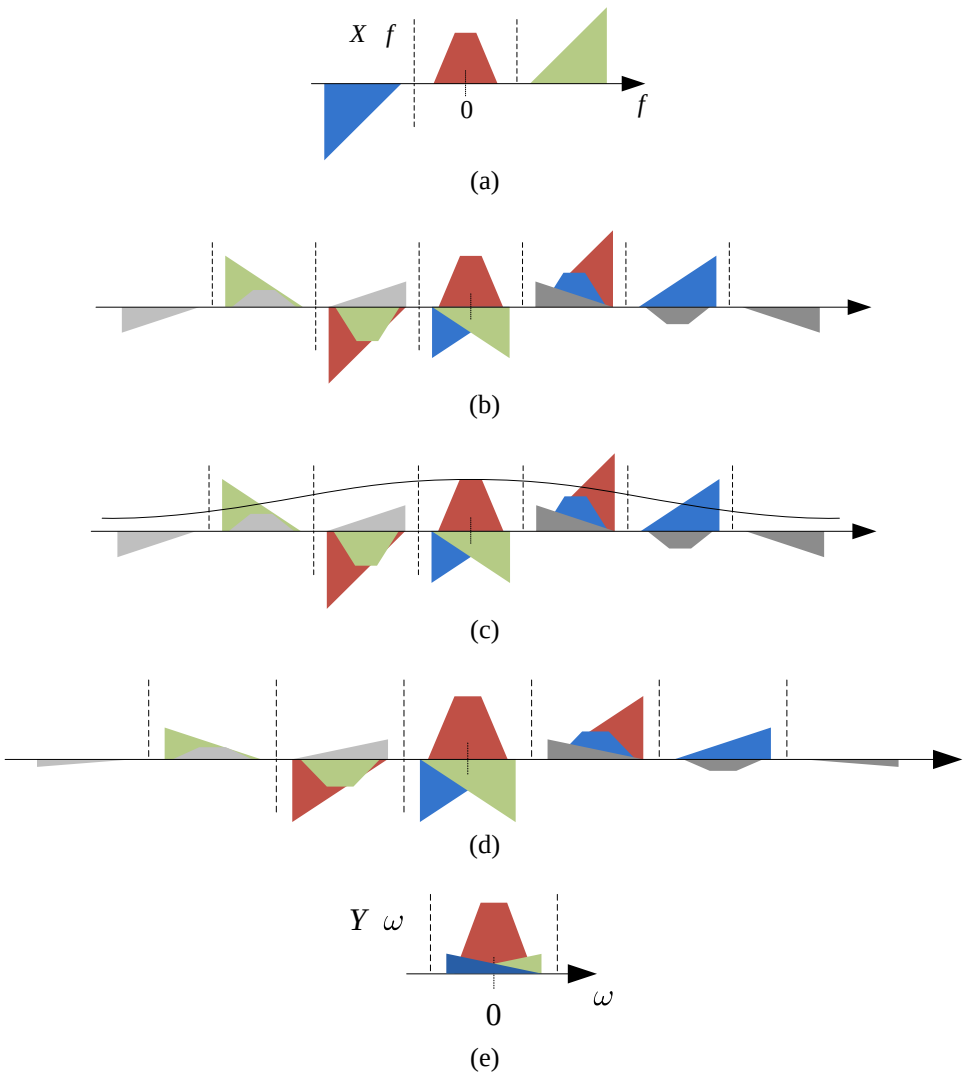

Fig. I-3 An illustration of the frequency-domain operations of the variable-gain FA system. (a) Input spectrum, (b) spread spectrum, (c) showing filter $H(f)$, (d) after filtering, (e) aliased.

The Concept of the Apparent Filter

Since the spreading and aliasing operations both cause frequency translations over intervals of size F_S, the sampled output $y[n]$ in Fig. I-2 cannot differentiate the operations. We can thus talk about an apparent filter $g(\tau)$ that the input $x(t)$ *appears* to be subjected to

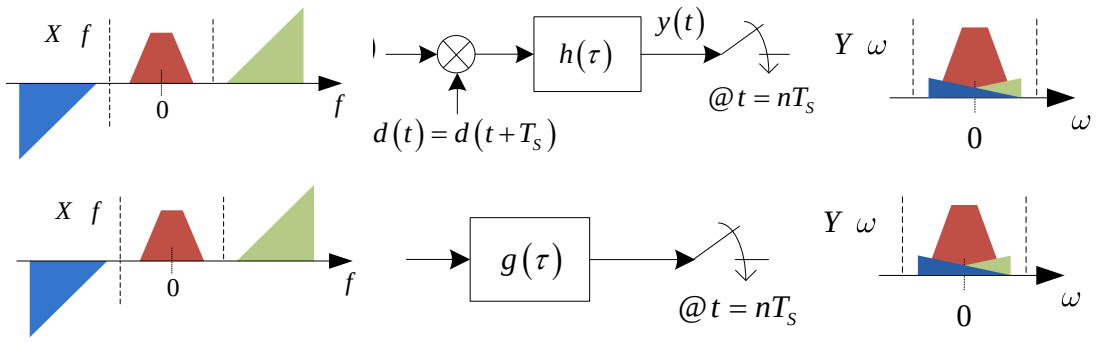

Fig. I-4 (top) the variable-gain FA system transforms a CT input spectrum to a DT output spectrum using periodically-time varying operations. (bottom) As far the DT output is concerned, the CT input appears subjected to an apparent LTI filter $g(\tau)$ prior to sampling.

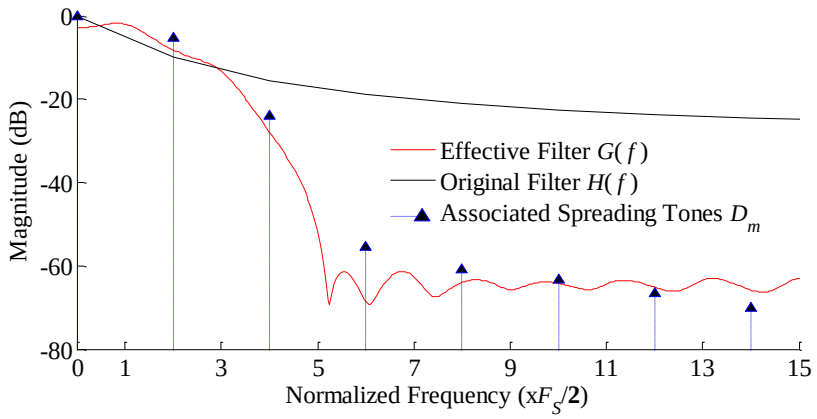

Fig. I-5 An example scenario where the variable-gain FA system employs a single-pole filter $H(f)$ and a spreader $d(t)$ whose frequency content is D_m. The input appears filtered by the strong anti-aliasing response $G(f)$.

before sampling. As illustrated in Fig. I-4, the apparent filter $g(\tau)$ is equivalent to the Filtering by Aliasing system *from the perspective of the sampled output*.

As will be described in chapter II of this dissertation, the apparent filter $g(\tau)$ can be designed to be significantly sharper than the physical filter $h(\tau)$. For example, Fig. I-5

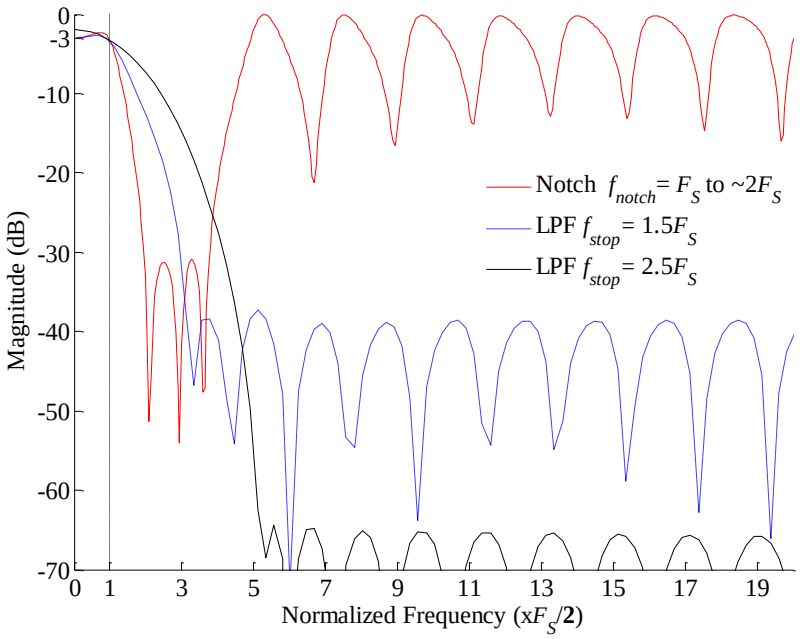

Fig. I-6 Four examples of the different anti-aliasing responses $g(\tau)$ obtained using a variable-gain FA system with a single-pole filter $h(\tau)$.

shows a single-pole $H(f)$, optimally chosen spectral tones D_m of $d(t)$, and the resulting apparent filter $G(f)$ seen by the output of the sampler.

As shown in Fig. I-6, the apparent filter $G(f)$ is also reconfigurable in terms of type, bandwidth, and roll-off using the periodic spreading signal $d(t)$. As will be discussed in chapter II, the ability to provide sharp digitally-programmable filtering from simple analog structures is highly desirable for future communication systems such as Software-defined Radios (SDRs) and Cognitive Radios (CRs).

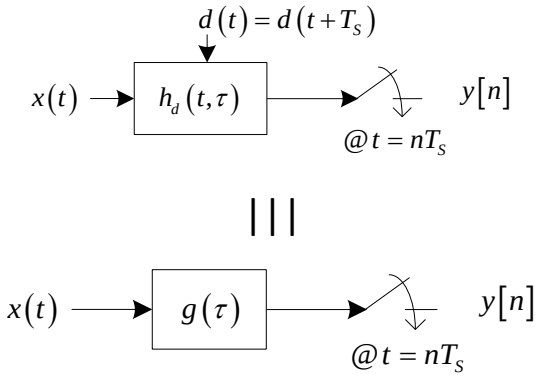

Fig. I-7 General Filtering by Aliasing. (top) A LPTV system $h_d(t,\tau)$ of the same period as the sampler appears to behave like an apparent LTI system $g(\tau)$ followed by a sampler (bottom).

I.2 General Filtering by Aliasing

Taking a closer look at the FA operation in the time domain, we realize that the FA system essentially focuses on the system's behavior between the continuous-time (CT) $x(t)$ and the discrete-time (DT) $y[n]$. In other words, involving the aliasing operation in the anti-aliasing task is equivalent, in the time-domain, to ignoring the system's behavior at all instances but the sampling instance. This leads us to consider a more general Filtering by Aliasing system where, instead of explicitly performing spreading and filtering prior to sampling, we can perform arbitrary linear periodically time-varying (LPTV) operations of period T_S. Similarly then, as illustrated in Fig. I-7, we can talk about an *apparent* linear time invariant (LTI) filter $g(\tau)$ that appears to operate on the input $x(t)$ from the perspective of the sampled output $y[n]$.

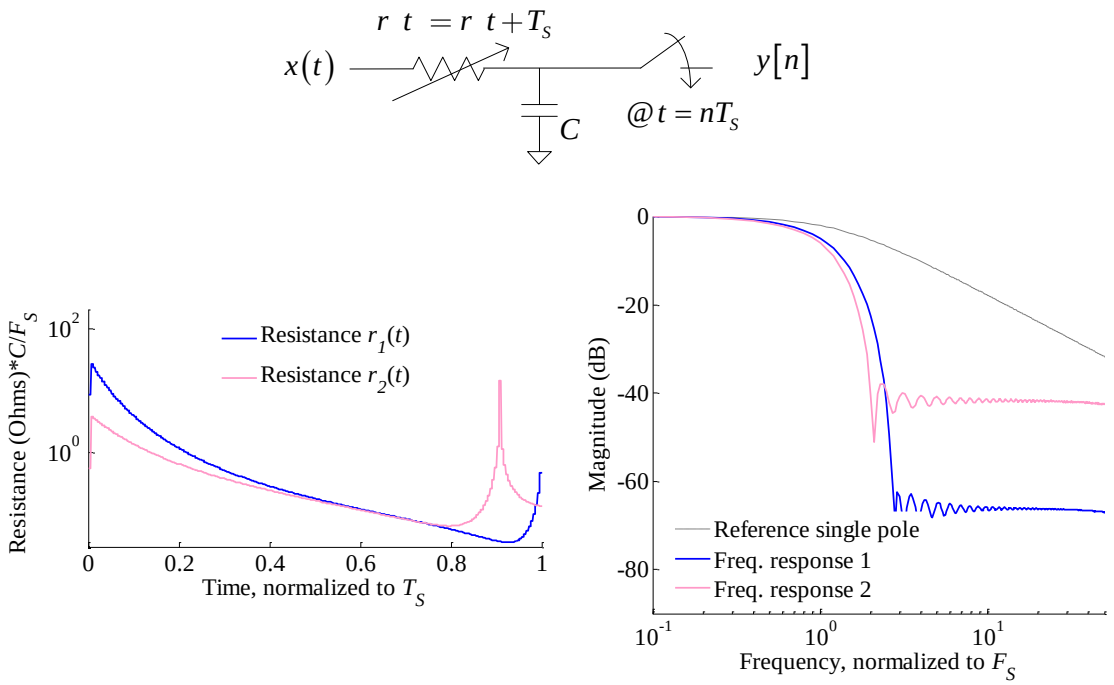

Fig. I-8 (top) An example of a general FA system where the pole of a single-pole passive filter is periodically varied by adjusting the resistance. (bottom) Two example resistance waveforms and the associated apparent filtering responses.

We can thus state the *general* Filtering by Aliasing (FA) concept as:

> *To periodically vary one or more components of a physical filter h(τ) prior to a sampler, at the sampling rate, such that a desired LTI filter response is apparent at the sampling instances.*

As an example, chapter II will describe how the resistor in a simple single-pole passive RC filter, such as that in Fig. I-8 (top), can be periodically modulated to obtain very sharp apparent filters such as shown on the bottom of Fig. I-8. In such a Filtering by Aliasing system, the resistance value $r(t)$ represents the digital control signal $d(t)$ from Fig. I-7. Unlike the variable-gain FA system then, the generalized FA concept allows us

to obtain drastic improvement in filter sharpness and reconfigurability without the need for active stages in the signal path. This result is unintuitive and potentially game-changing because of its potential use at the radio frequency (RF) front-end of a communications receiver, the details of which are left to chapter II.

I.3 Random Filtering by Aliasing

Given that Filtering by Aliasing can provide reconfigurable anti-alias filtering profiles by manipulating some control sequence $d(t)$, the FA technique can be readily applied towards random filtering by using, for example, a random ±1 sequence $d(t)$. Random filtering is desirable in compressive sensing (CS) [1], [2] or compressive sampling applications where a signal that is sparse in some domain can be recorded at a lower rate than its Nyquist rate. To make a compressive measurement, multiple measurement paths record the input signal at a sub-sampling rate after applying different, preferably random, transformations on the high-rate signal.

Chapter III will describe the existing techniques for the compressive sensing of spectrally-sparse continuous-time (SSCT) signals, a particularly tough CS problem, and how the use of a proposed Random Filtering by Aliasing (RFA) approach can address the drawbacks of existing schemes. RFA is a system-wide approach –built around FA analog measurements –that exhibits lower cost, higher spectral resolution, and better reconfigurability than existing schemes.

II. SELECTIVE FILTERING BY ALIASING

II.1 INTRODUCTION

It is well known that sampling a continuous-time signal at some rate F_S causes the aliasing of the spectral content lying at frequencies equal to or higher than $F_S/2$. Accordingly, to acquire a signal of interest with sufficient signal integrity, anti-alias filtering of unwanted spectral content should be performed. Traditionally, anti-alias filtering is performed using explicit filtering prior to sampling. In communications receivers in particular, anti-aliasing is traditionally jointly achieved using strong fixed off-chip filters as well as one or more filtering stages in the on-chip analog receive path, prior to the analog-to-digital converter (ADC). This solution becomes impractical as radio technology migrates towards wideband software-defined [3], [4], multi-standard [5], and cognitive [6] radios.

A software-defined radio (SDR), as the name suggests, is one that can communicate multiple types of signals of interest on demand and at runtime. Such a radio can potentially reduce the cost of realizing multiple communication standards on mobile devices and provide forward compatibility with changing protocols and spectrum allocation. A cognitive radio (CR) is another hypothetical radio that can detect signals of interest in a wideband spectrum, such as that of Fig. II-1, for the purpose of decoding such signals or identifying unused spectrum and utilizing it for secondary-user communication. The feasibility of a wideband CR will depend on the realization of efficient SDRs, which in turn requires efficient reconfigurable radio receivers.

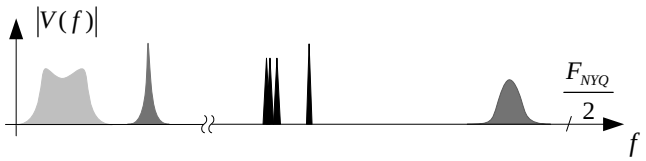

Fig. II-1 An illustration of a spectral scenario suitable for cognitive radio where multiple signals with different spectral location, bandwidth, power, and interference environment might exist. CR could be interested in decoding such signals or finding unoccupied spectrum.

A. The Need for Reconfigurable Integrated Filtering

Reconfigurable radio receivers for SDRs and CRs are expected to dynamically acquire or monitor one or more signals of interest whose frequency locations and bandwidths are not known apriori. Moreover, the interference and blocker profiles around these signals of interest are also dynamic and unknown at the time of manufacturing. The ideal reconfigurable radio places a wideband ADC at the antenna [3], [7], as shown in Fig. II-2a, and performs all the required reconfigurable operations in the digital domain. ADC technology, however, will not be able to simultaneously support the wide bandwidths and high dynamic range of interest in the near future [4], [7]-[9]. As such, the analog front-end is expected to stay as a signal conditioning stage prior to a practical ADC. Traditionally, as shown in Fig. II-2b, fixed off-chip filters are used to take care of strong out-of-band interference and simplify the task of the on-chip analog circuitry. Off-chip filters, however, are costly and inflexible, making them unsuitable for an integrated reconfigurable radio [4], [7], [8], [10]. Accordingly, for SDRs and CRs to become a reality, (a) sharp, (b) programmable, and (c) high dynamic range analog filtering is required from the integrated analog front-end, as illustrated in Fig. II-2c.

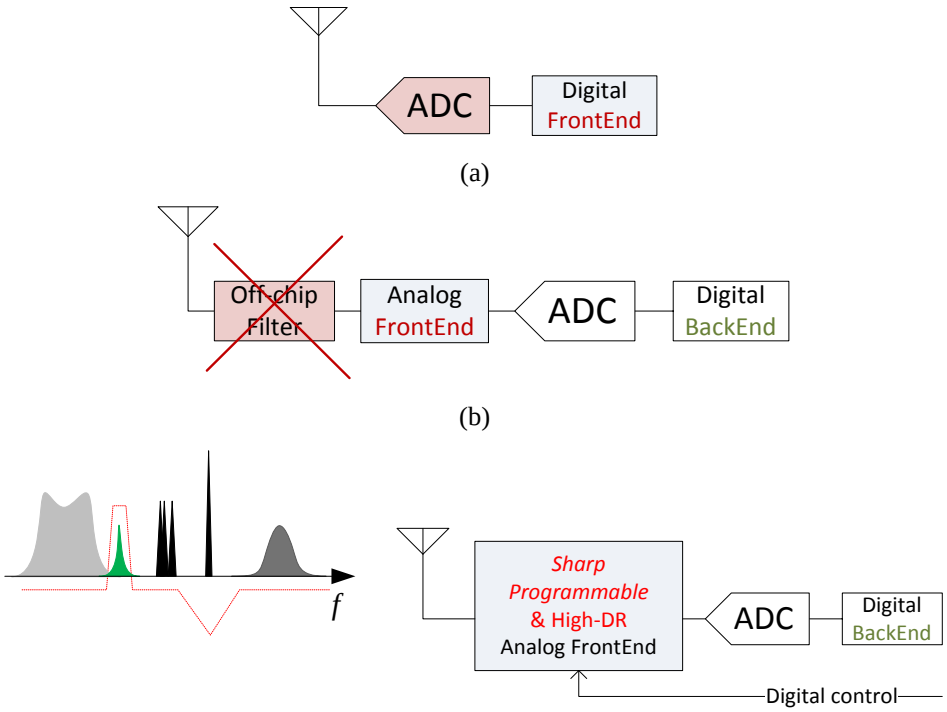

Fig. II-2 An illustration of why reconfigurable on-chip analog filtering is required for SDRs and CRs. (a) High speed high resolution ADCs are not practical. (b) off-chip channel selection filters are not suited for reconfigurable integrated radios. (c) the on-chip circuitry has to handle challenging filtering requirements.

B. Existing Baseband Filtering Solutions

A traditional CT anti-alias filter, or a cascade of them, with a sharp-enough roll-off is costly to integrate on chip. Such a filter also incurs further area and power consumption cost only to provide limited bandwidth reconfigurability [7]-[9], [11]-[13]. The work in [14] describes a list of recent analog filters in the literature showing strong trade-offs between filter order and reconfigurabilty.

Responding to the need for reconfigurable anti-alias filters, some works in the literature considered anti-aliasing techniques that rely on a particular type of sampling, termed integration (or charge) sampling (IS) [15]. IS realizes an implicit sinc-response filter by integrating the signal over a time window prior to the sampling operation. Systems based on IS were proposed in [4], [11]-[13], [16]-[18] as possible candidates for reconfigurable filtering owing to their programmable bandwidth. As they stand in the literature, however, IS-based anti-aliasing filters suffer from a prohibitively slow roll-off requiring a high oversampling rate and several additional stages for proper operation, e.g. [4], [18]. As a result, IS-based anti-aliasing similarly incurs high power consumption and non-linearity penalties. Fig. II-3 plots on the left the linearity, in terms of the input referred intercept point (IIP3), of several recent CT and IS-based baseband filters in the literature against their power consumption normalized to the number of poles or zeros.

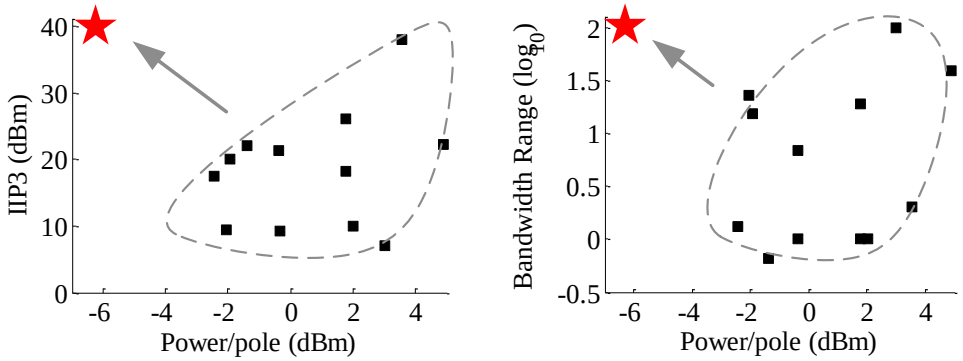

Fig. II-3 (left) The linearity of recent baseband filtering examples in the literature plotted in black as the IIP3 vs. normalized power consumption. Assumed knee frequency of 10MHz. (right) The bandwidth tuning range of the same pool of baseband filters. The red stars signify where one would like to operate for a reconfigurable filtering solution to be of value.

Fig. II-3 also plots, on the right, the reported or expected bandwidth tuning range of the same filters against the normalized power consumption. As the figures show, there is an apparent trade-off between the power consumption and performance (linearity and/or bandwidth tuning range) of existing filtering techniques. Of course, there also exists similar trade-offs between power consumption and filter sharpness, as well as between chip area and performance. For reconfigurable radio receivers to have a significant advantage over dedicated receivers, a novel solution that breaks away from the traditional trade-offs is required.

C. Existing Radio Frequency Solutions

As was mentioned earlier, an off-chip RF band-selection filter is traditionally used to suppress out-of-band interference. In an integrated wideband receiver, however, the on-chip active circuitry is exposed to unwanted blockers that desensitize the front-end.

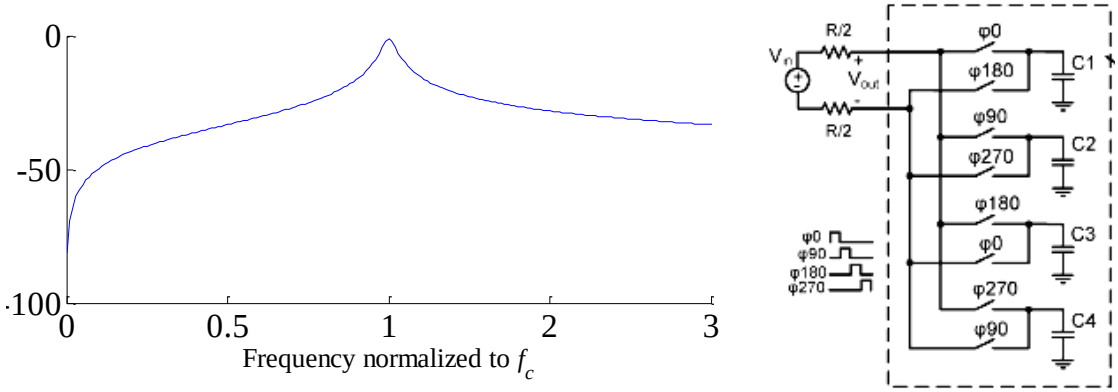

Fig. II-4 (right) An example from [20] of a 4-path differential N-path filter that translates the low-pass RC response to a bandpass filter at the switching frequency. (left) The ideal N-path filter with infinite *N* behaves like an RLC filter, suitable for faraway blockers only.

Recent RF techniques for integrated receivers combat such blockers using current-mode passive mixing followed by single-pole filtering, e.g. [10], [19]-[20], or single-pole impedance translation at the output of a low-noise amplifier (LNA), e.g. [20]. Impedance translation is the creation of an inductor-less "high-Q" bandpass filter at a tunable carrier frequency by employing an N-path [21] structure such as shown in Fig. II-4 on the right [20]. The achievable selectivity using such techniques is dictated by the impedance placed at the "baseband" end of the mixer and is thus limited to the single-pole behavior of an RC filter in typical implementations (e.g. Fig. II-4 on the left).

D. Filtering by Aliasing – A Potential Solution

This chapter will dive into the details of the proposed Filtering by Aliasing concept, particularly for the sharp reconfigurable filtering application.

We will show that, using a simple single-pole physical filter $h(\tau)$, *programmable* low-pass or band-pass anti-alias filters with more than 65dB of stop-band suppression can be achieved by periodically modulating the gain or the pole of the filter. When considering transition-band size, such anti-alias filters are comparable to 4^{th} order and 7^{th} order Elliptical and Butterworth filters respectively. We will also show that the FA approach can yield sharper anti-alias filtering by increasing the number of FA sampling channels or employing a higher-order (sharper) $h(\tau)$.

To demonstrate the practical realization of the proposed solution, we will also address the relevant hardware sensitivities and show that the variable-gain FA system (Fig. I-2) can be realized using similar components as in IS-based filters [4], [11] and other proven techniques in the literature, e.g. [22] while breaking off from the power

consumption vs. performance trade-off. More importantly, we will also show how the variable-pole FA system (Fig. I-8) can be realized using similar passive components as state-of-the-art N-path filters employed at RF frequencies. Despite the structural similarities, when compared to N-path filters, the FA solution exhibits significantly better selectivity and programmability, potentially bringing the sampler and ADC ever so close to the antenna.

Section II.2 will deal with the variable-gain FA system. The derivation of the apparent filter, its design, and its performance will be described and analyzed. Most of the qualitative analysis and the optimization technique details apply similarly to the variable-pole FA system and are thus omitted from section II.3 where the latter is discussed. Section II.4 examines the proposed FA technique within the context of prior art in terms of fundamentals and results.

II.2 VARIABLE-GAIN FILTERING BY ALIASING

Consider the proposed analog acquisition hardware of Fig. II-5a. The input $x(t)$ is mixed with a periodic signal $d(t)$ of period T_S and filtered by a filter having the impulse response $h(\tau)$. The output $y(t)$ of the filter, as well as its frequency response $Y(f)$, are given by (2.1) and (2.2) respectively. In (2.2), we made use of the periodicity of $d(t)$ to reduce $D(f)$ to tone values D_m at locations $mF_S = m/T_S$.

$$y(t) = \int_{-\infty}^{+\infty} x(\tau) d(\tau) h(t-\tau)\, d\tau \tag{2.1}$$

$$\begin{aligned} Y(f) &= H(f) \int_{-\infty}^{+\infty} D(\lambda) X(f-\lambda) d\lambda \\ &= H(f) \sum_{m=-\infty}^{+\infty} D_m X\left(f - mF_p\right) \end{aligned} \qquad f \in [\quad , -] \tag{2.2}$$

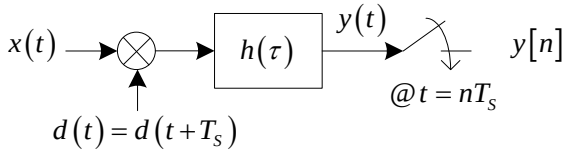

(a)

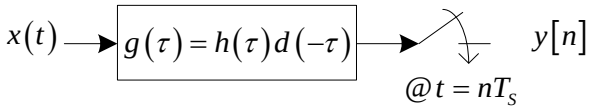

(b)

Fig. II-5 (a) The variable-gain FA analog acquisition system. (b) The apparent filter as seen by the sampled output.

After sampling at $t = nT_S$, the output $y[n]$ of the sampler, and its frequency response $Y(\omega)$ are given by (2.3) and (2.4) respectively. Note that $F_S = 1/T_S$.

$$y[n] = \int_{-\infty}^{+\infty} x(\tau) d(\tau) h(nT_S - \tau) d\tau \tag{2.3}$$

$$Y(\omega) = \frac{1}{T_S} \sum_{k=-\infty}^{+\infty} H\left(\frac{\omega F_S}{2\pi} + kF_S\right) \sum_{m=-\infty}^{+\infty} D_m X\left(\frac{\omega F_S}{2\pi} + (k-m)F_S\right) \qquad \omega \in [-\pi, \pi] \tag{2.4}$$

Notice from (2.2) that the output $Y(f)$ of the filter, prior to sampling, already contains aliases of the spectrum $X(f)$. This is intentional because all spectral slices of $Y(f)$, i.e. the spectral content of bandwidth F_S centered at multiples of F_S, are anticipated to alias at the sampler, as given by (2.4). The FA system thus attempts to take advantage of the anticipated aliasing to improve the apparent filter seen by the sampler output.

A. On the Apparent Anti-Aliasing Filter

As far as the output samples $y[n]$ are concerned, the input $x(t)$ appears to be filtered with a time-invariant impulse response $g(\tau)$ given by (2.5) and shown in Fig. II-5b. To arrive at $g(\tau)$, we replace $d(t)$ in (2.3) by a circularly shifted version $d(t\text{-}nT_S)$, identical to $d(t)$. The frequency response $G(f)$ is the Fourier transform of $g(\tau)$ given by (2.6). The apparent filter $G(f)$ is thus a linear combination of copies of the original filter response $H(f)$ shifted by mF_S and weighted by D_m.

$$g(t) = h(t)\,d(-t + nT_S) = h(t)\,d(-t) \tag{2.5}$$

$$G(f) = \mathbb{F}\{h(\tau)\,d(-t)\} = \sum_{m=-\infty}^{+\infty} D_m H(f + mF_S) \qquad f \in [-\infty, +\infty] \tag{2.6}$$

It is informative to also derive $G(f)$ from (2.4) in order to highlight the effect of the aliasing operation and intuitively understand the apparent filter. Performing a change of variables $l = k - m$ in (2.4) yields (2.7).

$$Y(\omega) = \frac{1}{T_S} \sum_{l=-\infty}^{+\infty} X\left(\frac{\omega F_S}{2\pi} + lF_S\right) \sum_{m=-\infty}^{+\infty} D_m H\left(\frac{\omega F_S}{2\pi} + lF_S + mF_S\right) \qquad \omega \in [-\pi, \pi] \tag{2.7}$$

It is clear from (2.7) that the apparent filter $G(f)$ is correctly given by (2.6). Equation (2.7) also shows that the variable-gain FA system subjects each spectral slice $X(f + lF_S)$, $f \in [-F_S/2, +F_S/2]$, of the input signal to a complex-weighted linear combination of the filter slices $H(f + kF_S)$, $f \in [-F_S/2, +F_S/2]$, of the filter. As such, the FA system can *modify*, and *enhance*, the selectivity in $H(f)$ to approach desired filtering profiles.

The sum of filters $H(f + mF_S)$ in (2.6) suggests that, to create a filter $G(f)$ that exhibits a low-pass or a low frequency band-pass response, all tones D_m that place copies of the pass-band(s) of $H(f)$ in the stop-band of the desired filter $G(f)$ should be suppressed. Accordingly, if the pass-band(s) of $H(f)$ lie(s) at relatively low-frequencies, the significant tones D_m would be similarly limited to low-frequencies. We can then

describe a more compact formulation of $G(f)$ where the frequency content of $d(t)$ is limited to $[-MF_S, +MF_S]$ as in (2.8).

$$G(f) = \sum_{m=-M}^{+M} D_m H(f + mF_S) \qquad f \in [-\infty, +\infty] \tag{2.8}$$

As we will show in the examples of section II.2.D, for a simple low-pass $H(f)$, the frequency content of $d(t)$ can be limited to the bandwidth of $G(f)$. A band-limited $d(t)$ has very important implications:

1- The FA solution does not impose additional bandwidth requirements on the analog circuitry such as the mixer and sampler beyond the anti-aliasing filter bandwidth.

2- The value of $G(f)$ at a given frequency f can be numerically evaluated (finite sum in (2.8)) allowing us to design the anti-aliasing filter in the frequency domain.

3- The FA system can be *approximated* by a discrete-time counterpart during the filter design and/or hardware realization stages (see section II.2.B below).

B. On the Discrete-Time (DT) Approximation

Since we can describe the FA response for a band-limited $d(t)$ in (2.8), we can also describe it for a discrete-time variant. Fig. II-6a shows the block-diagram for a discrete-time FA system. To generate the discrete-time analog signal $x[\eta]$ at the input to the FA system in Fig. II-6a, pre-sampling at rate NF_S should be performed. We use sample index η to correspond to signals at rate NF_S as compared to index n corresponding

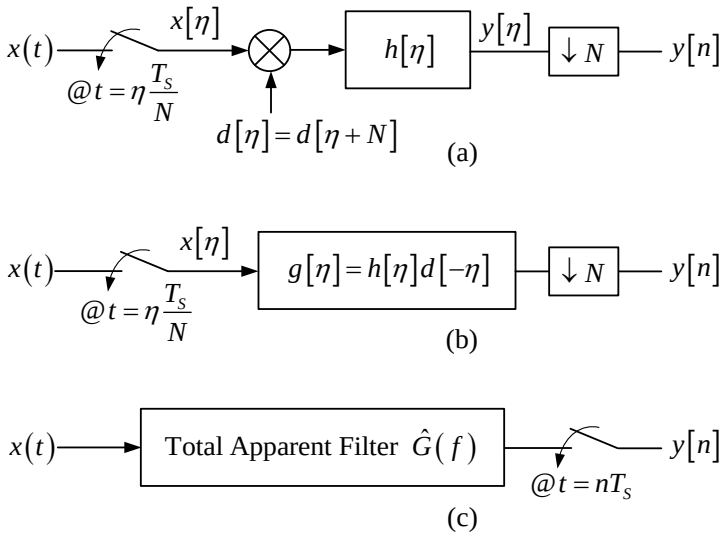

Fig. II-6 (a) The discrete-time variant of the FA system under pre-sampling with rate NF_S. (b) and (c) show the effective anti-alias filtering as seen by $y[n]$. N will in practice be bounded by the desired location NF_S of the first image pass-band of the effective filter.

to the final sampling rate F_S. Similar to the continuous-time case in (2.6), the apparent discrete-time filter $G(\omega)$ or $g[\eta]$ seen by the output $y[n]$ of the decimator in Fig. II-6b can be described by (2.9).

$$G(\omega) = \sum_{m=-M}^{+M} D_m H\left(\omega + \frac{2\pi m}{N}\right), \qquad \omega \in [-\pi, +\pi], \qquad D_m = D\left(\frac{2\pi m}{N}\right) \tag{2.9}$$

Because of the initial pre-sampling, the total apparent anti-aliasing filter $\hat{G}(f)$, illustrated in Fig. II-6c, would exhibit copies of the response $G(\omega)$ at multiples of NF_S. When the discrete-time FA variant is used to approximate the continuous-time FA

system, such as for the filter design in section II.2.C, N is made very large and NF_S would lie well above the frequency range of interest. For a discrete-time hardware implementation, however, a very large N is not desirable. The choice of N will be attended to along with other implementation concerns in section II.2.E. For the time-being, when referring to a discrete-time FA implementation, we will assume that all frequency content beyond $NF_S/2$ is sufficiently suppressed.

C. Designing the FA apparent Filter

We will describe in this section the design of the apparent filter $G(f)$ to meet a general target filtering profile. We describe the design problem in the frequency domain and over the frequency range of interest, independent of the choice of a continuous or discrete-time implementation.

The Filter Design Problem

For brevity, we will only go through the design scenario where the signal of interest, and accordingly the filter pass-band is centered at baseband. In a direct-conversion receiver [23], the signal of interest would have probably gone through proper down-conversion by the radio-frequency front-end. A band-pass version of the designed low-pass filters, centered at cF_S, can simply be created by multiplying $d(t)$ by $\cos(j2\pi ctF_S)$ and $\sin(j2\pi ctF_S)$ for the in-phase and quadrature demodulation paths respectively. The mixer of the variable-gain FA system in such a configuration would also serve as a tuner for a multichannel system [24] or software-defined radio, or as the second down-conversion stage of a super-heterodyne receiver [23].

To describe a general desired filtering response with a low-frequency pass-band we can use the set of conditions given by (2.10). The top row describes the pass-band while other rows describe the stop-bands.

$$\begin{aligned} 1-\delta_{pass} \le \left|G(f)\right| \le 1 \qquad & f \in \left[0, f_{pass}\right] \\ \left|G(f)\right| \le \delta_{stop-1} \qquad & f \in \left[f_{stop-1}, f_{stop-1}+\Delta_1\right] \\ \vdots \qquad & \vdots \\ \left|G(f)\right| \le \delta_{stop-B} \qquad & f \in \left[f_{stop-B}, f_{stop-B}+\Delta_B\right] \end{aligned} \tag{2.10}$$

The design problem can then be written as a minimization (2.11) of the stop-band magnitude under pass-band ripple constraints. $G(f)$ is described in (2.11) in terms of $H(f)$ and $d(t)$.

$$\begin{aligned} \min_{D,H} \quad & \delta_{stop} \\ \text{s.t.} \quad & 1-\delta_{pass} \le \left|\sum_{m=-M}^{+M} D_m H\left(f+mF_S\right)\right| \le 1 \qquad f \in \left[0, f_{pass}\right] \\ & \frac{\delta_{stop}}{\delta_{stop-1}}\left|\sum_{m=-M}^{+M} D_m H\left(f+mF_S\right)\right| \le \delta_{stop} \qquad f \in \left[f_{stop-1}, f_{stop-1}+\Delta_1\right] \\ & \vdots \qquad\qquad\qquad\qquad\qquad\qquad \vdots \end{aligned} \tag{2.11}$$

Note that, as in typical filter design problems, we would solve the continuous-frequency problem (2.11) on a discretized frequency axis with sufficient resolution. Also, using the same reasoning as in (2.8), if $H(f)$ is not high-pass and $d(t)$ is band-limited,

the response of $G(f)$ will decrease at higher frequencies and (2.11) need only be solved on a finite set of frequencies. The filter design problem in this general form is non-convex due to the lower-bound on the pass-band. A non-linear programming solver, e.g. the Newton Method [25], is required to solve it. As we show next, however, when $H(f)$ is a simple filter the optimization (2.11) can be re-written as a convex problem and solved optimally.

Simple $H(f)$ and the Convex Problem Formulation

We recall from section II.1 that strong on-chip filters incur high power and area costs and exhibit limited reconfigurability. It thus behooves us to use a simple physical filter $H(f)$ and embed the FA programmability into the spreading signal $d(t)$. The simplest choice for $H(f)$ is a low-pass single-pole filter, well known to be the simplest continuous-time analog filter. Single-pole filtering requires a single memory element and is usually embedded in the frequency response of general analog circuit components. Discrete-time analog single-pole filters, as described in [12], are also very simple to implement.

This section will take a closer look at the design problem (2.11) tailored to simple filters $H(f)$ such as the single-pole filter. To *describe the design problem in convex form* we will use some sufficiently oversampled $H(\omega)$ and $G(\omega)$, as discussed in section II.2.B, to represent the low frequency behavior of $H(f)$ and $G(f)$ respectively.

Let $H(\omega)$ be the single-pole filter $1/(1-\alpha \exp(-j\omega))$. Equation (2.12) shows the resultant $G(\omega)$ from (2.9).

Table II-1

Examples of simple $H(\omega)$ resulting in a simple closed form $G(\omega)$.

	$H(\omega)$	$G(\omega)$
Finite Impulse Response	$\sum_{k=0}^{2M} \exp(-j\omega k)$	$\sum_{k=0}^{2M} d[-k] \exp(-j\omega k)$
Single-pole	$\frac{1}{1-\alpha \exp(-j\omega)}$	$\frac{\sum_{k=0}^{2M} d[-k]\alpha^k \exp(-j\omega k)}{1-\alpha^{2M+1} \exp(-j\omega(2M+1))}$
Double-pole *	$\frac{1}{1-\alpha \exp(-2j\omega)}$	$\frac{\sum_{k=0}^{2M} d[-k]\alpha^{2k} \exp(-2j\omega k)}{1-\alpha^{4M+2} \exp(-j\omega(4M+2))}$

*This particular double-pole filter applies only to discrete-time FA because $H(\omega)$ is not strictly low-pass or band-pass.

$$G(\omega) = \frac{\sum_{m=-M}^{+M} D_m \prod_{\mu \neq m} \left(1-\alpha \exp\left(j\omega + j\frac{2\pi\mu}{N}\right)\right)}{\prod_{m=-M}^{+M} \left(1-\alpha \exp\left(j\omega + j\frac{2\pi m}{N}\right)\right)} = \frac{\sum_{k=0}^{2M} b_k e^{j\omega k}}{1-\alpha^{2M+1} e^{j\omega(2M+1)}} \qquad \omega \in [-\pi, +\pi] \qquad (2.12)$$

As can be seen from (2.12), the denominator of $G(\omega)$ is completely determined by the variable α of $H(\omega)$. Also, the number of coefficients b_k in the numerator of $G(\omega)$ is equal to the number of spreader variables D_m. The apparent filter $G(\omega)$, for a single-pole $H(\omega)$, can thus be mapped using a one-to-one function to the variables D_m and α. In fact, Table II-1 shows the explicit formulation of $G(\omega)$, for several simple filters $H(\omega)$, in terms of the $2M+1$ time domain samples $d[k]$ of the spreader.

The design of $G(\omega)$ can now be solved as a convex programming [26]-[27] problem

(2.13) in terms of $R(\omega)$, the squared magnitude of $G(\omega)$, $R(\omega) = |G(\omega)|^2$. After solving (2.13), spectral factorization [28] can be used to find the minimum phase $G(\omega)$ that satisfies the optimal $R(\omega)$. The optimal $d[k]$ and α, if set as a variable, would be derived from the closed form of $G(\omega)$. D_m is then the Discrete Fourier Transform (DFT) of $d[k]$ and $d(t)$ the continuous-time interpolation of $d[k]$.

$$\begin{aligned} \min_{D_m, H(\omega)} \quad & \delta_{stop} \\ \text{s.t.} \quad & \left(1-\delta_{pass}\right)^2 \le R(\omega) \le 1 \qquad \omega \in \left[0, \omega_{pass}\right] \\ & \frac{\delta_{stop}}{\delta_{stop-1}^{\ 2}} R(\omega) \le \delta_{stop} \qquad \omega \in \left[\omega_{stop-1}, \omega_{stop-1} + \Delta_1\right] \\ & \vdots \qquad\qquad\qquad\qquad \vdots \\ & 0 \le R(\omega) \qquad\qquad\qquad \text{all} \quad \omega \end{aligned} \tag{2.13}$$

In the following section we will present several examples of achievable anti-aliasing filters that underscore the versatility of the FA technique.

D. Variable-gain Filtering by Aliasing Responses

For brevity, and in line with using a simple filter $H(f)$, we limit our discussion to *FA responses based on a single-pole* $H(f)$. We will also discuss an enhanced interleaved variant of the FA system. The results for this scenario apply similarly to an FA system implemented in the discrete-time or continuous-time domains –provided that pre-sampling is performed at a sufficiently high rate for the discrete-time case. For uniformity, all frequency axes are normalized to $F_S/2$.

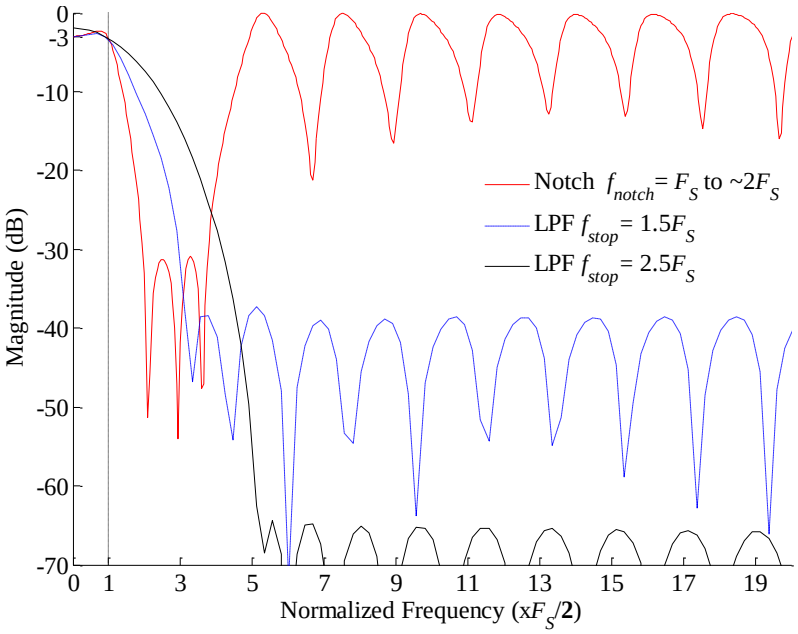

Fig. II-7 Sample effective filtering responses $G(f)$ achievable using a single-pole $h(t)$ or $h[\eta]$.

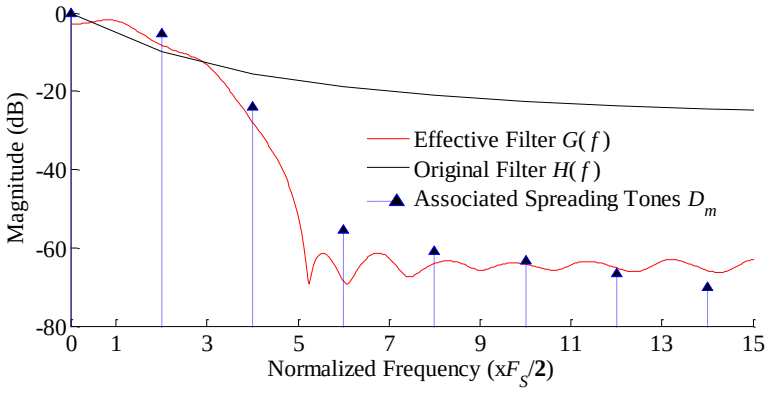

Fig. II-8 Combined plot of the original filter $H(f)$, the spreading tone values D_m and the resulting variable-gain FA response $G(f)$.

Standard Achievable Anti-Aliasing Responses

Fig. II-7 shows possible achievable $G(f)$ for a filter having a -3dB frequency of $F_S/2$. The responses show that the FA system can be programmed to suppress adjacent

narrowband interference as well as very wideband interference at varying offsets from the desired signal.

On one hand, the low-pass filters with 40dB and 65dB suppression illustrate how the stop frequency and stop-band suppression can be traded-off. On the other hand, the notch filter trades-off stop-band size and location exhibiting -30dB suppression between F_S and $2F_S$. It is important to note that the notch filter takes advantage of the higher frequency poles of $G(f)$.

Following our discussion on the bandwidth of $d(t)$ in section II.2.A, we show in Fig. II-8 an apparent low-pass filter $G(f)$, the original single-pole filter $H(f)$, and the associated spreading tone values D_m. As can be seen from the figure, the bandwidth of the spreading signal is comparable to the desired filter bandwidth.

Beyond the examples in Fig. II-7, a wide range of responses can be achieved by tailoring the stop-band requirements according to the optimization problem (2.13). In addition to designing multiple stop-bands with different suppression requirements, the problem (2.13) could be modified to allow a rolling-off stop-band such that suppression increases with frequency. Furthermore, as one would expect, improved roll-off can be achieved by using a higher order $H(f)$ filter. Appendix A, for example, describes the achievable apparent anti-aliasing filter responses using a double-pole discrete-time filter realization. Appendix B on the other hand shows apparent variable-gain FA responses for arbitrary interference profiles. For the bulk of this dissertation, however, we will continue to focus on a single-pole $H(f)$ and explore sharper filtering by the interleaved FA structure presented in subsection below.

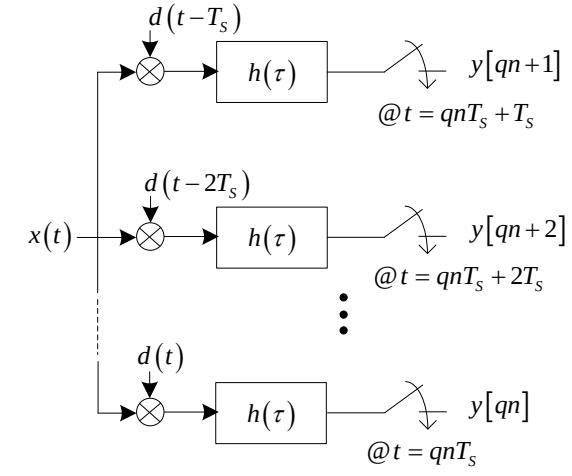

(a)

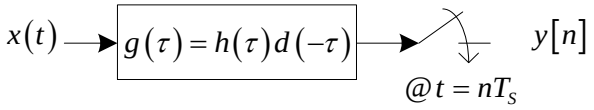

(b)

Fig. II-9 (a) An illustration of sampling with multiple phases of the FA system. Increasing the number of channels, q, is equivalent to increasing the period of $d(t)$ (referred to here as T_p) with respect to the total sampling period T_S and thus creating sharper FA responses. (b) The time-invariant effective filter.

Interleaved Filtering by Aliasing

Consider the FA structure in Fig. II-10a where q parallel FA channels operate simultaneously on the input $x(t)$. The k^{th} channel employs a T_p-periodic sequence $d(t-kT_p/q)$, a cyclical shift of $d(t)$ by kT_p/q, and samples at time instances $t = nT_p+kT_p/q$ with rate $F_p = 1/T_p$. The de-interleaved output $y[n]$, at rate $F_S = qF_p$, appears filtered by the same time-invariant filter $g(t)$ at all instances, as shown in Fig. II-10b.

Such an interleaved FA structure allows us to increase the period T_p of $d(t)$ (decrease the tone separation F_p) by q times while maintaining a constant total sampling rate F_S. A

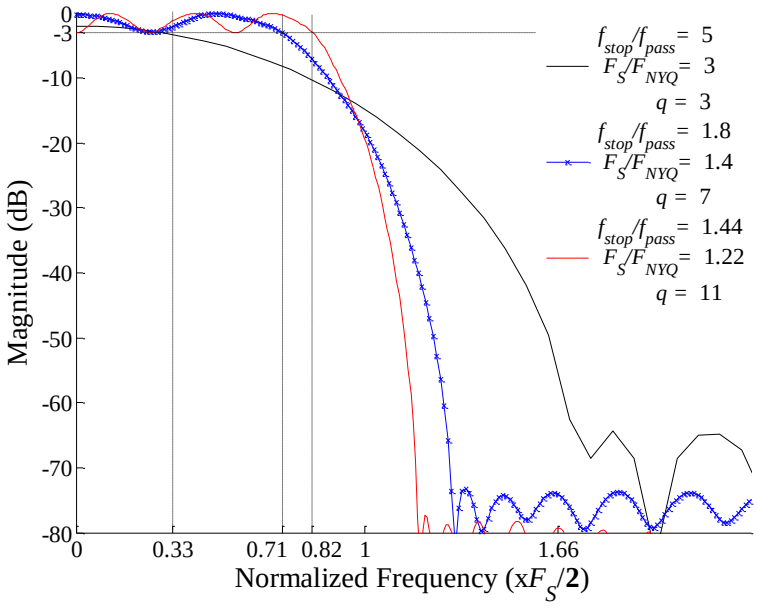

Fig. II-10 Sample apparent responses achievable using a single-pole $h(t)$ and a transition band of size $2F_p$. The figure shows how increasing the number of channels q allows for a sharper filter or equivalently a larger pass-band for the same sampling rate. The labeled q values alias the transition-band over itself.

larger $q = F_S/F_p$ increases the number of degrees of freedom available to the filter design described in (2.13). The result is sharper filter transitions for the same desired pass-band.

Fig. II-10 shows sample FA low-pass responses with interleaved channels using the single-pole $H(f)$. As the figure shows, increasing q allows for sharper filtering. We recall here that the frequency axes are normalized to $F_S/2$ for better illustration. In practice, we would increase q such that, for a constant pass-band size, the required sampling rate F_S is reduced. The legend on Fig. II-10 displays the ratio of the stop-frequency f_{stop} to the pass-frequency f_{pass} for each filter as well as the required sampling

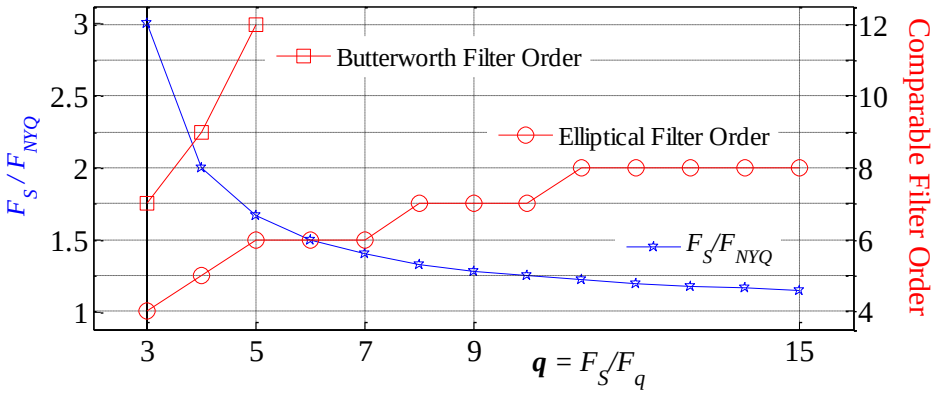

Fig. II-11 The sampling efficiency (F_S/F_{NYQ}) for different values of q plotted against the left vertical axis. Also shown, on the right vertical axis, are the Butterworth and Elliptical filter orders required for comparable performance. The results apply to filters such as in Fig. II-10 with >65dB suppression and a transition-band of size $2F_p$.

rate F_S with respect to the minimum Nyquist rate F_{NYQ} of the signal of interest. F_{NYQ} is calculated as $2f_{pass}$ while F_S is chosen as $(f_{stop}+f_{pass})/2$ allowing the transition band to alias over itself in the sampled output.

Equation (2.14) describes the trend in the required sampling rate F_S versus the desired signal Nyquist rate F_{NYQ} as q changes. The relationship (2.14) applies only to FA low-pass responses whose transition band is of size $2F_p$. Such filters, as shown in Fig. II-10, exhibit 65dB of stop-band suppression or more. If only 40dB of stop-band suppression is desired, the transition band can be of size F_p (refer back to Fig. II-7) and (2.14) would change accordingly.

$$\frac{F_S}{F_{NYQ}} = \frac{q}{q-2} \tag{2.14}$$

Fig. II-11 plots the sampling efficiency F_S/F_{NYQ} given by (2.14) on the left vertical axis versus q. It can be seen that increasing the number of channels q achieves diminishing returns in the sampling efficiency. In particular, the diminishing reduction in the sampling rate F_S with respect to F_{NYQ} might not justify the cost of extra FA channels in the analog domain. The optimal choice of q will depend on the existence of strong interferers very close to the signal of interest and the relative cost of ADC resolution and FA channel hardware.

Note that Fig. II-11 also plots, on the right vertical axis, the equivalent order of a Butterworth or Elliptical filter exhibiting the same sampling efficiency and stop-band suppression as the FA response for a given q. For example, for $q = 5$, the FA filter is comparable to a 6^{th}-order Elliptical filter or a 12^{th}-order Butterworth filter. We will return to the performance evaluation of the FA approach and comparison with other analog filtering techniques in section II.4.

E. Practical Considerations

Before presenting simulation results of a possible implementation of the variable-gain FA system, we will attend to possible high-level concerns regarding the practical realization. In particular, we will discuss the effects the amplitude resolution and switching rate of the spreading signal $d(t)$ on the apparent filter $G(f)$. In-depth concerns such as timing jitter and noise are beyond the scope of this work and will be attended to in future research. It is important to note that the discussions on resolution and switching rate apply similarly to the variable-pole FA system described in section II.3.

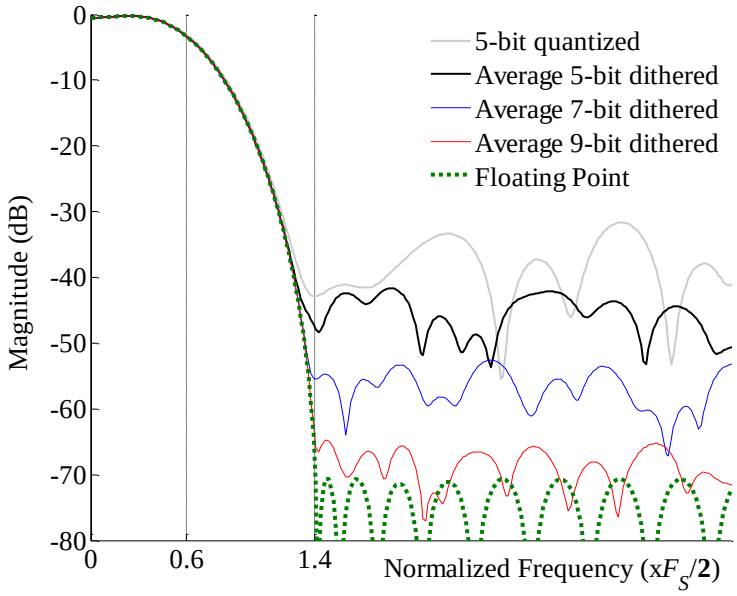

Fig. II-12 The effect of quantizing as well as dithering $d(t)$ on the stop-band of the apparent filter $G_Q(f)$.

Effect of Spreader Resolution

Independent of the details of the circuit implementation, we look to analyze the effect of a non-ideal mixing signal $d(t)$ on the apparent filter $G(f)$. In line with digital programmability of the mixing signals we assume an N_b-bit uniform quantization $d_Q(t)$ of $d(t)$.

Because $d(t)$ is T_p-periodic, its quantization creates a periodic quantization error $e(t) = e(t+T_p)$. The associated non-ideal apparent filter $G_Q(f)$ is given by (2.15) where E_m are the frequency tones associated with $e(t)$. The quantization error tones E_m disrupt the frequency response by placing unwanted copies of $H(f)$ in the stop-band of $G_Q(f)$.

$$G_Q(f) = \sum_{m=-\infty}^{+\infty} D_m H(f + mF_p) + \sum_{m=-\infty}^{+\infty} E_m H(f + mF_p) \qquad f \in [-\infty, +\infty] \tag{2.15}$$

Fig. II-12 shows, for example, the apparent filter for a 5-bit quantized $d_Q(t)$ showing about 30dB stop-band suppression, as compared to an original 70dB suppression. The figure also shows the improved apparent filter for a 5-bit quantized $d_Q(t)$ after dithering, showing 40dB of suppression. Dithering [29] is the process of intentionally adding noise prior to quantization to randomize the quantization error. In our FA application, dithering helps remove the periodicity of $e(t)$ thus flattening the stop-band gain variations caused by quantization. Fig. II-12 shows that, with dithering, a 9-bit quantized $d_Q(t)$ is sufficient to achieve more than 65dB of stop-band suppression. The achievable stop-band suppression shown in Fig. II-12, for a given quantization depth, can be directly extended to any filter $G(f)$ if the quantization error is randomized by dithering.

We remind the reader that if, in certain applications, weak stop-band suppression (low resolution) can be tolerated then the FA response can be redesigned, as shown in Fig. II-7, to exhibit sharper roll-offs.

Effect of Spreading Rate (NF_S)

In addition to finite resolution in the amplitude of $d(t)$, a practical spreading operation is typically realized as a digitally-switched gain stage, as in digital-to-analog converters [30], [31], that exhibits finite bandwidth or switching rate.

Fig. II-13 illustrates the digitally controlled gain stage within the continuous-time

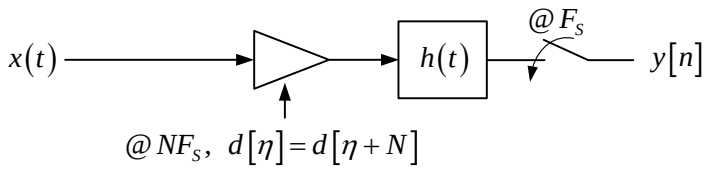

(a)

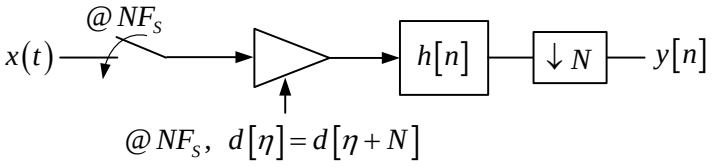

(b)

Fig. II-13 (a) The continuous-time FA system with rate-NF_S spreading sequence. (b) The discrete-time FA system with rate-NF_S input and spreading sequence. Note that in (a) the sequence $d[\eta]$ controls a continuous-time gain stage, the actual continuous-time gain value would practically be some shaping of the discrete-time $d[n]$.

variable-gain FA system (a) or the discrete-time variable-gain FA system (b). Notice the spreading sequence $d[\eta]$, a rate-NF_S discretization of $d(t)$.

As demonstrated throughout this manuscript, the spreading signal $d(t)$ is actually band-limited and does not impose high switching rate requirements on the variant $d[\eta]$. In particular, if the spreader switching rate NF_S is higher than the Nyquist rate of $d(t)$, it will not interfere with the low-frequency filtering response. Fig. II-14a shows two examples of a step-wise switched $d(t)$. It can be seen that the low-frequency content of $d(t)$ is maintained and accordingly the low-frequency spreading tones D_m and frequency response $G(f)$ would not be modified.

In a similar fashion to the discrete-time pre-sampling rate in section II.2.B, the switching rate of $d(t)$ determines *only* the locations of image pass-bands in the apparent

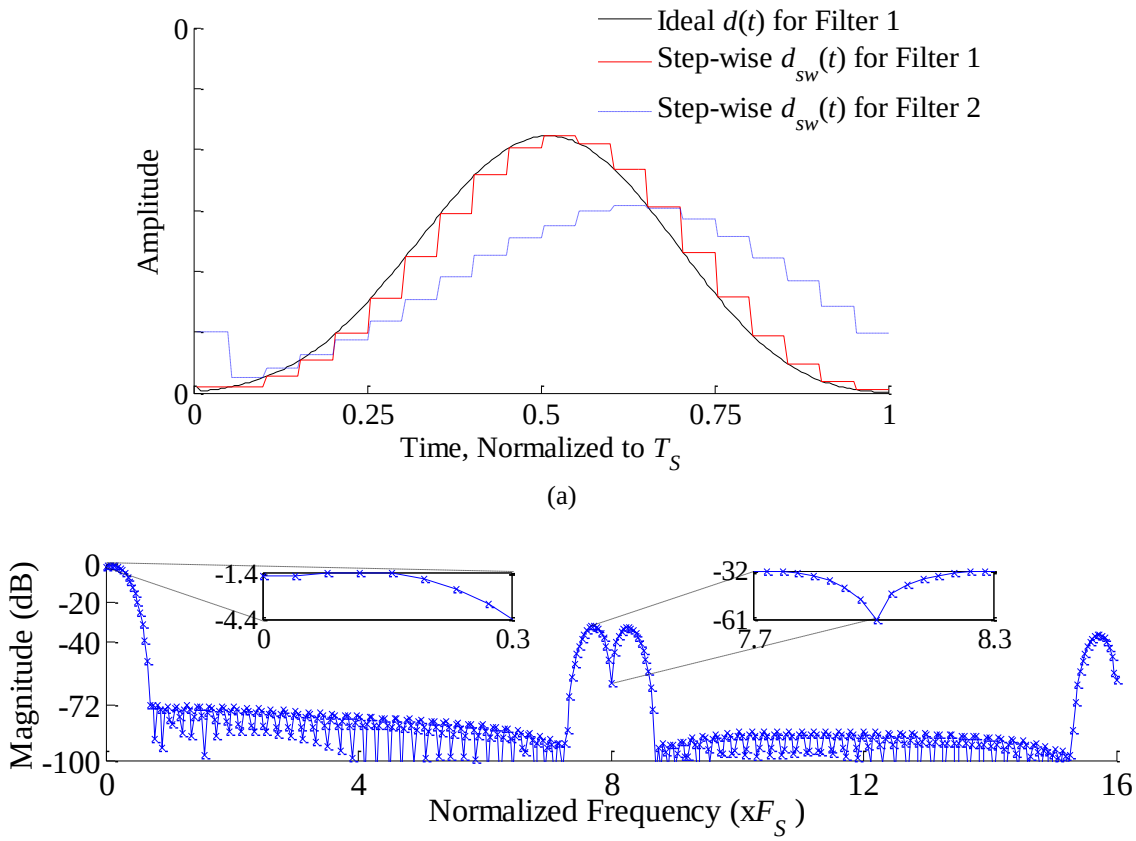

Fig. II-14 (a) Examples of a step-wise realization of the spreading signal $d(t)$. (b) A sample low-pass apparent filter $G(f)$ obtained using an $N = 8$ times switched $d(t)$ and a variable gain FA system with $q = 5$. Notice how the step-wise realization of $d(t)$ does not affect the low frequency response of $G(f)$ and creates suppressed pass-band images at multiples of NF_S.

filter. In practice, for a reasonably high N, the image pass-bands at multiples of NF_S lie far from the band of interest and are significantly suppressed by the inherent low-pass nature of preceding analog circuitry as well as the low-pass shaping of the pre-sampler or DAC itself, such as when using integration sampling [15] in [4], [11], [22]. Fig. II-14b shows an example $G(f)$ with a step-wise $d(t)$ running at $8F_S$. For completing we note that

the filter response corresponds to an interleaved FA realization with $q = 5$.

To put the FA spreader requirements in context, we note that digital-to-analog resolution and switching rates well beyond the described requirements are widely demonstrated in the literature, e.g. [22], [30]-[31], allowing for -40dB to -70dB FA suppression at NF_S of 3.2GS/s to 600MS/s.

F. Active-RC Example Implementation

As a straightforward realization of the variable-gain FA system, the circuit in Fig. II-15a was built and simulated in Cadence using true CMOS circuit models. The circuit is essentially a single-pole active-RC filter whose input resistance is controlled by binary control signals $\{b_0, b_{N-1}\}$.

High speed switching of the resistors of an active-RC filter is used in the literature, e.g. [32], according to the technique in [33], to control the average value of such resistors for gain and bandwidth control. In the circuit of Fig. II-15a then, we employ similar switching circuitry as already investiaged in the literature but we control the low-frequency variation of the resistor instead of its average value, thus achieving significantly sharper filtering responses than the single-pole performance obtained by a switched-resistor or a traditional active-RC filter.

In particular, using 7 binary-weighted resistors at the input, and switching them at $20F_S$ for a sampling rate F_S = 20MHz, we obtain the filtering responses shown in Fig. II-15b. The latter figure, showing both the expected Matlab results and the non-ideal Cadence simulation results, demonstrates around 50dB of stop-band suppression and a slightly suppressed image pass-band at 400MHz.

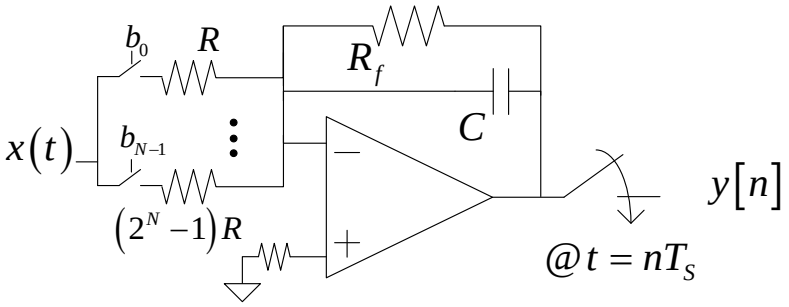

(a)

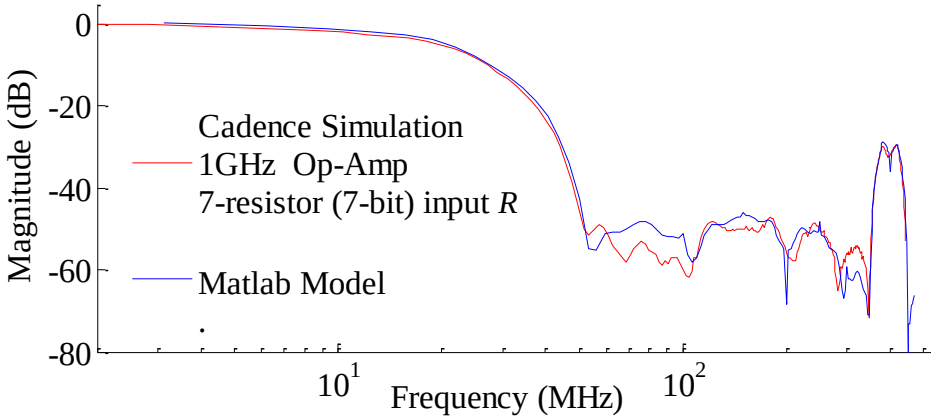

(b)

Fig. II-15 (a) An active-RC realization of the variable-gain FA system with a single-pole $h(t)$. the spreading sequence $d(t)$ is given by the inverse of $R(t)$, the input resistance. (b) Matlab design results and cadence simulation results using true CMOS Op-amp and switches.

We note here that, using a standard 1GHz GBW Op-amp, the simulated IIP3 was around 30dBm for 3mW of power consumption.

Even though such active-RC FA results are very appealing, the variable-pole FA system described next in section II.3 can be applied to a fully-passive RC filter, thus exhibiting even lower power consumption and higher linearity.

II.3 VARIABLE-POLE FILTERING BY ALIASING

As introduced in chapter I, the Filtering by Aliasing concept can be extended to general systems where the physical filter prior to a sampler can be periodically modified, not necessarily in terms of its gain, to obtain enhanced programmable responses at the sampling instances. This section discusses the application of the FA techniqe to the simplest filtering element, the passive RC filter, by periodically varying the pole. We omit the Filtering by Aliasing introduction and the details of the DT approximation and optimization problem. We instead focus on the strengths of the variable-pole FA system as a powerful filtering solution for SDR.

We propose that sharp, passive, and programmable anti-alias filters, highly suitable for SDR, can be created by periodically varying the resistor of a simple RC filter. The modulated resistor shapes the impulse response of the single-pole filter, otherwise monotonically decreasing, creating enhanced frequency selectivity at periodic instances. The concept can be similarly applied to RL, RLC or other filters.

Because of its passive nature, the proposed RC filter can be used to capture

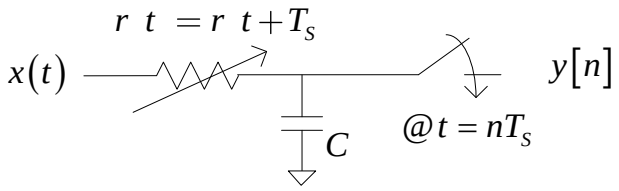

Fig. II-16 The enhanced RC anti-aliasing filter. The resistor is periodically modulated with the same period as the sampler. We refer to this system as variable-pole Filtering by Aliasing.

desired signals among large blockers. To that end, we will demonstrate how the variable resistor can also emulate harmonic-rejection mixing at band-pass frequencies without active gain stages. To our knowledge, no similar technique exists in the literature.

A. The Enhanced RC Filter

Consider the linear time-varying RC filter shown in Fig. II-16. If the resistance is varied periodically at the same rate $F_S = 1/T_S$ as the sampler, the CT filter exhibits a linear periodically time varying response. At the sampling instances, however, the complete system appears to exhibit a linear and time invariant (LTI) response which we refer to as the apparent filter. We are interested in analyzing and designing such apparent LTI response at the sampling instances such that it is sharp and programmable.

From the perspective of the sampled output, the voltage $y(t)$ across the capacitor at any time $t = nT_S + \tau$ is then given by (2.16).

$$y\left(nT_S + t\right) = x\left(nT_S + t\right) - r\left(t\right)C\frac{dy\left(nT_S + t\right)}{dt} \tag{2.16}$$

As we did for the variable-gain FA scenario, we resort to a DT approximation to design the apparent filter. In this non-intuitive variable-pole scenario however, it behoves us to use DT approximation to also formulate the apparent filter. Assuming that $r(t)$ is band-limited, the CT filter and sampler in Fig. II-16 can be approximated by the highly oversampled DT filter and decimator in Fig. II-17. The sample index η corresponds to the

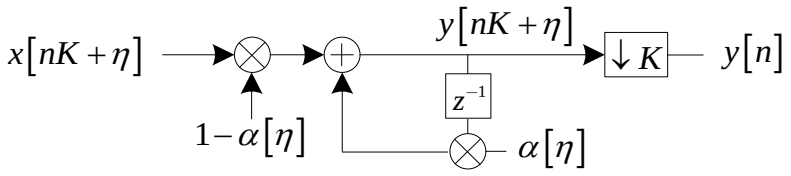

Fig. II-17 The DT approximation of the RC filter and sampler from Fig. II-16.

approximation rate $KF_S >> F_S$ while sample index n corresponds to the final sampling rate F_S. Equation (2.16) can thus be approximated by (2.17), where α is given by (2.18).

$$y[nK+\eta] = (1-\alpha[\eta])\,x[nK+\eta] + \alpha[\eta]\,y[nK+\eta-1] \tag{2.17}$$

$$\alpha[\eta] = \frac{r(\eta\Delta t)C}{\Delta t + r(\eta\Delta t)C}, \qquad \Delta t = \frac{T_S}{K} \tag{2.18}$$

By expanding (2.17), we derive the LTI impulse response $g[\eta]$ seen at the decimation instances, its frequency domain $G(\omega)$ is given by (2.19).

$$G(\omega) = \frac{\sum_{\eta=0}^{K-1}\left[(1-\alpha[\eta])\left(\prod_{i=1}^{\eta-1}\alpha[i]\right)\exp(-j\omega\eta)\right]}{1-\left(\prod_{\eta=1}^{K}\alpha[\eta]\right)\exp(-j\omega K)} \tag{2.19}$$

From (2.19) we can see that the apparent filter exhibits replicated poles at multiples of the sampling rate F_S, creating a uniform pass-band, and a multitude of zeros

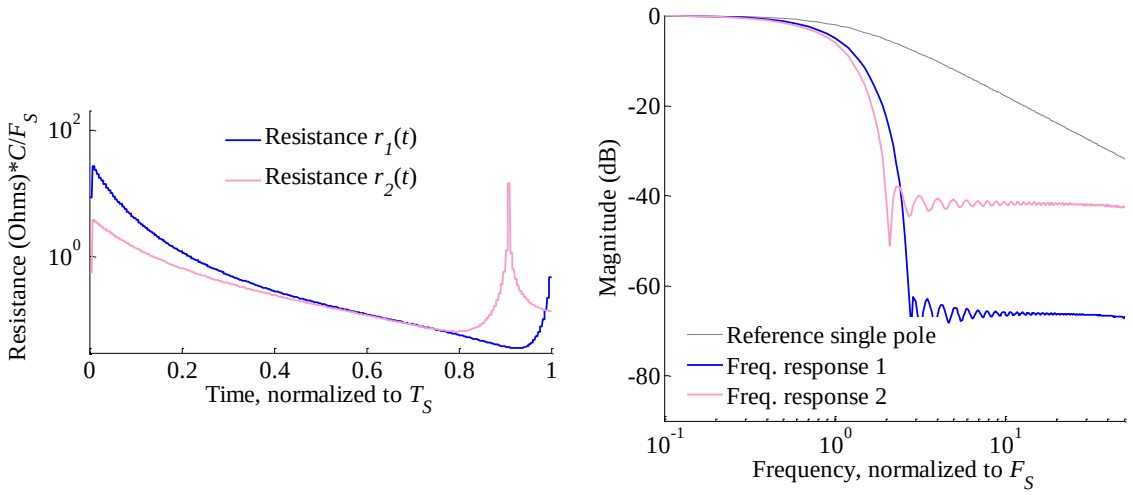

Fig. II-18 (left) Optimally chosen $r(t)$. (right) The corresponding frequency responses.

that will shape the filter stop-band. It is apparent that $G(\omega)$ is significantly more selective than the original single-pole filter response.

By designing the DT filter (2.19) for a desired frequency response we can compute the corresponding $r(\eta\Delta t)$ using (2.18). Similar to the variable gain case described in section II.2.C, we use convex optimization [27] and spectral factorization [28] to find the optimal squared magnitude $|G(\omega)|^2$ and the corresponding minimum phase $G(\omega)$ respectively. The numerator of $G(\omega)$ is scaled such that all $\alpha[\eta]$ belong to [0,1). $r(\eta\Delta t)$ is then computed and interpolated.

Fig. II-18 shows two examples of optimal $r(t)$ (over one period) on the left and the corresponding anti-aliasing frequency responses of the enhanced RC filter on the right. The plotted values of $r(t)$ are scaled by the sampling rate F_S and the capacitance C such that they can be extended to any desired application. Notice the very sharp achievable

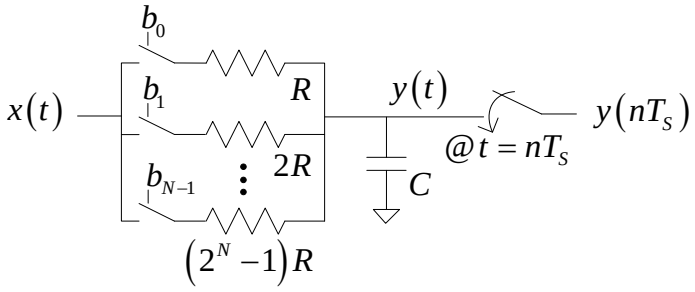

Fig. II-19 Binary-weighted switched realization of the system from Fig. II-16.

filtering compared to the reference unmodified RC response as well as the trade-off between stop-band suppression and filter roll-off. Note that frequency response 1, showing > 60dB of suppression very close to the sampling rate F_S, exhibits < 1.5dB of loss at $F_S/2$.

B. Practical Realization: Switched Binary-Weighted Resistors

Within the clock-programmable environment of SDRs, the variable $r(t)$ can be practically realized as a parallel network of binary-weighted resistors as shown in Fig. II-19. It is important to differentiate this setup from switched-resistor active-RC filters [32] that try to emulate a constant low-frequency resistance for tuning/variability rather than a variable $r(t)$ for filter shaping.

Also as discussed for the variable-gain scenario in sections II.2.E and II.2.F, if the binary control voltages b_i are switched at some rate LF_S, the resistance $r(t)$ appears sampled and held during intervals $t \in [lT_S/L, (l+1)T_S/L]$. The filter impulse response,

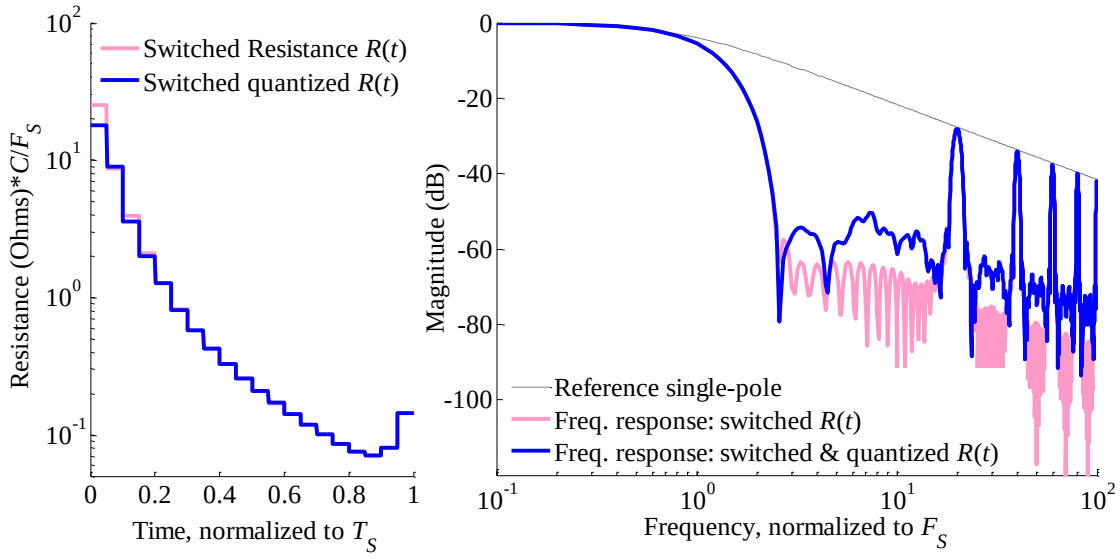

Fig. II-20 (left) Non-ideal $r(t)$: $20F_S$ switching rate and 8-bit quantization (255 levels), (right) the corresponding frequency responses.

however, exhibits single-pole decay for constant $r(t)$ and as such appears sampled and shaped by a low-pass filter. Accordingly, even though rate-LF_S switching creates images of the filter pass-band at multiples of LF_S, such images are suppressed at a 20dB/decade roll-off.

For example, Fig. II-20 shows, on the left, an ideal $r(t)$ as well as a non-ideal version sampled at $20F_S$ and quantized to 8-bits. The corresponding frequency responses are shown on the right. Notice that an 8-bit resolution in $r(t)$ is sufficient to provide >50dB of stop-band suppression. Stop-band suppression can be further improved by dithering prior to quantization as shown for the variable-gain scenario in section II.2.E.

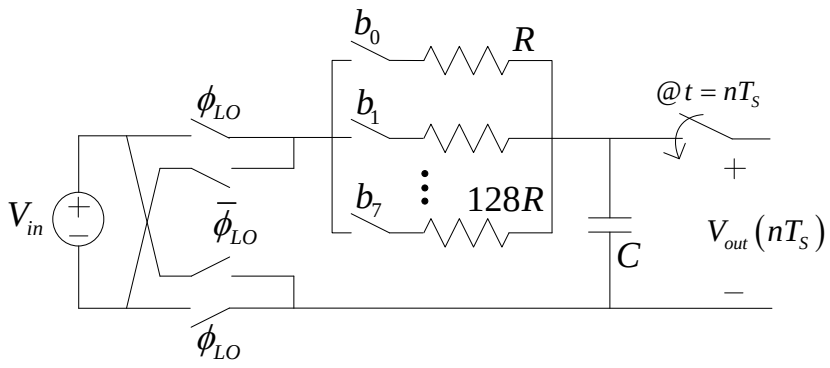

Fig. II-21 Band-pass enhanced RC filter. Φ_{LO} is a square-wave mixer clock.

C. Band-pass Filtering

If $G(\omega)$ in (2.19) is designed as a band-pass filter, one would expect the filter impulse response to essentially comprise a low-pass filter impulse response that is modulated by the carrier frequency. A passive ±1 mixer can then be used, as shown in Fig. II-21, to provide the sign alternation of the band-pass impulse response while $r(\eta\Delta t)$, computed from $|\alpha[\eta]|$, can be designed to provide both the filter low-pass impulse response as well as the envelope of the carrier sinusoid. Accordingly, if $r(t)$ is switched at M times faster than the carrier frequency, M-phase harmonic rejection would be inherently provided by the RC bandpass filter without any active gain components. Harmonic rejection is the ability to suppress the harmonics of a carrier frequency by emulating the true sinusoidal time-domain waveform.

As a numerical example, Fig. II-22 shows an 8-bit quantized and 8-phase switched $r(t)$ designed for a 20MHz-bandwidth filter at a center frequency of 500MHz.

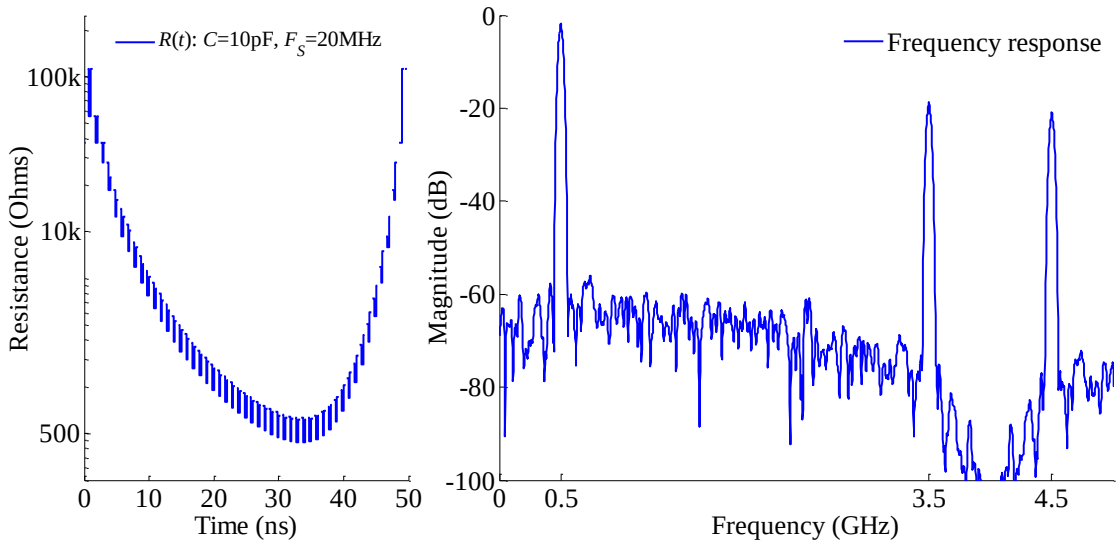

Fig. II-22 (left) $r(t)$ modulated for bandpass filtering, quantized and switched (right) the frequency response with and without quantization.

Given a 10pF capacitor, $r(t)$ ϵ [500Ω, 100kΩ]. On the right, Fig. II-22 plots the corresponding frequency response showing sharp band-pass filtering with 60dB suppression up to ~3.5GHz, the 7th harmonic of the carrier signal.

Of course, we have limited our discussion here to the use of the enhanced RC filter prior to a rate-F_S sampler. It is important to note, however, that multiple enhanced RC filters can be used in an interleaved fashion to provide multi-phase outputs, similar to the manner explained in section II.2.D for variable-gain FA. The filter sharpness can be further enhanced in such an interleaved setup.

II.4 IN THE CONTEXT OF PROIR ART

A. FA in the context of Reconfigurable Base-band Anti-alias Filtering

Referring back to Fig. II-11, the right vertical axis shows the equivalent order of a Butterworth or Elliptical filter that achieves similar performance to some FA responses of choice. We note here that similar results can be achieved using the variable-pole FA system as with the variable-gain FA system. We can see from the Fig. II-11 that the FA technique can easily match or exceed the filtering capabilities of traditional on-chip continuous-time (CT) analog filters, such as those listed in [13], using $q = 5$ for example. It is interesting to note here that, unlike traditional CT filters, the FA response can be enhanced by replicating the number of FA channels incurring only linear implementation cost in terms of q.

Beyond traditional continuous-time filters, we address in this section the relative performance of the proposed Filtering by Aliasing technique as compared to other anti-aliasing methods in the context of re-configurable anti-alias filtering for future communication systems. When compared against analog finite-impulse-response (AFIR) filters, e.g. [34], [35], or other DT analog filtering techniques [7], [8], [36] the FA might appear similar in the sense that the impulse response of the filter is clock-programmable and shaped by the switching circuitry. Existing DT analog filtering techniques attempt to mimic digital filters in the analog domain by using complex switched circuits that sample, share, accumulate, and amplify charges, resulting in insufficient dynamic range, high power consumption, high area, and limited band-pass capabilities. The FA technique

Table II-2

A comparison of anti-aliasing approaches under **65dB stop-band** suppression **low-pass** filtering.

	Order	q	Resol. (bits)	F_S/F_{NYQ}	$\boldsymbol{min\ qF_S/F_{NYQ}}$	Further Notes
Butterworth (for reference)	7, 9			3, 2		Continuous-time IIR filters, e.g. [13], exhibit only limited bandwidth programmability.
Successive Analog Decimation [4],[11]	-	3 serial stages	-	2 (30dB)* 90 (90dB)		✓ Programmable bandwidth. × Strong adjacent interference a bottleneck [11]. × Slow filter roll-off or high in-band droop. × High power/linearity cost.
MSHAWI † [13], [17],[18]	-	7, 9, 11	12 (9) ‡	15, 8, 6	66	✓ Prog. bandwidth and suppression FIR filter. ✓ Provides similar band-pass performance × Strong adjacent interference is a bottleneck. × High power/linearity cost.
FA single-pole	-	3, 4, 5	9	3, 2, 1.7	8.5	✓ Prog. bandwidth and suppression IIR filter. ✓ Provides similar band-pass performance. ✓ Passive implementation => low power, better linearity
FA double-pole	-	2, 3, 4	9	2, 1.3, 1.2	4.8	

*[4] and [11]targeted specific standards. Different results might apply for 65dB suppression.
†Multi-channel Sample-and-Hold Amplifier with Weighted Integration.
‡ Resolution reported in [13] in steps of 4-bits, should be reducible to 9 bits with dithering.

however, does not try to build "digital-style" filters in the analog domain but instead can be applied to simple passive analog components as demonstrated by the variable-poleFA in section II.3.

Table II-2 provides some insight into the relative complexities of various approaches for a candidate low-pass anti-alias filter with 65dB stop-band suppression or more. The order of a Butterworth filter with comparable stop-band suppression and transition-band size is maintained for reference. In the table, the minimum number of paths q multiplied by the corresponding sampling efficiency F_S/F_{NYQ} is calculated as a comparison metric.

As listed in the table, we are interested in the performance of two existing and recent techniques, both of which rely on integration sampling (IS) [15]:

1- As part of a software-defined radio solution, the works in [4], [11] proposed and built an anti-alias filter based on successive windowed-integration and decimation with intermediate discrete-time filtering. The filter is clock programmable for bandwidth and includes some control knobs for integration orders. Because of its reliance on successive sinc-type filtering, the filter suffers from slow roll-off or high in-band droop to provide strong suppression, unless more stages are added.

2- To take advantage of the FIR filtering provided by an integration sampler [15], the works in [13], [17]-[18] proposed preceding it with a mixer that shapes the original rectangular window impulse response. The result is a programmable FIR filter. To increase the order of the FIR, the authors proposed the multi-channel sample-and-hold amplifier with weighted integration (MSHAWI). As Table II-2 shows, the MSHAWI requires more paths while sampling at a higher rate than the proposed FA technique with a minimum qF_S/F_{NYQ} 7.7x higher than the single-pole FA system. As the order of the FA filter $h(t)$ increases, the gap also increases. It is important to note that the hardware cost associated with a higher

interleaved sampling rate, as well as the ADC resolution cost associated with large transition bands, do not scale linearly. It is also important to note that, when implemented in the variable-pole fashion from section II.3, the FA system exhibits significantly less hardware cost per path then the MISHAWI solution which requires high-cost integration stages.

Interestingly, the MSHAWI system based on preceding an integration sampler with a mixer is a special case of the variable-gain FA technique where the impulse response of the filter $h(t)$ is a rectangular time window. The derivation, design, and results of the two systems are however different.

In addition to the performance advantages, we should recall here that, unlike IS-based solutions, the FA technique does not depend on an accurate realization of $h(t)$. The resolution of the mixing sequence, discussed in section II.2.E and II.3.B, is the primary hardware requirement for FA. Given sufficient resolution, the sequence $d(t)$ or $d[\eta]$ can potentially be re-designed for variations in $h(t)$.

B. FA in the Context of Filtering at RF Frequencies

Section II.3.C proposed the use of the variable-pole FA solution as a band-pass anti-aliasing filter at RF frequencies. One might also find similarities between the variable-pole FA circuitry and commutation based filtering techniques, such as N-path filters [20]-[21], due to their use of periodic switching and their applicability to passive components. The similarities, however, are only superficial. The passive N-path filter, such as in Fig. II-4, relies on periodic switching between identical RC filters to tune the center frequency of the band-pass RLC-type filter. The proposed FA technique on the

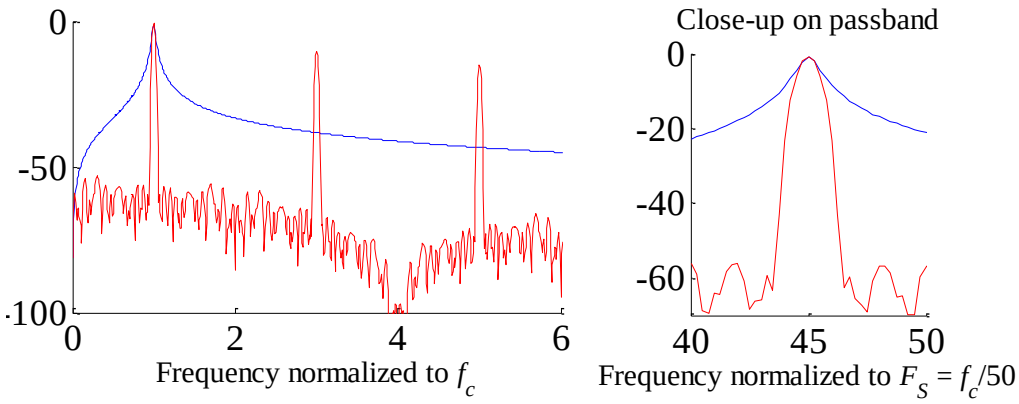

Fig. II-23 Comparison of the performance of the variable-pole FA response (red) against the popular N-path approach (blue) for band-pass filtering at RF frequencies. Both techniques rely on similary passive components and high speed switches. The 4-phase FA response (thus image pass-bands at harmonic frequencies) is plotted against the ideal infinite-phase N-path response (equivalent to an RLC filter) to show the drastic improvement in filtering performance.

other hand, while also providing such center frequency tuning, enhances the filter sharpness and provides filter response reconfigurability. The FA system from Fig. II-21 achieves this by essentially modifying the pole (feedback gain) of the filter while sharing the accumulated charge on the same capacitor.

To compare the achievable filtering responses of the FA and N-path approaches, the passive RC FA circuit in Fig. II-21 is simulated for an 8-bit resolution resistor and a 4-phase LO clock (8-phase plotted in Fig. II-22). The periodically varying resistance is optimized for a low-pass filter mixed with the envelope of the carrier, thus inherently providing 4-phase harmonic cancellation without gain stages as previously explained.

Fig. II-23 plots the achievable FA band-pass filter response, for a center frequency at $50F_S$, against the ideal RLC response for the same signal bandwidth. An ideal differential N-path filter with identical capacitance to the FA response can theoretically achieve such an RLC response using an infinite number of clock phases. We choose to plot the best achievable N-path filter to avoid the lengthy discussion of the trade-off between harmonic suppression and folding in N-path filters for finite N [20]. It is apparent from the figure that the selectivity of the FA solution far exceeds that of the best possible N-path solution, especially when considering near-band blockers.

III. RANDOM FILTERING BY ALIASING

III.1 INTRODUCTION

Recent works on Compressive Sensing (CS) [1]-[2], [37] theory described methods and conditions for the successful reconstruction of a sparse signal from compressive measurements. The compressive measurements can be described in matrix-vector format, where the sparse signal is a vector **s** with mostly zeros and the measurement operator is a matrix **Θ** with more columns than rows. The measurement result is a short vector **y = Θs** that can potentially hold sufficient information about the sparse **s**.

If the sensing matrix **Θ** satisfies the "restricted isometry property" (RIP) [37] with respect to the sparsity of **s**, several reconstruction algorithms [38]-[39] can be used to obtain the sparse **s**, or any non-sparse projection **v** of **s**, from **y**. The key aspect of CS theory is that a sensing matrix **Θ** drawn at random, from a Gaussian or Bernoulli distribution for example [40], can satisfy the RIP and be used for compressive measurements. This aspect drew the attention of researchers from various fields hoping to save acquisition, storage, or processing costs in their respective applications by making random measurements of large, but sparse, signals of interest.

Continuous-time spectrally sparse signals are one important class of signals whose acquisition can benefit from compressive sensing. In cognitive radio (CR) [41] for example, a large wideband chunk of spectrum, such as that of Fig. III-1, might contain a few narrowband signals of interest whose locations are unknown. CR might be interested

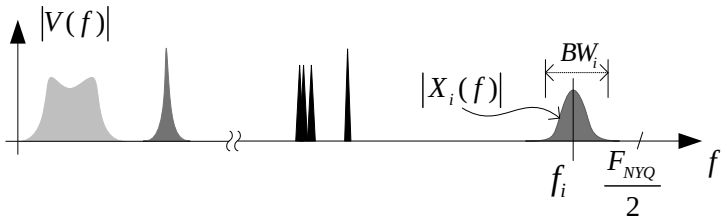

Fig. III-1 A general example of a sparse spectrum associated with a continuous-time band-limited signal $v(t)$. Note that $v(t)$ in this example is real-valued.

in the detection, acquisition, and/or decoding of such signals without sampling at the Nyquist rate of the wideband spectrum.

To enable reconstruction, however, CS theory requires that signals be sparse in a finite-dimensional space. If a signal is sparse in a continuous space, the sensing matrix $\mathbf{\Theta}$ would be infinitely wide, prohibiting practical reconstruction, which relies on solving the underdetermined linear set of equations given by $\mathbf{\Theta}$. A signal that is sparse in a continuous space will thus have to be sufficiently represented, and similarly sparse, on a *discretized* version of the continuous space.

So far, two main techniques have been proposed for the application of CS to the acquisition of spectrally sparse continuous-time (SSCT) signals:

The Random Demodulator (RD) [42]-[44] is only suitable for detecting synchronized multi-tone signals with known and uniformly allocated tone locations. Fig. III-2 (top) illustrates such a suitable spectral scenario. To differentiate the basis functions of different tones, the RD employs a windowed integrator in the analog domain, shown on the bottom of Fig. III-2, and thus requires high timing accuracy in the analog circuitry

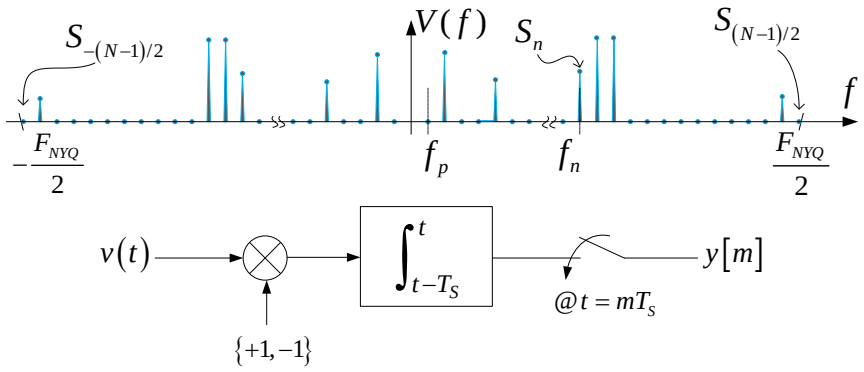

Fig. III-2 (top) An illustration of the uniform multi-tone spectral scenario suitable for the RD acquisition technique. (bottom) The RD employs a random mixer and windowed integrator to obtain compressive measurements.

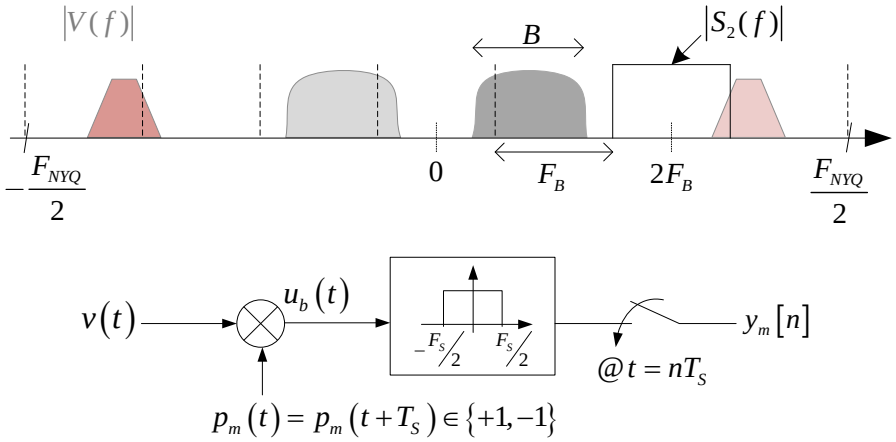

Fig. III-3 (top) An example of the MWC-type discretization of the spectrum into $N = 7$ slices of size $F_B \geq B$. (bottom) The MWC employs a random periodic spreader followed by a brick-wall filter (or a practical approximation) to obtain a compressive measurement.

[45], [48]. Multi-path versions of the RD have been discussed [46] and fabricated [47] for use in CR. Such structures inherit the disadvantages of the RD.

The Modulated Wideband Converter (MWC) [48]-[50] can acquire a more general signal that is "multi-band" in the frequency domain. The MWC divides the

spectrum into slices of bandwidth comparable to that of the largest expected signal band, as illustrated on the top of Fig. III-3, thus measuring the number of signal bands instead of the true spectral sparsity. To perform the rectangular slicing of the spectrum, the MWC requires brick-wall type selectivity and thus employs high-order analog filters [50], [51], incurring impractical cost in the analog hardware and limiting reconfigurability. Fig. III-3 (bottom) shows an individual measurement path of the MWC. In addition to the high analog hardware cost and constrained spectral scenarios, both the RD and MWC incur high CS reconstruction costs because they collapse the full spectrum into a single CS problem, which becomes impractical at high resolutions.

In this dissertation, we propose a novel compressive acquisition technique for SSCT signals termed *Random Filtering-by-Aliasing* (RFA)[1]. The RFA simultaneously addresses the analog measurement complexity, the CS problem size, and the inflexible sparsity models of existing schemes.

A high-level block diagram of the RFA is shown in Fig. III-4. The RFA performs arbitrary random filtering and aliasing in the analog domain according to our proposed *Filtering by Aliasing* (FA) concept. As illustrated in Fig. III-5, the RFA uses a variable-gain FA system with periodic spreader $d_m(t)$ followed by an arbitrary filter $h(t)$ to create an apparent filter $g_m(t)$ from the perspective of the sampler output of the m^{th} measurement. For the purposes of the proposed RFA, the periodic $d_m(t)$ can be drawn from a Bernoulli ±1 distribution, as in other CS schemes. Unlike the RD and MWC

[1] Not to be confused with Random Filters [52], the DT predecessor of the RD.

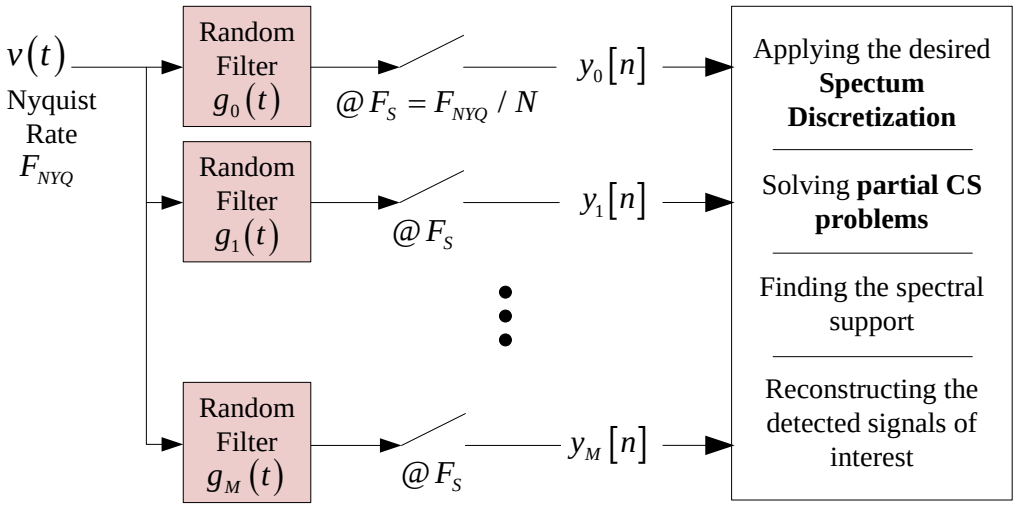

Fig. III-4 A block diagram of the proposed Random Filtering by Aliasing (RFA) system. The analog front-end performs coarse measurements by random filtering according to Filtering by Aliasing (Error! Reference source not found.). The digital back-end performs reconfigurable spectral discretization at a desired resolution and solves a set of small CS problems to reconstruct the signal or its support.

however, the filter $h(t)$ can be any type of simple filter, such as a single-pole filter, avoiding the complexity and/or accuracy requirements of the existing CS-for-SSCT techniques.

To enable such simple, and general, analog circuitry, the RFA (a) moves all spectral resolution and sparsity-model decisions to the low-cost, easily reconfigurable digital domain and (b) exploits the fact that measureFment accuracy can be relaxed when acquiring realistic noisy spectra. Moreover, as the digital RFA backend discretizes the spectrum according to a desired resolution, the wideband spectrum sensing problem is inherently partitioned into several small CS measurements. This significantly reduces the reconstruction complexity and enhances the partial detection probability.

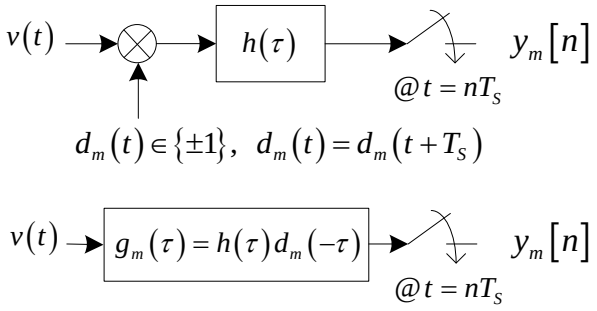

Fig. III-5 Each apparent random filter $g_m(t)$ of the RFA is achieved using a ±1 random periodic spreader $d_m(t)$ and an *arbitrary* filter $h(t)$. For the results in this manuscript, $h(t)$ is a single-pole filter, the simplest analog filter.

We will start with the matrix-vector CS formulation in section III.2. In section II.3, we will provide a unified description of how CS has been applied to SSCT signals and motivate the RFA. Section II.4 will describe the RFA concept and system. Section III.5 will dive into particular design trade-offs of the RFA. Section III.6 will provide simulation results while Section III.7 discusses the RFA complexity and programmability and compares it to the state-of-the-art.

III.2 Compressive Sensing Recap

Consider a signal vector **v** of length *N*. We say that **v** is *k*-sparse on some basis if there is a vector **s** with $k << N$ non-zero entries and an *N*x*N* matrix **Ψ** such that **v** = **Ψ s**.

Given that **v** is sufficiently represented by a sparse **s**, we can potentially make a measurement vector **y** of **v** that holds all information about **v** such that the length of **y** is $M < N$. The associated matrix-vector operation is given by (3.1), where **y** is expressed as

a transform $\boldsymbol{\Theta}$ of **s**. The reconstruction of the signal **v** from **y** follows from reconstructing the sparse **s**.

$$\mathbf{y}_{M\times 1} = \boldsymbol{\Phi}_{M\times N}\mathbf{v}_{N\times 1} = \boldsymbol{\Phi}\boldsymbol{\Psi}_{N\times N}\mathbf{s}_{N\times 1} = \boldsymbol{\Theta}_{M\times N}\mathbf{s} \tag{3.1}$$

Compressive Sensing (CS) theory [1]-[2], [37] tells us that the reconstruction of **s** from **y** is possible with high probability if the sensing matrix $\boldsymbol{\Theta}$ satisfies the *Restricted Isometry Property* (RIP) [37] with order *k*, where *k* is the number of nonzero elements in **s**. The RIP is given by (3.2) and can be intuitively seen as a requirement that the matrix $\boldsymbol{\Theta}$ preserve any *k*-sparse vector **s**, i.e. that any *k* columns of $\boldsymbol{\Theta}$ be sufficiently orthogonal.

$$(1-\delta_k)\|\mathbf{s}\|^2 \le \|\boldsymbol{\Theta}\mathbf{s}\|^2 \le (1+\delta_k)\|\mathbf{s}\|^2, \qquad \delta_k < 1 \tag{3.2}$$

It turns out that if $\boldsymbol{\Theta}$ is drawn from independent and identically distributed (i.i.d.) Gaussian or Bernoulli random variables [40], $\boldsymbol{\Theta}$ would satisfy the RIP with small δ_k for a number of measurements *M* in the order of O(k log(N/k)). Similarly, if measurement matrix $\boldsymbol{\Phi}$ is drawn from such random variables and $\boldsymbol{\Psi}$ is a unitary transformation then the sensing matrix $\boldsymbol{\Theta}$ would satisfy the same RIP conditions. Reconstruction of **s** from **y** can then be performed using several approaches. The optimal solution finds the sparsest **s** s.t. $\mathbf{y} = \boldsymbol{\Theta}\,\mathbf{s}$. The notion of "sparsest" or min$\|\mathbf{s}\|_0$ is practically approximated by min$\|\mathbf{s}\|_1$ and solved using linear programming [38]. With a slight penalty in the number of measurements, lower-complexity greedy pursuit algorithms [39] can be used.

We are not concerned, in this manuscript, with the particular CS reconstruction algorithms but instead with the CS measurement or acquisition method.

III.3 CS FOR SSCT SIGNALS: THE NEED FOR A NEW APPROACH

The class of band-limited continuous-time signals can be generally described by $v(t)$ in (3). The spectrum $V(f)$ of $v(t)$ contains envelopes $X_i(f)$ of arbitrary signals $x_i(t)$ centered at positions f_i on the frequency axis.

$$v(t) = \sum_i x_i(t) \exp(-j2\pi f_i) \tag{3.3}$$

Fig. III-1 shows an example spectrum $V(f)$ for a real-valued v(t). One can qualitatively describe the spectrum $V(f)$ as sparse by noticing that the spectral occupancy in Fig. III-1, the ratio of occupied spectrum to total available spectrum, is low. It is beneficial under such low spectral occupancy to take advantage of compressive techniques to avoid the cost of the sampling, analog-to-digital conversion, and digital processing of $V(f)$ at the Nyquist rate.

To apply CS theory to the measurement of SSCT signals, the spectral sparsity should be described on a discrete space such that $v(t)$ and/or its spectral support can be recovered from compressive measurements.

Multi-tone Signals: If the type and center frequencies of the signals $x_i(t)$ that can potentially exist in the spectrum were known, the spectral discretization can be tailored to such knowledge. For example, the RD [42]-[44] assumes that the signal $V(f)$ is "time-

frequency" sparse, i.e. is made of time-synchronized tones at multiples of some f_c, as in (3.4).

$$V(f) = \sum_{i=-(N-1)/2}^{+(N-1)/2} s_i \delta(f - if_c) \tag{3.4}$$

The sparsity is then described by the vector of tone values s_i for a total of N possible tones. The RD compressive measurements are obtained using a random ±1 spreader followed by a windowed integrator and sub-sampler. The RD is only suitable for detecting synchronized multi-tone signals with known and uniformly allocated tone locations, such as OFDM[2]. Any spectral content that lies outside the specific tone locations assumed by the RD disrupts the sparsity model and the associated recovery [45]. It also follows that the RD is sensitive to timing accuracies in the analog hardware, which describe the tone locations. Such sensitivity is theoretically discussed in [53] and [54] and numerically evaluated in [45].

Multi-carrier Signals: One can also envision a multi-carrier case where we know that the spectral envelope of $X_i(f)$ at $f = if_c$ is given by a pulse-shaping filter $B(f)$. Over one symbol period of the multi-carrier signal, $V(f)$ can be expressed as (3.5). The sparsity is similarly described by the vector of symbols s_i.

[2] In OFDM, each tone carries one value over one OFDM symbol period. The signal thus appears band-limited multi-tone to a synchronized receiver.

$$V(f) = \sum_{i=-(N-1)/2}^{+(N-1)/2} s_i B(f - if_c) \tag{3.5}$$

Multi-band Signals: If the possible carrier locations are unknown or belong to different systems, spectral discretization can be performed by re-slicing the spectrum as is done by the MWC [48]-[50]. The MWC assumes that $V(f)$ is made up of k signal bands $X_i(f)$, each with a maximum bandwidth B. $V(f)$ can thus be referred to as "block-sparse" [45] since the non-zero frequency components are grouped together. The MWC discretizes the Nyquist spectrum into N slices of size $F_B \geq B$ as given by (3.6), resulting in at most $2k$ non-zero bands $S_i(f)$ out of N. Fig. III-3 (top) illustrates the discretization of the MWC for a hypothetical $2k = 4$, $N = 7$ scenario.

$$V(f) = \sum_{i=-(N-1)/2}^{+(N-1)/2} S_i(f - iF_B) \tag{3.6}$$

At any time instant, the sparsity is thus described by the vector of time samples from each band as in (3.7).

$$v(t) = \sum_{i=-(N-1)/2}^{+(N-1)/2} s_i(t) \exp(j2\pi iF_B) \tag{3.7}$$

Each MWC compressive measurement is obtained using a random periodic ±1 spreader followed by a brick-wall filter and a sub-sampler. Because brick-wall filters are

not realistic, the board demonstration of the MWC [50] uses high-order analog filters followed by oversampling. The required digital equalization for such non-ideal MWC filters was discussed in [51]. Realizing high-order analog filters in integrated receivers is very difficult and incurs high power, area, and nonlinearity costs [7], [14]. Fixed high-order filters also prohibit system re-configurability by enforcing a certain sampling rate. This renders the MWC sensitive to the number of occupied bands [45] even under low spectral occupancy.

In addition to the hardware cost and limited re-configurability, the MWC and RD are also limited in terms of achievable spectral resolution because they collapse the full spectrum of interest into a single CS problem, which becomes impractical to solve when high spectral-resolution is desired.

The Need for a Better Approach: We find that, for a CS measurement to be of use in CR, higher spectral resolution, lower measurement cost, and better programmability than existing solutions are simultaneously required:

1- *Lower measurement cost:* Simply put, CS is intended to lower the cost of sampling and data processing. It accordingly is undesirable to incur high analog hardware cost prior to sampling, especially in a world where digital processing is significantly cheaper. It is also undesirable to solve large CS problems in the digital domain.

2- *Higher Spectral Resolution:* Realistic signal bands exhibit a wide range of bandwidths. For example [7], a GSM signal has a bandwidth of 200 kHz while an 802.11 signal has a bandwidth of 20-40 MHz. For a software-defined or cognitive

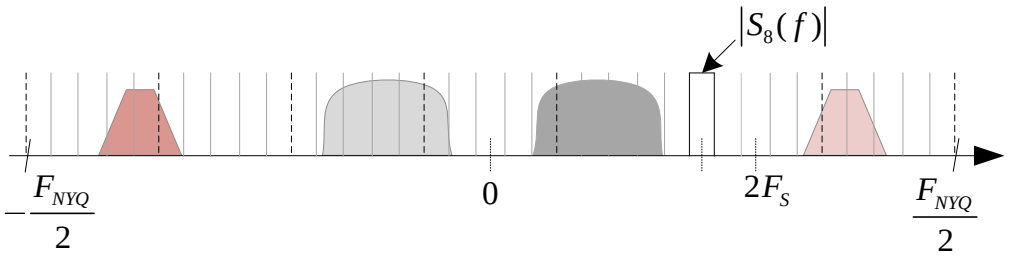

Fig. III-6 Higher resolution discretization allows for a better measure of sparsity and better acquisition of narrowband signals. Compare with Fig. III-3 (top).

acquisition system to replace traditional dedicated radios, it has to correctly and efficiently handle such a large range of bandwidth. If we were to perform spectrum slicing according to the largest expected signal bandwidth, as in the MWC, we would (a) fail to measure the true spectral occupancy for cognitive access, (b) jeopardize the success of the CS measurement scheme by misrepresenting the true sparsity, and (c) possibly bury narrowband signals in the noise of large spectral slices. Consider, for example, the illustration in Fig. III-3 where the MWC-type spectral slicing for a maximum bandwidth *B* shows 100% (7 out of 7) occupied bands whereas an increased resolution slicing in Fig. III-6 shows ~52% (18 out of 35) occupied bands.

3- *Programmability:* A reliable CS measurement technique should be able to support (a) a varying spectral sparsity by adjusting its sampling rate and (b) a variable spectral content distribution by varying its sensing resolution.

As we describe the Random Filtering by Aliasing concept, design, performance, and complexity, it will become apparent how the RFA approaches the seemingly contradictory cost, resolution, and programmability targets.

III.4 RANDOM FILTERING BY ALIASING

In agreement with the motivation presented in section III.3, the RFA makes low-complexity analog measurements, performs programmable high-resolution digital spectral discretization, and solves the large CS problem in small partitions. As we describe the various RFA aspects, we will refer to the unified set of parameters summarized in Table 1.

A. Analog-domain measurements

Consider the RFA block diagram in Fig. III-4. The m^{th} analog measurement branch of the RFA performs mixing with a random ±1 signal $d_m(t)$, filtering with a weak low-pass filter $h(t)$, and sampling at rate F_S. According to Filtering by Aliasing from chapter II, if the mixing sequence $d_m(t)$ is periodic of period $T_S = 1/F_S$, i.e. $d_m(t) = d_m(t + T_S)$, the system exhibits a time-invariant filter $g_m(t)$ at the sampling instances nT_S. The random filter $g_m(t)$ is given by (3.8) and its frequency response is given by (3.9). Because $d_m(t)$ is periodic of period T_S, the frequency response $D_m(f)$ is completely given by the tone values $D_{m,\mu} = D_m(\mu F_S)$ that appear in (3.9).

$$g_m(t) = h(t)\, d_m(-t) \tag{3.8}$$

Table III-1

LIST OF RFA PARAMETERS

	Significance	*Relationship*
F_{NYQ}	Nyquist bandwidth of $V(f)$	
F_S	Sampling rate of each measurement $y_m[n]$	
N	Ratio of F_{NYQ} to F_S. The number of aliasing spectral slices	$N = F_{NYQ} / F_S$
η	Index of an aliasing slice	$\eta \in \left[(1-N)/2, (N-1)/2\right]$
M	Number of measurements	
m	Index of a measurement	$m \in [0, M-1]$
F_{res}	Resolution of the RFA	
L	Ratio of F_S to F_{res}. The number of bins per spectral slice	$L = F_S / F_{res}$
l	Index of a bin in a spectral slice	$l \in \left[(1-L)/2, (L-1)/2\right]$
k	The sparsity in the spectrum. The number of non-zero elements for a given discretization.	$k \le LN$

$$G_m(f) = \sum_{\mu=-\infty}^{+\infty} D_{m,\mu} H(f + \mu F_S) \tag{3.9}$$

To give an example, if $h(t)$ is a windowed integrator, $g_m(t)$ is a random continuous-time FIR filter given by $d(-t)$, $t \in [0, T_S]$. If $h(t)$ is a simple single-pole filter, $g(t)$ is a random FIR filter followed by a deterministic IIR filter. Until we address the

design parameters of the RFA in section III.5, we will assume that $d_m(t)$ and $H(f)$ are properly chosen such that $G_m(f)$ is an arbitrary random filter.

After sampling, the output $y_m[n]$ exhibits a frequency response $Y_m(\omega)$ that is, effectively, a filtered and aliased version of $V(f)$ accordingly to (3.10).

$$Y_m(\omega) = \frac{1}{T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} V\left(\frac{\omega F_S}{2\pi} + \eta F_S\right) G_m\left(\frac{\omega F_S}{2\pi} + \eta F_S\right) \tag{3.10}$$

The summation in (3.10) is limited to the Nyquist bandwidth F_{NYQ} of $V(f)$, where $F_{NYQ} = NF_S$. It is important to note here that even though $G_m(f)$ extends over the whole Nyquist band of interest, the sampler need not have such a wide bandwidth because the down-conversion of the high frequency signals occurs at the ±1 mixer prior to filtering and sampling (see section II.2.A).

Given $G_m(f)$ from (3.8) or (3.9), we can see that the value of $Y_m(\omega)$ in (3.10) is a weighted sum of the values of $V(f)$ at offsets $\omega F_S/2\pi$ from the center of every aliasing slice. Using the M available measurements, we can write the CS problem (3.11) for any $\omega \in \{-\pi, \pi\}$, where $\mathbf{y}$, $\mathbf{G}$, and $\mathbf{v}$ are given by (3.12)-(3.14).

$$\mathbf{y}(\omega)_{M\times 1} = \mathbf{G}(\omega)_{M\times N}\, \mathbf{v}(\omega)_{N\times 1} \tag{3.11}$$

$$\mathbf{y}(\omega) = \begin{bmatrix} Y_0(\omega) & \cdots & Y_{M-1}(\omega) \end{bmatrix}^T \tag{3.12}$$

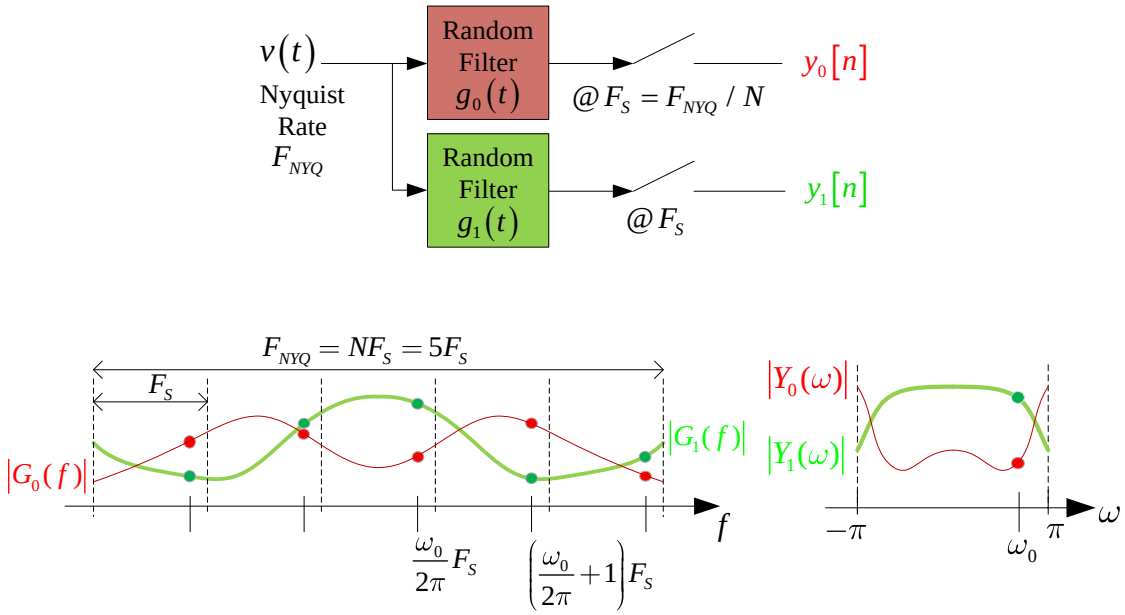

Fig. III-7 An illustration of how the analog acquisition stage (top) provides CS measurements. (bottom left) The magnitude plot of the M = 2 random filters across N = 5 spectral slices. (bottom right) The corresponding measurement results: for each frequency ω_o, 2 measurements of 5 original spectral locations are available.

$$\mathbf{G}(\omega)=\frac{1}{T_S}\begin{bmatrix} G_0\left(\frac{\omega F_S}{2\pi}-\frac{N-1}{2}F_S\right) & \cdots & G_0\left(\frac{\omega F_S}{2\pi}+\frac{N-1}{2}F_S\right) \\ \vdots & & \vdots \\ G_{M-1}\left(\frac{\omega F_S}{2\pi}-\frac{N-1}{2}F_S\right) & \cdots & G_{M-1}\left(\frac{\omega F_S}{2\pi}+\frac{N-1}{2}F_S\right) \end{bmatrix} \tag{3.13}$$

$$\mathbf{v}(\omega)=\left[V\left(\frac{\omega F_S}{2\pi}-\frac{N-1}{2}F_S\right) \quad \cdots \quad V\left(\frac{\omega F_S}{2\pi}+\frac{N-1}{2}F_S\right)\right]^T \tag{3.14}$$

For a graphical interpretation, Fig. III-7 sketches two random filters $G_0(f)$ and $G_1(f)$ associated with M = 2 measurements $Y_0(\omega)$ and $Y_1(\omega)$. In this example, N = 5 spectral slices. The figure annotates the measurements $\mathbf{y}(\omega_o)_{2x1} = \mathbf{G}(\omega_o)\mathbf{v}(\omega_o)$ as well as

the filter gains $\mathbf{G}(\omega_o)_{2x5}$ of the 5 measured frequency locations at offsets ω_o in each spectral slice. By extrapolating this example to any number L of CS problems for L locations ω_o, one can visualize how an arbitrarily high resolution measurement of the spectrum can be obtained from a single set of coarse analog measurements $\mathbf{y}(\omega)$, provided that the sensing matrix $\mathbf{G}(\omega)$ satisfies the RIP (3.2) for those ω_o.

For non-periodic signals, it is not meaningful to talk about a frequency "point" ω_o. Instead, we refer to a frequency "bin" of nonzero bandwidth centered at ω_o, which we discuss next.

B. Digital Spectral-Discretization:

As illustrated in Fig. III-7, the analog measurement stage of the RFA is intended to provide coarse slicing of a randomly weighted spectrum. The high, and reconfigurable, spectral resolution of the RFA is then achieved by a digital filtering stage that operates on the measurements $y_m[n]$ or $Y_m(\omega)$.

We will refer to $U_{m,l}(\omega)$ as the l^{th} bin of the m^{th} measurement band $Y_m(\omega)$. $U_{m,l}(\omega)$ is obtained from $Y_m(\omega)$ through a filter $B(\omega)$ centered at $2\pi/L$, as in (3.15).

$$U_{m,l}(\omega) = Y_m(\omega) B\left(\omega - \frac{2\pi l}{L}\right) \tag{3.15}$$

Fig. III-8 sketches the ideal $B(\omega)$ for $L = 11$ on the left and a practical alternative on the right. By performing slicing in the digital domain, we can accurately design a practical filter $B(\omega)$ to preserve all information about $Y_m(\omega)$ using, for example, a Nyquist

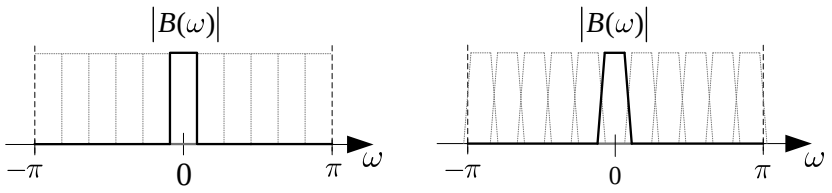

Fig. III-8 (left) Ideal filter that slices that DT spectrum into L = 11 bins. (right) A practical alternative that can be designed to provide full reconstruction.

filter-bank [55]. Without loss of generality, we assume throughout this manuscript that L is odd and that the bin centers are separated by $2\pi/L$. We accordingly refer to $F_{res} = F_S/L = F_{NYQ}/LN$ as the resolution of the RFA.

To connect the measurement bins $U_{m,l}(\omega)$ to the signal spectrum $V(f)$, we define $S_{\eta,l}(f)$, the l^{th} bin of the η^{th} spectral slice of $V(f)$, in (3.16). The reader might find it useful to look ahead to Fig. III-10 for a graphical interpretation of $S_{\eta,l}(f)$. Using (3.10), $U_{m,l}(\omega)$ can then be written in terms of $S_{\eta,l}(f)$ in (3.17).

$$S_{\eta,l}(f) = \begin{cases} V(f+\eta F_S)\, B\left(\dfrac{2\pi f}{F_S} - \dfrac{2\pi l}{L}\right) & , f \in \left[-\dfrac{F_S}{2}, \dfrac{F_S}{2}\right] \\ 0 & , \quad \text{otherwise} \end{cases} \tag{3.16}$$

$$U_{m,l}(\omega) = \frac{1}{T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} S_{\eta,l}\left(\frac{\omega}{2\pi} F_S\right) G_m\left(\frac{\omega}{2\pi} F_S + \eta F_S\right) \tag{3.17}$$

For a highly selective filter B(ω), such as those in Fig. III-8, the meaningful content in the measurement bin $U_{m,l}(\omega)$ is limited to a bandwidth of ~$2\pi/L$ centered

around $2\pi l/L$ of the measurement band $Y_m(\omega)$. The meaningful content of the signal bin $S_{\eta,l}(f)$, as given by (3.16), is thus similarly limited to a bandwidth of $\sim F_{res}$ centered around $\eta F_S + lF_{res}$.

Assuming that, by design, the random analog filter $G_m(f)$ exhibits *small* variation within one bin width, i.e. within a bandwidth of $\sim F_{res}$, we approximate it by its value at the center of the bin. We denote the value of $(1/T_S)G_m(f)$ at the center of the l^{th} bin in the η^{th} slice by $\theta_{m,n,l}$, as given by (3.18). We can then re-write (3.17) as the approximation in (3.19), which approaches equality as L increases (F_{res} decreases). We will attend to implications of this approximation in the section III.5.C.

$$\theta_{m,\eta,l} = \frac{1}{T_S} G_m\left(\eta F_S + lF_{res}\right) \tag{3.18}$$

$$U_{m,l}\left(\omega\right) \simeq \sum_{\eta=-(N-1)/2}^{+(N-1)/2} \theta_{m,\eta,l} S_{\eta,l}\left(\frac{\omega}{2\pi} F_S\right) \tag{3.19}$$

The frequency-domain relationship (3.19) can be similarly written in (3.20) for the time-domain samples $u_{m,l}[n]$ of measurement bins $U_{m,l}(\omega)$.

$$u_{m,l}\left[n\right] \simeq \sum_{\eta=-(N-1)/2}^{+(N-1)/2} \theta_{m,\eta,l} s_{\eta,l}\left(nT_S\right) \tag{3.20}$$

By combining all M measurements for bin l, we obtain the matrix-vector formulation (3.21), where $\mathbf{u}_l[n]$, $\mathbf{\Theta}_l$ and $\mathbf{s}_l(nT_S)$ are given by (3.22)-(3.24) respectively.

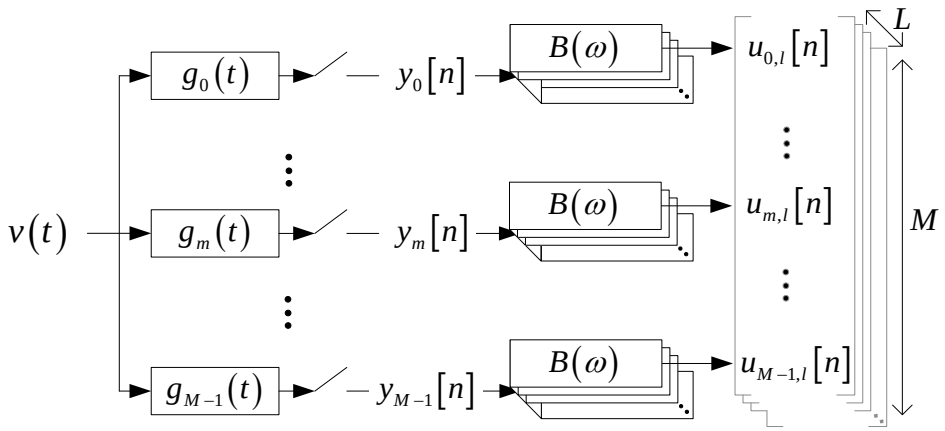

Fig. III-9 The RFA low-cost high-resolution discretization: Random non-selective analog filters and low-bandwidth samplers provide coarse slicing. Decimated digital filter banks provide accurate programmable discretization.

$$\mathbf{u}_l[n] \simeq \mathbf{\Theta}_l \mathbf{s}_l(nT_S) \tag{3.21}$$

$$\mathbf{u}_l[n] = \begin{bmatrix} u_{0,l}[n] & \cdots & u_{M-1,l}[n] \end{bmatrix}^T \tag{3.22}$$

$$\mathbf{\Theta}_l = \begin{bmatrix} \theta_{0,-\frac{N-1}{2},l} & \cdots & \theta_{0,+\frac{N-1}{2},l} \\ \vdots & \theta_{m,\eta,l} & \vdots \\ \theta_{M-1,-\frac{N-1}{2},l} & \cdots & \theta_{M-1,+\frac{N-1}{2},l} \end{bmatrix}_{M\times N} \tag{3.23}$$

$$\mathbf{s}_l(nT_S) = \begin{bmatrix} s_{-\frac{N-1}{2},l}(nT_S) & \cdots & s_{\frac{N-1}{2},l}(nT_S) \end{bmatrix}^T \tag{3.24}$$

The relationship (3.21) is a CS measurement with a sensing matrix $\mathbf{\Theta}_l$ and sparse vector $\mathbf{s}_l(nT_S)$. The sparsity of $\mathbf{s}_l$ and the RIP of $\mathbf{\Theta}_l$ will be discussed in section III.5.

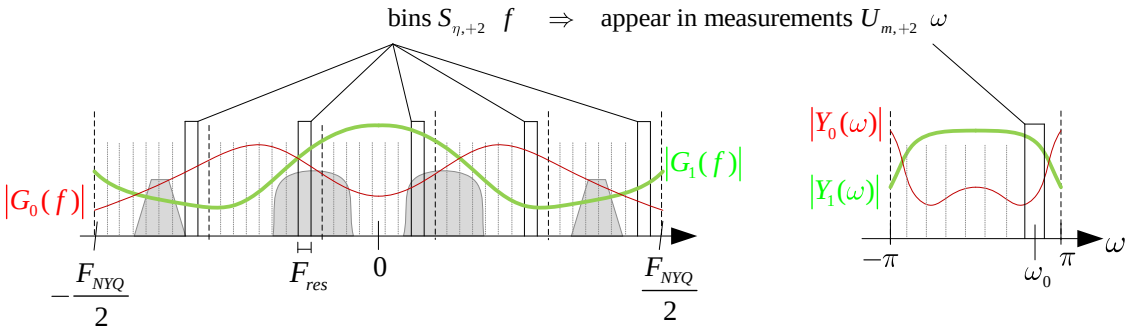

Fig. III-10 The RFA discretization with $N = 5$, $L = 9$ applied to the hypothetical spectrum of Fig. III-3 and Fig. III-6. The bins $S_{\eta,l}(f)$ for $l = 2$ are labeled for illustration.

Fig. III-9 illustrates the complete path from the wideband signal $v(t)$ to the bin measurements $u_{m,l}[n]$ highlighting the creation of the L vectors $\mathbf{u}_l[n]$. In a practical realization, the outputs $u_{m,l}[n]$ would be decimated and, along with the digital filters $B(\omega)$, would run at an appropriate low-rate comparable to F_{res}. We attend to such complexity matters in section III.7.

C. On solving the Partitioned CS Problems:

After the digital discretization stage, a total of L CS problems of the form (3.21) are available. Each CS problem is particular to the set of signal bins that lie at center frequencies $lF_{res} + \eta F_S$ for all η, as given by (3.24). To help visualize this partitioning, Fig. III-10 re-labels the spectral discretization example from Fig. III-6 according to the RFA indexing with $N = 5$ and $L = 9$. Note that, for ease of illustration, the bins $S_{\eta,l}(f)$ in the figure are sketched as rectangular, implying the use of the ideal $B(\omega)$ from Fig. III-8, which is not necessary.

Assuming valid sparsity and RIP conditions, solving (3.21) would identify the indices $\hat{\eta}$ and the corresponding waveforms $s_{\hat{\eta},l}\left(nT_S\right)$ that contain signal information.

Although we are not concerned in this work with the particular CS reconstruction approaches, we note that performing costly CS reconstruction at each time sample nT_S is inefficient because realistic spectral supports do not change at a sample by sample basis. Instead, as proposed by the authors of the MWC in an earlier work [55], it behooves us to perform CS reconstruction only frequently enough to track the spectral support, i.e. the indices $\hat{\eta}$ for the RFA. Doing so using accumulated information also helps combat noise [61]. Knowing the spectral support, matrix multiplication (3.25) is sufficient to reconstruct $s_{\hat{\eta},l}$ from $\mathbf{u}_l[n]$ at a desired rate. $\mathbf{\Theta}^{\dagger}_{\hat{\eta},l}$ is the pseudo-inverse of $\mathbf{\Theta}_{\hat{\eta},l}$, the matrix of columns $\hat{\eta}$ of $\mathbf{\Theta}_l$.

$$\mathbf{s}_{\hat{\eta},l}\left(nT_S\right)=\mathbf{\Theta}^{\dagger}_{\hat{\eta},l}\mathbf{u}_l\left[n\right] \tag{3.25}$$

We note again that the rate of $\mathbf{u}_l[n]$, and accordingly the rate of reconstruction of $\mathbf{s}_{\hat{\eta},l}$, would in practice be much lower than F_S since the bandwidth of $S_{\hat{\eta},l}\left(f\right)$ is ~F_{res}. Using the solutions for all L CS problems, the RFA can reconstruct signals that span one or more of the signal bins in $V(f)$.

It is important to recall here that, despite using existing methods to solve each CS problem, the RFA partitions the spectrum sensing problem of resolution $F_{res} = F_S/LN$ into L CS problems with a sparse vectors $\mathbf{s}_l$ of length N instead of LN. We will discuss the

complexity implications of such reduction in the CS problem size combined with the decimated processing rate and the simple analog hardware in section III.7.

III.5 RFA Design Considerations

As discussed so far, the RFA moves the spectral discretization task to the digital domain to simultaneously simplify the analog measurement and partition the CS problem. The analog measurement in particular is simplified because spectral bins that belong to the same CS problem, as shown in Fig. III-10, are not adjacent in the frequency domain. As a result, the analog random filtering, prior to sampling, does not need to provide brick-wall selectivity. $G_m(f)$ can be made to vary slowly across f such that it is incoherent over a bandwidth of F_S but highly coherent over a bandwidth of F_{res}. This idea is qualitatively illustrated in Fig. III-11 for $L = 11$. The incoherence in $G_m(f)$ over F_S allows us to obtain sufficient randomness in the sensing matrices Θ_l while the high coherence over F_{res} allows us to apply the approximation in (3.19) to describe the matrix-vector CS problem.

In this section we will investigate the coherence of $G_m(f)$ and its consequences. We will also investigate the effect of the partitioned CS problem on the number of CS measurements.

A. The coherence of $G_m(f)$ and the choice of h(t) or H(f)

From the treatment of multi-path channels in wireless communications theory [57], we know that the frequency-domain coherence of an impulse response is inversely proportional to its delay spread, a measure of the concentration of power in the impulse

response. Recall from (3.2) and Fig. III-5 that the apparent filter impulse response $g_m(t)$, according to Filtering by Aliasing, is given by $h(t)d_m(-t)$. We refer to $h(t)$ as the "physical" filter as compared to the "apparent" filter $g_m(t)$. Given that $d_m(t)$ is a random ±1 sequence, the delay spread of $g_m(t)$ is given by the impulse response $h(t)$.

In the spirit of compressive sensing, we are only interested in very simple filters $h(t)$ that keep the analog measurement cost to a minimum. The single-pole filter, inherent in most any analog circuit, is the simplest and most natural candidate. The impulse response of a single-pole $h(t)$ with time constant τ_p is given by (3.26). Fig. III-12a plots such exponentially decaying impulse response for $\tau_p = T_S$ as well as a sample resulting $g_m(t)$. We will assume a single-pole $h(t)$ for the remainder of the manuscript.

$$h(t) = \frac{1}{\tau_p}\exp\left(-t/\tau_p\right) \tag{3.26}$$

B. The incoherence of $G_m(f)$ and the RIP of $\boldsymbol{\Theta}_l$

The sensing matrix for CS problem l is given by $\boldsymbol{\Theta}_l$ in (3.23). Using (3.18) and (3.9), the elements of the sensing matrix can be expanded as in (3.27) in terms of $H(f)$ and $D_{m,\mu}$. Because $H(f)$ is a low-pass filter, we can assume that $|H(f)|$ is negligible beyond some RF_S and the summation in (3.27) is practically dominated by $2R+1$ terms as shown. $\boldsymbol{\Theta}_l$ can thus be described as the matrix multiplication of $\mathbf{D}$ and $\mathbf{H}_l$ in (3.28).

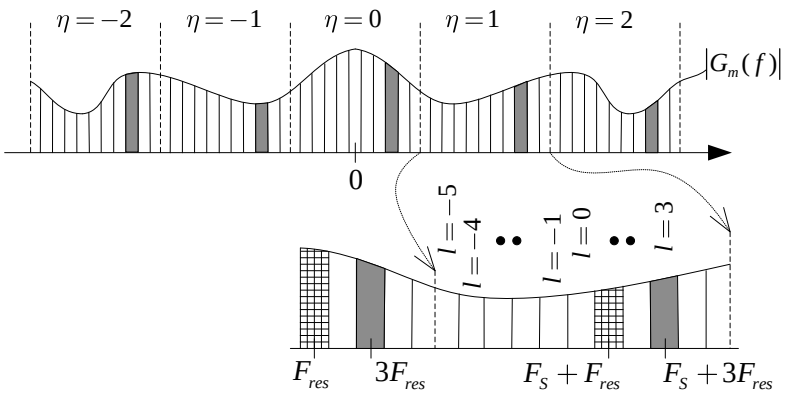

Fig. III-11 An illustration of how the RFA allows a slowly varying $G_m(f)$ to provide incoherence across different bins $S_{\eta,l}(f)$ belonging to a CS problem l while providing coherent gain over the bandwidth F_{res} of an individual bin.

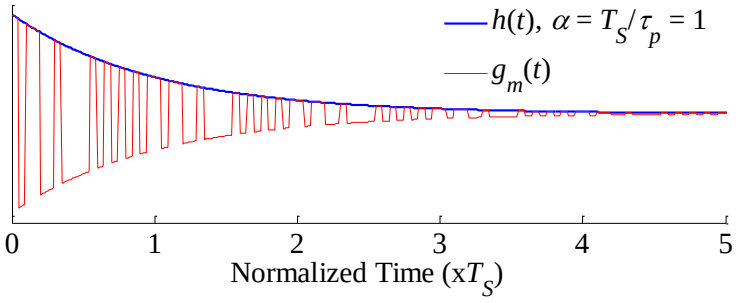

(a)

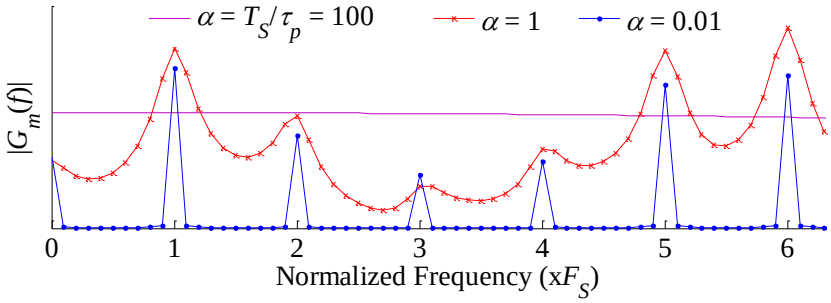

(b)

Fig. III-12 (a) The impulse response of a single-pole $h(t)$ with time constant $\tau_p = T_S$ and a sample resulting $g_m(t)$. (b) The magnitude plot of $G_m(f)$ for different $\tau_p = T_S/\alpha$, showing the effect of τ_p on the frequency coherence of $G_m(f)$.

$$\theta_{m,\eta,l} = \sum_{\mu=-R+\eta}^{+R-\eta} D_{m,\mu} H\left(lF_{res} + \eta F_S + \mu F_S\right) \tag{3.27}$$

$$\boldsymbol{\Theta}_l = \mathbf{D}_{M\times(N+2R)} \mathbf{H}_{l(2R+N)\times N} \tag{3.28}$$

The m^{th} row of **D** corresponds to the Fourier transform of the periodic signal $d_m(t)$ at frequencies μF_S, $\mu \in [-R-(N-1)/2,\ R+(N-1)/2]$. Because of the unitary nature of the Fourier transformation, **D** inherits the RIP characteristics of the M random ±1 signals $d_m(t)$. By allowing $d_m(t)$ to switch at least at a rate F_{NYQ}, **D** would match the RIP characteristics of a Bernoulli matrix of size MxN, which is what we desire.

The RIP of the sensing matrix $\boldsymbol{\Theta}_l$ is accordingly determined by the coherence of the matrix $\mathbf{H}_l$. The coherence of $\mathbf{H}_l$ [58] is defined as (3.29) where $\mathbf{h}_{l,\eta}$ is the η^{th} column of $\mathbf{H}_l$, given by $H(lF_{res} + \eta F_S + \mu F_S)$, $\mu \in [-R-(N-1)/2,\ R+(N-1)/2]$.

$$\operatorname{coh}\left(\mathbf{H}_l\right) = \max_{\eta_1<\eta_2} \frac{\left|\left\langle \mathbf{h}_{l,\eta_1}, \mathbf{h}_{l,\eta_2} \right\rangle\right|}{\left\|\mathbf{h}_{l,\eta_1}\right\|_2 \left\|\mathbf{h}_{l,\eta_2}\right\|_2} \tag{3.29}$$

The only parameter of the filter $H(f)$ is the time constant τ_p. The columns of $\mathbf{H}_l$, for any F_S, are thus a function of the ratio $\alpha = T_S/\tau_p$. Fig. III-13 plots, on the left, the coherence of H_l as a function of α for different bin indices l. As one would expect, as α decreases, the coherence of the matrix H_0 decreases towards zero because $H(f)$ approaches a frequency domain impulse. The coherence of H_l for $l \neq 0$, however, exhibits a minimum beyond which decreasing α is futile.

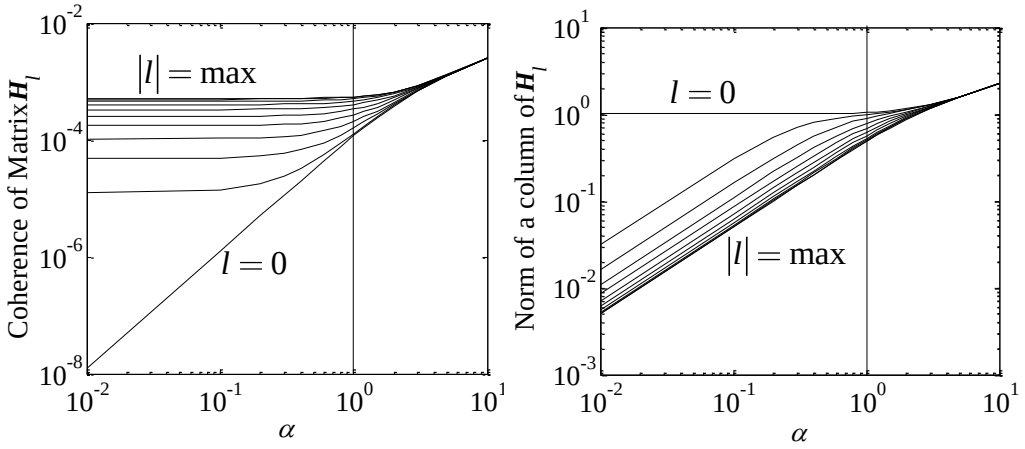

Fig. III-13 The coherence of $\mathbf{H}_l$ (left) and the norm of the columns of $\mathbf{H}_l$ (right) plotted against the inverse of the single-pole time-constant normalized to T_S. Plots show that reducing τ_p below T_S negatively affects $\mathbf{H}_l$ for $l \neq 0$.

Fig. III-13 also plots, on the right, the norm of the columns $\mathbf{h}_{l,\eta}$ of $\mathbf{H}_l$ as a function of α for different bin indices l. As expected again, the norm of columns $\mathbf{h}_{0,\eta}$ converges to unity as α decreases while the norm of columns $\mathbf{h}_{l,\eta}$ for $l \neq 0$ continues to decrease. This is an undesired effect because the signal bins belonging to CS problem $l \neq 0$ would exhibit lower gain.

To illustrate the effects of α on the gain and coherence of the apparent filter $G_m(f)$, Fig. III-12b plots three apparent filters $G_m(f)$ obtained using the same random sequence $d_m(t)$ but different time constants for $h(t)$. Notice how α = 100 causes undesired frequency coherence while α = 0.01 causes diminished gain away from the multiples of F_S.

Because all L CS problems use the same number of measurement M, the results in Fig. III-13 tell us that the performance of the RFA is limited by the RIP of $\mathbf{\Theta}_l$ for

$|l| = (L\text{-}1)/2$ which exhibits a floor below $\alpha = 1$. Accordingly, $\alpha < 1$ is not desired due to the loss in the relative gain of different bins (Fig. III-13 right) as well as the associated loss in measurement accuracy, which we discuss next.

C. The coherence of $G_m(f)$ and the measurement accuracy

To understand the effect of the approximation in (3.19) on the operation of the RFA, we look at the error caused by this approximation. From the perspective of the reconstruction stage, the approximate measurement (3.19) or (3.20) is treated as the ideal CS measurement. We rename the measurement from (3.20) as $\ddot{u}_{m,l}[n]$ and write it as (3.30) by replacing $\theta_{m,n,l}$ with the time-domain integral of $g_m(t)$.

$$\ddot{u}_{m,l}[n] = \frac{1}{T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} s_{\eta,l}(nT_S) \int_0^\infty g_m(\tau) e^{j2\pi(lF_{res}+\eta F_S)\tau} d\tau \tag{3.30}$$

The actual RFA measurement, however, as given by (3.17), translates to the time-domain relationship (3.31).

$$u_{m,l}[n] = \frac{1}{T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} \int_0^\infty g_m(\tau) e^{j2\pi(lF_{res}+\eta F_S)\tau} s_{\eta,l}(nT_S - \tau) d\tau \tag{3.31}$$

We can then write $\mathbf{u}_l[n]$ in terms of the ideal measurement vector $\ddot{\mathbf{u}}_l[n]$ and some measurement error in (3.32). An estimate of the elements of $\mathbf{e}_l$[n] is given in (3.33) and derived in Appendix C for a $g_m(t)$ with short delay spread.

$$\mathbf{u}_l[n] = \ddot{\mathbf{u}}_l[n] + \mathbf{e}_l[n] \tag{3.32}$$

$$e_{m,l}[n] \simeq \frac{j}{2\pi T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} s'_{\eta,l}(nT_S) G'_m(lF_{res} + \eta F_S) \tag{3.33}$$

It is interesting to note here that instead of treating $\mathbf{e}_l[n]$ as a measurement error, we can potentially include it in the reconstruction process since it is given by a random measurement of the derivative of $\mathbf{s}_l(nT_S)$, which has the same support as $\mathbf{s}_l(nT_S)$. This, however, is not necessary.

As will be demonstrated in section II.6.D, a reasonably low measurement error will have negligible effect on a realistic noisy sparse spectrum. In fact, for a CS problem with M measurements and k nonzero elements out of N, the work in [57] showed that the in-band signal-to-noise ratio (in-band SNR) of the original signal $v(t)$ will appear N/M times worse in the final reconstructed-SNR (RSNR) of an oracle CS solver. In contrast, the measurement-SNR (MSNR), defined similar to (3.34), appears enhanced by M/k in the RSNR.

$$\text{MSNR} = \frac{E\left[\left\|\ddot{\mathbf{u}}_l[n]\right\|^2\right]}{E\left[\left\|\mathbf{e}_l[n]\right\|^2\right]} \tag{3.34}$$

Eq. (3.35) gives an estimate of the MSNR, derived in Appendix D, for a $g_m(t)$ generated using a single-pole $h(t)$ with a small time-constant τ_p and a random ±1 $d_m(t)$. For a general analysis we use the variable $\alpha = T_S/\tau_p$ to describe τ_p.

$$\mathrm{MSNR} \simeq \frac{6/\pi^2}{\left(F_{res}\tau_p\right)^2} = \frac{6}{\pi^2}\alpha^2 L^2 \tag{3.35}$$

The square dependence of the MSNR on L is intuitive and encouraging. The MSNR can be made increasingly better by moving more of the spectral discretization into the digital domain, using the same gm(t) and for the same desired Fres. The dependence on α is also intuitive since increasing α enhances the frequency-coherence of Gm(f), which can be seen in Fig. III-12b. As is demonstrated in section III.6.C, the MSNR estimate (3.35) is accurate for large L, where (3.33) is valid.

D. The partitioned CS problem & the minimum required M

Assuming that a total k out of the LN bins in the spectrum are occupied, the partitioning of the high resolution problem distributes the k occupied bins across L CS problems. It is important then to address the effect of such re-distribution on the required number of measurements M.

We can describe the combined L CS measurements (3.21) as one large measurement given by (3.36), where the sensing matrix is LMxLN block-diagonal with diagonal entries $\boldsymbol{\Theta}_l$.

$$\begin{bmatrix} \mathbf{u}_{-\frac{L-1}{2}}[n] \\ \vdots \\ \mathbf{u}_{\frac{L-1}{2}}[n] \end{bmatrix} \simeq \begin{bmatrix} \boldsymbol{\Theta}_{-\frac{L-1}{2}} & & \\ & \ddots & \\ & & \boldsymbol{\Theta}_{\frac{L-1}{2}} \end{bmatrix}_{LM \times LN} \begin{bmatrix} \mathbf{s}_{-\frac{L-1}{2}}(nT_S) \\ \vdots \\ \mathbf{s}_{\frac{L-1}{2}}(nT_S) \end{bmatrix} \tag{3.36}$$

It is evident that if (3.36) can be solved, for any k occupied bins, with LM measurements, then the RFA would require M measurements. We are interested, however, in how M compares to the number of measurements M_{full} required for solving (36) using a dense sensing matrix. We know that a sensing matrix drawn from a random distribution would require M_{full} in the order of $O(k \log(LN/k))$.

For our partitioned approach, if we are interested in successful reconstruction across all L CS problems and for arbitrary located signal bins, we have to ensure that M satisfies the scenario where all the occupied signal bins occur in a single CS problem. If $\boldsymbol{\Theta}_l$ were to exhibit similar RIP characteristics as the random matrix, M would be on the order of $O(k \log(N/k))$, similar to M_{full}, indicating that our partitioning reduces the M_{full}xLN problem to L problems of size ~M_{full}xN.

In realistic sparse spectra, however, the k occupied signal bins of bandwidth ~F_{res} would be grouped together to form signal bands of bandwidths $\geq F_{res}$. The number of occupied bins appearing in a single CS problem would thus be practically much lower than k and M would be much smaller than M_{full}. Furthermore, an intelligent use of the RFA would adaptively alter the choice of F_S and L to redistribute the occupied bins more uniformly over the CS problems if necessary.

III.6 SIMULATION RESULTS

We remind the reader that the RFA is not associated with a particular CS reconstruction algorithm or support recovery schedule. To demonstrate the performance of the RFA as a method of acquiring CS measurements, however, we provide spectral support detection results as well as signal reconstruction results from Monte Carlo simulations.

A. The Simulation Setup

The results are generated using time-domain MATLAB simulations of the RFA. We generate an oversampled input signal with a pre-designed spectral occupancy and contaminate it with noise. As illustrated in Fig. III-9, the noisy signal $v(t)$ is mixed with M signals $d_m(t)$ (switching at $2F_{NYQ}$), filtered with $h(t)$, sampled (decimated), and passed through a Nyquist filter bank that provides the L bin signals $u_l[n]$ corresponding to L CS problems. All rates are normalized to F_S such that the results can be scaled to any frequency. The time constant of $h(t)$, unless noted, is $\tau = T_S$.

As motivated in section III.4.C, we solve each of the L CS problems using the CTF approach from [55] to accumulate signal information followed by Orthogonal Matching Pursuit (OMP) [59] to identify the strongest signal locations. The identified signals from all L CS problems are then sorted by their energy to find the strongest k signal bins $S(f)$ where k is the number of occupied bins $S(f)$ out of the total LN available bins. We accordingly call k the sparsity and k/LN the spectral occupancy. Were applicable, individual signal bins are reconstructed using (25) and combined for SNR calculation.

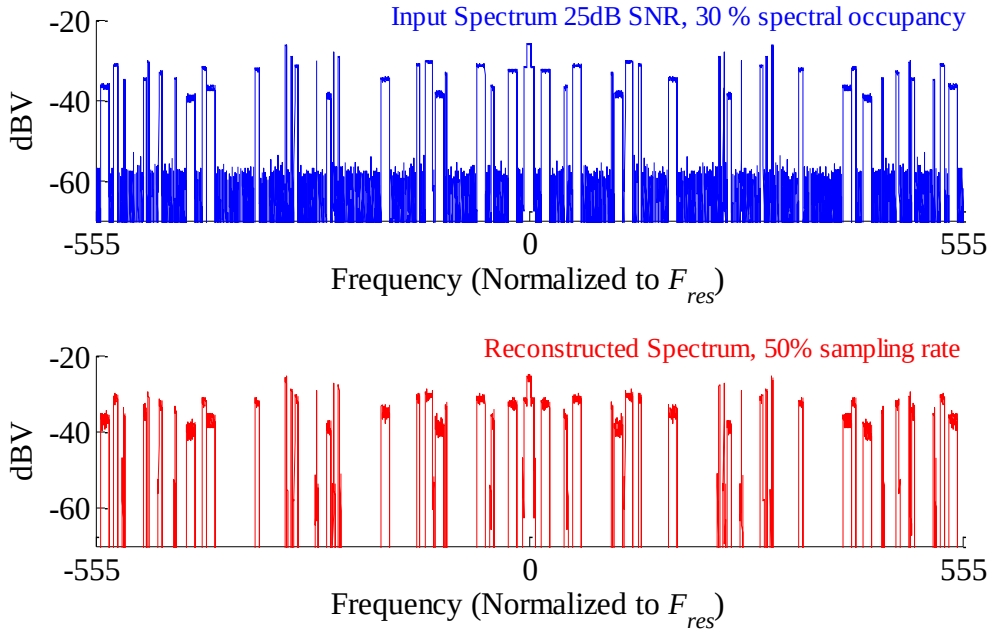

Fig. III-14 (top) An example input spectrum containing variable-bandwidth, variable-amplitude signals with a total spectral occupancy of 30% and SNR of 25dB. (bottom) The reconstructed spectrum by the RFA using L = 11, N = 101, M = 50 effectively sampling at ~50% of the Nyquist rate.

It is important to clarify that, in line with the convention in CS, the reported input SNR is the total SNR of the sparse spectrum, which approaches in-band input SNR at high k/LN.

B. Spectral Support Reconstruction

To emulate a general scenario we plot results for sparse spectra with variable signal bandwidths. Each spectrum realization includes signals of bandwidths in the range ($2F_{res}$, F_S) such that the total spectral occupancy is k bins out of LN. The signals are

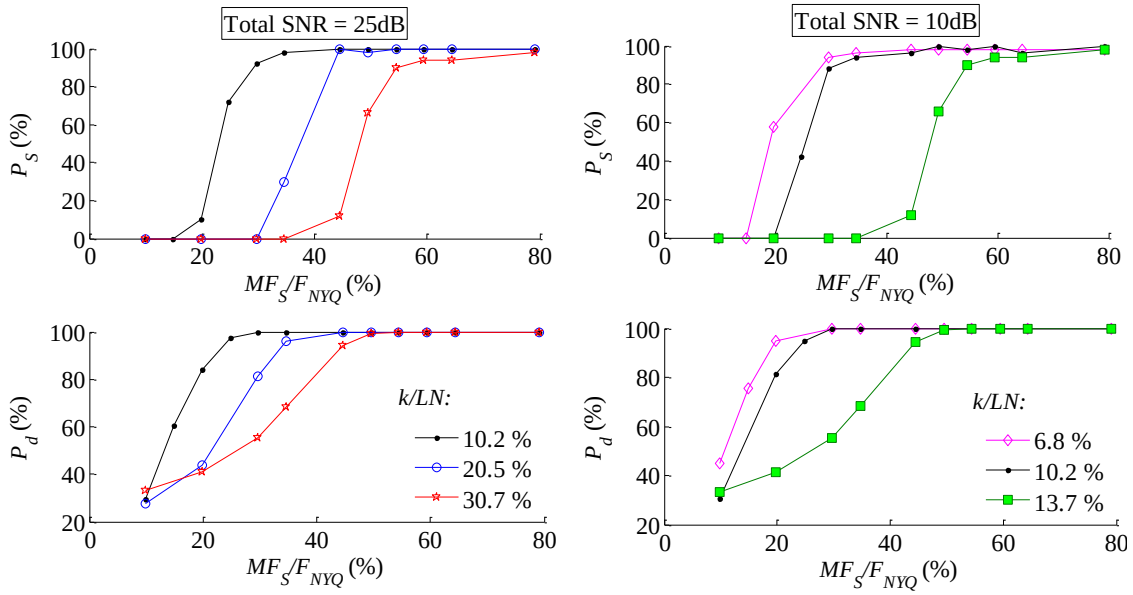

Fig. III-15 (top row) The rate of full support detection P_S for an input SNR of 25 and 10dB and different spectral occupancies. (bottom row) The average probability of detection P_d of a signal bin for the same scenarios. Across the plots, L = 11 and N = 101, with a total LN = 1111 bins in the spectrum.

placed at random carrier locations with random phase shifts and a random gain within a 10dB range.

An example input spectrum with ~30% occupancy (k/LN ~ 0.3) and 25dB of SNR is plotted in Fig. III-14 along with the reconstructed spectrum using ~ 50% sampling rate ($MF_S/F_{NYQ} = M/N$ ~ 0.5). Naturally, when sampling at lower than the Nyquist rate, the original noise in the sparse spectrum is folded and appears distributed over the reconstructed bands [57].

Fig. III-15 provides support reconstruction and signal detection rates versus the ratio of total sampling rate (MF_S) to the wideband Nyquist rate F_{NYQ}. Across the different

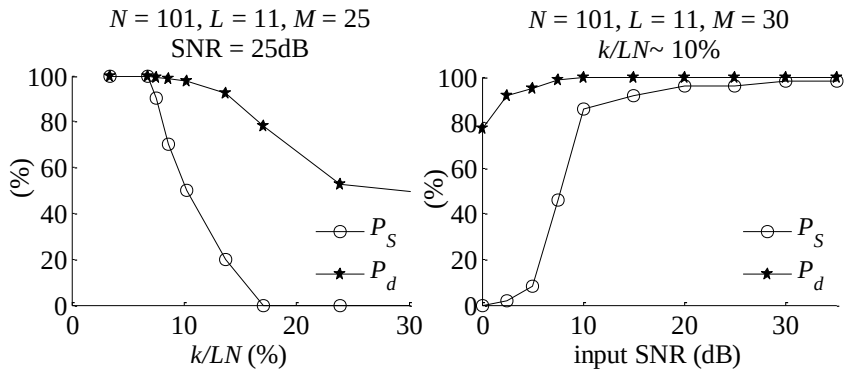

Fig. III-16 The full support and bin detection rates plotted against spectral sparsity (left) and input SNR (right). Sampling rate is M/N of F_{NYQ}.

input SNR and spectral occupancy scenarios, $N = 101$ and $L = 11$ are maintained constant for a resolution $F_{res} = F_{NYQ}/1111$. The top row plots P_S, the average rate at which the full support is correctly detected, i.e. the k detected bins are those which contain signal. The bottom row plots P_d, the average number of occupied signal bins among the k detected bins. P_d can be thought of as the true probability of detection and is a more useful parameter in noisy, high dynamic-range spectra.

Fig. III-16 provides two other perspectives on the support reconstruction performance by fixing the measurement sampling rate and plotting P_S and P_d against the spectral occupancy and input SNR.

C. Signal Reconstruction Performance

Fig. III-17a plots the average reconstructed SNR (RSNR) across different values of L. To investigate the effect of L, the input $v(t)$ is made virtually noiseless and $N = 101$ is maintained constant as L is increased, emulating a higher-resolution sensing of the

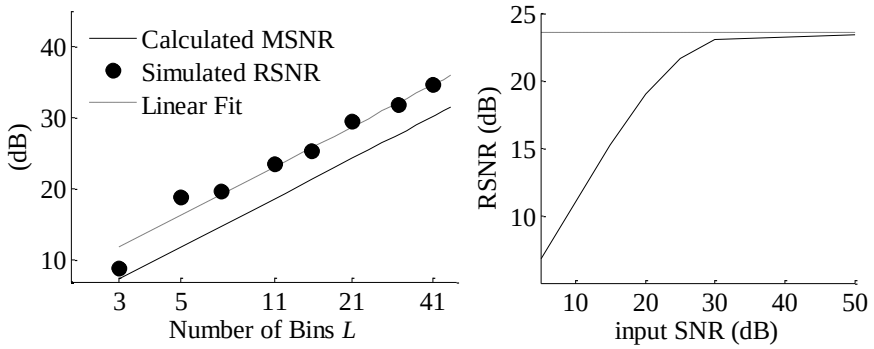

(a) (b)

Fig. III-17 (a) A plot of the simulated RSNR and the calculated MSNR across different values of L and a constant $N = 101$, $M = 50$, and high input SNR. The figure shows the quadratic growth of the MSNR and RSNR with L. (b) A plot of the RSNR versus input SNR showing the negligible effect of the measurement accuracy at low input SNR. For this plot, $L = 11$.

spectrum. The spectral sparsity k/LN is ~10.2% and the total sampling rate MF_S is ~ $0.5F_{NYQ}$. Recall that we described in section III.5.C how the number of bins L can affect the measurement SNR (MSNR) and provided an estimate of the MSNR for high L in (35). Fig. III-17a demonstrates the quadratic growth in RSNR as a function of L and the accurate fit at high L. Note that because we employ sub-optimal sensing matrices and use pseudo-inversion to reconstruct the signal bins, the RSNR/MSNR gain, as shown in Fig. III-17a, is less than that described in [57] for an oracle CS solver with ideal sensing matrices.

The discussion in section III.5.C also made the claim that the measurement accuracy need not be perfect in realistic noisy scenarios. Fig. III-17b demonstrates this by plotting the RSNR as a function of the SNR of the input signal $v(t)$. As shown by the

figure, the effect of the MSNR (dotted line) is negligible at low input SNR. In a realistic noisy sparse spectrum then, the RSNR would be dominated by the input SNR.

III.7 Complexity and Programmability

To complete our discussion of the proposed RFA, we evaluate the complexity and programmability improvements we set out to achieve. The resolution enhancement is inherent in the complexity and programmability.

A. Hardware Complexity Reduction:

Since the RD cannot capture general signals outside the multi-tone model, we evaluate the complexity reduction provided by the RFA by comparing its hardware requirements to those of the MWC.

As we noted in section III.5.A, the single-pole filter $H(f)$ employed by the RFA is implicitly provided at the output of typical analog circuit blocks, including mixers. The RFA thus completely avoids the M dedicated high-order analog filters required by the MWC, reducing the analog hardware cost by $M*C_F$, where C_F is the cost of a high-order analog filter. On the digital end, as is the case in typical multi-band receivers, an L-band decimated filter bank [62] can be utilized to implement the L filters $B(\omega)$ for spectral discretization, incurring approximately the cost of an equivalent single filter $B(\omega)$ at rate F_S per measurement. Assuming that the cost of a digital operation is C_{op} and the filter

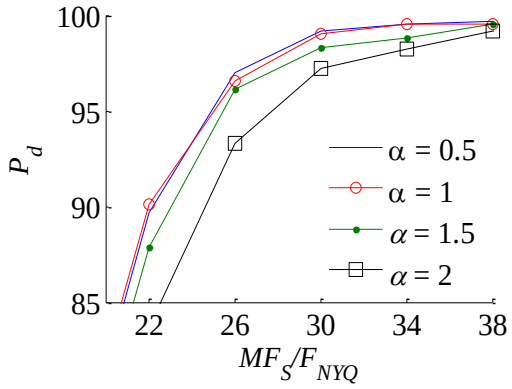

Fig. III-18 Plotting P_d for different values of α shows that the performance is not sensitive to the accuracy in $h(t)$. Figure also shows that $\alpha < 1$ achieves little gain in sensing capabilities, as illustrated by the coherence plot in Fig. III-13.

$B(\omega)$ has W_B taps, the RFA incurs an additional[3] digital filtering cost of $\sim M*W_B*2*C_{op}$. Comparing state-of-the-art digital chips and analog filters, one finds that digital circuitry is significantly more power and area efficient than the analog alternative. In terms of power consumption for example, digital circuit cost C_{op} in measured in nW whereas analog filter cost C_F is measured in mW, making the cost of the digital $B(\omega)$ filters negligible compared to the discarded analog filters. Understanding the implementation tradeoffs then, we have designed the RFA to achieve drastic savings in hardware cost by moving spectral discretization to the digital domain.

Furthermore, as discussed so far, the RFA reduces the single $M_{full}xLN$ CS problem to L decimated MxN CS problems, where $M < M_{full}$. The cost of performing OMP, using

[3] We have also ignored the $M*W_{Eq}*2*C_{op}$ cost of digital equalization (W_{Eq} taps) required by the MWC [51] as well as the increased sampling rate per each channel of the MWC to account for finite-order analog filters.

most efficient QR-decomposition approach, can be estimated as $Nk + Mk + k^2$ operations and $N^2 + NM + Mk + k^2$ memory elements [59]. Assuming that M and k are approximately L times smaller than M_{full} and k_{full} respectively, the RFA thus reduces both the number of operations and memory elements for CS reconstruction by L times. Assuming additionally that the L CS problems are solved at the appropriate L-times decimated rate per problem, the digital circuit cost would accordingly be reduced by L^2.

It is also important to note, within the context of hardware complexity, that the time constant of h(t) need not be accurately designed for $\alpha = 1$. Fig. III-18 plots Pd for different $\alpha \in \{0.5, 2\}$, ~10.2% spectral occupancy and 25dB input SNR. The figure, while demonstrating that α has no destructive effects on the performance, also supports the theoretical results based on the coherence of $\mathbf{H}_l$ in section III.5.B. The figure shows that increasing α beyond unity results in detection degradation while reducing α even as low as 0.5 achieves negligible detection improvement.

B. Programmability and Insensitivity to Spectral Occupancy

Because the RFA performs fine spectral discretization in the digital domain, its resolution and measurement accuracy are digitally controlled. We will also demonstrate that the RFA is able to adapt to an increase in spectral occupancy by varying the total sampling rate while using the same number of measurements. This is possible because, by design, the RFA relies on random analog filters $G_m(f)$ that enforce no constraints on the sampling rate.

As Fig. III-19 (top) shows, adjusting F_S for the same number of measurements and the same desired spectral resolution can provide enhanced support detection

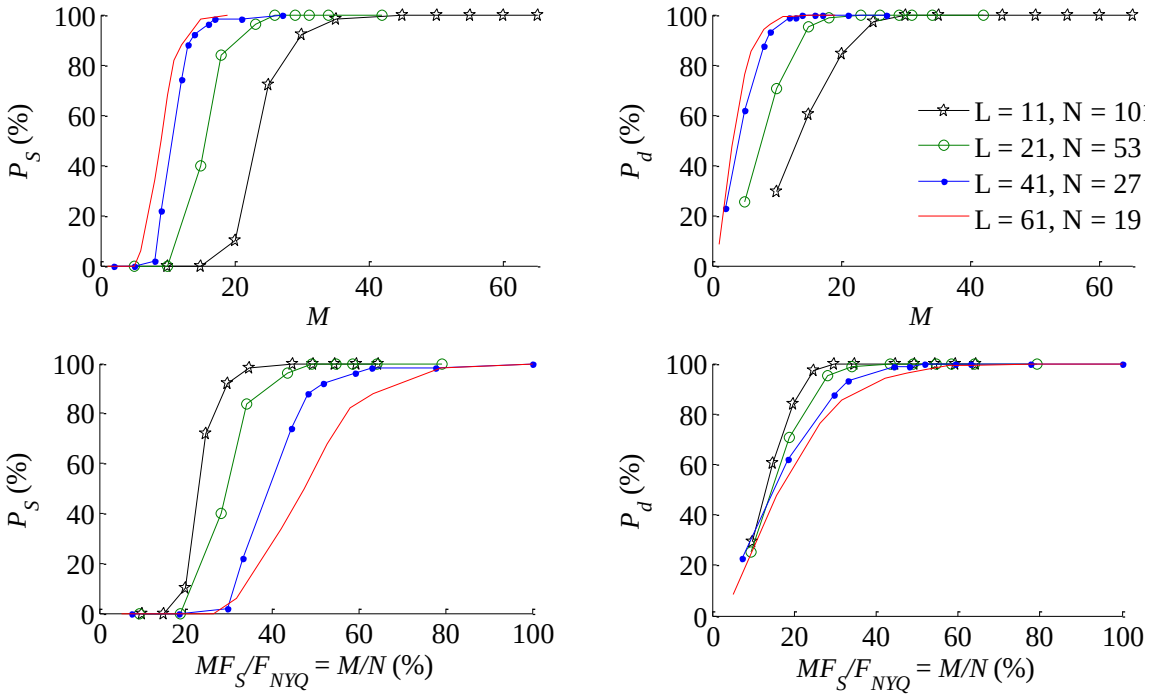

Fig. III-19 (top) Spectral support reconstruction rate P_S and probability of detection P_d versus the number of measurements M for a constant $LN \sim 1111$ and different L. The curves show that the RFA can vary the sampling rate F_S per channel (or equivalently the number of bins) for the same resolution F_{res} and Nyquist bandwidth F_{NYQ}, thus tracking a varying sparsity. (bottom) The same curves plotted against sampling rate for a different perspective. All plots correspond to a spectral occupancy of ~10.22% and SNR of 25dB.

performance. Increasing F_S translates to increasing the number of bins L while reducing the width of the sensing matrix N. It is interesting to note that, as the curves show, increasing F_S (or L) achieves diminishing returns in performance since M scales logarithmically with a decreasing N. For reference, Fig. III-19 (bottom) plots the performance curves for the same scenarios against the total percent sampling rate. To give an example, although $M = 20$ measurements are insufficient to achieve high P_S or P_d

using (L,N) = (11,101), the total sampling rate can be increased from ~20% to ~40% using the same M = 20 measurements to achieve near 100% P_S and P_d at (L,N) = (21,53).

The results in Fig. III-19 also demonstrate that, for a realistic spectrum with variable signal bandwidths such as in Fig. III-14, M does not correspond to the number of measurements M_{full} required to solve (36) using a dense sensing matrix. If the latter were the case, then increasing L would not be able to provide the enhanced performance seen in the figure. Instead, as L is varied, the k occupied bins in the spectrum are re-distributed rather uniformly across the different CS problems allowing MF_S to dictate the performance.

IV. CONCLUSION

This dissertation took a fundamentally novel look at the design of anti-alias filters. We showed that the proposed approach of involving aliasing in the filtering task allows for drastic improvements in the performance and programmability when compared with existing schemes of comparable complexity. We showed how the FA systems can be realized using proven digitally-controlled circuitry and discussed some of the associated implementation sensitivities. The proposed scheme is very well suited for the needs of future software-defined and cognitive communication systems. Beyond the scope of this dissertation, particular applications and implementation techniques that exploit the programmability, scalability, and low-sensitivity of the FA technique should be investigated.

This dissertation also proposed the novel Random Filtering by Aliasing technique for the CS acquisition of spectrally sparse continuous-time signals. By understanding the implementation trade-offs and the requirements on measurement accuracy, we designed the RFA to perform spectral discretization in the programmable digital domain while significantly reducing analog hardware cost and CS problem size. As a result, the RFA achieves better resolution, complexity and programmability than the state-of-the-art approaches without enforcing particular spectral sparsity models.

APPENDIX

A. FA: Double-pole Physical Filter

As discussed in section II.2.D, the FA technique can be applied to higher order filters $h(\tau)$ such that the resulting apparent filter is sharper. We provide here an interesting scenario where a 2nd order double-pole IIR DT filter is used in a DT FA implementation from Fig. II-6. The double-pole DT filter $H(\omega)$ and the resulting apparent filter $G(\omega)$ are given in Table 1. The physical filter is straightforward to implement and requires similar hardware as the single-pole DT filter but places two poles at the opposite ends of the unit circle in the z-plane.

By using the FA technique, one can visualize how, according to (2.9), the

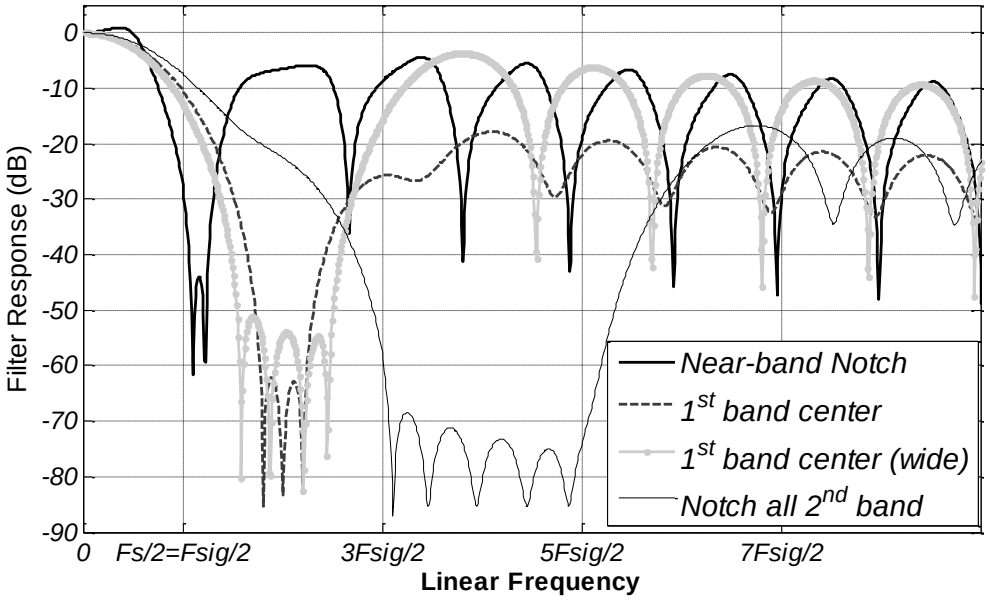

Fig. A-1 Notch anti-alias filtering using the FA DT double-pole technique with traded-off notch width, depth and proximity to the signal of interest.

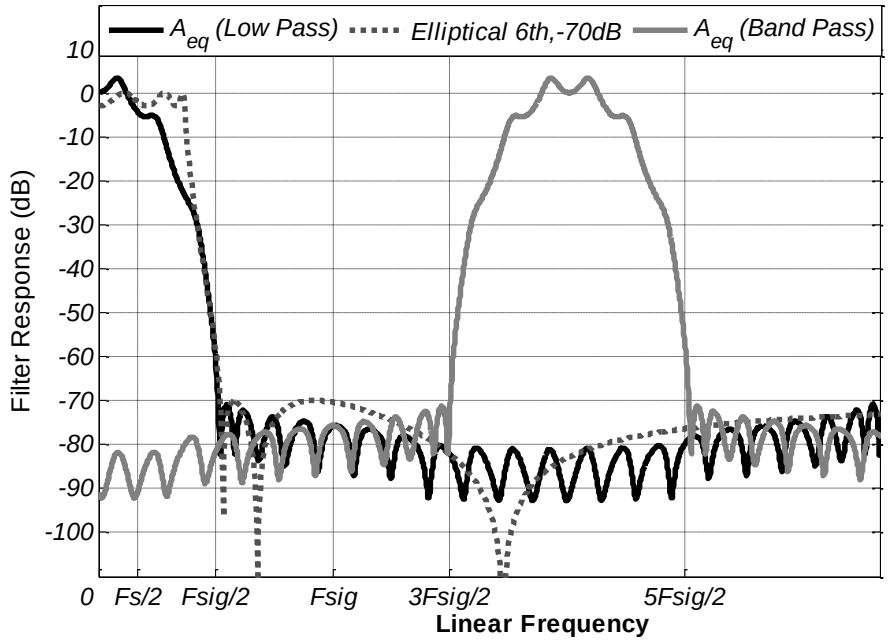

Fig. A-2 Effective Low- and Band-pass filtering using the proposed techniq. A 6th order elliptical filter is shown for reference.

apparent filter can exhibit low-pass or band pass behavior by moving the double poles around. The FA behavior can also be visualized in the time domain where the length K sequence $d[n]$ corresponds to a length $2K$ modulation of the impulse response due to the double delay, thus creating sharper filters.

Fig. A-1 and Fig. A-2 show some possible apparent filters created using the FA DT double-pole system. Note that such filters exhibit image pass-bands at half the switching rate, unlike single-pole implementation that exhibits images at the switching rate. To obtain the same sharper responses with a farther image pass-band, a 2-path single pole FA solution can be employed in the manner described in section II.2.D.

B. FA: Arbitrary Interference Profiles

This section shows a few examples of how an FA system can be used to provide reconfigurable filtering responses according to a given interference scenario.

The results in Fig. A-3 show apparent filters obtained using the double-pole DT FA implementation from Appendix A. The black curves show the optimal responses (given some constraints) for a given spectral content in grey. The optimal responses are chosen so as to maximize the signal to interference ratio (SINR) by describing the stop-band constraints of the filter design problem (2.10). In these results, the problem is solved directly using non-linear Newton Methods [25]. The signal of interest is assumed to be the grey region at frequencies below $B = F_S/2$.

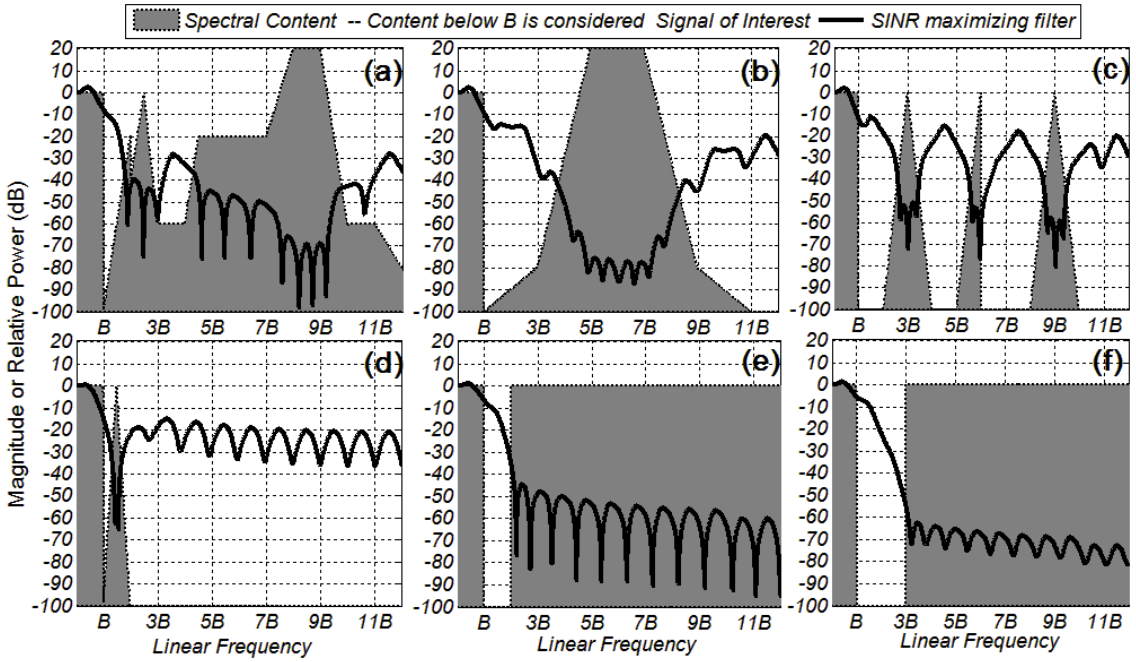

Fig. A-3 Examples showing the design of the FA apparent filter to maximinze the SINR given an interference profile. (a), (b), and (c) are arbitrary, (d) is a notch filter, (e) and (f) are low-pass filters that trade-off suppression with transition width.

C. RFA: Deriving the Measurement Error

We show here the derivation of $e_{m,l}[n]$, the error in measurement m for bin l, as given by (3.33). We start with subtracting $\ddot{u}_{m,l}[n]$ from $u_{m,l}[n]$ in (A1).

$$e_{m,l}[n] = \frac{1}{T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} \int_0^{+\infty} g_m(\tau) e^{j2\pi(lF_{res}+\eta F_S)\tau} \left(s_{\eta,l}(nT_S - \tau) - s_{\eta,l}(nT_S) \right) d\tau \tag{A1}$$

For a useful $G_m(f)$ that is designed to vary slowly within a bandwidth of F_{res}, the length of significant content in $g_m(t)$, it's power-delay profile [57] should be much smaller than T_{res}. Given that $s_{\eta,l}(t)$ is limited to a bandwidth of F_{res}, the difference in (A1) can then be simplified as the negative of the derivative of $s_{\eta,l}(t)$ at nT_S multiplied by τ, as in (A2).

$$e_{m,l}[n] \simeq \frac{1}{T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} \int_0^{+\infty} g_m(\tau) e^{j2\pi(lF_{res}+\eta F_S)\tau} \left(-s'_{\eta,l}(nT_S)\tau \right) d\tau \tag{A2}$$

The integration is now given by the negative of the derivative of $G_m(f)$ at $lF_{res} + \eta F_S$, resulting in (A3) or (3.33).

$$e_{m,l}[n] \simeq \frac{j}{2\pi T_S} \sum_{\eta=-(N-1)/2}^{+(N-1)/2} s'_{\eta,l}(nT_S) G'_m(lF_{res} + \eta F_S) \tag{A3}$$

D. RFA: Deriving the $MSNR_{m,l}$

We show here the derivation of MSNR in (3.32). We assume that α is high enough such that $G_m(f)$ exhibits similar coherence over different indices l as discussed in section II.5.B. Since all measurements are obtained in identical manners from uncorrelated measurement filters, the MSNR is then given by $\text{MSNR}_{m,l}$ of measurement m for bin l, shown in (A4).

$$\text{MSNR} = \frac{E\left[\left(\ddot{u}_{m,l}[n]\right)^2\right]}{E\left[\left(e_{m,l}[n]\right)^2\right]} \tag{A4}$$

It is fair to assume that $s_{\eta,l}(nT_S)$ and $G_m(lF_{res} + \eta F_S)$ are uncorrelated with each other, $s_{\eta,l}(nT_S)$ is uncorrelated over different η, and $G_m(f)$, by design, is incoherent over F_S. We can accordingly write the $\text{MSNR}_{m,l}$ as shown in (A5).

$$\text{MSNR}_{m,l} = 4\pi^2 \frac{E\left[\left(s_{\eta,l}(nT_S)\right)^2\right] E\left[\left(G_m(f)\right)^2\right]}{E\left[\left(s'_{\eta,l}(nT_S)\right)^2\right] E\left[\left(G'_m(f)\right)^2\right]} \tag{A5}$$

Assuming that the energy in $S_{\eta,l}(f)$ is uniformly distributed over $[-F_{res}/2, F_{res}/2]$ with PSD S_o, and that the power-delay profile of $g_m(t)$ exhibits single-pole decay with time constant τ_p (as for a single-pole $h(t)$ and a ± 1 $d_m(t)$), we obtain (A6).

$$\mathrm{MSNR}_{m,l} \simeq \frac{4\pi^2 \int_{-F_{res}/2}^{+F_{res}/2} S_o df}{\int_{-F_{res}/2}^{+F_{res}/2} S_o 4\pi^2 f^2 df} \frac{\int_0^{+\infty} \exp\left(-2t/\tau_p\right) dt}{\int_0^{+\infty} 4\pi^2 t^2 \exp\left(-2t/\tau_p\right) dt} \tag{A6}$$

$$= 4\pi^2 \frac{3}{\pi^2 F_{res}^2} \frac{2}{4\pi^2 \tau_p^2} = \frac{6}{\pi^2 F_{res}^2 \tau_p^2} = \frac{6}{\pi^2} L^2 \alpha^2$$

Note that (A6) is accurate for high L and α, where the assumptions on correlation and coherence hold.

REFERENCES

[1] Donoho, "Compressed sensing," IEEE Trans. Inf. Theory, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[2] J. Tropp, "Signal recovery from random measurements via orthogonal matching pursuit," IEEE Trans. Inf. Theory, vol. 53, no. 12, pp. 4655–4666, Dec. 2007.

[3] J. Mitola, "The software radio architecture," IEEE Communications Mag., vol. 33, no. 5, pp. 26-38, May 1995.

[4] R. Bagheri et al., "An 800-MHz–6-GHz software-defined wireless receiver in 90-nm CMOS," IEEE J. Solid-State Circuits, vol. 41, no. 12, pp. 2860–2876, Dec. 2006.

[5] A. Baschirotto, R. Castello, F. Campi, G. Cesura, M. Toma, R. Guerrieri, A. Lodi, L. Lavagno, and P. Malcovati, "Baseband analog front-end and digital back-end for reconfigurable multi-standard terminals," IEEE Circuits Syst. Mag., vol. 6, no. 1, pp. 8–28, 1st Quarter, 2006.

[6] J. Mitola and G. Q. Maguire, "Cognitive radio: making software radios more personal," *IEEE Personal Communications*, vol. 6, pp. 13-18, Aug. 1999.

[7] A. A. Abidi, "The path to the software-defined radio receiver," *IEEE J. Solid-State Circuits*, vol. 42, no. 5, pp. 954–966, May 2007.

[8] R. Bagheri *et al.*, "An 800-MHz–6-GHz software-defined wireless receiver in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2860–2876, Dec. 2006.

[9] V. Giannini *et al.*, "A 2-mm 0.1–5 GHz software-defined radio receiver in 45-nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3486–3498, Dec. 2009.

[10] Z. Ru, N. A. Moseley, E. Klumperink, and B. Nauta, "Digitally-enhanced software-defined radio receiver robust to out-of-band interference," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3359–3375, Dec. 2009.

[11] R. Bagheri, A. Mirzaei, S. Chehrazi, and A. Abidi, "Architecture and clock programmable baseband of an 800 MHz-6 GHz software-defined wireless receiver," *Inter. Conf. VLSI Design*, vol., no., pp.135-140, Jan. 2007.

[12] A. Mirzaei, S. Chehrazi, R. Bagheri and A. Abidi, "Analysis of first-order anti-aliasing integration sampler," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 10, pp.2994 -3005 2008.

[13] Y. Poberezhskiy and G. Poberezhskiy, "Sampling and signal reconstruction structures performing internal antialiasing filtering and their influence on the design of digital receivers and transmitters," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 1, pp.118-129, January 2004.

[14] M. Oskooei, N. Masoumi, M. Kamarei, H. Sjoland, "A CMOS 4.35-mW +22-dBm IIP3 continuously tunable channel select filter for WLAN/WiMAX receivers," *IEEE J. Solid-State Circuits*, vol.46, no.6, pp.1382-1391, June 2011.

[15] J. Yuan, "A charge sampling mixer with embedded filter function for wireless applications," *Proc. 2nd Inter. Conf. Microwave and Millimeter Wave Technology*, pp.315-318, 2000.

[16] A. Mirzaei, R. Bagheri, S. Chehrazi, and A. A. Abidi, "A second-order anti-aliasing prefilter for an SDR receiver," *Proc. Custom Integrated Circuits Conf.*, pp. 629-632, September 2005.

[17] Y. Poberezhskiy and G. Poberezhskiy, “Sampling technique allowing exclusion of antialiasing filter,” *Electronics Letters*, vol. 36, no. 4, pp. 297–298, Feb. 2000.

[18] Y. Poberezhskiy and G. Poberezhskiy, “Some aspects of the design of software defined receivers based on sampling with internal filtering,” *Proc. IEEE Aerospace Conf.*, pp. 1-20, Mar. 2009.

[19] D. Murphy *et al.*, “A blocker-tolerant wideband noise-cancelling receiver with a 2dB noise figure,” *Int. Solid-State Circuits Conf. (ISSCC)*, pp.74-75, Feb. 2011

[20] A. Ghaffari *et al.*, "Tunable High-Q N-Path Band-Pass Filters: Modeling and Verification," *Solid-State Circuits, IEEE Journal of* , vol.46, no.5, pp.998-1010, May 2011

[21] L. E. Franks and I. W. Sandberg, “An alternative approach to the realization of network transfer functions: The N-path filters,” *Bell Sys. Tech. J.*, vol. 39, pp. 1321–1350, Sep. 1960.

[22] E. O'hAnnaidh, E. Rouat, S. Verhaeren, S. Le Tual, and C. Garnier, “A 3.2GHz-sample-rate 800mHz bandwidth highly reconfigurable analog FIR filter in 45nm CMOS,” *Int. Solid-State Circuits Conf. (ISSCC)*, pp.90-91, Feb. 2010.

[23] A. Abidi, “Direct-conversion radio transceivers for digital communications,” *IEEE J. Solid-State Circuits*, vol.30, no.12, pp.1399-1410, Dec 1995.

[24] I. Vassiliou, K. Vavelidis, N. Haralabidis, A. Kyranas, Y. Kokolakis, S. Bouras, G. Kamoulakos, C. Kapnistis, S. Kavadias, N. Kanakaris, E. Metaxakis, C. Kokozidis, and H. Peyravi, “A 65 nm CMOS multistandard, multiband TV tuner for mobile and multimedia applications,” *IEEE J. Solid-State Circuits*, vol. 43, no. 7, pp. 1522–1533, Jul. 2008.

[25] R. Fletcher, Practical methods of optimization, 2nd ed., New York: John Wiley & Sons, 1987, p.113.

[26] M. Grant and S. Boyd, CVX: Matlab software for disciplined convex programming [Online], version 1.21, Available: cvxr.com/cvx , April 2011.

[27] M. Grant and S. Boyd, Graph implementations for nonsmooth convex programs, *Recent Advances in Learning and Control (a tribute to M. Vidyasagar)*, V. Blondel, S. Boyd, and H. Kimura, editors, pages 95-110, *Lecture Notes in Control and Information Sciences*, Springer, 2008.

[28] S. Wu, S. Boyd, and L. Vandenberghe, "FIR filter design via spectral factorization and convex optimization," [Online], Available: http:// www.ee.ucla.edu/~vandenbe/publications/magdes.pdf

[29] L. Roberts, "Picture coding using pseudo-random noise," *IRE Trans. Info. Theory*,, vol.8, no.2, pp.145-154, Feb. 1962.

[30] S. Y. Chen, N. S. Kim, J. Rabaey, "A 10b 600MS/s multi-mode CMOS DAC for multiple Nyquist zone operation," *2011 Symposium on VLSI Circuits (VLSIC)*, vol., no., pp.66-67, June 2011.

[31] W. T. Lin, T. H. Kuo, "A Compact Dynamic-Performance-Improved Current-Steering DAC With Random Rotation-Based Binary-Weighted Selection," *IEEE Journal of Solid-State Circuits*, vol.47, no.2, pp.444-453, Feb. 2012.

[32] A. Jiraseree-Amornkun *et al*., "Theoretical analysis of highly linear tunable filters using switched-resistor techniques," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 12, pp. 3641–3654, Dec. 2008.

[33] S. Xiao, J. Silva, U. Moon, and G. Temes, "A tunable duty-cycle-controlled switched-R-MOSFET-C CMOS filter for low-voltage and highlinearity applications," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2004,pp. 443–436.

[34] H. Jin and E. Lee, "A 350MHz Programmable Analog FIR Filter Using Mixed-Signal Multiplier", *ESSCIRC*, pp. 152-155, Sept., 2000.

[35] G. Xu and J. Yuan, "A CMOS Analog FIR Filter with Low Phase Distortion," *ESSCIRC*, pp. 747-750, Sept., 2002.

[36] Z. Ru, E. Klumperink, and B. Nauta, "Discrete time mixing receiver architecture for RF -sampling software-defined radio," in *IEEE J. Solid- State Circuits*, vol. 45, no. 9, pp. 1732–1745, Sept. 2010.

[37] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb.2006.

[38] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. Part II: Convex relaxation," *Signal Process.*, vol. 86, no. 3, pp. 589–602, 2006.

[39] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. Part I: Greedy pursuit," *Signal Process.*, vol. 86, no. 3, pp. 572–588, 2006.

[40] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Const. Approx., 2007.*

[41] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 2, pp. 201–220, Feb. 2005.

[42] S. Kirolos, J. Laska, M. Wakin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. Baraniuk, "Analog to information conversion using random demodulation," in *Proc. 2006 IEEE Dallas/CAS Workshop on Design, Appl., Integr. Software*, Oct. 2006, pp. 71–74.

[43] J. Laska, S. Kirolos, M. Duarte, T. Ragheb, R. Baraniuk, and Y. Massoud, “Theory and implementation of an analog-to-information converter using random demodulation,” in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2007, pp. 1959–1962.

[44] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk, “Beyond Nyquist: Efficient sampling of sparse bandlimited signals,” *IEEE Trans. Inf. Theory*, vol. 56, no. 1, pp. 520–544, 2010.

[45] M. A. Lexa, M. E. Davies, and J. S. Thompson, “Reconciling compressive sampling systems for spectrally sparse continuous-time signals,” *IEEE Trans. Sig. Process.*, vol. 60, no. 1, pp. 155-171, 2012.

[46] Z. Yu, S. Hoyos, and B. M. Sadler, “Mixed-signal parallel compressed sensing and reception for cognitive radio,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Apr. 2008, pp. 3861–3864.

[47] X. Chen, Z. Yu; S. Hoyos, B. M. Sadler, J. Silva-Martinez, "A Sub-Nyquist Rate Sampling Receiver Exploiting Compressive Sensing," *IEEE Trans. Circuits and Sys. I*, vol.58, no.3, pp.507-520, March 2011.

[48] 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, 2010.

[49] M. Mishali, Y. C. Eldar, and A. Elron, Xampling: Signal Acquisition and Processing in Union of Subspaces EE Dept., Technion-Israel Inst. Technol., CCIT #747 Oct-09, Oct. 2010 [Online].Available: arXiv 0911.0519, EE Pub No. 1704

[50] M. Mishali, Y. C. Eldar, O. Dounaevsky, and E. Shoshan, "Sub-Nyquist acquisition hardware for wideband communication," *2010 IEEE Workshop on Signal Process. Sys. (SIPS)*, pp.156-161, 6-8 Oct. 2010.

[51] Y. Chen, M. Mishali, Y. C. Eldar, and A. O. Hero, "Modulated wideband converter with non-ideal lowpass filters," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Mar. 2010, pp. 3630–3633.

[52] J. A. Tropp, M. B.Wakin, M. F. Duarte, D. Baron, and R. G. Baranuik, "Random filters for compressive sampling and reconstruction," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2006, pp. 872–875.

[53] Y. Chi, A. Pezeshki, L. Scharf, and R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, Mar. 2010, pp. 3930–3933.

[54] M. Duarte and R. Baraniuk, Spectral Compressive Sensing [Online]. Available: dsp.rice.edu/cs Preprint 2010.

[55] M. Mishali and Y. C. Eldar, "Reduce and boost: Recovering arbitrary sets of jointly sparse vectors," *IEEE Trans. Signal Process.*, vol. 56, no.10, pp. 4652–4702, Oct. 2008.

[56] John G. Proakis, Digital Communications. McGraw-Hill Book Co., 1995.

[57] A. J. Goldsmith, Wireless Communications. New York: Cambridge University Press, 2005.

[58] M. Davenport, J. Laska, J. Treichler, and R. Baraniuk, "The pros and cons of compressive sensing for wideband signal acquisition: Noise folding vs. dynamic range," Preprint, 2011.

[59] E. J. Candes, Y. C. Eldar, and D. Needell. Compressed sensing with coherent and redundant dictionaries. Technical report, 2010. Preprint available at http://arxiv.org/abs/1005.2613.

[60] B. Sturm and M. Christensen, "Comparison of orthogonal matching pursuit implementations," *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, pp.220-224, Aug. 2012

[61] M. Mishali and Y. C. Eldar, A sub-Nyquist system for sparse wideband analog signals [Online], Available: www.technion.ac.il/~moshiko/software.html

[62] P. P. Vaidyanathan, Multirate Systems and Filter Banks, Pearson Education India, 1993.