

Nanoseconds for the Masses

by

Eric VanWyk

B.S. Electrical and Computer Engineering

Franklin W. Olin College of Engineering, 2007

SUBMITTED TO THE DEPARTMENT OF MEDIA ARTS AND SCIENCES IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS OF THE DEGREE OF

MASTER OF SCIENCE IN MEDIA ARTS AND SCIENCES
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2017

©2017 Massachusetts Institute of Technology. All rights reserved.

The author hereby grants MIT permission to reproduce and to distribute publicly paper
and electronic copies of this thesis document in whole or in part in any medium now
known or hereafter created.

Signature of the Author: _____

Department of Media Arts and Sciences

May 24, 2017

Certified By: _____

Neil Gershenfeld

Professor of Media Arts and Sciences

MIT Center for Bits and Atoms

Accepted By: _____

Pattie Maes

Academic Head

Program in Media Arts and Sciences

Nanoseconds for the Masses

by
Eric VanWyk

SUBMITTED TO THE DEPARTMENT OF MEDIA ARTS AND SCIENCES IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS OF THE DEGREE OF

MASTER OF SCIENCE IN MEDIA ARTS AND SCIENCES

Abstract

High precision voltage measurement has become a hidden part of our daily lives. Our phones, our wearables, and even our household appliances now include precision measurement capability that rivals what was once only available in laboratory grade test equipment. Converters with 6 digits of resolution and nanovolt noise floors cost less than a dollar and fit into our watches.

In contrast, measurement of fast phenomena remains out of our daily reach, as it requires equipment too expensive and too unwieldy to be found outside the hands of specialists. Commoditization of sub-nanosecond measurement would improve our ability to process the information from spectral sensors, which in turn would impact portable medical diagnostics, environmental monitoring, and the healthy maintenance of the infrastructure we rely on.

Here a novel measurement architecture is presented that enables cost effective measurement of these sub-nanosecond phenomena, and is easily integrated into existing digital processes. It is built on the same founding premises that the sigma delta architecture uses to dominate low cost precision measurement:

- 1) Precise measurement with imprecise components
- 2) Digital logic replacements for analog components
- 3) Trade time for accuracy

A prototype unit constructed from existing digital communication components is shown to achieve 11 equivalent bits of resolution at 3GHz of analog bandwidth, with repeatability better than 1 millivolt and 3 picoseconds. Timing uncertainty is shown to be better than 1 picosecond.

Several use cases are presented: Differential dielectric spectroscopy, LIDAR, and USB 3 SuperSpeed channel sounding.

Thesis Supervisor: Neil Gershenfeld
Title: Professor of Media Arts and Sciences

Nanoseconds for the Masses

by
Eric VanWyk

The following people served as readers for this thesis:

Thesis Reader: _____
Shuguang Zhang
Principal Research Scientist
MIT Center for Bits and Atoms

Thesis Reader: _____
Fadel Adib
Assistant Professor of Media Arts and Sciences
MIT Media Lab

Acknowledgements

Neil - Thank you for making me question each and every assumption, twice. Thank you for continuously forcing me out of my comfort zone. I've gained so many new perspectives.

Shuguang and Dorrie - Thank you for welcoming me into your ever expanding web of friends.

Fadel - Thank you for your insights and inputs.

Alex – It's been great working with you. Thanks for all of your designs and experiments.

Jamie, Joe, Ryan, Tom, and John - You are the best kind of enablers. Thank you for keeping CBA running smoothly.

Yoko, Bonita, Eddy and Roomba - Thank you for keeping CBA running sanely.

Sam, Will, Amanda, Ben, Prashant, Grace, and Nadya – This has been quite the adventure; thank you for sharing it with me.

Noah, Lisa, Pranam, and Thras – Thank you for teaching me your magic.

Chris, Steve, and Snowflake – Thanks for being the best EE backup team beer can buy.

Mel, Kristen, Ben, and Eric – Thank you for keeping the bubble open for over a decade.

Jonathan – Thank you for the best part of 2016. Welcome!

Jenny, Emily, Mom, and Dad – Thank You.

Table of Contents

ABSTRACT	2
ACKNOWLEDGEMENTS	4
1 INTRODUCTION	8
1.1 APPLICATIONS AND NEEDS	8
1.1.1 <i>Food Quality and Safety: Dielectric Spectroscopy</i>	8
1.1.2 <i>Medical Imaging: Electrical Impedance Tomography</i>	9
1.1.3 <i>Power Efficiency: High Frequency Switching</i>	9
1.2 EFFECTS OF FREQUENCY ON MEASUREMENT.....	10
1.2.1 <i>Taking Advantage of Prior Knowledge</i>	10
1.2.2 <i>Sample and Hold Limits to Frequency Range</i>	10
1.3 ENABLING TECHNOLOGIES	10
1.3.1 <i>High Precision Digital Frequency Synthesis</i>	11
1.3.2 <i>Jitter Reduction and Low Phase Noise Oscillators</i>	11
1.3.3 <i>High Bandwidth Comparators</i>	11
1.3.4 <i>High Speed PHY</i>	11
2 SYSTEM DESIGN AND THEORETICAL ANALYSIS	12
2.1 TIME GENERATION ARCHITECTURE AND ANALYSIS	13
2.1.1 <i>Time Dilation from slow Phase Drift</i>	13
2.1.2 <i>Dual Numerically Controlled Oscillators</i>	13
2.1.3 <i>Alternative Generation Methods</i>	15
2.1.4 <i>Effects of Jitter</i>	15
2.2 VERTICAL RECONSTRUCTION AND ANALYSIS.....	17
2.2.1 <i>Architecture</i>	17
2.2.2 <i>Relation to sigma delta and delta converters</i>	17
2.2.3 <i>Ripple</i>	19

2.2.4	<i>Signal to Noise</i>	19
2.2.5	<i>Slope Overload and Slew Rate vs Granularity</i>	20
2.2.6	<i>Secondary Filtering and Oversampling</i>	21
2.2.7	<i>Idle Tones and Limit Cycles</i>	22
2.3	DIGITIZATION	24
2.3.1	<i>Digital Interface Requirements</i>	24
2.3.2	<i>Analog to Digital Converter Selection</i>	24
3	IMPLEMENTATION AND PERFORMANCE TESTING	25
3.1	SAMPLER CONSTRUCTION AND LAYOUT.....	25
3.1.1	<i>Sampling Comparator</i>	26
3.1.2	<i>Secondary Filter</i>	28
3.1.3	<i>Frequency Requirements</i>	28
3.1.4	<i>Digitizer</i>	29
3.2	STIMULUS GENERATION	29
3.2.1	<i>Digital Outputs</i>	29
3.2.2	<i>SRD / PIN comb generator</i>	30
3.3	PERFORMANCE TESTING	32
3.3.1	<i>Repeatability</i>	32
3.3.2	<i>Traditional Jitter Measurement</i>	34
3.3.3	<i>New Jitter Measurement</i>	34
3.3.4	<i>RF Device Characterization</i>	36
4	APPLICATIONS	37
4.1	DIELECTRIC SPECTROSCOPY AT HOME: EPOXY	37
4.1.1	<i>Results</i>	37
4.2	LIDAR	39

4.2.1	<i>Method and Setup</i>	39
4.2.2	<i>Results</i>	40
4.3	FUTURE WORK: USB SUPERSPEED™	40
4.3.1	<i>Channel Sounding and Training</i>	41
4.3.2	<i>Sounding Configuration</i>	41
4.3.3	<i>Use as a Stand-Alone Sensor</i>	43
5	CONCLUSION AND FUTURE WORK	44
6	TERMS	45
7	BIBLIOGRAPHY	46

1 Introduction

Analog to digital converters are inside most digital devices that respond in some way to the physical world. They give our processors access to sensors and user interfaces, and thereby occupy a hidden but crucial role in our daily lives. High precision measurement has advanced to the point that what used to be laboratory grade test equipment is now found inside smart watches that automatically tweet about our workouts with nanovolt precision.

However, high speed measurement has not advanced in accessibility in the same way. Accessible fast transient measurement could impact our in-home medical diagnostics, our understanding of the environment, and the healthy maintenance of the infrastructure we rely on. Unfortunately, it is still limited to equipment too expensive and unwieldy to be found outside of specialists' hands.

To make these measurements accessible, this thesis presents a novel combination of synchronous equivalent time undersampling and delta oversampling techniques. Two precisely controlled oscillators drift across each other to produce an extremely high equivalent time sample resolution in excess of 10^{18} Hz: An Exahertz. In lieu of a standard sampler, a modified delta converter oversamples many millions of these binary sample events to reconstruct the input waveform: 10^{18} Hz worth of events oversample down to 10^{12} Hz worth of virtual samples, which in turn achieves $\approx 10^{10}$ Hz of equivalent analog bandwidth. The result is the ability to slowly trace out fast, sub-nanosecond transients with higher precision than any individual component in the system. The approach is given the name "Exasampling".

Exasampling aims to do for nanoseconds what sigma delta did for nanovolts.

1.1 Applications and Needs

The Exasampling measurement approach is optimized for low cost, portable, and embeddable measurement of sub-nanosecond phenomena. Although any individual sub-nanosecond event is millions of times faster than human senses are capable of perceiving, in aggregate they direct and inform the macroscopic world.

In general, increased measurement bandwidth and measurement frequency has the following benefits:

- 1) Enhanced separation of independent phenomena by their spectral signature
- 2) Perception of physically smaller effects
- 3) Avoidance of low frequency interference sources

Three specific applications of these principles follow:

1.1.1 Food Quality and Safety: Dielectric Spectroscopy

Dielectric spectroscopy is a fast, non-destructive method of electrically measuring an object's material properties. Since its invention in 1901, it has been used in laboratories to measure unknown properties, and in industrial settings to classify known materials. However, the equipment costs are still prohibitive for home use.

The food industry uses dielectric spectroscopy to detect contaminants and pests. For example, bottles labeled “olive oil” or “cooking oil” sometimes contain cheaper filler materials, such as other oils, pig lard or even “gutter oil” [1]. Recent work combining dielectric methods with modern data analysis techniques has shown success in detecting these adulterations [2], but the methodology relies on the Agilent (Keysight) N5230A PNA Network Analyzer, a \$65k piece of equipment that weighs 25 kg. The electric signature of rice weevils between one and ten gigahertz can be used to detect their contamination of grain [3], and the ripeness of fruits and vegetables can be measured below 1.8 GHz [4]. The equipment used in these two studies is \$40k and larger than a microwave oven; these approaches are not yet ready to protect the home consumer from gutter oil and rice weevils.

The limiting factor for the deployment of these sensors to the home is the cost and size of sub-nanosecond measurement.

1.1.2 Medical Imaging: Electrical Impedance Tomography

The electrical properties of the human body vary by tissue type, blood perfusion, and fluid content. For example, cancerous masses tend to be more conductive than the surrounding tissue, whereas necrotic tissue is less conductive. This correlation has been successfully used to electrically classify types of skin cancers [5].

Electrical Impedance Tomography (EIT) forms electrical images of the tissue through the use of arrays of electrodes. Its originally use was to monitor localized lung function and perfusion [6]. Multi-frequency EIT, or EIS, is a candidate to replace mammograms in screening for and localizing breast cancer [7], and is being evaluated for non-invasive transrectal prostate screenings [8]. These methods have the potential to give surgeons the ability to detect and localize tissue boundaries during the surgery without the delay of sending out a biopsy.

EIS’s imaging resolution is a function of the number of electrodes, and its ability to differentiate tissue is limited by its electrical measurement bandwidth and resolution. The next generation of EIS based surgical and diagnostic tools will require higher bandwidth and channel count.

1.1.3 Power Efficiency: High Frequency Switching

Switched Mode Power Supplies (SMPS) efficiently convert and control electrical power by rapidly pulsing power transistors. Timing parameters are critical to the efficiency and safe operation of SMPS, but the optimal tuning changes with time, temperature, and loading conditions. They must be chosen conservatively when information is unavailable, sacrificing potential efficiency. [9] Online measurement and continuous optimization can save power and increase controlled power density. [10]

Recent developments in GaN power transistors are pushing SMPS switch times sub-nanosecond and switching frequencies above 25MHz [11]. This allows these GaN stages to have greatly improved power densities, but at the cost of making it more difficult to collect real-time data on their operation. The transients have severe frequency content at 500 MHz, and EMI out past 1 GHz [12], and therefore currently require a high end oscilloscope for proper characterization.

Integrating the proposed approach into an SMPS system would allow for online self-monitoring of the relevant waveforms. This has the potential to increase system efficiency, decrease radiated emissions, and increase system longevity.

1.2 Effects of Frequency on Measurement

The frequency content of a signal has a large impact on the measurement architecture used. To break this down, there are three commonly conflated frequency metrics for sampling measurements: the frequency range, the sampling rate, and the information rate. The Nyquist-Shannon sampling theorem states that if the sampling rate is at least twice the frequency range (bandwidth), the signal can be measured regardless of the underlying information content [13].

1.2.1 Taking Advantage of Prior Knowledge

A pure sinewave with a one nanosecond period (1 GHz) would require at least 2 gigasamples per second to reconstruct with a Nyquist sampler. However, continuing measurements do not provide novel information if the nature of the signal source is known. Super-Nyquist measurement techniques use prior knowledge in order to be able to measure signals with bandwidths above their sample rate's Nyquist limit.

The examples given in Section 1.1 all have transients that are faster than a nanosecond, but the information contained in those signals change slowly. Nyquist measurement is an inefficient option, because it would require a sample rate much higher than the information rate. The stability of the signals over time provides the prior knowledge necessary for super-Nyquist measurement. Of the existing super-Nyquist techniques, Exasampling is most closely related to Equivalent Time sampling, or synchronous undersampling. The stable signal is driven to be repetitive, which provides the prior knowledge necessary to successfully measure frequency content above the Nyquist limit.

1.2.2 Sample and Hold Limits to Frequency Range

The maximum frequency that an architecture can reliably measure is also limited by its front end circuitry. A common frequency limiting component is the Sample and Hold Amplifier (SHA). It samples an incoming waveform and holds its value in order to provide the converter a stable voltage to work with. Interleaving many ADC/SHA pairs can increase the sample rate of the aggregate converter, but cannot extend the maximum sample-able frequency. [14] These SHA carry the performance burden of the entire converter; the system input bandwidth and accuracy are a direct function of the performance of an individual SHA. [15] [16]

The root of the direct correlation is that a SHA does not assume any prior knowledge about the incoming waveform, and is assumed to be able to accurately sample even a random signal. Exasampling relies on a given sample being very similar in value to the sample immediately before it. This assumption allows the system an extra degree of freedom with which to optimize performance.

1.3 Enabling Technologies

Exasampling relies on several key enabling technologies from the high speed digital communication industry. The immense market pressures on wired and wireless communication

have commoditized frequency synthesis and high frequency digital. These have not only made this architecture possible, but have also made it very cost effective.

1.3.1 High Precision Digital Frequency Synthesis

Frequency synthesis is at the heart of radio communication. Multi-channel radio systems rely on digital frequency synthesis techniques to create independent communication channels. Improving the resolution of the synthesizer increases the number of channels available, and increasing the accuracy decreases the amount of bandwidth that is wasted between channels.

These requirements create market pressure from the telecommunications, networking, and physical interconnect space, and have led to the creation of a new type of synthesizer necessary for the proposed system. So-called “rational” synthesizers are capable of frequency precision better than one part in ten billion [17]. This jump allows a synthesizer to compensate for accuracy with resolution.

Dual synthesizers are further capable of generating closely spaced frequencies that maintain a stable phase drift relationship and a digitally stable frequency ratio. This allows for creation of beat frequencies tighter than their individual raw accuracy allows, which can be several orders of magnitude higher than previous approaches.

1.3.2 Jitter Reduction and Low Phase Noise Oscillators

High speed wired communication relies heavily on low jitter, low phase noise clock sources, as timing uncertainty directly contribute to error rate in these systems. Jitter-blockers are devices that are able to produce a low jitter clock source using a noisier reference as an input. These can now easily achieve picosecond jitter floors, with some devices capable of achieving tens of femtoseconds worth of RMS jitter.

1.3.3 High Bandwidth Comparators

Ultra-high speed latched comparators now have several gigahertz of analog input bandwidth, and specialized comparators now reach as high as 26GHz. These new comparators favorably compete with track-and-hold amplifiers.

1.3.4 High Speed PHY

The transmitters used in high speed digital physical layer components (PHY) now have sufficiently fast rise and fall times to generate several gigahertz of frequency content. For example, a modern LVPECL output can readily achieve a 35 picosecond rise that contains 10GHz of frequency content. These PHYs are often integrated into FPGAs and network controllers.

2 System Design and Theoretical Analysis

The Exasampling architecture creates a continuous time reproduction of a high speed transient signal that has been slowed down sufficiently for a slower analog to digital converter to digitize. A sample and hold amplifier presents discrete samples to its converter; this method presents a continuous waveform instead. Exasampling is a time dilating analog to analog converter, and can be thought of as a cross between a vector network analyzer, a sample and hold amplifier, and a delta converter.

Time base generation is so crucial to the Exasampling architecture that it draws its name from the equivalent time sample rate that it uses internally. One Exahertz is a quintillion, or 10^{18} , cycles per second. This is roughly one million times faster than the current world record for transistor switching speed [18]. If it were realized as a real time system, the clock would radiate as a hard X-Ray laser, making the direct approach ill-advised for the home consumer market. A new take on a classic undersampling strategy described in Section 2.1 creates the time base necessary to achieve this rate with existing, low cost components.

The signal is reconstructed by a novel take on the classic delta converter. Delta and sigma delta converters achieve high precision at low cost, but are bandwidth limited. The modifications described in Section 2.2 use the extreme undersampling to bring the benefits of these slow, precise converters to much higher frequencies. A primary advantage is that the loop has a single comparator that operates at the input signal bandwidth, with all other components operating at much lower frequencies. Decoupling in this manner allows higher input bandwidths for a given process technology than an equivalent SHA based approach.

The high-level system architecture is shown below in Figure 1. In this figure, the frequency response of a device under test (DUT) is being measured. The configuration shown uses two reconstruction channels and a single stimulus generator; other use cases may use different arrangements.

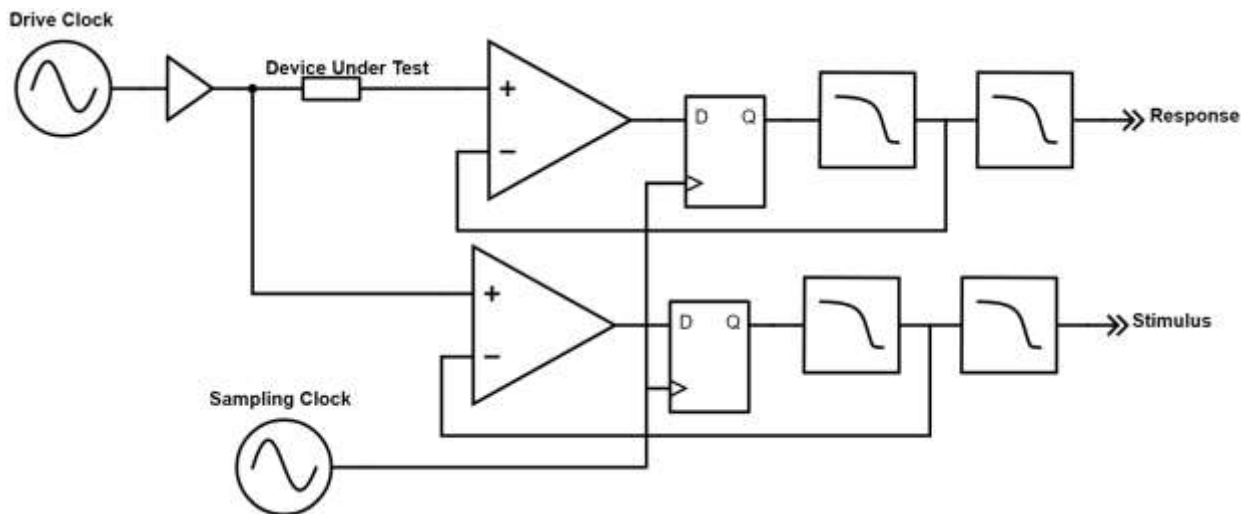


Figure 1: System Architecture

The drive generator produces a repetitive stimulus waveform synchronous to the drive clock (left) from the time base generator. Parameters relevant to the selection of the drive generator are presented in Section 3.2. The DUT's response waveform enters the sampling comparator's input and is reconstructed by the top channel. The stimulus waveform is similarly reconstructed by the bottom channel. The delta reconstructed signal(s) are passed to a secondary low pass filter, and then finally digitized and analyzed as described in Section 2.3.

2.1 Time Generation Architecture and Analysis

The time base system is responsible for rapidly generating measurement opportunities for the reconstruction system to work with. Whereas a traditional undersampling architecture uses relatively few but high fidelity measurements, Exasampling uses relatively many measurements with comparably lower fidelity. The key metrics here are repetition rate and phase offset consistency, i.e., many samples that are tightly spaced.

2.1.1 Time Dilation from slow Phase Drift

The sample clock f_S and drive clock f_D , with periods T_S and T_D , are separated by a very small frequency offset Δf . The two clocks slowly drift across each other with a relative phase drift $\Delta\phi$ each sample period that accumulates to a total phase offset ϕ . As ϕ accumulates, the time base sequentially sweeps through the entirety of the input signal.

$$\Delta\phi = T_S - T_D$$

The slow drift produces a time dilation factor, D , which is the ratio of the periods of the original signal and its output reconstruction. The value of D is typically on the order of 10^6 to 10^{10} .

$$D = \frac{T_D}{T_S - T_D}$$

As an example, an Exasampling system configured with $f_S = 100\text{MHz}$ and $f_D = 100\text{MHz} + 0.1\text{Hz}$ produces a relative phase drift $\Delta\phi \approx 10^{-17}$ seconds, or 10 attoseconds. A full reconstruction cycle takes 0.1 seconds with a dilation factor $D = 10^9$. The reconstruction collects and uses 10 megasamples during this period.

2.1.2 Dual Numerically Controlled Oscillators

Creating stable, small Δf frequency spacing has historically been a difficult task. However, the wireless communication industry relies on well-defined channel spacing for successful communication in the presence of other radios, and has therefore developed several technologies that make this now possible.

The key is in using two numerically controlled oscillators that share a single reference time base, and then providing sufficient time for these oscillators to stabilize. A simplified overview is shown in Figure 2.

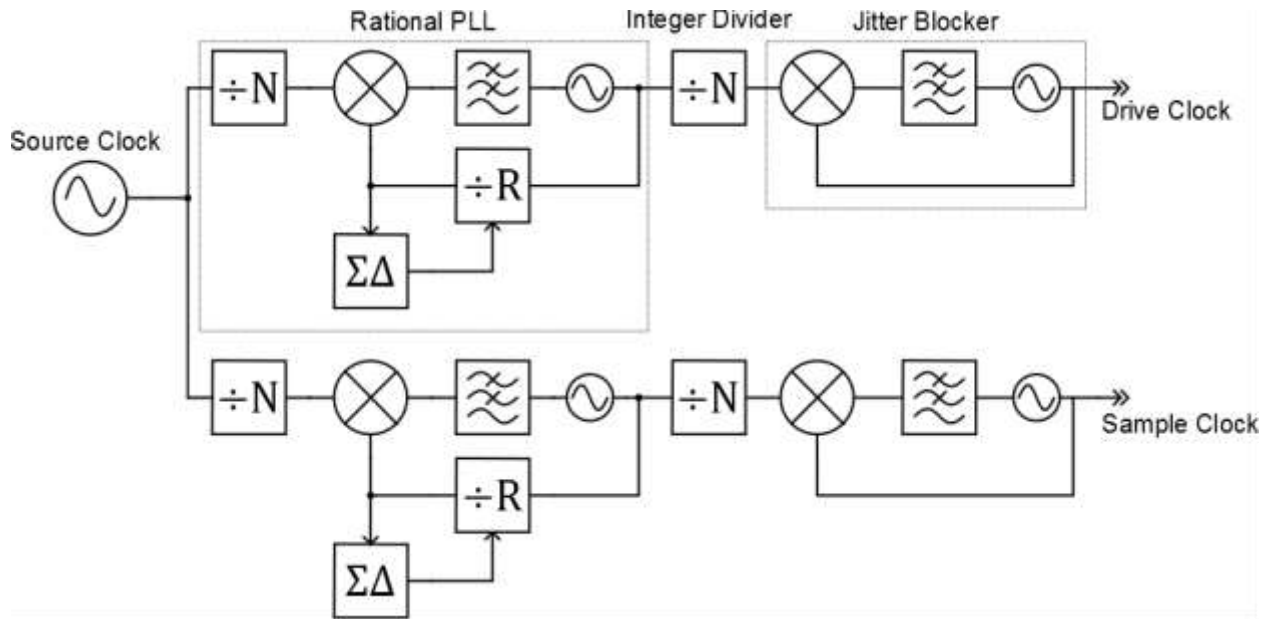


Figure 2: Frequency synthesis block diagram

A standard crystal oscillator provides a single time base to both synthesizers. Each synthesizer is based on a “rational” phase locked loop (PLL). These devices provide very high ratiometric precision, at the expense of added jitter. An integer divider provides some jitter reduction, but more importantly, it allows the two intermediate frequencies to be dissimilar and well-spaced. As discussed in Section 2.1.4, this decision is intended to reduce crosstalk between the two synthesizers. Finally, optional jitter blocking PLLs provide the final cleanup to produce the output clocks f_S and f_D .

The input clock is nominally 49.152 MHz. The Drive Rational PLL upconverts this to $3.3\text{GHz} + 3.3\text{Hz}$, whereas the Sample Rational PLL upconverts to 3.0GHz nominal. Both of these intermediate frequencies are chosen so they are not an integer multiple of the input clock to avoid integer boundary spurs. The Drive Integer Divider downconverts to $100\text{MHz} + 0.1\text{Hz}$, and the Sample Integer Divider produces 100MHz.

Each step in the clock chain is digital and ratiometric to the input frequency. The time dilation factor D is therefore “digitally perfect” over a sufficiently long averaging window if there is no “cycle-slip”. The combined system accuracy is completely independent of the individual absolute accuracies.

The absolute accuracy is limited primarily by the accuracy of the original crystal oscillator, which is typically tens of parts per million at birth and drifts strongly with temperature. This is five orders of magnitude less accurate than what would be required of an independent synthesis approach. Error in the shared oscillator produces error in the drive frequency, but does not produce error in the reconstruction of the resulting signal. This can be calibrated out by adjusting T_D and T_S , limited by the accuracy of the host’s own time base.

The oscillators are intended to run uninterrupted for many seconds at a time at any given frequency. This allows the jitter blocker filters to be quite aggressive, sacrificing settling time for further jitter reduction.

The combination of dual rational synthesis, low phase noise oscillators, and long running stability provides the Δf and jitter combination necessary for Exasampling.

2.1.3 Alternative Generation Methods

Undersampling time base generation can be broadly categorized as either controlling the sample phase directly (Sequential or Coherent Equivalent Time), or by sampling at an uncontrolled phase and measuring where it happened to occur (Random or Interleaved Equivalent Time [19]). The time-to-digital or digital-to-time converter determines the maximum equivalent time sample rate and contributes to the sample jitter of the device.

The most closely related method is the Vernier method of time interval measurement, which is often found in systems with random sample phases. The Vernier method measures a time interval by using a pair of startable oscillators with differing periods and a coincidence detector. In 1993 this method achieved sub-picosecond resolution with 2.5 ps jitter. [20]. More generally this method's resolution is determined by the beat frequency of the two oscillators, and its accuracy is determined by the jitter of the coincidence detector and the oscillator start behavior.

A key difference is that whereas Vernier measurement relies on startable oscillators, the oscillators in Exasampling run continuously. This allows the PLLs to be tuned primarily for stability rather than for the start transient, which allows for better jitter performance. Secondly, the absence of a coincidence detector removes its impact on system performance. Lastly, the Vernier method uses many clock periods to measure a single time interval, whereas the proposed method generates a time interval every single clock period. This removes the inherent linear tradeoff between sample rate and timing resolution, allowing for both a faster real time sample rate and a higher time dilation ratio.

Dual slope or Time Stretching methods are able to achieve a time resolution of roughly 10 picoseconds [21]. The resolution of this approach is directly scaled by the time stretch ratio, hence these methods were rejected due to their low repetition rate.

Traveling Wave or Tapped Delay Line [22] methods rely on large numbers of repeated digital elements. These approaches were rejected due to being unable to obtain suitably low delay elements to achieve the necessary resolution. However, this may be an appropriate solution for an ASIC based approach.

Delay Locked Loops (DLL) are sometimes used in interleaving ADCs to generate phase offsets between their separate track-and-hold channels, and can keep jitter below 600fs [23]. ADCs using this approach may calibrate out DLL with an aperture adjustment.

2.1.4 Effects of Jitter

Jitter is the all-encompassing term for short term variability in the time base generation system. These fluctuations in the relative timing between the clocks limit the effective system analog

bandwidth and distort the reconstructed signal. The effects of random jitter J and phase dependent jitter $J(\phi)$ modify the ideal steady state response to $Y \approx \langle X(\phi + J + J(\phi)) \rangle$.

Random jitter is the probability distribution of $\Delta\phi$ around its nominal value given as $J = f(\Delta\phi)$; the cycle to cycle variations in exactly when the sample is taken. It is by definition the portion of the jitter that is independent of both phase and signal content, and can therefore be independently analyzed. The steady state response becomes $Y \approx \langle X(\phi + J) \rangle$.

$$Y \approx \langle X(\phi + J) \rangle$$

$$= \frac{1}{N} \sum_{n=1}^N X(\phi + j_n)$$

Expanded by definition

$$Y \approx \int X(\phi + J)P(J)dJ$$

Which is by definition equivalent to

$$Y \approx (X * P)(\phi)$$

Writing it in this convolutional sum format allows us to treat the random jitter's probability distribution as the equivalent of an FIR filter applied to the input waveform.

Truly random jitter is typically assumed to be Gaussian in nature and measured in femtoseconds RMS. The effects of random jitter on an Exasampler can therefore be estimated as a Gaussian low pass filter with a corner frequency given by $f_c = 1/2\pi\sigma$, where σ is the RMS value of the total system jitter.

The effect of jitter is typically given as a signal to noise SNR limit equal to $-20 \log(2\pi f_{in}\sigma)$. It has been previously shown that oversampling can exceed this limitation in bandpass receivers. [24]. The reduction in the effect of aperture jitter on an Exasampling system is roughly equivalent to the reduction shown by Patel et al. taken to an extreme case.

Phase dependent deterministic jitter produces $Y \approx \langle X(\phi + J(\phi)) \rangle$. This can be viewed as a nonlinear mapping between real time and equivalent time, thereby causing distortions as ϕ "wobbles". Oscillator "Lock-In" or "Injection Locking" is the tendency for an oscillator to shift its frequency or phase to synchronize to another oscillator of similar frequency, and is a likely source of phase dependent jitter. Assuming that synchronization of the two oscillators did occur, digital error would accumulate and eventually break synchronization temporarily. The oscillator would over-correct until the digital error was compensated for and then synchronization would occur again. The system would then continue to alternate between these states, producing deterministic jitter and spurs.

2.2 Vertical Reconstruction and Analysis

Vertical reconstruction of the signal is based on an extreme version of oversampling, which in turn relies on having an equivalent time sample rate many times faster than the input signal bandwidth. ϕ sweeps through the input signal slowly enough that many of the analyses can assume it to be stationary.

2.2.1 Architecture




Figure 3: Reconstruction Block Diagram

The block diagram of the reconstruction architecture is shown in Figure 3. The input signal, X , is synchronous to the drive clock f_D . It is attached to the positive input of the sampling comparator. The reconstructed signal, Y , is attached to the negative input. The comparator's d-flip-flop samples $X(\phi) > Y(t)$ synchronous to the sample clock f_S . This serial stream has values V_+ and V_- , and is filtered to generate Y .

A low speed analog to digital converter digitizes the time dilated and reconstructed signal for further processing. This ADC has a Nyquist pre-filter to further eliminate noise from the sample clock. The digital stream is then filtered and/or demodulated as necessary.

2.2.2 Relation to sigma delta and delta converters

The core feedback loop and oversampling method is similar to that of the seldom used delta converter, which is related to the extremely popular sigma-delta converter [25]. This family of converters can achieve high resolutions, but by nature are limited to input bandwidths much smaller than would be implied by the Nyquist limit of their internal sample rate.

The delta converter is an analog to digital converter invented in 1946 and then “reinvented” in 1952, primarily for the purpose of encoding voice signals. [26] It produces a high rate binary stream representing either a positive or negative delta in the signal's value. Reconstruction is achieved with a simple integrator.




Figure 4: Delta Converter Architecture

The sigma delta converter is very similar to the delta converter, differing primarily in the placement of the integrator. It has greatly improved low frequency accuracy and noise performance, and it is currently the dominant architecture for high precision conversion. This approach was patented in 1954 first as an analog to digital encoder, and then expanded on by several contributors in the 1960s. [27] [25].




Figure 5: Sigma Delta Converter Architecture, 1st order

The Exasampling architecture's cost advantage comes in part from having as few high frequency components as possible; only the sampling comparator is required to handle the full input bandwidth. This necessitates moving the integrator from the sigma delta position to the delta position, meaning that it cannot perform noise shaping in the same way a sigma delta can.

The feedback filter in an Exasampler is not as constrained as it is in a delta converter, because it does not have to match a separate reconstruction filter in a receiver. This eliminates the impact of the low frequency drift error common to the delta architecture. Additionally, this feedback filter does not even need to be linear, it only has to be stable. This becomes important in the physically realized system discussed in Section 3.

2.2.3 Ripple

Each sample period increases or decreases the current estimate by ΔY_{\pm} . Unlike a purely integrating delta converter, this is not symmetric and is dependent on the current value of Y . Given a first order feedback filter with a corner frequency of f_F , ΔY_{\pm} is:

$$\Delta Y_+ = 2\pi(V_+ - Y)f_F/f_S$$

$$\Delta Y_- = 2\pi(V_- - Y)f_F/f_S$$

$$\Delta Y_{\pm} = \Delta Y_+ - \Delta Y_- = 2\pi(V_+ - V_-)f_F/f_S$$

In the ideal case, Y would never deviate by more than $1 \times \Delta Y_{\pm}$. This assumes that the comparator always makes the correct decision, and pushes the output back towards the correct value.

Comparator hysteresis adds directly to this peak to peak ripple, as it creates a region of uncertainty around the target value. Samples that fall within the hysteresis window will register positive or negative with unpredicted behavior.

Pure delay through the feedback loop can also be a significant contributor to ripple. With $f_S = 100$ MHz, there is only 10 ns of timing budget for the result of the sampling comparator to propagate through the comparator, the digital hold element, and the output D/A. The output alone is a significant contributor because there is a trade-off between $|V_{\pm}|$ and propagation delay, with modern CMOS outputs achieving delays and slew times on the order of 3-6ns. [28]

To first order, the total pure delay through this path forces the comparator to work with information from some number of cycles N in the past. Worst case, this linearly increases the peak output ripple by $N\Delta Y$. This value remains relatively flat with respect to f_S above a critical value, as N is proportional, and ΔY is inversely proportional, to f_S . Increased sample speeds decrease ripple up to a limit dictated by the propagation delay through the loop filter.

$$\Delta Y_{pp} \approx Hys + N\Delta Y_{\pm}$$

2.2.4 Signal to Noise

The impact of the ΔY_{pp} ripple on system performance can be quantified in its Signal to Noise Ratio (SNR). The following calculations assume flat frequency and probability distributions in the ripple noise. This is a conservative assumption for the probability distribution, and an optimistic assumption for the frequency distribution, as discussed in Section 2.2.7.

The peak SNR (PSNR) can be measured using either peak to peak noise or RMS noise, a choice that is primarily driven by the marketing department of the company selling the device. The calculations here use peak to peak, resulting in the “flicker free” equivalent resolution. Using RMS noise for PSNR gives results typically 2.7 bits better than peak to peak.

$$PSNR = 20 \log \left(\frac{\text{noise}}{\text{fullscale input}} \right)$$

$$PSNR = 20 \log \left(\frac{\Delta Y_{PP}}{V_+ - V_-} \right)$$

$$\Delta Y_{PP} \approx N 2\pi (V_+ - V_-) f_F / f_S + hys$$

$$PSNR \approx 20 \log \left(2\pi N f_F / f_S + \frac{hys}{V_{\pm}} \right)$$

A filter frequency to sampling frequency ratio of 10,000 gives a limit of -64dB PSNR. This implies any single sample has roughly 11 equivalent bits (ENOB) before oversampling.

2.2.5 Slope Overload and Slew Rate vs Granularity

Delta converters must balance their resolution (granularity) and the maximum input slew rate they can correctly digitize. The condition of an input waveform slewing too quickly for the delta converter is referred to as slope overload. An increased step size increases the maximum slew, but also increases the granularity. Some delta converters address this by dynamically adjusting their step size. [29]

A similar effect is present in Exasampling. A rapid change in the input waveform may be incorrectly reconstructed if it is too fast for the feedback filter at the current dilation factor.

The slew rate is the continuous time form of the loop filter response, and is given as:

$$\frac{dY}{dt} = \Delta Y f_S = 2\pi f_F (Y - V_{\pm})$$

This can be applied to the input waveform X with the dilation factor as:

$$\left| \frac{dX}{dt} \right| \leq D 2\pi f_F (X - V_{\pm})$$

The slope overload condition can be detected by monitoring the sampling comparators output for a long stream of 0s or 1s. If necessary, D can be increased, in turn increasing the maximum slew rate at the cost of also increasing measurement time. This provides better control over the bandwidth vs resolution trade off than a traditional delta converter.

The effect is shown in Figure 6. In this simulation, an overload condition occurs at $0.5\mu s$ where the ideal reconstruction has an instantaneous step change. The output of the sampling comparator, $V(Q)$ remains high for $110ns$ until the reconstruction output finally slews up to the ideal value.

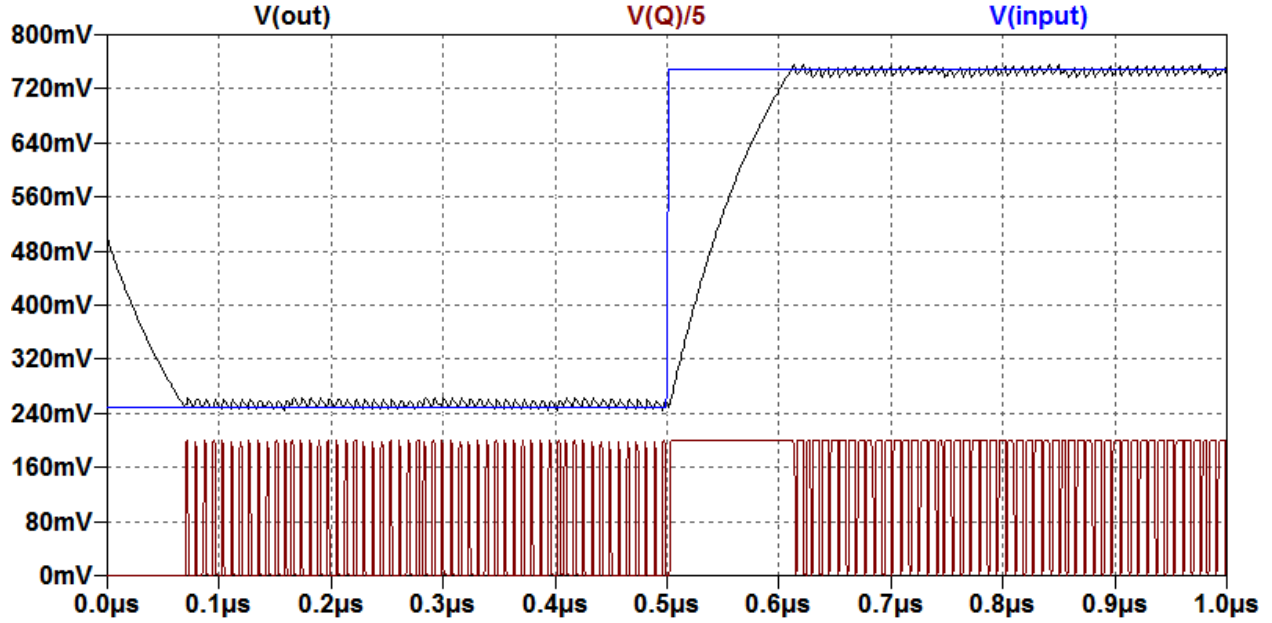


Figure 6: Simulated slope overload - $f_D = 500$ MHz, $D = 10^4$, $f_c = 1.6$ MHz

2.2.6 Secondary Filtering and Oversampling

The secondary filter decouples the ΔY_{\pm} in the feedback path from the ripple seen by the end converter. This decoupling provides several benefits including lessened hysteresis effects and shorter idle tone sequences. Additionally, keeping the primary filter first order avoids the conditional stability issues associated with higher order sigma delta and delta converters.

The ΔY_{pp} ripple noise is spread over a huge bandwidth compared to the reconstructed signal of interest. For example, a 10GHz signal reconstructed with $D = 10^9$ and $f_S = 100$ MHz has 7 decades difference between f_Y and f_S .

The upper limit for SNR improvement from oversampling is given as the ratio of f_S/f_Y . For the example above, this is nearly 12 effective bits of improvement. However, this assumes a perfectly flat frequency spectrum of the noise, a brickwall oversample filter, and does nothing to improve the linearity of the system. This is a very optimistic upper bound.

The secondary filter directly affects the measured Y waveform. Passband distortion here can be reflected to the input waveform time domain as Df_{F2} . Since D can vary with use case, it is recommended to set the corner frequency of the secondary filter sufficiently high such that at the lowest expected value of D , $Df_{F2} \gg BW$.

Filter agility is then achieved in the digital domain. The secondary filter effectively becomes a Nyquist filter for the analog to digital converter.

2.2.7 Idle Tones and Limit Cycles

Limit Cycling is the condition in which a modulator produces a repetitive output bit stream with a fixed output length. Idle tones are a semi-related but less well understood behavior that causes similar distortions, but does so with bit stream outputs that are not perfectly repetitive. [30].

Both of these effects result in distortion at subharmonics of the internal sample rate. The quantization noise is no longer spectrally flat, and forms peaks. This increases the requirements of the secondary filter beyond what was calculated in 2.2.3 above. A simulated exaggerated idle tone spectrum is shown in Figure 7. A peak 40dB above the noise floor is found at 15.625 MHz with $f_s = 500$ MHz. This is the 32nd subharmonic.

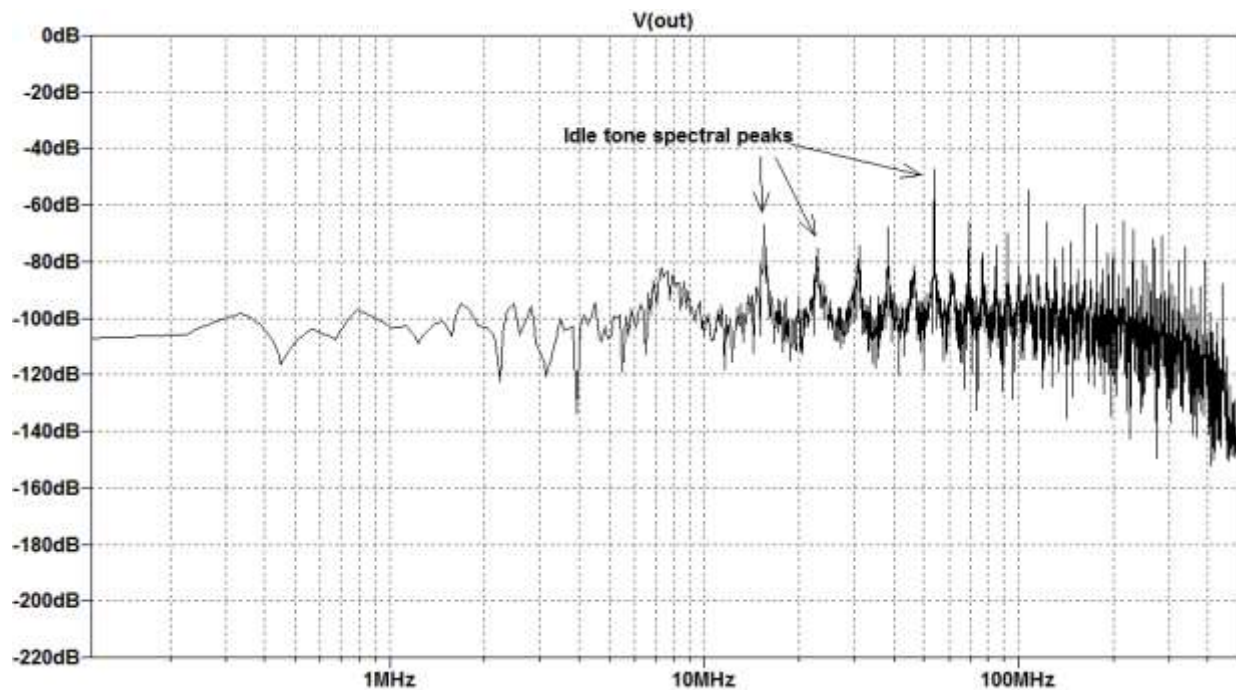


Figure 7: Idle tone spectrum

An Exasampler produces similar artifacts dependent on the input value. Figure 8 shows a simulated and exaggerated limit cycle behavior. The output stream repeats the 12 bit sequence "...101001010100..." for a fixed input of 40% of full range. Adding noise to the input breaks the limit cycle at the cost of increasing peak to peak ripple, as shown in Figure 9.

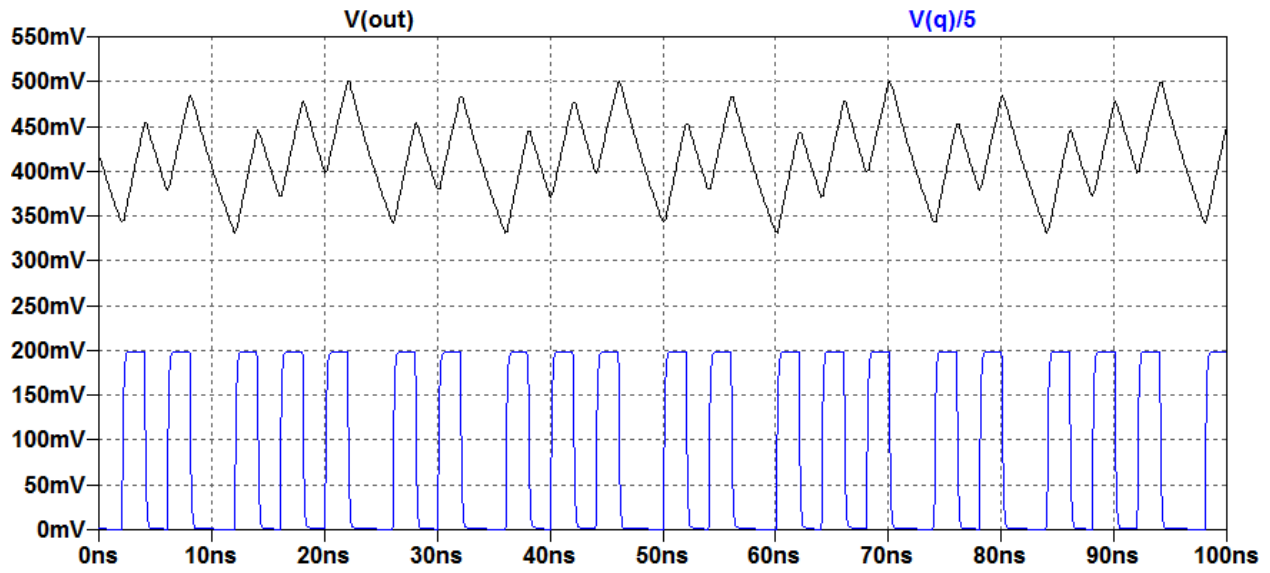


Figure 8: Limit Cycle




Figure 9: Dithering prevents limit cycle

2.3 Digitization

The Exasampling system is an analog to analog converter, and therefore requires an analog to digital converter to bring the signal into the processor. As discussed above, the resultant signal can carry 10 to 16 bits of useful equivalent resolution, and the resultant bandwidth can be arbitrarily low to fit the system requirements.

2.3.1 Digital Interface Requirements

One of the hidden costs of direct high frequency measurement is the rate at which data is generated and processed. Direct digitization of gigahertz signals generates gigabytes of data per second, which creates a hard coupling between the bandwidth of the signal and the cost of the associated digital logic. In contrast, Exasampling and traditional undersampling decouple the analog and digital bandwidths by their dilation factor. Again, the only component in the entirety of the Exasampling architecture that carries the full bandwidth requirement is the input to the comparator; all other components operate at slower, cheaper frequencies.

For reference, consider the processing requirements of sampling a 700 MHz signal at 8 bits and 3GS/s, creating a data deluge of 3 gigabytes per second. The ADC083000 from Texas Instruments achieves real time transmission of this using a total of 66 digital logic outputs arranged as 33 differential pairs, each of which must be routed with controlled impedance, terminated, and length matched. This typically requires an FPGA to process. The ADC08B3000 is a sister part that uses a piece of high speed memory to buffer up to 4,0000 samples and then slowly dole them out over a much slower 20 pin 200 MB/s interface. This still requires the digital logic inside the converter to be operating at 3 GB/s, but is now slow enough for a high end microcontroller.

An equivalent Exasampling system can push the same bandwidth and resolution over a single pin with an arbitrarily slow data rate, such that an arbitrarily low cost microcontroller can be used. For example, the MSP430FR2111 from Texas Instruments sells for \$0.45 each is sufficient for processing a 10 GHz signal at 12 bits of resolution when $D = 10^9$. Increasing the performance of the microcontroller allows the Exasampler to reduce its dilation factor and increase the effective data rate; the analog and digital bandwidth decoupling allows for a smooth tradeoff between cost and quality.

2.3.2 Analog to Digital Converter Selection

The choice of analog to digital converter can be largely driven by the bandwidth and resolution. The major noise sources in the system are very high frequency compared to the reconstructed signal; ripple from the delta converter, power supply noise from the stimulus generator, and blow-through from the signal(s) being reconstructed. Active filters may run out of gain bandwidth at these frequencies.

Rather than filtering these sources out, it is possible to use analog to digital converter architectures that are inherently less susceptible to them. Architectures with wide apertures and/or slow settling times are preferred over tighter sampling windows. This odd anti-requirement is similar to that found in coulomb counters. For example, the input integrators in dual-slope and sigma delta converters make them preferable to a successive approximation converter with a sample and hold.

3 Implementation and Performance Testing

An Exasampler was constructed with 4 reconstruction channels and 2 differential stimulus channels.



Figure 10: Prototype Sampler PCB

The prototype is composed of three separable parts: the time base generator, the sampler, and the digitizer. The time base generator is an AD9578 evaluation board lightly modified to provide f_D and f_S as differential ac-coupled LVPECL signals. The sampler PCB consumes these clock signals to drive the hosted stimulus generators and reconstruction channels. The time dilated reconstructions are then passed to the digitizer for processing and analysis.

At single digit production volumes the total cost was less than \$200 per device. Without further cost-down efforts, increasing production volume to 10k units brings the figure under \$100 for this 4+2 channel device. Over 3/4ths of this cost is in the choice of comparator; Chapter 4.3 examines using PHYs instead of sampling comparators to further reduce this cost.

3.1 Sampler Construction and Layout

The sampler PCB layout is shown in Figure 11. It measures 130mm by 65mm. The drive clock enters on the left through J18 and J19 to drive the two LVPECL stimulus generators. The four reconstruction channels are vertically arrayed in the center, receiving their inputs on the SMA connectors to their left. The sample clock enters on the right, is buffered by a Silicon Labs Si53301 clock buffer, and sent to the reconstruction channels. These clock lines have matched skew, but this proved to be an unnecessary precaution. Power management and termination voltage generation is located in the upper left.

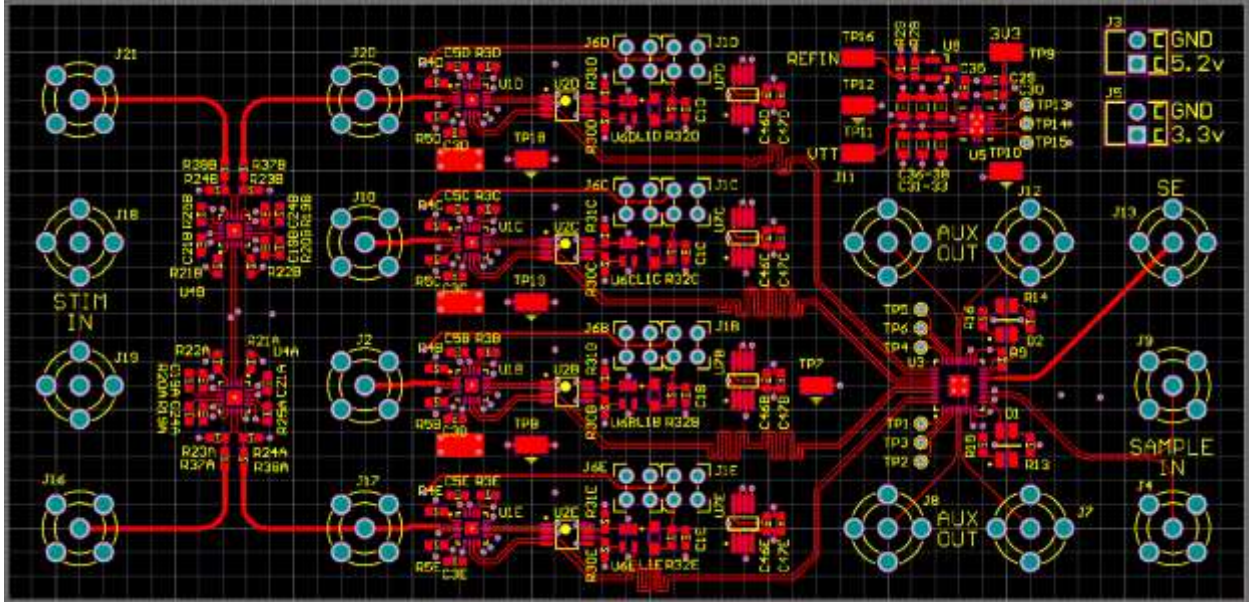


Figure 11: Sampler layout




Figure 12: Exasampling unit cell layout

3.1.1 Sampling Comparator

The sampling comparator is constructed from three independent ICs. Most importantly, the ADCMP573 from Analog Devices is a SiGe based high speed sampling comparator with 8 GHz of analog bandwidth. [31]. It is the only component in the system that directly interacts at the full bandwidth.

Unfortunately, commercially available sampling comparators with comparable bandwidth are typically latching and use a low swing output stage, such as LVDS, LVPECL, or CML. Exasampling requires a registered output and a full swing output. An On-Semi MC100EP52DTG d flip-flop registers the latched output, and an Analog Devices ADCMP600BKSZ expands the output range to fully cover the input range.

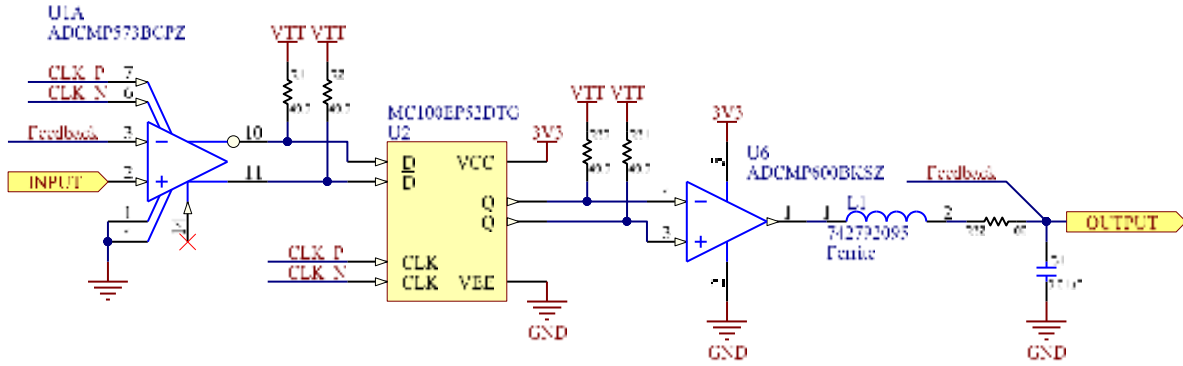


Figure 13: Sampling comparator and primary filter schematic

This introduces several propagation delays through the feedback loop, as shown in Figure 14. Setup and Hold time of the flip flop is achieved by clocking it on the opposite edge of the latch, and routing the clock line to provide 45 picoseconds of clock skew. The physical impact of this skew requirement can be seen in the physical separation between U1* and U2* in Figure 12.




Figure 14: Feedback loop timing diagram

The performance penalty of having to translate from latched LVPECL to registered CMOS is roughly 5 nanoseconds plus half a sample period. An equivalent monolithic sampling comparator would increase resolution by roughly one equivalent bit per the calculations in Chapter 2.2.

The primary filter was initially set with $f_c = \frac{1}{2\pi}$ MHz. Additional filtering was provided by the slew of the output DAC, and by a series ferrite bead with a notch frequency equal to $f_D = 100$ MHz. The total frequency response of this system is shown in Figure 15. The additional notching is intended to cut high harmonics out before they have a chance to rectify against the secondary filter's ESD protection diodes, though system evaluation later showed this protection to be an unnecessary precaution.

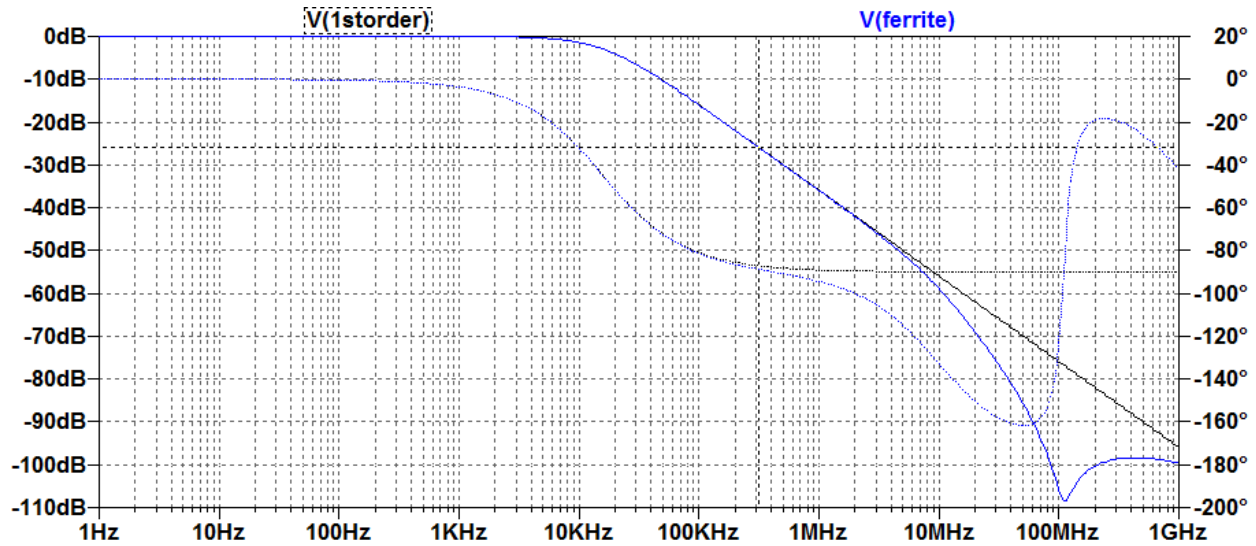


Figure 15: Primary filter with and without ferrite bead

3.1.2 Secondary Filter

The secondary filter is footprinted as a 4th order low-pass with double corners at 100 kHz and 1 MHz, as shown in Figure 16. In retrospect, this approach was overkill, as the additional filtering provided here did not significantly improve system noise performance.

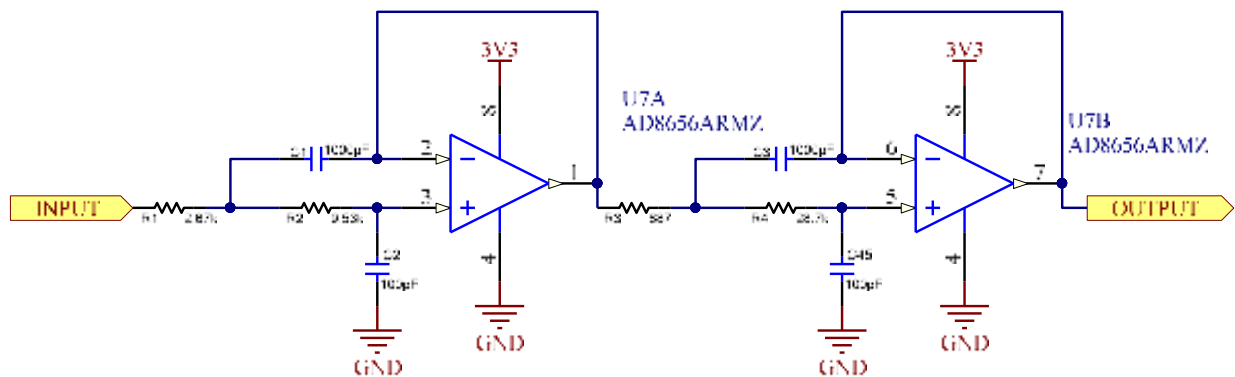


Figure 16: Secondary filter schematic

3.1.3 Frequency Requirements

Figure 17 highlights the relative frequencies of the signals running through the unit cell. The only signal that carries the full bandwidth of the input signal is X , shown in green between the input SMA connector and the comparator. Blue signals operate at f_s (nominally 100 MHz); this is only the clock and the communication lines of the d flip flop. The signals shown in yellow are the portion of the signal path that operate at the slow, time dilated reconstruction with only a few kilohertz of bandwidth.

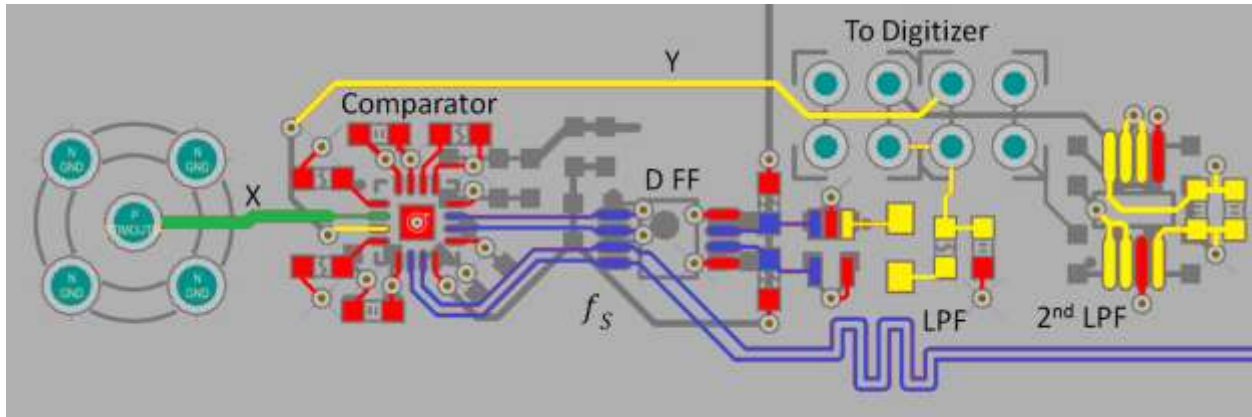


Figure 17: Unit cell frequencies. Green is 10GHz. Blue is 100MHz. Yellow is 10kHz.

The high frequency section is extremely small and localized; it may be possible to use a coplanar waveguide layout for X instead of a microstrip in order to remove the need to control the dielectric stack-up. In higher precision applications it may be prudent to place a ferrite choke along the Y feedback path to further reduce high frequency coupling.

3.1.4 Digitizer

Digitization of the reconstructed signal is performed by a Cypress PSoC 5LP programmable system on chip. This integrates programmable analog filters, a 48 ksps 16 bit $\Sigma\Delta$ ADC, digital filter logic, an ARM M3 Cortex operating at 80 MIPs, and a USB 2.0 device port. [32] The four channels were filtered, digitized, and passed via the USB port to a MATLAB program for further processing. Filters optimized in MATLAB can be ported back to the PSoC.

3.2 Stimulus Generation

The Exasampling system architecture provides a deal of flexibility in what waveform is sent to the device under test. The only mandatory requirements are that the drive stimulus waveform must be synchronous to the drive clock and stable over many repetitions. The choice of stimulus generator is therefore dependent on the specifics of the device under test and the intended measurement.

Several measurement modes use the drive clock directly without an intervening stimulus generator. Other use-cases have the equivalent circuits built into the device-under-test.

3.2.1 Digital Outputs

High speed digital communications use output buffers that produce a square stimulus waveform with sharp, well defined edges and have well known output impedances. For example, LVPECL is 50Ω and can achieve a $>35\text{ps}$ rise time. For a well matched DUT, this is essentially a step response measurement. Pairing them with a sampling comparator built on the same standard often implies that they have matching input and output voltage ranges, and that their analog bandwidths are comparable.

One important note is that many output stages are strongly dependent on their load. For example, ECL uses a differential current steered output stage. If one of the outputs is allowed to

float, then the current source will overload and collapse when that output is enabled. At the next clock edge the current source will take time to recover, compromising the turn on transient.

The spectrum of an LVPECL based stimulus generator is given in Figure 18, showing that the output stage provides energy out to the 25th harmonic of the drive clock. This can measure the frequency response of a device out to 3 GHz, without any RF rated components.

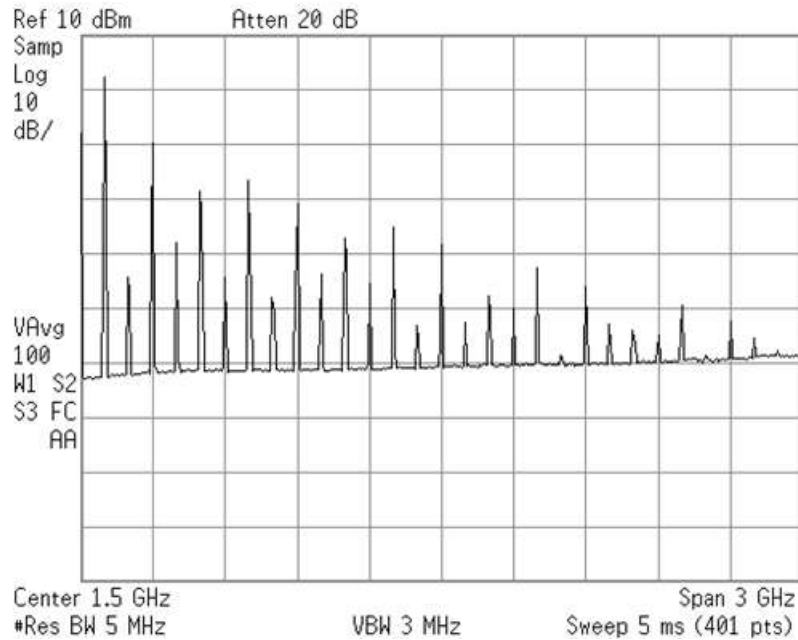


Figure 18: LVPECL spectrum

3.2.2 SRD / PIN comb generator

Diode based comb generators historically use a step recovery diode's recovery "snap" action that occurs when its bias changes abruptly from forward to reverse [33]. This creates a pulse of comparatively high current with a very short duration, pushing energy to high harmonics of the drive clock.

Traditional SRD comb generators intend to produce a waveform that cleanly approximates an impulse. This allows the measured response to be interpreted directly as an impulse response. This cleanliness requirement is not strictly necessary in Exasampling use cases, as the impulse response can be obtained mathematically if input and output are both known. Relaxing this requirement gives design flexibility to optimize other parameters.




Figure 19: PIN Comb Generator simplified schematic

A comb generator was adapted from David Bowman's (G0MRF) Source2 [34] to function as a stimulus generator. As shown in Figure 19, the drive clock feeds a MAR3+ wide-band RF amplifier, which then drives in to a HSMP-3822 dual PIN diode. The snap action of this diode pair produces the waveform shown in Figure 20, and the spectral response shown in Figure 21. Note that this was tuned not for a clean impulse, but rather for strong high frequency content. This created harmonic content out to the limits of the spectrum analyzer; -42.7dBc at 6.63 GHz.

This entire subsystem costs less than \$2 USD in 1000 piece quantities.




Figure 20: PIN Comb Generator waveform




Figure 21: PIN Comb Generator spectrum

Both of these plots were taken with 20dB of broadband attenuation. When this generator is improperly terminated, it radiates strong broadband interference well outside of FCC guidelines and is sufficient to affect most of the lab’s test equipment. The attenuation provides the necessary termination when the DUT is not well matched. This is less energy efficient than reducing the amplifier gain, but has the advantage of being more robust to load pulling effects.

3.3 Performance Testing

A major difficulty in characterizing the Exasampling prototype is that several of its performance metrics exceed that of the available test equipment. The main metrics of concern are repeatability, jitter, and (vertical) signal to noise.

3.3.1 Repeatability

The first measure of performance is how well the unit agrees with itself over many repetitions. These tests do not inform the true accuracy of the system, but it does set bounds on the quality of the information provided.

The output of a PECL buffer was intentionally incorrectly terminated in order to produce a large amount of overshoot and ringing. The response was reconstructed with $D = 10^9$ and $f_D = 100\text{MHz}$. Six separate reconstructions were phase aligned and then overlaid to generate Figure 22.

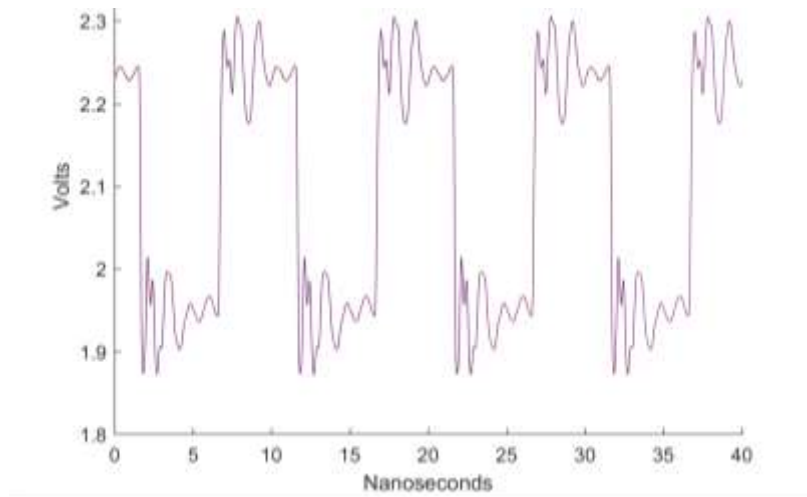


Figure 22: PECL with mistuned termination

The reconstructions match sufficiently well that they are indistinguishable at a zoom level that shows the entire signal. Figure 23 is a tight zoom in on a single representative nanosecond. It shows timing repeatability of 3 picoseconds and vertical repeatability of 1 millivolt.

For a 400mV full-scale signal, 1 millivolt peak-to-peak uncertainty equates to an SNR of 52dB or 8.6 flicker-free bits. The system can handle a 1.2V full scale signal anywhere in its 3.2V input range, so the PSNR is as high as 13 effective bits.

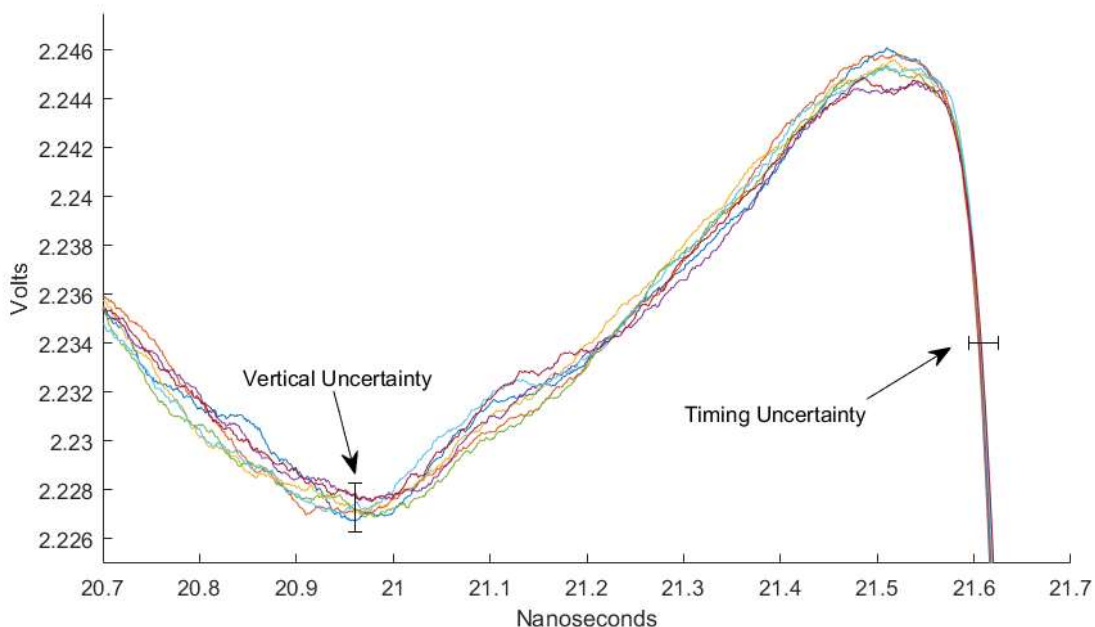


Figure 23: Reconstruction overlaid detail zoom

3.3.2 Traditional Jitter Measurement

The first attempt to measure total system jitter was made with a 7300A Teledyne Lecroy Oscilloscope. The drive and sample clocks were fed directly into the oscilloscope, which then measured and trended the edge to edge timing. The measured jitter with this method matched the Oscilloscope's own channel to channel jitter spec of 2.5ps RMS, and similarly matched its single channel jitter spec of 1.0ps RMS. This strongly implies that the Exasampling system's time base generator is at least matching the performance of the oscilloscope, but it does not allow us to confidently put a number on that performance.

Jitter was also measured with a E4403B Agilent spectrum analyzer. This only has a single input, so it can only measure jitter independently. It cannot measure the relative jitter. The width of the main lobe was >10 kHz, suggesting a picosecond jitter figure.

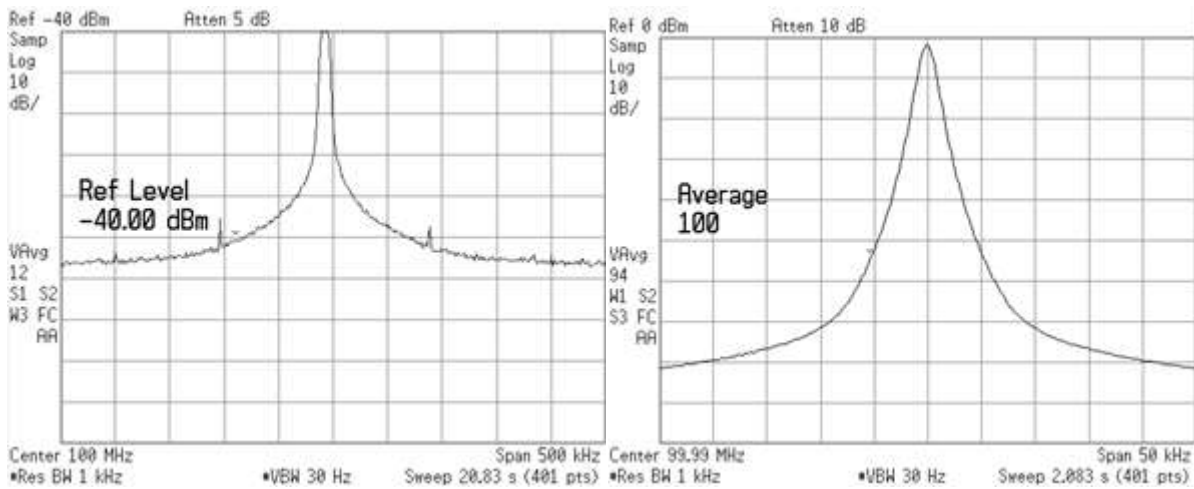


Figure 24: Spectral Spurs

A pair of narrow band spurs are visible at ± 80 kHz in Figure 24. These are only present with both clocks enabled, and are only barely peaking above the Spectrum Analyzer's noise floor. The symmetry of the spectrum implies that the PLL is appropriately compensating for the pulling effect of the other oscillator [35].

3.3.3 New Jitter Measurement

An adapted version of the beat frequency method was used to produce a reconstruction of the probability distribution function of the sample phase offset drift, $J = f(\Delta\phi)$. This method relies on the Exasampling system's time dilation to measure its own jitter indirectly.




Figure 25: Configuration for Jitter Self-Test

In this self-test mode, the main feedback loop is broken open-loop and configured as shown in Figure 25. The sampling comparator latches high if the drive clock arrives before the sample clock. This captures the cumulative total system jitter effects from the drive clock, sample clock, and the comparator. Each sample period is a single test of $f(\Delta\phi) < \phi$, and the averaging from the lowpass filter produces a running estimate of the cumulative distribution function $F(\Delta\phi)$. The entirety of this distribution is sampled as ϕ drifts from negative to positive.

Note that the effects of the low pass filter are irrecoverably included in this measure and subsequently interpreted as additional system jitter. This is not an issue when measuring total system performance, because this effect is also present during reconstruction. However, it does make this a conservative measure when attempting to characterize jitter alone.

Raw results are shown in Figure 26. Time base is 500 microseconds per division with $D = 10^9$, giving an equivalent time base of 500 femtoseconds per division. The cumulative probability distributions on the rising and falling edges are both shown. Analog persistence captures the distribution of these measured distributions; horizontal alignment is largely a function of the oscilloscope's trigger.




Figure 26: Jitter probability as a function of nominal phase offset in picoseconds

This data shows a Gaussian distribution of jitter 732 femtoseconds RMS, which aligns well with the calculated jitter budget and has an equivalent filter corner at $f_c = 1/2\pi\sigma \approx 200$ GHz. This is well above the filter imposed by the comparator's input bandwidth, and is therefore assumed to be good enough to be ignored in further calculation.

3.3.4 RF Device Characterization

The “gold standard” measurement equipment to characterize RF devices is the Vector Network Analyzer. To directly compare performance, a calibrated VNA and the prototype Exasampler both measured the frequency response of a 1200 MHz low pass filter. The results are shown in Figure 27 below.

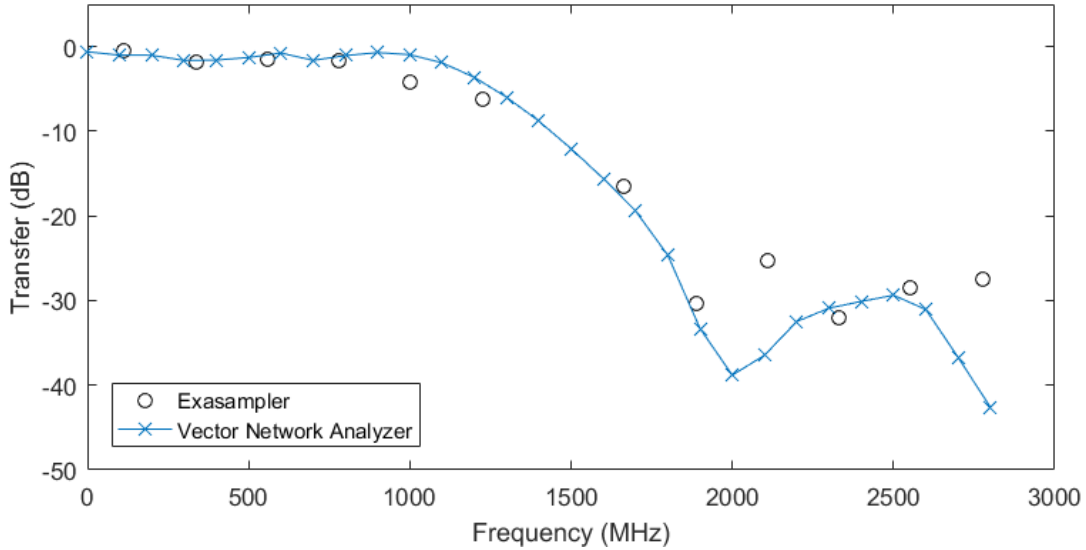


Figure 27: Characterization of 1200 MHz Low Pass Filter

There is good match at lower frequencies, but diverges as the filter response rolls off past -30dB and misses the notch at 2 GHz. The divergence is likely due to the difference in vertical dynamic range between the VNA and the prototype Exasampler. The VNA is very well optimized for dynamic range by measuring a single frequency at a time, and digitizes logarithmically. In contrast, the Exasampler measures the entire set of harmonics at once, and digitizes linearly.

4 Applications

Exasampling excels in applications where the system under test is relatively stable; the rate of information generated is much slower than the frequency of the signal. Three example applications are given.

Dielectric Spectroscopy highlights the Exasampling's ability to look into the chemical properties and changes of household materials. LIDAR not only shows the raw temporal resolution of the system, but also that good measurement can compensate for bad sensors.

Finally, the improvement to the USB 3 SuperSpeed protocol shown in Section 4.3 shows how these techniques can be integrated into existing technologies. If the proposed improvement is adopted by the USB IF, then as a side-effect all future USB 3 devices and hosts would be able to make Exasampled measurements natively.

4.1 *Dielectric Spectroscopy at Home: Epoxy*

Dielectric spectroscopy provides a non-invasive, non-destructive method of evaluating the composition and state of a sample. The cost per test is very cheap, as it usually does not require any consumable or disposable components. However, the upfront cost is prohibitive for the home market. [36]

Exasampling makes spectral sensing affordable for the home, consumer, and personal health markets. This accessibility can bring better understanding and more fulfilling interactions with the chemical processes that affect our daily lives. Potential applications include refrigerators that can detect milk souring before a human nose can, pill bottles that ensure the efficacy of their contents, and cars self-diagnose lubricant degradation. [37]

Dielectric monitoring of epoxy curing is a Double/Bubble® Red Extra Fast Setting Epoxy Adhesive is a two part epoxy with an advertised work time of 3 to 5 minutes. The Exasampling system was used to monitor this chemical reaction in real time.

4.1.1 *Results*

The raw data collected is shown in Figure 28. The dilation factor was chosen to make the cure time occur over ≈ 30 reconstructions, making the evolution of the response easier to visually follow.




Figure 28: Epoxy Cure Test

The chemical change can be seen more clearly by plotting the harmonics vs time, as shown in Figure 29. These have been normalized to more clearly show the trends. The transmission response at 100MHz changes by three-fold over the curing process.

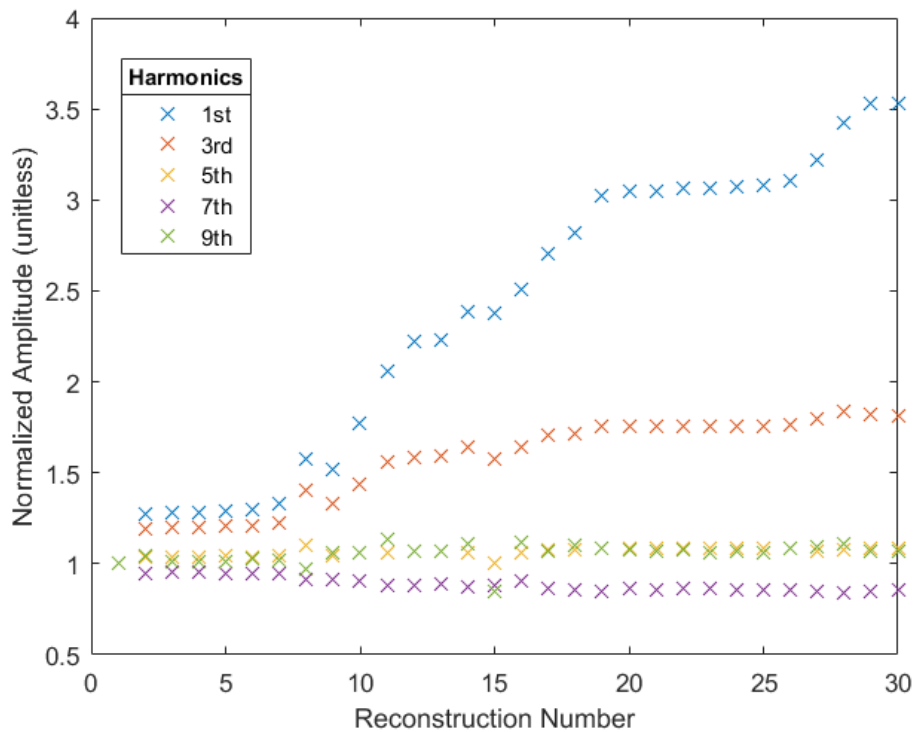


Figure 29: Transmission harmonics over time

4.2 LIDAR

Time-of-flight imaging uses the time delay between the light source and receiver to measure distance, examine sub-surface phenomena, and even see around corners. [38] This section describes measurement of sub-millimeter changes in the path length of a light beam using Exasampling and an optical Ethernet PHY. Although not a full LIDAR, this test case highlights both the temporal resolution improvement that Exasampling provides to time-of-flight systems, and its ability to measure phenomena with greater fidelity than the underlying sensor natively supports.

4.2.1 Method and Setup

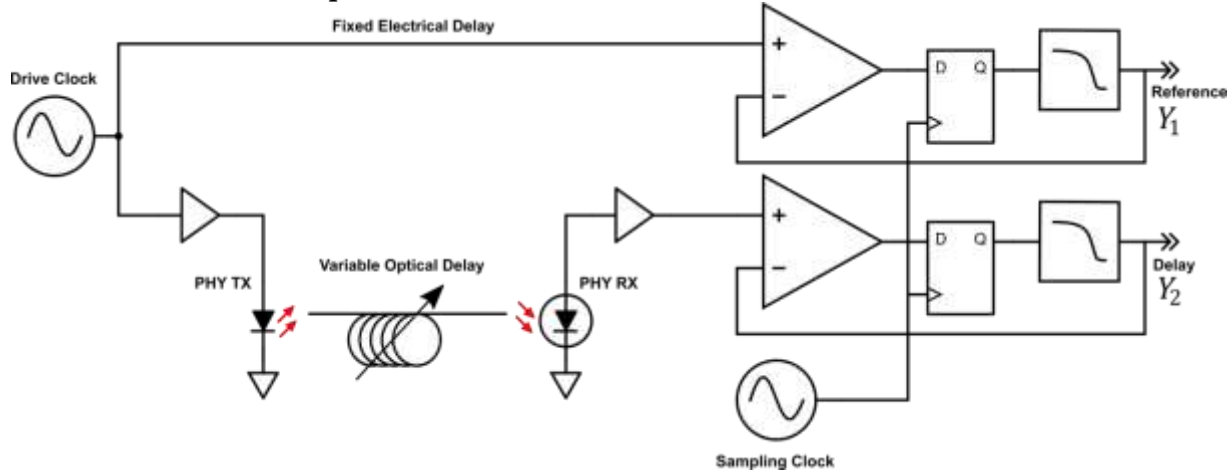


Figure 30: Exasampled LIDAR Architecture

The time-of-flight measurement is performed using a two channel Exasampler, as shown in Figure 30. The primary channel measures a variable delay optical path with an optical Ethernet PHY and produces result Y_1 . The second channel directly converts the drive clock through a fixed delay line to produce result Y_2 .

Cheap Optical Ethernet PHYs are easy to electrically interface with an Exasampling system. The Broadcom AFBR5972 and AFBR5803 optical Ethernet PHYs provide LVPECL interfaces to 125Mbaud optical transmitters and receivers. Critically, they provide raw interfaces with no internal clocking logic. [39]

The relative delay of the two channels is a measure of the relative path lengths, modulo the drive frequency. Since the delay of one channel is fixed, a change in the relative delay ΔT indicates a change in the delay of the channel being measured. These times are all scaled by D , giving the measure of change in distance as:

$$\Delta d = \frac{c\Delta T}{D}$$

For a nominal $D = 10^9$, this gives a scaling factor of 0.3m/s, or 300 nanometers per microsecond. A 1 MSPS conversion of this result should therefore give 300nm horizontal resolution, but this can be oversampled further.

4.2.2 Results

A typical result waveform is shown below in Figure 31. The down converted drive clock is shown in blue, and the receiver response in yellow. Note the receiver response has significant slew; the PHY is rated for several nanoseconds of jitter, duty cycle distortion, and rise/fall times. The nominal Δd was measured with a linear positioning micrometer.




Figure 31: Sample LIDAR waveform. 2 seconds per division

The highest dilation factor achieved produced a change in timing delay of 16 seconds per meter, which is equivalent to slowing the apparent speed of light by a factor of 4.8 billion.

4.3 Future Work: USB SuperSpeed™

The latest update to the omnipresent USB communication standard greatly increased its data transfer rate with the addition of two additional high speed communication lanes. Exasampling techniques can be used to improve the characterization and training of these lanes. Additionally, the proposed modification upgrades a standard USB port into a multi-gigahertz analyzer, and opens the door to dielectric spectroscopy for anyone that owns a laptop and half of a USB cable.




Figure 32: USB SuperSpeed Topology

4.3.1 Channel Sounding and Training

The SuperSpeed channels operate at 10 gigabits per second of data throughput, and 12.5 gigabaud of raw symbol rate. This requires tunable equalizers in both the transmitters and receivers to compensate for cable loss and length to ensure reliable operation over the full range of conditions.

The PHYs are currently incapable of true channel sounding; they are unaware of the true cable parameters. They instead use a more holistic approach, channel training, to find an equalizer tuning that is sufficient to compensate for the unknown channel. A known training sequence, “TSEQ” is repeated at least 2^{16} times to allow the receiver to explore the equalizer’s tuning space. Candidate tunings are evaluated on their raw error rate and the types of errors they generate.

If the actual cable parameters were known, a more optimal equalizer tuning could be directly computed, potentially increasing reliable data throughput. An adaptation of Exasampling can directly measure those parameters using the hardware resources that are already present in the PHY.

4.3.2 Sounding Configuration

One proposed configuration is shown in Figure 33. This approach requires minor hardware changes in the sounding PHY, and can be performed with an unmodified partner PHY with test modes already present.

The hardware change is limited to the sounding PHY’s receiver. It gains a sampling clock, and provides a selectable feedback path to the existing sampling comparator’s input. Alternately, two additional sampling comparators can be added that are always in this mode. The configuration shown has two sampling comparators in an inverting configuration.

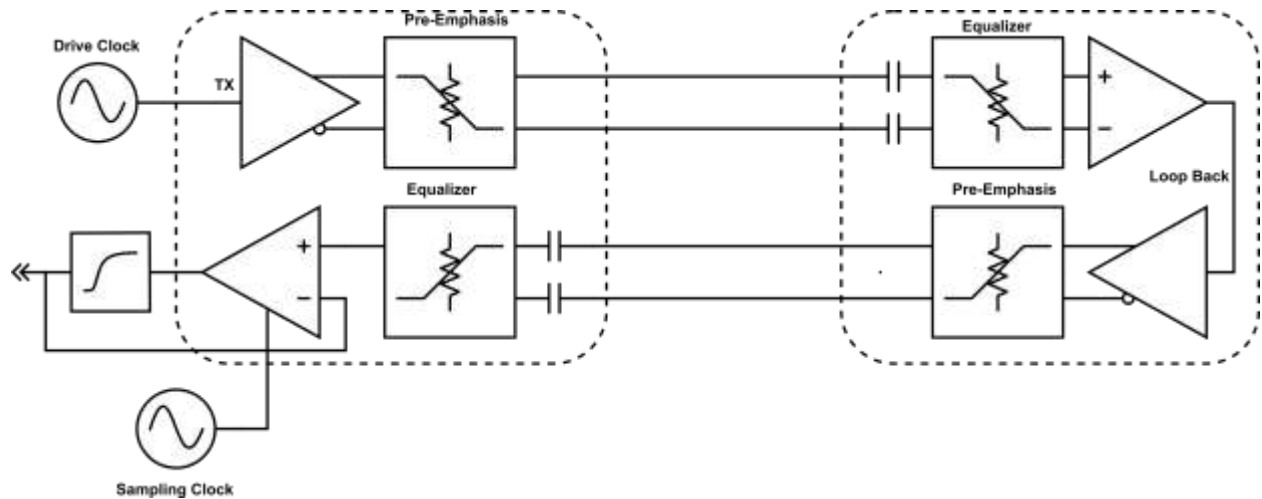


Figure 33: Channel Sounding Configuration

Data is clocked out of the sounding PHY at f_D or $f_D/2$, received by the partner PHY and immediately retransmitted back. This is achieved with a test mode sometimes referred to as “analog loop-back” because it bypasses the digital processing and re-timing systems. It is important to note that the remote PHY’s transmitter and transmitter equalizer are in play in this mode. The USB 3.1 spec allows for up to 1 picosecond of added jitter here.

This is now an Exasampler, albeit with a split architecture. It produces a pair of single-ended reconstructions of the differential waveform. They include the effects of the remote PHY’s transmitter equalizer, the cable, and the training PHY’s receiver equalizer.

To account for potentially non-linear inter-symbol interference, the training PHY can emit a training sequence clocked at the sample clock rate. A short 3-bit rotary grey sequence such as 00011101 is recommended.

Compliance testing of a USB 3.1 channel requires an oscilloscope with at least 20 GHz of bandwidth. This approach can perform most of those lab tests as self-tests.

4.3.3 Use as a Stand-Alone Sensor

The sounding PHY is now a fully functional Exasampler, and can be repurposed for many of the applications above. For example, it can perform Dielectric Spectroscopy by replacing the partner PHY with the material of interest. A USB DS probe would simply be a connector and conductors in the desired geometry.




Figure 34: Measuring milk freshness with USB dielectric spectroscopy

5 Conclusion and Future Work

This thesis shows that high frequency does not necessarily mean high complexity or high cost. Digital communication components can be used to make high frequency, high fidelity analog measurements of stable systems.

The next major step for Exasampling is to integrate it more tightly on silicon. However, this isn't a necessary step. These techniques are ready for immediate application and deployment. It is a cheap and cheerful method to reduce cost and complexity for several classes of sensor systems.

A surprising result is that the imperfections in cheap digital components actually become beneficial. Overshoot, ringing, and other EMI generating artifacts allow these components to push energy into high harmonics. Intentionally mistune matching networks. Use low quality components. Break things. If it is repeatable, Exasampling can pull useful information out of these typically undesirable effects.

The hard work has already been completed by the communication industry; they have already commoditized the gigabaud. Now we can commoditize the gigahertz.

Nanoseconds for the masses.

6 Terms

f_D, T_D	Frequency, period of the drive clock
f_S, T_S	Frequency, period of the sample clock
f_X, T_X	Frequency, period of the synthesis clock
f_F	Corner frequency of the feedback filter
Δf	Difference in frequency between the drive and sample clocks
D	Time Dilation factor, unitless
ϵ_X	Error of the synthesis clock, unitless
ϕ	Sample Phase Offset. The current phase delay between the drive and sample clocks.
X	Input waveform to the Exasampler
Y	Output waveform of the Exasampler
f_X	Bandwidth of the input waveform
f_Y	Bandwidth of the output waveform. Equal to f_X/D
Q	Quantization noise
V_{\pm}	Output voltage of the sampling comparator's D/A for a positive or negative result.
J	Random jitter
$J(\phi)$	Phase dependent jitter
VNA	Vector Network Analyzer
SHA	Sample and Hold Amplifier, similar to track and hold amplifier
DUT	Device Under Test
SNR	Signal to Noise Ratio
PSNR	Peak Signal to Noise Ratio

7 Bibliography

- [1] M. Marino, "Executive arrested in taiwan cooking oil scandal," *Olive Oil Times*, 24 October 2014.
- [2] M. A. Sairin, S. A. Aziz, N. N. A. Nizar, N. A. A. Latiff, A. Ismail, D. M. Hashim and F. Z. Rokhani, "Lard Detection in Edible Oil Using Dielectric Spectroscopy," in *Sensors for Everyday Life*, 2016, pp. 245-271.
- [3] S. O. Nelson, "Agricultural Applications of Dielectric Spectroscopy," *Journal of Microwave Power and Electromagnetic Energy*, pp. 75-85, 39:2 2004.
- [4] S. O. Nelson, "Dielectric Spectroscopy of Fresh Fruit and Vegetable Tissues from 10 to 1800 MHz," *Journal of Microwave Power and Electromagnetic Energy*, pp. 31-47, 40:1 2005.
- [5] P. Mohr, U. Birgersson, C. Berking, C. Henderson, U. Trefzer, L. Kemeny, C. Sunderkötter, T. Dirschka, R. Motley, M. Frohm-Nilsson, U. Reinhold, C. Loquai, R. Braun, F. Nyberg and J. Paoli, "Electrical impedance spectroscopy as a potential adjunct diagnostic tool for cutaneous melanoma," *Skin Research & Technology*, pp. 75-83, 27 January 2013.
- [6] R. P. Henderson and J. G. Webster, "An Impedance Camera for Spatially Specific Measurements of the Thorax," *IEEE Transactions on Biomedical Engineering*, pp. 250-254, May 1978.
- [7] T. A. Hope and S. E. Iles, "Technology review: The use of electrical impedance scanning in the detection of breast cancer," *Breast Cancer Research*, pp. 69-74, 13 November 2003.
- [8] R. J. Halter and Y.-J. Kim, "Toward Microendoscopic Electrical Impedance Tomography for Intraoperative Surgical Margin Assessment," *IEEE Transactions Biomedical Engineering*, pp. 2779-2786, November 2014.
- [9] L. Hoffmann, C. Gautier, S. Lefebvre and F. Costa, "Thermal measurement of losses of GaN power transistors for optimization of their drive," in *Power Electronics and Applications*, Lille, France, 2013.
- [10] D. Berning and A. Hefner, "IGBT half-bridge shoot-through characterization for model validation," *Industry Applications Conference*, pp. 1491-1499, 6 October 1996.

- [11] Peregrine Semiconductor, "PE29100 Product Specification," 01 08 2016. [Online]. Available: <http://www.psemi.com/pdf/datasheets/pe29100ds.pdf>. [Accessed 5 11 2016].
- [12] H. C. P. Dymond, "Reduction of oscillations in a GaN bridge leg using active gate driving with sub-ns resolution, arbitrary gate-impedance patterns," *IEEE ECCE*, 2016.
- [13] C. E. Shannon, "Communication in the Presence of Noise," *Proceedings of the IRE*, pp. 10-21, January 1949.
- [14] S. v. T. E. Louwsma and B. Nauta, *Time-interleaved Analog-to-Digital Converters*, Springer, 2011.
- [15] B. Razavi, "Design of Sample-and-Hold Amplifiers for High-Speed Low-Voltage A/D Converters," *CICC*, pp. 59-66, 1997.
- [16] P. J. Pupalais, B. Yamrone, R. Delbue, A. S. Khanna, K. Doshi, B. Bhat and A. Sureka, "Technologies for Very High Bandwidth Real-time Oscilloscopes," *IEEE BIPOLAR / BICMOS CIRCUITS AND TECHNOLOGY MEETING*, p. 9, 2014.
- [17] Analog Devices, "AD9578 Dual Precision PLL Synthesizer," [Online]. Available: <http://www.analog.com/en/products/clock-and-timing/clock-generation-distribution/ad9578.html>. [Accessed 6 12 2016].
- [18] E. Yüce, "All-optical switching of a microcavity repeated at terahertz rates," *Optics Letters*, pp. 374-376, 2013.
- [19] P. J. Pupalais, "Random Interleaved Sampling," [Online]. Available: http://cdn.teledynelecroy.com/files/whitepapers/wp_ris_102203.pdf. [Accessed 3 11 2016].
- [20] T.-i. Otsuji, "A Picosecond Accuracy, 700MHz Range, Si-Bipolar Time Interval Counter LSI," *IEEE Journal of Solid-State Circuits*, pp. 941-947, 1993.
- [21] J. Kalisz, "Review of Methods for time interval measurements with picosecond resolution," *Metrologia*, pp. 17-32, 2003.
- [22] M. Zieliński, "Architecture of the multi-tap-delay-line time-interval measurement module implemented in FPGA device," *ACTA IMEKO*, pp. 77-81, 2015.
- [23] K. Poulton, "Architectures and Issues for Gigasample/second ADCs," *AACD*, pp. 1-18, 2006.
- [24] M. Patel, I. Darwazeh and J. O'Reilly, "Bandpass sampling for software radio receivers, and the effect of oversampling on aperture jitter," in *Vehicular Technology Conference*,

Birmingham, 2002.

- [25] W. Kester, "ADC Architectures III: Sigma-Delta ADC Basics," 1 10 2008. [Online]. Available: <http://www.analog.com/media/en/training-seminars/tutorials/MT-022.pdf>. [Accessed 5 11 2016].
- [26] F. de Jager, "Delta Modulation: A Method of PCM Transmission Using the One Unit Code," *Philips Research Reports*, p. Vol 7, 7 1952.
- [27] C. Culter. United States of America Patent 2,927,962, 1960.
- [28] Analog Devices, "ADCMP600/ADCMP601/ADCMP602," 14 February 2017. [Online]. Available: http://www.analog.com/media/en/technical-documentation/data-sheets/ADCMP600_601_602.pdf.
- [29] J. E. Abate, "Linear and Adaptive Delta Modulation," *Proceedings of the IEEE*, pp. 298-308, No 3 Vol 55 1967.
- [30] "Idle tone behavior in Sigma Delta Modulation," in *Audio Engineering Society*, Vienna, 2007.
- [31] Analog Devices, "ADCMP572/ADCMP573 Ultrafast 3.3 V/5 V Single-Supply SiGe Comparators," [Online]. Available: http://www.analog.com/media/en/technical-documentation/data-sheets/ADCMP572_573.pdf.
- [32] Cypress Semiconductor Corporation, "PSoC® 5LP: CY8C58LP Family Datasheet," San Jose, CA.
- [33] Hewlett Packard, "Application Note 918: Pulse and Waveform Generation with Step Recovery Diodes," Palo Alto, 84.
- [34] D. Bowman, "Source 2," 2005. [Online]. Available: <http://www.g0mrf.com/source2.htm>.
- [35] B. Razavi, "A Study of Injection Locking and Pulling in Oscillators," *IEEE JOURNAL OF SOLID-STATE CIRCUITS*, pp. 1415-1424, 2004.
- [36] Agilent, "Basics of Measuring the Dielectric Properties of Materials," 2005.
- [37] Agilent, "Microwave Dielectric Spectroscopy Workshop," 2004.
- [38] A. Kadambi, H. Zhao, B. Shi and R. Raskar, "Occluded Imaging with Time-of-Flight Sensors," *ACM Trans. Graph*, vol. 35, no. 2, p. 12, 2016.

- [39] Broadcom, "Broadcom," 2013. [Online]. Available: <https://docs.broadcom.com/docs/AV02-2701EN>.
- [40] D. Rytting, "ARFTG 50 Year Network Analyzer History," in *ARFTG Microwave Measurement Conference*, Atlanta, GA, 2008.
- [41] S. O. Nelson, "Agricultural Applications of Dielectric Spectroscopy," *Journal of Microwave Power and Electromagnetic Energy*, pp. 75-85, 2004.
- [42] USB IF, "USB 3.0 SuperSpeed Equalizer Design Guidelines," 10 6 2011. [Online]. Available: http://www.usb.org/developers/docs/whitepapers/USB_Superspeed_Equalizer_Design_Guidelines_2011-06-10.pdf.
- [43] J. Williams, *Analog Circuit Design: Art, Science, and Personalities*, Butterworth-Heinemann, 1991.
- [44] Avago Technologies, "HSMP-382x, 482x Surface Mount RF PIN Switch and Limiter Diodes," 9 April 2017. [Online]. Available: <https://www.broadcom.com/products/wireless/diodes/pin/hsmp-3822>.
- [45] USB IF, "Universal Serial Bus 3.1 Specification," 2013.
- [46] D. L. J. W. J. D. N. M. D. P. S. J. H. B. H. S. H. C. P. Dymond, "Reduction of oscillations in a GaN bridge leg using active gate driving with sub-ns resolution, arbitrary gate-impedance patterns," *IEEE ECCE*, 2016.
- [47] G. C. J. C. E. D. R. D. B. B. d. R. A. P. M. J.-M. G. a. W. L. V. Emre Yüce, "All-optical switching of a microcavity repeated at terahertz rates," *Optics Letters*, pp. 374-376, 2013.