

Moore or Less? A Critical Comparison of Bandwidth vs. Resolution

Benjamin Egg (Cubic Defense Applications & SDSU, San Diego, CA; benjamin.egg@gmail.com); fred harris (San Diego State University, San Diego, CA; fred.harris@sdsu.edu); Chris Dick (Xilinx, Inc., San Jose, CA; chris.dick@xilinx.com)

Abstract

Moore's law is evident in the fantastic success of the microprocessor, the variety of inexpensive RF devices, and wide range of tiny radios and phones now available. On the other hand, mixed signal devices, like data converters seem to have fallen behind the *Moore-like* growth of their semiconductor counterparts. An analysis of data converter performance, considering both bandwidth and resolution, indicates radio architectures are historically uncorrelated with data converters of the same time period. Data converters, which are generally fabricated from two distinct semiconductor materials, diverge from Moore's law; however, we illustrate a clear connection between particular engineering personalities and data converter trends. Two generational and professional engineering styles are discussed as *b-type* and *r-type* personalities—bandwidth and resolution. Bandwidth and resolution are controllable variables from a semiconductor fabrication stance, and bound a link's theoretical information capacity—Shannon's theory. Analog-to-digital and digital-to-analog converters (ADC and DAC, respectively) are judged by bits of resolution and sample rate. Taken individually, digital logic and analog RF have experienced exceptional growth, but together have underperformed. The performance predictions of Moore's law state that the density of transistors in an integrated circuit doubles every 18 months, which inversely relates to speeds; however, data converters haven't adhered to that log-linear performance—converters appears uncorrelated at first glance. Recent developments in parallel signal processing devices, such as the Virtex series FPGA from Xilinx, have stimulated the converter market to make up for lost time, but stylistic preferences are strong. A generational and occupational model reveals allegiances on both sides of the resolution-bandwidth divide. Ultimately, we clarify and define these biases, reveal their roots, and confirm their impact.

I. Introduction

Data converter performance lags behind the fantastic advancements of its two subsystems—analogue RF and high

speed digital semiconductors. Moore accurately predicted the trends in transistor densities in 1965 [1]; nevertheless, composite mixed signal semiconductor technology underperforms. Trends in RF integrated circuits reveal that the analog portion in a converter historically trails available RF specific devices, and digital circuitry reveals a similar unbalance. The lagging performance of mixed signal devices is analyzed from a unique angle—engineering trends forged from the microprocessor success.

The versatility of application and predictable performance growth of the microprocessor, and subsequent application specific DSP, have broadly influenced the data converter market and thus performance—independent of converter's unique technological potentials and/or limitations. More recent parallel processing devices, like the field programmable gate array (FPGA), have influenced the ADC and DAC market with evident converter performance spikes and performance increases.

The balance of this paper looks at: a quick background on transistor densities, RF semiconductors trends and related performance benchmarks to quantify our assumptions, section II. Then in section III, microprocessor developments are compared against data converter trends to graphically reveal their relationship. Section IV discusses current state-of-the-art converters, and the field programmable gate array's (FPGA) role in recent data converter performance gains. A very general personality profile is formed in section V, which will aide the reader to identify his/her bias toward R or B type decisions—resolution or bandwidth. We look at a case study of two hardened R and B type engineers and discusses their approach to solving the same technical problem, from very different angles. Finally, we briefly review the various conclusions and present the unexpected results.

II. Digital Logic and RF semiconductors

Transistor densities and clock speeds in digital logic semiconductors have increased year to year as predicted by Moore; similarly, RF silicon devices have decreased in size and price, while increasing in frequencies, albeit at a slower

pace. Not unexpectedly, combining two unique and dissimilar semiconductor technologies— analog and digital, results in a slower development cycles. This is a reasonable result since research monies advancing digital logic, and similar investments into RF semiconductors, has provided healthy returns on investment. Unfortunately, converters have gone casually along the middle-ground between its analog and digital counterparts. The potential performance of converters is put into perspective by developing weighted formulas to quantize the individual performance of digital and analog semiconductors, independently. The figure of merit for RF integrated circuits (ICs), analyzes the Noise Figure (F), Bandwidth (BW), and power consumption (P) to quantize performance. We combine these values into a single figure that balances contribution of each.

Noise Figure: is the ratio of input SNR over output SNR. A noiseless device has a 0 dB noise figure, while a good quality, low noise amplifiers has a 2 dB noise figure.

$$F = \frac{SNR_{in}}{SNR_{out}} \geq 0dB \quad (2.1)$$

Bandwidth: is the range from lowest to highest frequencies for specified operation, where the input to output signal undergoes no (or minimal, i.e. -3 dB) distortion.

$$BW = 20 \cdot \log(f_{upper} - f_{lower}) \quad (2.2)$$

Combining these equations we arrive at a figure of merit:

$$MF = \frac{BW}{F} \quad (2.3)$$

The figure of merit, MF, helps us develop a general formula for a hard to parameterize analog technology.




Figure 2.1 Amplifier bandwidth (y-MHz) over time(x-year).

Digital electronics are measured against only toggle rates. While power consumption and cost are significant, they distract from our goal of identifying potential growth due available technology. Core clock frequency establishes a logic fabric boundary, and the log function allows for clearer graphical illustration.

$$f_{clk} = 20 \cdot \log(f) \quad (2.4)$$

Both figures 2.1 and 2.2 use a log scale for the y axis, illustrating analog bandwidth and digital logic toggle frequencies, respectively [2]. As an example, in 1976, Analog bandwidths for amplifiers were greater than 50 MHz, and digital logic could toggle at 25 MHz, or 1968 logic toggled at 500 MHz. A Nyquist sample rate for a 50 MHz analog signal is 100 MHz; therefore, it is technologically feasible that a 100 MHz converter could have been developed around 1976.




Figure 2.2 Logic toggle rate, in MHz per year.

Without diving too deep into widely available details of RF amplifier performance and digital logic speeds, we can get the general idea that of analog bandwidth, and digital frequencies grow with time. IT is worth noting that some milestones in analog or digital technology are reached early, while not cost effective or reliable, so they have not been considered, at the authors' discretion.

III. Microprocessor – Converter Correlations

Serial processing machines like the microprocessor and DSP have helped data converters achieve early technological milestones, while unintentionally stifling a 'Moore's law' like logarithmic growth later on. The enemy

of *great is good enough*. Specifically, ADC and DAC performance was furthered as a result of the fantastic developments in semiconductor technology fueled by the microprocessors success.

When the 4004 Microprocessor hit the market in 1969 with a 4 bit bus and 108 KHz clock rate, converters were uniquely military, hardwired to hardware, and substantially faster than 4004 could manage. With the 8080's release in 1974, the bus width increased to 8-bits and the operating frequency was 2 MHz. The converter market was still incompatible with the microprocessor. The leading converter in 1974 was the AD7570 with 10-bit resolution and 50 MHz conversion rate. Clearly, the bus widths were incompatible and the converters operating frequency was 25 times faster than the processors'. We see that data converters had too much resolution and too high data rate for the general microprocessor to manage.




Figure 3.1a Microprocessor CPU speeds follow a log-linear curve, as predicted by Moore.

The famous 8086 microprocessor was released and widely adopted in 1978, sporting a 16 bit bus and 4.47 MHz clock. Bus resolution was now sufficient, and many lower frequency converter operations became manageable. The 8086's was widely accepted and software tools were quickly developed around the General Purpose Processor's (GPP) success. Consequently, computer aided design, and automated manufacturing were facilitated and led to rapid developments that further increased clock speeds and transistor densities. By 1985 the bus widths had expanded to 32-bits and processor clocks reached 33 MHz. A short time later in 1989, the clock frequencies surpassed 100 MHz. Following this log-linear growth trend to recent technology,

the Pentium 4 operates with a 32 bit bus at 3 GHz internal clock rate.

We should also mention that around 1982, while the speed of the GPP was 4.42 MHz (8086), the application specific digital signal processing (DSP) chip entered the data processing market. Differing from the microprocessor, which could perform a variety of General Purpose Processes (GPP), the DSP was specifically designed to multiply and accumulate (MAC), which is the main function in DSP. Figure 3.1b illustrates the performance of the DSP ASIC from 1982 to 2003 [3]. While the GPP required several operations (4-16) to access data and cached coefficients, then multiply them and accumulate (MAC) the results, the DSP's pipelined architecture performed the MAC operation in only 1 or 2 clock cycles—a 4 to 16 fold advantage.




Figure 3.1b DSP benchmarks from 1982 to 2002, are very similar to Intel GPP performances, divided by 4-16.

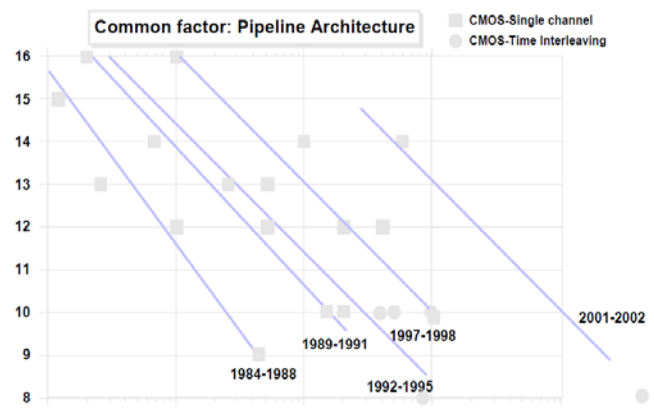


Figure 3.2 Converter trends spanning 1984-2002. Y-axis is resolution in bits; x-axis is sample rate in MHz.

As previously mentioned, early converter bandwidth's far exceeded the GPP throughput capacity. Even 2nd generation devices like the 8086, with a 16 bit bus and core frequency of 4.47 MHz, trailed the concurrent MOD-815 and the MOD-1020 ADCs with their 15 MHz and 20 MHz data rates, respectively. Considering the processing cycles required to compute even a minor filtering operation, the 8086 was 4 to 64 times deficient in processing capability. Figure 3.2 illustrates some converter development trends regarding resolution and sample rates from 1984 to 2002. Superimposing analog, digital, CPU and converter peak performances on a log axis graph, figure 3.3 helps contrast the technological advancements. The graphical representations are: analog bandwidth—blue diamonds; digital logic—magenta dots; converter performances—red line with black circles; and DSP and microprocessor performance—green line with yellow triangles.

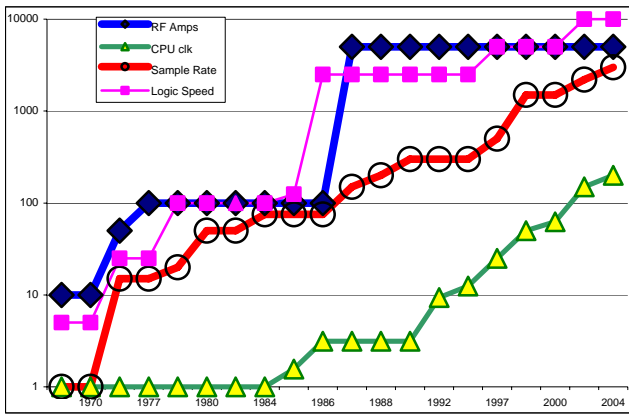


Figure 3.3 Analog, Digital, CPU clock speeds and Converter speeds, 1968 to 2004.

An additional 3rd dimension helps conceptualize the resolution and sample rate relationship over time. We gain insight into converter trends from figure 3.4's scatter like illustrations, color coded by decade (70's-blue, 80's-red, 90's-magenta, 2000's-cyan). In addition, the 3D points are collapsed to the back wall (removing resolution information), and the DSP performance curve is superimposed in bold-green.

In a final push for clarity, a log scale comparison of DSP performance (and similar GPP, CPU performance) versus converter bandwidths is illustrated in figure 35. Particularly clear is the performance lag of 10-100 times or

10-15 years. This performance lag illustrates the historical disconnect of processors and converters.

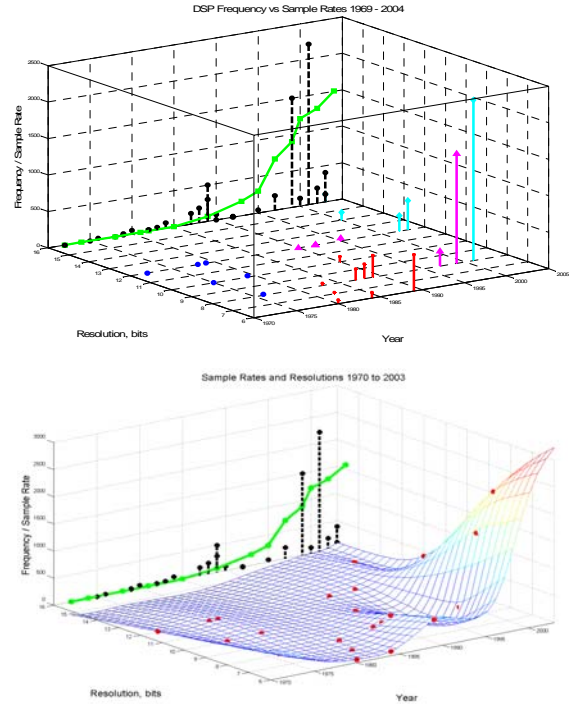


Figure 3.4a and 3.4b illustrates resolution and sample rate versus year, with a collapsed to 2D view on the back side, superimposed by the DSP performance (green).

A quick tangent is needed here to formulate a computational rule of thumb for comparative analysis. We will assume that our signal processing task, against which all are processes will be compared, is a 16 coefficient FIR filter. The appropriate equation is

$$f_s \leq \frac{f_{clk}}{C_{MAC}} \cdot \frac{1}{N_{taps}} \quad (3.1)$$

Where f_{clk} is the GPP clock rate, C_{mac} is the Computational load for a single Multiply and accumulate operation (MAC). N_{taps} is the number of taps, or the number of MACs and f_s is the maximum converter data rate or sample rate.

$$f_s \leq \frac{4.47 \text{ MHz}}{8} \cdot \frac{1}{16} = 34.9 \text{ KHz} \quad (3.2)$$

The result in 3.2 indicates that the 8086, with a 4.47 MHz clock can process the prototype filtering operation with a maximum ADC throughput of 34.9 KHz. That rate is sufficient to handle audio signal processing where the analog bandwidth is less than 15 KHz. The results of the GPP / DSP curve in figure 3.5 is based on a 16 tap filter with C = 1.

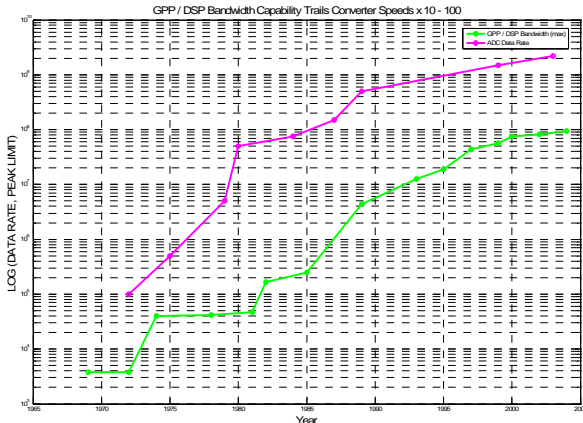


Figure 3.5 clearly shows that converter bandwidths are 10-50 times greater than historically available DSPs, or that converters lead by 10-15 years.

Data converters appear to have been neglected, which may not be a logic assumption at first. While data converters appear to lead processors due to bandwidth, the reality is that the converter development is driven by the micro processor and DSP performance curve. This is evident by the very similar growth of the data converter to the microprocessor (fig 3.5), while the true indicators of converter performance—digital and analog benchmarks, don't have correlated performances (fig 3.3).

IV. Modern Converters and the FPGA

Over the last several years, data converters seem to have experienced performance leaps in both resolution and bandwidth; however, high resolution converters have been available from 12-16 bits since the 16-bit bus of the 8086 microprocessor in 1978. Throughput, or sample rate has been the historical obstacle, and recent triumph for data converters. Even the industries' fastest GPPs and DSPs, available today with 1 GHz processing clocks, actually have bus throughput rates an order of magnitude slower. Giga-sample data rates are not feasible; further more, when

filtering and other signal processing operations are involved, the throughput is reduces even more. Cutting edge serial processors are limited to less than 10-100 MHz throughput [4, 5]

The field programmable gate array (FPGA) has enabled growth and influenced converter market and technology advancements since the early 1990's. However, the ushered in faster converter designs when they took on large signal processing jobs, previously performed by the GPP or DSP. Particularly influential was the Virtex series FPGA's massive logic resources, and subsequent embedded multipliers, in 1998-2000. A quick graphic shows FPGA development and converter growth since 1998, in figure 4.1.




Figure 4.1 illustrates available converter rate, MHz (background) and FPGA I/O rates, MHz (foreground).

Prior to 1999, CPU and DSP bandwidths were limited to 100-300 MHz; furthermore, few to no additional cycles remained at those rates to perform signal processing operations. Figure 4.1 illustrates the IO toggle rates of the FPGA, and recent data converters [7, 8]. Most impressive about the FPGA, and a primary reason for its success, and the spike in converter sample rates, was the immense processing capability of the FPGA. The Xilinx Virtex-4, SX55 FPGA contains 512 18x18 MAC engines, operating at 400+ MHz. That's over 200 Billion MAC operations per second, enabling generous real-time digital signal processing [6]. It is important to note that all of the converters listed have internal demultiplexers to divide the throughput by 2 or 4 (channels); see table 4.1 below.

Year	FPGA	IO Speed	Converter	IO rate	Channels
1999	Virtex-E	622	MAX106	1000	2
2001	Virtex-II	750	Max108	1500	2
2003	Virtex-2 pro	840	AT84AS003VTP	2000	4
2005	Virtex-4	1000	ADC081500	1500	2
2006	Virtex-5	1250	ADC083000	3000	4

Table 4.1 FPGA I/O rates and ADC throughput.

V. Engineer Profiling

Resolution and bandwidth biases stifle innovation. We are going to breeze through an actual development project approached from hardened R-type and B-type engineers, who are tasked to solve the exact same problem, but come up with two different solutions to the same problem. Furthermore, their preferences are entrenched, resolution on one side, and bandwidth on the other. Furthermore, they seem unaware of any impediments that might limit their design approaches to the exact same problem. The goal here is to show the reader both sides of the divide, and facilitate clear understanding of where preferences come from, and the cost of continued ownership. Serial processing and parallel processing preferences heavily influence engineering dispositions.

The Specs:

Improvised explosive devices (IEDs) are often detonated via remote RF triggers, with RF carriers spanning a large spectrum with unpredictable power levels and modulations. Triggers can range from a simple car alarm remote, to a radio controller for a model car or plane, a modified cell phone, unlicensed band ‘walkie-talkies’, etc. Without radiating a massive amount of jamming power over the whole 2.5 GHz spectrum, prohibiting all communications--friend and foe, the system should detect an ‘unfriendly signal’, determine its spectral residence and jam it before a successful message to detonate is received. The general requirements are:

- | | |
|--|--|
| a) Bandwidth | DC-2.5 GHz |
| b) Modulation (unfriendly) | Unknown/all |
| c) Power (unfriendly) | Unknown |
| d) Other signal Powers | Radio stations, own radios, cellular phones. |
| e) Time to detect and jam unknown signal | Minimum |
| f) Cost | Generous |
| g) Development and Delivery | 5 Months |

R-Type Solution:

The challenge, as seen by the resolution-leaning engineer, is to detect the ‘unfriendly signal’ when it is just developing, or at the nascent minute power level as the trigger of a remote detonating device is initializing, and the Power Amplifier and message that will trigger an IED is beginning to form. The keys to success are:

- Resolution: detects the signal at very low power, allowing more time to transmit the jammer.
- Dynamic Range: More resilient to interference like radio stations and personal radios, avoiding saturation.
- FFT time: A DSP capable of analyzing a spectral region and providing a result very fast.
- VCO: A fast settling VCO/PLL and Mixer to tune the narrow bandwidth through the entire spectrum to analyze and detect new signals in minimal time.

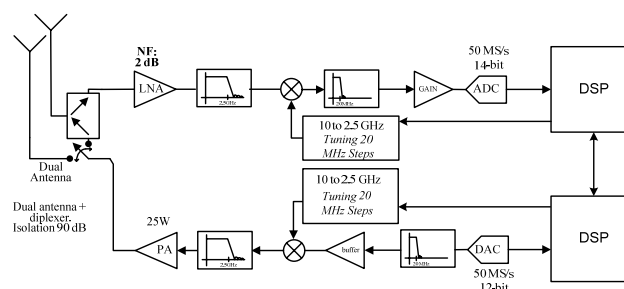


Figure 5.1 The resolution heavy approach tunes the VCO through 20 MHz steps, performs the FFT in a DSP device and can quickly tune to and jam any spectral region.

This solution analyzes a 20 MHz window by performing the FFT and comparing the spectral result with a previous result, searching for a ‘new’ signal. If no differences are detected, the VCO step to the next 20 MHz increment and performs the FFT and compare operation again. If a change is detected, the Transmit VCO tunes to that frequency, if not already there, and the PA is ramped to full power wherein an appropriately offensive jammer is unleashed to distort any IED detonation transmissions.

One of the clear benefits to this design is the wide dynamic range—within the 20 MHz band, the 14 bit converter can detect very small signals, early. Also, the 20 MHz bandwidth is narrow enough to filter out any large, local signals that might violate the receiver’s dynamic range. Overall this is a great and ingenious solution.

B-type solution:

The challenge as seen by a bandwidth-focused engineer is to detect the ‘unfriendly signal’ quickly, by analyzing almost the whole spectrum, simultaneously. This approach cannot detect the ultra-weak signals, but can detect almost the whole bandwidth simultaneously. The resolution approach has a 2.5 mS search time (Table 5.1), which is a worrisome amount of time in which a signal could appear and disappear and never be detected. The bandwidth approach, while suffering from less dynamic range, seems to offer a statistical advantage, by being in the right place at the right time, but with less sensitivity.

The keys to success are:

- a) Bandwidth: tuning to many narrow bandwidths consumes time. Analyze 700 MHz regions.
- b) FFT time: The parallel resources of the FPGA are unmatched in computation time.
- c) Fixed VCO: A fixed VCO mixes the spectrum from 1.75 GHz to baseband.
- d) Baseband Low pass filter 0-700 MHz
- e) Band pass filter 700-1400 MHz

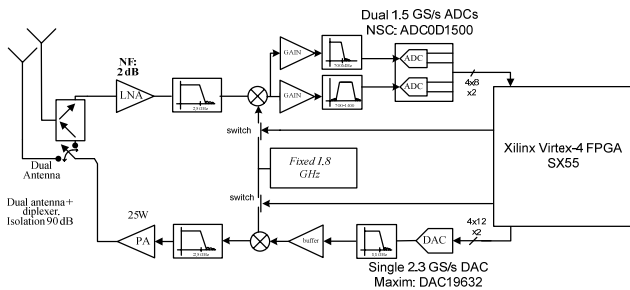


Figure 5.2 Bandwidth approach using fixed mixer frequency at 1.8 GHz, dual 1.5 GS/S ADCs and single 2.3 GS/S DAC.

The bandwidth focused design handles massive amounts of spectrum, over 1 GHz of bandwidth simultaneously. Lower 8-bit resolution of the high speed ADCs is compensated by quicker detection times since tuning is not necessary. Table 5.1 lists some details comparing each design.

The data and examples show that both engineering solutions are reasonable--there are several ways to solve the same problem with resolution and bandwidth as variables. Table 5.1 indicates that the resolution technique may be at a disadvantage, but considering the benefit of detecting the signal earlier via better dynamic range, that is debatable.

Furthermore, what is the value of the lower probability of being saturated by local signals (radio stations, own radios, etc)? The provided specification don't provide trigger transient or start up time, which leaves a lot of open design questions. Often the decision to choose one architecture or the other is developed by past preferences and ‘comfort-zones’ rather than available technology. In this case, both engineers straight-away went to their comfort zones—bits or bandwidth.

	Resolution	Bandwidth
Dynamic range	14 bits	8 bits
Sample Rate	50 MHz	2 x 1500 MHz
FFTs	1	2
FFT Resolution	49 KHz	185 KHz
FFT Length	1024	2 x 8192
Processing Time	20uS*	21uS
Processor	DSP TI	FPGA Xilinx SX55
Overall Search	2.5 mS	.042 mS
Overall Jam	2.55 mS	.092 mS

Table 5.1 Resolution and Bandwidth parameters. *estimate, PENTEK.

Conclusion:

Moore’s law is evident and validated in the microprocessor’s fantastic success. Surprisingly, converters *are* correlated with Moore’s law, but in a distant manner. Converter technology appears to be leading the microprocessor curve by some 10-15 years, as seen in figure 3.5. But, the term leading is erroneous. The technology was consistently available for faster converters in any reported year, much faster! But the success of the microprocessor obliged the converter to ‘wait-around’ for the masses to catch up. In this high speed converter study, and general conversation, I have heard asked many times, “what are you going to do with all that bandwidth anyway?” That query explains a lot about the styles, or preferences that have influenced a modest converter market, until the recent push generated by the powerful FPGA. If today’s engineers, surrounded by all the evidence of Moore’s law, query about the need for a faster converter, then it’s a fair assumption that the same question has caused curious bandwidth hungry engineers to stop searching for that extra MHz. So, Moore predicted the microprocessor’s success, but digital and analog logic have lead the way—on their own. Converters

have 'lead from behind' as they are correlated with the microprocessor, and thus adhere to Moore's law as a function of microprocessor performance. Ultimately, they have the potential to perform much faster, as a function of current semiconductor technological limits—2 to 10 times faster.

References

- [1] Gordon Moore, <http://www.intel.com/technology/mooreslaw/index.htm>.
- [2] Walt Kester, "Analog-Digital Conversion", Analog Devices, Inc. www.analog.com, 2004.
- [3] FFT, <http://www.xilinx.com/ipcenter/catalog/logiccore/docs/xfft.pdf>.
- [4] TI FFT, <http://www.pentek.com/deliver/TechDoc.cfm?Filename=wp021802.pdf>.
- [5] TI DSP, Table 2., <http://focus.ti.com/lit/an/spraa03/spraa03.pdf>.
- [6] Xilinx Virtex-4, <http://direct.xilinx.com/bvdocs/publications/ds112.pdf>.
- [7] Virtex Series, http://www.xilinx.com/products/silicon_solutions/fpgas/virtex/index.htm.
- [8] See: www.national.com, www.atmel.com, www.maximic.com.