Fast Digitization and Digital Receiver Technology

Ralph C. Kimball

Echotek Corporation, 555 Sparkman Drive, Suite 400, Huntsville, AL 35816 USA

Abstract. The potentially lucrative wireless market has led to technological advances in mixed signal devices such as high speed, high resolution A/D and D/A converters. This same market has also driven the development of high performance multi-channel digital receiver and digital transmitter ICs. Similarly, advances in semiconductor processes, coupled with the need for reduced time-to-market, has led to the development of large, enhanced performance, in-circuit programmable logic devices. A review of the key characteristics of these mixed-signal, signal processing and programmable logic devices is presented. The application of these devices and technologies to the instrumentation of Accelerators and Storage Rings is discussed and presented by way of examples. Issues relating to the requirements associated with real-time processing, I/O throughput, reconfigurability, reliability, maintainability and packaging requirements are also addressed.

INTRODUCTION

The instrumentation requirements associated with Linear Accelerators and Storage Rings involves a broad mixture of analog and digital technologies. While the process monitoring sensors and control loop feedback driven correction elements are inherently analog, the data analysis and control loop response is increasingly dependent upon digital signal processing (DSP). This architecture leads to the need for an Analog-to-Digital (A/D) converter (or ADC) to translate the input analog sensor data into a digital data stream for processing, and for a Digital-to-Analog (D/A) converter (or DAC) to translate the digital feedback corrections to analog, thereby completing the control loop. Clearly, the characteristics of these mixed-signal devices are therefore critical to the performance assessment of the overall instrumentation subsystem.

Similarly, the computational and throughput requirements placed on instrumentation subsystem digital processing elements must be well understood. This is necessary in order to properly define DSP architectures that will provide sufficient resources and flexibility at a reasonable cost. Consideration must also be given to both centralized and distributed processing configurations in order to embrace the full breadth of instrumentation requirements, from the subsystem level to the total system level. Also, early obsolescence due to the rapid rate of advancements that are associated with today's digital technology is an ever-

CP648, *Beam Instrumentation Workshop 2002: Tenth Workshop*, edited by G. A. Smith and T. Russo

growing concern to the instrumentation engineer. Accordingly, product roadmaps and reliance on internationally recognized "Standards" such as ANSI, IEEE, etc. must be entered into the instrumentation design equation to mitigate the probability or impact of such events.

In view of the fact that the current trend, for both the mixed-signal and DSP elements of the above instrumentation architecture, is toward reduced geometry with increased performance, additional consideration must be given to the "packaging" of these devices. Since shrunk semiconductor process geometries result in reduced power rail voltages, most of the newer devices require multiple different DC voltages to power the analog section, the "core", and the input/output (I/O) section. Also, the reduced process geometries support smaller physical packaging with fine-pitch pins or balls, and can operate at significantly higher clock rates. These characteristics translate into the need for fine-geometry, multi-layer, stripline Printed Circuit Board (PCB) designs with special consideration given to manufacturability and testability of the end product. Furthermore, sufficient documentation must be provided for each element of the design to further ensure the utility and longevity of that product design.

A/D CONVERTER ARCHITECTURES

A number of different A/D architectures have evolved over the years, primarily driven by application requirements, semiconductor technology and cost. In general, the most popular A/Ds of today can be categorized as being a Flash, Pipelined, Successive Approximation, or Sigma-Delta design. The term "sampling A/D" is commonly used to indicate that the A/D front-end has a Sample-and-Hold (S/H) circuit to sample (or track) the input signal during one phase of the encode clock and to hold that sampled level throughout the other phase of the encode clock. Sampling A/Ds that are designed for digitizing Intermediate Frequency (IF) inputs typically have a minimum encode rate specification. This is primarily due to optimizing the S/H circuit and "droop" will result from lower frequency encode rates, owing to charge "bleed-off" in the hold circuit. Similarly, the term "under sampling" (also referred to as "super Nyquist" sampling) is used to categorize the use of an A/D to also serve as a down-conversion mixer. For a given encode clock frequency F_s, the A/D will alias the digitized signal down in bands, or "Nyquist zones" (i.e., $F_s/2$ bands), with spectral inversion occurring for the second Nyquist zone ($F_s/2$ to F_s), the fourth Nyquist zone ($3*F_s/2$ to $2*F_s$), etc. While most A/Ds are optimized for first Nyquist zone applications (i.e., DC to $F_s/2$), a number of the newer A/D converters are targeted at higher Nyquist zone applications to permit direct digitization of common communication system IFs.

Flash A/Ds utilize 2^n -1 comparators to produce an n-bit wide output in a single clock cycle. These devices are typically very fast and limited to about 8 to 10 bits, owing to the need for large numbers of closely matched comparators and the associated input loading. While Flash A/Ds were often used without a S/H, or the

S/H was provided externally, the current trend is to include very broad bandwidth S/H circuits integrated into the Flash A/D front-end to improve performance and to permit under-sampling high frequency inputs. Today, Flash A/Ds are typically employed in high input frequency digitization applications with clock rates from 500 MHz to >1.5 GHz, and can also be typified as being physically large (i.e., >1 square inch) and having relatively high power dissipation (i.e., 4 to 5 watts).

The Pipelined architecture takes the Flash architecture a step further by dividing the digitization process into a series of pipelined stages. Each stage employs an input S/H, an m-bit A/D, a D/A to convert the A/D digital value of that stage to an analog voltage, and a difference circuit with gain that produces a residual output. This output represents the difference of the input signal to that stage with the D/A output of that stage. This "residual" is then input to the next stage's S/H and the process is repeated. Since each stage is pipelined, new digital output values will be available at the encode clock rate once the pipeline (number of stages) is filled. The sum of the bit-widths of the "sub-ranging" stages is typically a couple of bits wider than the actual output sample bit-width "n", and this bit-width reduction is accomplished by a pipelined digital "correction" logic section. In view of the above processing, this architecture is also sometimes referred to as a "Sub-Ranging" A/D and dominates the wireless A/D market. These devices are typically small (i.e., about 0.5 square inches), relatively low power (i.e., <1 to 1.5 watts), and currently support clock rates of >100MHz at a 14-bit resolution, to >200 MHz at a 12-bit resolution.

The Successive-Approximation A/D utilizes a single comparator in conjunction with a register-driven D/A and a difference circuit. For each new conversion, the register is loaded with a ½ of full-scale value and the comparator output then represents the Most Significant Bit (MSB). Based on this MSB value, the register is then updated by a value, appropriately signed and of a ¼ of full-scale value. The process is then repeated and this is done a total of n times, utilizing smaller binary weights in progression each time, as required to produce an n-bit output. This process therefore requires a full n-clocks to produce an n-bit output, rather than the single clock per output required by the Flash A/D, or the Pipelined A/D once the pipeline is filled. This implies that the output sample data rate for a Successive-Approximation A/D is, at most, 1/n that of either of the other two architectures. Also, it should be noted that S/H droop would be significant if the circuit was not designed for long hold times relative to the A/D clock rate. Consequently, these converters are not well adapted for higher input frequency analog signal digitization applications.

The Sigma-Delta A/D architecture is significantly different in that it employs a very simple front-end consisting of a difference circuit, an integrator followed by a comparator, and a 1-bit D/A. A digital filtering logic section then follows this front-end. The D/A output is differenced with the input being digitized, and this difference is then input to the comparator. Consequently, the comparator output consists of a string of ones and zeroes, at the clock rate, that are input to both the

D/A and the digital filter logic. This bit stream of ones and zeroes is digitally filtered and decimated to produce the n-bit output sample. The combination of a very high clock rate relative to the output sample rate, with digital filtering to shape the passband, makes this architecture desirable for low frequency, narrow bandwidth, high-resolution A/D applications.

A/D SPECIFICATIONS

The theoretical Signal-to-Noise Ratio (SNR or S/N) for an n-bit A/D is the ratio of the rms full-scale, digitally reconstructed, analog input "V", to its rms quantization error (i.e., V*LSB/√(12), where LSB is the Least Significant Bit. Since the rms value of the full-scale input is 0.5V/√(2), the theoretical maximum SNR = 2^n √(3)/√(2) = 1.225*2^n. Since SNR is typically specified in decibels, the numerically simple and therefore commonly used expression is SNR_{db} = 20Log(2^n √(3)/√(2)) = 20Log(2)*n + 20Log(√(3)/√(2)) = 6.021*n + 1.763 (dB). This is also the Dynamic Range (DR) for an A/D, since DR is defined as the ideal SNR for that A/D (i.e., quantization noise limited). The "data sheet" SNR specifications for an A/D are typically extrapolated values given as a function of input signal frequency at a specific level below full-scale and for a given A/D encode rate. Similarly, the Signal-to-Noise and Distortion (SINAD) ratio specifications simply include the first N harmonics of the Total Harmonic Distortion (THD), where N is normally specified on the data sheet. The Effective Number of Bits (ENOB) is simply the reversal of the SNR calculation but is typically based on SINAD to reflect harmonic distortion (i.e., ENOB (bits) = (SINAD-1.763)/6.021).

For IF-sampling A/Ds, Aperture Uncertainty or "jitter" in the S/H circuit contributes to the noise floor owing to sample-to-sample timing variations. The theoretical SNR for an A/D limited by aperture uncertainty is SNR = -20 Log(2π F_{in} T_{au}), where T_{au} is the aperture uncertainty in seconds rms. Obviously, the effects of this jitter are more pronounced with increasing input signal frequency due to increasing dV/dt. The aperture uncertainty specification for a sampling A/D is commonly given in picoseconds rms. Typical values for examples of today's high speed, high resolution sampling A/Ds are <0.25 ps rms for the Analog Devices AD9430 (12-bits at a >200 MHz encode rate and a 700 MHz input bandwidth) and <0.1 ps rms for the Analog Devices AD6645 (14-bits at a >100 MHz encode rate and a 270 MHz input bandwidth). While these aperture uncertainties seem very small, it should be noted that, for an A/D like the AD9430 that could be used to directly digitize the typical RFs associated with Accelerators and Storage Rings, the 0.25 ps rms specification translates into a theoretical SNR of 96 dB for a 10 MHz input, to 76 dB for a 100 MHz input, and to 62 dB for a 500 MHz input. It should also be noted that the S/H associated with a sampling A/D acts essentially like a mixer in that the input signal is multiplied by the encode clock. Consequently, the above noted performance is further degraded by the presence of

jitter on the encode clock. For this reason, it is imperative that the encode clock be generated by a very low phase noise source.

The Spurious-Free Dynamic Range (SFDR) of an A/D is probably the most frequently misunderstood of all A/D specifications and, as such, deserves special attention. This is sometimes due to confusing the determination of the SFDR for an A/D with that of an amplifier. An engineer at Avantek first introduced the concept of Intercept Point as an indicator of the SFDR of an amplifier in 1964. When an amplifier is operating in the linear range (i.e., below the 1 dB compression point), the levels of the spurious responses can be estimated accurately with a simple equation: SFDR = 2/3(P1-P0-10Log(BW)-NF). P1 is the input Intercept Point obtained by subtracting the amplifier gain from the output (third order) Intercept Point; P0 is the effective input noise power with no signal (i.e., -114 dBm for a 1 MHz bandwidth in a room temperature 50 ohm environment) ; BW is the system noise bandwidth in MHz; and NF is the amplifier noise figure in dB. For example, an amplifier with a third order Intercept Point of +38 dBm, a noise figure of 4 dB and a gain of 16 dB would have a 1 MHz bandwidth SFDR of: 2/3(+38-16-110) = -88 dBm, or SFDR = 88 dB. It should be noted that increasing the bandwidth will reduce the SFDR and decreasing the bandwidth will increase the SFDR.

The SFDR of an A/D is typically specified as the ratio of the rms signal amplitude to the rms value of the peak spurious spectral component, and the peak spurious component may or may not be a harmonic. SFDR may be specified as dBc (i.e., dB below the carrier), in which case it will degrade as the input signal level is lowered, or it may be specified as dBFS (i.e., always referred back to the converter full-scale input level). The important point here is that there is no mention of bandwidth in the A/D SFDR definition, unlike that for the amplifier SFDR definition. Consequently, the A/D SFDR specification, as given on an A/D's data sheet, is simply based on the highest spur level observed on a DFFT associated with a given data set or sets, for the specified input and operating conditions (i.e., for the specified Nyquist zone).

It should also be noted that the SFDR specification for an A/D could be significantly greater than the A/D's ideal SNR (i.e., DR). For example, the Analog Devices AD6644 is a 14-bit, 65 MSPS (Million Samples Per Second) A/D with a DR = 6.021 * 14 + 1.763 = 86.06 dB, but has single-tone and multi-tone SFDR specifications of >90 dB and >100 dB, respectively. This "phenomenon" is most easily explained by the concept of "processing gain". For today's highly linear A/Ds, like the AD6644, input signal variations that are less than the resolution (an LSB step) can be accurately extracted from the noise by digitally processing multiple samples. The signal of interest is essentially buried in the noise that is randomizing the A/D states, such that the periodicity of the signal (i.e., correlated signal) can be differentiated from the (uncorrelated) "white" noise.

It is for this reason that "dither" noise is sometimes added to the A/D input signal to improve the A/D's linearity, and therefore improve its SFDR. This is typically accomplished by adding white noise, band-limited to a few hundred kHz, to the

input signal to be digitized. This practice is primarily limited to digital receiver applications or other applications where sharp digital filtering of the digitized sample stream is accomplished by post-processing. This is due to the fact that this noise must be placed in a portion of the spectrum such that it can be subsequently removed by digital filtering. The minimum input power level for the noise being added is normally dependent on the specific A/D and its architecture. Dither is most frequently used with sub-ranging A/Ds and the minimum power level is generally just enough to cause the input signal to cross one or more sub-ranges, thereby mitigating localized nonlinearities. Dithering by adding white noise is most advantageous when dealing with relatively pure single-tone inputs, since multi-tone and noisy single-tone inputs are essentially self-dithering.

Processing gain is realized in the Frequency domain by the DFFT imparting a gain of 10Log(number of points) (dB) for complex samples, thereby dropping the noise floor in each frequency bin and exposing the low level correlated signals. Increasing the A/D sampling rate also lowers the noise floor, because the noise spreads out over more frequencies, while the total integrated noise remains constant. In the time domain, decimating digital filters can therefore impart a processing gain of up to 3 dB for each halving of the input noise bandwidth. This is accomplished by digitally lowpass filtering the input data sample stream at rate F_{in} and then reducing the output data sample rate to F_{out} (note: F_{out} must be consistent with the new output bandwidth sample rate requirement). The digital lowpass filter and data rate decimation process is generalized as yielding a S/N improvement equal to 10Log(decimation), but it should be noted that this assumes that the input noise bandwidth is being reduced proportional to the decimation factor. For example, the digital lowpass filter must have a (passband + transition band) ± 0.125 F_{in} to permit a decimated output sample rate $F_{out} = F_{in} / 4$, thereby reducing the noise bandwidth by 4 and realizing a processing gain (S/N improvement) of up to 6 dB.

While there are numerous other important A/D specifications, they will not be addressed here in the interest of brevity. Instead, the focus is on today's high speed, high-resolution A/Ds, since these devices are most applicable to direct digitization of the input RF and IF signals associated with Accelerators and Storage Rings. There is, however, one observation with regard to these devices that relates to the above discussion of noise floor and SFDR. These A/Ds can, in the absence of an active input signal, output digitized noise floor sample streams that contain just two or three distinct values. Accordingly, some correlation can exist within these digital sequences, and are manifested as "spurs" appearing on an otherwise white noise floor, when viewed on the associated DFFT. These spurs are easily identified since, inputting a signal that often may even be at a level below one LSB, will cause that signal to appear and the noise floor to then be white.

D/A CONVERTERS

It is apparent from the previous discussion of A/D converter architectures that D/As are an integral part of A/Ds, and critical to the performance characteristics of those devices. Additionally, the use of stand-alone D/As is required to complete the digital processing based control loop topology. Again, it is tacitly assumed that high-resolution devices utilized in relatively low frequency correction applications are well understood and that the emphasis is to be placed on high frequency Accelerator and Storage Ring analog control requirements.

Most of today's high speed, high resolution, D/As are current-based designs, whereby the digital input code is translated into a balanced differential output current-pair that is linearly related to that input code. The resolution of an n-bit D/A is then a function of the full-scale current setting (typically about 20 ma) divided by 2^n -1. A number of the newer high-resolution D/A designs employ multiple switched current segments, somewhat analogous to the sub-ranging A/D architecture, to improve linearity and increase dynamic performance. The differential current outputs can be converted to voltage outputs by resistive terminations if signal reconstruction down to DC is required, or more commonly for the IF applications being addressed, by transformers configured to match a given load impedance (normally 50 ohms) and to perform a differential to single-ended voltage conversion.

The "glitch" energy (i.e., narrow unwanted spikes on the output occurring at the update clock rate) used to be a D/A specification of major concern and designers often placed a S/H after the D/A to sample the output, after the glitch interval, and to hold that level during the next sample update. The more recent D/A designs can be typified as having very low glitch energy (i.e., <5 pV-s) and it is increasingly common for this specification to be omitted from D/A data sheets. This is due in part to the fact that these devices are typically being used in band-limited IF output applications whereby this relatively low glitch energy is well filtered. Accordingly, these device data sheets opt for dynamic performance characterization and emphasize SFDR, SNR, etc., much the same as for the high speed, high resolution A/Ds. However, it should be noted that, unlike for an A/D, the SFDR for a D/A is specified as the difference, in dB, between the rms amplitude of the output signal and the peak spurious signal over the specified bandwidth.

Since the D/A-output reconstructed waveforms consist of a series of discrete samples that are updated at the clock rate, F_{out}, the output analog response exhibits a sinx/x roll-off that, for the first Nyquist zone, will vary from 0 dB at DC, to -4 dB at $F_{out}/2$. Additionally, the output waveforms will be replicated in each successive Nyquist zone, with spectral inversion occurring for the second Nyquist zone, the fourth Nyquist zone, etc., as was the case for A/Ds. Unlike the A/D's, where the level of the replicated images was determined by the A/D's front-end analog input response, the sinx/x roll-off continues for the replicated D/A-produced images. Hence, the response will be -13 dB at $3F_{out}/2$, -18 dB at $5F_{out}/2$, etc. which explains

why D/As are normally associated with relatively low IF output frequencies. In fact, the maximum output frequency to be reconstructed by a D/A is commonly designed to be $<0.25F_{out}$ in order to keep the sinx/x roll-off to <0.22 dB.

In view of the fact that the digitally reconstructed output waveform sinx/x response is a function of the D/A output sample update rate, the D/A manufacturers have placed major emphasis on increasing the supported update rate, while maintaining good dynamic device performance characteristics. Low power, high resolution (12-bit to 16-bit) devices, such as the Analog Devices AD97xx family of D/As, are currently available with output sample update rates to 400 MSPS. However, these devices would impose serious I/O bandwidth requirements were it not for their inclusion, within the devices, of interpolating digital Finite Impulse Response (FIR) filters. This digital processing consists of increasing the sample clock rate and inserting samples of zero value between the valid input samples, lowpass filtering the resultant sample stream, and applying gain to correct for the drop in signal level owing to the inserted zero values.

For example, interpolation by 4 would increase the sample stream data rate "F" by 4 to F_{new} and insert 3 zero values between each original F sample rate input value. The new sample stream would then be lowpass filtered with a FIR filter design that provides a passband of DC to $<\pm 1/8F_{new}$ and a stopband from $\pm 1/8F_{new}$ to $\pm(F_{new}/2)$. Since the insertion of 3 zero values along with each original input value reduces the resultant output level to one fourth that of the original level, a gain of 4 would be applied to the output sample data stream. This is accomplished by shifting each binary sample value left two bits toward the MSB (i.e., multiplication by 4). These D/A devices are referred to as "interpolating" D/As, and this architecture is dominating the more recent D/A designs because it mitigates the I/O bandwidth requirements associated with achieving a flat output response. In fact, these designs, particularly when employing higher interpolation rates, even target applications at the common communications IF (70 MHz) market. This is possible because reconstructed output waveforms, with frequency content as high as 100 MHz, still fall within the 0.22 dB flat output response at D/A output update sample rates of <400 MHz.

DIGITAL RECEIVERS

The classic analog heterodyne (or super-heterodyne) receiver typically employs a low-noise amplifier followed by a mixer that is driven by a Local Oscillator (LO) to downconvert the input RF to a lower IF. The mixer is then typically followed by one or more analog filters to band-limit the IF output (i.e., eliminate the unwanted mixing product terms or "images"). To obtain a high SFDR, these mixers typically must have a high IP3 and therefore require high LO drive levels. Also, depending upon the conversion loss of the mixers and insertion loss of the filters, additional IF amplification stages may be required. The IF signal A(t) is then converted to Inphase (I) and Quadrature (Q) signals via a 90-degree (quadrature) hybrid (i.e.,

produces a 0 degree and a 90 degree phase-shifted output signal pair) and downconverted to baseband. The quadrature components are: I (t) = A (t) cosine (2πF (t)) and Q (t) = A (t) sine (2πF (t)), where F is the downconversion frequency. These quadrature signals are then Square-law (diode) detected and lowpass filtered. For single-sideband applications, the 90-degree hybrid and mixers are more commonly replaced by an I&Q demodulator. This device combines two balanced mixers and a 90-degree hybrid in a single device that permits down conversion with quadrature outputs that are selectable as either the upper-sideband, or the lower-sideband, with rejection of the carrier (LO) and the unselected sideband.

This conversion to quadrature signals is necessary because a Square-law detector is simply a device whose output voltage is proportional to the square of its input voltage and, therefore, contains no phase information. Consequently, two signals in quadrature are required, each Square-law detected, to produce a Baseband complex signal (i.e., "video output pair") that preserves both the amplitude and the phase information, as required for subsequent processing. When that subsequent processing involves digital demodulation or information processing, these quadrature Baseband signals are digitized by A/D-pairs to produce quadrature (I and Q) digital data streams, as required.

A digital receiver (DR) performs the exact same function, but does it digitally, and potentially, much better. In an ideal world, a DR would utilize a single A/D to directly digitize the input RF signal. In the real world, the A/D is utilized at the earliest possible point in the classic heterodyne receiver chain, typically at an IF stage. However, it should be noted that advancements in A/Ds, as discussed above, make the delineation of what is IF and what is RF fuzzy, at best. Consequently, let's assume that the RF associated with Accelerators and Storage Rings can be directly digitized by an appropriately selected high speed, high-resolution A/D.

Digitization Requirements

It is first necessary to realize that the A/D must be digitizing a band-limited signal so that out-of-band signals will not alias into the desired information bandwidth. Consequently, some means of active or passive bandpass filtering must exist within the input path prior to the A/D. Next, it must be assured that the input signal will not exceed the full-scale input to the A/D, or clipping and distortion will result. Also, the peak anticipated input signal level should approach the A/D's full-scale input level with acceptable "headroom", in order to utilize the full dynamic range of the A/D. This may or may not require gain. These A/Ds typically have differential inputs, high differential input impedance (i.e., >1 kohm), and an input DC offset requirement. In view of this, wideband transformers having turns-ratios of 1-to-1 or higher are commonly used to drive the A/D differential inputs and, therefore, some gain may be realized via the transformer turns-ratio selection. Any remaining gain requirements would then have to be met with external amplification.

Digital Mixing

The digital "real" data stream out of the A/D at a sample rate F_s is directly converted to a Baseband "complex" quadrature data stream via a digital downconversion process (i.e., quadrature hybrid and analog mixers, or I&Q demodulator, equivalent). In the past, this downconversion process often used a digital processing "trick", owing to logic resource and speed limitations. Note that if the desired downconversion frequency is $-F_s/4$ (i.e., centering the desired frequency span at DC), the mixing process only requires four digital samples at that downconversion frequency. If these digital samples corresponded to the phase points (0, -90, -180, -270) degrees (i.e., negative frequency rotation sequence), the resultant cosine sequence would be (1, 0, -1, 0), and the associated sine sequence would be (0, -1, 0, 1). Therefore, only the sign of the input data must be changed periodically to generate the desired I and Q output data streams and, since the sine and cosine waveforms used for the mixing process are "perfectly" reconstructed by the alternating 1's and 0's sequences, the SFDR is beyond question. Similarly, using the sine and cosine sequences for counterclockwise rotation will result in a $+F_s/4$ upconversion. This is useful when the lower image is to be centered at DC, since that image will now alias to DC. The resultant quadrature data are then lowpass FIR filtered and decimated. It is common for the $\pm F_s/4$ up/downconversion process to be combined with the FIR filtering as a single process.

The $F_s/4$ up or downconversion requirement that the sample frequency be exactly four times the downconvert frequency is very restrictive and, therefore, not very desirable. Nevertheless, these "tricks" eliminate the need for true digital multiplication, are logic-efficient and support higher clock rates than would otherwise be possible, and are therefore still commonly included in today's Commercial digital receiver devices. Also, one other digital processing "trick" is worth mentioning at this time, because it is frequently used to invert a real data stream spectrum. Simply stated, mixing the input real data stream with a frequency at $F_s/2$ inverts the real spectrum and, for a digital mixer, this can simply consist of a (1,-1) multiplication sequence. This sequence then reduces to simply changing the sign of every other input data sample. This spectral inversion is normally used to reverse a previous spectral inversion that occurred during the digitization process (i.e., second Nyquist zone digitization, etc.). However, it should be noted that any DC component of the input data stream will mix with $F_s/2$, resulting in replicated terms at $\pm F_s/2$, and all multiples thereof.

A more flexible digital downconversion process can be realized by the use of a Numerically Controlled Oscillator (NCO). An NCO performs the LO down-conversion to quadrature without a fixed sample frequency to downconversion frequency requirement. The NCO consists of a digital sine and cosine generator, with the computations being done by a Coordinate Rotation Digital Computer (CORDIC) algorithm that does the sine and cosine computations necessary to

generate 0 and 90 degree phase-shifted LO-equivalent digital waveforms. The computations are based on an iterative half-angle process, are pipelined, and take a few more clocks than the number of bits of precision specified. However, once the pipeline is filled, new digital sine and cosine outputs then result for each and every clock.

The input parameter to the CORDIC is a two's complement (i.e., signed) binary value between zero and $\pm 2^{(n-1)}$, where n is the number of bits of precision, that represents the desired phase for which the sine and cosine values are to be computed. Thus, an input stream of digital phase values, at a given clock rate, translates into digitally synthesized sine and cosine waveforms that are the desired upconversion or downconversion LO frequency. The input two's complement digital number to the CORDIC algorithm is computed, for each clock cycle, as a "Phase Value" (PV) that is the running n-bit sum of a "Center Frequency Step" (CFS) and a "Phase Offset" (PO). The CFS, which is actually a phase increment to be added each clock, is computed as CFS = $(F_s/F_c)\ 2^n$, where F_s is the desired LO frequency to be synthesized and F_c is the digital clock update rate. Also, remember that the 2^n value is signed, so the range is $\pm 2^{(n-1)}$, and that values ranging from 0 to $(2^n - 1)$ represent a complete rotation around a unit circle. The sign of this phase increment determines the direction of phase rotation around the unit circle, clockwise or counterclockwise, corresponding to a negative or a positive LO frequency, respectively.

The NCO Phase Offset value is typically zero, but can be used to impart a fixed (starting) phase offset on the synthesized output. It is commonly used to "initialize" the PV accumulator via a "sync pulse", as might be required at the start of a coherent pulse or sweep. Also, the PO feature can be used to time-interleave one or more NCOs, thereby permitting the synthesis of even higher LO frequencies. Since the PV is the sum of the CFS value and the PO value, with 2^n bit precision, this accumulator performs modulo math, and the highest frequency that can be synthesized by a single NCO is $\pm F_c/2$. However, this is not a problem since the A/D clock rate and NCO clock rate are normally equal, and so the A/D data will alias at $F_c/2$. If the required LO frequency is higher than one half the NCO clock rate (i.e., A/D clock rate is a multiple of the NCO clock rate), then two NCOs can be time-interleaved to synthesize LO frequencies up to the NCO clock rate, three NCOs can be time-interleaved to synthesize LO frequencies up to one and one half the NCO clock rate, etc.

The digitally synthesized sine and cosine outputs from the CORDIC algorithm are then used to multiply the input A/D real data stream, A_{in} to generate the I and Q quadrature output data streams I (t) = A_{in} cosine (PV (t)) and Q (t) = A_{in} sine (PV (t)). The SFDR associated with the digitally synthesized sine and cosine waveforms is a function of how "pure" these waveforms are. The purity is a function of the numerical precision employed and, for an n-bit binary word size, ranges from perfect (i.e., as would be the case for $F_s/4$ phase increments), to the limit of the CORDIC precision. These variations from correct sine and cosine values, to values

that are numerically limited approximations to the correct values, gives rise to spurious spectral components. Therefore, there is a direct relationship between the SFDR and the numerical precision associated with an NCO design. Some NCO designs within Commercial digital receiver devices include provisions for amplitude and/or phase "dither" to mitigate the effects of repetitive accuracy errors, but for NCOs implemented in FPGAs, it is usually easier to simply increase the numerical accuracy employed to increase the SFDR.

The CORDIC output precision for the sine and cosine values is typically set to about one half that of the PV value used to generate them. For example, CORDICs utilizing a 32-bit PV value, with 16-bit sine and cosine output values, can typically provide a worst-case SFDR of >96 dB. Similarly, utilization of a 36-bit PV value, with 18-bit sine and cosine output values, can typically provide a worst-case SFDR of >108 dB, or better. These examples are given because 16-bit by 16-bit and 18-bit by 18-bit "hardware" multipliers, optimized for speed and resource utilization, are common in today's programmable devices. Also, it should be noted that these are "signed" multiplies and that the A_{in} data should therefore be two's complement values, MSB-justified. If the A_{in} data precision is less than the sine or cosine precision (i.e., multiplier input precision), then the unused multiplier bits associated with the A_{in} input should be the LSBs and should be tied low (input zeroes).

Digital Filtering

The NCO output I and Q Baseband quadrature data streams contain both sum and difference frequencies and must therefore each be lowpass filtered to remove the unwanted sum frequencies. Since it is very desirable to highly over-sample the wanted information bandwidth to achieve substantial processing gain, the overall (Nyquist) bandwidth may be substantially greater than the information bandwidth of interest. For multi-channel wireless applications, this translates into the need for lowpass filters with high decimations and very sharp transition bands. These needs can be met through the use of one, and typically more than one, decimating digital filter. Commercially available single-channel, and more recently, quad-channel digital receiver chips from Vendors such as the Analog Devices AD6640, Graychip (now TI) GC4016, and Intersil (formerly Harris) HSP50216, typically employ one or more cascaded integrator-comb (CIC) filters, followed by one or more FIR filters, to meet these requirements.

CIC filters do not require multipliers and have very limited logic requirements. Their structure consists of N digital integrator stages operating at the input data rate f_s, followed by N digital comb stages that have a differential input sample delay M, that is normally 1 or 2 clocks, operating at the decimated (output) data rate $f_o = f_s/R$, where R is the decimation factor. The CIC filters associated with Commercial digital receiver devices typically have a minimum decimation of from 2 to 8, and can provide very high decimations (i.e., up to 65536 for the HSP50216). However,

CIC filters also typically have a large transition band, owing to the (unity gain) Power (magnitude squared) frequency response: $P(f) = (1/(M*R)^N)^2 (\sin(\pi M f R/f_s) / \sin(\pi f/f_s))^{2N}$, where f is the input signal frequency.

This expression is commonly rewritten as: $P(f) = (\sin(\pi M f) / \sin(\pi f/R))^{2N}$ and is usually plotted as: P (f) dB = 10 Log (P(f)), for frequencies from DC to one half the output sample rate (f_o). However, some Vendors prefer to display this form of the frequency response as a plot, from DC to f_o that then mirrors about $f_o/2$, or from $-f_o/2$ to $+f_o/2$ that is then symmetrical about DC. Also, these frequency responses are commonly plotted vs. f/f_o, since the actual frequency response for digital filters is essentially a normalized function of the output sample rate and so it doesn't matter what the actual output sample rate is. This is an important point since, for any given digital filter design (i.e., not only CICs, but also FIRs, etc.), the frequency response of the filter scales with the sample rate.

Commercial digital receiver devices have primarily been designed for wireless communications applications, where a relatively narrow bandwidth single channel is to be extracted from a large, multi-carrier frequency band. For such applications, the high decimations that can be provided by CIC filter utilization translates into a significant reduction in logic resource requirements in conjunction with potentially high processing gains. The one or more FIR filters that follow the CIC filter section are then utilized to provide the sharp transition bands that are required in order to extract the narrowband channel information, while rejecting adjacent channel information that would otherwise alias into the passband. These FIR filters also typically provide additional decimation and, therefore, additional processing gains.

A FIR filter is the digital equivalent to an analog transversal filter and, as such, outputs the sum-of-products for a number of weighted taps, spaced across a delay line. The tap weights are the filter coefficients, and the tap spacing (delay) is determined by the input data clock rate (note: symmetrical coefficient FIR filters are phase linear). The decimation by "m" can consist of simply throwing away (m-1) out of every m output samples (i.e., the resultant sum-of-products generated for each clock cycle). Alternatively, the additional (m-1) available clock cycles can be used to facilitate either the processing of additional taps, or to permit a reduction in the logic resources that are required to process the given number of taps. FIR filter hardware design tools are available from numerous sources such as the programmable logic Vendors (i.e., Altera, Xilinx, etc.), as well as hardware-independent FIR filter design (coefficient computation) tools from Vendors like MATLAB and organizations such as IEEE. Basically, the "flatness" of the passband, the width of the transition band, and the attenuation provided in the stopband, are functions of the filter design methodology, in conjunction with the number of taps and the numerical precision employed. A common FIR filter design technique uses window functions (i.e., Blackman, Hamming, etc.) to reduce sidelobes. Another popular method, Remez (or Parks-McClellan) algorithm, is based on optimum approximations for passband and stopband cutoff frequencies, that allows tradeoffs between transition band and passband ripple.

FIR filter logic design can be typified as being either a "polyphase", or a "fully parallel" (transversal) implementation. Polyphase simply means that different coefficients and different data are used at different times. These filters typically employ a fixed number of taps with fixed coefficients, fixed decimation, and are normally designed with the FPGA Vendor's tools. Polyphase FIR filter designs are popular because their implementation requires a minimum of logic resources, but at the price of being rather inflexible. A fully parallel FIR filter design can actually utilize any one of a number of different logic implementations in order to support either fixed or User-programmable coefficients, allow for selection of a coefficient set from multiple such coefficient sets, permit a variable number of taps, support variable decimations, etc. In general, these designs are User-created and then compiled for the target FPGA, require significantly more logic resources than polyphase implementations, but offer as much application flexibility as required.

It is important to note that large numbers of very fast "built-in" hardware multiplier/accumulators have been included as an integral part of the more recent FPGA device designs (i.e., the "Stratix" family from Altera and the "Virtex-II" family from Xilinx). This trend is due to the level of hardware processing parallelism that is possible within FPGA designs, and that cannot be attained by existing "traditional" DSP and CPU devices with fixed-architectures and instruction-sets. Also, these FPGA-implemented hardware multiplier/accumulator logic elements are based on highly efficient core logic designs that combine to form 18-bit by 18-bit (or higher) precision units with little overall performance degradation (i.e., speed reduction or additional peripheral logic requirements). This, in combination with the increased logic, routing, and memory resources offered by these newer devices, makes the utilization of fully parallel FIR filter designs more practical to implement. This is very appealing from an applications point-of-view, owing to the above noted increase in flexibility that can be realized by such FIR filter designs. FPGA filter implementations can result in increased performance at significantly reduced cost, even if the FPGA is operated at clock rates well below that of today's GHz DSPs and CPUs, not to mention the additional potential system-wide benefits such as simplified real-time processing requirements, reduced I/O requirements, etc.

DIGITAL TRANSMITTERS

The classic single-sideband analog transmitter is the reverse process of that for the classic heterodyne receiver. Therefore, it consists of band limited I and Q Baseband signals that are to be upconverted to a final output RF, via one or more intermediate stages (i.e., IFs). In its simplest form, where the RF output is also the first IF, a single I&Q modulator is driven by both quadrature inputs, configured for selection of either the upper-sideband or the lower-sideband, and an LO that is at the appropriate output RF frequency. When the input I and Q data is digital, two

D/As are required to convert these input data streams to analog quadrature signals. Consequently, analog lowpass filtering of both quadrature signals is then required to remove the replicated images associated with a sampled process.

For a digital transmitter, the I&Q modulator is replaced by an NCO configured to perform an "imageless" upconversion. The upconversion is imageless because the multiplication is accomplished entirely in the complex frequency domain (i.e., I and Q complex digital input data streams multiplied by NCO complex LO waveforms). Thus, the entire Baseband is simply translated up in frequency by the programmed LO frequency. At this point, additional signal processing requirements become dependent upon the final RF output requirements. If higher RF output frequencies are required than can be reasonably provided directly by a D/A, then two D/As are usually employed to convert the IF I&Q data streams for utilization by a subsequent analog (I&Q modulator) upconversion stage. In this event, both D/A outputs must be lowpass filtered to remove the sampling-related spectral replications.

If the RF output frequency band is not beyond that which can be directly supported by a D/A, then either of the digitally (imageless) upconverted quadrature data streams can provide the desired information bandwidth. The selected (I or Q) output sample stream will therefore be centered at the desired RF (LO) frequency, and will be output as a real data stream for reconstruction by that D/A. However, it should be noted that the complex multiplications are no longer necessary since only the real outputs need be computed (i.e., the required functionality is now equivalent to that of the analog I&Q modulator). Having to only compute the real output: IF (t) = I (t) cosine (PV (t)) - Q (t) sine (PV (t)), where PV is the NCO Phase Value (i.e., LO), reduces the number of digital multipliers required for the output data sample computation by one half (i.e., from 4 to 2). This essentially reduces the transmitter NCO logic requirements to that of the digital receiver, but with an LO and quadrature pair (I and Q) as the inputs producing a real IF output, as opposed to the receiver's LO and real IF inputs producing quadrature pair (I and Q) outputs.

The digital filtering requirements associated with digital transmitters is significantly different from those associated with digital receivers. As previously noted, the reconstructed analog output from a D/A exhibits a sinx/x roll-off that is output sample rate dependent. Therefore, the highest possible output sample rate should be employed to mitigate the roll-off over the desired output information bandwidth. Since the output spectrum is replicated at every multiple of the sample rate and must therefore be band-limited, increasing the output sample rate also simplifies (analog) filtering requirements.

The input digital sample rate F_s can be increased by a factor R (i.e., output sample rate $F_{out} = R\ F_s$) via interpolation. Commercially available digital transmitter devices, that are the counterpart to the digital receiver devices, typically utilize a combination of one or more interpolating FIR filters to increase the data rate. The FIR filters provide sharp transition bands to maximize the stopband and provide good out-of-band rejection. One or more interpolating CIC filters, with relaxed

stopband requirements, then follow the FIR filters in order to efficiently support high interpolation factor requirements. However, it should be noted that, while the decimating CIC filters are unconditionally stable, interpolating CIC filters, if disturbed, can go unstable, resulting in broadband, high level white noise. This problem is avoided if there is no rounding employed within the integrator stages, since it is the unbounded growth of small errors in the integrator stages that cause the instability. However, the restriction on rounding translates into a word size growth requirement that, in turn, increases logic resource requirements. Also, it should be noted that noise (i.e., bit errors) within the integrator stages can still cause instability. One Commercial four-channel digital transmitter device (the GC4116 from Graychip) has a patented feature ("auto flush") that detects CIC instability and automatically re-initializes the filter.

The sample rate increase must also be sufficient to support the desired output IF sample rate requirements. In view of the limitations on today's D/A clock rates, the IFs that can be directly supported by digital transmitters are typically <100 MHz (i.e., $<0.25F_{out,}$ and F_{out} <400 MHz), and the lower the IF frequency, the better the performance. In some cases, the designs take advantage of the spectral replications associated with a sampled process. In these cases, the NCO usually is programmed to place the desired output frequency in a higher Nyquist zone, with consideration given for the sinx/x D/A output response (i.e., the peak of the response will be 13 dB down at 1.5 F_{out} and 18 dB down at 2.5 F_{out}).

Another common digital transmitter D/A feature is "zero-stuffing", whereby the D/A output section is clocked at twice the input sample rate with "zero" (mid-scale) value outputs being inserted between each input sample. This stretches the sinx/x response at the expense of 6 dB in output power level at DC. Similarly, these D/As may also include an inverse sinx/x filter that is usually an 11-tap, symmetric, FIR filter that employs binary-weighted coefficients, so it is logic efficient and fast. This filter flattens the D/A output sinx/x response to < 0.1 dB) from DC to 0.45 F_{out} at the expense of a 3.815 dB insertion loss. Also, it should be noted that changing the sign of the interpolated data stream, commensurate with the $F_s / 2$ digital mixing sequence previously discussed, results in spectral inversion. This sequence is sometimes combined with an interpolation filter to create a highpass response as a D/A device feature (i.e., AD9772A, etc.) that is useful in direct IF output applications.

PRODUCT DESIGN CONSIDERATIONS

The transformation of a design concept into a viable Accelerator or Storage Ring instrumentation product, using today's technology, is worthy of discussion. Such products must be reasonable to manufacture, reliable, and maintainable, in order to be cost effective. Additionally, consideration should be given to multi-use product designs to afford commonality, not only within a given facility, but also across multiple facilities. This is a far more reasonable goal today than it was in the past,

owing to the increased bandwidth of today's mixed-signal devices, combined with the dynamic in-circuit programmability and burgeoning resources of current and emerging FPGAs.

Such product designs must start with schematic capture, via one of the numerous Commercial products such as ORCAD, etc., utilizing standardized and verified parts libraries and associated mechanical footprints. Notes should be present on the schematic to reflect operational limits, special component placement requirements, critical signal routing constraints, controlled impedance requirements, decoupling, thermal considerations, etc. Since these designs will employ programmable logic devices (i.e., EPLDs and FPGAs), and may also include programmable ASICs (i.e., digital transmitter or receiver chips, DSPs, etc.), register maps and supporting documentation requirements then arise in support of the software development associated with hardware debug and test, applications program interface (API) development, etc.

The schematic capture package also generates a Bill of Materials (BOM) that lists part designators, vendor part numbers, quantities, etc., as well as a "netlist" file that provides the component routing (interconnect) information and identifies the mechanical footprint for each component. The netlist file is imported to a PC Board Layout package, such as PCAD, etc., that creates the component footprints for subsequent placement by the PC Board layout person, and "rubber bands" the required interconnects (routes), based on the netlist interconnect data. The layout person first creates the mechanical board outline, reads in the netlist data, assigns the layers (multi-layer boards are required for the designs being addressed), and then proceeds with initial component placement.

For purely digital boards, it is fairly common practice for the layout person to assign routing "rules" to be used by an "auto-router" that actually routes the board design. This practice is not generally acceptable for densely routed layouts because these routers typically drop too many vias (i.e., small plated-thru holes that provide between-layer routing paths) to thereby simplify the routing from point A to point B. However, the vias actually reduce the available PC board routing area by virtue of being holes that block routing paths, and so the routers keep asking for more layers in order to create more routing paths. There are limits on the number of layers that are "practical" for PC board designs, based on controlled impedance requirements, board thickness (i.e., to fit a card guide or connector), cost, etc. In general, PC board design is one of the remaining areas where it is hard for an auto-router software package to beat a skilled PC board layout person. This is especially true when it comes to dense, high speed, mixed-signal designs.

However, these board layout packages do typically offer features that greatly enhance the efficiency with which such designs can be "hand routed", making hand routing a viable, though more expensive alternative, but one that can yield a significantly better product. Additionally, the design verification features (tools), that are an integral part of such packages, are essential to the development of a quality, well controlled, and well documented product. The resultant output

(Gerber) files, in conjunction with board lay-up drawings, are then available for electronic transfer to a "board house" (i.e., PC board manufacturing facility).

The selection of a board house for a specific design is not simply cost-based, because different board houses have different capabilities. The small physical sized, lead pitch, ball spacing, etc. of today's components commonly translate into dense boards with fine line-spacing, very small (or micro) vias, buried and blind vias, very flat surface features, etc. These requirements may well narrow the field of board houses that can meet the requirements without, or with acceptable up-charges. For high speed, mixed signal designs, one must add controlled impedance to this list of requirements. Similarly, the potentially high component cost and typical low manufacturing volume associated with the type of products being addressed here, is such that a "net list" board test (i.e., Gerber file based) capability like that of a "Probot" becomes an essential board house requirement in order to ensure receipt of "good" boards.

Next, the requirements associated with "board-stuffing" must be addressed. Again, the small geometries encountered in dealing with today's components mandates the use of "pick-and-place" equipment, re-flow ovens (and re-flow profiles), and fine-tipped or hot air soldering equipment for touch-up (re-work) under suitable magnification. Board inspection is also somewhat more involved. For example, Ball Grid Arrays (BGAs) must either be x-rayed, optically inspected (ERSA Scope, for example), or tested via a JTAG boundary scan. Board stuffing houses (i.e., contract manufacturers) capable of meeting these requirements often refuse work orders below a minimum quantity that may well be in the hundreds to thousands. However, there are fully qualified houses that will service "prototype" or "limited production" quantity orders, albeit at an up-charge.

EXAMPLE BEAM INSTRUMENTATION PRODUCT DESIGN

So, how does one put this all together? Perhaps the best approach would be to do an example Storage Ring Beam Instrumentation product design. Let us assume that the product is to support two requirements, the first being bunch-by-bunch current monitoring, and the second being beam steering. For the purpose of generality (i.e., to be in about the middle of common Storage Ring RFs, we will assume that the RF is 360 MHz, the Storage Ring diameter is about 80 meters, and we therefore can have a maximum of 300 current bunches. Thus, the revolution frequency F_{rev} for a bunch will be 1.2 MHz. For sensors, we will assume that we have four BPMs, spaced evenly about a cross section of the ring, and the outputs have been amplified, if required, and band-limited via filters. A "trigger" will also be required to identify the "bunch number 1" time. We will first address the bunch-by-bunch current monitoring requirement.

To provide bunch-by-bunch current monitoring, a sample must be taken at the appropriate time, relative to the start of a RF cycle, for each and every bunch, with

these bunches being spaced at the RF interval. For this example, that requirement translates into the need for a 360 MHz A/D encode rate, sampling a 360 MHz analog (RF) BPM output, and thereby creating sets of 300 digital samples, corresponding to the maximum of 300 discrete bunches that can be distributed about the Storage Ring. These sets of 300 samples will then repeat at the revolution frequency (i.e., 1.2 MHz x 300 samples = 360 MSPS). It is assumed that provision for co-adding up to "n" bunch samples for each bunch is required (i.e., adding bunch 1 sample for revolution 1, with bunch 1 sample for revolution 2, etc., and similarly for bunch 2 samples, etc.). It is also assumed that all four BPMs must be processed such that the data from the four BPMs can be used in support of beam steering.

First, we must select an appropriate A/D, one that provides the best resolution and can directly digitize a 360 MHz input. This is beyond the capability of existing 14-bit A/Ds, and the best 12-bit A/D choice is the Analog Devices AD9430. This device supports a maximum encode rate of >210 MHz (rated performance is to 210 MSPS, but it works well beyond this), and has a full-power analog input bandwidth of 700 MHz. Two of these devices must be time-interleaved to support the 360 MSPS requirement. This is accomplished by inputting the 360 MHz reference clock (sinusoidal waveform at +4 dBm) to a differential PECL receiver followed by a PECL flip flop with a preset control (required to support initial time alignment). The virtually perfect 50% duty cycle differential outputs then drive the A/D encode inputs 180 degrees out of phase with respect to each other (i.e., both A/Ds encode at 180 MHz, but on opposite edges of that clock). The estimated SINAD at a 180 MHz sample rate for a 360 MHz input is >58 dB. The Aperture Uncertainty (jitter) is 0.25 ps rms, so the jitter-limited S/N is 65 dB.

The first (in time) AD9430 will output an LVDS sample stream at 180 MSPS that is translated by the FPGA into two LVTTL data streams, corresponding to samples for bunches 1 and 3, then 5 and 7, etc. Similarly, the second AD9430 outputs samples for two data streams with bunches 2 and 4, then 6 and 8, etc. These four 90 MSPS data streams will be processed by an FPGA that is configured to have four dual-port memories, corresponding to the four input bunch sample data streams. Each of the four dual-port memories will be organized as 128 by 36 bits and support the co-adding of samples for 75 bunches (i.e., memory 1: address 1 is bunch 1, address 2 is bunch 3, etc.) so four memories support sample data for the 300 bunches. At the start of each "n" co-add sequence, the samples are simply written to the appropriate memory addresses. For each of the remaining (n-1) sample data sets, the new sample is added to that read out of the memory, and the new sum is then written back to that same memory address. These four concurrent processes therefore execute at a 90 MHz rate. Upon completion of the "n" co-add sequences, the sample sum data is read out of memory and the new sample data is written into memory, and the process repeats. The output sum data are written to a FIFO buffer within the FPGA for subsequent output via the selected Bus Interface (i.e., VME, PCI, RACEway, etc.).

One or more dual-port memories can also be configured to accept individual (selected) bunch sample data streams in support of a true "decimating digital filter" feature, as opposed to a simple co-add. Thus, as the sample data from a given bunch is written cyclically through the memory, it can be concurrently processed by a decimating FIR filter configured to have as many taps, and as much precision, as required for the desired passband, transition band, and stopband. This can be taken one step further. A "digital receiver" feature can also be added via NCO downconversion of an individual bunch sample data stream. This frequency translation (up or down) can be by any frequency up to one half the effective sample rate (i.e., half the revolution frequency), and the quadrature bunch sample data can then be FIR-filtered and decimated, as appropriate.

This A/D and FPGA architecture can be replicated four times, on a single PC board to process the data from all four BPMs. This permits data processing, within an FPGA, that utilizes all four sources of bunch sample data in support of beam steering. Intellectual Property (IP) FPGA core designs are readily available and simplify the implementation of complex logic functions (i.e., Logarithms, Complex DFFTs, Rectangular-to-Polar, etc.). Thus, the "standard" beam position algorithms can be implemented with relative ease and both time-domain and frequency-domain processing can be supported.

Since the 360 MSPS time-interleaved A/Ds alias at multiples of 180 MHz, a signal at 361 MHz will alias to 1 MHz, as will a signal at 359 MHz. If this is a problem, then a bandpass filter with an asymmetric response can be used to band-limit the BPM output such that a band can be passed that is from a known "clean" portion of the spectrum. A digital receiver can be included in the design to extract the desired high frequency "span of interest" from the A/D-pair output sample stream. Time-interleaved NCOs perform the programmed downconversion and the output quadrature Baseband data is then digitally filtered and decimated, as required. At this input data rate, the NCO, CIC, and FIR filter functionality must be implemented within the FPGA, because this data rate exceeds that of currently available commercial digital receiver devices. However, these real-time parallel processing requirements are easily implemented within an FPGA.

Since the FPGAs are in-circuit programmable and can be reconfigured from memory (i.e., normally "Flash devices"), a given hardware design can be tailored to multiple applications. Flash memory devices can provide sufficient storage capacity for multiple designs and can also be rewritten via the Bus Interface. Consequently, all features of the above example design could be implemented individually, or in combination, providing an application-flexible common design. The processing afforded by such high speed, mixed signal, FPGA-based designs, can significantly reduce the external processing and I/O requirements that are associated with Accelerator and Storage Ring instrumentation. Additionally, this approach can minimize the instrumentation system complexity, yielding increased reliability, all while providing enhanced performance, flexibility, and commonality, at a potential cost savings.