

Digital Low-Level RF Control Using Non-IQ Sampling

Lawrence Doolittle (LBNL, Berkeley, California),
Hengjie Ma, Mark Stuart Champion (ORNL, Oak Ridge, Tennessee)*

Abstract

The success of digital feedback with synchronous IQ sampling for cavity field control in recent accelerator projects make this LLRF control scheme a popular choice. This short-period synchronous sampling does not, however, average out well-known defects in modern ADC and DAC hardware. That limits the achievable control precision for digital IQ LLRF controllers, while demands for precision are increasing for future accelerators such as International Linear Collider. For this reason, a collaborative effort is developing a digital LLRF control evaluation platform to experiment using coherent sampling with much longer synchronous periods, on the order of the cavity closed-loop bandwidth. This exercise will develop and test the hardware and software needed to meet greater future RF control challenges.

INTRODUCTION

The current generation of accelerators such as SNS[1] specified RF control accuracy on the scale of 1° and 1%, and achieved that with digital LLRF.

The upcoming generation of accelerators, such as LUX[2], LCLS[3], ILC[4] and XFEL[5], will require accuracies as tight as 0.01° and 0.01%. Reaching such goals will require deep understanding of hardware limitations, careful construction, and sophisticated software.

As in any feedback system, the ultimate error is dominated by the measurement process. Error terms fall into one of the following categories: systematic error, accuracy, linearity, repeatability, stability, resolution, noise. Non-IQ sampling primarily addresses linearity, which in turn improves the ability of calibration processes to address repeatability and stability.

Low latency digital feedback processing is performed in a Field Programmable Gate Array (FPGA). When the (usually downconverted) RF samples come every 90° (or 270° , or 450°) apart, the design is known as IQ. Then the digital processing techniques are easy to understand, and can be accomplished even in relatively small chips. Handling non-IQ samples involves a conceptually small change to the math, but increases the DSP resources needed on the chip (roughly 8 instead of 3 multipliers), and requires more attention to the effects of rounding error in the computation.

*ORNL/SNS is managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract DE-AC05-00OR22725.

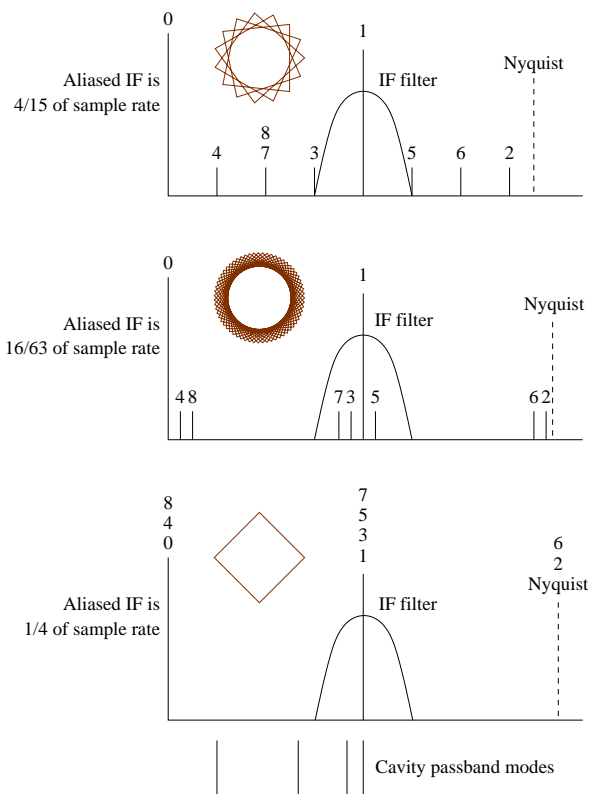


Figure 1: Aliased harmonic frequencies

WORKING WITH ADCS

Modern ADC chips in the 80 MS/s, 14-bit class[6][7] are designed, optimized, and characterized for use in communications gear. Some of the parameters that are critical for accelerator LLRF control are not considered important in that context: in particular, their differential linearity and temperature stability of gain. IQ processing, with a coherent sampling period of four samples, is unusually sensitive to differential nonlinearity.

The critical ADC of an LLRF system is used to digitize a nearly modulation-free CW signal. Small amounts of differential nonlinearity generate high harmonics of an input carrier, and those harmonic contributions depend on the input signal's amplitude and phase. In an IQ system, all odd harmonics of the signal alias to the same frequency as the carrier itself, as shown in the bottom frame of figure 1. As the input phase and amplitude changes (in response to either operator commands, or higher level software correcting for cable temperature changes), the distortion changes and corrupts the measurement. This can be thought of as a

consequence of only exercising the ADC at four points in its range.

By changing the sampling rate even slightly, say to 63/16 of the carrier frequency, most harmonics no longer line up with the carrier. This situation is depicted in the middle frame of figure 1. In this example, the 62nd and 64th harmonics are the first to overlap the carrier. Digital processing can distinguish the carrier signal from all lower harmonics, at the expense of long group delay filtering. Now the ADC is exercised at a constellation of 63 points in its range.

A slightly larger shift in sampling rate, say to 15/4 of the carrier frequency (top frame of figure 1), will keep most harmonics (including the third harmonic, which may also be generated in mixers or preamplifiers) farther from the carrier. This change also lowers the first harmonic numbers to overlap with the carrier (14 and 16 in this case). While the constellation of points exercised (15) is smaller, and the linearity improvement therefore not as great as the previous case, the relaxed filter requirements make this situation desirable for low latency feedback processing.

On the output side, it is helpful to keep all harmonics out of the passband of the IF filter, and therefore away from the high power amplifier. It would be especially unfortunate if such an aliased harmonic lined up with one of the nearby cavity modes.

Keeping the carrier near the middle of the Nyquist band, as these examples have done, simplifies the required signal processing, as will be seen shortly.

DSP

The traditional $\alpha + \beta z^{-1}$ filter element from IQ processing is still effective at providing programmable phase shift. For a phase step θ , the filter coefficients to provide a gain of A and rotation of ϕ are found by

$$\begin{pmatrix} 1 & \cos \theta \\ 0 & \sin \theta \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = A \begin{pmatrix} \cos \phi \\ \sin \phi \end{pmatrix}$$

Note that when $\theta = \pi/2$ for IQ sampling, the matrix above becomes the identity matrix. Angles near 0 or π give a matrix with near zero determinant, and consequent worsening of signal to noise ratio in the output.

A pole at the carrier frequency is generated with the IIR filter

$$\frac{1}{1 + pz^{-1} + z^{-2}}$$

where $p = 2 \cos \theta$. Note how this filter reduces to a multiplier-free I and Q accumulator when $p = 0$. A non-zero p not only adds a multiplier, but requires an unrealizably short feedback path within the filter, and exacerbates the effects of finite precision arithmetic. One way to address the (internal to the filter) feedback latency problem is to recast the polynomial as $1 + bz^{-1} + z^{-2} + cz^{-3} + dz^{-4}$, where $b \approx p$, and c and d are chosen to place two of the roots at the desired $e^{\pm i\theta}$. If b is represented by only one or two binary bits, multiplication by b can be implemented

without resorting to long-latency multipliers. The shortest path remaining that involves multipliers is three cycles long. This approach works well when $|p| < 1$, and therefore $1/6 < f_{IF,aliased}/f_S < 1/3$.

It is not impossible to make a processing chain for θ far from $\pi/2$, but they are more complex and have inherently larger group delay than the modified IQ design.

Normalized (to the sampling rate) frequencies such as 4/15 cannot be expressed exactly in the binary counters that are typically included in single-chip DDS products. This is not a problem for designs based on the familiar block diagram[8], since generalized modular arithmetic is easy to program in an FPGA.

Highly decimated RF vector waveforms are used for everything from operator comfort displays to automated self-tuning of these LLRF systems. Either CORDIC[9] or table-and-multiplier-based techniques are needed to convert these carrier frequencies to DC for averaging, an additional complexity relative to IQ systems.

SYSTEM ARCHITECTURE

To control both system cost and phase noise, the chain from master oscillator to cavity measurement must be kept simple. Phase noise added between the MO and each station (from DC to the closed loop bandwidth of the cavity) becomes actual error on the fields seen by the beam. Low frequency phase noise of the MO itself is ‘‘common mode’’ to all parts in the accelerator, and therefore irrelevant for most purposes.

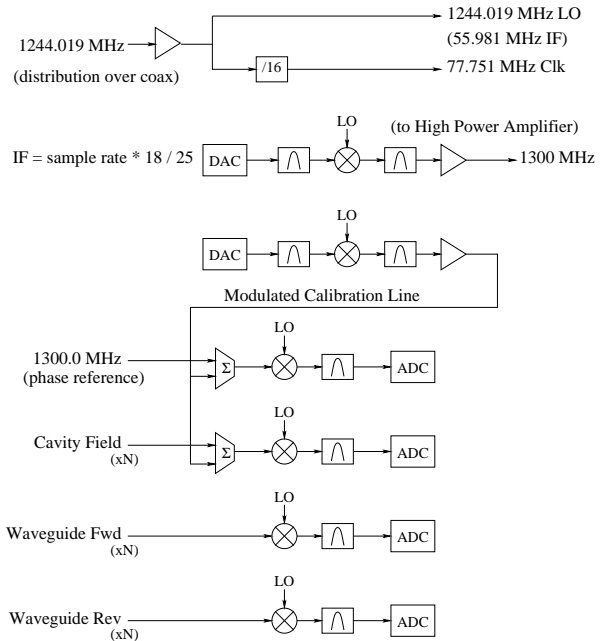


Figure 2: Example application to ILC frequency handling

As long as infrastructure is in-place to distribute a low-phase noise LO to every RF station, it makes sense to reuse that signal to derive the digital sampling clock. When com-

bined with non-IQ sampling strategies, this leads to unusual frequencies for the LO. In the ILC example shown in figure 2, where the sampling clock is 1/16 of the LO, and the IF is 18/25 of the sampling clock, the low-side LO is 200/209 of the RF. Such infrastructure design needs to be coordinated with other accelerator instrumentation and control.

The highly linear RF measurement that non-IQ sampling offers will make calibration and phase transfer processes more accurate. The simple and effective pulsed reference line demonstrated by SNS[10] (and incorporated in figure 2) should be considered the baseline technique. A similar technique can be applied to CW machines, using frequency offsets instead of time offsets.

A more complex calibration scheme could measure the S_{11} of the cavity probe and its cable *in-situ*, and use symmetry to compute the exact amplitude and phase at the probe.

The phase calibration process design for large machines is tightly coupled to their phase reference distribution system design. A transfer of standards must take place many times per second. The starting point is the machine's reference oscillator, and after many transfers that phase reference information must be connected to the field in each active cavity.

EVALUATION BOARD

An evaluation board is under construction[11] (see figure 3) to test some of these concepts, using technology that could scale to large machines like ILC. It is designed for sample rates up to 80 MS/s, based on LTC2249 ADCs and an XC3S1000 FPGA. The board is 96.5 x 127 mm, will dissipate about 5.5 Watts, and interface to a host computer over USB 2.0.

Four RF or IF input channels can be used for traditional forward/reflected/transmitted/spare RF control of a single cavity, or to compute the vector sum of four cavities.

The board's input clock is routed through an AD9512, so the board can derive its sampling clock from an LO signal. If component data sheets can be believed, the board's additive clock phase noise could be as low as 0.3 ps rms.

A high speed LVDS link can connect two such boards, to experiment with vector sum control of eight cavities. This link is also a surrogate for the links needed for the two-level designs proposed for handling 32-cavity vector sums in ILC.

To support research on piezoelectric tuner control of SC cavities, and/or control of fast ferrite phase shifters, the board includes twelve channels of medium speed baseband output, with update rates as high as 520 kS/s.

ACKNOWLEDGEMENTS

The authors wish to thank Brian Chase of FNAL and Alex Ratti of LBNL for their support. Thanks also to the




Figure 3: Evaluation board under construction

gEDA community for the software used in this evaluation board design.

REFERENCES

- [1] Champion, M., "Low-level RF Control System Requirements for RFQ, LINAC, and HEFT," SNS 104010300-SR0002-R00, October, 2002.
- [2] Corlett, J. N. et al., "A Recirculating Linac-based Facility for Ultrafast X-Ray Science," PAC 2003, Portland, Oregon.
- [3] Akre, R., et al., "SLAC Linac RF Performance for LCLS," Linac 2000, Monterey, California.
- [4] Chase, B., "LLRF Specs and Topics," ILC RF/Timing workshop, FNAL, Jan. 2006.
- [5] S. Simrock, "Overview of XFEL RF Specifications and Comparison to ILC Specs," ILC RF/Timing workshop, FNAL, Jan. 2006.
- [6] Linear Technology LTC2249, <http://www.linear.com/> .
- [7] Analog Devices AD6644, <http://www.analog.com/> .
- [8] Doolittle, L. R., "Embedded Networked Front Ends," ICALEPCS 2003, Pohang, Korea.
- [9] Volder, J., "The CORDIC Trigonometric Computing Technique," IRE Trans. Electronic Computing, Vol EC-8, pp330-334, Sept. 1959.
- [10] Piller, C., "Multiplexed Cavity RF and Reference RF Signal Processing for SNS LLRF," SNS Tech Note, March 7, 2003.
- [11] <http://recycle.lbl.gov/llrf4/>