

## PROGRESS ON THE SNS RING LLRF CONTROL SYSTEM\*

K. Smith, M. Blaskiewicz, J. M. Brennan, J. DeLong, F. Heistermann, A. Zaltsman  
 BNL, Upton, New York 11973, USA  
 T. L. Owens  
 ORNL, Oak Ridge, TN 37830, USA

### Abstract

The SNS Ring RF System[1,2] will comprise three  $h=1$  ( $f_{rev} = 1.05$  MHz) cavities and one  $h=2$  cavity, each with individual digital LLRF control electronics. During each 1ms accumulation of 1 GeV protons in the SNS ring cycling at 60Hz, circulating intensity increases to  $1.5E14$  particles. This intensity translates to an average circulating current (at the end of accumulation) of 35A and a peak  $h=1$  current of 50A. The LLRF system primary task is to regulate the phase and amplitude of the RF gap voltage in order to maintain a smooth bunch with minimum peak current and a sufficient beam free gap to accommodate the extraction kicker rise time. Maintaining stable control of the cavity-beam system with such intense beam loading is non-trivial, and to do so, the LLRF system will use a combination of techniques including cavity voltage I&Q feedback, beam current feed-forward compensation, dynamic tuning and cycle to cycle adaptive feedback. This paper describes the progress on the LLRF control system to date.

### INTRODUCTION

The SNS Ring LLRF Control System is being designed to provide individual digital LLRF control for each of the four cavities in the ring. To date, a number of prototype and first article boards have been designed and tested, and the first integrated test of the cavity IQ control loop with the production HRF cavity and PA was recently completed. This paper will discuss the concept and architecture of the ring LLRF system and present performance data acquired for the particular case of the cavity IQ control loop. Details of the RF related machine parameters, the theory of the ring LLRF control and the HRF system design are found in the references [1] and [2].

\* SNS is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U.S. Department of Energy. SNS is a collaboration of six US National Laboratories: Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL), Thomas Jefferson National Accelerator Facility (TJNAF), Los Alamos National Laboratory (LANL), Lawrence Berkeley National Laboratory (LBNL), and Oak Ridge National Laboratory (ORNL).

### SYSTEM ARCHITECTURE

Discussion will be limited to the core of the digital LLRF system, configured from three basic components: the Cavity Controller Card, the VME Carrier Card, and Carrier Daughter Cards providing application specific IO functions. The design goal has been to develop an architecture which satisfies all current requirements, but is flexible and easily reconfigurable in order to meet changing demands, both during system development and throughout commissioning and operation. To accomplish this, the system is designed to be modular, and different control loops are configured by simply selecting appropriate carrier daughter cards and writing the firmware to implement a particular control algorithm. A simplified block diagram of the implementation of a representative loop, the cavity IQ control loop, is shown in Figure 1. Similar channels provide the other LLRF functions such as feed-forward compensation and dynamic tuning control.

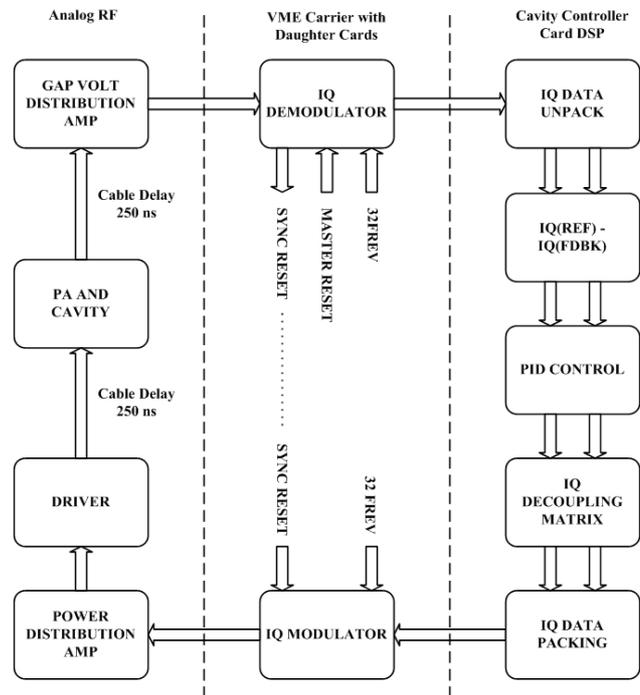


Figure 1: Cavity IQ Control Loop.

### Cavity Controller Card

The Cavity Controller Card utilizes all commercial off the shelf (COTS) hardware. A standard SNS VME based Input Output Controller (IOC), the Motorola MVME-2100, hosts a Bittware Hammerhead-PMC+ module. The Hammerhead-PMC+ is a PCI Mezzanine Card (PMC) with four ADSP-21160 processors. Table 1 outlines some basic specifications.

Table 1: Some Hammerhead PMC+ Specifications

DSP	4 x ADSP-21160, 32 bit fixed and floating point
DSP Clock, Speed	80 MHz, 480 MFLOPS
DSP Core Architecture	Single Instruction Multiple Data (SIMD)
DSP Internal SRAM	4 Mbit dual port
DSP External IO Bandwidth	640 MB/s uP port (shared) 6 x 80 MB/s link ports Autonomous IO Controller
Board Shared SDRAM	64 MB (256 MB max)
Board FLASH	2 MB
Board External IO Bandwidth	64bit, 66 MHz PCI (Host to DSP, Host to SDRAM) 4 link ports (F.P.) 4 link ports (VME P2)

The ADSP-21160 provides a combination of good processing power and as important, the considerable external IO bandwidth required for real-time feedback control. In particular, each DSP has six 80 MB/s “link ports” (two to each of two nearest neighbor DSPs and two external) providing high speed point to point IO. While active, the LLRF loops use only these links to provide minimum delay, deterministic connections between the DSPs and carrier daughter cards. Because the DSPs have large internal SRAM, the loops need not make access to DSP external resources (e.g. uP bus and SDRAM) during the accumulation cycle. For each cycle, all necessary loop parameters, function tables, and diagnostic history data are stored in DSP internal SRAM. Access to external resources only occurs during the roughly 10ms dead-time between accumulation cycles when loops are inactive.

Via a direct Ethernet connection to the controls network and a PCI local bus connection to the DSP PMC card, the MVME-2100 IOC provides the necessary bandwidth to exchange all necessary diagnostic and system configuration data on a cycle to cycle basis.

### VME Carrier Card

The VME Carrier Card [3] is a custom designed VME host board which provides power, a VME interface and other resources for up to four daughter cards. These include an on-board ADSP-21160, FPGAs, SRAM and FLASH to provide extra processing power and functionality as needed. The board design also supports the use of a mating VME P2 rear transition module providing a convenient means of supplying external clock and trigger signals as well as DSP link port connections for both the on-board DSP and the daughter sites.

### Carrier Daughter Cards

The Carrier Daughter Cards provide an easily configurable interface between the digital and analog portions of the LLRF system. While original prototype system testing was based on daughter cards providing single channel functionality (e.g. one A/D channel), the latest architecture employs just two daughter cards (in development), a four channel 14 bit A/D and a four channel 14 bit D/A card, each sharing a common FPGA based digital section designed around an Altera Stratix FPGA. The Stratix FPGA provides tremendous processing power and functional flexibility to the cards. We have already developed A/D FPGA cores to provide direct A/D and IQ demodulation, and a digital receiver core is in development. Similarly, D/A cores have already been developed for direct D/A, IQ modulation and Direct Digital Synthesis (DDS). Commercial IP cores are of course an option as well. The analog front ends of the cards can be configured to use either baseband or RF signals, DC or AC coupled as needed. The VME Carrier Card can host two each of either four channel card (“double width”), or four single channel cards (“single width”).

## CAVITY IQ CONTROL LOOP

### Configuration

The recently completed Cavity IQ control loop test was the first test of a fully integrated loop utilizing all the core LLRF and HLRF hardware together. The cavity was driven closed loop, ramping from 0 to 10kV in 0.5ms and holding at 10kV for a 1.5ms flat-top. No attempt was made to include dynamic tuning or simulated beam loading, though it is planned to do so as soon as practical. For this test, an IQ demodulator card, a single DSP on the Cavity Controller card, and an IQ modulator card were configured as shown previously in Figure 1.

The IQ demodulator samples the h=1 cavity RF to 14 bits at 32 x frev, and correlates the samples with 32 sample h=1 sine and cosine sequences to produce baseband I & Q values updating at frev. 16 bit I & Q values and the turn number (tracked from reset time) are packed into 32 bit words and sent to the DSP via a link port. The link ports are 8 bit parallel point to point links capable of clocking bytes at 80MHz, but were only clocked at 32 frev for simplicity during the test.

The DSP unpacks the I & Q data and based on a stored I & Q function table (I & Q vs. turn number), processes the I & Q errors in parallel PID algorithms. For simplicity, the decoupling matrix was implemented as a simple rotation accounting only for pure loop delay at the cavity center frequency. No attempt was made to correct coupling of the I & Q modulations arising from the cavity RLC response[4] or the change in cavity center frequency vs. drive level. Efforts are underway to include these corrections. The corrected I & Q drive values are then sent to the IQ modulator via a second link port and the loop closes through the analog signal chain.

A key constraint on the cavity IQ control loop bandwidth is the loop delay. Total pure loop delay is about 2.6 $\mu$ s, contributed from: cable delays (500ns), driver, PA and analog filters (250ns), correlation group delay (500ns), other FPGA pipeline delays (250ns), link ports (375ns), and DSP loop code execution (750ns). With a loaded Q of about 25 at 1.05 MHz, the phase slope of the cavity response corresponds to a group delay of about 6.25 $\mu$ s, yielding a total effective loop delay of about 9 $\mu$ s. However, since the cavity group delay only represents the response from the cavity complex poles, it can be compensated. The pure delay of 2.6 $\mu$ s is fundamental in limiting the control bandwidth achievable. Thus, much of the effort aimed at developing the latest hardware includes moving to higher clock rates to reduce pipeline delays in the FPGAs and the link ports. The daughter cards currently in development will easily clock at 128 x freq reducing processing pipeline delays by a factor of four. It is also planned to investigate other IQ demodulation algorithms (4 x RF IQ sampling, etc.) to increase throughput. DSP code optimization continues in an effort to minimize latency in the loop processing algorithm, though it is not likely that the processing time could be reduced below 0.5 $\mu$ s (40 DSP clock cycles). With higher FPGA clock speeds and optimizing of the IQ demodulation algorithm, pure loop delay should ultimately be reduced to about 1.5 $\mu$ s. Although the DSP offers certain advantages over FPGA, particularly when making frantic code changes during initial testing and debugging, the penalty paid in terms of processing delay can easily outweigh this. Thus, the LLRF system architecture will also permit bypassing the IQ loop DSP if this becomes desirable.

**Test Results**

As a preliminary attempt to characterize the loop, a test of the closed loop impulse response was performed. The cavity was brought to 10kV, and then a single turn I command for 9kV was issued to the loop, followed by a return to 10kV. Figure 2 shows the I & Q transient responses and Figure 3 shows the corresponding closed loop frequency response for the I loop. The loops were “optimized” by adjusting gains for the P, I and D terms of the controller while observing the transient response in the I channel. Both the I and Q loops were active with the same gain coefficients used on corresponding I and Q gain terms.

In Figure 2, the I command and I & Q responses have been normalized to the 10kV steady state command. The Q response has been offset to 1.0 for clarity.

In Figure 3, the horizontal axis is shown as the index of a 512 point FFT of the data in Figure 3. The 1dB bandwidth is approximately 70kHz = 1MHz \* (38/512).

Though the closed loop bandwidth observed is likely sufficient for our system, consideration must be given to the effect of beam cavity interaction under the high beam loading conditions that will exist.

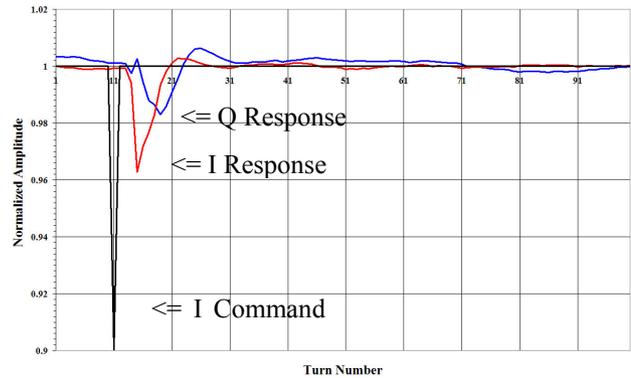


Figure 2: Cavity IQ Control Loop CL Transient Response.

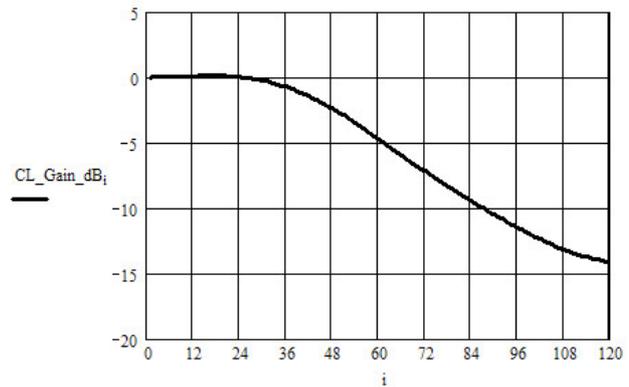


Figure 3: IQ Loop CL Frequency Response (Icav/Icmd).

The relative beam loading parameter ( $Y=I_b/I_o$ ) rises to around 3 at the end of accumulation. Even for optimal detuning, the system operates close to the high current Robinson unstable area [5][6][7]. Direct RF feedback around the cavity was considered in the early stages of the ring RF system design. However, cost and space constraints would not permit the installation of feedback amplifiers in the ring or a nearby structure, and the long cable delays between the LLRF and cavity make it otherwise unfeasible. Options for dealing with this include using somewhat higher than optimal detuning to move further away from the Robinson unstable area, beam current feed forward compensation to dramatically reduce the effective cavity shunt impedance seen by the beam, and investigating IQ loop gain bandwidth tradeoffs. Current plans include using all these techniques and we hope to begin tests in the near future.

**REFERENCES**

- [1] M. Blaskiewicz, et al., Proc. 2001 PAC, pp 490-494.
- [2] A. Zaltsman, et al., these proceedings.
- [3] J. DeLong, et al., these proceedings.
- [4] R. Garoby, “Low Level RF And Feedback”, CERN PS/RF Note.
- [5] F. Pederson, IEEE, TNS, Vol. NS-32, No. 5, pp 2138-2140 (1985) and references therein.
- [6] D. Boussard, CERN 92-03, Vol. II (1992).
- [7] F. Pederson, private communication.