OPERATIONAL PERFORMANCE OF THE SNS LLRF INTERIM SYSTEM*

Lawrence Doolittle†, Alessandro Ratti, Marco Monroy, LBNL, Berkeley, CA 94720, USA
Mark Champion, Hengjie Ma, ORNL, Oak Ridge, TN 37830, USA

Abstract

A new approach has been taken to develop and build the Low-Level RF Control System for the SNS Front End and Linear Accelerators, as reported in a separate paper in this conference[1]. An interim version, based on the proven LBNL MEBT design, was constructed to support short-term goals and early commissioning of the Front End RFQ and DTL accelerators, while the final system[2] is under development. Additional units of the interim system are in use at JLab and LANL for concept testing, code development, and commissioning of SNS SRF cryomodules. The conceptual design of the MEBT system had already been presented elsewhere[3], and this paper will address operational experiences and performance measurements with the existing interim system hardware, including commissioning results at the SNS site for the Front End and DTL Tank 3 together with operational results from the JLab test stand.

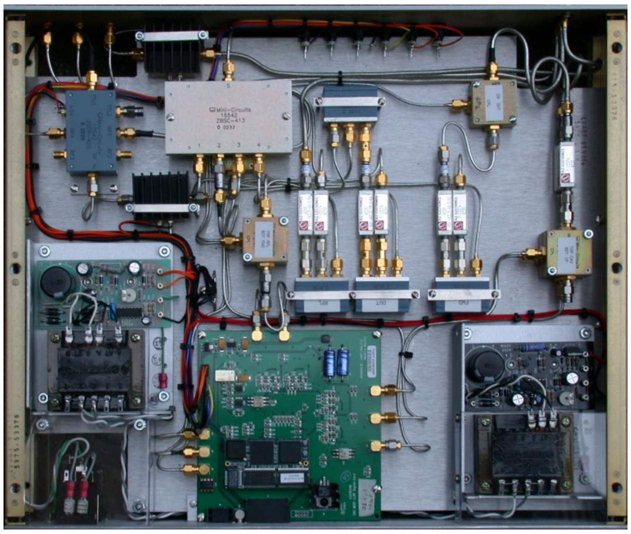

Figure 1: Chassis photograph.

INTRODUCTION

The LLRF hardware, firmware, and software described here all involve evolutionary changes from the SNS MEBT design[3]. The hardware in particular was modified to be simpler, easier to assemble, and more rugged and reliable in the field. The 19" rack-mount chassis changed from 4U high × 500 mm (20") deep to 2U high × 350 mm (14") deep.

It takes very little laboratory electronics to turn on a chassis and operate a cavity with it. Besides power and Ethernet, it needs: an external TTL trigger/gate, a +5V interlock permit, and +5 dBm of the relevant LO (e.g., 352.5 MHz). The frequency of operation is 50 MHz above the LO; three passive connectorized filters included in the chassis have to match the RF. The available output power level of +10 dBm will normally need amplification to drive a cavity. The down conversion mixer for the cavity signal is supplied externally; in the SNS installation, that mixer is temperature controlled and matched to a mixer for the phase reference signal.

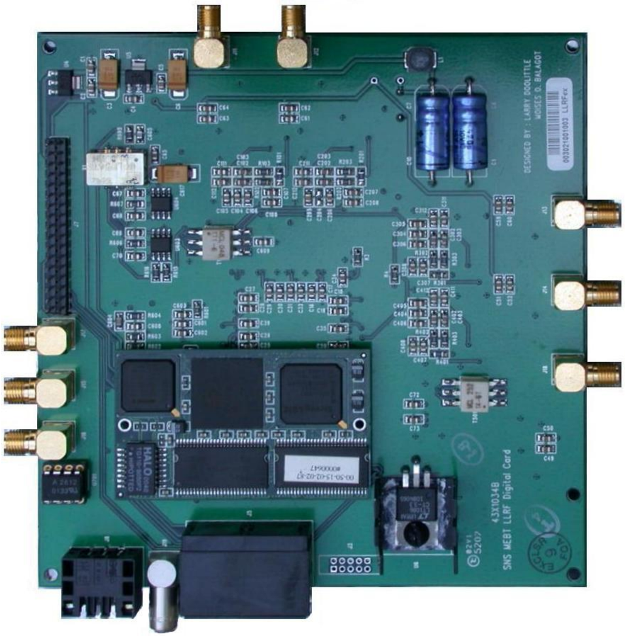

Figure 2: Photograph of main LLRF circuit board.

DESIGN AND CONSTRUCTION

The heart of the chassis is a 137×132 mm circuit board that holds the four 12-bit ADCs (ADS808) operated at 40 MS/s, and a single 12-bit DAC (DAC902) operated at 80 MS/s. A (now obsolete) Xilinx XC2S150 FPGA pro-

* This work is supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. The SNS project is being carried out by a collaboration of six US Laboratories: Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL), Thomas Jefferson National Accelerator Facility (TJNAF), Los Alamos National Laboratory (LANL), E. O. Lawrence Berkeley National Laboratory (LBNL), and Oak Ridge National Laboratory (ORNL). SNS is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U.S. Department of Energy.

† ldoolitt@recycle.lbl.gov

vides the signal processing path, and connects to an embedded single board computer, the Bright Star Engineering nanoEngine[4]. Four of that computer's digital I/O pins are used to program the FPGA via its JTAG interface. The FPGA also provides access to housekeeping functions: a phase lock loop, 4 channels of slow DAC, 8 channels of slow ADC, a serial number, and a digital thermometer. The chassis as a whole dissipates 28 Watts, using linear 5V and 15V open frame power supplies.

The phase lock loop circuitry, based on an ADF4001, can lock the on-board voltage controlled crystal oscillator (VCXO, 80 MHz Connor-Winfield UPLD54TE) to an external source. In our case, that source is a 50 MHz phase reference that is used to synchronize cavities across the accelerator site. For most bench testing purposes, that capability can be ignored, and the system free runs at a frequency that is typically 6 kHz low.

Of the four high speed input channels on the board, one was originally labeled "spare." That channel has now been pressed into service recording the input phase reference signal. Measurement of its phase, and performing software corrections based on it, corrects for both the four-fold state held in the PLL, and most of the phase drift involved in bringing the cavity field pickup signal from the tunnel to the control rack[5].

The original SNS MEBT chassis generated its 50 MHz IF output from a distinct 50 MHz source and an analog I/Q modulator. To eliminate the need for that separate (external) signal source, a high speed DAC was pressed into service to directly generate a vector controllable 50 MHz signal. A nominal 30 MHz waveform is updated at 80 MS/s using a DAC with a current rise time of a nominal 1 ns. This raw output has "spurs" at 50 MHz, 110 MHz, 130 MHz, and higher, with power in each line proportional to $1/f^2$. A pair of passive analog filters selects the 50 MHz output for up conversion to the final frequency. We have tested use of DDS (Direct Digital Synthesis) computations inside the FPGA to adjust the output frequency by up to 625kHz. No additional hardware is required to implement this adjustment, which is in concept limited only by the 2.5 MHz bandwidth IF filters used.

Hard-wired interlocks provide positive cutoff of the RF drive in case of external fault conditions. This feature is purposefully independent of, and not over-ridable by, the FPGA and embedded computer.

BENCH CHARACTERIZATION

Careful measurements on the input ADC channel show a white noise spectrum from 0.2 to 5 MHz of 0.22 bits/$\sqrt{\mathrm{MHz}}$ (1.0 bit rms of each sample). Clock jitter adds an additional broadband noise term of 0.0075°/$\sqrt{\mathrm{MHz}}$ (1.3 ps rms on each sample). This noise by itself, integrated over the 300 kHz-ish noise bandwidth of a closed loop cavity system, contributes only about 0.012% and 0.004° of rms cavity field noise. There is additional low frequency noise that doubles the phase noise.

FIRMWARE

The firmware inside the FPGA handles all timing and computation within the 1 ms SNS beam pulse, without intervention from the host computer. Digitized cavity input samples first have their DC offset and the 10 MHz (nominal) output of the setpoint DDS subtracted, to create an error signal. This error is passed through a digital filter

$$(\mathrm{Re}(A_P) + \mathrm{Im}(A_P)z^{-1}) \cdot (1 + A_I/(1 + z^{-2}))$$

to generate the output drive signal. The $\mathrm{Re}(A_P) + \mathrm{Im}(A_P)z^{-1}$ term provides gain and phase rotation based on the complex proportional gain coefficient A_P. The $1/(1 + z^{-2})$ term is the integrator, and A_I (always real and positive) sets its relative strength. The three multipliers are constructed as dynamic constant coefficient multipliers, for speed and reduced FPGA logic footprint. This filter and the DDS account for most of the 1829 FPGA logic elements used in the design (53% of the chip).

The FPGA only has 48 kbits of on-board RAM, in 12 banks of 4096 bits each. It is allocated as follows:

512×8	Feedforward table
1024×12	Decay wave (reflected or transmitted signal)
1024×11	Forward wave
1024×11	Reflected wave
1024×10	Transmitted wave

Each table is filled half with Real (I) and half with Imaginary (Q) components of the RF waveforms. The latter three history buffers are intended for troubleshooting and operator comfort displays; they have programmable decimation from the raw signal rate, so a full 1 ms pulse can be viewed. Firmware to replace the decimation with 16-bit resolution averaging has been written but not tested. The decay wave data is used by curve fitting routines that determine cavity detune frequency.

With a few exceptions, the signal processing path is expressed as pure synchronous RTL (Register Transfer Level) Verilog in the 40 MHz clock domain. All logic related to the host interface, including support for the low speed housekeeping devices, is expressed as pure synchronous RTL Verilog in the domain of the 25 MHz bus clock provided by the embedded single board computer. Data handoff between these two clock domains happens at the 60 Hz (max) pulse rate, so classic asynchronism issues such as metastability are minimized.

SOFTWARE

Software inside the embedded computer handles all network (100 Mbit Ethernet) traffic, EPICS, and all the computation that happens between beam pulses spaced at least 15 ms apart.

Recent work on EPICS[6] has made it portable to any software environment with POSIX compatibility. We use this capability to run EPICS R3.14.1 on an ordinary Linux (kernel 2.4) system. We combine about 1600 lines of application and hardware specific C code, with another 900 lines of custom "glue" and the existing EPICS code base, to make the final EPICS server.

A key application specific routine run on this processor is curve fitting to determine the cavity detune frequency. In most cases this information is passed over the network to run a stepper motor or water temperature control, to keep the cavity on resonance.

The 200 MHz StrongARM processor boots Linux from Flash and NFS mounts application-specific files in nine seconds. Four seconds later, the FPGA is programmed and the EPICS server is ready for client connections.

FRONT END BEAM TESTS

This module was first tested with beam on the SNS Front End RFQ in January 2003. During the limited time available, The RF system as a whole had some large (15°), repeatable phase fluctuations, although it could never be isolated to single piece of equipment. No such noise has been seen in the LLRF system before or since. Further tests are planned when this system runs both the RFQ and DTL 1 during the next beam commissioning run.

SRF TEST RESULTS

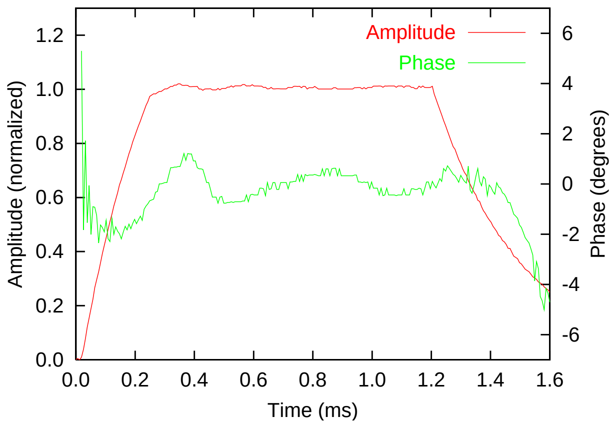

Figure 3: Cavity waveforms from a JLab SRF cavity test.

A prototype SRF cavity was successfully controlled by an interim LLRF system at JLab. An example of such operation is shown in figure 3. The system did an excellent job of removing the phase and amplitude motion within the pulse caused by Klystron power supply droop. The largest error terms were repetitive from pulse to pulse. Careful hand adjustment reduced these terms to $\pm1\%$ and $\pm1^\circ$ peak excursion. Plans exist for software to make automated and more flexible adjustments to the feedforward tables to reduce that error further. The pulse-pulse phase fluctuations during these initial tests appeared to be in the $\pm0.3^\circ$ range.

This hardware was used to demonstrate the following capabilities and features in the SRF environment:

- Control of slow cavity tuners, using curve fits to the trailing edge decay waveform.
- Single shot (burst mode) determination of cavity frequencies detuned by up to 8 kHz (15 bandwidths).
- Fast piezoelectric tuner compensation of Lorentz force detuning.
- Calorimetric (cryogenic) cavity Q measurement.

DTL TEST RESULTS

This system has run SNS DTL (Drift Tube Linac) tank 3 in open loop mode at various duty factors for conditioning the structure. Its resonance calculation was used to track the temperature changes of the tank during warm-up, with the frequency adjusted using an external synthesizer. Work is underway to replace that synthesizer adjustment with the local DDS capabilities.

CONCLUSIONS

Using FPGA technology for low-latency signal processing has immediate benefits for many applications in accelerator electronics, such as instrumentation, interlocks, and controls. Even FPGAs that are relatively small (by today's standards) can do an impressive amount of work, simplifying hardware design enormously.

The implementation presented here has shown the benefits of keeping hardware simple. The combination of an FPGA and a networked computer has made it easy to add new features, test hardware performance, and explore application of this gear to other accelerator development projects. We look forward to collaborating with other labs on further development of this technology.

More documentation on this project is posted online[7].

An ongoing, collaborative effort with ORNL and LANL[2] takes the concepts and architecture demonstrated here and engineers a system that can be fielded in quantity $\sim$100 for the long term needs of the SNS accelerator.

REFERENCES

[1] M. Champion *et al.*, "The Spallation Neutron Source Accelerator Low Level RF Control System," these proceedings.

[2] A. Regan *et al.*, "Newly Designed Field Control Module for the SNS," these proceedings.

[3] L. Doolittle *et al.*, "The SNS Front End LLRF System," Linac 2002, Gyeongju, Korea, August 2002.

[4] Bright Star Engineering, Andover, Mass., USA, http://www.brightstareng.com/arm/nanoeng.htm

[5] H. Ma, "SNS LINAC RF Control: Reference System and Phase Measurement," SNS LLRF Tech Note 3, May 2002.

[6] http://www.aps.anl.gov/epics/

[7] http://recycle.lbl.gov/~ldoolitt/llrf/