

FPGA-BASED INSTRUMENTATION FOR THE FERMILAB ANTIPROTON SOURCE*

W. Ashmanskas[#], S. Hansen, T. Kiper, D. Peterson, Fermilab, Batavia, IL, 60510, USA

Abstract

We have designed and built low-cost, low-power, Ethernet-based circuit boards to apply DSP techniques to several instrumentation updates in the Fermilab Antiproton Source. Commodity integrated circuits such as direct digital synthesizers, D/A and A/D converters, and quadrature demodulators enable digital manipulation of RF waveforms. A low cost FPGA implements a variety of signal processing algorithms in a manner that is easily adapted to new applications. An embedded microcontroller provides FPGA configuration, control of data acquisition, and command-line interface. A small commercial daughter board provides an Ethernet-based TCP/IP interface between the microcontroller and the Fermilab accelerator control network. The boards are packaged as standard NIM modules. Applications include Low Level RF control for the Debuncher, readout of transfer-line Beam Position Monitors, and narrow-band spectral analysis of diagnostic signals from Schottky pickups.

INTRODUCTION

Improvements to antiproton production, collection, and cooling are an important component of Fermilab's Run 2 luminosity campaign. Updated instrumentation can support these improvements by providing enhanced beam diagnostics, by providing finer control and monitoring of RF signal generation, and by relaxing constraints in the operation of the antiproton source. We have built and are now commissioning several instrumentation updates based on a common, FPGA-based hardware design, with minor variations specific to each application. We discuss applications—both existing and anticipated—and then implementation.

APPLICATIONS

Debuncher / Main Injector phase jump

Our first application relaxes the constraint that the Debuncher [1] rotator/adiabatic RF cavities operate at the Main Injector flattop frequency, while preserving the bucket-to-bucket transfer between 120 GeV protons leaving the MI and 8 GeV antiprotons reaching the Debuncher from the production target. Among other benefits, this added degree of freedom simplifies the planned recalibration of “8 GeV” across the accelerator complex. Instead of phase-locking to the MI RF signal, the modified Debuncher RF source digitally measures the phase of the MI RF with respect to an internal 53.1 MHz reference, then jumps the phase of its synthesized

waveform 200 μ s ahead of MI extraction, before the Debuncher cavities are energized. The Debuncher RF and incoming beam match in phase (averaged over bunches), while Debuncher and MI frequencies can differ by a few kHz. The phase-jump circuit was initially prototyped by reprogramming a digital damper board [2] and is now implemented in the hardware described below.

Debuncher LLRF upgrade

We are gradually adding features and replacing legacy modules in the Debuncher LLRF system. For example, the synthesized frequency can be ramped, allowing the cavities normally used for adiabatic debunching to prepare reverse protons at programmable momenta, to characterize chromatic effects in the AP2 transfer line (between target and Debuncher). We have programmed another module's D/A converter as an arbitrary waveform generator to precompensate amplifier distortion in the Debuncher barrier-bucket RF system. Finally, we are designing modules that will drive each cavity with its own synthesized waveform and digitize its fanback signal, allowing precise phase alignment with reduced human intervention.

Transfer line Beam Position Monitors

The 34 BPMs in the AP2 line are unusual in that they see only 53 MHz bunch structure, they see relatively small (10^{10}) numbers of particles per spill, and many signal cables see significant crosstalk from the Debuncher injection kicker. Reliable BPM measurements for reverse protons are needed to characterize the AP2 lattice, as part of an acceptance upgrade project. Readout during normal operation is also useful for monitoring orbit stability and recording secondary particle flux along the line. A prototype using Ethernet-capable digital oscilloscopes demonstrated that narrow-band processing at 53 MHz adequately suppresses the (largely 1–10 MHz) kicker noise, so we built digitizer modules—described below—with the goal of matching or surpassing the oscilloscope-based measurement. In a related effort, after demonstrating—again with an oscilloscope-based prototype—that 10^8 antiprotons can be usefully measured in a single pass when bunched at 53 MHz, we opted to instrument the Debuncher-to-Accumulator transfer line (7 BPMs) in the same manner.

Enhanced diagnostics

Many future antiproton source improvements must be implemented with minimal disruption to ongoing collider operations. By increasing the number of diagnostic measurements routinely recorded while stacking, we hope both to reduce the need for dedicated study time and to

*Work supported by the U.S. Department of Energy
[#]ashmanskas@fnal.gov

quantify more quickly the effects of changes made while running.

We implemented a Debuncher antiproton intensity monitor, by adiabatically re-bunching the beam for 20 ms toward the end of each Debuncher cycle (nominally 2 seconds) and observing the 53 MHz signals on a BPM. The readout, prototyped with a spectrum analyzer, has already been useful for scanning Debuncher apertures while stacking. A BPM board will replace the analyzer shortly. A spare channel of the same board can monitor an Accumulator BPM during each cycle's RF manipulations, providing a measurement of Debuncher-to-Accumulator transfer efficiency.

Once investigations of candidate preamplifier circuits are complete, we aim to instrument the complete set of Debuncher BPMs for 53 MHz antiproton signals, allowing orbit measurements to be made while stacking. The present system works as designed for reverse protons bunched at 2.5 MHz, but runs only during dedicated study shifts.

We have also considered using down-converters and embedded CPUs—similar to those used for BPM processing—to analyze Schottky spectra in order to report emittances, momentum distributions, and perhaps stack-tail figures of merit. The signal-processing demands are modest: typically tens of kHz of bandwidth centered at 70–300 MHz. In exchange for some programming effort, one could liberate for ad-hoc use by operators and students the costly general-purpose laboratory instruments that currently perform these measurements. So far, the only form of this concept that we're pursuing is to use an RF DDS board as a stable frequency synthesizer (as LO for a mixer) with a straightforward control-system interface, for Accumulator emittance measurement.

IMPLEMENTATION

To date, we have built boards of two types, and two additional types of boards are in the design stage. The designs are largely identical, with only the analog front and back ends varying between applications.

Common digital framework

We decided to build NIM modules, to minimize demands on rack space and infrastructure. Modules communicate with the accelerator controls system via Ethernet and with other nearby modules via an ad-hoc 4-pair LVDS bus.

At the center of each module design are an Altera Cyclone EP1C6 FPGA and a TI MSP430F149 CPU. The FPGA controls a 32MB SDRAM to capture diagnostic data and to supplement the CPU's 2 KB RAM and 60KB flash memory. A Wiznet IIM7010A daughter card (1" × 2") provides a full TCP/IP implementation that appears to the CPU as a few control registers and a circular memory buffer. A 256KB serial flash allows the CPU to store IP and MAC addresses, board serial number, persistent register contents, and the FPGA configuration image. A USB-to-UART bridge chip provides a serial command-

line interface between the CPU and a PC for initial configuration and debugging. Once the network numbers are configured, the same command-line interface is available to remote TCP/IP clients. The interface enables flash memory updates, board resets, and read/write access to registers and memories implemented by the FPGA.

The FPGA makes its internal contents, such as data capture buffers, results of signal processing algorithms, and parameters controlling waveform acquisition, processing, or synthesis, accessible to the CPU (and hence to remote clients) via an emulated A16/D32 "bus" (inspired by the VME backplane) that lives entirely within the FPGA. For example, address 0x0459 may map to a 32-bit register controlling the delay between an incoming trigger pulse and the start of data acquisition, while addresses 0x1000-0x10ff may map to a RAM holding 256 recent ADC samples. The CPU manipulates the "bus" via a dozen I/O pins connecting CPU and FPGA. The remote client manipulates the "bus" via simple commands such as "WR 0459 0000fab4" and "RD 1040" whose replies look like "R 1040 0000024c". This simple protocol, originally developed for a digital damper project [2], allows arbitrary numbers of FPGA registers to map straightforwardly into accelerator controls "devices" by means of a Java-based "Open Access Client."

The FPGA's internal "bus" can also be manipulated via 4 LVDS pairs bussed between modules, so that one network connection can control a full crate of boards.

Each board's signal processing is clocked by a 53.1 MHz VCXO with ± 20 ppm tuning range. A feedback loop can lock the VCXO to an external reference, such as the GPS-synchronized 10 MHz available in many service buildings. If a wider tuning range is needed, a VCO can replace the VCXO.

The FPGA can access about 8 front-panel TTL inputs and outputs, via buffer chips, for trigger inputs, to trigger an oscilloscope, or to decode the real-time accelerator clock event signals found in every service building.

The FPGA also drives an AD9751 10-bit 300 MSPS DAC to synthesize diagnostic waveforms.

RF DDS board

We synthesize 53MHz RF drive waveforms using AD9953 DDS chips, with frequency, phase, and amplitude under control of the FPGA. We measure the phase of two RF inputs (MI reference or cavity fanback) by sampling at 4/7 of 53 MHz with the AD9238 dual 12-bit ADC. At this sampling rate, one obtains in-phase and quadrature data on alternate samples, allowing the FPGA to calculate the phase with a divider and an arctan lookup.


This board implements the phase-jump feature, and will, after a second board production run, implement individual digitally controlled drive of each cavity. The board's diagnostic DAC serves as an orbit-synchronized arbitrary waveform generator for the RF barrier bucket.

Exploiting the $\sin(x)/x$ envelope of the DDS output, one can easily generate frequencies above the Nyquist rate by modifying the output filter. We provide a 300

MHz output from a 212 MHz clocked DDS by high pass filtering the output to pick the Fclock to $2 \times \text{Fclock}$ range.

BPM downconverter board

The 53 MHz signal from each BPM plate is demodulated to baseband using AD8348 quadrature demodulators. To suppress odd harmonics, a 65MHz low-pass filter precedes the demodulator. After a 5 MHz anti-alias filter, the in-phase and quadrature baseband signals are digitized with AD9201 20 MSPS 10-bit dual ADCs. The FPGA subtracts pedestals, computes $\sqrt{I^2 + Q^2}$, and integrates over the 1.6 μs bunch train for each of four plates (two complete BPMs). The AD8348 includes 40 dB of programmable gain, allowing a variety of signal magnitudes to be matched to the ADC span.



Swiss army knife

This design will mix and match features of the BPM and RF boards: one direct ADC input, one DAC, one demodulator channel (with ADC pair), and one DDS. It will include a more full-featured CPU such as an ARM, Coldfire, or Blackfin running Linux. We can then consider doing FFT analysis on the board, as well as

making the processor a bona-fide ACNET or EPICS front end, filling the role played by MVME processors in today's ACNET.

As cell phone and wireless technology has reached the frequency range of interest for stochastic cooling systems, we may also experiment with microwave synthesizers and demodulators on a future version of this board.

Design tradeoffs

One could imagine that a system such as this would consist only of A/D converters, a large FPGA, and D/A converters. The addition of dedicated direct digital synthesizers and analog downconverters reduces the speed, power consumption, input/output, and processing requirements of the FPGA. The use of a DDS is appropriate when generating sinusoidal outputs since the sine conversion, sinc compensation, and D/A are built into the DDS chip. The analog quadrature downconversion was chosen for similar reasons. It is possible to implement this as an entirely digital system, but the decision was made to use low cost commodity parts to offload some of the processing and to permit the use of a lower speed FPGA and A/D converter. In exchange for reduced generality, we gain more instrumented channels per dollar, reduced infrastructure demands (e.g. cooling), faster FPGA compile times, and simplified board assembly. More general-purpose instruments were essential during the prototyping phase.

ACKNOWLEDGEMENTS

We are indebted to B. Chase, G. Dugan, G.W. Foster, K. Gollwitzer, D. McGinnis, S. Pordes, M. Tigner, and S. Werkema for advice and inspiration.

REFERENCES

- [1] Tevatron I design report, Fermilab, 1984.
- [2] P. Adamson, et al., these proceedings.