

UPDATES TO THE LOW-LEVEL RF ARCHITECTURE FOR FERMILAB *

J. A. Einstein-Curtis[†], B. E. Chase, E. Cullerton, P. Varghese, Fermilab, Batavia, IL, USA
 S. Biedron¹, S. Milton, Colorado State University, Fort Collins, CO, USA
 D. Sharma², RRCAT, India

¹also at Faculty of Electrical Engineering, University, Ljubljana, Slovenia

²also at Fermilab, Batavia, IL, USA

Abstract

Fermilab has teamed with Colorado State University on several projects in LLRF controls and architecture. These projects include new LLRF hardware, updated controls techniques, and new system architectures. Here we present a summary of our work to date.

BACKGROUND

The Fermilab multi-cavity field controller (MFC) is a VXI-based LLRF controller designed for precision vector control of accelerating cavities [1]. The design has been used on several test stands at Fermilab, as well as in several parts of the main accelerator network. The design and its variants consist of a digital signal processor (DSP) for handling scaling values sent to a control system, a complex programmable logic device (cPLD) to handle the bus communications and board management, and a field programmable gate array (FPGA) to handle signal processing.

A newer design was proposed to replace the MFC as several components were being obsoleted and better performing components became available. One of the major updates was switching to a system-on-chip (SoC) module, where the control system front-end software can be run on the same chip as the signal processing logic. This leads to significant possible savings in interconnect, power, and board space usage, but also leads to other design challenges.

This new design requires a multi-disciplinary development effort, as we are developing for systems at every level of the systems design process, from firmware simulation and development up through new techniques for data transfer to the Fermilab control system. The control system is based on older technologies, so interfacing with newer developments in controls and data transfer requires a careful balance between the capabilities of the control system and the requirements of the project.

While moving away from a crate architecture provides for more flexibility and opportunities in high-performance design, it also requires more design effort up front to create a standardized system that can easily be iterated and debugged. Creating the interfaces between each internal subsystem is the greatest challenge, as they all need to be tested and verified along with their required dependencies.

* Operated by Fermi Research Alliance, LLC under Contract No. De-AC02-07CH11359 with the United States Department of Energy.

[†] jeinstei@fnal.gov, also at Colorado State University, Fort Collins, CO, USA



Figure 1: An example of a SoC MFC Board.

HARDWARE DESIGN

The SoC-based MFC design, as shown in Fig. 1, introduces several levels of complexity to an already complex system. When developing a system with discrete components, work is easily parsed out to individual developers. The SoC complicates such design flows as every function resides on the same chip. The traditional FPGA development flow relies on having well-defined boundaries and interfaces between different levels, such that each level builds on the previous.

For the Fermilab SoC design, the design builds on existing experience, but still needed to be integrated in to a new control architecture as a network-attached device (NAD) instead of a card in a crate.

Additional thought went in to using a system-on-module (SoM) for the FPGA and embedded processor component to allow for future upgrades using a standardized platform. Several companies now make SoMs that incorporate an SoC, including Critialink and Novtech. A Novtech SoM design is shown in Fig. 2.

The SoC design provides a Linux-based development framework for controls system interfacing and low-level systems debugging, which is much more flexible than a traditional crate-based architecture. In addition to the high-level protocols, the SoC architecture allows for standard interconnect technologies to be used inside the chip for data transfer, such as AXI and Altera's Qsys interconnect.

PIP-II LLRF SYSTEM DESIGN

The Proton Improvement Plan-II (PIP-II) is Fermilab's plan for improving beam intensity, generate more neutrinos, and upgrade the Fermilab accelerator complex. At the heart of this plan is a new superconducting linear accelerator. The Fermilab LLRF group has been working on new

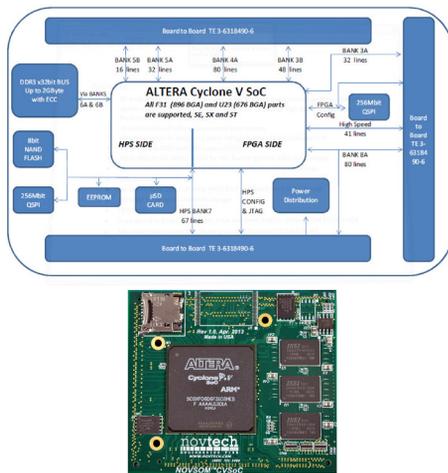


Figure 2: Novtech SoM Module.

architectures based on the use of NAD devices, integrating the control system front-end and signal processing in to one system.

The plan includes lower noise components with a flexible digital architecture to allow for ease of upgrades and deployment on a large scale. The hardware design is currently in the design phase and includes provisions for various forms of communication interfaces, including high-speed fiber and copper links.

Separating the RF components in to discrete chassis allows the LLRF group to design a minimal noise system while also compartmentalizing development on each component. In particular, links to the machine protection system (MPS) and events and timing system are simplified through the use of the standardized interfaces available in modern SoC architectures. Much like the LCLS-II design [2], the PIP-II LLRF architecture hopes to build on the LCLS-II development experience and follow newer standards for best practice. A diagram of the proposed LLRF architecture can be seen in Fig. 3.

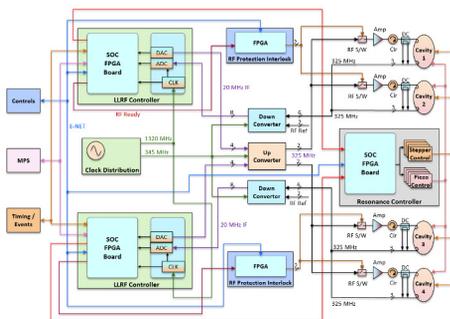


Figure 3: PIP-II LLRF Design.

DEVELOPMENT PARALLELIZATION

In addition to hardware and systems design work, additional effort has been spent on simplifying the development process itself. Firmware development can take a significant

amount due to the resources necessary for synthesis (code compilation), fitting (resource placement and utilization), and timing analysis. In addition, firmware is usually used in situations where a deterministic or real-time response is necessary, so simulation is necessary to guarantee correct operation.

We have looked in to ways to ease the burden of a firmware developer by distributing compilation in to a cluster, relieving the need for significant resources on a local machine.

This work is still in the initial design phases, as integrating with multiple vendor toolsets while not losing functionality when distributing across multiple nodes requires thought. When adding automated verification and simulation, the process can get even more complicated.

This work ties in closely with a distributed simulation effort for FEL simulation that is discussed in a previous conference paper [3]. With the wide availability of computing resources at facilities across the USA, many opportunities are open now that were not before.

FUTURE PLANS

Fermilab has several major projects that are ramping up now that LBNF/DUNE is beginning construction and integration. The LLRF group is active in several of these projects, as precision control over the accelerating field is essential. This requires tight timing requirements, accurate phase control, and an understanding of all the noise sources inherent in a system.

Digital systems add an additional source of complexity, requiring an understanding of the interactions on the borders between the analog systems that add noise and complexity to a design. Future projects are currently looking at all these concerns and are developing better performing systems for future accelerator control systems.

CONCLUSION

As new experiments and upgrades are introduced to the Fermilab facility and campus, the LLRF group is working closely with academic and corporate partners to be able to continue to meet and exceed design goals. Modern accelerator systems require a high-level of monitoring and control that is not possible without considering a full systems design, including everything from new electronics hardware to new control architectures.

REFERENCES

- [1] Fermilab Internal Document, "Multi-cavity Field Control (MFC) Module Description". http://beamdocs.fnal.gov/AD/DocDB/0036/003676/001/MFC_Description.pdf
- [2] G. Huang *et al.*, "High precision RF control for the LCLS-II" presented at NA-PAC'16, Chicago, IL, USA, Sept 2016, paper FRA2IO02, this conference.
- [3] J. Einstein *et al.*, "FEL Simulations Using Distributed Computing" IPAC'16, Busan, Korea May 2016, paper MOPMW037. doi:10.18429/JACoW-IPAC2016-MOPMW037