

EVOLUTION OF THE CERN POWER CONVERTER FUNCTION GENERATOR/CONTROLLER FOR OPERATION IN FAST CYCLING ACCELERATORS

D. Calcoen, Q. King, P.F. Semanaz, CERN, Geneva, Switzerland

Abstract

Power converters in the LHC are controlled by the second generation of an embedded computer known as a Function Generator/Controller (FGC2). Following the success of this control system, new power converter installations at CERN will be based around an evolution of the design – a third generation called FGC3. The FGC3 will initially be used in the PS Booster and Linac4. This paper compares the hardware of the two generations of FGC and details the decisions made during the design of the FGC3.

FGC ORIGINS

LHC power converters are current sources composed of three elements:

- A *voltage source* consisting of the components needed to convert electrical energy from a given source into the desired voltage on a given load. This includes protections and analogue voltage regulation.
- Calibrated *current measurement* devices to acquire the measurement needed for the current regulation. In some cases these require very high levels of absolute accuracy, stability and very low noise.
- A *Function Generator/Controller (FGC)* which integrates the voltage source and current measurement system to become a current source by providing various services including function generation, current regulation and state control.




Figure 1: FGC1 developed from 1996-1999.




Figure 2: FGC core architecture.

The FGC project was launched in 1996 [1] with a specification that included:

- Resolution and tracking error down to 1 ppm on the current.
- Timing synchronisation to better than 1 ms.
- Radiation tolerance of up to about 2 gray per year with a life expectancy of 20 years.
- Use of WorldFIP as the synchronous fieldbus for remote operation and including support for real-time control at up to 50 Hz for accelerator-level regulation loops (i.e. orbit, tune and chromaticity).
- Mechanical and electronic design by CERN and manufacture of 2000 units by industry.

More specifically the FGC is responsible for:

- Generating the synchronised function of current versus time to be used as the reference, as required by operations at the system (accelerator) level.
- Measurement and regulation of the current.
- Control and monitoring of the power converter state.
- Integration with system level interlocks, in this case the LHC Powering Interlock Controller (PIC).
- Remote (WorldFIP) and local (RS232) communication.
- Collection and reporting of FGC and power converter diagnostics.
- Logging for post mortem analysis.
- Component identification and temperature measurement using Dallas 1-wire devices.

The first milestone for the project was to deliver 75 FGCs by the end of 1999 to run the converters in the LHC magnet test facility [2][3]. This milestone was met with a first non-radiation tolerant version called the FGC1, shown in figure 1. Seventy-five were produced and were operated in various test facilities from 2000 until 2010.

Figure 2 shows the core architecture of the FGC. It is based on an integer microcontroller (MCU) and a floating point DSP. This architecture was kept unchanged in the evolution from the FGC1 to the final design used in the LHC, the FGC2.



Figure 3: FGC2 developed from 2000-2003.

FGC2 OVERVIEW

At the core of both FGC1 and FGC2 is a Motorola M68HC16Z1 integer microcontroller (MCU) clocked at 16 MHz and a Texas TMS320C32 floating point digital signal processor (DSP) clocked at 32 MHz. Both devices date from the early 1990s and are still in production today. Each processor has its own error-corrected 512KB SRAM. Furthermore, 16KB of the MCU's SRAM is visible to the DSP to enable communication between them. Programs are stored in flash memories connected to the MCU.

Aside from the addition of error detection and correction (EDAC) chips on the SRAM, the biggest change in the FGC2 was mechanical. The circuits were reformatted from four 6U x 160 mm cards to six 6U x 80 mm cards which are mounted on a passive motherboard and enclosed inside a metal cassette, as shown in figure 3.

Glue logic was moved from a Spartan FPGA to eleven smaller flash-based Xilinx 95000 series CPLDs, all mapped into the MCU memory space. Being flash-based these CPLDs do not suffer from corruption of their programming under radiation and they were hoped to be sufficiently tolerant to support the radiation anticipated in the LHC. Unfortunately recent tests have shown that they have a small cross section for latch-up in a high-energy particle field and this will become a problem when the LHC reaches nominal intensity. A replacement design is in development for deployment in 2015 for the systems affected by radiation.

The redundant acquisition of the circuit current is performed by two ADS1201 delta-sigma ADCs. These were the first delta-sigma ADCs from Burr Brown and they do not contain a digital filter. This was implemented externally in Spartan 20 FPGAs which are vulnerable to radiation, so a second analogue interface was developed for the FGCs using MAX337CWI SAR ADCs. The noise and linearity are worse than the ADS1201 but the tunnel-

based systems require less precision and these ADCs are sufficiently radiation tolerant to operate next to the LHC.

So far, the FGC2 has been a very successful design for the LHC which is currently operating at about 5% of nominal radiation levels [4]. Its strong points are, in particular:

- Use of the robust WorldFIP fieldbus for both data and synchronisation [5].
- Use of the Dallas 1-wire bus allowing identification of components and measurement of the temperature.
- Metal cassette providing a robust mechanical enclosure.

The FGC project delivered 1750 controllers for operation in the LHC in 2007 at a cost of 7 million Swiss Francs and 50 FTE years. Hardware development took six years and the software took 8 years so there was a strong motivation to re-use as much as possible when developing the next generation FGC.

FGC3 DESIGN

An evolution of the FGC2 was the chosen for the control of new power converters in the PS Booster and the new LINAC 4. The primary objective was to improve performance to allow a higher current regulation bandwidth while reducing costs by re-using as much as possible from the FGC2 development. A key choice was to not maintain the radiation-tolerant specification of the FGC2 since the FGC3 will not be deployed in radiation areas; this has a big impact on the cost of development and production.

Mechanics

The specification for the new controls electronics included the requirement that the crate be only 3U high, compared to 6U for the FGC2. Fortunately the increased density of modern components allowed the FGC3 to fit within a cassette that is half the height and width of the FGC2. The length increased from 160 to 220mm so the volume is reduced by a factor of 3. Figure 4 shows an FGC3 inside its cassette. It can run with passive ventilation and dissipates 8W. The rear connector is a




Figure 4: FGC3 developed from 2007-2011.




Figure 5: FGC3 motherboard, FGC_Ether and analogue interface daughter boards.

5-row 160-pin DIN connector (as used for VME 64 cards).

The FGC2 structure of a passive motherboard and six daughter boards is replaced by a densely integrated motherboard with two mezzanine daughter boards, as shown in figure 5.

Daughter boards

Anticipating an evolution of the network and analogue interfaces of the FGC3, these functions were moved onto two daughter boards.

When the FGC3 was designed, WorldFIP was the only network choice available, so the first network daughter board provides identical behaviour to the FGC2's WorldFIP interface. Software support for WorldFIP will be maintained indefinitely to allow an FGC3 to operate in the non-radiation areas of the LHC. Alstom stopped developing WorldFIP in 2001 so in 2010 we adopted Ethernet as our future fieldbus (see below). Thus a second variant of the network daughter board now exists (shown in figure 5) and this will be used in all non-LHC installations.

The first use of FGC3s is with small ($\pm 10A$, $\pm 30V$) power converters in the PS Booster. For these an analogue daughter board has been produced with four high precision ADC channels (ADS1274) and two 16-bit DACs (MAX5541). This ADC is a sixth order delta-sigma design with a signal bandwidth up to 50 kHz. We implement our own low-latency digital filters in the FPGA on the motherboard, running at up to 50 kbps. The next application of the FGC3 will be with fast-pulsed converters based on capacitor discharge. These require two higher speed acquisition channels (500 kbps for up to 4 ms) so a second type of analogue daughter board has been designed with two ADS1675 delta-sigma ADCs as well as two slower channels using an ADS1274.

Processors

Modern MCUs or DSPs can easily outperform the combined processing power of the MCU and DSP in the FGC2. However, maintaining the architecture of an MCU

and a floating point DSP does not present major disadvantages and allows the re-use of the FGC2 software. The required boost in performance comes from the Renesas RX610 MCU and Texas TMS320C6727b DSP. In the first design in 2007 the little-endian Renesas M32C/87 MCU was selected because of our long experience working with the Renesas M16C and M32C families. This choice required some low-level modifications in the software due to the endian difference with the FGC2's big-endian Motorola M68HC16Z1.

In 2010 the Renesas RX family became available and the Renesas roadmap presented it as the future for the M32C family. As well as having a higher clock rate and a floating point unit, the RX has selectable endianness so we could choose big endian and maintain compatibility with the HC16.

The change in MCU family introduced an eight month delay in the FGC3 project just as we were about to launch pre-series production. But the cost of this short-term delay was considered well worth the benefit given the long-term perspective of the design. More than 1000 units are expected to be produced over the next ten years.

Programmable Logic

The FGC3 logic is concentrated in one Xilinx Spartan 3AN/700 FPGA. Many of the designs were migrated from the FGC2's CPLDs with only minor modifications. This particular FGPA includes an EEPROM to hold its configuration which reduces the chip count by one and is better value than an FPGA plus external EEPROM.

One weakness in the FGC2 design was the inability to reprogram the CPLD logic remotely via the WorldFIP network, since the logic was needed for the processors to run. To avoid this restriction in the FGC3, the MCU and network interface can work without the FPGA. A new logic program can be received via the network and written by the MCU to the EEPROM inside the FPGA.

Synchronisation and communications

Synchronisation is a key requirement for the FGC. Both FGC2 and FGC3 need a 50 Hz sync pulse to discipline a phase-locked loop. In the FGC2 this is a digital PLL based on a fixed quartz oscillator while in the FGC3, clock purity is more important so a VCXO is driven by a 14-bit DAC. Thus the PLL has digital phase detection and a software PI algorithm combined with an analogue oscillator; but it still needs the 50 Hz sync.

When using the WorldFIP network, the sync pulse is simply the IRQ signal from the MicroFIP interface device. This works because WorldFIP is deterministic



Figure 6: Ethernet switch and sync pulse injector.

and the jitter is $\pm 1.5 \mu\text{s}$. Ethernet is not deterministic so a separate 50 Hz sync signal is routed to the FGC on a spare pair of wires in the UTP cable carrying the 100 Mbps Ethernet. The jitter on this sync signal is much less than 100 ns. To make this work, a 24-port sync pulse injector box is mounted under each 24-port Ethernet switch and each link is routed via the box. This is shown in figure 6.

The 50 Hz comes from a timing receiver in the gateway PC and is routed on a coax cable running in parallel with the gigabit Ethernet backbone which links the gateway to the switches. The extra cost and space required for this custom designed box is minimal and the reward is considerable: completely standard low-cost Ethernet infrastructure and components can be used. The gateway is a rack-mounted PC with two Ethernet interfaces built-in. The unmanaged switches are cheap commodity items and 100 Mbps LAN chipsets as used in the FGC3 are also widely available from numerous suppliers.

The FGC communication protocols are the same running on WorldFIP packets or raw Ethernet packets. The cycle time remains 50 Hz in both cases. Throughput is far higher with Ethernet because packets are up to 1500 bytes instead of 120 bytes and multiple Ethernet packets can be sent in the same cycle. The gigabit backbone has enough capacity to allow up to 64 FGC3s to be linked to each gateway (instead of 30 over WorldFIP) and for all to communicate at maximum rate simultaneously.

Network address

The network address of an FGC on WorldFIP (1-30) is encoded on a small passive kaptan circuit board soldered into the DB9 connector. Thus FGC cassettes may be changed but the network address (and thus the system's identity to the control system) remains with the connector.

This has proven extremely effective and reliable and to allow a similar solution with Ethernet the bottom byte of the MAC address is coded on a passive circuit board soldered into a DB9 dongle that must be plugged into the front panel of the FGC3 in order to release the system from reset. At boot up, the FGC3 reads the network address (1-64) from the dongle and uses it to form the MAC address that it writes into the SMSC LAN9221 Ethernet interface chip.

SOFTWARE TOOLS

The FGC2 software environment is based on a commercial IDE from RistanCASE GmbH called DAC and two commercial compilers; the HC16 HI-CROSS+ compiler from Hiware and the TMS320C3x/4x C compiler from Texas Instruments. Both were developed for Windows 98 and have been frozen since 1999. So despite being commercial products, neither is supported and compatibility with future versions of Windows is not assured. This is a concern given the anticipated 20 year lifetime for the FGC2.

The FGC3 software environment is based on open source tools. The IDE is Eclipse with Git for version control. The GNU gcc toolchain generates code for the

Renesas RX610. It can also generate code for the TMS320C6x family but for the moment we are using a proprietary compiler from Texas.

Some bugs were detected in gcc for the RX and the M32C (when it was being used) but each time the availability of the source code and the large community that supports the gcc toolchain allowed us to fix the bugs and generate a new compiler in less than five days. This experience encourages us to migrate to gcc wherever possible.

We have chosen NewLib as the ANSI C library for the RX because it supports multi-threaded applications and is perfectly suited to small embedded systems.

The real-time OS used on the MCU of both the FGC2 and FGC3 is called NanOS. It is an in-house development derived by simplifying $\mu\text{C}/\text{OS-II}$ [6].

CONCLUSIONS

Despite the completely different MCUs, DSPs and fieldbus technologies, the FGC3 is running the software developed for the FGC2. Both programs share a common code base and share about 90% of their source code. By maintaining the core architecture combined with the flexibility given by the programmable logic and the separation of functionality into different modules, both generations of FGC appear almost identical at the conceptual level.

The porting and debugging of the low-level functions to the new FGC3 took about 1 year, while for the higher level functions it took another 6 months. By comparison the development of the current operational FGC2 software took twelve years. By inheriting technology from the FGC2, the project saved an estimated 30 FTE years compared to starting from scratch.

As hardware design tools (PCB and FPGA) become more powerful, it has become clear that it is the software which becomes the hardest and most time consuming part of an embedded converter controller development project.

REFERENCES

- [1] J.G. Pett et al, CERN, "A Strategy for Controlling the LHC Magnet Currents", EPAC'96 p. 2317.
- [2] Q. King et al, CERN, "Developments in the High Precision Control of Magnet Currents for LHC", PAC'99 p. 3743.
- [3] Q. King et al, "The All-Digital approach to LHC power converter current control", ICALEPCS'01, THBT004
- [4] Q. King, "Status of the LHC power converter controls", ICALEPCS'09, MOB003
- [5] Q. King, "Advanced uses of the WorldFIP fieldbus for diverse communications applications within the LHC power converter control system", ICALEPCS'05, WE4B.3-20
- [6] Jean J. Labrosse, "MicroC/OS-II: The Real Time Kernel", 2002, CMP Media LLC.