

CONTROL OF FAST-PULSED POWER CONVERTERS AT CERN USING A FUNCTION GENERATOR CONTROLLER

R. Murillo-Garcia, Q. King, M. Magrans de Abril, CERN, Geneva, Switzerland

Abstract

The electrical power converter group at CERN is responsible for the design of fast-pulsed power converters. These generate a flat-top pulse of the order of a few milliseconds. Control of these power converters is orchestrated by an embedded computer, known as the Function Generator/Controller (FGC). The FGC is the main component in the so-called RegFGC3 chassis, which also houses a variety of purpose-built cards. Ensuring the generation of the pulse at a precise moment, typically when the beam passes, is paramount to the correct behaviour of the accelerator. To that end, the timing distribution and posterior handling by the FGC must be well defined. Also important is the ability to provide operational feedback, and to configure the FGC, the converter, and the pulse characteristics. This paper presents an overview of the system architecture as well as the results obtained during the commissioning of this control solution in CERN's new Linac4.

INTRODUCTION

CERN has nine accelerators, numerous experiments and various test areas. One commonality of such infrastructure is the need of a power system to energize the magnets (dipoles, quadruples, septums, etc.) and the various radio frequency cavities. In total, CERN has around 5,600 power converters, accounting for a typical consumption during exploitation of 1.2 TWh per year [1].

From an operational point of view, power converters fall into three categories:

- **Cycling:** the field, current or voltage reference value can synchronously be modified on a cycle-by-cycle basis. At CERN each cycle is a multiple of a basic period, which has a fixed duration of 1.2 seconds.
- **DC:** the reference is asynchronously modified as required during operation to reach a desired field, current or voltage.
- **Aperiodic:** operators trigger aperiodic cycles to ramp the current to maintain the required beam stability. This is the case of the LHC.

Fast-pulsed converters are used with cycling circuits to generate a pulse with a flat-top duration of only a few milliseconds and synchronous to an event such as beam injection or extraction.

This paper describes the use of the third generation of the CERN-designed Function Generator/ Controller embedded platform (FGC3) to control the four types of

fast-pulsed power converters currently being commissioned in CERN's new Linac4 (see Table 1).

The Linac4 beam is 400 μ s long, whilst the flat-top pulse is stable solely during a few milliseconds. As these two events must be synchronized to ensure the correct functioning of the accelerator, the FGC3 must guarantee time accuracy and inter-FGC precision better than 10 μ s.

Table 1: Fast-pulsed Power Converters in Linac4

Name	Peak Power	Pulse	Number*
Mididiscap	30 kW	5 ms	50
Maxidiscap	4 kW	2 ms	48
Modulator	5.5 MW	1.8 ms	14
HMinus	150 kW	1.2 ms	3

FUNCTION GENERATION/CONTROLLER

Most power converters currently being installed at CERN are controlled with FGC3s, shown in Fig. 1. This embedded device includes a mainboard, an Ethernet-based communication card[†] and an analog card with four high precision ADC channels (ADS1274) and two 16-bit DACs (MAX5541) [2]. The mainboard includes a Renesas RX610 microcontroller (running a small footprint real-time operating system called NanOS), a TI TMS320C6727 floating point DSP with a 10 kHz interrupt-driven task and a Xilinx FPGA for glue logic and peripheral handling.



Figure 1: Function Generator Controller: mainboard, Ethernet and analog cards.

The FGC3 is integrated within a CERN-designed chassis called RegFGC3 [3], shown in Fig. 2. A backplane links the FGC3 with a variety of cards dedicated to state control, analog interlock, digital interlock, measurement conditioning, etc. Communication buses over this backplane include:

- **SPIVS:** serial bus used by the FGC3 to send the reference value in digital form.
- **SCIVS:** serial bus used by the FGC3 to exchange parameters with the cards.

* Number of converters on completion of Linac4.

[†] The FGC3 can also be fitted with a WorldFIP-based communication card.

- QSPI: a serial bus supporting detailed diagnostics.

Although the number and type of cards connected to a RegFGC3 chassis depends on the needs and specifications of a given power converter, the state control card is always present. It implements the state machine for the particular power converter. It is partially driven by the FGC3 output commands such as power on, power off, reset and timing pulses (described in the section Timing Pulses), and it relays back to the FGC3 status and fault information.



Figure 2: RegFGC3 chassis for the Modulator converter.

CONTROL OF OPERATIONAL POWER CONVERTERS

The three-tier control system of the FGC3-based power converters is illustrated in Figure 3. The equipment tier encompasses the pulsed converter, RegFGC3 crate, Ethernet switches, and sync pulse injectors, whilst the top tier includes the controls applications, alarms service, etc. mainly used by the machine operators. Between the FGC3s and these applications lie the front-end computers known as FGC_Ether gateways. The gateways[‡] is a rack-mounted Linux-based computer with two network interface cards: one connected to the CERN Technical Network for access to the upper control layers and the second connected to a gigabit Ethernet network with enough capacity to manage a cluster of up to 64 FGC3s. Each FGC3 integrates a 100 Mbps LAN chipset to manage the raw Ethernet-based communication.

Control applications can perform three operations on a device: get a property, set a property and subscribe to a property. A property is a device-specific parameter through which the underlying hardware – a power converter in this case – can be configured, controlled and monitored.

Gateways receive these commands over the Technical Network using an in-house distributed communication protocol known as Controls Middleware (CMW) [4]. The commands are reformatted and forwarded to the appropriate FGC3. A task associated with the command is then executed in the FGC3 and a packet with the command response is transmitted upstream through the gateway to the control application.

[‡] In the FGC lingo and by extension in this paper the term ‘gateway’ refers to the front-end computer responsible to forward to and fro commands and responses between the upper control layers and the FGCs, thus acting as a gateway to these messages and hence the name.




Figure 3: Power converter control.

TIMING DISTRIBUTION

The correct operation of the accelerator depends on the accurate timing of the pulse generation. To this end, the General Machine Timing (GMT) plays a preeminent role by distributing timing events and telegrams over a dedicated communication network to timing receiver cards, driven by a UTC-synchronous 40 MHz clock with small jitter and known delay [5,6]. One such card installed in all the gateways is a PCI variant known as the CTRI. These timing packets provide information on the machine state; post-mortem; the cycle selector of the next cycle within the super-cycle; and the delay in milliseconds of the next event, be it a start of cycle, beam injection or extraction.

In addition, the CTRI cards in all the gateways are configured to generate a 50 Hz synchronization signal on one of its four outputs. This is linked through a coax cable to the first Pulse Injector. This CERN-designed 24-port box injects the sync signal onto a spare pair of wires of the 100 Mbps Ethernet cable. Thus, a single Cat7 cable from the Ethernet switch to the FGC3 device coalesces the data communication link and the synchronization signal (Fig. 3), which saves significantly on cables and connectors. The combination of an Ethernet switch and a Pulse Injector is known as an FGC_Ether star point [7]. Up to three star points can be daisy chained using a backbone of Gigabit Ethernet and a coax cable to enable one gateway to support up to 64 FGC3s.

In the FGC3, the 50 Hz sync signal is utilized to discipline the PI-based Phase Locked Loop (PLL), which synchronizes a 25 MHz Voltage Controlled Crystal Oscillator (VCXO) through a 14-bit DAC. The resulting clock has a jitter below 40 ns and accuracy better than 2.5 μs [8]. The clock is fanned out to the various integrated circuits in the FGC3 (MCU, DSP, FPGA,

Pre-Press Release 23-Oct-2015 11:00

Copyright © 2015 CC-BY-3.0 and by the respective authors

ADCs, and DACs) making the firmware and the interrupt-driven software tasks in the MCU and DSP fully synchronous with the GMT.

FGC3 SOFTWARE

The first software class written for the FGC3 devices provides function generation and regulation [9], property management, event and analog signals logging, ADCs and DACs calibration, monitoring and diagnostics of the converter, and an interface with the interlock system.

This code base has been reused as the foundation for a new class of software used to control fast-pulsed power converters. Function generation and regulation is not needed and thus have been removed. Instead a mechanism to synchronize the converter with the GMT has been developed to achieve the required time constraints. Due to size restrictions, this section only presents two of the features specific to this software.

Timing Pulses

The FGC3 commands the electronics of the fast-pulsed power converter with timing pulses. These control the use of the capacitors stocking the energy needed to generate the current or voltage pulses.

Figure 4 provides a typical sequence of events starting with the GMT transmitting a timing event over the timing network to the CTRI in a gateway. The software running in the gateway formats the event and broadcasts it to all the FGC3s connected to the FGC Ethernet network. This information includes the current UTC and millisecond time, the type of event and the remaining time in milliseconds until the event will occur. The FGC3 retrieves this last value and writes it into the FPGA register TIME_TILL_EVENT. This register implements a free-running counter that starts down counting at 1 MHz at the start of the next millisecond. Thus, the exact event time coincides with the register reaching the value zero.

The FPGA also provides a peripheral consisting of a set of GPIOs that can generate timing pulses. Each output channel has two associated registers: PULSE_ETIME defines the time of the rising edge of the pulse with respect to the TIME_TILL_EVENT and PULSE_WIDTH specifies the pulse duration in microseconds.

As represented in the example of Fig. 4, the timing event is received 875 ms in advance. The timing pulse A0 controls the capacitor charging time, which starts 410 ms prior to the arrival of the beam and lasts for 400 ms. The timing pulse A1 controls the capacitor discharging time spanning from 4 ms before the event to 1 ms after the event. Finally, A2 is used to synchronize the acquisition measurement sampled precisely at the event time. This measurement is then made available to the operators.

To complete the above description it should be noted that the Modulator converter, employed to power the radio-frequency klystrons, requires a fourth timing signal to trigger an active bouncer used to stabilize the pulse.

The FPGA registers PULSE_ETIME and PULSE_WIDTH for each output channel (A0, A1, A2) are initialized from FGC3 properties, which are




Figure 4: Simplified chronogram of timing pulses for fast-pulsed converters.

configured based on the requirements of the fast-pulsed power converter. With this strategy, converters with different topologies can be controlled homogeneously.

Pulse State Machine

The main power converter state machine encoded in the FGC3 [10] is driven by external stimuli (inputs, commands, etc.) and accounts for various working modes: direct, cycling, economy, etc. Fast-pulsed converters operate in the cycling state, allowing operators to specify a different pulse reference for each cycle. This state has been extended with a pulse state machine (Fig. 5) to sequence the generation of pulses. The following functionality is implemented within each pulse state:

- Waiting: waits for a timing event.
- Preparing: verifies if the pulse for the current cycle is enabled and latches the reference and timing parameters.
- Setting: configures the timing pulses as described in the previous subsection; sends the reference value to the converter; and if a polarity switch is present and its position does not match the sign of the reference, the output command to change the polarity switch is activated.
- Reporting: gathers the current and voltage measurements and verifies if a fault condition has occurred. This information is published and made available to the operators.
- Fault: logs the error condition.

Some of these states have an associated timeout. If the timeout expires before the onwards transition is asserted, the state machine defaults into the Fault state, where the appropriate information is logged. This is then followed by a transition to the Waiting state to handle the next

timing event. This prevents the state machine from locking up, allowing subsequent timing events to be handled and thus improving the software resilience.




Figure 5: Pulse state machine.

PERFORMANCE

Since the FGC3s are synchronized with the General Machine Timing, the timing pulses and consequently the current or voltage pulses are generated synchronously. This can be observed in Figure 6, which shows pulses with various references sampled from six circuits installed in Linac4. Note how the overlapping flat-top stability is of merely 1.5 millisecond. Longer pulses would require an over-dimensioning of the power system, resulting in higher costs.




Figure 6: Six Maxidiscaps pulsing synchronously.

The FGC3 timing accuracy is demonstrated in Fig. 7, which includes three oscilloscope snapshots. The left and middle ones were taken in the laboratory. The signals probed comprise: the pulse; the exact event time as output by the CTRI card in the gateway; and the acquisition timing pulse (A2) defining the event time from the FGC3's perspective. It can be observed how the difference between these latter two signals is less than 200 ns, well below the requirement of 10 μs. The right image shows the acquisition timing pulses (A2) of three different Maxidiscap converters connected to one gateway and a Modulator connected to a different gateway in Linac4. The inter-FGC timing precision among the three Maxidiscaps is negligible whilst between these and the Modulator is only 186 ns, due to differences in cable length. In any case these values are also below the 10 μs requirement.

FUTURE WORK AND CONCLUSIONS

Two important software features are yet to be developed:

- Fast sampling: the software must be adapted to provide support for a new analog interface which is based on four LTC2378 SAR ADCs sampling at a rate of 500 ksp/s. This will allow acquiring the pulsed signals with a better time resolution.
- Slow regulation: the Modulator converter works in open loop, i.e. the flat-top pulse is not regulated. The FGC3 shall implement an algorithm that trims the reference value on a pulse by pulse basis to minimize the error.

Despite the above, the core functionality of the software is fully operational and has been integrated with the tools in CERN's new Linac4 accelerator to control the currently installed. When complete, there will be 115 circuits in Linac4 controlled by this software.

Working in collaboration with the converter designers has proved essential to successfully reach this goal. Their feedback was crucial to provide an intuitive interface and to correctly parameterize each circuit. In return, an unprecedented level of remote diagnostics and monitoring is now available, which will shorten the intervention times during machine exploitation.




Figure 7: Precise synchronization between the pulse and the event proving time accuracy and precision.

Pre-Press Release 23-Oct-2015 11:00

Copyright © 2015 CC-BY-3.0 and by the respective authors

REFERENCES

- [1] CERN Engineering Department:
http://ern.ch/en-dep/CERN_Energy_Consumption/2012/ElectricityFlyer2012.pdf
- [2] D. Calcoen et al., “Evolution of the CERN power converter function generator/controller for operation in fast cycling accelerators”, ICALEPCS, Grenoble, France (2011).
- [3] M. Di Cosmo et al., “The new modular control system for power converter at CERN”, ICALEPCS, Melbourne, Australia (2015).
- [4] K. Kostro et al, “The Controls Middleware (CMW) at CERN – Status and Usage”, ICALEPCS, Gyeongju, Korea (2003).
- [5] J. Serrano et al., “Nanosecond level UTC timing generation and stamping in CERN’s LHC”, Proceedings of ICALEPCS, Gyeongju, Korea (2003).
- [6] J.C.Bau et al., “Managing the real-time behavior of a particle beam factory: the CERN Proton Synchrotron complex and its timing system principles”, IEEE Transactions on Nuclear Science 45, 1998.
- [7] S. Page et al., “Migration from WorldFIP to a Low-Cost Ethernet Fieldbus for Power Converter Control at CERN”, ICALEPCS, San Francisco, USA (2013).
- [8] M. Magrans et al., “Clock and Event Synchronization of the Function Generation/ Controller”, ICALEPCS, Melbourne, Australia (2015).
- [9] Q. King et al., “Function generation and regulation libraries and their application to the control of the new main power converter (POPS) at the CERN CPS”, ICALEPCS, Grenoble, France (2011).
- [10] Q. King et al., “CCLIBS: The CERN Power Converter Control Libraries”, ICALEPCS, Melbourne, Australia (2015).