MIGRATION FROM WORLDFIP TO A LOW-COST ETHERNET FIELDBUS FOR POWER CONVERTER CONTROL AT CERN

S. Page, Q. King, H. Lebreton, P. Semanaz, CERN, Geneva, Switzerland

Abstract

Power converter control in the LHC uses embedded computers called Function Generator/Controllers (FGCs), which are connected to WorldFIP fieldbuses around the accelerator ring. The FGCs are integrated into the accelerator control system by x86 gateway front-end systems running Linux. With the LHC now operational, attention has turned to the renovation of older control systems as well as a new installation for Linac 4. A new generation of FGC is being deployed to meet the needs of these cycling accelerators. As WorldFIP is very limited in data rate and is unlikely to undergo further development, it was decided to base future installations upon an Ethernet fieldbus with standard switches and interface chipsets in both the FGCs and gateways. The FGC communications protocol that runs over WorldFIP in the LHC was adapted to work over raw Ethernet, with the aim to have a simple solution that will easily allow the same devices to operate with either type of interface. This paper describes the evolution of FGC communications from WorldFIP to dedicated Ethernet networks and presents the results of initial tests, diagnostic tools and how real-time power converter control is achieved.

EVOLUTION OF REQUIREMENTS

From LHC to CERN's Cycling Accelerators

Power converters in the LHC are controlled by the second generation of Function Generator/Controller (FGC2), which are connected to the controls network via 2.5 Mbps WorldFIP fieldbuses and gateway computers. The converter control requirements of the cycling accelerators at CERN are different and needed to be considered when selecting an appropriate fieldbus for the third-generation FGC (FGC3) [1].

Physical Requirements

Many of the power converters in the LHC are exposed to some level of radiation due to their proximity to the accelerator. In order to operate within that environment, the FGCs that control them had to be constructed using components that were radiation-tolerant or which implemented countermeasures to allow their use (e.g. error-corrected memory). The MicroFIP interface chipset met this requirement, allowing the use of the WorldFIP fieldbus.

By contrast, power converters and their controllers within the cycling accelerators are located away from the beam lines so there is no requirement for radiation tolerance of fieldbus components.

The installation of power converters along the LHC tunnel made a linear fieldbus, such as WorldFIP, attractive. A single WorldFIP cable can run up to 500m along the tunnel, extendable further using copper and fibre-optic repeaters.

In the cycling accelerators, power converters are mostly grouped together in dedicated equipment halls, for which a star topology is more convenient.

Logical Requirements

The LHC is a storage ring and the magnet currents remain stable for long periods, so it is not necessary to take high-frequency acquisitions during normal operation. As a result, the 2.5 Mbps speed of the WorldFIP fieldbuses is sufficient.

The power converters in CERN's cycling accelerators pulse every few seconds and high-frequency acquisitions must be available to allow their performance to be monitored and tuned. Thus, a fieldbus with significantly increased throughput is required.

In addition, the synchronisation of the current reference functions in the cycling accelerators is more demanding. In the LHC, 100 µs is acceptable while in the cycling accelerators, the requirement is 1 µs.

CANDIDATE FIELDBUSES

There is a clear trend towards Ethernet-based fieldbuses and a review of international standards in 2009 identified Ethernet Powerlink and EtherCAT as potential candidates. However, both require specialised interface devices if precise synchronisation is to be achieved and neither is well matched to the FGC protocol running over WorldFIP.

CERN has lead the development of an extension to gigabit Ethernet known as White Rabbit [2], that allows nodes to be synchronised to better than 1 ns over a large geographical area. When the technology was reviewed in 2009, it was not ready for operation and the bandwidth and synchronisation capabilities far exceeded the requirements. White Rabbit is also expensive: if 64 FGC3s were connected to one gateway, the estimated cost per node, excluding cabling, would be 450 Swiss Francs.

In 2013, a CERN working group on Ethernet-based fieldbuses published its report on the state of all fieldbuses at CERN and identified candidate fieldbuses for future applications [3].

In the end, the simplest "low-tech" solution was also the cheapest. Standard 100 Mbps copper Ethernet was selected for communication with the FGCs, while a separate 50 Hz sync pulse is fanned out from the gateway to provide synchronisation. The same FGC protocol used over WorldFIP is encapsulated in raw Ethernet frames allowing significant code reuse.

ISBN 978-3-95450-139-7

The only optimisation is the use of one spare pair of wires in the Ethernet cable going to the FGC to transport the 50 Hz sync signal (only two of the four pairs are required for 100 Mbps Ethernet). This saves significantly on cabling and connectors at the cost of developing a simple Ethernet Pulse Injector (EPI), which introduces the sync signal into the Ethernet cable.

Taken together, the use of FGC protocols over standard 100 Mbps Ethernet and the distribution of a 50 Hz sync signal using the EPI are called "FGC_Ether".

FGC_ETHER ARCHITECTURE

The architecture of an FGC_Ether fieldbus is illustrated in Figure 1. The gateway computer includes a CERN-designed CTRI timing receiver, which is configured to provide interrupts for the software and the 50 Hz sync signal for the FGCs. This is transmitted on a coax cable in parallel with the 1 Gbps Ethernet backbone that links the gateway to the switches. Each switch is mounted above a pulse injector to form a star point supporting up to 22 FGCs. Up to 64 FGCs can be connected to one FGC_Ether fieldbus, requiring a minimum of three star points. More than three can be used if the geographical distribution of converters favours such a topology.

Table 1 gives the 2012 costs of the main FGC_Ether components. The cost per node for a fully populated fieldbus, excluding cabling, is 90 Swiss Francs.

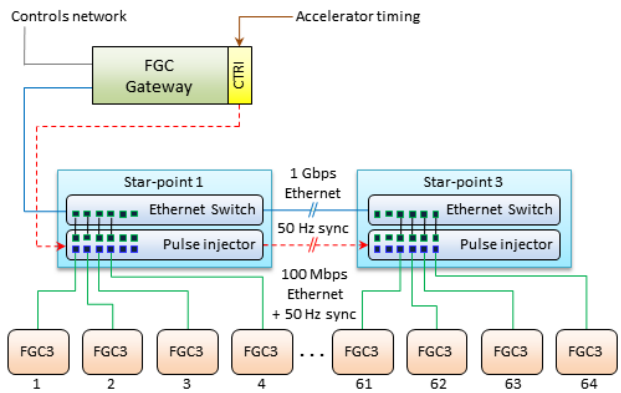

Figure 1: Overview of FGC_Ether fieldbus architecture.

Gateway

The FGC_Ether gateway computers are Kontron PCI760 2U industrial PCs with two external Ethernet interfaces, the first connected to the CERN controls network and the second to the FGC_Ether fieldbus. The gateways run the 64-bit Scientific Linux CERN 6 (SLC6) operating system. The MRG real-time kernel is used, however the Ethernet driver does not exploit any real-time features. In order to synchronise with both UTC and events within the rest of the control system, each gateway has a CTRI PCI timing receiver.

The gateways' software is based upon a framework known as the Function Generator/Controller Daemon (FGCD). A significant quantity of code is shared between the WorldFIP and FGC_Ether variants of FGCD.

Table 1: 2012 Costs in Swiss Francs (CHF)

Equipment	CHF	Comment
Gateway computer	2500	Includes CERN-designed CTRI timing receiver
Ethernet switch	170	Cisco 24+2 port unmanaged switch (SF 102-24)
Ethernet pulse injector	240	Designed by CERN and manufactured by industry
FGC3 chipset	30	SMSC LAN9221 controller and supporting components

Ethernet Switch

The Cisco SF 102-24 unmanaged Ethernet switch was chosen for the first deployment of FGC_Ether. It has a relatively low cost and features 24 100 Mbps ports that are used for the links to the FGCs and two 1 Gbps ports that are used to connect the gateway to the first switch and to provide a high-bandwidth backbone between switches. It should be noted that it is a standard Ethernet switch with no specific real-time features. Generally, only 22 of the 24 100 Mbps ports are used in order to allow the possibility of migrating to a switch with 24 1 Gbps ports in the future.

Ethernet Pulse Injector

The CERN-designed Ethernet Pulse Injector (EPI) is a fan-out module for the 50 Hz sync signal with the particularity that it uses pairs of RJ45 connectors that allow 100 Mbps Ethernet to transit the unit while adding the sync signal on a spare pair of the outgoing cable. The EPI is a 1U rack-mounted box that is installed directly below the switch, as shown in Figure 2.

Figure 2: FGC_Ether star-point with switch and EPI.

Short patch cables link the switch ports to the input ports of the EPI and longer cables link the EPI output ports to the FGCs in the vicinity. Four BNC outputs on the rear panel allow EPIs to be daisy-chained. The latency of a sync pulse transiting the EPI is 60-120 ns and the latency for 100 m of cable is about 500 ns.

FGC3 Chipset

The FGC3 network interface is based on a COTS 100 Mbps LAN controller [4]. There are many similar products from a range of vendors; the SMSC LAN9221 was selected because it has good support for big endian processors. The interface card also has an optocoupler to receive the 50 Hz sync pulses, plus a DAC driving a 25 MHz VCXO for the local clock.

ISBN 978-3-95450-139-7

Table 2: Limitations of WorldFIP and FGC_Ether

	WorldFIP	FGC_Ether
Max. FGCs per fieldbus	30	64
Max. cable length per link	500 m	100 m
Fieldbus cycle time	20 ms	20 ms
Max. throughput: Gateway → FGC	3 KB/s	75 KB/s
Max. throughput: FGC → Gateway	3-6 KB/s	750 KB/s

Addressing

Up to 64 FGCs may be connected to an FGC_Ether fieldbus. Since raw Ethernet frames are used, addressing is by MAC address.

The gateway computer uses the physical MAC address of its Ethernet interface that the FGCs can detect when they are started by waiting for the reception of a broadcast time frame.

It would be undesirable to use individual MAC addresses for every physical FGC since this would require a mechanism for the gateways to discover them and they must be replaceable with identical units should they fail in operation. Instead, addresses are associated with logical devices within the accelerator control system. This is achieved using a small dongle with the last byte of the address (1-64) encoded in a DE-9 connector. The rest of the address is zero, with the exception of a bit within the high byte to indicate that the address is locally managed. The dongles can be exploited in three different ways:

1. For modular converters, which can be quickly repaired by replacing a module, the dongle is fixed to the power converter with a small chain and connected to the FGC when it is installed, ensuring that the controller for the power converter is always at the same address.
2. For converters where there is an operational spare, the dongle may be moved to the FGC in the spare converter at the same time as the intervention to connect the spare to the circuit.
3. Finally, the PS Booster is a special case because only a subset of the magnet circuits is needed at any time, so there are fewer power converters than circuits. To support this, an address (and therefore a dongle) is associated with each circuit, rather than each power converter. This allows data associated with a particular circuit to be maintained within the control system. When a converter is connected to a circuit, the associated dongle is connected to its FGC, allowing it to be addressed with the corresponding device name within the control system.

Limitations

The key design objectives of FGC_Ether were to provide adequate communications throughput and sub-microsecond synchronisation at low cost, and to be available for operation in mid-2011. It is not intended to be a general-purpose fieldbus, though some ideas and techniques may be applied to other similar applications. The limitations of FGC_Ether are compared with WorldFIP in Table 2.

FIELDBUS CYCLE AND PROTOCOL

In order to allow as much common software as possible, the FGC_Ether protocol is very similar to that used over WorldFIP. Both cycles have a period of 20 ms and each includes broadcast transmission of UTC time and events from the gateway and a status structure from each FGC. Additionally, several milliseconds of the cycle in both cases are reserved for the transfer of aperiodic commands and responses.

The Maximum Transmission Unit (MTU) of Ethernet is significantly larger than WorldFIP (1500 vs. 122 bytes). This is exploited to include real-time current references with the transmission of the time structure, reducing the number of transmissions necessary per cycle. It also allows commands received from the control room and their responses to be transferred in a reduced number of frames.

PLL AND SYNCHRONISATION

In the LHC, the synchronisation requirement is undemanding and was easily achieved using a digital phase-locked loop (PLL) implemented in the FGC2's microcontroller. It is disciplined by the time of arrival of the broadcast time packet sent by the gateway at the start of each 20 ms WorldFIP cycle [5], which has a typical jitter of 3 µs.

With the fast-cycling accelerators, power converter synchronisation needs to be much better. This is achieved using the 50 Hz sync signal, which is fanned out from the gateway to the FGCs and which has very low jitter (< 1 ns). If this signal becomes unavailable, then a degraded synchronisation mode is activated automatically, using the time of arrival of the Ethernet broadcasts. These depend upon the gateway's Ethernet stack, which is not real-time, so the jitter in the FGC's network interrupts (NET IRQ) is 20 µs. It is also far from Gaussian; it is typically around 4 µs, but with random broadcasts delayed by up to 15 µs.

The PLL software attempts to filter out the delayed NET IRQs, and can deliver surprisingly good behaviour in the short term, but over a period of hours, the phase of the local clock will inevitably wander up to 20 µs. For some power converters, this will be unacceptable for operation while for others it can be tolerated.

Figure 3 shows the FPGA logic created to support the PLL software. The FGC's network interface supplies two sync signals to the FPGA logic: EXT SYNC (the external high-quality 50 Hz sync signal) and NET IRQ (the arrival time of Ethernet frames). The software can calculate the phase error by reading the EXT SYNC TIME and NET IRQ TIME registers and taking the difference with the INT SYNC TIME register. The PLL module produces two synchronised clocks for the rest of the system:

ISBN 978-3-95450-139-7

25 MHz and 50 Hz. The INT SYNC TIME register is written once by the software at start-up to launch the production of the 50 Hz clock, synchronised with the external 50 Hz. After that, it is only responsible for controlling the phase using a standard PI algorithm to set the DAC that drives the VCXO. It was found that the DAC's 14-bit resolution was inadequate, so it is oversampled at 625 kHz by the FPGA logic, which adds at least another four bits of effective resolution.

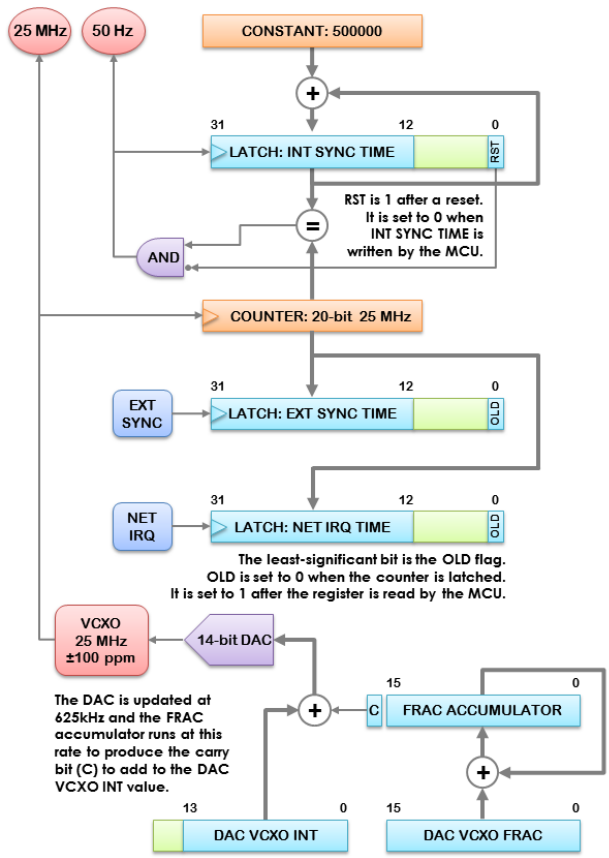

Figure 3: PLL architecture.

FGC TO FGC DATA EXCHANGE

Until now, data has only been exchanged between the gateway and the FGCs, however, the larger bandwidth of Ethernet opens up the possibility to exchange data between FGCs at 1 kHz. There are two use cases for this: master/slave FGCs and remote measurement FGCs.

There are several applications in which multiple converters will need to work together, either in series for increased voltage, or in parallel for increased current. These could be managed by multiple FGCs in a master/slave configuration provided they can exchange information reliably in real-time. The most demanding use case will be the main power converters of the SPS accelerator in which twelve 12 MW converters operate in series with the magnets.

The other application for FGC to FGC data exchange will use an FGC as a remote measurement device. An application at CERN requires the regulation of magnetic field measured by hall probes in several magnets, about 50 m from the power converters. FGC to FGC data exchange will allow an FGC to be used as a remote measurement unit on behalf of the regulating FGCs, thus avoiding long analogue cables. The measurement FGC has four ADCs, so it will be able to measure the field in four nearby magnets. By broadcasting its measurement, it will act on behalf of four regulating FGCs.

At the time of writing, FGC to FGC data exchange is under development.

DIAGNOSTICS

As FGC_Ether is based upon standard Ethernet, a significant number of diagnostic tools are readily available. A script has been written that remotely invokes tcpdump on the gateway computer to capture and filter Ethernet frames from the fieldbus, before returning the resulting data to a local instance of Wireshark. This uses minimal resources on the gateway computer, while allowing complex analysis. An FGC_Ether module extends Wireshark, decoding frames from the fieldbus and presenting their contents to the user in a clear and easily filterable fashion.

Finally, the script was integrated into the main expert tool used for manipulating FGCs within the control system, allowing any device to be selected and its fieldbus data captured.

These diagnostic tools have proven invaluable during the development of the fieldbus and represent a significant improvement over the FIPWatcher tool that is used to analyse communication issues on the WorldFIP in the LHC.

CONCLUSIONS

Early experience from operational use of FGC_Ether is encouraging. Up to 3000 nodes are expected to be deployed over the next decade, so the fact that FGC_Ether is based upon standard low-cost COTS components will provide significant savings and provide security against component obsolescence.

The compatibility with the FGC protocols used with WorldFIP allowed significant code reuse, accelerating development and reducing maintenance costs for the future.

REFERENCES

[1] Q. King, "High Speed Communications for FGC3", August 2010, https://edms.cern.ch/document/1083320/

[2] J. Serrano et al., "The White Rabbit Project", ICALEPCS'2009, Kobe, Japan, October 2009.

[3] Q. King et al, "Fieldbuses for Control Systems at CERN", June 2013, https://edms.cern.ch/document/1262875/

[4] SMSC, "High-Performance 16-bit Non-PCI 10/100 Ethernet Controller with Variable Voltage I/O", http://www.smsc.com/media/Downloads_Public/Data_Sheets/9221.pdf

[5] Q. King, "Advanced uses of the WorldFIP fieldbus for diverse communications applications within the LHC power converter control system", ICALEPCS'2005, Geneva, Switzerland, October 2005.

ISBN 978-3-95450-139-7