Technical Article

The use of MicroTCA-based data acquisition systems in linear accelerators

Nutaq® INNOVATION TODAY FOR TOMORROW®

The use of MicroTCA-based data acquisition systems in linear accelerators

Written by Tristan Martin

Table of Contents

1 Introduction 3

2 MicroTCA and FMCs 4

3 Implementing fpga-based low-level RF control loops 8

4 Global orbit feeback 15

5 Beam position monitor (BPM) 16

6 Orbit controller (OC) 17

7 Orbit correction magnets (OCM) 18

8 Loop requirements 19

9 Conclusion 20

10 References 21

List of Figures

Figure 1: A Nutaq MicroTCA data acquisition system 3
Figure 2: Perseus601x block schematic 4
Figure 3: Perseus 611x block schematic 5
Figure 4: Nutaq's MI125-32 ADC FMC 6
Figure 5: Double-stacked FMCs 7
Figure 6: Key components of an LLRF control system 8
Figure 7: A Nutaq MicroTCA system 9
Figure 8: LLRF conceptual design of ALBA synchrotron light source 10
Figure 9: Direct RF signal sampling 11
Figure 10: Nutaq's Model-Based Design Kit (MBDK) and Xilinx® System Generator for DSP™ 12
Figure 11: Nutaq solutions for accelerated development 14
Figure 12: Global orbit feedback system 15
Figure 13: Functional block diagram of Nutaq's QSFP(+)/SFP(+) transceiver modules 19

1 Introduction

FPGA-based data acquisition systems may seem very specific in nature. In fact, they can be used in a wide variety of applications, including:

- Pre-clinical imaging (PET, MRI, MicroPET)
- Radio astronomy
- Geolocation
- Cargo security inspections
- High-energy physics
- MIMO radar

In this paper, we investigate the use of MicroTCA-based data acquisition systems in high-energy physics applications, namely linear accelerators (linacs). MicroTCA data acquisition systems are used for two key functions within linacs: low-level radio frequency (LLRF) control systems and beam position monitoring (BPM). LLRF control systems manage the amplitude and phase of the electromagnetic fields inside an accelerator's RF cavities. Proper electronics are essential to ensure the highest possible beam quality. Similarly, BPM systems are essential for maintaining beam stability.

Manufacturers like Nutaq combine advanced mezzanine cards (AMCs) with VITA 57.1 FPGA mezzanine cards (FMCs) to create a wide variety of MicroTCA-based data acquisition systems. These systems are known for their high channel density and low cost.

Figure 1: A Nutaq MicroTCA data acquisition system

2 MicroTCA and FMCs

FPGA boards are normally used to implement LLRF digital control loops and beam positioning algorithms. Nutaq's Perseus 601x AMC, for example, is based on the Virtex-6 FPGA and offers high-performance, high-bandwidth, low-latency processing. Its MicroTCA software tools are highly suitable for the development of control algorithms for linacs and other high-energy physics applications.

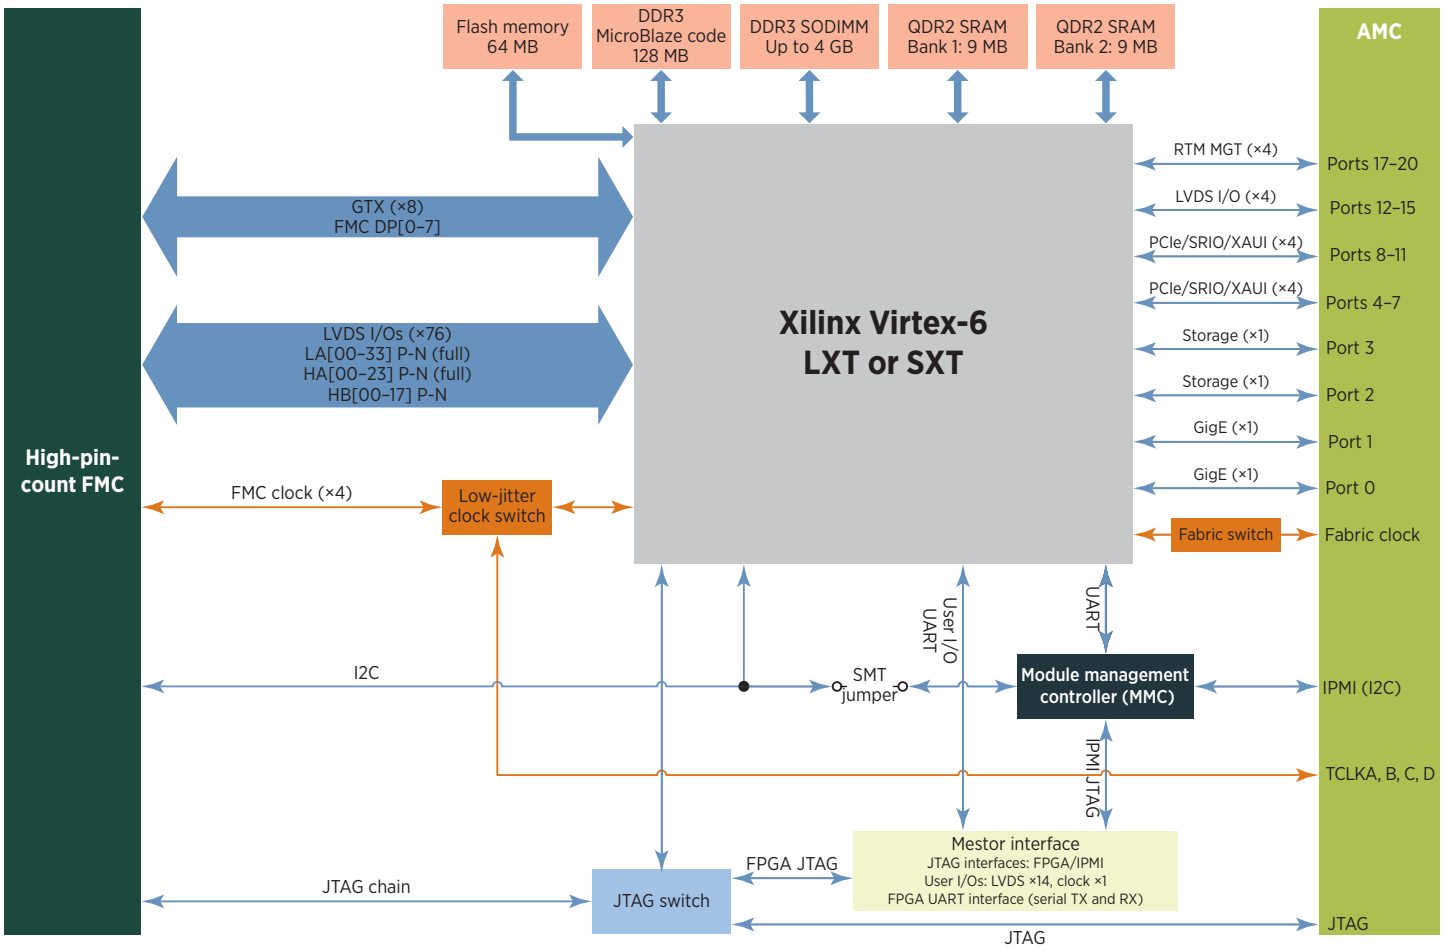

Figure 2: Perseus601x block schematic

High-pin-count (HPC) FMC sites enable the addition of high-speed and high-channel-count analog or digital input/output (I/O) cards to the FPGA carrier. Observed trends in the last few years indicate the need for higher density platforms. The term "density" is used to refer to the number of I/O channels (both analog and digital) per board or per FPGA. The need for increased density, along with the participation of many contributors involved in high-energy physics research, led to the creation of the PICMG MicroTCA.4 R1.0 standard. This standard defines a new set of specifications that provide shelf and AMC carrier support for rear transition modules (RTM) in order to add additional I/O expansion capabilities.

Nutaq's Perseus 611x was developed to address the need for increased I/O density and high-speed interconnects. The Perseus 611x is a double-width AMC featuring a Virtex-6 FPGA, two HPC FMC sites, and RTM expansion. When carriers are designed with RTM capabilities, it is useful to route as many of the FPGA's GTX high-speed transceivers (and optionally Fat Pipes Regions 1 and 2) as possible to the RTMs. The flexible interconnect architecture makes it possible to establish dedicated high-speed and low-latency communication links between several boards. These links can also be used to route data to and from external subsystems using different transport protocols, including PCIe, SRIO, SAS, SATA, 10-GbE (XAUI), SFP, and XFP.

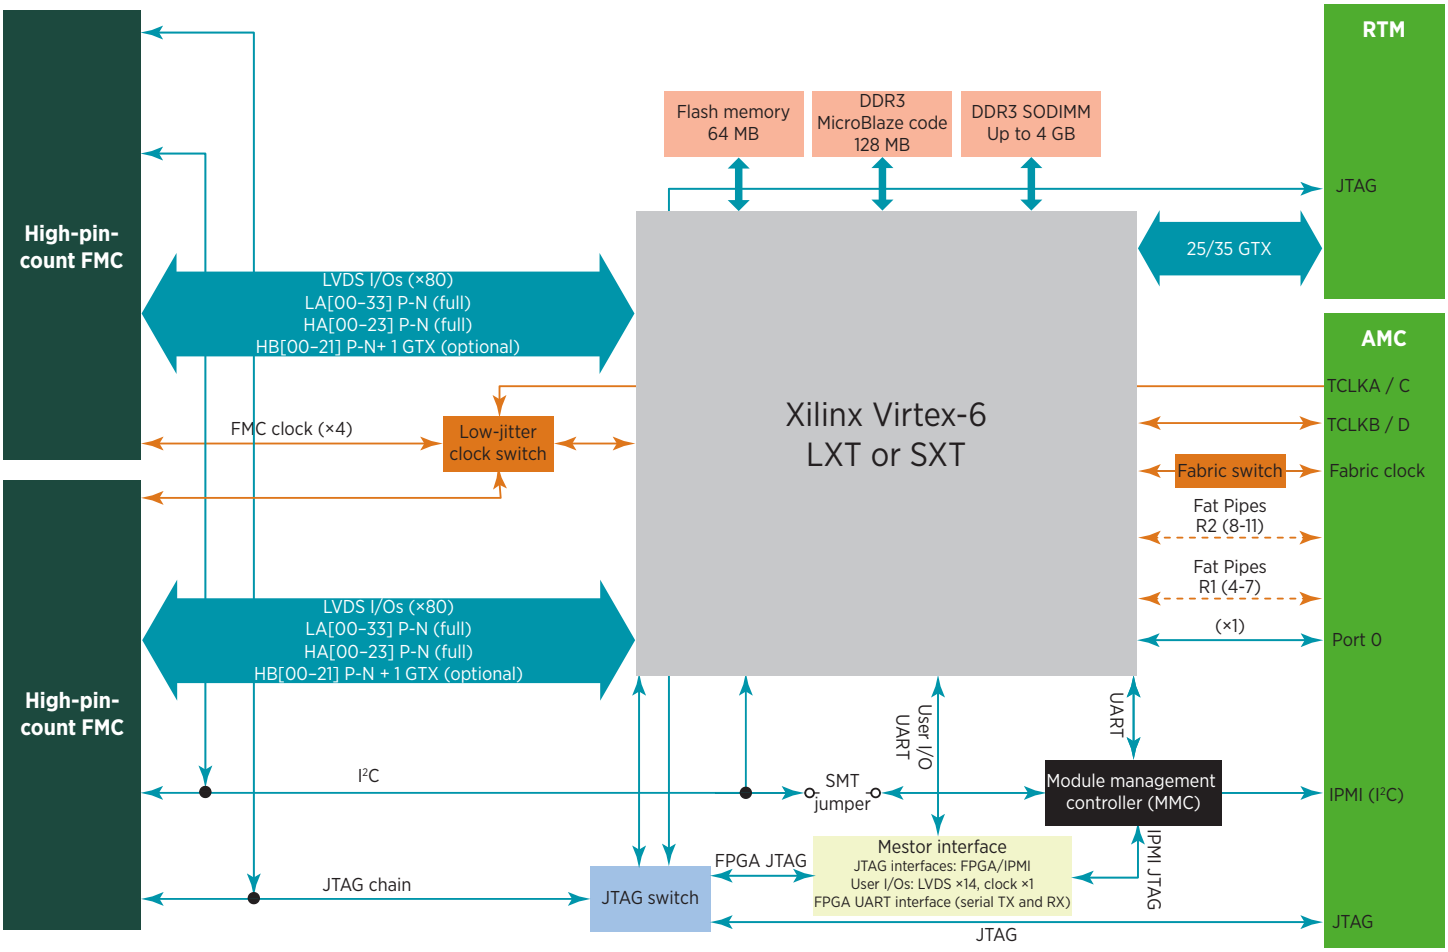

Figure 3: Perseus 611x block schematic

Additionally, a double-width card like the Perseus 611x doubles the number of FMC sites and doubles the front-panel I/O density when compared to traditional single-width AMCs like the Perseus 601x. By interfacing twice as many channels to a single FPGA, large-scale acquisition systems like those used in linacs can benefit from lower per-channel costs, as well as higher system density. Manufacturers offer a wide range of FMC cards equipped with different channel configurations for their digital-to-analog converters (DACs) and analog-to-digital converters (ADCs).

Nutaq's engineering team was recently put to the challenge of designing a high-speed ADC card with the lowest cost-per-channel in the industry. A few months later, the MI125-32, an HPC FMC with a 32-channel ADC (125MSPS, 14 bits) was born.

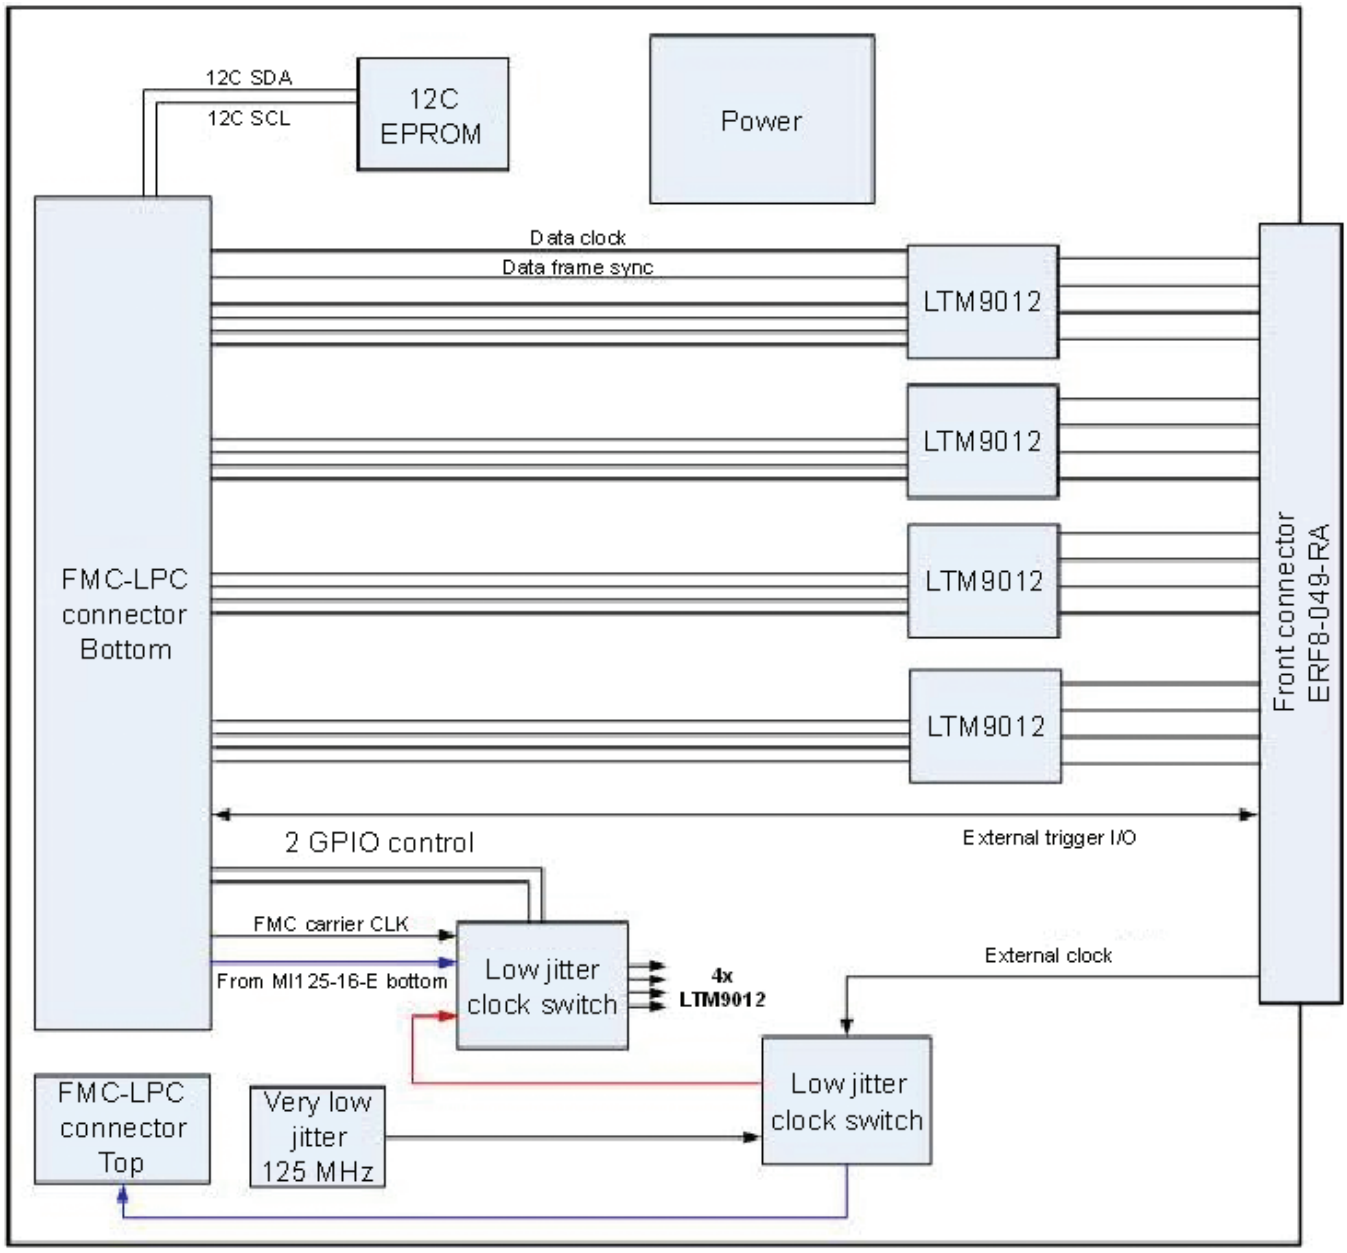

Figure 4: Nutaq's MI125-32 ADC FMC

With channel density in mind, engineers have leveraged Nutaq's innovative "double-stack" concept. "Double-stack" takes advantage of the MicroTCA's full-size form factor, enabling the use of two FMCs on a single AMC carrier. The trick to implementing a double-stack FMC is to design a board that only requires the signals from a low-pin-count (LPC) connector to function. Instead of using an LPC connector to connect the board to a carrier, an HPC connector is used instead. Signals destined to both FMCs in the stack go through this connector to the extender FMC; signals for the extender FMC use the LPC signals, and signals for the top FMC use the extra signals available from the HPC. These extra signals from the HPC connector are then rerouted to an LPC connector on the extender FMC to be connected to the top board.

Figure 5 shows how signals are routed in double-stacked FMCs.

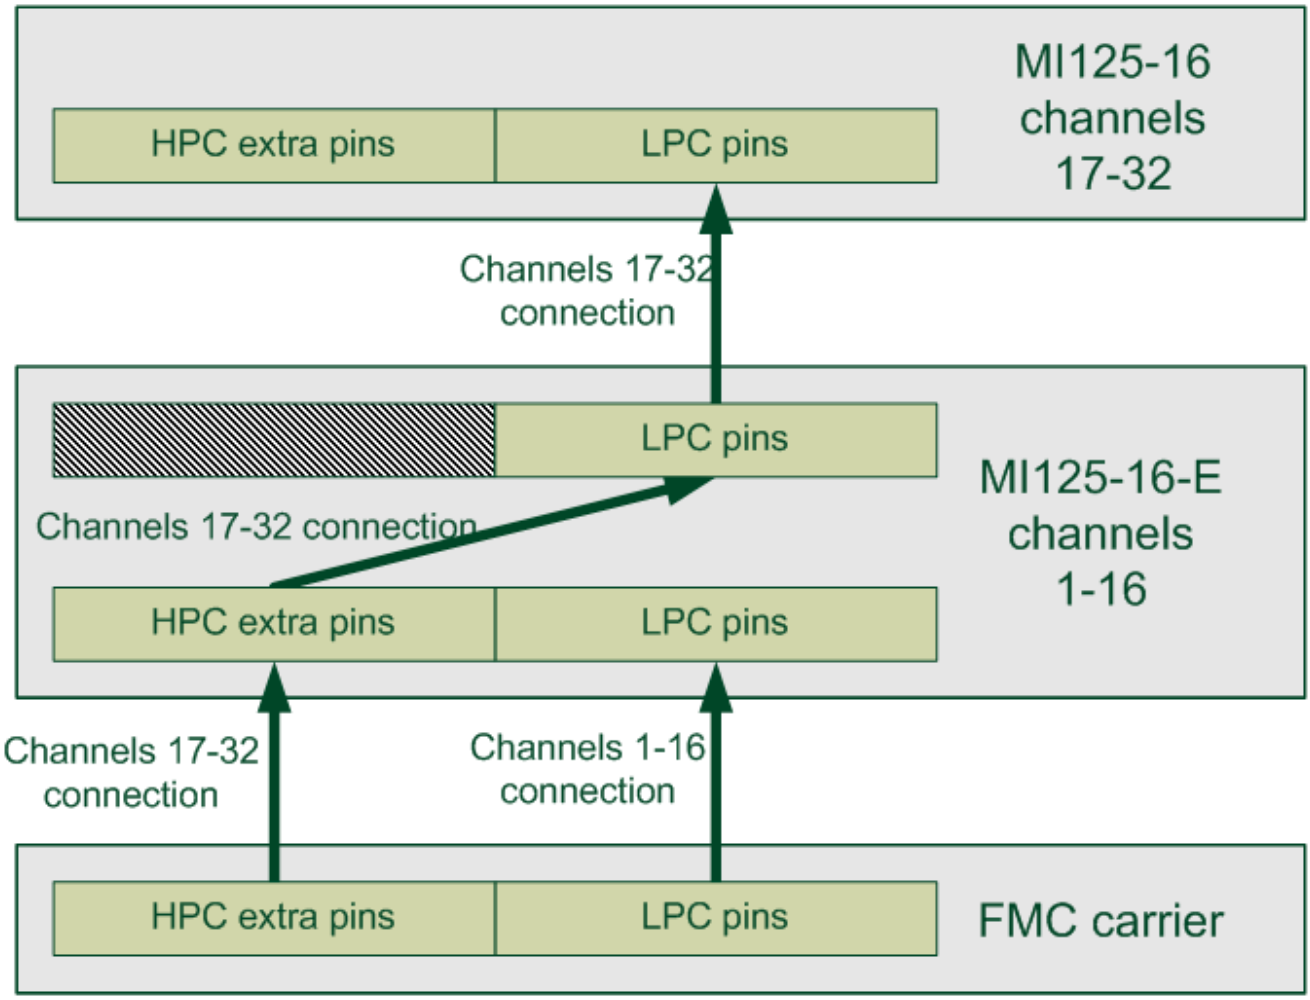

Figure 5: Double-stacked FMCs

The following sections discuss how the MI125-32 FMC can be integrated in an FPGA-based MicroTCA data acquisition system and used in a linac.

3 Implementing FPGA-based low-level RF control loops

"Radio frequency (RF) cavities are metallic chambers spaced at intervals along linear accelerators (Linac) and are shaped to resonate at specific frequencies." [1] When particles pass through the electric field in a cavity, some of the energy from the radio waves is transferred to the particles. To maintain the RF field in the cavities, both amplitude and phase control loops are required.

Figure 6 shows the key components in an LLRF control system:

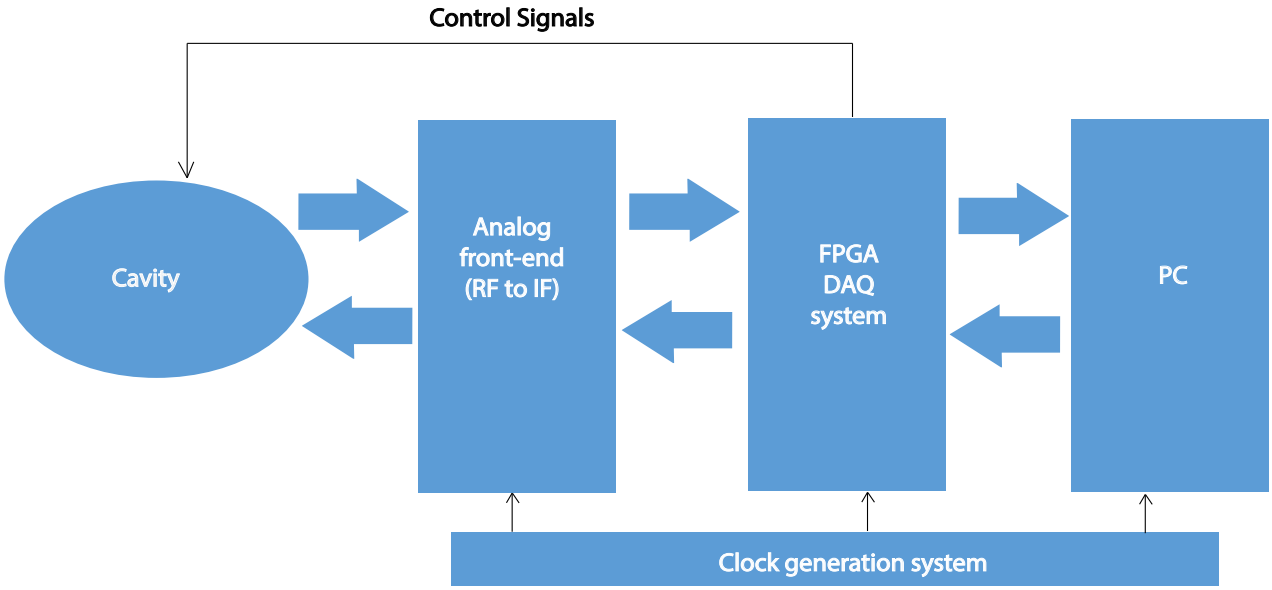

Figure 6: Key components of an LLRF control system

As we can see, an LLRF system is composed of five major elements:

- The cavity in which the particle is accelerated
- An analog front-end to perform the up/down conversion of the RF signal into an intermediate frequency (IF) signal for interfacing with the FPGA data acquisition system.
- The FPGA data acquisition system. In many ways, the 'brain' of the entire LLRF system, the FPGA data acquisition system is a digital board with all the ADCs and DACs required for signal conversion between the analog and digital world, along with an FPGA processor on which the control algorithms are implemented.
- A host PC, for communication with the other linac systems
- A clock generation system, to provide the clocks required by the ADC/DAC converters and RF front-end modulators. All the clocks are derived from a single reference clock to ensure perfect synchronization across all the elements.

Researchers from the Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT) in Madrid used a Nutaq platform to implement their LLRF firmware. The tasks were broken down as follows [3]:

- Under sampling + IQ digital demodulation
- Amplitude and phase loops
- Mechanical tuning loop
- Fast and slow diagnostics
- Fast interlocks
- Beam loading compensation
- Automatic conditioning

If you are new to LLRF but are familiar with system-level designs of wireless communication systems, you can see that the two are actually quite similar. These similarities include the operations performed inside the FPGA, which typically include digital filtering and baseband IQ processing.

In Nutaq's MicroTCA platforms, the FMCs inserted between the RF board and the MicroTCA digital board are used to digitize RF signals at sampling rates ranging anywhere from 1 MHz to thousands of MHz.

Figure 7 shows a typical Nutaq MicroTCA system based on the analog IQ approach.

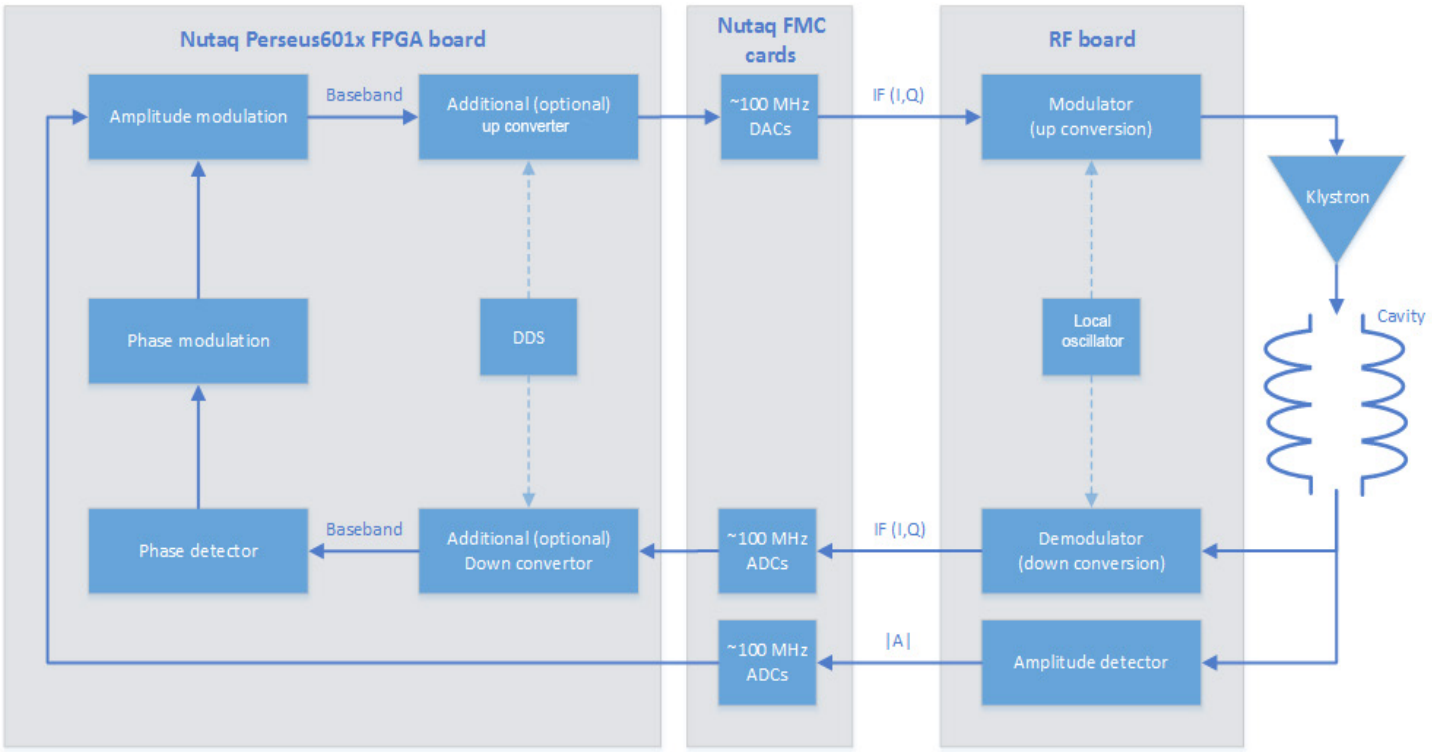

Figure 7: A Nutaq MicroTCA system

It is common to use FMCs with additional analog channels to provide interfaces to standard equipment like oscilloscopes for diagnostics and monitoring. Moreover, the analog channels can be used to receive readings from sensors. State machines, for example, are used to drive the interlocks accordingly to ensure safe operation of the linac. LLRF systems often include a computer to perform system management functions like configuring the digital boards and accessing and monitoring the main parameters and signals.

Nutaq's MicroTCA platform includes an SAMC-514 Intel Quad-Core i7 embedded CPU. The CPU implements the local control system, which includes a graphical user interface (GUI) for managing the LLRF digital board over Ethernet or PCI Express. The C/C++ EAPI and other tools from Nutaq's Board Support Development Kit (BDSK) helps to simplify the control system implementation by providing the main building blocks.

The cavity's RF parameters can vary from one accelerator to another and will most likely differ depending on its type and purpose. Let's use ALBA as an example. ALBA is a third-generation synchrotron light source with a 268 m circumference and is located near Barcelona, Spain. It has six 160 kW RF plants, each operating at 500 MHz with two inductive output tube (IOT) transmitters (klystrode) per RF cavity. The power of two 80 kW IOTs is fed to the RF cavity through a Cavity Combiner (CaCo). The total combined RF power for the six RF plants is 960 kW [5]. In ALBA, the main loops for controlling amplitude, phase, and tuning of the RF cavities are implemented on a Nutaq VHS-ADC CompactPCI digital board equipped with 16 ADCs, eight DACs and a Virtex-4 FPGA. The flexibility offered by this FPGA-based solution enabled the loops resolution and bandwidth parameters to be adjustable.

The following table shows the typical parameter ranges:

	Resolution	Bandwidth (kHz)	Dynamic Range
Amplitude Loop	< 0.1% rms	[0.1, 50]	30 dB
Phase Loop	< 0.1º rms	[0.1, 50]	360º
Tuning	< ±0.5º		< ±75º

In this system, a 520 MHz (500 + 20 MHz) analog RF front-end is used for down-conversion (RF to IF) and up-conversion (IF to RF). The 520 MHz for down-conversion is synchronized with an 80 MHz clock for digital acquisition.

Figure 8 shows the design:

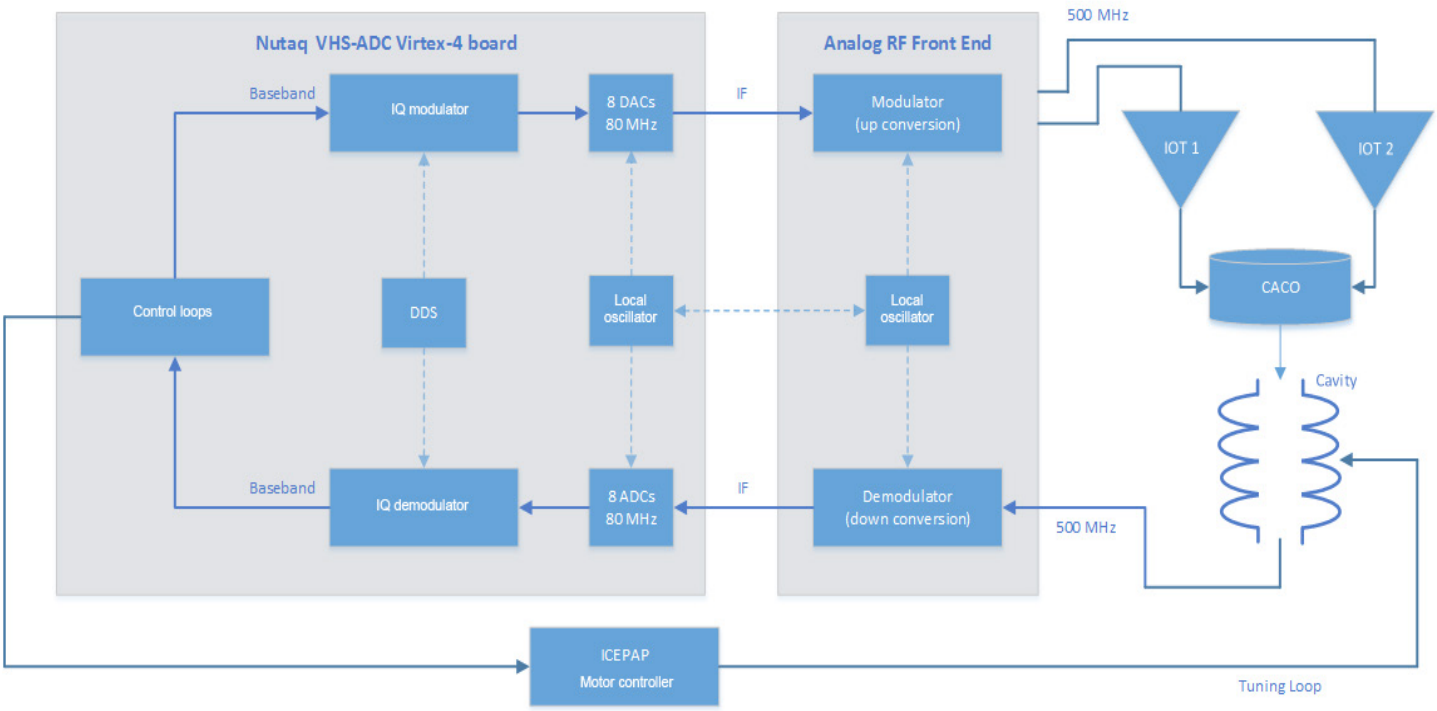

Figure 8: LLRF conceptual design of ALBA synchrotron light source

Once the signals are down-converted to baseband, phase modulation is used to switch from the IQ domain into the phase domain [17]. This is done with a Coordinate Rotation Digital Computer (CORDIC), an algorithm that converts the Cartesian I/Q vectors to polar (phase and amplitude) coordinates in order to simplify the phase controllers.

Research has been conducted to determine if the LLRF timing complexity could be significantly reduced by using the RF-baseband conversion method. In such an experimental setup, the conventional RF-IF front-end is replace by a direct conversion (RF-baseband conversion) front-end. This approach is inspired from a technology trend observed in the world of telecommunication where, more and more, two-mixer stage analog front-ends are slowly being replaced by direct conversion chips capable of taking an RF signal as an input and performing the full down-conversion to baseband in a single mixer. RF front-ends are becoming increasingly integrated – many chips on the market today even perform the analog-to-digital conversion, outputting digital I and Q based signals.

It's interesting to note that, not long ago, most LLRF systems were purely analog-based. The innovation drivers that came into play over the last few years in the FPGA data acquisition system business (arrival of faster ADCs, increased processing power of Xilinx FPGAs, and arrival of FMC and AMC form factors) certainly had an impact on the way LLRF systems are now designed, supporting a transition from the analog to the digital world. Similar to what has happened with the arrival of "Zero-IF" converters from the wireless industry, new innovations are likely to push the use of digital electronics even closer to the RF cavity. Ultimately, this will allow direct sampling of the RF signals and simpler system architectures, as shown in Figure 9.

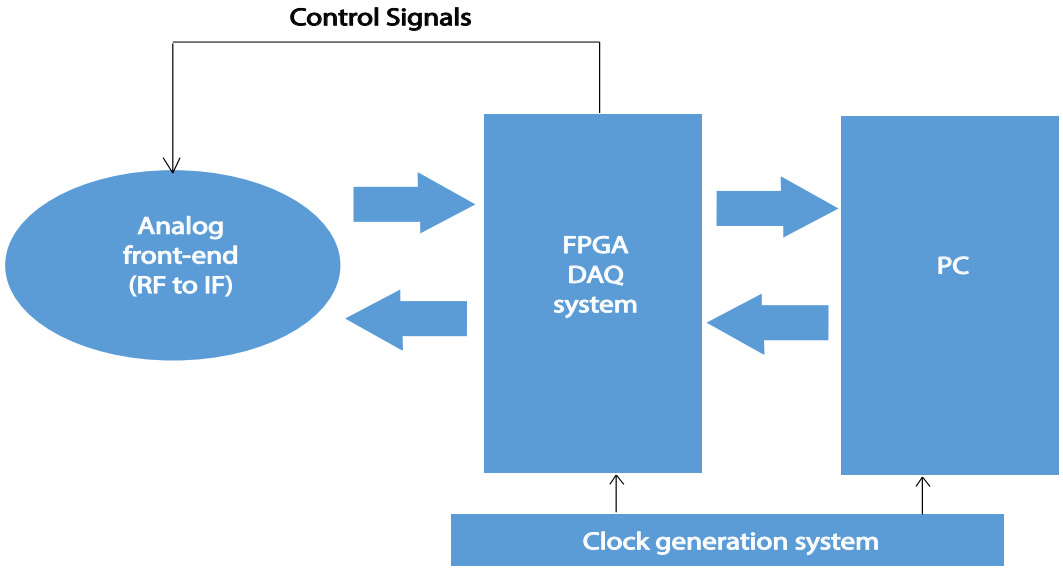

Figure 9: Direct RF signal sampling

In their paper, "Beam Position Monitor System of the ESS Linac", H. Hassanzadegan et al state:

> "Sampling in RF has the advantage of increasing the measurement bandwidth and simplifying the design of the analogue front-end. The main drawback, is that, when implemented on an RTM, some signal degradations may occur due to the bandwidth limitation and cross talk at the connection point of the RTM to the digital module. Also, with this method, jitter requirements for the ADC clock become more stringent due to the higher ADC speed. Sampling in IF, on the other hand, eases the ADC sampling, because of the lower IF signal frequency and less bandwidth requirements. The compromise, however, lies in the additional complexity of the analogue front-end due to the RF-IF conversion stage." [6]

Certainly, there are challenges with this approach. But LLRF systems, once purely analog, are now becoming more and more digital, up to the point where analog electronics could play a very minor role.

It's not clear when the transition to a fully digital LLRF system implementation will take place. For now, analog technology clearly has its place in reducing risk when deploying control systems for multi-million dollar linacs. Innovations such as faster ADCs/DACs and more powerful FPGA devices are likely to drive FPGA data acquisition systems and manufacturers like Nutaq will help in this transition. Software tools will also be required to support engineers in charge of LLRF digital signal processing, as their responsibility in the LLRF control loop grows over time.

There are two common approaches for baseband digital processing in LLRF systems. In the first approach, the cavity probe voltage is first down-converted to an intermediate frequency (IF) and then sampled. In the second approach, an analog IQ demodulator is used to directly convert the probe voltage into in-phase and quadrature components. The main advantage of the second approach is that the need for a precise synchronization of separate I and Q ADC sampling is eliminated. This leads to a more versatile design valid for a wide variety of RF frequencies [7].

A transition towards a fully digital LLRF means that the FPGA data acquisition system is getting something of a “promotion” – it requires an ADC/DAC solution that samples even faster and FPGAs that can perform signal processing tasks once handled by the analog circuitry. Traditionally, these components would be programmed in the VHDL programming language. Now the LLRF electronics engineer may be forced to move towards higher level block-diagram programming tools like Xilinx® System Generator for DSP™ in order to cope with the additional complexity. The CORDIC IP by Xilinx can be used to target the FPGA of Nutaq’s Perseus 601x and Perseus 611x AMC carriers [4]. Alternative custom implementations can also be realized using Nutaq’s Model-Based Design Kit (MBDK) or Board Support Development Kit (BSDK).

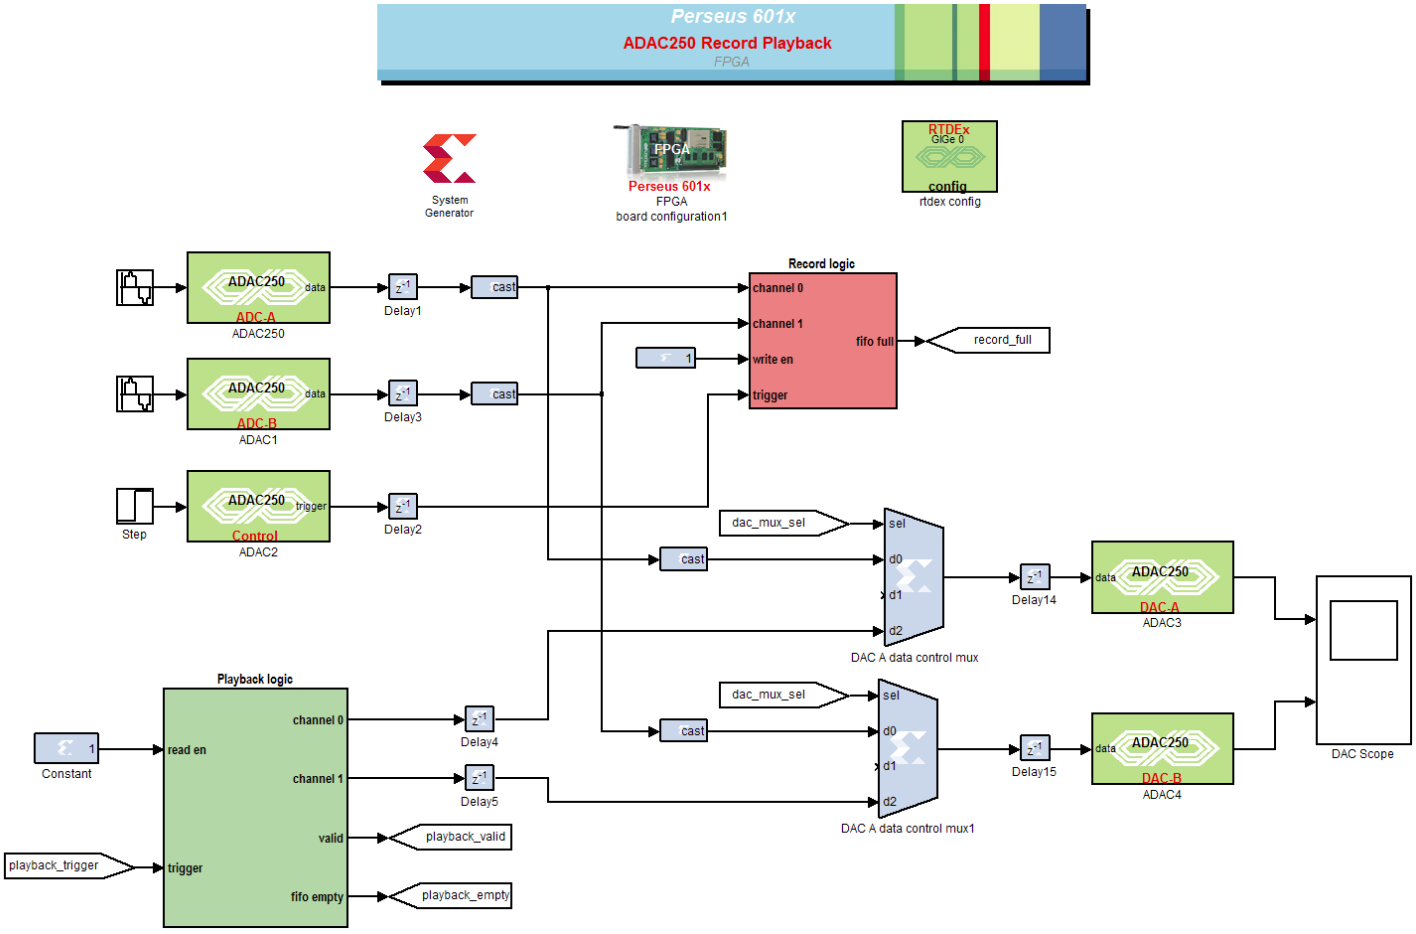

Figure 10: Nutaq’s Model-Based Design Kit (MBDK) and Xilinx® System Generator for DSP™

The improved flexibility offered by digital control systems enables the creation of more complex control loops [2]. For example, the FPGA board can be programmed with different algorithms for performing comparisons. Digital solutions also enable the use of redundant diagnostic channels for displaying the parameters of the control loop. Nutaq's real-time data exchange (RTDEx) interface supports a real-time readout of digital signals from within the FPGA. Reading and writing from custom registers is also easy, which is useful for control loop optimization, tuning, and diagnostics. Researchers involved with the ESS Bilbao Ion Source used a similar experimental setup for the study of a pulsed digital LLRF system (amplitude, phase, and tuning loops) for the Radio Frequency Quadrupole (RFQ) for the ion source test stand in Zamudio, Spain. The same Virtex-4 FPGA unit for signal processing and PI regulation was used and programmed with Nutaq's MBDK.

> "With model-based approach, the implementation of the system was eased and the number of bugs in the FPGA algorithms was reduced compared to the VHDL method. The LLRF hardware was co-simulated in MATLAB-Simulink using an aluminium mock-up cavity operated at low power. The results of these tests verified the ability of the LLRF system to meet all the LLRF requirements in addition to providing a fast response and a large phase margin for loop stability." [8]

Users of FPGA data acquisition systems can benefit from the years of development done by software tool manufacturers. Most of this development was done to address the movement toward a fully digital architecture driven by the wireless communications industry and its big push toward software-defined radio (for which high-speed ADCs/DACs and FPGAs are key components).

On top of industry-standard board-level libraries and APIs, the Nutaq BSDK includes a complete and efficient transport layer that enables both remote control and high-speed data exchange between Nutaq's FPGA-based hardware and standard processor blades or remote computers as well as a MicroBlaze soft processor, instantiated within the FPGA fabric and running a commercial, embedded Linux distribution. A RISC processor, MicroBlaze runs Nutaq's Central Communication Engine (CCE) as well as user-defined tasks and applications.

The CCE is an implementation of a Remote Procedure Call (RPC) library. It exposes the function of the Nutaq software libraries to the network. The CCE receives and handles TCP (Gigabit Ethernet) or PCIe commands, answers to requests in real-time, and supports multi-user connections. The External Application Programming Interface (EAPI) interacts with the CCE and enables host-based applications to remotely control Nutaq FPGA boards through an Ethernet connection or PCIe interface. The EAPI can be run under a Windows or a Linux operating system.

The command-line interface (CLI) is the basic client interface for Nutaq FPGA boards. It provides a shell where users can type commands and interact with the board's FPGA. The CLI offers many useful features, like programming an FPGA bitstream in the on-board flash memory, reading or writing specific addresses on the AXI bus, loading data to a specific address in DDR3 SDRAM, etc.

The main objective of Nutaq's RTDEx IP core is to provide developers with a framework to exchange data with a host computer device through GigE or PCIe links with the highest bandwidth (greater than 750 MBPS sustained data rate) and lowest latency possible (less than 1 ms average round trip latency).

Nutaq's FPGA SDRAM recording module enables you to store bursts of data in the onboard SDRAM. These can then be transferred to a host device for storage and analysis. The FPGA SDRAM playback module enables transmission of large portions of data from a host device such as a computer to the FPGA SDRAM, which can then be read by the FPGA at a very high speed (5.7 GBPS sustained data rate).

Figure 11 shows a graphical representation of these functional blocks along with their respective implementation location (some on the host PC, some in the FPGA):

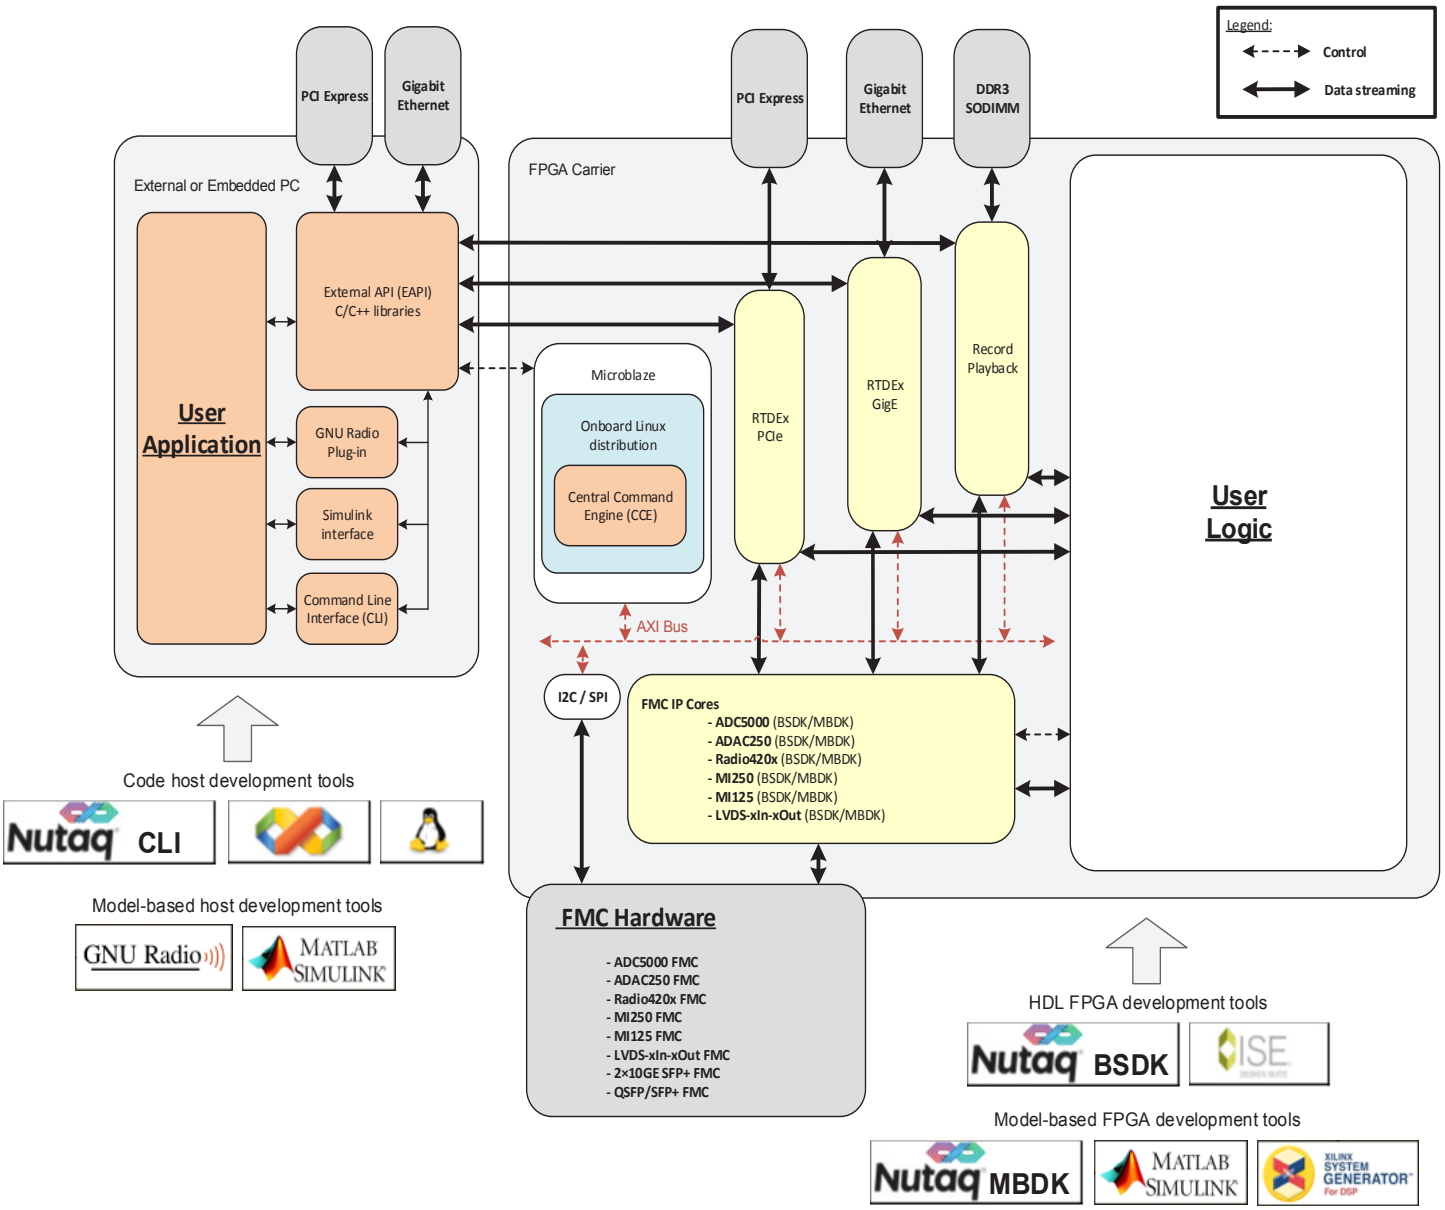

Figure 11: Nutaq solutions for accelerated development

4 Global orbit feedback

In linacs, the trajectory of the beam, which is comprised of accelerated particles, is stabilised and bent to follow a circular path or ring. This meticulously controlled path is referred to as the “orbit”.

The reference orbit to which the beam is to be steered may be either a “design orbit” or a “golden orbit”. A design orbit is an ideal orbit, flat except for where intentional bumps are applied. Bumps are used in colliders to separate the beams at the non-colliding interaction regions. In electron accelerators, the bumps are used for the generation of x-rays (as a result of bending the beam). A “golden orbit” is a previously acquired orbit recreated for experimentation purposes.

A linac’s global feedback loop is comprised of many components. First, a beam position monitoring (BPM) system, consisting of several hundreds of strategic locations along the ring, measures the position of the beam in the flat plane. Secondly, a global BPM data distribution network sends the measurements to a central orbit controller where a digital controller computes the corrections to be applied to the beam. These corrections are sent back to the orbit correction magnets (OCM) spread along the ring. Finally, corrector magnets are electronically driven to steer, stabilize, bend, and focus the beam.

The global feedback system uses as input position measurements from over one hundred BPMs, per plane (X-Y), per ring. The use of multiple BPMs and correctors enables a centralized algorithm to determine the applied corrections. The overall performance of the feedback system depends on the type of the digital controller and the rate of the loop. Target rates can vary from as low as 1 Hz (slow orbit feedback) up to a few kHz (fast orbit feedback). Figure 12 shows a typical global orbit feedback system.

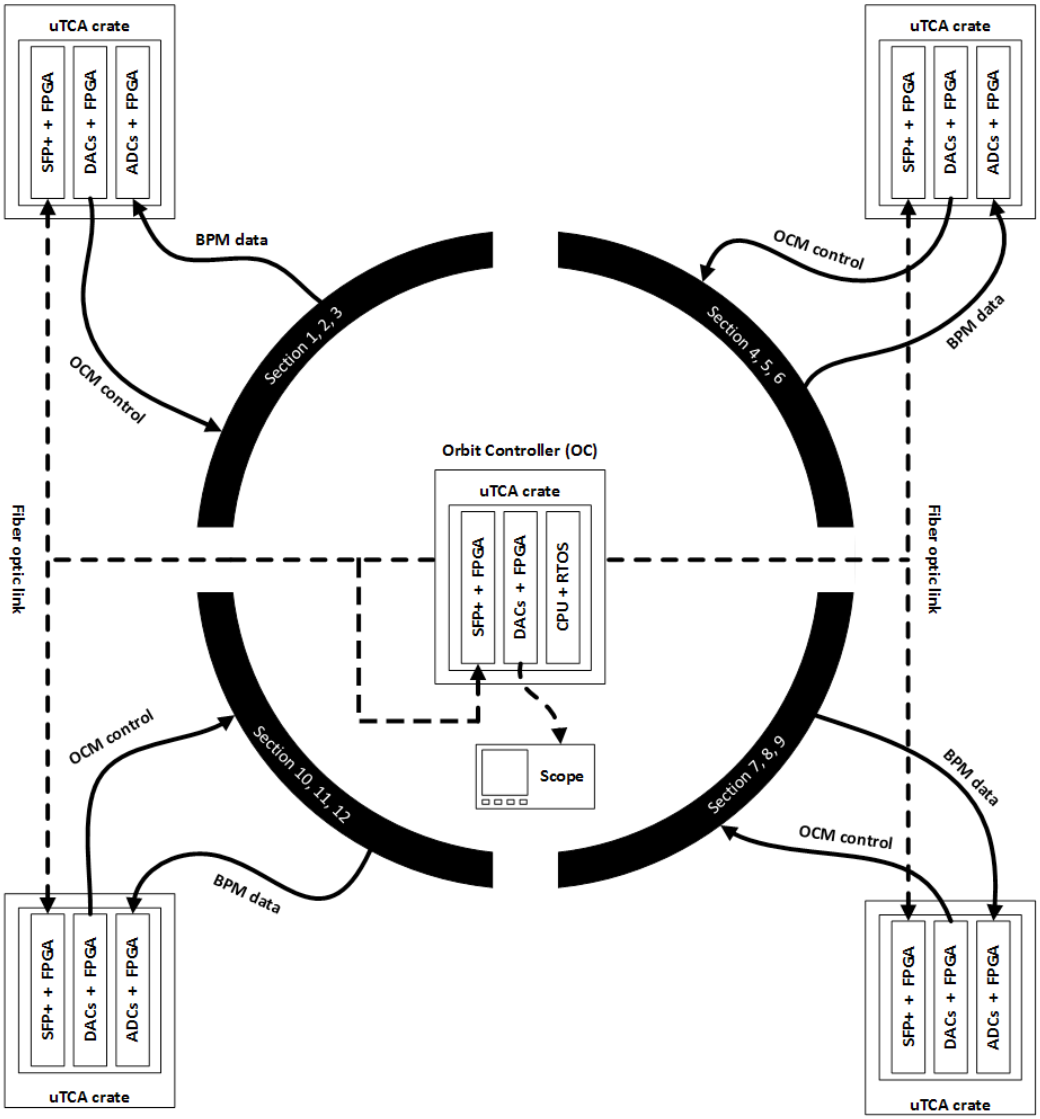

Figure 12: Global orbit feedback system

5 Beam position monitor (BPM)

BPM modules are electronic modules that perform fast analog processing of beam pickup signals. They consist of wideband analog modules [9] and are capable of non-interceptive beam position measurements.

BPM modules are found in linacs, microtrons, and transfer lines. They can measure the position of three types of beams: single bunch, macro-pulse, and continuous wave (CW). Note the following characteristics with BPM modules:

- Can work in the S-band, L-band, and X-band. L-band and S-band beams can be processed, provided the bunch groups are short (<3 ns).
- Bunches at any repetition rate up to 500 MHz can be measured.
- A 5 MHz repetition of individual bunches can be distinguished from one another [10].
- Macro-pulses or single bunches up to few MHz repetition rate are measured individually.
- CW beams can be measured continuously.
- Beam position motions up to 5 MHz can be observed.

The front-ends' X and Y outputs are typical analog 2V signals. The precision of each position reading is approximately 10 µm. The analog X and Y coordinates from the BPM modules are continuously digitized with high speed analog-digital converters (ADCs). The digitized signal is referred to as the "beam position monitor data" (or simply "BPM data").

Nutaq's MI125 FPGA mezzanine card (FMC) provides up to 32 ADCs and can be used to digitize the X-Y data from up to 16 BPMs. Multiple MI125 cards can be combined in a MicroTCA chassis. BPM data acquisition on the orbit controller is either interrupt or timer-driven, and data on all the channels is acquired simultaneously. Nutaq's RTDEx and Mestor interfaces, combined with the MI125's external clock input capability, provide a host-driven synchronous external triggering system for synchronised acquisitions across all nodes.

6 Orbit controller (OC)

The orbit controller serves as the system's control center. It has the following responsibilities:

- Orbit display and BPM/corrector disabling functionality
- BPM data acquisition
- Response matrix measurement and building the constraint vectors
- Performing inversion and solving the system to calculate the orbit corrector values
- Calculating the correction using Singular Value Decomposition (SVD) calculation [11]
- Running the orbit PDI control algorithm [12]
- Computing and distributing the set points to the Horizontal and Vertical Corrector Magnets (HCM/VCM), to over one hundred corrector dipoles per plane per ring.

Nutaq's MicroTCA platform includes an SAMC-514 Intel Quad-Core i7 embedded CPU to perform system's management functionalities such as orbit display, enabling/disabling BPM/corrector, and logging data to a SATA hard disc.

Nutaq's MicroTCA platform uses the Nutaq Perseus 601X Virtex-6 FPGA carrier card for performing the real-time computations part of the control algorithm. Rahmati et al. (FPGA Based Singular Value Decomposition for Image Processing Applications, IEEE, 2008) shows that an FPGA implementation of the Jacobi-SVD algorithm is possible using a reasonable amount of FPGA logic resources. They observed a 3:1 computation time reduction on 20x20 and 30x30 SVDs when compared to optimal general purpose processor (GPP) implementations.

7 Orbit correction magnets (OCM)

Three types of configurations in linacs use magnets: corrector, aligner, and sweeper. To counter the tendency of particles to travel in a straight line, dipole magnets are used to bend the path of the particle beam into a ring shape There can be several thousand dipoles along a multi-kilometer ring. Only the first magnet configuration (corrector) is part of a fast orbit feedback loop and will be discussed here.

Electromagnets are commonly used to generate the controllable magnetic fields. Possessing a current of several thousand amperes, they produce powerful magnetic fields which the dipole magnets use to enable the beam to handle tighter turns. The more energy a particle has, the greater the magnetic field needed to bend its path. High-speed digital to analog converters (DACs) are used to control the correction magnetics' power supply. The electromagnets use a superconducting coil that enables a high current to flow without losing any energy to electrical resistance.

The first fully digitally controlled magnet power supplies were commissioned at Paul Scherrer Institute (PSI) in 1999 [14]. Today, most correction magnets use a self-optimizing power supply control system to enable more complex control techniques at higher sampling rates.

Other type of magnets are also used:

- Insertion magnets (Quadrupole magnets) – Acting like lenses to focus a beam, they gather the particles closer together. Three quadrupoles are used to create a system called an inner triplet. Inner triplets tighten the beam, making it narrow by 10 times or more, depending on the type of accelerator, down to few micrometres across.

- Lattice magnets – Thousands of "lattice magnets" bend and tighten the particles' trajectory. They keep the beam stable and precisely aligned. Beam motion is reduced to stay within a fraction of a micrometer in both the horizontal and the vertical plane.

8 Loop requirements

Feedback rates can vary from as low as 1 Hz (slow orbit feedback) up to a few kHz (fast orbit feedback). Corrections applied through slow orbit feedback suffer from poor synchronization, which causes undesirable orbit perturbations. A gain in beam stability and reproducibility of the experiments are observed when the rate of the feedback loop is increased. With slow orbit feedback systems, the maximum corrector change has to be limited in each iteration. In the faster feedback systems, the corrector's set points are simultaneously applied around the ring and it is necessary to limit the changes on an iteration less often. The increasing sensitivities of the experiments create the need for faster orbit feedback system [15].

Roundtrip latency, from BPM to OCM, is also a key factor when designing or selecting hardware platforms to implement a complete orbit feedback system. Most orbit controllers are designed to operate with a latency of approximately 100 µs. Most recent studies on this matter suggest a hardware architecture that can operate under very low latency constraints, in the range of approximately 20 µs [16].

The following Nutaq components implement the low-latency framework:

- The MI125 takes no longer than 9 clock cycles (90 ns @ 100 MHz) to digitize the X-Y analog output of the BPM module.
- Aurora BSDK/MBDK cores provide ready-to-use implementations of the Xilinx Aurora communication protocol for the Perseus 601x AMC carrier board. The cores can serialize and transmit 128 bits of data from one FPGA to another (including user FIFOs) in only 45 cycles (450 ns @ 100 MHz).
- The second FPGA, equipped with Nutaq's QSFP(+)/SFP(+) transceiver modules, provides connectivity to 6 multi-gigabit transceivers (MGT) for rapid serial transmission over fiber optic links.

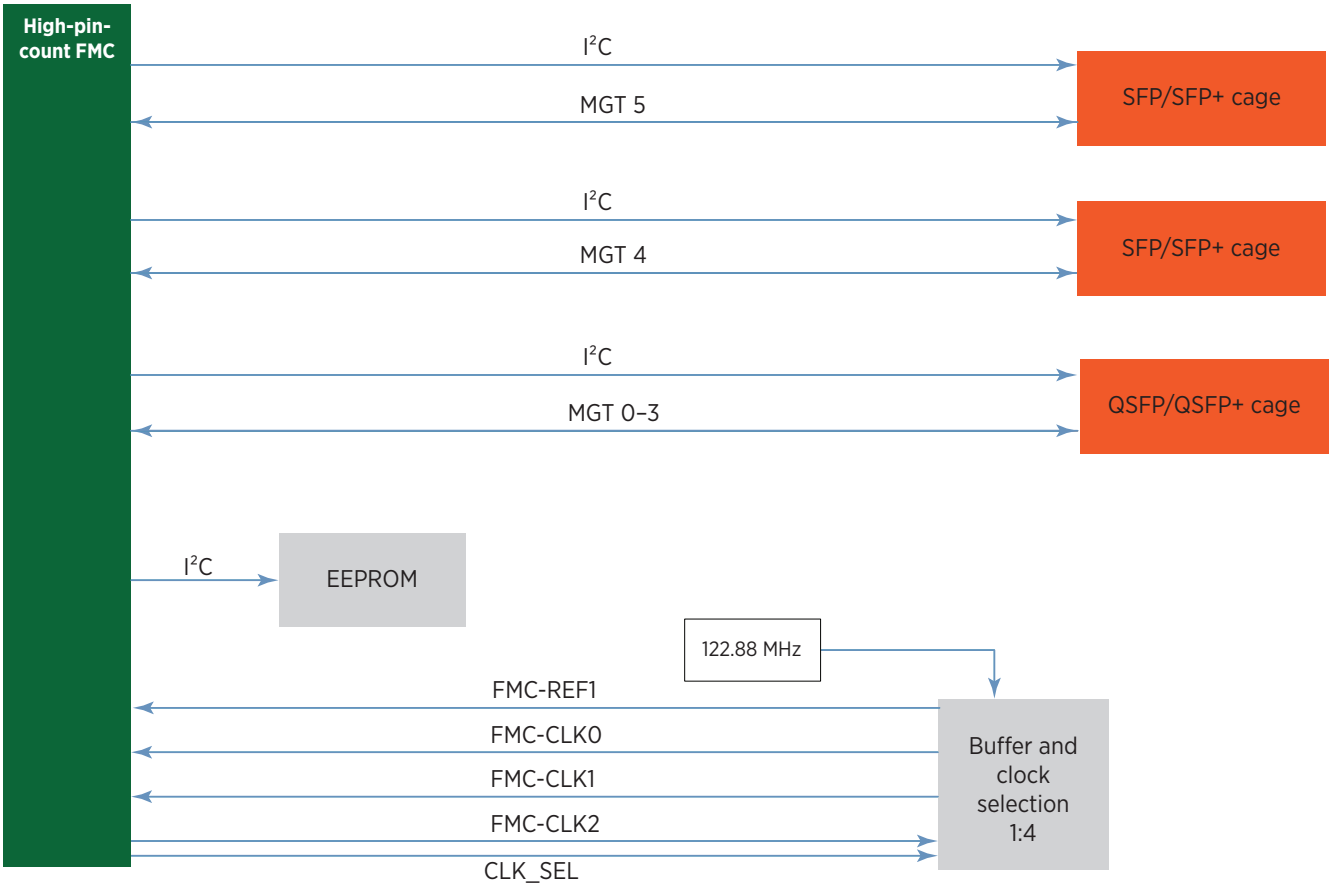

Figure 13: Functional block diagram of Nutaq's QSFP(+)/SFP(+) transceiver modules

Combined, these components are the perfect framework for implementing the most "low latency demanding" orbit digital control algorithms.

9 Conclusion

Working with commercial off-the-shelf (COTS) hardware has its benefits. A wide variety of highly mature and well-documented AMC carriers, FMC interface cards, and MicroTCA chassis can be found around the globe from well-known and established providers, who offer warranty of provision, technical support, and hardware warranties. For developers focused on signal processing in the area of high-energy physics, Nutaq's solutions accelerate the design, testing and deployment of innovative ideas. Our hardware platforms optimize programmability, processing power, flexibility, and cost, while our model-based design and open source software environment enable projects to be delivered with reduced development cycles and lower costs.

There are many advantages when working with COTS products, out-of-the-box drivers, and flexible software development environments; engineers can start prototyping their application on day one and benefit from dedicated technical support resources to ensure no down time. Engineers and researchers are left with more time to push the limits of science with their innovative ideas implemented on cutting-edge technologies.

10 References

[1] http://home.web.cern.ch/about/engineering/radiofrequency-cavities
[2] http://accelconf.web.cern.ch/Accelconf/SRF2009/papers/thppo048.pdf
[3] http://accelconf.web.cern.ch/accelconf/IPAC2011/papers/mopc160.pdf
[4] http://www.xilinx.com/products/intellectual-property/CORDIC.htm
[5] https://conferences.lbl.gov/getFile.py/access?contribId=18&sessionId=21&resId=1&materialId=slides&confId=27
[6] http://eval.esss.lu.se/DocDB/0002/000271/001/MOP240.pdf
[7] http://www.ehu.es/argitalpenak/images/stories/tesis/Ciencia_y_Tecnologia/Design,%20analysis%20and%20implementation%20of%20a%20versatile%20low%20level%20radio%20frequency%20system%20for%20accelerating%20cavities.pdf
[8] http://essbilbao.org:8080/ESSBilbao/en/publications/thp071.pdf
[9] http://www.bergoz.com/index.php?option=com_content&view=article&id=37&Itemid=481
[10] http://www.gmw.com/beam_diagnostics/Bergoz/bpm/lr_overview.html
[11] http://accelconf.web.cern.ch/accelconf/IPAC10/papers/mopec029.pdf
[12] http://accelconf.web.cern.ch/accelconf/p03/PAPERS/FPAB041.PDF
[13] http://accelconf.web.cern.ch/accelconf/IPAC2011/papers/thpo011.pdf
[14] http://accelconf.web.cern.ch/accelconf/IPAC2011/papers/thpo011.pdf
[15] http://iopscience.iop.org/1742-6596/425/4/042009/pdf/1742-6596_425_4_042009.pdf%5d
[16] http://www.slac.stanford.edu/econf/C0508141/proc/pres/ILCAW0908_TALK.PDF
[17] http://accelconf.web.cern.ch/Accelconf/SRF2009/papers/thppo048.pdf

Nutaq products are constantly being improved; therefore, Nutaq reserves the right to modify the information herein at any time and without notice.

Nutaq® INNOVATION TODAY FOR TOMORROW®

2150 Cyrille-Duquet, Quebec City (Quebec) G1N 2G3 CANADA
T. 418-914-7484 | 1-855-914-7484 | F. 418-914-9477
info@nutaq.com