

# A NEW REAL-TIME PROCESSING PLATFORM FOR ELETTRA 2.0 STORAGE RING

G. Gaio<sup>†</sup>, A. Bogani, M. Cautero, L. Pivetta, G. Scalamera, I. Trovarelli  
Elettra-Sincrotrone Trieste, Trieste, Italy  
L. Anastasio, University of L'Aquila, L'Aquila, Italy

## Abstract

Processing synchronous data is essential to implement efficient control schemes. A new framework based on Linux and DPDK will be used to acquire and process sensors and actuators at very high repetition rate for Elettra 2.0. As part of the ongoing project, the current fast orbit feedback subsystem is going to be re-implemented with this new technology. Moreover, the communication performance with the new power converters for the new storage ring is presented.

## INTRODUCTION

At Elettra, a novel hardware/software platform [1] has been embraced to facilitate real-time interfacing and processing for devices equipped with Ethernet interfaces. This architectural paradigm, based on DPDK [2], an industrial-grade package primarily implementing network stack bypass techniques, was deployed last year in the FERMI free electron laser [3] as part of a comprehensive control system upgrade. Furthermore, it has recently been employed in the upgrade of the legacy Global Orbit Feedback (GOF) [4] for the Elettra synchrotron.

While this upgrade may not be deemed strictly necessary, it serves as an ideal testbed for the evaluation of hardware and software technologies over the next two years. These technologies are essential for the real-time control of the equipment slated for installation in the forthcoming particle accelerator, scheduled to replace the Elettra synchrotron in 2026.

## ELETTRA 2.0

The most critical devices installed in the new Elettra 2.0 accelerator will feature a dual Ethernet connection to the control system: the first one dedicated to configuration and supervision, the second for real-time control.

To justify the additional cost of a real-time connection, the communication performance in terms of latency and jitter with these devices must be at least a couple of orders of magnitude better than a mildly loaded conventional control system (10 ms), e.g. greater than 10 kHz. Although systems managing high-speed data do not improve the absolute performance of the accelerator (excluding bunch by bunch systems), they could be profitably used to enhance overall machine stability, speed up optimization and augment diagnostic capabilities.

For Elettra 2.0, equipment featuring dual interfaces will include 1344 magnet power converters (PS) and 171 beam

position monitors (BPM). Furthermore, high-speed dedicated network links will be extended to low-level RF systems for the radiofrequency cavities, insertion devices, and photon diagnostics in the beamlines, even though the hardware for these systems has not yet been completely defined.

Due to the heterogeneous nature of connected devices and the opportunity to develop applications capable of harnessing highly synchronized data (LOCO [5], beam based alignment, post-mortem analysis, instability detection, electron-photon correlations, electron/photon beam optimizations, power converter prognostics and beam position monitor fault detection), the decision was made to centralize all fast interfaces at a unique point, a dedicated Intel based server. This server will be underpinned by a system based on DPDK and coded in C, ensuring the flexibility to implement advanced control schemes surpassing mere orbit stabilization.

## GOF UPGRADE FOR ELETTRA

### Hardware

Currently, the Elettra storage ring hosts 96 BPM detectors (Libera-Electron) that transmit data via Ethernet at 10 kHz to twelve Motorola MVME6100 CPUs, with each receiving data from eight BPMs using four 48-port Extreme X440-G2 switches. Each CPU, equipped with digital to analog cards (DAC), can superimpose a maximum current equal to one-fortieth of the total strength onto the corresponding seven horizontal/vertical correctors.

To implement a global correction, the entire orbit data is shared among the CPUs via reflective memory with each one handling one-twelfth of the feedback processing at Libera transmission rate.

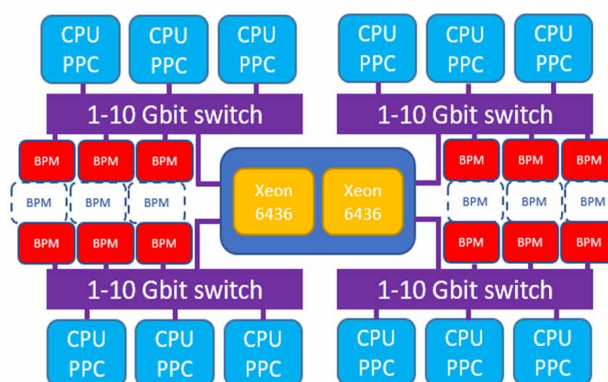


Figure 1: Network layout of the upgraded GOF.

<sup>†</sup> giulio.gaio@elettra.eu

A dual-socket DELL PowerEdge R750 rack-mount server equipped with Xeon 6436 Gold processors, 384 GB of memory and eight 10Gbit ports (Intel X710-DA4 model) is connected to the same switches. These switches can be reconfigured by means of a dedicated Tango device server to redirect BPM traffic from the MVME6100 boards to the first four network ports of the DELL server (see Fig. 1).

Furthermore, the legacy GOF software running on the VME CPU boards can be instructed, by the same Tango device server, to halt feedback calculation and process only a UDP packet, containing the DAC settings, that is transmitted by the DELL server at the feedback repetition frequency (see Fig. 2).

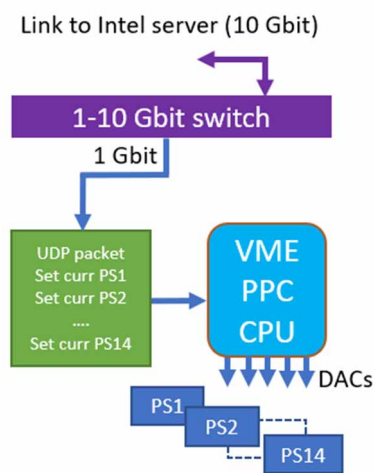


Figure 2: PS setting scheme of the upgraded GOF.

In addition to receiving data from Libera (0.96 Mpps), a prototype of the BPM detector for Elettra 2.0 [6], connected to two out of the loop BPMs (low gap BPM), sends packets to the same server at a turn-by-turn frequency (1.15 Mpps) for online orbit analysis.

### Software

At the operating system level, multiple strategies have been adopted to isolate real-time processes from the Linux scheduler and from any other process that could introduce delays in execution time.

All real-time processes have been assigned to the second socket of the server, which is completely isolated from the Linux scheduler. In addition to disabling power-saving, kernel un-maskable routines or uninterruptible housekeeping processes, such as Machine Check Exceptions (MCE) and NUMA balancing, which in earlier versions of the platform generated periodic jitter of 20-30  $\mu$ s every 5 minutes, have been halted.

The network cards interfacing with the equipment have been installed in PCIe slots to ensure SMP affinity with the second socket, minimizing DMA transfer times and reducing hardware-level jitter.

The operating system is Voltumna Linux [7], an in-house built distribution based on Yocto, compiled with optimizations specific to the processor architecture, running Linux kernel version 5.10 and DPDK 22.11. The

server operates in diskless mode, booting from the network, to eliminate potential jitter caused by performing I/O with hard disks.

The Linux kernel command line for this setup is as follows:

```
BOOT_IMAGE=(tftp,192.168.205.26)/ccd-sde-6346-2.x/ccd-sde-d-6346-06v45n-gph-2.4-0-1-g4ea9f39-dirty/boot/bzImage norandmaps mitigations=off systemd.gpt_auto=offconsole=ttyS0,115200n8 console=tty0 rootfstype=nfs ip=:::eno8303:dhcp::: mount.usr=/voltumna/ccd-sde-6346-2.x/ccd-sde-d-6346-06v45n-gph-2.4-0-1-g4ea9f39-dirty quiet pti=off isolcpus=1,3,5,7,9,11,13,15,17,19,21,23,25,27,29,31,33,35,37,39,41,43,45,47,49,51,53,55,57,59,61,63 nohz_full=1,3,5,7,9,11,13,15,17,19,21,23,25,27,29,31,33,35,37,39,41,43,45,47,49,51,53,55,57,59,61,63 rcu_nocbs=1,3,5,7,9,11,13,15,17,19,21,23,25,27,29,31,33,35,37,39,41,43,45,47,49,51,53,55,57,59,61,63 irqaffinity=0,2,4,6,8,10,12,14,16,18,20,22,24,26,28,30,32,34,36,38,40,42,44,46,48,50,52,54,56,58,60,62 intel_idle.max_cstate=0 processor.max_cstate=0 idle=poll nosoftlockup nmi_watchdog=0 nowatchdog default_hugepagesz=1G hugepagesz=1G hugepages=8 mce=off numa_balancing=disable
```

To accelerate software development, we partly reuse the code of “*l3fwd*” example provided within the DPDK development framework. Following initialization, in which each of the eight 10 Gbit ports is assigned to one core for network traffic control, a separate process responsible for managing feedback is initiated on a ninth core.

Below is the command line for executing this process:

```
./build/gof2 -l 3,5,7,9,11,13,15,17,19 -n 8 --force-max-simd-bitwidth=512 --p0xfff --config="(0,0,3),(1,0,5),(2,0,7),(3,0,9),(4,0,11),(5,0,13),(6,0,15),(7,0,17)" -P
```

The time required to acquire all the 96 packets in each feedback cycle typically falls within the 3 to 4  $\mu$ s range, with occasional outliers extending to 5  $\mu$ s. Jitter is limited to 1-2 microseconds.

DPDK optimize packet processing in multiples of the dimensions of the active vector instructions (SSE4, AVX, AVX512), thus optimal jitter performance is achieved when the quantity of received packets per port aligns with a multiple of 4. At the start of the real-time program, it is mandatory to prevent data streaming from interfering with port initialization; otherwise, the jitter between the first and last packets in a feedback cycle can grow from 5  $\mu$ s to 25  $\mu$ s. In order to stress the server and decrease processing latency, settings are calculated and transmitted to the VME CPU boards at a frequency of 20 kHz, notwithstanding that BPM data reception occurs at 10 kHz.

The feedback calculation time fluctuates between 2  $\mu$ s when exploiting AVX512 vector instructions and 50  $\mu$ s when no optimizations are applied. Presently, the correction algorithm mirrors that of the legacy system (SVD+PID), with the exception of the forthcoming implementation of notch filters centered at 50 Hz harmonics (see Fig. 3).

To ease configuration and data monitoring, a shared memory system manages the data transfer between the second socket, where the feedback process runs, and the first socket, where a dedicated Tango server oversees remote feedback control (see Fig. 4).

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2023). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

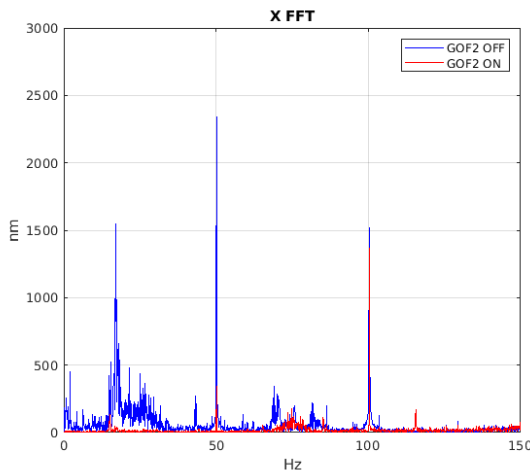


Figure 3: BPM noise with or without the upgraded GOF running at 20 kHz (corrector cut-off frequency 70 Hz).

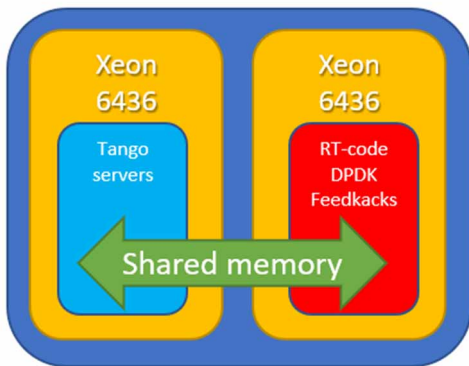


Figure 4: Communication scheme between the RT and non-RT domain.

Beyond enabling the adjustment of feedback parameters, the Tango device allows to retrieve BPM and corrector data from circular buffers operating at the feedback frequency, and computes the response matrix for both planes within 9 seconds.

### ELETTRA 2.0 POWER SUPPLIES

As previously mentioned, for Elettra 2.0, all power supplies will be equipped with high-speed links. While magnets, such as the bending magnets, do not need changing the setpoint in standard user operations, having fast feedback driving the current/voltage of the power supply is valuable for conducting fast diagnostic tests capable of detecting any deviation in power supply performance. Additionally, ensuring perfectly synchronized settings for an entire family of magnets (in Elettra 2.0 each magnet has a dedicated PS) can prove beneficial, considering unforeseen operational scenarios that were not initially accounted for in the design.

Elettra developed two prototypes of power converters, capable of handling currents of 20 A and 100 A, respectively [8]. These converters will be employed to drive multipole magnets and three types of correctors. The power supplies are equipped with a controller developed

by CaenELS [9], offering interfaces such as an Ethernet port connected to the ARM microprocessor and a dual gigabit SFP port directly connected to the FPGA for fast feedback point to point or daisy chained setting (see Fig. 5).

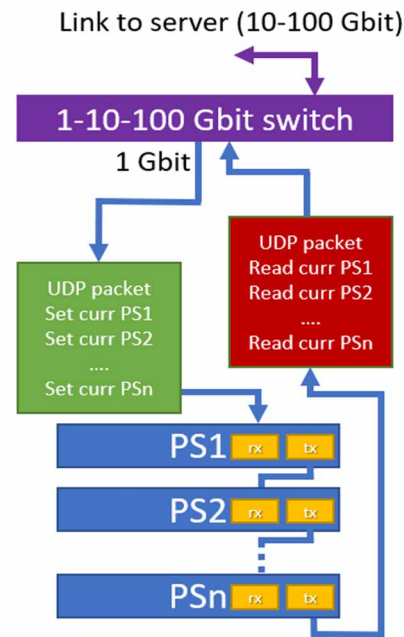


Figure 5: TX/RX scheme for the Elettra 2.0 power supply control.

To evaluate the maximum set/read speed of these power converter prototypes, a server configured in a manner similar to the one used for the GOF upgrade transmits a UDP packet containing the current setpoint at a rate of 100 kHz to the power supply's SFP interface. This frequency aligns with the speed of the internal regulation loop within the FPGA.

A network analyzer, the Anritsu MT1000 (see Fig. 6), in a pass-through configuration, captures the reply packet from the power supply, containing the current readout.

Based on the data recorded by the analyzer, the time measurement between packets amounts to 10  $\mu$ s, with a rms well below the microsecond (see Fig. 7), confirming the desired jitter performance in terms of packet processing.




Figure 6: 10 Gbit network analyzer.

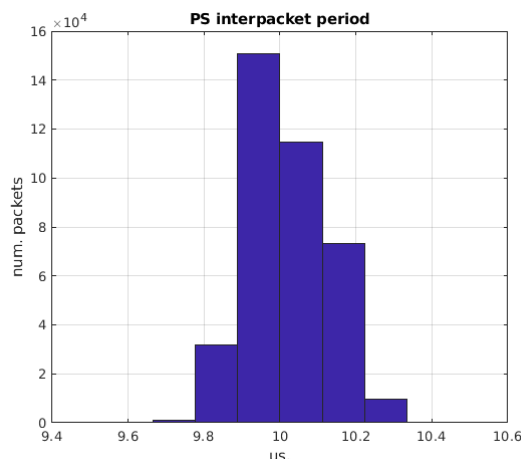


Figure 7: power supply tx/rx period distribution (100 kHz setting rate).

## CONCLUSION

Starting from the beginning of the next year, the GOF upgrade will become fully operational, facilitating the assessment of components for Elettra 2.0. These components include the processing system itself, the prototypes of BPMs and power supplies.

In terms of sheer performance, a single Skylake core of a Xeon processor, like the one employed in our tests, can effectively manage up to 9 million packets\* per second when coupled with an Intel X710-DA4 10Gbit port [10]. Using a Mellanox ConnectX-6 100Gbit adapter, processing capacity increases to 80 million packets per second [11].

Considering that each BPM box for Elettra 2.0 will acquire two BPMs, and that the power supplies will be organized in one daisy chain per cabinet (with 8 cabinets per section), a total of 84 real-time connections to the BPMs and 96 connections to the PSs will be established, resulting in a combined total of 30 million packets per second, at a repetition rate of 100 kHz.

Given that the final system acquisition is planned for the latter part of 2025, the forthcoming server will be based on the incoming Xeon Sapphire Rapids processor. This processor will encompass twice the number of cores compared to the current server, double the memory access bandwidth and support PCIe 5.0 bus.

While the learning curve for DPDK usage may be notably steep, the advantage of developing code in the C language, within a system guaranteeing minimal jitter levels, provides enhanced flexibility and minimizes development effort.

## REFERENCES

- [1] G. Gaio, G. Scalamera, “Development of Ethernet Based Real-Time Applications in Linux Using DPDK”, in *Proc. ICALEPCS’19*, New York, NY, USA, Oct. 2019. doi:10.18429/JACoW-ICALEPCS2019-MOPHA044
- [2] DPDK, <https://www.dpdk.org/>
- [3] E. Allaria *et al.*, “Highly coherent and stable pulses from the FERMI seeded free electron laser in the extreme ultraviolet”, *Nat. Photon.*, vol. 6, no. 10, Sep. 2012, pp. 699–704. doi:10.1038/nphoton.2012.233
- [4] A. Bulfone *et al.*, “Status of the Elettra Global Feedback project”, in *Proc. EPAC’06*, Edinburgh, Scotland, United Kingdom, Jul. 2006, paper THPCH091, pp. 3003-3005.
- [5] J. Safranek, “Experimental determination of storage ring optics using orbit response measurements”, *Nucl. Instr. and Meth. A*, vol. 388, no. 1–2, Mar. 1997, pp. 27-36, doi:10.1016/S0168-9002(97)00309-4
- [6] G. Brajnik *et al.* “Elettra 2.0 eBPM: complete system overview”, in *Proc. IPAC’23*, Venice, Italy, 7-12 May 2023. doi:10.18429/JACoW-IPAC2023-THPL077
- [7] L. Pivetta *et al.* “Voltumna Linux: a custom distribution for (embedded) systems”, presented at ICALEPCS’23, Cape Town, South Africa, Oct. 2023, paper THPDP026, this conference.
- [8] M. Cautero *et al.* “Elettra 2.0: magnet power converter strategy”, in *Proc. IPAC’23*, Venice, Italy, May 2023, pp. 3743-3745. doi:10.18429/JACoW-IPAC2023-WEPM076
- [9] CaenELS, <https://www.caenels.com/>
- [10] DPDK performance report 20.11, [http://fast.dpdk.org/doc/perf/DPDK\\_20\\_11\\_Intel\\_NIC\\_performance\\_report.pdf](http://fast.dpdk.org/doc/perf/DPDK_20_11_Intel_NIC_performance_report.pdf)
- [11] DPDK performance report 23.03, [http://fast.dpdk.org/doc/perf/DPDK\\_23\\_03\\_NVIDIA\\_NIC\\_performance\\_report.pdf](http://fast.dpdk.org/doc/perf/DPDK_23_03_NVIDIA_NIC_performance_report.pdf)

\* 64 byte packet