ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA033

TIMING, SYNCHRONIZATION AND SOFTWARE-GENERATED BEAM CONTROL AT FRIB*

E. Daykin†, M. Konrad, Facility for Rare Isotope Beams, East Lansing, Michigan

Abstract

The Facility for Rare Isotope Beams, once completed, will require hundreds of devices throughout the machine to operate using synchronized timestamps and triggering events. These include, but are not limited to fault timestamps, time-dependent diagnostic measurements and complex beam pulse patterns. To achieve this design goal, we utilize a timing network using off-the-shelf hardware from Micro Research Finland. A GPS time base is also utilized to provide client timestamping synchronization via NTP/PTP. We describe our methods for software-generated event and beam pulse patterns, performance of installed equipment against project requirements, integration with other systems and challenges encountered during development.

TIMING NETWORK TOPOLOGY

The FRIB timing network consists of two main segments: a stable, high-accuracy fiber event and time distribution network, as well as Precision Time Protocol (PTP) [1] and Network Time Protocol (NTP) [2] distribution over the facility LAN. A schematic of this configuration can be seen in Fig. 1. Devices are classified as high-, medium-, or low-accuracy depending on their sensitivity to time drift and jitter. These requirements are detailed below in Table 1:

Table 1: Timing Requirements

Class	Accuracy	Examples
High	$\pm 1\ \mu s$	Beam instrumentation
Medium	± 1 ms	Devices with timestamped interlocks
Low	± 1 s	Most other network devices

Fiber Distribution

Fiber-optic event signals are suitable for devices classified as high-accuracy and/or require event triggers. Events and timestamps are propagated over a three-level, tree-like network of inexpensive Micro Research Finland (MRF) [3] distribution fan-outs. The master fan-out transmits events to nine level 2 nodes, each in turn serving a handful of level 3 nodes within their respective 'region'. Level 3 nodes finally distribute these events to client devices in their immediate operating area.

Network Time Distribution

Network timing is suitable for devices classified as medium or low-accuracy. For more modern devices in both classes, PTP is preferred due to its superior accuracy, more robust master selection, and more graceful handling of leap seconds.

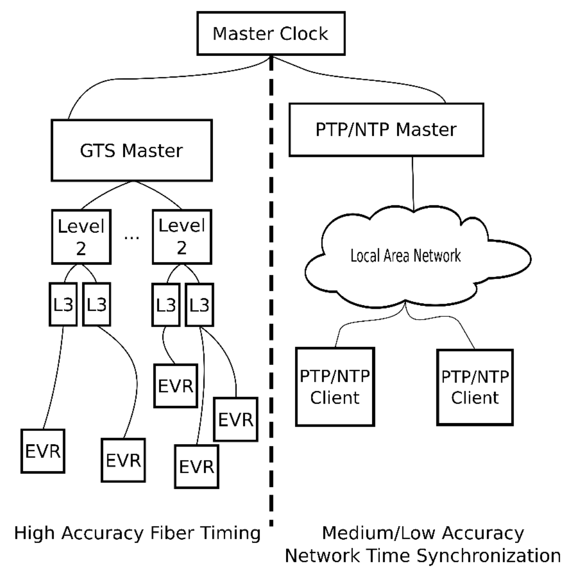

Figure 1: Simplified layout of FRIB's timing network.

MASTER DESIGN

The master time's stability is maintained via widely available timing hardware. This hardware generates precise 1 pulse per second (1 PPS) and 10 MHz signals, which are used for stable event distribution over MRF timing hardware, in addition to providing a phase-locked oscillator (PLO) reference signal for the facilty-wide RF clock. The fine design of this system is detailed in Fig. 2.

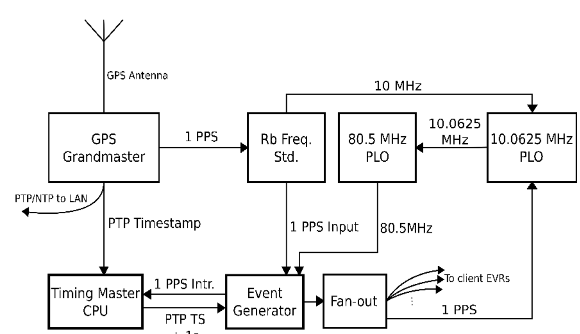

Figure 2: Synchronization and event generation machinery.

* Work supported by the U.S. Department of Energy Office of Science under Cooperative Agreement DE-SC0000661
† daykin@frib.msu.edu

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA033

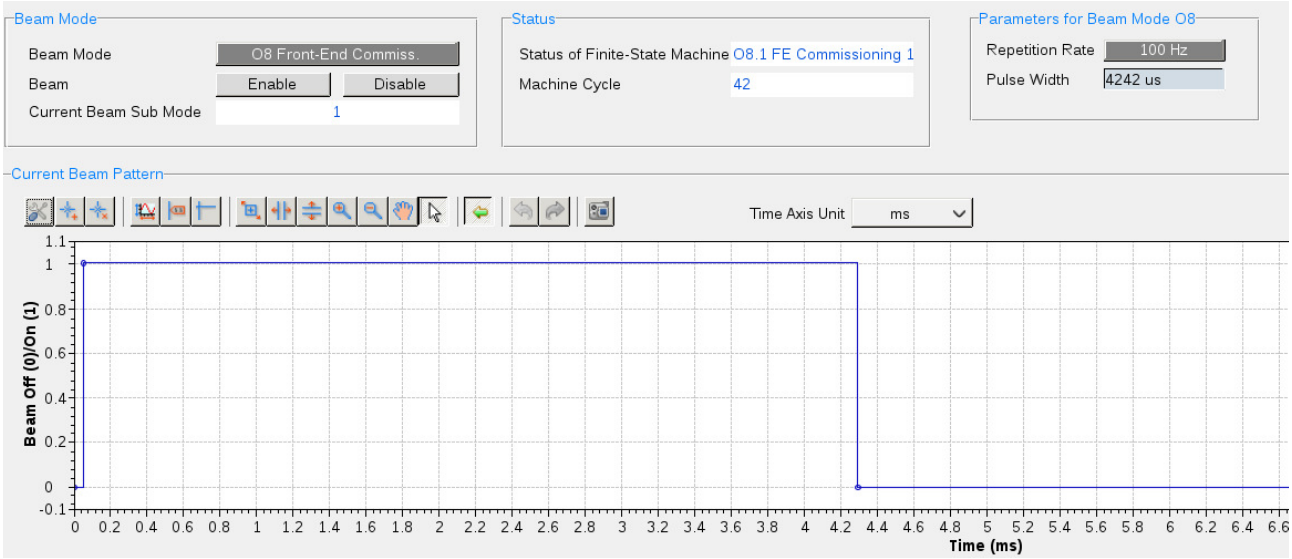

Figure 3: Generating a beam pattern with the scheduler interface.

Time Synchronization

An off-the-shelf grandmaster [4] receives a current timestamp from the Global Positioning System (GPS), which is forwarded to its embedded NTP and PTP daemons for LAN distribution, as well as a 1 PPS coaxial output to a rubidium frequency standard [5]. The rubidium standard's very low jitter is supplemented by the 1 PPS input from the grandmaster to maintain the long-term stability of the standard's 1 PPS and 10 MHz RF outputs. The 10 MHz output is used to discipline the facility-wide RF reference clocks, while the stable 1 PPS signal is forwarded to the fiber event generation equipment.

Timestamp and Event Transmission

Fiber Event Link Event generation is performed using the EPICS 'mrfioc2' driver [6] on a CPU running a real-time linux kernel and an MRF cPCI-EVG-300 event generator. Upon receipt of a 1 PPS pulse, the EVG raises a PCI interrupt to the CPU. The CPU, in turn, forwards the PTP timestamp back to the EVG. The EVG's internal PLL is locked to an 80.5 MHz signal from the facility-wide RF reference clock, which it uses as the event-carrier clock signal. Event triggers and timestamps are then distributed to client devices subscribing to them. Furthermore, the 1 PPS start-of-second event is used to synchronize the RF reference clock, closing the synchronization loop between it and the timing master.

Network Time For devices that do not support fiber timing, the GPS grandmaster also runs embedded implementations of PTP and NTP. The NTP application acts as a stratum 1 clock, while PTP acts as the best master (priority 0) within the FRIB LAN [1]. Whenever supported, PTP timestamps are attached at the physical layer, providing sub-μs accuracy, otherwise offsets can reliably be expected within ± 1 ms . NTP's ± 100 ms accuracy is generally acceptable in any other devices where PTP is unsupported.

SOFTWARE-GENERATED BEAM CONTROL

Once completed, FRIB will have the capability to deliver roughly 400 kW of beam power on target. To protect the target from thermal stress and accommodate for a wide variety of experiment use cases, FRIB's beam pattern is required to be dynamic and highly user-configurable. To this end, the *beam scheduler* is able to manipulate the beam chopper to generate single pulses of variable length, 'ramping' pulsed waveforms of linearly-increasing duty factor, and near-CW operation with 99.5% duty factor.

Machine Cycle Anatomy

Timing-sensitive beam instrumentation requires event data to arrive in 10 ms intervals, beginning with a 50 μs 'notch' (beam off period) and ending with a 'beam off' event, ensuring the beam will always turn off if events stop arriving (*e.g* due to a software crash). These machine cycles are defined by their duration in *event cycles*, thus, events must be transmitted at one of 805,000 discrete ticks[2]. The data is structured into what essentially amounts to a $2 \times n$ array denoting events, and the times at which those events occur, in a monotonic range between the 0^{th} and $804{,}999^{\text{th}}$ tick in the cycle. The real-time EPICS IOC broadcasts these cycles using two 50Hz software sequencers running in tandem [7], which alternate arming and loading a machine cycle worth of events into the EVG hardware queue.

User-Facing Beam Control

The capabilities of the beam scheduling machinery are easily configurable via a graphical interface (Fig. 3). For

[1] Highest-priority in FRIB's use case; lower priorities are also configurable.

[2] 80,500,000 ticks per second $\div$ 10 ms per machine cycle = 805,000 ticks per machine cycle

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-MOPHA033

operators, a simple screen allows pulse width and repetition rate to be modified within design parameters. For experts, a more full-featured menu is available for modifying parameters outside of the written beam mode specification, such as ramp settings, notch width and machine cycle length.

PERFORMANCE

Fiber Timing

For valid measurements, diagnostics equipment requires machine cycles to be exactly 10 ms long, with an event car-rier running at exactly 80.5 MHz. Figure 4 shows the readout from a client event receiver. Had the event generator been using its internal fractional synthesizer instead of external RF reference, the field labeled 'Tics Per Second' would be seen drifting by tens or hundreds of event cycles, which would result in noisy or unusable measurements.

Start-of-cycle Timing	SoC Min	10.000 msec
	SoC Max	10.000 msec
	Tics Per Second	80500000

Figure 4: GTS monitoring data on a client EVR.

Network Time

Despite the presence of PTP-unaware network hardware between devices, timing synchronization of less than 1 μs is possible if hardware timestamping is supported and used. To test a reasonably typical scenario, a PTP-unaware switch was placed between the GPS grandmaster and a test device supporting hardware timestamping. After a 1-minute sample of time offsets polling at 1-second intervals, a mean error of 2.4 ns was reported, with a standard deviation of 232.6 ns, outperforming the 1 ms requirement by a significant margin.

Beam Control

Experimental and machine-protection considerations require complex and varying beam parameters. With software event scheduling, tasks such as warming up charge strippers and targets are point-and-click operations, saving operating time and expense.

CHALLENGES

A few unexpected issues were encountered during the implementation of FRIB's timing system. With respect to network time synchronization, it is not guaranteed that any particular switch in the network is PTP-aware, Because of this, it is not possible to use peer-to-peer delay negotiation; however, this is perfectly acceptable if it is safe to assume reasonably symmetrical propagation delay. Additionally, synchronizing time across virtual LANs (VLANs) is inherently difficult by the nature of a VLAN; the problem is further compounded by the fact that PTP multicast packets are not routable.

With respect to the hardware itself, we found a small subset of fiber-optic cables to have connection issues which were resolved by re-certifying the quality of the cable. Additionally, the power supplies used in the Compact PCI chassis have a tendency to get quite hot and fail prematurely. This is mitigated by the fact that these power supplies are redundant and inexpensive.

The main unresolved issue with FRIB timing lies with the 100 Hz machine cycle rate. In practice, diagnostic event receivers are picking up 60 Hz line noise. The fifth harmonic of this noise, 300 Hz, is impossible to filter out of beam current monitor measurements. Dropping the machine cycle rate to 97 Hz is being explored.

CONCLUSION

FRIB's timing synchronization and distribution system is able to deliver timestamps and event triggers to a large number of devices at in a reliable and cost-effective manner. By employing a highly stable timing master in sync with the network time and RF reference, we can obtain accurate beam measurements and reliably reconstruct fault timestamps. Software-generated event scheduling provides high flexibility to user needs, as operators have the ability to construct long, complex beam patterns without firmware modification.

REFERENCES

[1] "IEEE standard for a precision clock synchronization protocol for networked measurement and control systems," *IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002)*, pp. 1–300, Jul. 2008. DOI: 10.1109/IEEESTD.2008.4579760.

[2] D. Mills, J. Martin, J. Burbank, and W. Kasch, "Network time protocol version 4: Protocol and algorithms specification," RFC Editor, RFC 5905, Jun. 2010, http://www.rfc-editor.org/rfc/rfc5905.txt

[3] *Micro Research Finland,* http://mrf.fi

[4] *EndRun Technologies Meridian II Precision TimeBase,* https://endruntechnologies.com/pdf/MeridianII-Time-Frequency-Standard.pdf

[5] *Stanford Research Systems FS725 Frequency Standard,* https://thinksrs.com/products/fs725.html

[6] M. Davidsaver *et al.*, *mrfioc2: EPICS driver for Micro Research Finland event timing system devices*, version 2.2.0, https://github.com/epics-modules/mrfioc2

[7] M. Konrad, "Timing and Synchronization at FRIB," in *Proc. of International Workshop on Personal Computers and Particle Accelerator Controls (PCaPAC'16), Campinas, Brazil, October 25 28, 2016*, (Campinas, Brazil), Geneva, Switzerland: JACoW, Sep. 2017, pp. 105–107, ISBN: 978-3-95450-189-2, http://jacow.org/pcapac2016/papers/thpoprpo10.pdf