TIMING AND SYNCHRONIZATION AT FRIB*

M. Konrad†
Facility for Rare Isotope Beams, Michigan State University, East Lansing, MI 48824, USA

Abstract

The Facility for Rare Isotope Beams (FRIB) requires a timing system for distributing a common time base for time-stamping data as well as for triggering actions of multiple devices distributed over the machine. Three different technologies are used to accomplish these goals. An event system based on commercial off-the-shelf hardware made by Micro-Research Finland provides time-stamps and triggers to more than 500 fast data-acquisition devices like beam diagnostics electronics, LLRF controllers, and machine protection nodes. This system is also used to distribute FRIB's complex beam pulse patterns with event rates of more than 50 000 events per second. For many hundred devices which require a lower timing accuracy like programmable logic controllers and computers the Precision Time Protocol (PTP) is used. Additionally the Network Time Protocol is used for legacy devices that do not support PTP, yet. We describe the architecture of the FRIB timing system and how the different timing subsystems are synchronized. We also describe how FRIB's beam pulse patterns are generated.

INTRODUCTION

FRIB [1] is a project under cooperative agreement between US Department of Energy and Michigan State University (MSU). It is under construction on the campus of MSU and will be a new national user facility for nuclear physics. Its driver accelerator is designed to accelerate all stable ions to energies >200 MeV/u with beam power on the target up to 400 kW [2]. The FRIB driver linac requires a timing system for distributing a common time base for time-stamping data as well as for triggering actions of multiple devices distributed over the facility. Commissioning of the front-end as well as the core components of the timing system is currently underway. The remaining parts of the accelerator and the timing system are planned to be commissioned over the next two years.

ARCHITECTURE

The client devices connected to the FRIB timing system have very different timing requirements. They can be divided into the three accuracy classes listed in Table 1. Due to the large number of timing clients in each of the classes the use of different technologies helps to reduce cost. Devices in the "high" accuracy class are connected to an event system that broadcasts its events over a fiber network. This system is based on commercial off-the-shelf products by Micro-Research Finland [3]. For the "medium" accuracy class

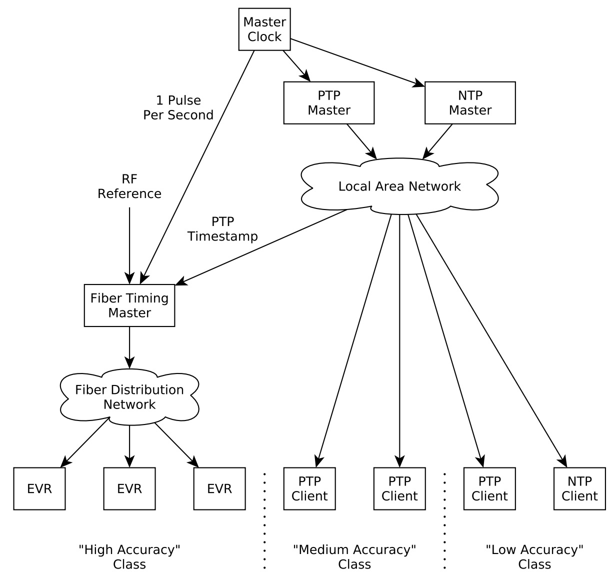

Figure 1: The FRIB timing system relies on three different technologies: A fiber-based event system as well as the Ethernet-based Precision Time Protocol (PTP) and Network Time Protocol (NTP). The three subsystems are synchronized to a common master clock.

the Precision Time Protocol (PTP version 2 as defined by the IEEE 1588-2008 standard) is used. In the last years this protocol has gained a lot of popularity for industrial automation applications and is widely supported by state-of-the-art programmable logic controllers. Additionally the Network Time Protocol (NTP) is supported. Compared to NTP the more modern PTP offers features like more efficient multicast messaging, cleaner handling of leap seconds as well as fault tolerance. This is why PTP is preferred over NTP even for clients in the "low" accuracy class. NTP is used only for devices that do not support PTP.

The event system, the PTP timing system and the NTP timing system can be considered three independent subsystems of FRIB's timing system. All three subsystems are synchronized to a common master clock to ensure data which has been time-tagged by devices connected to different timing subsystems, can be correlated (see Fig. 1).

EVENT SYSTEM

The event system is the most complex and most critical part of the FRIB timing system. In the following we will give a more detailed description of this subsystem.

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

ISBN 978-3-95450-189-2

Table 1: Accuracy Requirements of Timing Client Devices at FRIB. Devices that have low accuracy requirements but require triggers are considered part of accuracy class "high".

Accuracy Class	Typical Client Devices	Triggers Required	Estimated Number of Clients
"High" (≤ 100 ns)	Fast beam-diagnostics equipment, LLRF controllers	Yes	≈ 500
"Medium" (≤ 1 ms)	Programmable logic controllers, motor controllers	No	≈ 100
"Low" (≤ 1 s)	Computers	No	Hundreds

Table 2: Time Structure of the Beam during the Power Ramp up to 400 kW. The arrows correspond to a linear ramp.

Time (s)	Pulse Rate (kHz)	Pulse Length (μs)
0–30	2 ⟶ 25	0.6
30–60	25	0.6 ⟶ 5
60–100	25	5 ⟶ 20
100–500	25	20 ⟶ 39.4

Synchronization with RF Systems

The RF master oscillator is synchronized to the timing system (see Fig. 2). An Event Receiver (EVR) card connected to the event system generates a one pulse-per-second signal which is used to discipline a low-noise Rubidium frequency standard. This device provides a 10 MHz signal that is used as a reference signal for a phase-locked oscillator with an output frequency of 10.0625 MHz. This frequency is distributed to all low-level RF controllers which derive their operating frequency (depending on the type of cavity 40.25 MHz, 80.5 MHz, 120.75 MHz, or 322 MHz) from this reference signal using an internal phase-locked oscillator. In a similar way the clock signal for the Event Generator (EVG) is derived from the RF reference signal. This approach ensures low jitter of the event system relative to the RF reference frequency.

Time Structure of the Beam

The FRIB driver linac requires a wide variety of beam-pulse patterns. During nominal 400 kW operation it is operated with a pulse rate of 100 Hz and a duty cycle of 99.5%. The average beam power can only be ramped up slowly to 400 kW to reduce thermal stress to the target. The ramp-up procedure consists of four phases during which the pulse rate and pulse width are increased as defined by Table 2. As soon as a phase has been completed the next phase starts automatically. After 500 s the power ramp-up is complete and the system proceeds to nominal operation.

A variety of pulsed operation modes have been defined for commissioning and beam tuning. Some of them can generate either a single shot triggered by the operator or periodic pulses with adjustable frequency and pulse length (within certain constraints).

FRIB's timing patterns need to be expected to evolve over time. For example the power ramp-up procedure might get optimized once measurements of the actual target temperature are available.

Due to the complexity of FRIB's time structure we have decided to let the timing master broadcast "beam on" and "beam off" events. Thus timing clients do not need to have any knowledge about the pulse pattern. No large memories are required for storing pulse patterns in each timing client and no distribution system for pulse patterns is needed. On the other hand the large number of events that need to be played out in real time during the power ramp-up make this approach challenging on the timing master side.

Streaming of Event Data

The EVG provides two hardware sequencers that are implemented in its FPGA. They allow the EVG to play out a predefined series of events in real time. Each sequencer consists of a memory that holds up to 2048 time-stamp – event code pairs as well as control logic that compares the time-stamp of the next event in the list to the current time and sends out the event code if a match is detected. For most other accelerator facilities that are using the MRF timing system so far (mainly synchrotron light sources) 2048 time-stamp – event code pairs is sufficient but this is not enough to store the millions of time-stamp – event code pairs required for FRIB's power ramp-up. To work around this limitation both sequencers are used alternately ensuring continuous operation. The sequence of events is broken into 10 ms chunks containing one machine cycle each. These chunks are loaded into the sequencers alternately. Even during the power ramp-up mode which requires up to ≈ 25 000 "beam on"/"beam off" event pairs per second the maximum number of events per machine cycle is ≈ 500 which leaves enough headroom for additional diagnostics events and other extensions.

During initialization both sequencers are configured for single-shot operation. A multiplexed counter generates a 100 Hz trigger signal that starts both sequencers in case they are armed. The "start of sequencer 1" interrupt is set up to trigger filling and arming of sequencer 2 (in this order). The "start of sequencer 2" interrupt triggers filling and arming of sequencer 1. To start the process, the events for the first cycle are loaded into sequencer 1 and it is armed.

Note that the sequencers are configured for single-shot operation which means they need to be re-armed before they will play out a sequence again. Only one of them is armed at a time and the other one will thus ignore the trigger. Also note that this behavior also makes the system fail in a

ISBN 978-3-95450-189-2

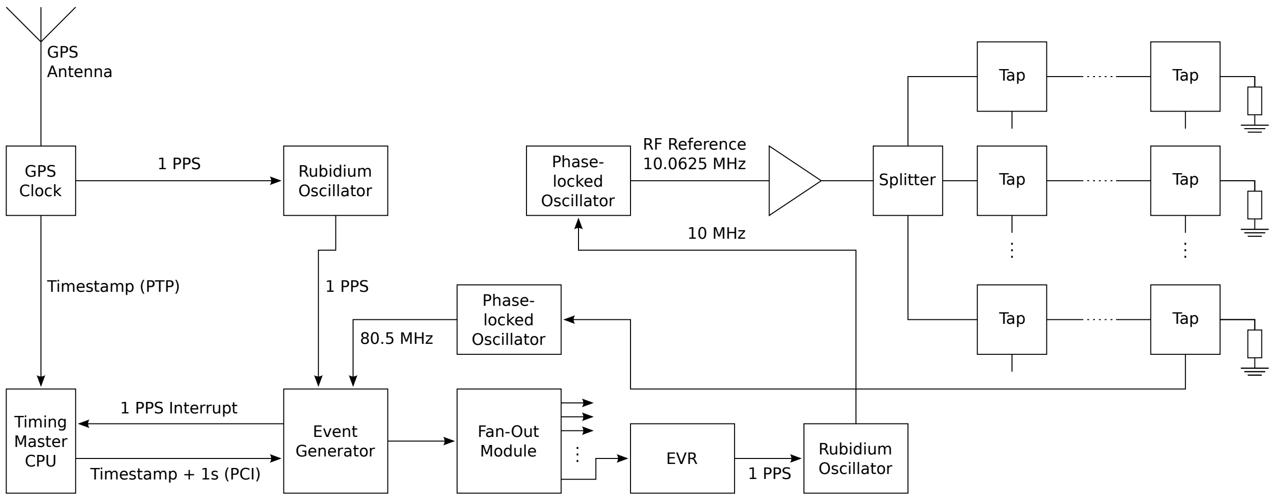

Figure 2: Synchronization of the event timing system to the GPS master clock as well as synchronization of the RF reference oscillator to the the event system.

clean way: In case the computer crashes the EVG will finish playing out the loaded machine cycles. When finished the sequencers will not be re-armed and the EVG automatically stops broadcasting events. Breaking up the event stream into machine cycles ensures that the beam is always off when the EVG stops broadcasting events.

The streaming of the event data has been implemented on the timing master EPICS IOC using the mrfioc2 driver [4]. Recent code contributions to this driver by FRIB and Paul Scherrer Institut have added support for the cPCI-EVx-300 hardware series and the start-of-sequencer interrupts. Leveraging these features, continuous data streaming to the EVG's sequencers can now be implemented on the EPICS database level. Development of C/C++ code is not required.

Note that each time-stamp – event code array needs to be computed and downloaded into the EVG in less than 10 ms. This real-time constraint is being met reliably by our EPICS IOC running on a standard Intel CPU. With Debian GNU/Linux we are using the same operating system that is running on all other FRIB controls computers. However for this application a real-time kernel is being used.

Beam Scheduler

The time structure of the beam ("beam schedule") is calculated in software by the EPICS Input/Output Controller (IOC) depending on the beam mode and the parameters selected by the operator. Depending on the beam mode the number of pulses as well as the pulse width are either fixed or provided by the operator directly or calculated dynamically for each machine cycle (e. g. during the power ramp-up). This behavior is implemented in the EPICS database to simplify online troubleshooting and to facilitate potential future modifications. For each machine cycle an aSub record generates the time-stamp – event code array based on the the number of pulses and the pulse width. This data is then downloaded into the hardware sequencers. The aSub record automatically distributes the beam pulses evenly over the allowed period to spread out the heat load evenly on the rotating target.

A finite-state automaton implemented in the EPICS state-notation language defines the sequence of states for each beam mode (e. g. the different phases of the power ramp-up). It also handles exceptions like a trip of the machine protection system.

SUMMARY

By using the MRF timing system, PTP, and NTP, the FRIB timing system is able to provide time-stamps and triggers to many hundreds of clients at low cost. By synchronizing the three subsystems data time-stamped by different subsystems can still be correlated. The RF reference frequency is also synchronized to the timing system.

The event stream for the fiber-based timing system is generated in real time by the EPICS IOC of the timing master. The event data is then streamed into the EVG for real-time distribution. This allows the timing master to play out event sequences of arbitrary length and complexity without the need for custom firmware or the client devices having knowledge about the time structure of the beam. Changes to the time structure of the beam require only changes to the EPICS record database of the timing master making this solution very flexible.

REFERENCES

[1] FRIB, `http://www.frib.msu.edu`

[2] J. Wei *et al.*, "FRIB Accelerator: Design and Construction Status," in *Proc. 13th Int. Conf. on Heavy Ion Accelerator Technology. (HIAT'15)*, Yokohama, Japan, September 2015, paper MOM1I02, pp. 6–10.

[3] Micro-Research Finland Oy, `http://mrf.fi/`

[4] mrfioc2 device driver for Micro-Research Finland Hardware, `http://epics.sourceforge.net/mrfioc2/`

ISBN 978-3-95450-189-2