

SNS 109020200-TD0006, R00

SNS Timing Master Functional System Description

April 2011



A U.S. Department of Energy Multilaboratory Project

SPALLATION NEUTRON SOURCE

Argonne National Laboratory • Brookhaven National Laboratory • Thomas Jefferson National Accelerator Facility • Lawrence Berkeley National Laboratory • Los Alamos National Laboratory • Oak Ridge National Laboratory

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

Eric Bjorklund, bjorklund@lanl.gov
Dave Thompson, dht@ornl.gov

Date Published: April 2011

Prepared for the
U.S. Department of Energy
Office of Science

UT-BATTELLE, LLC
managing
Spallation Neutron Source activities at
Argonne National Laboratory Brookhaven National Laboratory
Thomas Jefferson National Accelerator Facility Lawrence Berkeley National Laboratory
Los Alamos National Laboratory Oak Ridge National Laboratory
under contract DE-AC05-00OR22725
for the
U.S. DEPARTMENT OF ENERGY

Table Of Contents

1	Timing System Overview	1
1.1	Timing System AIP.....	1
1.2	Timing System Interfaces	1
1.3	Timing System Concepts	3
1.3.1	Machine Cycle Time Line	3
1.3.2	Turns and Sub-Revolutions.....	5
1.3.3	Macro Pulses, Mini Pulses, and Micro Pulses	6
1.3.4	Pulse Flavors.....	6
1.4	MPS Connections.....	7
2	Event Scheduling	9
2.1	Beam Scheduling	9
2.2	Rep-rate computation.....	9
2.2.1	Master Rep-Rate	9
2.2.2	Calculation of rep-rate patterns.....	10
2.2.3	Data Acquisition Trigger events	10
3	EPICS Software	11
3.1.1	EPICS database and device support.....	11
3.1.2	Timing Master Sequencer Programs.....	11
3.1.2.1	Beam-On Switch	12
3.1.3	Timing Master Drivers.....	12
3.1.4	Future Requirements.....	12
4	Event Link.....	13
4.1	SNS Event Definitions	15
4.1.1	Hardware Events.....	15
4.1.2	The Beam-Related Event Cluster.....	18
4.1.3	Software Events	19
5	Real Time Data Link.....	21
5.1	RTDL Frame Definitions	21
5.1.1	Time Stamp (Frames 1-3).....	21
5.1.2	Ring Revolution Period (Frame 4).....	22
5.1.3	MPS Mode (Frame 5)	22
5.1.4	60 Hz. Phase Error (Frame 6)	22
5.1.5	Beam Width (Frame 7)	23
5.1.6	Line Frequency (Frame 8)	23
5.1.7	IOC Reset (Frame 15).....	23
5.1.8	Pulse Flavor (Frame 17)	23
5.1.9	Last Cycle Veto (Frame 24).....	23
5.1.10	Cycle Number (Frame 25)	24
5.1.11	Master Rep-Rate (frame 26)	24
5.1.12	Chopper Control Frames (frame 28 to 32 and 36 to 38).....	24
5.1.13	Last Peak Beam Current (Frame 33)	26
5.1.14	Last Integrated Beam Current (Frame 34).....	27
5.1.15	Last Integrated Beam On Target (Frame 35).....	27
5.1.16	BLM Reprate (Frame 40).....	27
5.1.17	Stored turns Frame (Frame 41).....	27
5.1.18	Message CRC (Frame 255).....	27
6	Master Reference Generator	29
7	Additional Hardware in the timing master.....	31
7.1	Global Positioning System Interface.....	31

7.1.1	Software	31
7.1.2	Hardware.....	31
7.2	Reflective Memory Board.....	32
7.3	Line Sync Pulse Generator.....	32
7.4	Event Link Clock	32
7.5	MPS Interface	32
8	Bibliography	35
9	Contacts.....	36
10	Index	37

1 TIMING SYSTEM OVERVIEW

The SNS timing system consists of two data transmission links distributed to all accelerator systems. The *Event Link* transmits 8-bit timing events that define the SNS *Machine Cycle*. The events on the event link are synchronized with the accumulator ring RF frequency. The *Real Time Data Link* (RTDL) transmits a series of 24-bit data frames prior to the beginning of a machine cycle. These data frames contain information about the next cycle such as the time of day, the type of pulse to be delivered (pulse *flavor*), information on whether the preceding pulse was good etc. An SNS *Timing Gate Trigger* module translates timing events into TTL gates. An SNS *Utility Module* listens to both the event link and the RTDL. The utility module can interrupt the IOC on the occurrence of specified events. It can also provide the IOC with data from the RTDL. The Machine Protection System uses the RTDL to set the machine mode masks throughout the system.

This document is a description of the timing system and does not make a distinction between hardware and software functions. There are separate documents that describe the low level details of the hardware and software. References to the BNL timing master hardware have been removed where possible. (1)

1.1 TIMING SYSTEM AIP

The timing system AIP had several goals when it was approved.

- Increase reliability and availability.
- Migrate to up to date hardware that can be maintained for several years into the future while reducing the board count and complexity.
- Design the hardware to operate standalone with limited functionality without the EPICS software even running. This will allow software upgrades without impacting systems using timing.
- Eliminate a number of nagging problems that require constant monitoring and event retuning.
- Maintain complete compatibility with existing timing client hardware.

None of these goals can be stated in a functional system description that just deals with system inputs and outputs. The original system meets this FSD and in fact is the FSD. Much of the motivation for the details of the design of the hardware and software flows from these AIP goals.

1.2 TIMING SYSTEM INTERFACES

There are a small number of hardware interfaces to other systems in the accelerator and a lot of software interfaces. The timing system is closely coupled to the MPS by both hardware and software. The hardware interfaces include a reflective memory board to get RTDL data from beam instruments, the line and event link clocks, a GPS time of day clock, and links to MPS for beam interrupting events.

Software interfaces include several from the MPS, the timing soft IOC, ion source PVs, and RFQ PVs. These PVs support proper machine setup and limits and are an integral part of the MPS. The software on the timing master IOC is assisted by a soft IOC and both software packages work together.

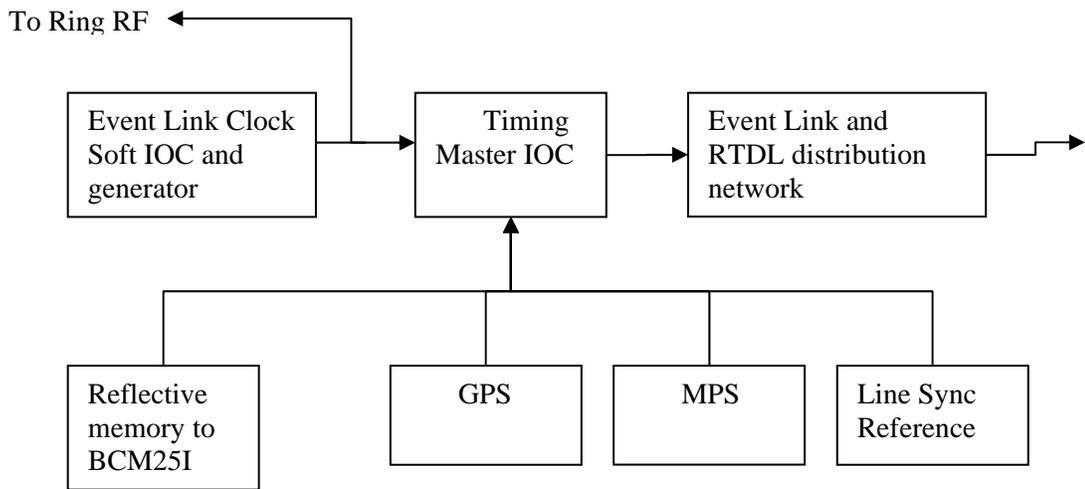


Figure 1 Hardware Interfaces

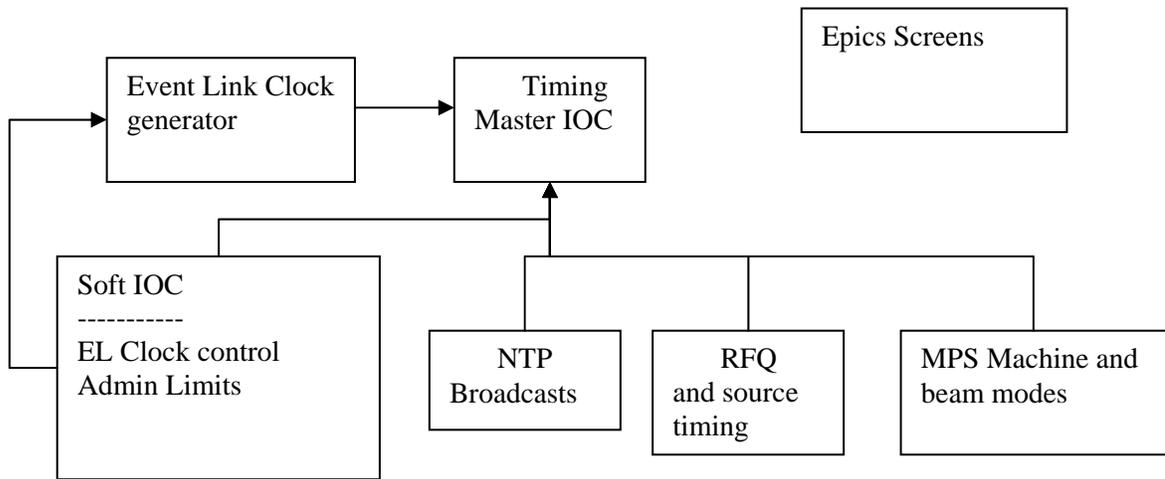
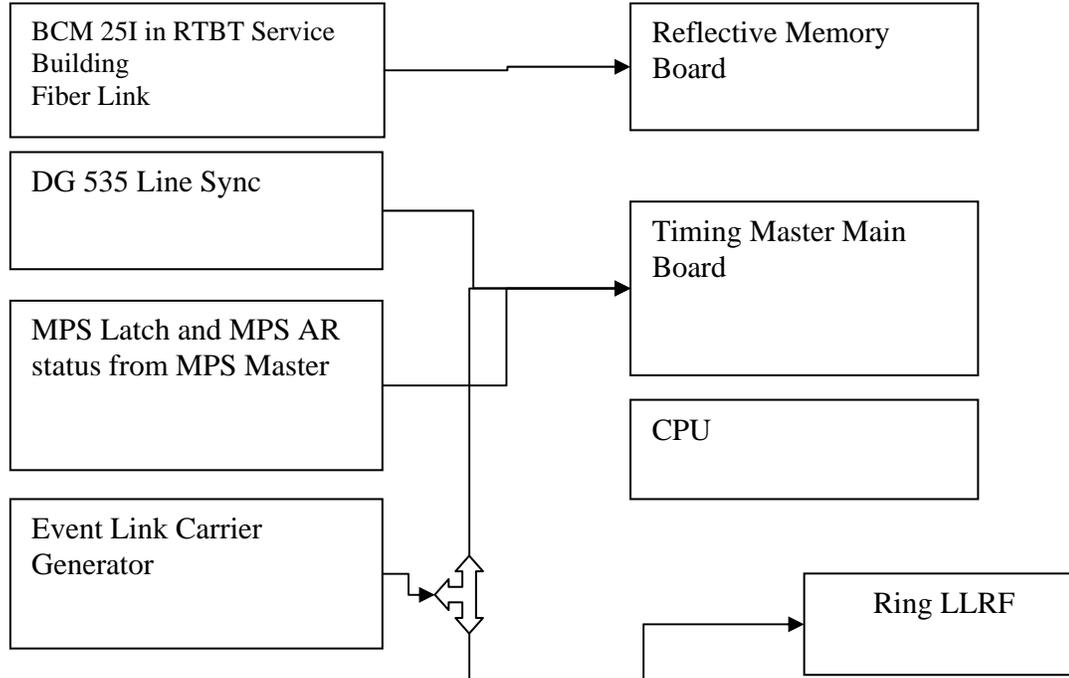


Figure 2 Software Interfaces



1.3 TIMING SYSTEM CONCEPTS

1.3.1 Machine Cycle Time Line

The fundamental repetition rate of the SNS timing system is 60 Hertz. The starting point for a machine cycle is a 60 Hz “smoothed” signal derived from the positive zero crossing of the AC line frequency. The maximum repetition rate for delivering beam to the first neutron production target is 60 Hertz. The maximum repetition rate for delivering beam to the second neutron production target is 20 Hertz. When the second target is implemented, there is some consideration that the 20 Hz of beam to target 2 might be delivered from the negative zero-crossing, allowing us to continue to deliver 60 Hz of beam to target 1. Consequently, the timing system design must allow for the possibility of 120 Hz operations, which means that the maximum length of a machine cycle can only be 8.3 milliseconds.

The SNS machine cycle is subdivided into two sections, a “Time Critical” section and a “Non-Critical” section. The time critical section begins at the start of the cycle, includes the acceleration of beam in the linac, the accumulation of beam in the ring, and ends with its extraction from the ring and delivery to the neutron production target. The parts of a machine cycle are defined by *Events*, which are broadcast on the SNS *Event Link*. Events can be either hardware or software generated. Software-generated events are not allowed to occur during the time-critical portion of the machine cycle, since they might interfere with the delivery of time-critical hardware events.

The most time-critical event in the machine cycle is the “Extract” event. This is the event synchronizes the *Extraction Kicker* magnets and the *Neutron Choppers*. The neutron choppers are located between the neutron production target and the experimental instruments. They come in two basic flavors, *T-Zero* and *E-Zero*. The job of the T-Zero choppers is to “chop” out the initial high-energy neutrons that get ejected immediately after the protons hit the neutron production target. These choppers

are high-inertia and cannot respond quickly to changes in their synchronization pulses. One of the jobs of the timing system, therefore, is to try to keep the Extract event as stable in time as possible, in spite of fluctuations in the AC line frequency. Another job of the timing system is to try to keep the machine cycle from getting too far out of phase with the AC line frequency. We discuss the management of these two conflicting goals in subsequent sections of this document.

It is not desirable to leave beam in the accumulator ring any longer than necessary. Therefore (and unlike more traditional accelerators), the injection of beam into the accumulator ring is timed such that it always ends slightly before the “Extraction” event. A special “End-Inject” event occurs just before Extract and defines the fixed point where the accelerator RF and beam gates end. If more beam is desired, the beam and RF gates get longer by growing toward Cycle-Start. If less beam is desired, the beam and RF gates compact toward the End-Inject event.

The picture below illustrates the basic layout of the SNS machine cycle.

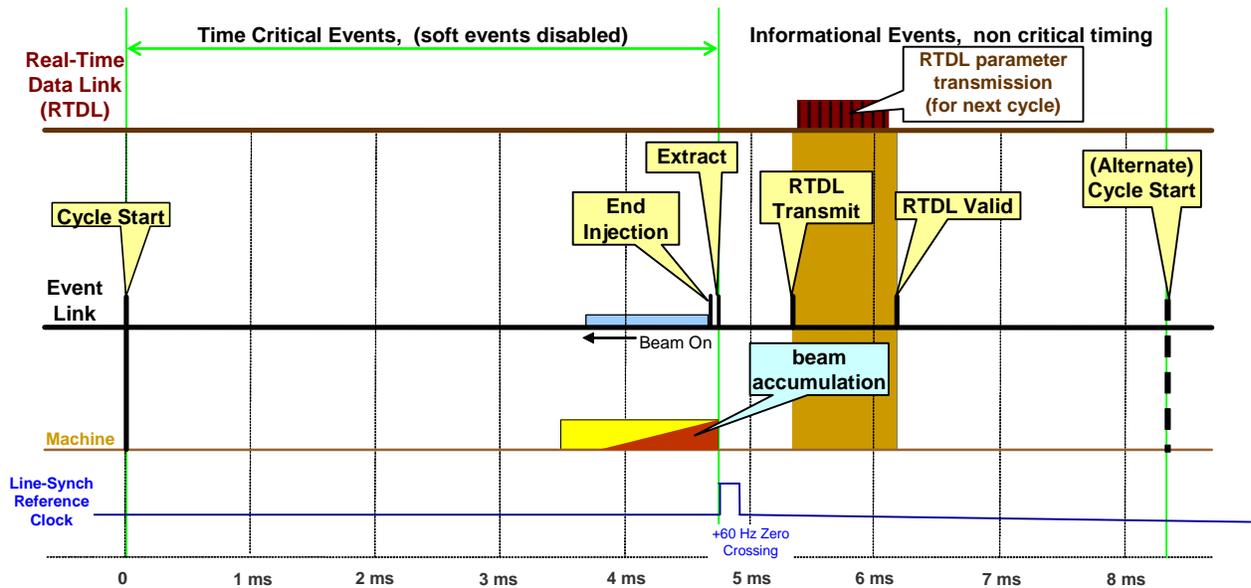


Figure 3
Basic Machine Cycle Time Line

The time-critical portion of the cycle begins with the “Cycle-Start” event and ends with the “Extract” event. Software-generated events are not allowed to occur within this period of time. If software does request an event during the time-critical portion of the machine cycle, it will be queued in a FIFO and not released until after the Extract event. Events that control the firing of the accelerator RF gates, front-end choppers, and data acquisition diagnostic gates also occur within the time-critical section.

The non-time-critical portion of the machine cycle begins after the Extract event and continues until the next Cycle-Start (or Alternate Cycle-Start, if running in 120 Hz mode). During this time, software-generated events are allowed to occur. Software-generated events announce non-time-critical activities such as global error-counter clearing, or the start or stop of a new rep-rate pattern generation. This is also the period when the *Real Time Data Link* (RTDL) frames for the next machine cycle are transmitted. Shortly after the Extract event, the timing master sequencer program computes the appropriate RTDL frame values for the next machine cycle and then generates the “RTDL Xmit” event. The RTDL-Xmit event triggers a gate that instructs the RTDL encoder module to transmit the current set of frames over the link. After the entire set of RTDL frames has been sent, the “RTDL Valid” event signals that all frames have been sent and are available for use by client IOCs or diagnostic NADs.

1.3.2 Turns and Sub-Revolutions

The fundamental unit of the SNS timing system is the *Turn*, which is defined as the amount of time it takes the beam to travel around the accumulator ring one time. For rough approximation purposes, one turn is approximately equivalent to one microsecond. In reality the value of a turn can vary between 911 and 974 ns depending on the energy of the beam and the size of the orbit. The dominant term is the beam energy which can vary between 842 MeV and 1.3 GeV. As shown above in Figure 1, the timing system clock and the RF signal that maintains the extraction gap in the beam are derived from the same source. The ring clock frequency is 32 times the revolution frequency, which gives the timing system clock a granularity of about 29.5 nanoseconds at 1 GeV. Each timing system clock tick is referred to as a *Sub-Revolution*.

Using the ring revolution frequency as the basis for the timing system clock allows the neutron choppers to be precisely synchronized with the extraction kick. It also allows us to tune the ring, set different energies, and adjust the size of the orbit without having to constantly change all the timing parameters.

The disadvantage, of course, is that if you have any “wall clock” timing constraints, you now have to deal with “error bars” around your settings. For example, if you want to guarantee that the High-Power RF gates turn on at least 100 microseconds before the Low-Level RF gates, you need to set the difference to 110 turns in order to accommodate the worst-case (fastest) clock speed at 1.3 GeV. In most cases the driver software at the receiver card makes this adjustment automatically.

The table below gives the relationship between beam energy, the ring RF frequency, and the turn and sub-revolution times: (2)

Beam Energy	Mini-pulse Length (ns)	Gap Length (ns)	Revolution Period (ns)	Revolution Frequency (MHz)	Clock Frequency (MHz)	Beta	Sub-Revolution Period (ns)
842 MeV	655.1	324	973.4	1.027,323	32.874,340	84.984 %	30.42
1.0 GeV	630.3	315	945.4	1.057,767	33.848,545	87.503 %	29.54
1.3 GeV	607.4	304	911.2	1.097,502	35.120,070	90.790 %	28.47

Table 1
Beam Energy and Revolution Frequency

The ring period calculation is based on the ring circumference of 248 meters¹ (depending on tuning) and the Beta of the beam. $F_{rev} = \beta C / 248$ $\beta = \sqrt{1 - E_p^2 / (E_k + E_p)^2}$ where E_p = rest mass of proton and E_k is the kinetic energy.

1.3.3 Macro Pulses, Mini Pulses, and Micro Pulses

An SNS beam pulse has three levels of structure. The highest level is the *Macro Pulse*, which is defined by timing system gates in the Machine Protection System (MPS) master. A macro pulse can be anywhere from 1 to 1060 turns (about 1 millisecond) long. A turn is defined by the length of time required for the beam to complete one revolution around the ring, about 240 meters. The desired size of the macro pulse determines how long the accelerator RF needs to be on. It also determines the intensity of the beam delivered to the neutron production target.

When the beam enters the accumulation ring, it is compressed from the approximately one millisecond macro pulse into an approximately one microsecond *Mini Pulse*. In order for this to work, the macro pulse has to have a structure imposed on it in which there is a “gap” of about 300 nanoseconds between each injected turn. This is known as the *Extraction Gap*. When it comes time to extract the beam from the ring, the extraction kickers turn on during the 300 nanosecond extraction gap. This allows the kickers to be at full power by the time beam arrives and prevents “beam spraying”. The mini pulse structure is imposed by the LEBT chopper, which gets its clock from the timing system.

The finest level of structure is the *Micro Pulse*, which reflects the accelerating structure imposed by the linac RF. The LEBT chopper can “ramp up” the beam intensity from the start to the end of the macro pulse by selectively “chopping out” selected mini pulses and by adjusting the width of each mini pulse.

The figure below shows the relationship between the macro pulse, mini pulse and micro pulse structures.

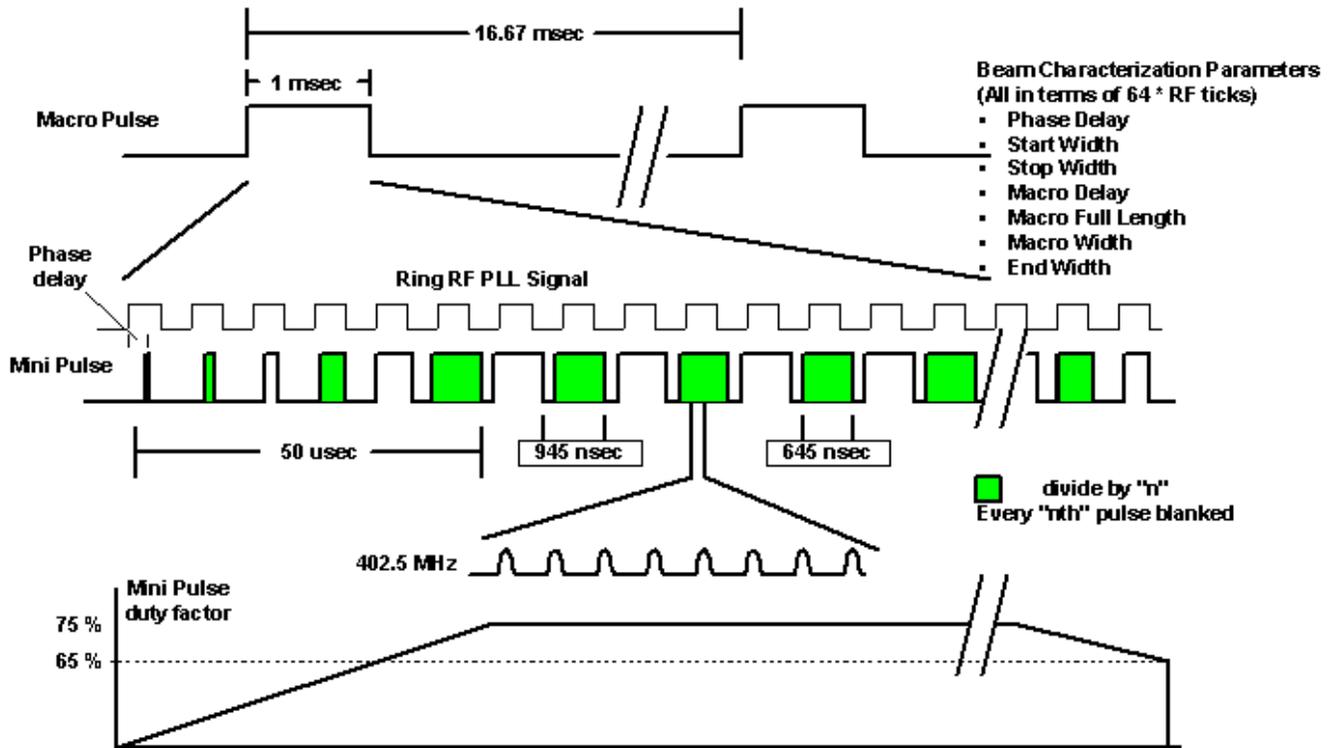


Figure 4
Macro, Mini, and Micro Pulses

1.3.4 Pulse Flavors

Each machine cycle has a *Pulse Flavor* associated with it. The pulse flavor for the next machine cycle is transmitted on the RTDL. There are eight possible flavors for a machine cycle. These are:

0	No Beam
1	Normal Beam (Target 1)
2	Reserved for Target 2
3	10 μ Second Diagnostic Pulse (not used)
4	50 μ Second Diagnostic Pulse
5	100 μ Second Diagnostic Pulse
6	Special Physics Pulse 1
7	Special Physics Pulse 2

Beam flavors are used by the LEBT chopper to determine what the chopping pattern should be for the next cycle. They could also be used by various systems, such as Low-Level RF, to do adaptive tuning based on what kind of beam is being delivered. Using the SNS utility module, an IOC application can set up records to process only when beam of the particular flavor they are interested in has been scheduled.

The first 6 flavors are called “Normal” flavors. They reflect the current type of beam being produced and are usually determined from the MPS *Beam Mode*. The last two flavors (6 and 7) are called “Cycle Stealing” flavors. If they occur at all, they only occur at low repetition rates – typically 0.1 Hertz. They “steal” cycles from the normal flavored pulses. The stolen cycles have special parameters set by the LEBT chopper and are used for on-line accelerator physics experiments.

“No Beam” (flavor 0) means that no beam is scheduled for the next cycle. This could be because beam has been turned off, or the MPS beam mode is either “Off”, “Standby” or “MPS Test”, or the MPS is not made up for the current Beam/Machine mode. It could also mean that the beam is running at a repetition rate under 60 Hz and the next cycle is not one scheduled for beam.

“Normal Target 1” (flavor 1) means that the full intensity beam is scheduled for the next cycle. In order for the flavor to be “1”, the MPS beam mode needs to be either “1 Millisecond” or “Full Power”. This does not mean that the actual beam gate needs to be 1 millisecond long.

Flavor 2 is reserved for the second target. It has the same properties as flavor 1, but reflects the parameters of the Target 2 beam. If the timing system is changed to operate at 120 Hz, then flavor 2 cycles are only scheduled on “alternate” (negative zero-crossing) cycles. Flavor 2 becomes the “other” normal flavor in this case. If the timing system continues to run at 60 Hz, then flavor 2 would steal cycles from flavor 1 (up to 20 Hz). The current plans for target 2 include delivering beam directly from the linac to the target since planned target 2 instruments do not need short pulses. It is also a possibility that flavor 2 will be un-chopped in order to increase the available power to target 2 by as much as 30%.

Flavors 3, 4, and 5 correspond to the MPS beam modes for 50 μ Second, 100 μ Second, and 100 μ Second, diagnostic modes respectively. These modes reflect restricted intensity conditions due to the use of intercepting diagnostics. Consequently, these flavors cannot be “shared” with each other or the other “normal” flavors. Neither can they steal pulses from the other normal flavors.

Flavors 6 and 7 are for intended for on-line accelerator physics experiments. They are intended to run at low rep-rates (0.1 Hz) so that they do not noticeably interfere with normal beam. They may steal cycles from any of the other “normal” flavors, provided their pulse width does not violate the current MPS operating mode.

Note that the present time only flavor 1 is used for production beam delivery.

1.4 MPS CONNECTIONS

The MPS has two connections to the timing master hardware, one for the MPS “Auto-Reset” signal and one for the MPS “Latched” signal, that go directly into two event encoder channels. These connections come from the MPS “Trigger Control Chassis” by fiber optic link from the front end and will cause an event to be broadcast on the event link whenever an MPS fault occurs. In addition, the MPS status signals are monitored on digital input channels. The timing master IOC monitors these signals and disables delivery of the “Beam-On” event if MPS is faulted. The difference between a “Latched” and an

“Auto Reset” fault is that the “Latched” fault requires an action from an operator to restore the beam once the fault has cleared. The timing master IOC can also drop MPS itself, if it detects a condition that would indicate that beam should not be delivered.

2 EVENT SCHEDULING

2.1 BEAM SCHEDULING

Beam scheduling is a major function of the timing master real time software. Scheduling starts with the computation of the pulse pattern over the 600 cycle super cycle. This computation must take into account the master rep-rate so that all beam cycles coincide with RF cycles at the master rate. Pattern computation should only take place when a new beam rep-rate is requested or when the master rate is changed.

The hardware is arranged so that beam-on can only be set to trigger one time so that the software has to re-trigger it on every cycle. This way the software cannot fail and leave the Beam-On uncontrolled. The software makes the beam scheduling decision between the extract event of the previous cycle and when the RTDL message is transmitted for the next cycle. Several other pre-conditions are required to generate a beam on event and flavor greater than zero as listed below.

Beam-On Pre-conditions
Kicker charge event issued at the end of the previous cycle
Beam-On Switch EPICS PV is ON
MPS AR Status is NOT FAULTED
MPS Latch Status is NOT FAULTED
Beam-On pattern bit for the current cycle is set
Beam is not in single shot OR a single shot is requested

Table 2 Beam On Pre-conditions

2.2 REP-RATE COMPUTATION

All rep-rate patterns except the beam-on and source-on pattern are calculated at startup and never modified thereafter. The only rep-rate that normally needs to be recalculated is the beam rate. The only constraint on all event rep-rate patterns is that cycle 599 has to be “on”, all events must trigger on cycle 599.

2.2.1 Master Rep-Rate

The master rep-rate is a selection of one of the even divisors of 60 Hz. The selected master rep-rate defines the maximum beam rep-rate and the minimum rep-rate for RF and ring kicker magnets. A set of events, one for each master rep-rate choice, are sent on the event link. The contract is that “If a beam event occurs it will occur on a cycle with the event corresponding to the master rep-rate.” The Master Rep-Rate is included in an RTDL frame and can be used to determine if a rep-rate selection in a subsystem is valid. For example, an RF system can use any of the RF events which have the same selection of rep-rates as the master. The RF system can select the same or a faster event as the master but some combinations are not usable either. For example, if the master is at 20 Hz, the 30 Hz event cannot be used by an RF system because the 30 Hz event does not occur on all cycles when there is a possible beam event. The master rep-rate will be set at 60 Hz. except during machine restart and physics study times when 1 Hz is typically used.

2.2.2 Calculation of rep-rate patterns.

There are two ways to calculate a pattern. One is based on an algorithm presented by Eric Bjorklund (3). In this case an algorithm tries to evenly distribute variable rep-rate patterns over the period of a super cycle. Part of the work in this assumed that patterns may have to be constrained by other patterns so that the events in a second pattern line up with events in a first pattern even though the second pattern may be slower. A related report detailed a way to measure the “evenness” of a pattern so that the smoothest, least ugly, pattern can be chosen. (4)

Unconstrained patterns similar in evenness to Bjorklund’s algorithms can be calculated using a technique similar to Bresenham’s line algorithm. (5) Using this technique the code iterates across all of the slots in a pattern while adding the rep-rate to an accumulator each time. When the accumulator overflows the slot is marked to enable the event on the corresponding cycle. Done with scaled integer math this always results in the last slot getting a 1, even if the pattern is for 0.1 Hz and the last slot is the only one set.

2.2.3 Data Acquisition Trigger events

The timing system provides several data acquisition events that are mainly used by beam instruments. There is a pair of “Fast” events, a pair of “Slow” events, a laser wire trigger event, a laser wire acquisition event, an “On Demand” event, and a “No Beam” event. The rep-rates of the fast and slow events can be adjusted; however they should normally be set to 5 Hz and 1 Hz, respectively while the laser events should be at 30 and 10 Hz.

Events named “DIAG_*” are events that are used for accelerator and ring diagnostics where the diagnostic trigger must be adjusted by the timing master with the beam gate width. Events named “DIAG_RTBT_*” are events intended for the RTBT and target systems where the timing for the event is relative to T_{ext}

The “Fast” events only fire on cycles that have beam events. The “Slow” gate only fire on cycles that have a “Fast” event (and, by implication, beam events). The “On Demand” event only fires on cycles that have both “Fast” and “Slow” (and “Beam”) events. The “On Demand” event is fired by pressing the “Trigger” button on the “Data Acquisition Triggers” section of the screen. The “No Beam” event fires at the same rate as the “Fast” event, but out of phase with the beam. Unlike the “Fast” event, the “No Beam” event will fire regardless of whether there is any beam scheduled for that cycle or not.

All diagnostic events are scheduled with RF cycles, so if the slow event is set to 1 Hz it will happen on cycles that also have 1 Hz events. Likewise if the fast diagnostic event is set at 5 Hz it will occur on cycles when there is a 5 Hz RF event.

3 EPICS SOFTWARE

The EPICS software for the timing master IOC consists of:

- The timing master sequencer program.
- EPICS databases for monitoring and controlling the timing events and RTDL frames.
- EPICS software running in related soft IOCs that monitor beam power and energy.

The timing system is about generating events on the Event Link and data on the RTDL and the software and hardware have to cooperate in real time to provide timing to the accelerator. Some EL events and RTDL frames are entirely managed by the FPGA firmware, some EL and RTDL frames are managed by the EPICS software layer and the sequencers and just transmitted by the FPGA, and some RTDL frames that are initialized on each cycle and updated by the sequencer before being transmitted.

The architecture of the software should be partitioned according to when the software module should run. The AIP timing master provides interrupts at specific points in the time line that can be used to schedule module execution. The legacy timing master does the same thing but depends on the specific features of that collection of hardware and so does things in a more ad-hoc way. In many cases the AIP timing master uses hardware to do functions that the legacy timing master did in software, so the software modules have less to do.

In the discussion below, each software module will be described as a sequencer. A sequencer would be implemented as a main loop that waits on some external event or interrupt. This has nothing to do with an EPICS sequencer.

3.1.1 EPICS database and device support

The EPICS interface for the timing system controls and supervises the timing system and provides much of the data for the RTDL. There are tight links to the MPS, the event link clock soft IOC, and to other systems that can affect beam status. Much of the interaction between the timing sequencers and EPICS in the legacy timing system is implemented through a common memory block and the EPICS “symb” device support. The sequencers pick up data from common memory and update status in common in a loose asynchronous fashion. The sequencers are the synchronous connection between EPICS and the FPGA.

There are drivers and device support for specialized subsystems such as the GPS receiver, the line sync module, the soft event FIFO, and the IOC rebooter. The drivers connect EPICS records directly to the subsystem both in and out.

3.1.2 Timing Master Sequencer Programs

There are three sequencer tasks coordinated by timing master board interrupts. These tasks run at different times in the cycle so that sequencer tasks are executed in a just in time flow. All three of the tasks are coordinated by events coming from the master board through interrupts to the interrupt service routine down to the tasks using binary semaphores. Despite the term “task” do not infer a literal vxWorks task as an implementation directive. There may be one or many tasks in the implementation; we are only specifying when things need to run relative to the timeline.

The cycle start task does the following: Updates the time of day from GPS, NTP, and the on board clock, chooses the best source, predicts the time for the next cycle start, and then stores that time for use on the RTDL in the cycle end task.

The line sync task does the following: Read from the timing master board the phase and elapsed time from the previous cycle. Calculate the input line frequency as a running average over an adjustable number of cycles. Calculate the timing, in seconds, between cycle start and the actual line crossing. Apply low pass filtering and scaling for the phase and frequency, apply slew rate limiting, and then calculate and set the frequency register of the master reference generator.

The cycle end task does the following: Decide if the next cycle is a beam cycle and set the Beam-On flag and diagnostic event flags as needed. Set the timestamp RTDL frames. Set the flavor and other RTDL frames. Set the LEBT and MEBT RTDL frames. Adjust the position of the Beam-On and diagnostic events in the event ram. Copy other enabled RTDL data to and from Common for EPICS access. Coordinate several informational soft events. Copy soft events from the soft event vxWorks FIFO to the hardware soft event FIFO. Last, get the beam on target data from BCM25I out of the reflected memory card and set the appropriate RTDL frames. Trigger RTDL send on the master board.

3.1.2.1 Beam-On Switch

The Beam-On Switch is an EPICS PV in the timing master database. A CalcOut type record forces the beam switch to OFF when one of several interlocks are not made up, including MPS Latch, front end statuses, Chief Operator Switch, RF and Ion Source events not turned on, and an invalid event link clock.

3.1.3 Timing Master Drivers

These EPICS support modules initialize and monitor the hardware in the timing master IOC. Writing these modules as drivers ensures that IOC configuration follows common practice, and that access to registers from EPICS records is supported and controlled. Drivers will either be written or adapted for the master board, the GPS board, and the reflective memory board installed in the master IOC.

3.1.4 Future Requirements

There are two major projects on the horizon that may impact the timing master, the power upgrade project which is now at CD2, and the second target station now at CD0.

The PUP project will increase the beam energy, beam pulse width, and beam current. The energy increase will increase the ring revolution frequency and hence the event link clock. The maximum number of injected turns may have to be increased slightly to provide for a 1 ms macro pulse.

The second target station will have a greater impact. The timing system will be called upon to schedule beam to two targets, 40 Hz to target 1 and 20 Hz to target 2. In addition, the second target is expected to receive an un-chopped beam directly from the linac. The timing system will be called upon to generate an extra trigger event for either the first or second target station or both so that kickers to steer the beam to the appropriate target may be triggered. Which controls may be provided for the second target will depend on operation's requirements at the time.

4 EVENT LINK

The SNS Event Link is modeled after the RHIC Beam Synchronous Link [4]. It is a differential PECL, bi-phase mark encoded, serial link that transmits 8-bit event codes synchronized with the ring revolution period. An SNS event actually contains 12 bits. These are a start bit, an 8-bit frame number, a parity bit, and two stop bits. The format of an event frame is show below:

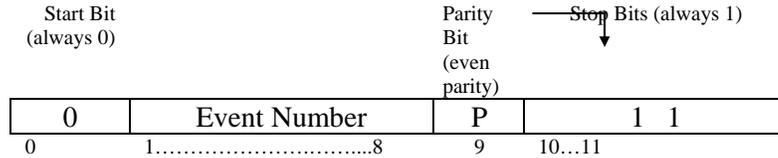


Figure 5
Event Frame Format

The frame is transmitted in “Big-Endian” format (i.e. most significant bit first).

Events are divided into two categories – hardware events and software events. As the name implies, hardware events are generated by the event link hardware. They have very precise timing and are prioritized based on the hardware port that generated them. Scheduled events have the highest priority followed by externally triggered events and soft events. The event encoder module contains an event mapping table that determines which event number actually gets transmitted when a particular hardware event is triggered.

Software events are generated (again, as the name implies) by software running on the timing master IOC. They do not have precise timing and are prioritized on a “First-In-First-Out” basis. Software events are usually either informational (e.g. the rep-rate patterns have been changed) or signal some non-time critical activity (e.g. clear all MPS error counters). Software events are not allowed to occur during the “Time Critical” portion of the machine cycle between Cycle-Start and Extract (see section 1.3.1). If a software event is requested during this time, it is stored in a FIFO and not released until after the Extract event. To meet this constraint all writes to the hardware soft event FIFO have to occur in the cycleEndTask.

SNS events have several functions:

- They define the fixed portions of the SNS machine cycle (i.e. Cycle-Start, End-Inject, Extract, or RTDL-Valid).
- They set relative locations for those variable delay and width gates that need to be synchronized across the entire accelerator (e.g. RF gates, start of beam, or certain diagnostic pulses).
- They control the repetition rates of variable-rate gates (e.g. ion source, RF gates, or start of beam).

Most of the events used to define gates with variable rep-rates, widths, and/or delays occur fairly far in advance of the actual gate being generated. This situation is shown below in the more detailed SNS machine cycle timeline diagram:

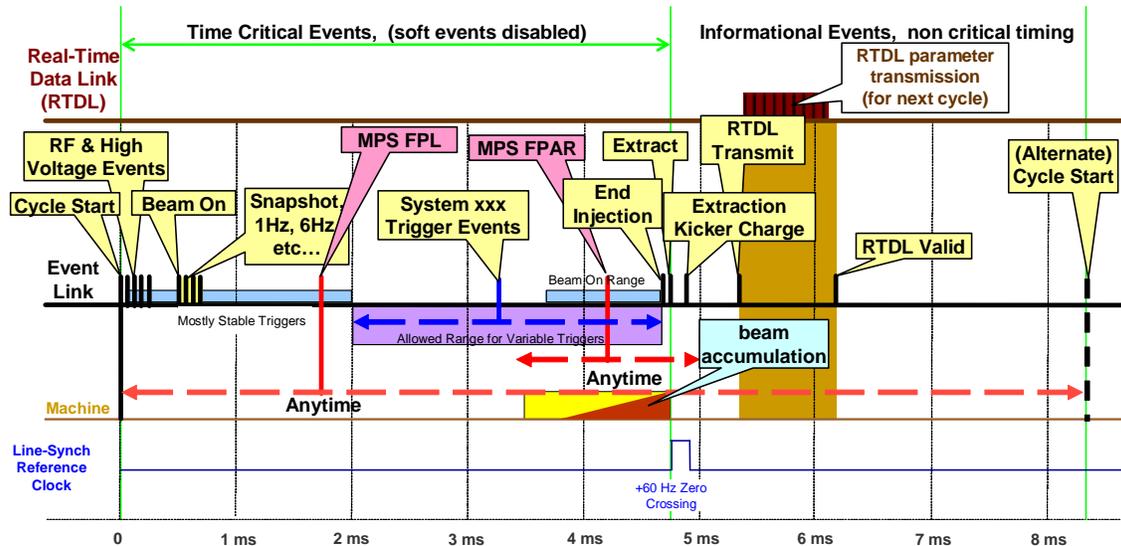


Figure 6
Detailed Machine Cycle Time Line

Notice that the variable rep-rate event definition events (RF, Beam-On, Diagnostic Triggers) all occur on the left hand side of the timeline, near Cycle-Start, in an area labeled “Mostly Stable Triggers.” The events they generate, however, all occur on the right hand side of the time-critical section of the timeline, in an area labeled “Beam-On Range.” This allows timing system client applications to define triggers or “pre-pulses” that can come up to about 2.5 milliseconds before the actual gate, if needed.

The label “Mostly Stable Triggers” means that while events can be moved around within this area, their positions will be stable during production beam delivery. The cycle end task can access event ram and adjust the positions of the mostly stable events when beam parameters change. All changes must occur in the cycle end task to prevent access collisions in the event ram.

The purple area labeled “Allowed Range for Variable Triggers”, with the “System xxx Trigger Events” box pointing to it. Is an area for system-specific single shot triggers events which will not interfere with the “Mostly Stable” events. No variable trigger events are in use at this time, the hardware provides for future implementation.

MPS trip events (Fast-Protect Auto-Reset and Fast-Protect Latch) can occur at any time during the machine cycle between the Beam-On event and cycle end and are transmitted without interfering with scheduled events.

The following table details the timing of the fixed portions of the machine cycle timeline. All times are specified in turns past cycle start.

Time-Critical Section	0-5050	Beam acceleration, accumulation, and extraction portion of the timeline. Software events are locked out during this time. There should be no event “jostle” during this period.
Mostly Stable Triggers	10-2110	Events for variable rep-rate and variable width beam and RF gates occur in this area. Although these events can move, the operational constraints of the accelerator should keep them from interfering with each other.
Variable Trigger Range	2111-5047	Single shot system-specific events occur during this period.
End Inject	5048	Ending time for all beam and RF gates.
Beam-On Range	2119-5048	This is the time period during which the gates triggered by events in the “Mostly Stable” region actually fire.
Kicker Charge	5062	Start charging extraction kicker for the next cycle. This event is actually part of the next cycle which is why the cycle counter is incremented just before this event.
Cycle-End	5150	This is the point, 100 turns after the end of the time-critical section, that the timing master sequencer computes the RTDL frames and enables the variable-rep-rate gates for the next cycle. Cycle-End can be delayed by 0 to 1000 turns when ring stored turns are required. This timing allows the beam on target IOC to transmit the most recent beam charge to the timing master.

**Table 3
Machine Cycle Timeline Definitions**

4.1 SNS EVENT DEFINITIONS

4.1.1 Hardware Events

The following table lists the hardware events currently implemented in the SNS timing system. The “SNS Event Number” is the number of the event that is broadcast on the event link.

SNS Event Num	Event Name	Position in Cycle*	Qualifiers	Comments, usage, and nominal time-in.
0	----	----	----	Not Used
1	Cycle-Start	0 Turns	60 Hz.	Defines the start of the time-critical section of the timeline.
3	MPS-Reset	Anywhere in cycle not conflicting with a fixed event	Triggered by MPS	Single-Shot. Triggered by MPS Auto-Reset signal.
4	MPS-Latch	Anywhere in cycle not conflicting with a fixed event	Triggered by MPS	Single-Shot. Triggered by MPS Latch signal.

SNS Event Num	Event Name	Position in Cycle*	Qualifiers	Comments, usage, and nominal time-in.
		event		
27	Source-On	2 Turns	Variable Rep-Rate.	Controls ion source rep-rate only. Event location remains fixed. Gate width and delay are set in the MPS Master IOC (which generates the Ion Source Gate). Use gate trigger delay = 5046-Beam Width Turns Always occurs on cycle 599 at any rep-rate.
36	Beam-On	2111-Beam Gate Width 1052-2111 Turns	Variable Rep-Rate.	Event controls the beam rep-rate. Can be independent, or made coincident with Source-On, or Warm-LLRF. Event will be inhibited if the Kicker-Charge event did not fire on the previous cycle. Use gate trigger delay= 2937 Turns. Always occurs on cycle 599 at any rep-rate when beam is on.
37	Beam-Ref	2109 – Beam Gate Width ----- 1050-2109 Turns	60 Hz.	Event controls where the start of beam occurs. This event enables “in time” source pulse selection in MPS trigger control. Use gate trigger delay=2939 Turns.
38	End-Inject	5048 Turns	----	60 Hz. Occurs 2 cycles before Extract. This event signals the fixed end point for all beam and RF gates.
39	Extract	5050 Turns	60 Hz.	Defines the end of the time-critical section of the timeline. Re-enables soft event delivery.
40	Kicker-Charge	5062 + storage turns ----- 12 Turns after Extract	Selectable master rates menu.	Selectable Rep-Rate from RF rates. Instructs the extraction kickers to begin charging. Must occur at least 13 milliseconds before extraction. Consequently, this gate must fire after the Extract event of the <u>previous</u> cycle. The Beam-On event will be inhibited if this event did not fire on the previous cycle.
41	Diag Laser Trigger	Beam-On + Chopper Delay + Chopper Ramp Up + 8 Turns	30 Hz. (Adjustable)	This triggers the laser. Sync to beam cycles. Use gate trigger delay = 2929 Turns to select the first full minipulse
43	RTDL-Xmit	After turn 5150 + number of stored turns	60 Hz.	Triggered after all the RTDL frames for the next cycle have been loaded. This is effect a soft event. The hardware ensures that this event has been sent before the RTDL starts transmitting.
44	RTDL-Valid	After all RTDL frames have been sent.	60 Hz	.Signals the end of the RTDL transmission. Implemented in the hardware.
45	Diag-Demand	Beam-On + Chopper Delay + 8 Turns	Manual request. Always occurs on same cycles as Diag-Slow event.	Single-Shot. Always occurs on same cycle as Diag-Slow and Diag-Fast events. Use gate trigger delay = 2931 Turns to align gate to first minipulse

SNS Event Num	Event Name	Position in Cycle*	Qualifiers	Comments, usage, and nominal time-in.
46	Diag-Slow	Beam-On + Chopper Delay + 6 Turns	1 Hz. (Adjustable) Always occurs on same cycle as Diag-Fast event.	Based on the corresponding RF event mask and Beam-On. Use gate trigger delay = 2933 Turns to align gate to first minipulse.
47	Diag-Fast	Beam-On + Chopper Delay + 4 Turns	(Adjustable) Always occurs on same cycle as Beam-On event.	Based on the corresponding RF event mask and Beam-On. Use gate trigger delay = 2935 Turns to align gate to first mini-pulse.
48	Diag-No-Beam	3 Turns after Extract	(Same rate as Diag-Fast)	Not necessarily coincident with anything. This is a nominal 6 Hz. event that can be used to test diagnostic triggers when there is no beam. Trigger and delay put the gate outside the beam acceleration/accumulation/extraction part of the machine cycle, so you are guaranteed not to have any beam. Use gate trigger delay = 2931Turns
49	Diag Laser	Beam-On + Chopper Delay + Chopper Ramp Up + 10 Turns	Up to 10 Hz (adjustable) Occurs only with Beam-On and Diag-Laser-Trigger events.	This is intended to trigger diagnostic laser acquisition. Based on the corresponding RF event mask and Beam-On. 20 Hz is not a valid selection. Use gate trigger delay = 2927 Turns to select the first full minipulse
50	15 HZ RF	28 Turns after cycle start	Occurs on cycles with 1 and 5 HZ	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 3,7....599.
52	60 HZ RF	21 Turns after cycle start	Occurs on all cycles	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation.
53	30 HZ RF	22 Turns after cycle start	Occurs on alternate cycles	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 1, 3....599.
54	20 HZ RF	23 Turns after cycle start	Occurs every 3 rd cycle	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 2, 5....599
55	10 HZ RF	24 Turns after cycle start	Occurs every 6 th cycle with 20 and 30 HZ events	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 5, 11....599.
56	5 HZ RF	25 Turns after cycle start	Occurs every 12 th cycle with 30,20,15, and 10 HZ events	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 19, 39....599
57	2 HZ RF	26 Turns after cycle start	Occurs every 30 th cycle with 30,20, and 10 Hz events	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 29, 59....599
58	1 HZ RF	27 Turns after cycle start	Occurs with all other RF events every 60 th cycle.	Event used to trigger HPRF, LLRF, and Injection Kickers as needed by operation. Present on cycle 59,119....599.

SNS Event Num	Event Name	Position in Cycle*	Qualifiers	Comments, usage, and nominal time-in.
59	Diag RTBT Slow	2129 Turns + stored turns	Trigger after Turns + stored turns	On cycles when Diag-Slow is triggered
60	Diag RTBT Fast	2127 Turns + stored turns		On cycles when Diag-Fast is triggered Trigger after 2923 Turns + stored turns
61	Diag RTBT	2125 Turns + stored turns	Same as Beam-On	Triggers RTBT and target instruments. Use gate trigger delay = 2925 Turns + stored turns.
63	Pre-Pulse	8 Turns before cycle start	60 Hz.	Just before cycle start to clean up last cycle faults.

Table 4
Hardware Event Definitions

4.1.2 The Beam-Related Event Cluster

The beam related event cluster is managed by the cycle end sequencer. These events have to be positioned and triggered depending on timing settings. The AIP timing master will move the events around in the event table. The sequencer will also have to provide the connection between the Beam-On switch and the hardware as well as work to set the beam and diagnostic event rep-rates.

The “Beam/Diagnostic Event Cluster” is worth mentioning at this point. This is a cluster of five events that are locked to each other and travel together as a fixed group. The cluster begins with the “Beam Ref” event. This is the only event that has a variable delay. It runs at 60 Hz and occurs 2939 turns before the start of the actual beam gate. It can be used to produce a diagnostic “fiducial” gate that tells you where beam would turn on if it were scheduled for this cycle.

The “Beam-On” event is triggered 2 turns after the “Beam Ref” event. This is the event that actually generates the “Beam Gate” (i.e. the gate that goes to the LEPT chopper and lets the beam out of the source). This event is subject to the beam rep-rate set in the Timing Master EDM screen. This event is also subject to the “Single-Shot” and “Beam Off” buttons on the Timing Master EDM screen. If beam is inhibited for any reason (e.g. MPS trip, or “Beam Off” selected), the Beam-On event will not be generated. The Beam-On event will also be inhibited if the Kicker-Charge event did not fire on the previous cycle. This provision allows you to change rep-rate patterns in the middle of a super cycle without accelerating beam before the extraction kickers have charged.

The behavior described above, is not the behavior specified in the timing system final design review. The above implementation was done to accommodate “dumb” diagnostics that only had a single trigger and no way of knowing whether or not there was any beam associated with that trigger. Thus, the philosophy was “No Beam, No Trigger”. Once the diagnostics start being able to listen to the event link and the RTDL, the diagnostic trigger event mechanism may need to be reconsidered.

The following figure illustrates the various relationships between which events trigger which other events, and when those events occur.

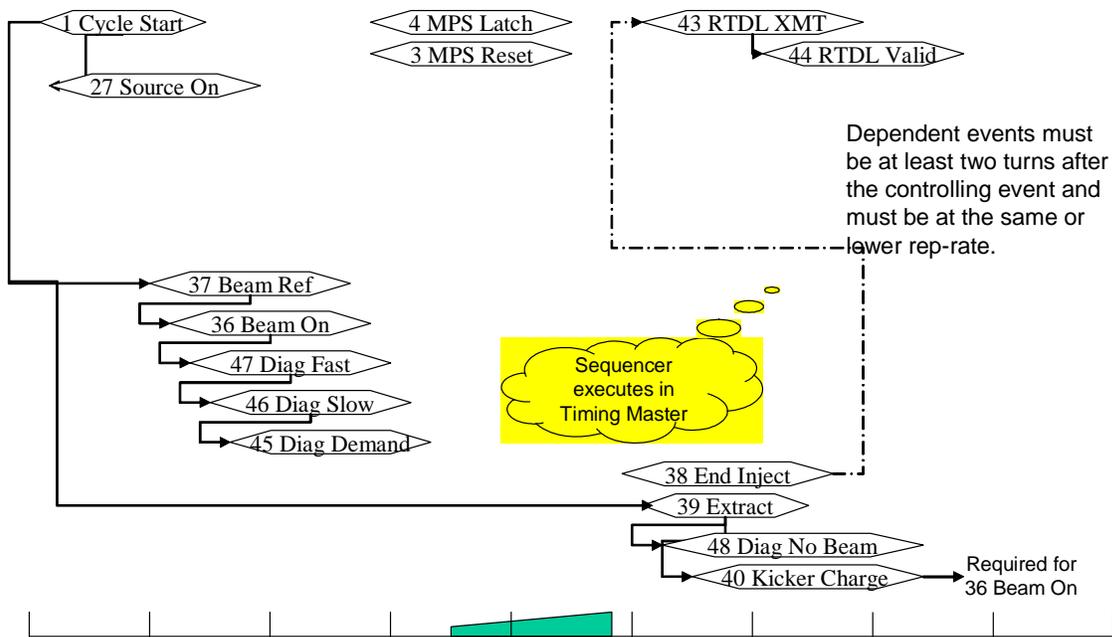


Figure 7
Hardware Event Relationships

4.1.3 Software Events

The following table lists the software events currently in use by the SNS timing system. All software events must ultimately be generated by the timing master IOC (because that's where the event generator lives).

Event Num	Event Name	Comments
232	Dev-network-marker	Sent just prior to RTDL VALID at 60 Hz by the development timing master located in the CLO lab
233	Spare-network-marker	Sent by the spare timing master in the ring service building.
234	Target-beam-veto	Sent at the same time as the RTDL message. Indicates that a veto bit is set in the veto frame.
235	Target-beam-good	Sent at the same time as the RTDL message. Indicates that no veto bit is set in the veto frame.
236	Beam-on-precursor	Sent at the same time as the RTDL message, indicates that the flavor frame is not zero.
240-247	Flavor events	Sent at the same time as the RTDL message, indicates which flavor is being sent. The event sent is 240+flavor.
249	Test-network-marker	Sent by the test network timing master

Event Num	Event Name	Comments
250	MPS Diagnostic Snapshot	Soft event to trigger waveform recording from all diagnostics after a fault. Will occur only on a DIAG_SLOW cycle and if either a latch or an AR fault occurs. Not currently implemented.
251	Compute-Rep-Rate	Signals that a new rep-rate pattern computation has begun in the timing master IOC. (Obsolete)
252	New-Rep-Rate	Signals that new rep-rate patterns have been set in the timing master IOC. (Obsolete)
253	MPS-Error-Reset	Reset error counters on all MPS chassis.
254	Util-Error-Reset	Reset error counters on all SNS Utility Modules.

Table 5
Software Event Definitions

5 REAL TIME DATA LINK

The SNS Real Time Data Link (RTDL) is modeled after the RHIC Real Time Data Link . It is a 10 Mhz, bi-phase mark encoded serial link.

An *RTDL Frame* consists of 51 bits. Its payload is a 24 bit data word representing an SNS accelerator parameter pertaining to the upcoming machine cycle. The rest of the RTDL frame consists of a start bit, an 8-bit frame number, an 8-bit frame CRC plus 10 '1's following the frame. (6) The format of an RTDL frame is:

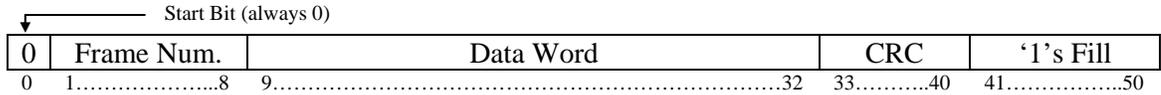


Figure 8
RTDL Frame Format

Before the start of each machine cycle, a series of frames about the upcoming cycle is transmitted on the RTDL. Up to 255 RTDL frames are possible and about 26 are currently implemented. The last RTDL frame of the transmission is always frame 255 and contains a 24-bit "Message CRC" which is computed by the timing master IOC.

5.1 RTDL FRAME DEFINITIONS

These are the frames that currently constitute the RTDL transmission.

5.1.1 Time Stamp (Frames 1-3)

The first three RTDL frames contain the 64-bit EPICS timestamp to be applied at the next Cycle-Start event. The 64-bit timestamp is divided into two 32-bit long words as follows:

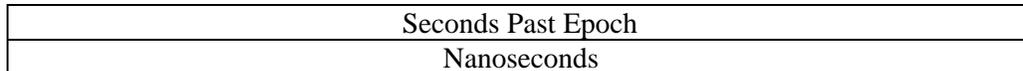


Figure 9
EPICS Timestamp Format

The 64-bit EPICS timestamp is transmitted on three RTDL frames. The RTDL timestamp encoding is shown below:

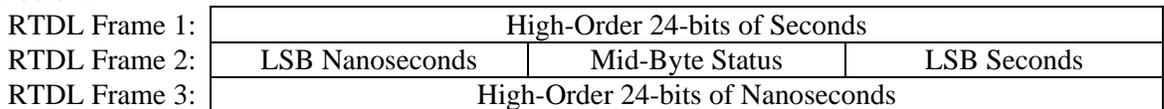


Figure 10
RTDL Timestamp Format

The timing master inserts several bits of status information into the middle byte of frame 2.

Bit	Definition
0x1	This is a timing master ID bit. The timing master on this network is the Development timing master.
0x2	Indicates that the ion source switch on the master screen is on and that the ion source events are present on the event link.
0x4	Indicates that the beam permissive state is on and that beam and diagnostic events may be present on the event link.

0x8	Indicates that the RF switch on the master screen is on and that RF events are present on the event link.
0x10	This is a timing master ID bit. The timing master on this network is the Spare timing master.
0x20	Indicates that the timing master is in single shot mode.
0x40	Indicates that a single beam pulse or a burst of pulses is pending or in progress.
0x80	This is a timing master ID bit. The timing master on this network is the Test timing master.

5.1.2 Ring Revolution Period (Frame 4)

RTDL frame 4 contains the ring revolution period in picoseconds. This value is useful for converting timing system units into real time. The ring revolution period is acquired either by measuring the frequency of the event link clock or by reading the frequency setting of the ring clock. The accuracy of this parameter must be better than 1 part in 10^6 .

5.1.3 MPS Mode (Frame 5)

The MPS Mode number is broadcast on RTDL frame 5. This frame is primarily of interest to the MPS system. MPS uses the value in this frame to set the masks of which conditions to ignore and which conditions are relevant for the current operating mode.

The MPS system wants to make sure that the mode information coming from the RTDL is correct and does not contain any transmission errors. For this reason, the mode information is encoded three ways in RTDL frame 5. The low order byte contains the unadulterated mode number. The middle byte contains the mode times two plus one. The high order byte contains the “ones compliment” of the mode. This encoding is show below in Figure 11:

One’s Compliment of Mode	$(\text{Mode} \times 2) + 1$	Mode
--------------------------	------------------------------	------

Figure 11
RTDL Mode Frame Format

The encoding of the beam mode and the machine mode is done in the MPS PLC based on the MPS key switches in the control room. The following table illustrates the coding of the modes in the mode frame. The darkened row and column are not used.

	Source	MEBT BS	CCL BS	LDump	IDump	Ring	EDump	Target
OFF	0							
Standby	1							
Test	2							
10 usec		3	8	13	18	23	28	33
50 usec		4	9	14	19	24	29	34
100 usec		5	10	15	20	25	30	35
1 msec		6	11	16	21	26	31	36
Full Power		7	12	17	22	27	32	37

Figure 12
RTDL Mode Frame Encoding

5.1.4 60 Hz. Phase Error (Frame 6)

This frame represents the distance in nanoseconds between the actual 60 Hz line frequency zero crossing and the “smoothed” zero crossing signal from the master reference generator (see section 6).

This signed value is read from the master reference generator every cycle. The time is measured from the extract event to the line crossing pulse. A positive value in Frame 6 means that the line crossing occurs after the extract event. This frame can be used to correlate line sync related phenomena and to improve background subtraction algorithms.

5.1.5 Beam Width (Frame 7)

This frame represents the width in turns of the current beam flavor. It is used to set the Beam-On gate in the trigger control chassis. If beam is scheduled to be on in the next cycle, this frame reflects the width of the particular flavor of beam scheduled for that cycle. If beam is off, this frame reflects the width of the primary flavor.

5.1.6 Line Frequency (Frame 8)

The frequency, in units of 100 micro Hz, of the 60 Hz events on the event link.

5.1.7 IOC Reset (Frame 15)

Each SNS Utility Module and timing receiver has a 24-bit “IOC Reset” address associated with it. This address is set by jumpers on the utility module board and hard coded to the serial number of the timing receiver. When the module firmware detects a value in RTDL frame 15 that matches its own IOC Reset address, it will assert the VME SYSRESET line, causing the IOC to reboot. This feature allows an IOC to be rebooted remotely, even if its software is completely hung up. It will not, of course, allow for the remote reboot of a powered-off IOC.

IOC reboots are requested by writing the reset address of the specified IOC to a special record maintained by the RTDL encoder driver. A second PV will accept several forms of the IOC name by reading a file based on the name and using the code read the file.

RTDL frame 15 is broadcast every cycle (as are all the RTDL frames defined here). Normally its value will be 0, indicating that no IOCs should be rebooted. The control software should reset the reset frame on the next cycle after sending a reset code.

The data files supporting this are in `/ade/epics/iocCommon/ics-tim-ioc/var/rebooter` on the accelerator server. The `learnRebootCodes.pl` script in `Support/timing/SupportApp/rebooter` should be executed after a utility card is moved.

5.1.8 Pulse Flavor (Frame 17)

RTDL frame 17 contains the “pulse flavor” of the next cycle. Pulse flavors are described in section 1.3.4.

5.1.9 Last Cycle Veto (Frame 24)

RTDL Frame 24 describes possible veto conditions to the experimental stations. Unlike the other RTDL frames, this frame pertains to the *previous* cycle rather than the *next* cycle. Each bit in this frame describes a reason why one or more experimental stations may want to veto data from the previous cycle. These bits are defined in the `snsTiming.h` file. 12 veto bits are currently defined. Their values and symbolic names are listed below:

Bit	Name	Description
0	VETO_NO_BEAM	No beam was delivered on the previous pulse.
1	VETO_NOT_TARGET_1	Beam was delivered to target 2 (not to target 1)
2	VETO_NOT_TARGET_2	Beam was delivered to target 1 (not to target 2)

3	VETO_DIAGNOSTIC_PULSE	Beam was a “reduced intensity” diagnostic pulse
4	VETO_PHYSICS_PULSE_1	Beam was one of the special physics study pulses
5	VETO_PHYSICS_PULSE_2	Beam was one of the special physics study pulses (the other type)
6	VETO_MPS_AUTO_RESET	Beam was interrupted by an “Auto Reset” MPS trip (fast protect)
7	VETO_MPS_FAULT	Beam was interrupted or not delivered because of a “Latched” MPS trip
8	VETO_EVENT_LINK_ERROR	Timing system detected corruption on the event link
9	VETO_RING_RF_SYNC	Timing system has lost synch with the Ring RF signal
10	VETO_RING_RF_FREQ	Measured ring RF frequency is outside acceptable range
11	VETO_60_HZ_ERROR	60 Hz line phase error is out of tolerance

Figure 13
Veto Bit Definitions

5.1.10 Cycle Number (Frame 25)

The cycle number within the *Super Cycle* is broadcast on RTDL frame 25. An SNS super cycle is 600 cycles long. Therefore, this number will always be between 0 and 599. This information is useful for client IOCs that implement local rep-rates.

5.1.11 Master Rep-Rate (frame 26)

This is the setting of the timing system master rep-rate. This will be one of the frequencies available on the RF events. The timing master schedules Beam-On events based on super cycle slots available at the master rate. Since Beam-On events are scheduled this way the master rep-rate is the maximum available Beam-On rep-rate. RF systems that are used to accelerate beam must operate at the master rep-rate or an even multiple of the master rep-rate. The master rep-rate frame is a contract between the beam and the accelerator. Accelerator systems, kicker magnets and RF can run at any frequency that covers all of the master rep-rate slots. The beam is then only scheduled in master rep-rate slots.

5.1.12 Chopper Control Frames (frame 28 to 32 and 36 to 38)

These frames control the LEBT/MEBT chopper operation on a cycle to cycle basis. When the flavor is a Beam-On flavor these frames define the shape of the chopped beam. The timing master has a set of PVs that define each flavor. Once the flavor of the next cycle is determined the timing master sequencer checks the values for consistency and then copies the set of parameters from the flavor setting PVs into these frames. The flavor may have to be modified by the timing master sequencer if the flavor PVs request more mini-pulses than are available in the macro-pulse defined by the beam width/beam gate.

The timing header file “snsTiming.h” defines macros for stuffing and un-stuffing these frames. In some cases the actual “stuffing” is done in EPICS databases.

Note that the details of these frames may change as driven by the interface requirements of the LEBT/MEBT chopper controller. Documentation for the LEBT/MEBT chopper controller should be consulted when using these definitions.

Frame details as of this writing:

RTDL_CHOPPER_P1_FRAME (28)			
High Bit	Low Bit	Name/MACRO	Description
23	12	Beam Delay RTDL_CHOPPER_BEAM_DELAY	The number of beam off turns generated between the Idle delay and before the first mini-pulse to the linac. During this phase all four LEBT vanes are driven to either +2500 Volts or -2500 Volts with one

			pair of vanes swapping polarity in the middle of each turn.
11	0	Dark Current Suppression RTDL_CHOPPER_RFQ_ON	<p>If the RFQ is on this cycle and the beam is off a 1 in this nibble commands the chopper to chop the beam off pattern for the full duration of the RFQ LLRF pulse to suppress dark current or for chopper testing.</p> <p>The Beam Delay parameter will also be extended to turn 5048.</p>

RTDL_CHOPPER_P2_FRAME (29)

High Bit	Low Bit	Name/MACRO	Description
23	12	Ramp Down RTDL_CHOPPER_RAMP_UP	The number of beam turns for the ramp down phase. The pulse width changes from PW On to PW Stop over this number of turns.
11	0	Ramp Up RTDL_CHOPPER_RAMP_UP	The number of beam turns for the ramp up phase. The pulse width changes from PW Start to PW On over this number of turns.

RTDL_CHOPPER_P3_FRAME (30)

High Bit	Low Bit	Name	Description
23	12	Beam Off RTDL_CHOPPER_BEAM_OFF	The number of turns of chopped beam off after the Idle Phase and before the ramp up.
11	0	Beam On RTDL_CHOPPER_BEAM_ON	The number of turns chopped with the output width equal to PW On

RTDL_CHOPPER_P4_FRAME (31)

High Bit	Low Bit	Name/MACRO	Description
17	12	PW Stop RTDL_CHOPPER_PW_STOP	The width of the last ramp down pulse at the end of the ramp down in units of 1/64 turns.
11	6	PW On RTDL_CHOPPER_PW_ON	The width of each Beam-On pulse in units of 1/64 turns.
5	0	PW Start RTDL_CHOPPER_PW_START	The width of the first pulse at the start of the up ramp in units of 1/64 turns.

RTDL_CHOPPER_P5_FRAME (32)

High Bit	Low Bit	Name/MACRO	Description
23	12	Num Off	Number of beam off between

		RTDL_CHOPPER_NUM_OFF	pulse groups.
11	6	Num On RTDL_CHOPPER_NUM_ON	Number of turns per group from the first beam pulse
5	0	Phase Delay RTDL_CHOPPER_PHASE_DELAY	Fine delay in 1/64 turn added to the entire cycle.
RTDL_CHOPPER_P6_FRAME (36)			
High Bit	Low Bit	Name/MACRO	Description
23	0	Idle RTDL_CHOPPER_IDLE	Number of turns from Cycle Start to the first ramp up or beam pulse
RTDL_CHOPPER_P7_FRAME (37)			
High Bit	Low Bit	Name/PV signal Name/VHDL Signal	Description
23	12	MEBT1/ MimicBeamOnRisingEdgeDelayTicks	Phase delay of the rising edge of the Mimic Beam On pulse used by the chumps beam in gap fault comparator. In units of 1/64 turns. Must not be more than 40.
11	0	MEBT 1/ lebtToMebtPhaseTicks	Phase delay LEBT to MEBT chopper output in 1/64 turns.
RTDL_CHOPPER_P8_FRAME (38)			
High Bit	Low Bit	Name/PV signal Name/VHDL Signal	Description
23	12	MEBT 2/ MimicBeamOnFallingEdgeDelayTicks	Phase delay of the falling edge of the Mimic Beam On pulse used by the chumps beam in gap fault comparator. In 1/64 turn.
11	7	MEBT 2/ mebtChopperCycleEnLow	Disables the chopper when set.
6	0	MEBT 2/ mebtPwOnWidthAdjustment	Adjusts the width of the MEBT chopper pulse in 1/64 turn.

Figure 14
Chopper Related RTDL Frames

5.1.13 Last Peak Beam Current (Frame 33)

This frame is not currently used.

5.1.14 Last Integrated Beam Current (Frame 34)

This frame is not currently used.

5.1.15 Last Integrated Beam On Target (Frame 35)

Charge from RTBT_Diag:BCM25 in units of 10 nC. The Data from the BCM is transmitted in real time over a fiber optic distributed memory network using VMIVME-5565 cards in both the timing master and the BCM.

The data is passed in a memory structure located at the base address of the shared memory block. There are 40 bytes each allocated for node ID strings and 4 bytes in big-endian order for the charge reading from the BCM. A very minimal protocol has been implemented to ensure proper data transmission. The timing master initializes its memory card in the first ID string block with the string "TimingMaster" and the BCM initializes the second ID block with "BeamOnTarget".

When the timing master builds the RTDL, frame 35 is the last item to be copied from the memory card to the RTDL hardware. This is a simple copy, no data checking or scaling is done in the timing master. Once the data is copied the timing master writes back a zero to the data location. On the BCM side a '1' is written to the memory location at the time of a beam trigger, and the value is written as soon as the data conversion is complete. The '1' value is excluded from the possible valid data values that the BCM can send. The BCM reads back the value on the RTDL and confirms that the sent data and the received data are the same. If a '1' shows up in the RTDL in frame 35 the BCM knows that the timing master missed the data. The accelerator cycle times this operation so no other handshaking between the master and the BCM is needed.

5.1.16 BLM Reprate (Frame 40)

Bits in this frame represent RF On cycles in local areas of the machine. These are used so that beam loss monitors can use beam off cycles for back ground subtraction only when RF generated X-Rays are present. The implementation of this frame depends on the requirements of the BLMs. This requirement must be reviewed prior to completing the implementation of the AIP timing master,

5.1.17 Stored turns Frame (Frame 41)

Stored turns are the number of turns of delay after Text that is added to the extraction kicker timing. Stored turns are used for physics studies and machine tuning. Neutron instruments and all RTBT beam instruments must add the stored turns to their trigger delay in order to maintain correct timing.

5.1.18 Message CRC (Frame 255)

The last RTDL frame transmitted is always frame 255. This frame contains a 24-bit CRC of the entire RTDL transmission. The CRC value is computed by the timing master. Details of the CRC computation can be found in [7].

For details about CRC generation see SNS-NOTE-CNTRL-45. (6)

6 MASTER REFERENCE GENERATOR

The purpose of the Master Reference Generator is to monitor the AC line frequency and produce a stable reference pulse that does not change faster than the T0 neutron choppers can keep up with, and that does not stray too far from the actual zero crossing. The current operating requirements are that the reference pulse should not change any faster than 1 mHz per second, and that it not stray more than ± 500 microseconds from the actual zero crossing. Operating experience and operating conditions may dictate that the reference generator change from time to time so the major parameters should be adjustable.

Several facts were discovered at the beginning of the development of the line sync system in the timing master. A DSP board offered by Brookhaven implemented a second order PLL. A lot of effort was committed by the BNL developer to determining the exact zero crossing point at the input. Investigating the behavior of a 2 pole PLL revealed that a PLL by itself cannot control the slew rate. The damping factor should be very close to critical in order to control overshoot. A 2 pole PLL cannot be tuned to be over damped; close to critical damping is the best that can be done. When the PLL is tuned to be slow enough for the choppers it cannot maintain a reasonable phase error. It was also found that when a hard slew rate limit is imposed on a simple PLL it enters chaotic behavior if a transient that exceeds ± 180 degrees phase error comes along. So it just isn't possible to employ a second order PLL that tries to maintain zero frequency error and zero phase error based only on a phase measurement to close the loop.

The second thing that was discovered was that the cycle to cycle phase jitter is important at the output and not the input, so the BNL effort to find a precise line crossing point was not needed. Instead we use a pulse generator on line trigger plugged into a clean power source. This clean power source is separate from all other loads in the timing master cabinet and is not provided from UPS backed up power. Configuration management must be in place to control how power is provided to this component. The third thing discovered was that any noise present in the signal at the point where the slew rate limit is being enforced reduces the effective slew rate below the actual limit setpoint. The slew rate limit is very close to the performance of the power utility and every bit of the slew rate is needed to be able to track the line. Therefore the limiting algorithm must be accurate and not be degraded by noise.

Note that at 60 Hz a 0.001 Hz change in frequency results in a change in period of 277 ns. On a per cycle basis a 0.001/60 Hz change in frequency is a change in the cycle period of 4.6 ns per cycle. The resolution of a $16 \times \text{Frev}$ clock is about 60 ns so actual changes to the cycle period will be about 12 cycles apart when the output frequency is under control of the slew rate limiter. The period may change back and forth on every cycle when the loop is not slew rate limiting.

The output of the master reference generator is a signal that will trigger the "Cycle-Start" event. Changes in the ring RF frequency will cause the machine cycle to increase or decrease in actual time, even though the "turn time" remains at constant 5,050. The difference between the ring revolution time at the highest and lowest design frequencies is about 62.2 nanoseconds per turn. Since there are 5,050 turns between Cycle-Start and Extract, the variation in the machine cycle length can be up to 314 microseconds.

The actual algorithm operates in a blend of two modes. Most of the time the input frequency is measured and the output frequency is set to match the input frequency as long as the slew rate requirement can be met. The phase stays close to zero when the slew rate limit is not being approached. When the slew rate limit mode is entered the phase increases exponentially based on the difference in the input and output frequencies. A key point here is that the phase is the frequency difference integrated over time so it is necessary to track the input frequency as closely as possible in order to best handle the sudden transient when it comes along. The line frequency is essentially random so if there is a frequency error at the start of a sudden transient, then half the time the offset will be advantageous and half the time the wrong initial offset will eventually make the phase error worse. We are not interested in the average phase error, only the amount of time when the phase error exceeds the limit. Maintaining a small frequency error then reduces the times when the phase error exceeds the limits by being close to home when the transient does come.

To re-zero the phase difference a small offset is added to the output frequency such that the phase will be driven toward zero. This is in effect a slow single pole PLL. A dead-band is used in the PLL loop to disable it when the phase is small and to damp out the overshoot after a large transient. A single pole loop, not trying to maintain zero frequency error, exhibits an inverse exponential transient response.

Two measurements are made, phase error and input frequency. The input frequency is calculated by measuring the time of the last 60 cycles at the reference input. The slew rate calculation needs a noise free signal in order to accurately control the slew rate so a tunable low pass filter has been added to the frequency feed forward path.

The output frequency is normally the sum of the filtered input frequency (gain=1) and the PLL offset. The process of tuning the algorithm tries to minimize the output phase error against the variation of the power grid. Note that the PLL loop is single pole so it is unconditionally stable. With frequency feed-forward there is no need for a PLL to try to maintain a low frequency error. When slew rate limiting is in effect the loop acts like it has two poles and can go unstable. This condition is handled by gradually allowing the slew rate to increase while the phase is out of its allowed range. Eventually the loop will reach a linear condition and will then be able to recover. A lab setup with a signal generator capable of producing small frequency steps while maintaining phase coherency is the best way to experiment with the algorithm.

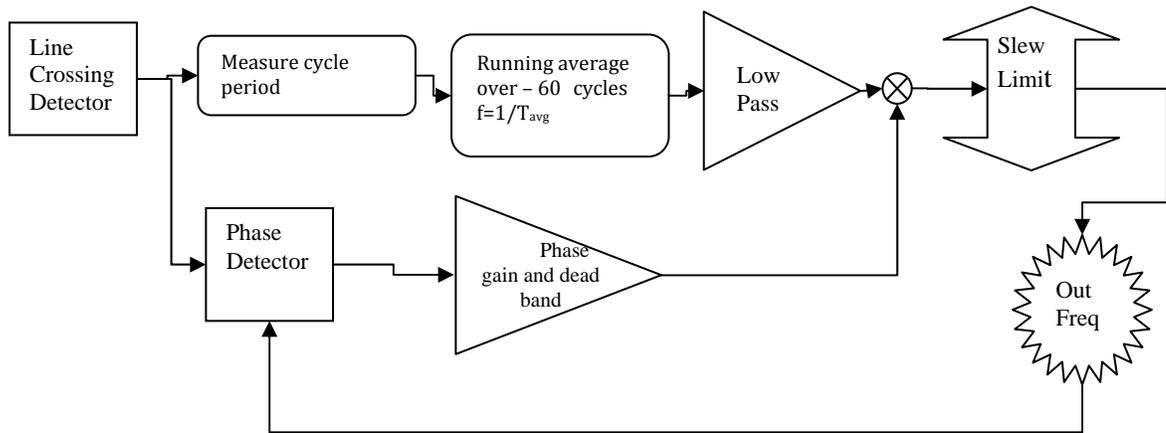


Figure 15 Line Sync Block Diagram

The rollover point where the phase jumps from -180 degrees to +180 degrees has to be handled carefully. Adjustments are made in the frequency low pass filters to make this rollover point act continuous.

There are improvements that could be made to the algorithm. First, adaptive filtering methods could be used to optimize the tuning to the ever changing dynamics of the power utility. Second, the slew rate limiting routine could be changed to match the dynamics of the slowest chopper's internal position control loops. In both cases advanced control theory offers improved phase error performance for the timing system. A cost to benefit assessment would only justify pursuing either of these ideas if the phase error causes significant veto rates for the neutron data acquisition systems.

7 ADDITIONAL HARDWARE IN THE TIMING MASTER

7.1 GLOBAL POSITIONING SYSTEM INTERFACE

The timing master IOC uses a global positioning system (GPS) to get the time of day that it broadcasts on the RTDL. (1) The GPS receiver also serves as the site NTP server.

7.1.1 Software

The GPS board has an EPICS driver that reads the timestamp from the registers on the card, validates the time, and provides the time to the cycle end sequencer in an EPICS time stamp format.

7.1.2 Hardware

The GPS receiver is a TrueTime XL-DC Time and Frequency Receiver. This is a 2U chassis that is mounted in the timing master rack. The interface to the timing master IOC is a TrueTime VME-SG2 VME card. The card interfaces to the receiver via the receiver's external IRIG-B output. The VME-SG2 card can also function as a stand-alone time source.

The VME-SG2 card has two sets of "Freeze" registers that can be used to capture and read the current time. The first set of "Freeze" registers can be triggered via a VME read operation. This set can be used by EPICS device-support routines to capture and display the current time in a string record. The second set of "Freeze" registers is triggered by an external TTL signal. A timing gate synchronized with the "Cycle-Start" event is used to trigger these registers. The captured value is then used to compute the time of the next "Cycle-Start" event and distribute this time on the "Real Time Data Link" (RTDL).

The GPS antenna and the IRIG-B signal to the VME-SG2 module connect to two BNC connectors in the rear panel of the GPS receiver module. These connections are show below in Figure 16.

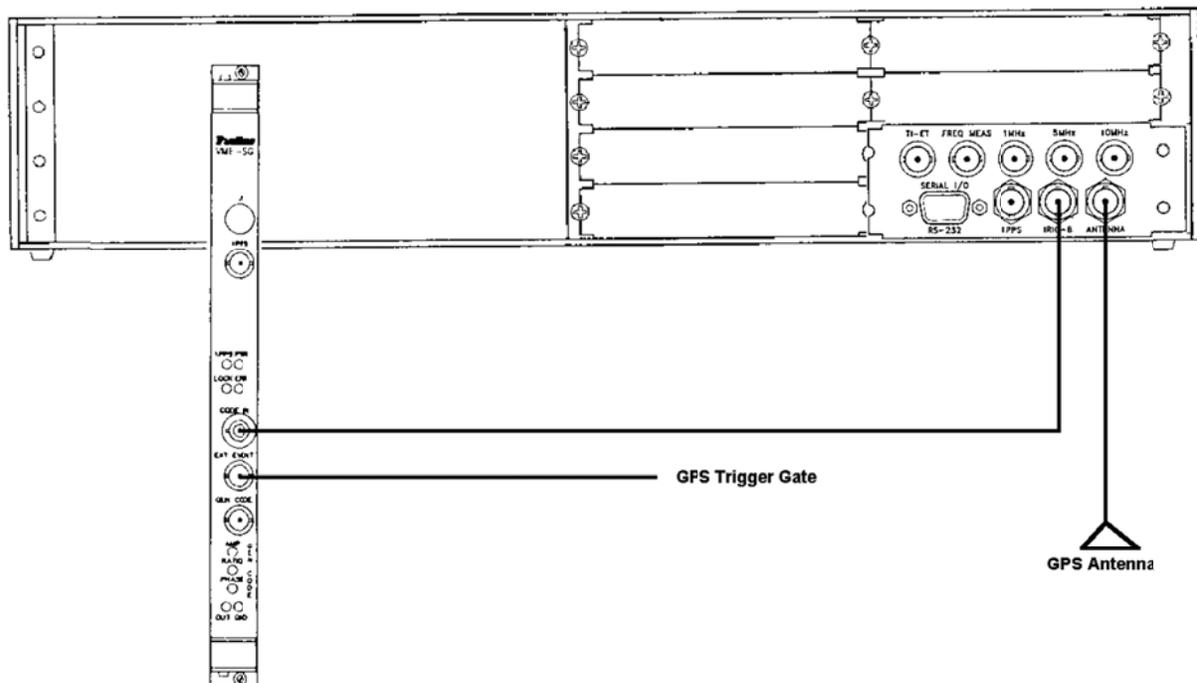


Figure 16
Connections to GPS Receiver and VME-SG2

The GPS antenna connects to the “ANTENNA” BNC port on the back of the GPS receiver. The “CODE-IN” port on the front panel of the VME-SG2 connects to the “IRIG-B” BNC port on the back of the GPS receiver.

The “EXT EVENT” port on the VME-SG2 connects to the “GPS Trigger” gate. The “GPS Trigger” gate occurs at the same time as the gate that triggers the Cycle-Start event. When this gate fires, the current GPS time is copied into an internal register in the VME-SG2. The timing master IOC reads this value, adds 16.667 milliseconds to it, and broadcasts it on the RTDL as the timestamp for the next Cycle-Start event.

NOTE:

There are two different GPS receiver models in use. The older XL-DC provides 5 Volts to the antenna to power the preamplifier in the antenna. The Xli receiver provides 12 Volts on the antenna connector and therefore cannot be connected directly to a 5 Volt antenna. When the Xli receiver is in use the antenna must be powered using an external 5 Volt power supply and isolated from the receiver power using a DC inserter block.

7.2 REFLECTIVE MEMORY BOARD

A VMIVME 5565 Reflective Memory Board is situated in the timing master IOC. This board shares memory with RTBT_Diag:BCM01. Beam on target data is written into a location in the card at the end of a machine cycle and copied to the RTDL by the timing master. The remote IOC must be allowed time to process the beam on target data so reading of the data from the card at the master must occur some 600us after Text and allow for any additional stored turns.

A mini-driver in the timing master main loop provides for managing the data from the BCM and detecting transmission errors.

The card in the timing master is jumpered as “NODE 1” and the card in the BCM is jumpered as “NODE 2”.

7.3 LINE SYNC PULSE GENERATOR

The line sync pulse generator interfaces the raw AC line to the timing master hardware at a logic level. The line crossing is detected and after an adjustable delay a pulse is sent to the event encoder.

7.4 EVENT LINK CLOCK

This is a signal generator that generates a sine wave at 32 times the revolution frequency. The frequency is adjustable from the control room based on a calculation that converts a beam setting to frequency. The output is split and is used in the ring low level RF system so the amplitude must be compatible with both users.

The event link clock is adjusted in small steps by the software driver so that PLLs in devices in the timing system can follow.

7.5 MPS INTERFACE

Two hard links come to the timing master from the trigger control chassis. These indicate when MPS Fast Protect Latch fault and MPS Auto Reset faults occur. As noted above, events should be generated

when either fault occurs. In addition, the presence of either fault should suspend beam scheduling for the duration of the fault. An MPS FPL fault must trip the Beam-On switch so that manual intervention is required to restart the beam.

8 BIBLIOGRAPHY

1. **B.Oerter, R.C.Sibley.** “*System Requirements Document for Timing System*”, SNS 109020000-SR0001-R01.
2. **Multiple.** *SNS Parameters List*, SNS-100000000-PL0001-R10.
3. **E.Bjorklund.** “*The Theory of Rep-Rate Pattern Generation in the SNS Timing System*”, SNS-NOTE-CNTRL-99.
4. —. “*A Metric for Measuring the Evenness of Timing System Rep-Rate Patterns*”, SNS-NOTE-CNTRL-100.
5. Bresenham's line algorithm. *Wikipedia*. [Online]
http://en.wikipedia.org/wiki/Bresenham's_line_algorithm.
6. **E.Bjorklund.** “*Cyclic Redundancy Check Codes Used in the SNS Real-Time Data Link*” SNS-NOTE-CNTRL-45.
7. **Inc., TrueTime.** “*Manuals for TrueTime XL-DC Time and Frequency Receiver and VME-SG2 GPS Interface*”, TrueTime Inc. 3750 Westwind Blvd., Santa Rosa, CA 95403.

9 CONTACTS

Oak Ridge:

Dave Thompson
Doug Curry

thompsondh@ornl.gov
curryde@ornl.gov

10 INDEX

Event Link, ii, 5, 7, 14, 16
Machine Cycle, ii, 5, 7, 8, 17, 18
The Real Time Data Link
RTDL, 5
