

Synchronization of GlueX Timing Systems

Kei Moriya

Arizona State University

Mark M. Ito

Thomas Jefferson National Accelerator Facility

GlueX Note 2686-v1

March 31, 2015

Abstract

Techniques for synchronizing time measurements from the various timing devices used by GlueX are discussed, in particular how synchronization can be achieved in the absence of an external, global timing signal. Some of these techniques were necessary during the Fall 2014 run, and all of them are useful as consistency checks on timing system operation when an external timing reference is used for synchronization.

Contents

1	Introduction	2
2	Timing Systems	3
3	CAEN TDCs	4
3.1	The Problem	4
3.2	The Solution	4
3.3	The Outlook	6
4	F1TDC TDCs	7
4.1	The Problem	7
4.2	The Solution	8
4.3	The Outlook	10
5	FADC-125	10
6	Conclusions	12
A	Configuring the F1-TDC	13
A.1	Practical Use	13
B	Acknowledgements	14

1 Introduction

In the Fall 2014 run GlueX had four devices that functioned as TDCs: the F1-TDCs, CAEN 1290 TDCs, Flash-250 ADCs, and Flash-125 ADCs. For each of these timing systems, measured time differences between channels within the system behave as expected: these differences are physically meaningful modulo fixed time offsets associated with each channel. These offsets are effectively constant in time, they depend on such things as cable lengths and propagation delays through electronics modules, change rarely if at all, and can be measured and applied as a per channel correction with only the physical time difference remaining after application. However time differences measured between channels in different systems are subject to non-fixed offsets, depending on the details of the device and they way they are operated.

The standard way to overcome this situation is to put a reference signal into each system where the time of arrival of the signal at each system varies only by a fixed offset, typically by fanning out a single signal and sending a copy to each system, or by doing something equivalent. The reference must be present in each system for every event. In addition at CEBAF the common practice is to use a reference signal with a fixed phase relation to the 1499 MHz RF frequency of the machine. In that case times can in principle be referenced to the arrival of beam on target. Note that the use of the RF clock is not necessary to synchronize timing between the systems, any fanned out signal will suffice, but using the RF has a clear advantage.

GlueX has always planned to use the RF signal to synchronize systems. Unfortunately, during the Fall 2014 run, the system for delivering this signal to the Tagger Hall and to Hall D was not ready. Lacking the RF, the fallback plan was to use a trigger time signal to synchronize the systems. The trigger system operates on a 250 MHz clock system clock, and so triggers only occur on edges separated by 4 ns. These edges have no fixed phase relation to the machine clock so the beam bucket time would have to be determined from detector signals (if at all). But as noted above, the trigger time would be sufficient to serve as a global time reference. All relative times within the various detector systems could be obtained once offsets were calibrated out.

As data taking started, however, none of the trigger time reference signals were installed. At some point part way through the run a trigger reference signal was implemented for the F1-TDCs. We did not succeed in implementing a reference signal into the CAEN TDCs (used only by the Time-of-Flight, ironically) at any time during the run.

Fortunately, referencing the F1-TDCs, the CAEN TDCs, and the Flash-250 ADCs to the trigger time can be achieved with no external trigger reference signal at all. The rest of this note will describe the techniques for doing so; a different technique must be applied for each system due to the vastly different designs of the devices. The methods for the TDC systems exploit the presence of a register in each Trigger Interface (TI) board that counts the number of 250 MHz clock cycles starting from

Timing System	Freq. (MHz)	Period (ns)	Detector Systems
CAEN 1290 TDC	41.67	24	TOF
F1-TDC	31.25	32	BCAL FDC PSC ST TAGH TAGM
Flash-250 ADC	250	4	BCAL FCAL PS PSC ST TAGH TAGM TOF
Flash-125 ADC	125	8	CDC FDC

Table 1: GlueX Timing Systems

the beginning of the run. These techniques were demonstrated during the run and turned out to be necessary to bring the Time-of-Flight TDCs into registration with the rest of the detector.

2 Timing Systems

GlueX uses two types of TDCs, the CAEN 1290A and the JLab F1-TDC. There are also two types of Flash ADC boards referred to as the FADC-250 and FADC-125. The FADCs act as timing devices, albeit with lower resolution than the proper TDCs. These devices are listed in Table 1. Each system has its own system clock, of a different frequency for each system. These are also shown in the table.

All of these clocks are derived from a single 250 MHz oscillator on the Trigger Supervisor board (TS), the “global system clock”. As such, the phases of each clock are stable with respect to each other. The data acquisition (DAQ) and trigger systems run on the global system clock. As mentioned earlier, triggers can only occur on an edge of this clock. So in general there is a phase ambiguity for when the event that caused the trigger occurred. This ambiguity is always present and none of the techniques described below address this uncertainty. This is the famous “4-ns uncertainty.” The problem then reduces to keeping all four timing systems in registration relative to one another. Given the special circumstances at the start of the Fall run, we needed to establish this registration in the absence of an external reference signal.

Since the Flash-250s run from the global system clock, as are the DAQ and trigger systems, one way to frame the problem is to ask that all timing systems are in registration with the FADC-250s. There should be an event-independent offset between times measured in any timing system and the FADC-250s for two signals that occur “at the same time.” We will discuss the registration techniques from this point of view.

One key feature of the DAQ system is that each Trigger Interface (TI) board (one per crate) keeps a count of the 250 MHz clock, where the counter is reset at the beginning of each run by the “synch” signal. The value of this counter at trigger time




Figure 1: Global system clock vs. CAEN TDC clock. The global system clock is represented by the TI counter and is shown modulo 6. The trigger can come on any edge of the global system clock.

can be read by the Read-Out Controller (ROC) for each event and should agree for all crates. The TI can be configured to use either 32 bits or 48 bits for this counter. We ran with the TI set to use 48-bits. This gives a roll-over period of about 13 days, longer than any run we took or are likely to take. Use of a 32-bit counters would have given a roll-over period of only 17 seconds and would make the counter unsuitable for measuring time from the synch at the beginning of the run.

3 CAEN TDCs

The CAEN 1290A TDC[1] is a traditional pipeline TDC with a least count of 25 ps. It is driven externally by a system clock with a 24 ns period. All times for each event are referenced to one of these clock edges, *i. e.*, the common TDC stop comes on this edge. The design was originally intended for use in a collider where the bunch crossing time would provide a system clock.

3.1 The Problem

For GlueX, the trigger comes on an edge of a 4 ns clock, but after that the TDC will not be stopped until the next edge of the 24 ns clock. This introduces a random component of the delay between the trigger and the TDC stop, event-to-event, of $n \times (4 \text{ ns})$ where n is an integer between 0 and 5 inclusive. We have been referring to this as the “six-fold ambiguity” of the CAEN TDCs. The situation is illustrated in Fig. 1.

3.2 The Solution

This ambiguity can be resolved by looking at the TI counter for each event, call it n_{TI} . Since it is counting 4-ns periods, the value of the counter, modulo 6, will choose the correct delay (strictly, the correct n), up to cyclic permutations. Which permutation to use depends on the relative phase of the 4 ns system clock and the 24 ns CAEN TDC clock at synch time at the beginning of the run. After synch time, the permutation




Figure 2: CAEN TDC hit timings divided into $n_{\text{TI},6}$. The timing can be seen to shift for increasing values of $n_{\text{TI},6}$, and takes minimum shift values for $n_{\text{TI},6} = 2$. The figure is for run 1802 of the fall 2014 run, summed over all files.

remains fixed until the next synch, in other words it is stable throughout the run. The two-dimensional histogram in Fig. 2 shows the distribution of raw TDC times from all channels (no reference signal subtracted explicitly) along the x -axis, and the value of n_{TI} modulo 6 ($n_{\text{TI},6}$), along the y -axis. The one-dimensional histogram on top shows projections onto the x -axis separately for each value of $n_{\text{TI},6}$. The particular cyclic permutation of the random delay (5,6,0,1,2,3,4) for this run is apparent in the 2-D plot and the 4-ns shift of each distribution is clear in the 1-D plot. The width of each individual distribution is due to summing all channels, uncalibrated and unreferenced, into a single histogram. In summary, knowing the permutation (for the run) and $n_{\text{TI},6}$ (for the event) determines n , and the correct time offset can be applied.

It turns out that for the Fall run, the phase of the synch with respect to the 24-ns clock was completely uncontrolled (*i. e.*, random). This had the result of randomizing the correct permutation from run to run. This situation meant that the permutation had to be calibrated on a run-by-run basis. It turns out that because of the details of how the global system clock is stepped down to a 24-ns period, only three of the six possible permutations are allowed. Nonetheless which of these three is the correct one varies randomly from run-to-run and the need to calibrate remains.

The calibration was performed by looking at the time difference between TDC times and ADC times for pulses **from a single detector element**. These signals only differ by a fixed delay: the difference in the “cable delay” after the split into discriminator input cable (for the TDC) and FADC input cable. The only source




Figure 3: CAEN TDC hit timings divided into $n_{\text{TI},6}$. The timing can be seen to shift for increasing values of $n_{\text{TI},6}$, and takes minimum shift values for $n_{\text{TI},6} = 2$. The figure is for run 1802 of the fall 2014 run, summed over all files.

of variance event-to-event is different time-walk between TDC and FADC and the intrinsic electronic resolution of each device. The two-dimensional histogram of Fig. 3 shows the difference between FADC and TDC times along the x -axis with the six values of $n_{\text{TI},6}$ along the y -axis. All channels corresponding the CAEN TDC are plotted. The one dimensional histograms shows the projections for each value of $n_{\text{TI},6}$. Note that these distributions are much sharper than those in Fig. 2. These distributions are used to determine the correct permutation for each run. Note that it would suffice to follow this procedure for a single channel; the correct permutation is the same for all channels in a given run. Also note that this calibration is done without using any external reference signal.

3.3 The Outlook

For the Spring 2015 run, the TS board has upgraded firmware that will control the phase of the 24-ns clock with respect to the 4-ns clock at synch time. This will not eliminate the six-fold ambiguity, but it will fix the permutation in the same way for each run. This will eliminate the need for a run-by-run calibration.




Figure 4: F1TDC hit timing distribution for Start Counter in run 1802. The hit times are in the TDC bin numbers.

4 F1TDC TDCs

The F1TDCs are based on a chip designed for the COMPASS experiment[2] and were developed at JLab by Ed Jastrzemski and others.

4.1 The Problem

The F1-TDC operates on a different principle than a traditional pipeline TDC. The TDC code comes from a 16-bit “TDC value counter” with an unusual property, described below. For each input pulse coming into a particular channel, the value of the counter at the time of arrival is recorded as the TDC code. Fig. 4 shows a typical distribution of raw hit times. Since hits come at random times with respect to the TDC value counter the distribution is uniform in time and populates all possible TDC values. In practical use, only TDC value differences are physically meaningful, and this behavior of the unreferenced times is not a problem as long as a suitable external reference is input to the system. As mentioned earlier, at the beginning of the Fall 2014 run, there was no such signal.

There is a system clock running with a 32 ns period (31.25 MHz). The clock driving the TDC value counter is derived from the system clock. The system clock is also used to determine the look-back time for read-out. Because the read-out window is tied to this clock, it can shift event-to-event with an eight-fold ambiguity depending on when the trigger arrives with respect to the system clock. However, the TDC value recorded for a hit is insensitive to this shift. This means that as far

as time registration with the FADC-250, this eight-fold ambiguity of the read-out window is completely irrelevant.

The unusual property of the TDC value counter mentioned above is that rather than running from 0 to 65535 ($10^{16} - 1$) and wrapping back to 0, it runs from a value a bit more than 0 to a value a bit less than $10^{16} - 1$, and then wraps back to the low value.

The size of the read-out window (“trigger window”) must be less than 40% of the wrap-around period (“TFRAME”). This ensures that TDC values will be unambiguous in that no more than one wrap-around event can happen in the read-out window. The beginning of the read-out window can be at most 90% of TFRAME before the trigger (“trigger latency”), where the terms in parentheses are those found in the F1-TDC document.[2].

Details of how the time bin size and wrap-around period are set are described in Appendix A.

4.2 The Solution

At synch time the TDC value counter is reset to its low value. We can now view the TI counter (n_{TI} , which we have already discussed in the context of the CAEN TDC above) as measuring the time of the trigger since synch time, *i. e.*, since the beginning of the run! We know the period of wrap-arounds of the TDC value counter, call it t_{frame} , so we know how many times it wrapped around since the beginning of the run, call it m . The design of the F1-TDC tells us that TDC codes recorded by the hits give the time since the last wrap-around, call it Δt . The time of the hit since the beginning of the run is then simply $t_{\text{hit}} = mt_{\text{frame}} + \Delta t$, all known quantities. The time of the trigger since the beginning of the run is given by the TI counter, $t_{\text{trig}} = n_{\text{TI}} \times (4 \text{ ns})$. In this way we can calculate the difference in time between the hit and the trigger $t_{\text{trig}} - t_{\text{hit}}$ unambiguously, *i. e.*, we have referenced times to the FADC-250 system. Again this procedure does not depend on any external reference signal.

In practice, what is done is to calculate the TDC code corresponding to the trigger time, *i. e.*, the code we would have obtained if the trigger had been input to a channel of the CAEN TDC, and use that representation of t_{trig} as a reference for each of the raw hits. Fig. 5 shows the correlation of all TDC hit times (t_{hit} , y -axis) vs. t_{trig} in this representation (x -axis). Fig. 6 shows the distribution of the difference of these two quantities, giving a physical time spectrum.

For the F1-TDC there is no need for a run-by-run calibration. The effect of an uncontrolled phase between the TDC system clock (32-ns period) and the FADC-250 clock (4-ns period) only effects the choice of the time window, as mentioned above.

For the Fall 2014 run, we did have an external trigger signal read-out by one of the F1-TDCs. If we use that *measured* signal for t_{trig} then we get Fig. 7 for the correlation between hit times and trigger time and Fig 8 for the distribution of their differences.




Figure 5: Correlation of TDC hit times against R for Start Counter in run 1802. Both axes have units of ns.




Figure 6: TDC hit times minus R for Start Counter in run 1802.




Figure 7: Correlation of TDC hit times against TD time for Start Counter in run 1802. Both axes have units of ns.

These distributions are essentially identical to those using the TI-based t_{trig} expect for a fixed offset.

Figs. 9 and 10 compare the two methods for obtaining t_{trig} directly. Fig. 9 shows the measured value (y -axis) vs. the TI-based value (x -axis) and Fig. 10 shows their difference. The inset histogram zooms in on the large spike just below 3,000 ns. The rms spread of the distribution is 70 ps. The source of the spikes at other times is still under investigation.

4.3 The Outlook

As of this writing that the measured t_{trig} signal has been used to synchronize the F1-TDCs with the rest of GlueX in all of the offline analysis. It is likely that most of the 70 ps spread seen above is due to the time resolution of the trigger signal measurement. If so, some of this resolution could be recovered by using the TI-based trigger time rather than the measurement.

5 FADC-125

At this writing the authors do not have a detailed understanding of the situation with the FADC-125s.

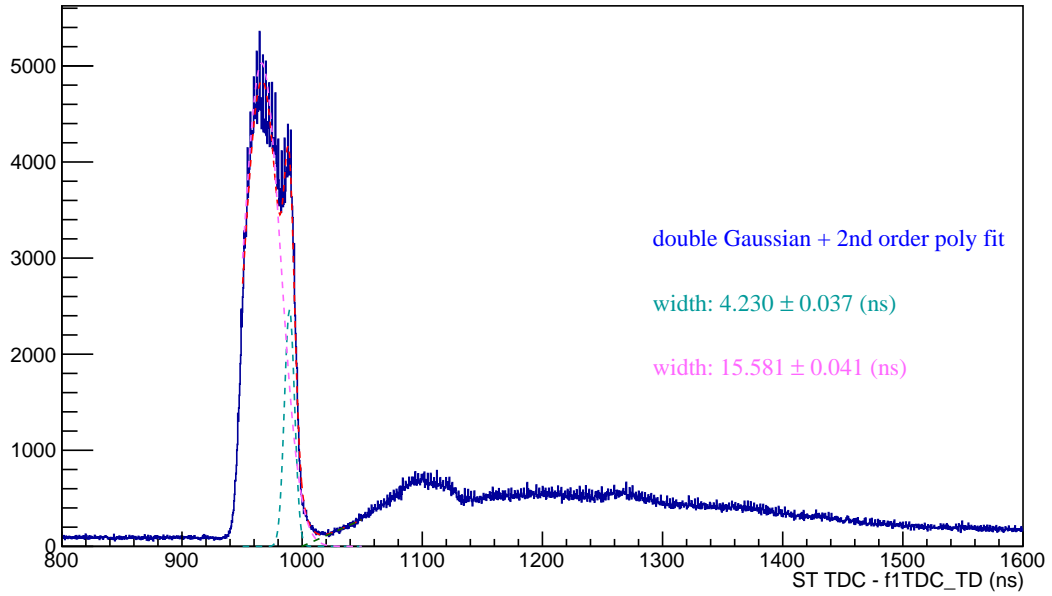


Figure 8: TDC hit times minus TD time for Start Counter in run 1802.




Figure 9: Correlation of TD time against R in run 1802. Both axes have units of ns.

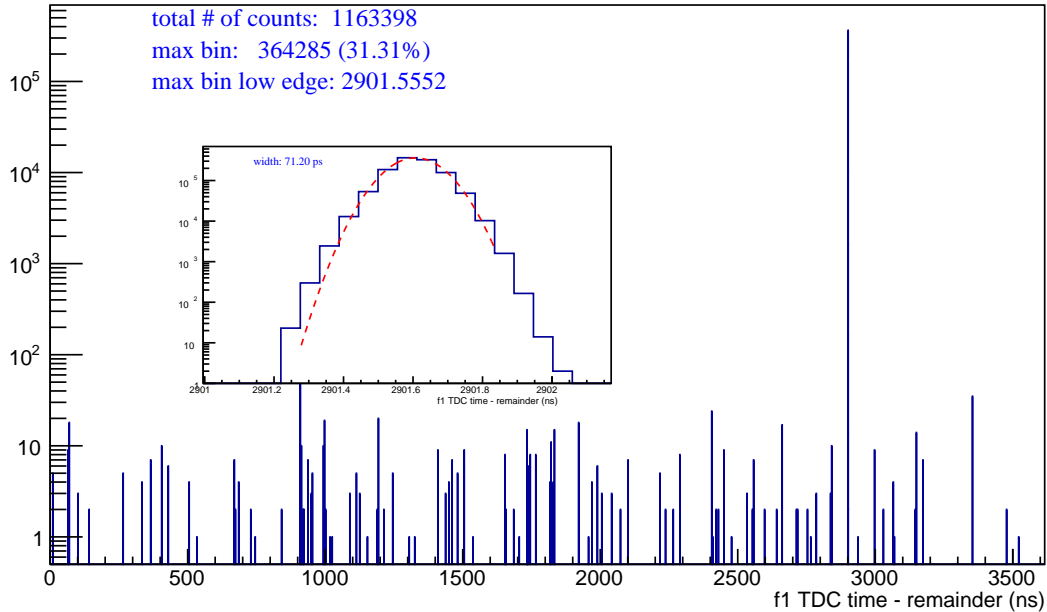


Figure 10: TD times minus R in run 1802.

6 Conclusions

For all of these timing systems, as long as times are referenced within the system, there is no problem with synchronization. The problems addressed in this note only arise when taking time differences between channels in different systems. In that case there may be event-to-event or run-to-run random shifts in the raw, unreferenced times. In all cases a global timing signal, input into each system suffices to synchronize the systems. If that global signal is derived from the accelerator RF frequency, there is the additional advantage that times can in principle be referenced to the arrival of the beam bucket of interest. Achieving this will depend on the GlueX detector as a whole having sufficient resolution to separate adjacent beam buckets.

Given the techniques described above however, it is possible to synchronize the F1-TDCs, the CAEN 1290 TDCs and the Flash-250s, even in the absence of a global timing signal, by using the system clock counter in the Trigger Interface boards. With a change to the Trigger Supervisor firmware it is possible to synchronize the F1-TDCs with the FADC-250s without run-by-run calibration of the relative phase of their clocks.

Going forward these properties, apart from any considerations from external global timing signals, should be checked to confirm correct system operation.

A Configuring the F1-TDC

When setting up the configuration, three parameters are specified. They are called `REFCNT`, `HSDIV`, and `REFCLKDIV`. There is an internal clock within the F1TDC that has a cycle of 32 ns, and the total time for the TDC to roll over is given by

$$T = 32 \times (\text{REFCNT} + 2), \quad (1)$$

where the additional two cycles are due to the system. For the GlueX F1TDCs, the parameter `REFCLKDIV` is always set to be $2^7 = 128$. The number of time bins to be used is proportional to the final parameter `HSDIV`, and since the rollover time T has been determined by `REFCNT`, the bin size for each time bin is inversely proportional to `HSDIV`. For completeness, the number of time bins N and the single bin size τ are given by

$$N = (\text{REFCNT} + 2) \times 152 \times \text{HSDIV}/\text{REFCLKDIV}, \quad (2)$$

$$\tau = \frac{T}{N} = \frac{32 \text{ REFCLKDIV}}{152 \cdot 2\text{HSDIV}}. \quad (3)$$

The number 152 is a constant within the F1TDC system. Since the value of `REFCLKDIV` is always fixed, the only parameters of interest are `REFCNT` and `HSDIV`.

A.1 Practical Use

When a run is started, the configuration parameters are passed to the F1TDCs. Starting from run XXX, this information has been written out to the data stream, so that by using the `DF1TDConfig` class, the configuration parameters are accessible from the data. There are several consideration when setting these parameters. Obviously, a larger value of `REFCNT` leads to a larger rollover cycle which is preferred if hits can come in a very large window of several μs . However, since the total number of bins N also increases, there is a limit imposed by the maximum value that N can take. Also, since the bin size of each bin τ is proportional to `REFCLKDIV`, while N is inversely proportional, a balance is needed between τ and N within the given constraints.

The configuration parameters for each run have been set as follows. Ed Jastrzemski has provided GlueX with a program that will give values of the configuration parameters based on input of what the rollover window size and bin size should be close to. For example, if the bin size is requested to be 56 ps with a rollover window of at least $3.5\mu\text{s}$, a possible set of configuration parameters is given by

$$\text{REFCNT} = 112, \quad (4)$$

$$\text{REFCLKDIV} = 128, \quad (5)$$

$$\text{HSDIV} = 482, \quad (6)$$

so that we have $T = 3648$ ns, $N = 65250$, and $\tau = 55.9$ ps. It is this convoluted way that the parameters are set that causes confusion.

There can also be an inconsistency since for Eq. (2), if all parameters are not chosen carefully, the value of N may be a non-integer value. This has actually happened to the earlier data before this possibility was pointed out to Ed...

To obtain the TDC bin size, the configuration information is needed. For run 1802 in Fig. 4, the configuration was

$$\text{REFCNT} = 224, \tag{7}$$

$$\text{REFCLKDIV} = 128, \tag{8}$$

$$\text{HSDIV} = 241, \tag{9}$$

so that the physical constants were given by $T = 3616$ ns, $N = 65250$, and $\tau = 55.9$ ps. Therefore, we can convert the hit times in Fig. 4 into units of ns using these constants; however, there is still no reference to these measurements, and for this we need either the F1TDC cable or the TI time.

For the cable passed from the TD to one F1TDC channel, the hit time which we call the TD time can be accessed within the offline software by requesting all DF1TDCHit objects from the event loop and accessing the correct channel. The code snippet below shows how this can be done for the fall 2014 data.

```

nf1TDC_TD = 0;
for(unsigned int i=0;i<f1tdchit.size();i++){
    if(f1tdchit[i]->rocid == 51
        &&
        f1tdchit[i]->slot == 17
        &&
        f1tdchit[i]->channel == 8){
        f1tdc_TD[nf1TDC_TD] = (ULong64_t)f1tdchit[i]->time; // *0.0559; // in ns
        nf1TDC_TD++;
    }
}

```

B Acknowledgements

The authors wish to thank Chris Cuevas, William Gu, Ed Jastrzembski, and Ben Raydo for explaining this stuff to us; Beni Zihlmann for cabling in the trigger reference signal for the F1-TDCs; and Aristeidis Tsaris for first “discovering” the six-fold ambiguity in the CAEN TDC times.

References

- [1] “V1290A-2eSST 32 Channel Multihit TDC (25 ps),”
<http://www.caen.it/csited/CaenProd.jsp?idmod=789&parent=11>

- [2] H. Fischer, J. Franz, A. Grunemaier, F. H. Heinsius, L. Hennig, K. Konigsmann, M. Niebuhr and T. Schmidt *et al.*, “Implementation of the dead time free F1 TDC in the COMPASS detector readout,” Nucl. Instrum. Meth. A **461**, 507 (2001), and references therein.