GPS Timing in Substations at Dominion
Where We are Now and Where We Need to Go

Introduction

Precise timing plays an increasingly essential role in the protection and monitoring systems being deployed throughout the grid at large. It is easily overlooked, however, even taken for granted, despite its growing necessity within our power infrastructure. Dominion, recognizing the emergence of timing as a potential facet in future (and even some current) protection schemes, sought to educate ourselves on the specifics of what makes the surprisingly complicated technology behind timing work.

Dominion has been using substation clocks since the mid-90s to synchronize our grid monitoring equipment. Our original clocks were GOES satellite and WWVB receivers. The GOES (Geostationary Operational Environmental Satellite) system effectively served as a precursor to modern day GPS. WWVB is a radio station dedicated to transmission of time, operated and run by the National Institute of Standards and Technology (NIST). In the late-90s, Dominion began transitioning to GPS based clocks and has remained with GPS since.

Accurate timing is required for a number of existing and upcoming power technologies. Metering is inherently associated with time. Post event fault analysis is enormously dependent on timing considerations. As schemes evolve to meet the growing complexity of our protection options, accurately sequencing events has become more and more critical. It also plays a major role in the traveling wave fault location method, and traveling wave systems are even being debuted as a protection technology in newer relays. Synchrophasor data is strongly timed as well, and again, PMUs are slated and intended to provide protection and operations functionality in the future. Finally, IEC 61850 is incredibly time sensitive, as it synchronizes protection schemes across a substation and even the grid. Many consider it to be the future of protection.

In this paper we'll address these issues in more detail and elaborate on how Dominion is diving into timing evaluation to make the best decisions going forward. We'll also cover the consequences of timing gone wrong and the surprising and unexpected problems it may cause. Timing is easy to cast aside as a given and treat as an afterthought, but the technology behind it is fascinating and deserving of attention and scrutiny.

Timing Basics and Common Timing Applications at Dominion

Satellite clocks used in the electric utility sector consist of several components, and understanding how these components work and interact is helpful in understanding how best to interpret your timing. The clocks consist of a GPS receiver chip that interfaces with custom electronics and an oscillator to generate the precise time. The GPS receiver chip is usually manufactured by a GPS specialty company. The output of the chip is then fed into the clock components to be translated in accordance with clock settings. It is important to understand that *the GPS receiver has a distinct firmware that is separate from the firmware used by the clock (the electronics that decode the GPS receiver's output)*. While they work in conjunction with each other as one system, they are independent entities.

Substation clocks, having interpreted the GPS signal, convert it into many types of outputs, such as frequency outputs, 1 PPS (Pulse per second), and Inter-range Instrumentation Group time code, more commonly known as IRIG. The specific type of IRIG used at Dominion, and indeed used by many power utilities across North America, is IRIG-B, which will be elaborated on later.

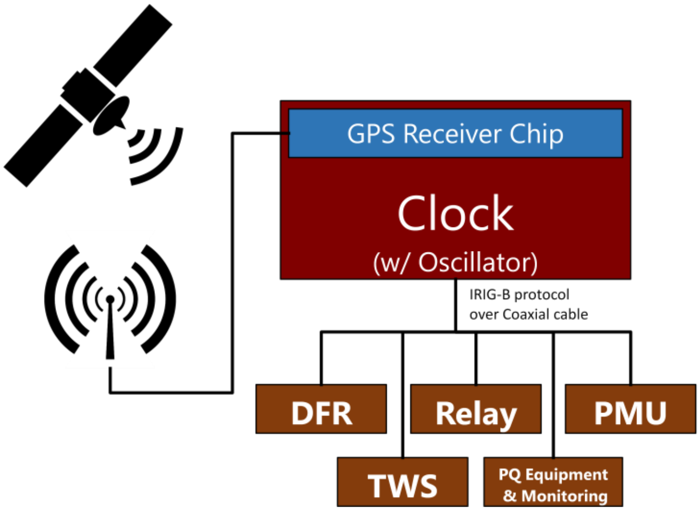

Figure 1 – Timing propagation from satellite down to individual devices

Dominion has GPS receivers installed at all transmission substations (100kV and above) and many distribution substations. The timing signals are provided to all protective relaying, Digital Fault Recorders (DFRs), and Phasor Measurement Units (PMUs), among other devices. With more time sensitive technologies on the horizon, IEC 61850 being a notable example, especially those with protection functionality, we anticipate the importance and necessity for accurate and dependable clocks to increase, but many of our existing applications, detailed below, already depend on accurate timing to maximize their utility.

Digital Fault Recorders (DFRs) and Protective Relays: *1ms accuracy required*

Virtually all new substation equipment today accepts a time sync signal and uses it to time stamp data. The data is principally of two types, sequence of event data and waveform or analog data. Sequence of event data (often called "SER" data) records the changing states of logic bits either external (via an input) or internal to the device, and each state change is accompanied by a precise time stamp of the transition. In the following example, a change of state of a circuit breaker is shown.

2017-02-19 17:56:22.561	*CB 372*	***Open***
2017-02-19 17:56:32.572	*CB 372*	***Close***

The time stamp can also be applied to point on wave measurements. When microprocessor relays or Digital Fault Recorders detect a fault, they can be programmed to record waveform data including voltages and currents measured before, during, and after a fault. This precise time stamp can be used to correlate events across a wide area. See *Figure 2* for an example of time-stamped waveform data.

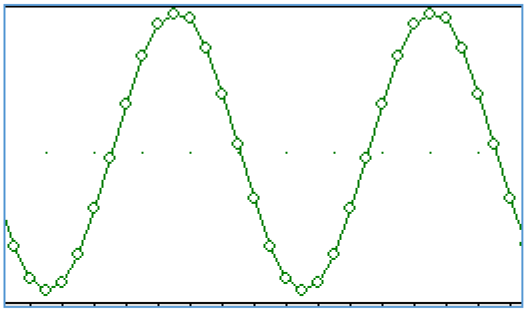

Figure 2 - Each point on the waveform has an associated precise time stamp.

In most cases, protective relays do not rely on precise time to make decisions to trip circuit breakers and clear faults from the system. Timing is, however, used for recorded data to simplify post-mortem event analysis. In fact, timing is so essential for post event analysis that NERC has issued a regulatory requirement, PRC-002-2, that mandates certain DFRs and/or relays have time synchronization accurate to within 2 milliseconds of the UTC time scale.

While timing isn't a staple of base-level modern protection today, it should be noted that future protection technologies, some of which are already here, do rely on precise timing to function properly. As mentioned previously, IEC 61850 is a protection standard that deals with communication of analog and digital quantities. When comparing data from across a substation or across the grid, having high resolution and accurate timing is critical. Likewise, traveling wave fault location, to be discussed below, can also be used for transmission line protection and is highly dependent on high resolution time data.

Lightning correlation: *1ms accuracy required*

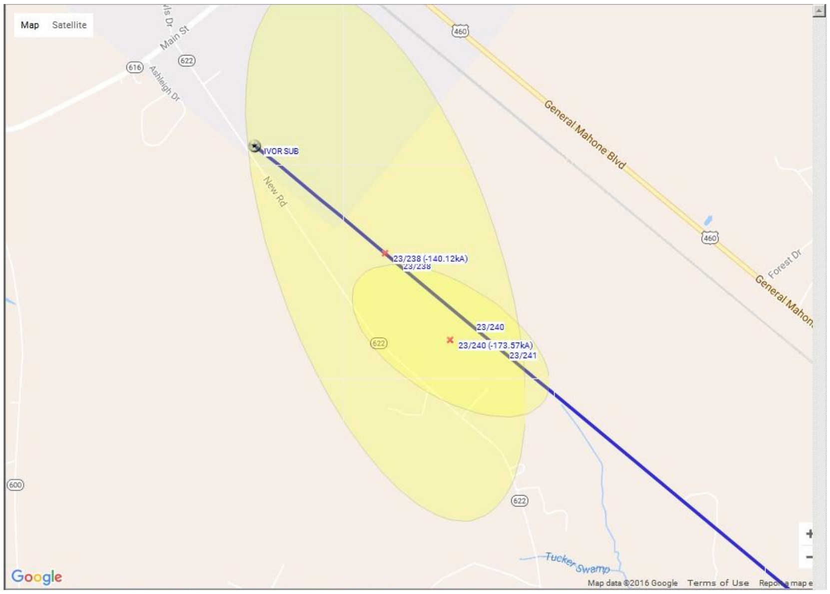

Figure 3 –Lightning Locations, highly dependent on precise timing information.

Lightning strikes are one of the most common causes of faults on the Dominion transmission system. While they generally only result in momentary interruptions, with the strike easily diminished before the first reclose attempt, they can cause physical damage to the equipment they impact, most notably insulators and arresters. As such, it is a good utility practice to inspect in the area of faults and correlated strikes. There are companies that provide lightning data for a fee.

These companies have sensors placed around a large geographical area and can detect location, intensity, and polarity of lightning strikes, all ascertained by triangulation of VHF interferometry in conjunction with LF magnetic direction finding. The strikes are time stamped to one millisecond and utilities can compare this strike data to the fault inception time measured from the point on wave timestamps. This comparison gives the utility a good location to inspect as well as an idea of what to look for based on the magnitude of the strike. Accuracy to one millisecond with respect to absolute time is sufficient for reliable lightning correlations.

Travelling wave fault location: *0.1 µs accuracy required*

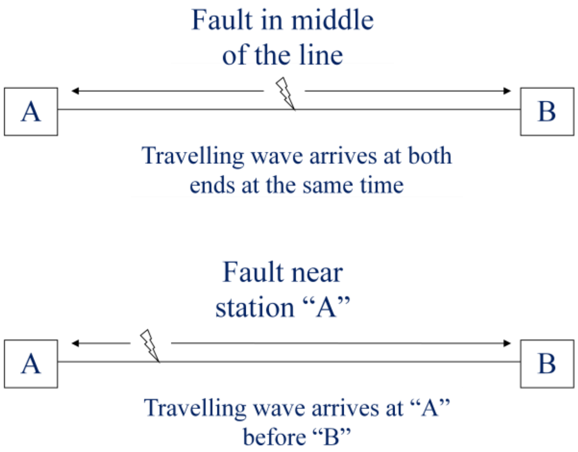

Figure 4 – Fundamental Principle behind Traveling Wave Location

Travelling wave fault location is a method of locating faults on transmission lines that is independent of voltage, current, and line impedance. This method is based on the principle that a travelling wave is generated at the point of the fault. The voltage collapse at the fault creates a high frequency pulse that travels down the transmission line at nearly the speed of light. The arrival time of this pulse is measured at each end of the line. With a precise arrival time, usually time stamped to 0.1 microseconds, the fault location can be determined within a few hundred feet. This method requires very accurate and precise absolute time at each terminal that measures the wave arrival time.

As mentioned before, traveling wave fault locations are now being used by some relay vendors as a protection technology. The relays on both sides of the line are in constant communication, and if both see the traveling wave high frequency pulse and determine it was generated within the line, they operate. The relay vendors claim it's one of the fastest interruption technologies available, requiring just a relatively simple algebraic calculation.

Synchrophasor data: *1 µs accuracy required*

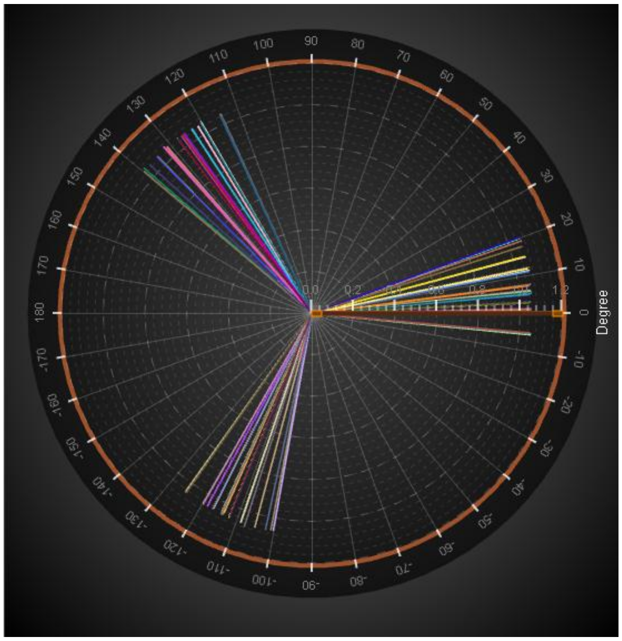

Figure 5 – Voltage Angles, as reported by PMU devices across Dominion's grid.

Synchrophasor data consists of very precisely time stamped voltage and current phasors sent at very high resolution to an operating center, typically at a resolution of 30 or 60 phasors per second. Data from many substations are sent to a centralized server, and the server/applications can compare readings from across the grid. These comparisons can give power system operators and engineers a near real time indication of the health of the grid. This same data can be used to provide inputs to wide area control systems that can make decisions and initiate control actions. When synchrophasor data is used to provide automatic control actions, accurate, reliable, and secure timing is essential. This data requires an absolute time accuracy of 1 microsecond.

Power Quality Monitors: *1ms accuracy required*

Power quality monitors are most often used on the distribution system to monitor the power being delivered to customers. If a customer is having issues with their power, having time synched data can help the utility correlate events with the customer's measurements to events recorded by nearby PQ monitors. It is also possible for transmission level events to create voltage problems on the distribution system. In this case, having time synched data from the PQ monitors makes it easy to correlate with the events recorded from the transmission DFRs or protective relays.

SCADA measurements

SCADA measurements provide data to grid operators about the condition and state of equipment in the field. Breaker status (open or close), transformer temperature readings, voltage, current, and power measurements all come to the operators via the SCADA system. Data sent to the operators can either be time stamped with the time of the actual event or with the time the reading arrived at the operating center. The scan time of SCADA systems can be anywhere from one or two seconds to 10 seconds or more. A lot can happen between scans, so it is helpful to know the time of each event rather than the time the data arrived.

Timing: You don't need it until you need it

The need for timing isn't always apparent until it's "too late." With post event analysis in particular, it's hard to appreciate how useful timing is until you try to piece things together without it. The Northeast Blackout of 2003 is considered by many to be one of the first motivators behind mandating high resolution, standardized timing across all protection devices and monitors. When trying to pin down the first in a long series of dominoes, timing is everything. When an event extends over a wide area, particularly across multiple time zones, decrypting what happened would be all but impossible without a solid timeline.

Timing should not be undervalued on smaller scales either. At Dominion we recently encountered a problem with power line carrier signals reaching their destinations. We have SER data for our carrier sets, identifying when blocking and tripping signals are sent and received, and for how long, but these devices did not have IRIG signals wired to them. A recent failure to effectively send line transfer trip drove us to dig into our carrier timing, but because our SER data wasn't synchronized, lining the events up proved very difficult. Had we not experienced this event, we'd have never assumed precise timing was necessary for carrier sets, but having struggled through this once, going forward we've decided to include IRIG inputs on all carrier devices, now having a better appreciation for the value it provides. Ultimately, if you have the option of precise timing, it's highly recommended to use it.

More Timing, More Problems

Time is such an integral part of everyday life, it's easy to forget how complicated it can be. A high functioning GPS driven clock is actually subjected to a variety of transitions over the course of the year, not all of which are strictly scheduled. Leap second adjustments, for instance, aren't on a schedule. A clock cannot know whether or not we'll have a leap second seven years from now, so they can't be programmed to account for them. Clocks depend on the GPS signal itself to anticipate such timing adjustments, and it doesn't always go smoothly. Clocks are likewise dependant on their antenna, just as a clock's down-line devices are dependent on the clock. There is room for error, error which Dominion has experienced firsthand.

The following is a partial list of the various critical events a clock must be capable of handling.

- Daylight saving time: The clock must be capable of appropriately gaining and losing an hour depending on the transition to or from DST.
- UTC Year roll over: UTC is the time against which all other global times are offset. The transition between years is one worth monitoring, especially because it may be accompanied by the loss or gain of a leap second.
- Regional Year roll over: The regionalized year roll over, Eastern Time in Dominion's case, can be prone to issues and should be investigated.
- June leap seconds added or removed: Leap second adjustments, if required, are typically implemented at two times during the year, in June and during the UTC year roll over, as previously mentioned.

- GPS Epoch week roll over: Weeks are identified within the GPS signal, but the space allocated allows a maximum entry of 1024. Therefore every 1024 weeks, a GPS week "epoch" roll over occurs, and the weeks are reset to 0.
- Holdover: Holdover takes place when the GPS signal is lost by a clock. The clock must maintain good time using its internal oscillator until the signal is reacquired.

Dominion has experienced a number of issues during these transitions, which is in fact what prompted us to investigate timing in the first place. During the transition from 2015 to 2016, an error occurred with our standard substation clock model that altered the time on hundreds of relays and ultimately resulted in the failure of most of them to appropriately register the year transition. Our clocks are set to report local time with the UTC offset. Many of the communications processors receiving this signal are Windows based, and as Windows based machines, they run natively at UTC and must translate the signal using the provided local offset. Inexplicably, and undesirably, the programmed Eastern Standard Time UTC offset changed within the clock immediately at midnight UTC (the moment of the UTC year transition) from -5 hours to +3 hours. Consequently, all relays receiving time from one of these communications processors found its time off by 8 hours. When local time hit midnight, the local start of the new year, the offset returned to -5, but because many of these relays infer the year rather than read it from the clock, they never transitioned. The relays needed to see a transition from 23:59 to 00:00 to register the new year, but with the temporary 8 hour offset, the clocks went from 15:59 to 00:00, eluding the year transition logic.

This issue compounded itself in surprisingly serious ways. Many of these relays were involved in remote control schemes and alarm reporting such as reclosing adjustments, abnormal potential alarms, and so on. The medium through which these control adjustments and alarms were transmitted, however, was rejecting information from the relays because the time-stamps claimed the wrong year. It is essential for operators to be aware of many of these alarms, and losing presumed control over remote relay behavior and information can have potentially dangerous consequences to field personnel on top of jeopardizing grid operability. Increasingly, the protective aspects of our grid are genuinely dependent on time.

Dominion had another more recent transitional issue, although this time with a leap second. Leap second additions are implemented as a 61st second in the final minute before midnight UTC. Dominion experiences them at 7PM local time in December and 8PM local time in the summer. The included leap second should register at 7:59:60PM. Our clocks, however, failed to recognize the leap second and transitioned to 8:00PM as usual. Five seconds later, internal error checking within the clock noticed we were a second ahead and broadcast the time 8:00:05PM for two consecutive seconds, effectively postponing the true leap second. While this sounds like a minor problem, it was enough to take down our synchrophasor server. The server rejects incoming data that doesn't fall within a window of accuracy, and data one second off easily qualified for filtering. The server failed to return to proper functionality even after time re-synced, so manual intervention was required. Had we been working in an environment where synchrophasor data provided critical support to either operations or protection, we would have had a major problem on our hands.

Proper leap second progression:	Improper leap second progression:
19:59:58	19:59:58
19:59:59	19:59:59
19:59:60	20:00:00
20:00:00	20:00:01
20:00:01	20:00:02
20:00:02	20:00:03
20:00:03	20:00:04
20:00:04	**20:00:05**
20:00:05	**20:00:05**
	20:00:06

Figure 6 – Example of appropriate leap second implementation vs. what we observed during a leap second in mid-2015

Both of these issues were remedied with clock firmware updates. It's important to recognize the distinction between clock firmware and GPS receiver chip firmware. In our experience, most issues stem from clock firmware, with no issues originating from the GPS chip, but either can require an upgrade. Only through monitoring and testing can users be sure their clock of choice is appropriately handling timing events, and to effectively monitor our clocks at Dominion, we first needed to learn their language, IRIG-B.

Demystifying the Signal: A Deep Dive into IRIG-B

Dominion, as a practice, tests any new firmware we implement, especially those pertaining to protection applications. Given the various functions our clocks provide, they are deserving of rigorous firmware testing, but the prospect of testing all devices that require clock signals against a new clock firmware is daunting. Rather than testing each individual device, we wondered, "What if we could just test the IRIG signal itself?" If the IRIG is correct, we know we can at least trust the clock to deal with timing events properly, and that's our principle concern for updated clock firmware.

How, then, do we check the IRIG signal itself? With a strong understanding of the IRIG-B data stream, interpreting the signal is fairly simple. The key is recording the data. To do that, we created an IRIG data monitor using an Arduino microcontroller. The microcontroller is programmed to interpret the signal and record it, in binary, to a connected laptop. The file can later be interpreted by a user. Altogether, the monitor is less than $50, but it's proved an incredibly useful tool for firmware and clock testing.

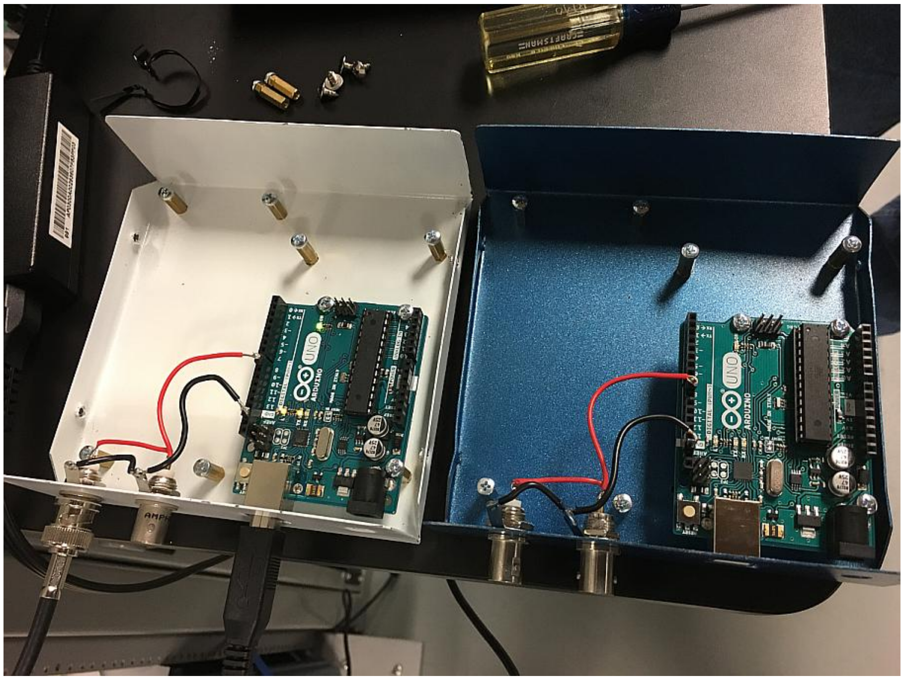

Figure 7 - $50 Arduino IRIG monitor

The IRIG-B signal itself, while not overly complicated, requires some decoding to understand. The format we use at Dominion is IEEE 1344 with SBS (straight binary seconds). IRIG-B sends a repeating signal format once a second, at 100 pulses per second. Within those 100 pulses, we decode both the time and additional timing information. Here's a quick overview of what each second contains.

- **Second Number:** *7 bits, 0-60*; leap seconds come through as second 60
- **Minute Number:** *7 bits, 0-59*
- **Hour Number:** *6 bits, 0-23;* 24 hour clock used
- **Day Number***: 10 bits, 0-366;* leap year accommodated
- **Year Number:** *8 bits, 0 – 99;* the century and millennium are assumed
- **Leap Second Pending:** *1 bit;* asserted up to 59 seconds before leap second insertion
- **Leap Second:** *1 bit;* leap second adjustments can either add or remove leap seconds, with 0 adding a leap second and 1 removing a leap second
- **Daylight saving changeover pending:** *1 bit;* asserted up to 59 seconds before daylight saving time change
- **Daylight saving time active:** *1 bit*
- **Time zone sign:** *1 bit;* sign of local time zone offset
- **Time zone offset:** *4 bits, 0-12*
- **Time zone offset (30 minutes):** *1 bit;* provides addition 30 minute offset for time zone
- **Time figure of merit:** *4 bits;* 0 indicates highest merit, 0xF indicates worst merit, used as indicator of time accuracy relative to UTC based on clock's internal parameters
- **Parity:** *1 bit;* provides parity of all preceding bits, used as an indicator of signal quality
- **Continuous Time Quality:** *3 bits;* similar to time figure of merit, although distinct in that it indicates maximum deviation in time accuracy even when the clock is synchronized to UTC *(Continuous Time Quality is exclusive to C37.118 and is not a part of the IEEE 1344 standard)*
- **Straight Binary Seconds:** *17 bits, 0-86399;* total number of seconds since midnight

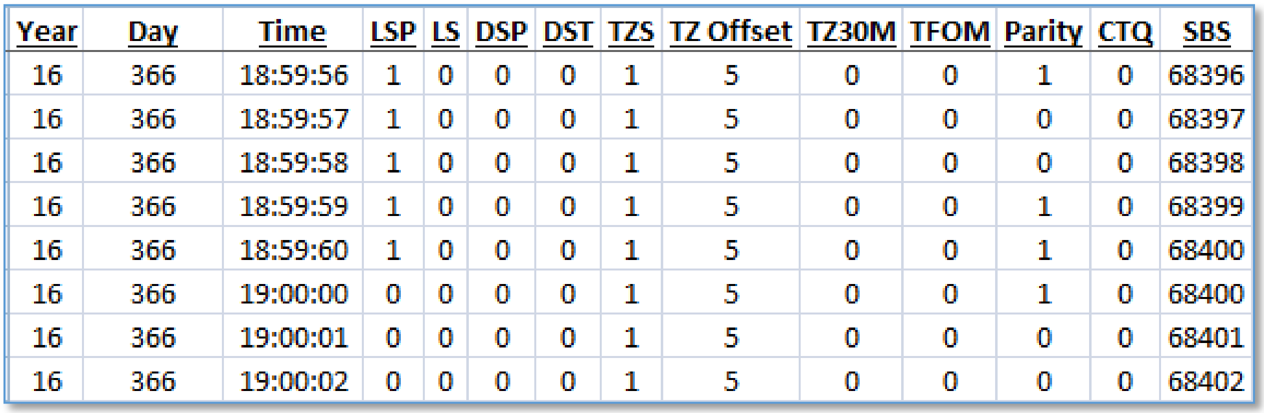

Year	Day	Time	LSP	LS	DSP	DST	TZS	TZ Offset	TZ30M	TFOM	Parity	CTQ	SBS
16	366	18:59:56	1	0	0	0	1	5	0	0	1	0	68396
16	366	18:59:57	1	0	0	0	1	5	0	0	0	0	68397
16	366	18:59:58	1	0	0	0	1	5	0	0	0	0	68398
16	366	18:59:59	1	0	0	0	1	5	0	0	1	0	68399
16	366	18:59:60	1	0	0	0	1	5	0	0	1	0	68400
16	366	19:00:00	0	0	0	0	1	5	0	0	1	0	68400
16	366	19:00:01	0	0	0	0	1	5	0	0	0	0	68401
16	366	19:00:02	0	0	0	0	1	5	0	0	0	0	68402

Figure 8 – Example of decoded IRIG-B signal on the final day of a leap year experiencing the addition of a leap second at midnight UTC

Moving Forward

Dominion is currently evaluating some of the latest clocks on the market to determine which clock(s) we intend to use going forward. With our newfound knowledge of the intricacies of timing, we have developed a comprehensive test plan against which our chosen options will be subjected. We've created a relational database to track results, storing everything from clock performance to manufacturer issues.

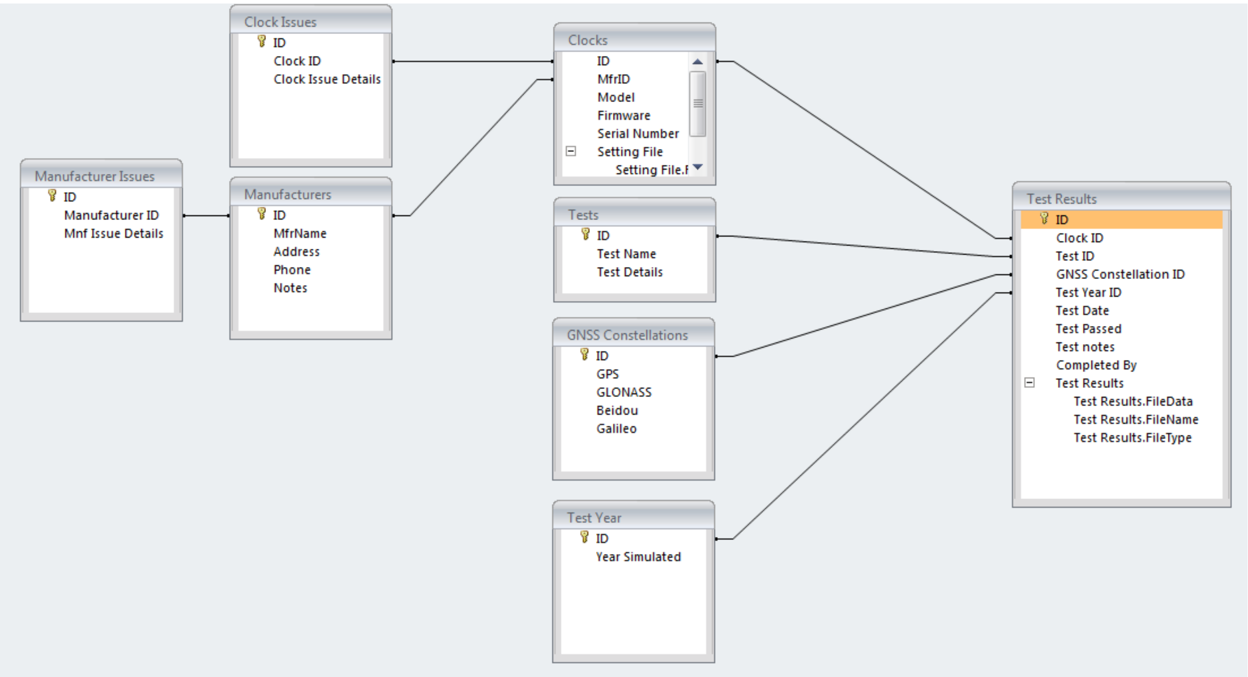

Figure 9 – Snapshot of relational database hierarchy, detailing the organization of our testing procedure and results

An initial set of ten tests will be run through each clock under a variety of GNSS satellite configurations. Furthermore, these tests may be run for any given year, the range of which currently runs from 2017 to 2030. Altogether, this sums to over 10,000 distinct tests, although we will obviously be prioritizing likely scenarios over obscure ones. GNSS constellation technologies being considered are GPS, GLONASS, Beidou, and Galileo, and these tests will be run under not only single constellation options, but any combination of the above, amounting to 16 varieties. Again, however, the priority will be on the common options.

The ten tests we currently plan on running, which may grow, is as follows. These exactly mirror the transition events mentioned in the "More Timing More Problems" portion of this paper.

- **Daylight Saving Time On:** transition from DST off to DST on
- **Daylight Saving Time Off:** transition from DST on to DST off
- **Year Rollover UTC:** no leap second: year transition at UTC midnight with no leap second
- **Year Rollover UTC:** leap second add: year transition at UTC midnight adding leap second
- **Year Rollover UTC:** leap second remove: year transition at UTC midnight removing leap second
- **Year Rollover Local:** year transition at local time
- **June Leap Second Add:** midyear leap second addition
- **June Leap Second Remove:** midyear leap second removal
- **GPS Epoch Week Rollover:** transition from week 1023 to week 0
- **Holdover:** ability of clock to retain time during loss of signal

A dependable GNSS satellite signal simulator is a critical element in running the aforementioned tests. In 2016, Dominion purchased a GNSS simulator and attended extensive training on functionality and use. The simulator is able to be programmed to achieve all variations of our testing arrangements, GNSS satellite configurations, leap second implementation, and of course specific time generation. A rack including each prospective clock and the simulator has been built in our test lab. The next steps are simply to pick our priorities and start testing.

Conclusions

When Dominion first started researching timing, none of us anticipated how far we would come. No one expected there to be so much to it, and because it was so unexpected, and because we were all, in fact, so surprised by its complexity, we realized a stronger knowledge of timing wasn't just important to us but also to the industry at large. One of the most important takeaways here is that this seemingly simple and routine aspect of our day to day lives, a modern luxury taken for granted, is incredibly powerful and needs to be treated as such. When and if timing is neglected and ignored, there have been and will continue to be serious consequences to the protection and operation of our grid. Timing must get the level of attention and scrutiny it deserves. Our main discoveries and perspectives are provided below.

- Accurate and dependable timing is increasingly necessary. Not only do more and more modern protection schemes depend on it, but improper timing has proven surprisingly problematic in unexpected ways.
- The complexity and importance of timing is easily underestimated. Utilities should be making a conscious effort to better understand the role timing plays in their protection and monitoring schemes. Recognizing issues stemming from timing problems is difficult if you don't have a strong foundation in timing fundamentals.
- There is value in being able to monitor and test your clock's performance. You can't manage what you can't measure. Understanding the signal itself helps you better understand how your equipment uses it, and it simplifies firmware testing. Rather than testing each piece of equipment's interaction with the clock, you need only test the clock itself.
- The more we learned about timing, the more we wanted to learn. We saw it as necessary and worthwhile to implement and execute a robust set of timing tests going forward, both to evaluate current clocks as well as new clocks worth pursuing. Timing systems should be given the same priority and thoroughness as any substation system.

Biographies

Patrick Hawks
804-393-8276 (mobile)
804-257-4716 (fax)
patrick.j.hawks@dom.com
Dominion Virginia Power
2400 Grayland Ave, Richmond, VA 23220

Robert Orndorff
804-396-5757 (mobile)
804-257-4716 (fax)
robert.m.orndorff@dom.com
Dominion Virginia Power
2400 Grayland Ave, Richmond, VA 23220

Kyle Thomas
804-316-2861 (mobile)
804-257-4716 (fax)
kyle.thomas@dom.com
Dominion Virginia Power
2400 Grayland Ave, Richmond, VA 23220

Patrick Hawks has worked at Dominion Virginia Power since 2008. He has a B.S. degree in Computer Engineering from Virginia Commonwealth University. His experience includes six years in system protection, wherein he participated in system modeling, relay settings, and protection standard development, implementing a substantial degree of automation, both in setting calculation and setfile generation. He has since moved into Fault Analysis and become involved in analyzing system events and web app development.

Robert Orndorff has worked at Dominion Virginia Power since 1984. He earned an A.A.S degree in Electronics in 1986 and spent 11 years as a field relay technician and in 1997 transferred to the Fault Analysis department where he currently works. His current responsibilities include maintaining and configuring Dominion's Digital fault recorders, event retrieval and analysis from smart relays and DFRs. Robert is an IEEE member and has been a member of the Transient Recorder's User Council (TRUC) since 2002.

Kyle Thomas received his M.S. degree in Electrical Engineering from Virginia Tech in 2011 and is currently pursuing his Ph.D. while Supervisor of the Dominion Virginia Power's Electric Transmission Operations Research group. He has technical expertise in power system protection/control, wide-area measurements, fault analysis, cascading analysis/physical security, and system simulations. Kyle is a technical lead of Dominion's synchrophasor installations, applications, and training, and is actively involved in the North American Synchrophasor Initiative (NASPI), IEEE, TRUC, and CIGRE organizations.

References

IEEE Standards Association, *IEEE Standard for Synchrophasor Measurements for Power Systems, C37.118.1, 2011*

Range Commanders Council, *IRIG SERIAL TIME CODE FORMATS*, IRIG Standard 200-16, published August 2016. Retrieved from http://www.wsmr.army.mil/RCCsite/Pages/Publications.aspx

Vaisala, *Vaisala Thunderstorm and Total Lightning Sensor TLS200*, Retrieved from http://www.vaisala.com/en/meteorology/products/thunderstormandlightningdetectionsystems/lightningsensors/Pages/default.aspx

Interface Control Working Group, *Global Positioning Systems Directorate Systems Engineering & Interface Specification IS-GPS-200, 9-Dec 2015*. Retrieved from http://www.gps.gov/technical/icwg/

US Naval Observatory, *Leap Second Information page*, Retrieved from http://tycho.usno.navy.mil/leapsec.html