

Timetagging and Delays in Phasor Measurement Systems

KE MARTIN
Electric Power Group, USA
martin@electricpowergroup.com

KEYWORDS

Synchrophasors, Phasors, Delay, Latency, PMU, PDC,

1 INTRODUCTION

Direct estimation of phasors from power signals is a concept that grew out of computer relaying in the 1980's [1]. The construction of a stand-alone Phasor Measurement Unit (PMU) in 1988 was made possible by the introduction of relatively high-power microprocessors for the high computational burden, and the Global Position System (GPS) for precise timing. Phasor measurement systems are now extensively deployed worldwide with applications only beginning to be developed. The most established use is data acquisition for event and system analysis. Real-time monitoring is also widely used, and at least a few control applications have been developed [2]. Each application has particular measurement process and overall system requirements. These include accuracy, resolution, reporting rate, and reporting delay as well as signal selections. This paper investigates the delay, or latency, in measurement and reporting, from the time the AC system is at a particular value until the data representing that value is available for use. The elements of delay in a phasor measurement system are broken down and described, and measured results in an operational system are presented. The terms "latency" and "delay" are used interchangeably, though usually refers to the overall effect.

2 ELEMENTS OF LATENCY

2.1 Phasor measurement systems

In a typical phasor measurement system, a phasor measurement unit (PMU) connects to voltage and current secondaries at a substation. It estimates the synchrophasor equivalent of the AC waveforms. It sends the data in real-time to a central site and may record it locally as well. At the central site, a phasor data concentrator (PDC) receives and correlates the data from several PMUs, sends it to other end use systems and may record the data (Figure 1).




Figure 1 : Typical phasor measurement system

This analysis examines the measurement and transmission to the PDC. The PDC will introduce certain delays; that activity and resulting latency is analyzed in the IEEE PDC guide C37.244-2013[3]. For analysis here, the measurement chain is broken down into separate elements as shown in Figure 2.

The time blocks somewhat representative of typical delays found in production systems. The following sections discuss the elements in some detail.




Figure 2: Elements of latency showing an approximate time contribution for each

2.2 Input filtering and A/D conversion

The measurement process starts with input of the AC voltage and current signals. In all current PMUs, initial over range protection and input filtering is followed by an A/D to convert the signals to a digital form. Input anti-alias filtering is needed to remove signal components above the Nyquist frequency for the A/D sampling rate. These delays are generally small, but can approach 1-2 ms.

2.3 Phasor windowing

An AC waveform cannot be represented by a single point. A span of points on the waveform is required to approximate what the waveform looks like. The more complex the waveform, the more points and time are required. That is, if it was known that a waveform was a pure sinusoid at a particular frequency, only two points in $\frac{1}{2}$ of a cycle determine the exact phase and amplitude (phasor) characteristics. But in general the sinusoid is not pure and the exact frequency is not known, so more points and time are needed. A measurement covering a complete cycle allows deriving a good representation of the waveform. With a one cycle window, data must be output faster than 2/cycle to prevent aliasing. If the output is slower than that, more filtering is required. Filtering expands the measurement window whether included in the estimation stage or applied to the phase estimates. Generally the filtering will require a window a little greater than twice the reporting interval. This is discussed in detail in Section 3.

2.4 Phasor processing

Phasor processing consists of estimation of phasors, frequency, and ROCOF as well as any other calculations done by the PMU. The processors used in modern PMUs are very fast and this delay is usually a small part of the overall latency. In two cases where it was possible to estimate the processing time, it took less than 0.5 - 2 ms for six 3-phase positive sequence phasors.

2.5 Data transmission serializing

When the data is sent from one system to another, it has to be converted from an internal memory format into an exchange protocol and sent between the systems. It will usually be converted into a serial protocol where data is sent out one bit at a time. The amount of time required for each bit depends on the clocking speed of the interface. The number of bits depends on how the protocol converts the bytes of the message and the number of bytes in the message. Taken together, the message time interval depends on the clocking speed and the bit length of the message. Note that the serializing delay is only counted once as it is de-serialized at the receiving end while it is serialized at the sending end.

It is important to consider all places where the data is serialized. The data may be sent and received across several links, and each will typically add its own delay. For example, data may be sent by Ethernet to a router that sends over a WAN and then over another router to another Ethernet. Data will be buffered, usually by message, in each router and then serialized onto the next link. Thus the serializing delays for each link need to be added.

Synchrophasor messages using the C37.118.2 [4] messaging are short, usually on the order of 40-70 bytes. When using RS-232 serial, 2 bits are added to each byte, a 25% protocol overhead.

Ethernet using TCP or UDP adds a fixed overhead which is about 50% overhead for messages of this size. Encryption and other security considerations can add much more than that.

Message protocol:	RS232	RS232	TCP/UDP	TCP/UDP	TCP/UDP
Rate (BPS)	38.4 K	256 K	T1 (1.544M)	10M	1 G
Bit time (μs)	26.04	3.9	0.65	0.1	0.001
Serialized time (μs)	18,229	2,734	544	84	0.83
Ave data rate (30/s)	21 KBPS		25.2 KBPS		

Table 1: Approximate serializing times for a 70 byte C37.118 message. Assumptions: RS232 with 10-bit/byte format (25% overhead) & TCP/UDP (50% overhead).

Example delays for typical serial, communication system, and Ethernet speeds are in Table 1. The average data rate illustrated at 30 data frames/s shows that the overall channel loading is small. Serial communication at 38.4 KBPS has 18.2 ms of delay, which is considerable. Even a relatively fast T1 circuit adds 0.65 ms with the added IP protocol. Note that long distance communication will often use a fractional T1, so the serializing delay will be at the fractional rate.

2.6 Modem, router, & other interface processing

Phasor measurements are transmitted in digital format. Analog modems are used when the data is sent over an analog communication system. With all digital systems there will be a CSU/DSU, digital modem, router, or some other interface device. They all perform the same functions: protocol conversion, buffering to handle differing rates, and physical media interface. Most can include retransmission for error correction as well.

In order to send data over an analog communication system the digital representation has to be translated into an analog signal. Simple coding schemes have limited bandwidth, so complex trellis coding schemes were developed to fit more information into a limited bandwidth. The down side of these more complicated methods is that the coding takes longer; it was found that a simple coded 9600 BPS link had a delay of 8 ms while a V.32bis modem delay was 60 ms.

Digital modems, CSU/DSU, and similar devices provide translation between terminal and communication system bit rates and usually asynchronous to synchronous translation. Some buffering delay is required to match rates. Data input may be in bursts as long as the average rate is lower than the channel rate. There may be some delays for compression and error checking where used, but most of the delay will probably be the communication channel serial rate. No digital modem (or CSU/DSU) system was available for testing, but a sub rate unit (SRU) in the fiber multiplex system performs the same function and was tested. It was found to introduce a 10-15 ms delay when measured the same way as the analog modem delay.

Modem latency can be eliminated by eliminating the modem. A data stream output by the PMU can be “sampled” by the synchronous sub-rate input to a SONET channel bank as long as the sub-rate is greater than twice the PMU data output rate (meeting Nyquist sampling criteria). This eliminates any interface delay, and was found to work very well. This approach reduced delay to only the serializing which in this case was 5.5 ms (70 byte message) at a 128 KBPS clocking rate.

Routers usually have one or more Ethernet network inputs and a WAN input. They not only provide protocol conversions, but filter and direct traffic as well. They also have buffering and prioritization capabilities as well. As long as the network is not busy, the delays will be limited to the messaging buffering time, generally less than 1 ms. If the network is busy and some links, such as the WAN, are overloaded, there will be delays as the excess traffic is buffered. Highest priority will be sent first, so if phasor traffic is set at a high priority, it will experience minimum delay. Router delays should be 0.1 ms to 10's of ms, but could be much longer if traffic is not controlled.

Retransmission for errors is provided in several RS232 and network protocols. This simplifies application functionality for error management. However it also increases the maximum delay that must be anticipated and can cause additional network backups that increase overall data loss. For applications that have no time criticality, such as archiving, using these protocols is beneficial. For applications that provide data where time of delivery is important, it is usually better to use an unsecured method like UDP to deliver data without repeat and design the network to eliminate

overloads. In that usage, a few lost samples cause less disruption than a delivery backup along with blocks of lost data.

2.7 Transmission distance

Transmission distance delay is time it takes to send a signal from one point to another. If transmitted by strictly electromagnetic waves through space, it would be the speed of light. In real terms, data will go from a digital signal into a communication system carrier. This carrier will propagate through the air as a radio signal, or over an artificial medium such as a fiber optics as a light beam. Signals travel on a radio or microwave carrier at close to light speed, perhaps 0.99% of light speed. Over a fiber optic cable the speeds are slower, about 67% of light speed. In either case it is usually a small contribution to the overall delay. For example, a signal traveling 300 km will have a 1 ms delay over a microwave system and a 1.5 ms delay over a fiber optic cable.

Communication using geostationary satellites is significantly longer. They orbit at an altitude of about 36,200 km above the equator of the earth. Sending a signal to the satellite and back covers 72,400 km for a total trip time of 0.242 sec, much longer than with terrestrial systems.

2.8 Communication system processing

This category includes processes that are not included elsewhere. Within a communication system, the overall signal may be demodulated to split out some signals before others are sent on. There may be buffering between elements of the overall communication system. The system may include backup routing such as SONET rings. With this technology signals are sent two directions in a ring and the receiving end can take either signal. The ring may operate continuously or the ring may come into play only when the primary route fails. In either case, the ring operates within the communication system so the phasor system will only see differences in delay. The same is true for buffering within the communication system. These delays and delay changes will usually be within a couple ms or less, so are not a major delay contributor. Since they are very specific to the particular system, the overall effect must be evaluated by the system designers.

2.9 System input processing (PDC or Application)

The system that is receiving the data has to de-serialize it and place it in some kind of internal storage or buffer. The receiving device will typically have to receive a whole packet in order to do error checking. Often the receiving system will be handling other processing tasks as well, so there may be some random delays involved. All timing tests used in this study were using a Phasor Data Concentrator (PDC) which has somewhat random delays due to multiple processing, though the typical delay is around 2 ms.

3 DELAY IN THE PHASOR MEASUREMENT PROCESS

3.1 Synchrophasor Estimation

A phasor is not a quantity like a voltage or current magnitude that can be measured directly from an electric or electromagnetic field. First it is the representation of an AC signal. The signal values vary with time so the signal must be observed over a period of time to determine the properties. The length of this “window of observation” depends on the parameters to be determined and the signal characteristics. A synchrophasor estimate requires enough of a complete cycle of the fundamental to accurately estimate the magnitude and phase angle. If the frequency were accurately known and there were no distortions to the sinusoid, this could be done in a small fraction of a cycle. However these conditions are generally not true, so something approaching one complete cycle is the minimum required. Second, phasor parameters are magnitude and phase angle at a fixed frequency, so “measurement” requires a reference frequency and a process to convert signal waveform samples from the time domain to the phasor parameters in the frequency domain. The most common way to do this is using some form of a Fourier transform, the most popular being a discrete Fourier transform (DFT) or a fast Fourier transform (FFT). There are other ways to estimate phasor values, but to simplify the following discussion, the phasor estimation process will be analyzed using digital samples of a 60 Hz AC waveform and a DFT to estimate the phasor.

3.2 Synchrophasor Timetag

The phasor estimate will be a kind of ‘average’ value across the window where it is computed. While the input data may change instantly, the estimate will change gradually as the change slides into the window. Using a sample rate of 2880/sec and a 1 cycle DFT with a calculation window (48 sample points), Figure 3 shows the response to a step change in the AC waveform. The phasor value is computed at each sample point, so each window overlaps the previous one by 47 points. As the window moves along the waveform, the estimate begins to change as soon as the stepped waveforms enter the window. The phasor value ramps up linearly for 1 cycle until the whole window is past the step. This complete response completes in one cycle (.0167 sec). This illustrates the effect of the window in the response to changes in the signal. If the window is longer, it will clearly take more time for the estimate to assimilate changes in the signal. However a longer window will minimize the effect of measurement errors caused by noise or other effects.




Figure 3: Step change in ac: 1-cycle phasor estimation window, timetag at end of window.




Figure 4: Step change in ac: 1-cycle phasor estimation window, timetag at middle of window.

The phasor value is an instantaneous value of the signal parameters at a particular point in time. In the example of Figure 3, the time of the last sample was used for the time of the estimate. This creates a normal appearing input and response. However the phasor value is an “estimate” of the phasor value within the window. Since the values in the window are changing, the phasor should be the “best estimate” of their value which will be some kind of average of what is in the window. To assess the measurement, we use TVE which is the complex difference (vector error) between the “true” value and the calculated one. The true phasor can be estimated in this case by inspection: before $t=0$ it has a magnitude of 0.7 and after the step 1.0. The phase angle does not change. The 4th plot in Figure 3 shows the error value. Before $t=0$ the error remains 0 but it jumps to 29% at the step because the phasor has not begun to respond. Then the error returns to 0 over the cycle as the estimate responds. Does this represent the “best estimate”?

If we use the timetag to the middle of the window, we get the measurement as shown in Figure 4. Note that the response ramp for the phasor is moved back to the phasor estimate amplitude is about .85, half way to the final value. The estimate and error now start ½ cycle before the step and end ½ cycle beyond. However the maximum error is now 21% and the error at all other points is commensurately reduced. Experimenting with the timetag position will show that using a timetag at the center of the window will minimize the error throughout the window. In most cases of power system measurement, as long as the windowing and estimation functions are symmetric across the measurement window, the estimated phasor best represents the value at the center of the window.

In this example the windowing is illustrated graphically with a step change. A more subtle example but one that has a more common impact is measurement at off-nominal frequencies. At off-nominal frequency, the phase angle will constantly rotate. For example, if $f = 60.1$ Hz, then the error is 0.1 Hz and the phase angle will rotate 360° in 10 s. In 10 s there are 600 cycles, so in one cycle the

phasor rotates 0.6° . The rotation is uniform, so the estimate will yield the angle in the center of the window. If the timetag is at the middle, the error is 0. If at either end, the error will be 0.3° . If the window is longer frequency further off nominal, the error is larger. Since this type of error is continuous, it clearly illustrates the need to use the time in the center of the window.

3.3 Filtering and output

In the illustrations a new phasor is computed for every data sample. This will certainly give the fastest and most accurate measurement but the amount of data thus created is impractical for communication and storage. Real phasor measurement units thus transmit a sub-rate of the sampled values, usually at a rate that is a submultiple of the power system frequency. This can lead to an aliasing problem. In any sampled data system frequencies greater than twice the sample rate will be aliased into the sampled data, appearing as a signal at a lower frequency. For example, if phasors are sent out (“reported”) at a 30/s rate, any signal content with frequency ≥ 15 Hz will be aliased. Filtering is required to prevent aliasing. Any filtering, whether applied with the original estimate or on the phasor estimates, effectively becomes part of the window and the timetag should be at the center of the total window. Generally the minimum filter length is a little over $1/(\text{Nyquist frequency})$, so anti-alias protection for the 30/s reporting rate requires a 4 cycle filter. (If there is no frequency content above Nyquist as determined other ways, the filtering is not necessary.)

In summary, a synchrophasor is estimated over an interval of time that is long enough to take into account frequency content of the signal and the measurement reporting rate. The timetag for the estimate is time of the center of the window. The PMU will have to take samples for the whole window and then make the estimate. Thus the phasor will be transmitted a small interval of time after the end of the window. If the since timetag is used as the reference time, the delay for transmission of the data from the PMU will be $\frac{1}{2}$ the window length and a small amount of time for the phasor computation. In modern PMUs, the computation time is generally 100’s of microseconds, so does not contribute greatly to delays. However $\frac{1}{2}$ the estimation window can be significant. For example, if phasors are sent at 10/s and anti-alias is required, the minimum delay will be 200 ms.

4 DELAY MEASUREMENTS

4.1 System Setup

This section presents results of delay measurements made both in the laboratory and on an operational system. Delay measurements require a start and end time, and it is necessary to find a consistent and appropriate time to use. In this case, we are looking for the time from when something occurs on the power system to when the measurement can be available to control system. The phasor timetag accurately reflects time a power signal value occurs, which can be used to determine an event, so the timetag is the appropriate delay start time. The end time is when the data has been received and available for use by an application. In this case we will use the time supplied by the PDC when the data is written into the internal table. The PDC function may be built into a controller; if an external PDC or other post processing equipment before the controller is used, that extra delay will have to be accounted as well.




Figure 5: Delay measurement scheme using GPS as the common timing reference

Figure 5 shows the scheme for these measurements. GPS supplies precise UTC time to both the PMU and the PDC for the start and end time stamps. This measurement time includes all delays including the PDC input processing as shown in Figure 2.

The phasor timetag should be accurate to within 1 μ s, and must be within 26 μ s (31 μ s for 50 Hz) to pass the C37.118.1-2011 [5] standard requirements. This PDC has a real-time operating system that minimizes program response to real time inputs, but the hardware is old enough that the typical input delay is 1 - 3 ms with some processing variability and interrupts. The local timetag is derived from a hardware timing clock synchronized to GPS that provides 1 μ s resolution. So the measurements have microsecond resolution but some unknown accuracy errors in the small millisecond range and interrupts approaching 10 ms. Most delay measurements are in the 10's of ms, so it is easy to visually differentiate effects. 3-5 minute files of data at a 30/s rate were used (5400-9000 measurements) so it is easy to see trends and exceptions as well as unusual phenomena.

The PMUs used in this test provided the timetag in various places. Two gave the timetag only at the end of the window, one at the beginning, and one was selectable. If the time stamp corresponds to the end of the estimation window, then the only delay in the PMU is the input and phasor processing. In two cases this delay was found to be less than 1 ms, and in another which had distributed PMU units, it was more than 20 ms. The timetag location will be noted in the following analyses. Note that current commercial PMUs only use middle of window timetag since that has been shown to produce minimum error.

4.2 Direct PMU to PDC connection

These tests used the PMU connected directly to the PDC without modem or other communication system between. An end of window timetag was used so the delay includes only phasor processing, output serializing, and input system processing; the normal $\frac{1}{2}$ window delay is not included.




Figure 6: Delay for a direct connection PMU-PDC connection with RS232 serial




Figure 7: Delay for a direct connection PMU-PDC connection with Ethernet

The original PMUs all used serial communication. The Macrodyne 1690 used in these tests outputs data at a 38.4 kbps rate. The time to output the first bit of a message is consistently about 0.5 ms. In Figure 6 the delay recorded in a 3 minute file is plotted. The message was 38 bytes, so based on Table 1 above the serializing delay is about 9.9 ms. The average delay in the plot is 12.3 ms, so the PDC input delay is approximately 1.9 ms. The time spikes at the start of plot and at the 5000 mark are due to the PDC writing data to the disc drive.

For contrast, a Figure 7 shows an average 4.4 ms delay for a 10 MBPS Ethernet connected PMU using TCP (with no other traffic to prevent packet delays). The serializing delay for the 112 byte packet is about 90 μ s, insignificant in this measurement. Previous tests on the PDC indicated the Ethernet input delay was about 3 ms. Since we were not able to measure PMU output delay directly, we assumed that the output delay was about 1.4 ms.

These two examples clearly show the serializing delay impact. This delay will add to the total no matter where introduced. For example, in a SONET or SDH system, the information may be carried over a fractional T1 (E1) which could have a rate as low as 64 KBPS. Even though both ends may be Ethernet, the data will slow through the channel link, adding to the overall delay.

4.3 PMU to PDC connection through modems

The next two examples use analog V.34 modems operating at 33, 600 BPS. One modem type used an RS-232 serial input operating at 38.4 KBPS over an analog microwave, since it was not possible to get a pair set up in the lab for this test. The signal distance was about 300 km which would add a rather insignificant 1 ms delay. The other modem had a 10/100BaseT Ethernet input and was tested back to back in the test lab.

The delay shown from the on-line example in Figure 8 shows about 5 ms continuous noise with some larger outliers. The average delay is about 94 ms. Contrast this with the back to back test in Figure 9. The noise is about 1 ms and the overall delay 61 ms. The higher noise related to added communication systems is not surprising, but the 33 ms difference was not expected. Since that is the time interval between data reports it probably relates to internal modem buffering. The serializing delay of about 20 ms sets the minimum, but the added 40-75 ms is attributable to buffering and coding. Many details of modem operation are internal to the modem operating system and not available to user analysis. Both PMUs use an end of window timetag.




Figure 8: Delay through RS232/V.34 analog modems over 300 km analog microwave.




Figure 9: Delay through Ethernet/V.34 analog modems back to back.

4.4 PMU to PDC using direct modulation

This method uses an over-sampling principle to convert the asynchronous RS232 serial PMU output to the synchronous serial used in the SONET communication system. A 128 KBPS sub-rate is $>2\times$ the 38.4 KBPS serial PMU output, by Nyquist theory we can use that to transmit the data. Surprisingly it works well and is practically error free. The delay is greatly reduced to around 19.5 ms as shown in figure 10. The serialized time is dictated by the actual message which was 50 bytes requiring 13 ms at the 38.4 BPS rate.

4.5 PMU to PDC using Ethernet over routers & WAN

This is the most common type of installation being used now since Internet Protocol based on Ethernet is almost. In this example the PMU sends data to a local router which forwards the data by WAN to another router that converts it back to Ethernet for the PDC. The WAN uses a 256 KBPS sub-rate carrier and travels about 300 km. The serializing and distance delays are only 3.3 ms and 1 ms, accounting for little of the observed delay. Careful setup and management of all router aspects would probably allow reducing this figure.

**Actual Trends in Development of Power System Protection and Automation
Yekaterinburg, 03.06 – 07.06, 2013**




Figure 10: Delay using direct sampling into 128 KBPS multiplex, 100 km fiber SONET.




Figure 11: Delay using Ethernet/256 KBPS WAN with routers, 300 km fiber/MW SONET.

5 CONCLUSION

The sources of delay in the measurement reporting chain have been described with both an indication of how significant they are and how they can be managed. Measured results quantify this analysis. These findings are summarized in Table 2 below.

Delay source	Range (ms)	Comments
Input filtering & A/D	0.1 to 2	Higher sampling rates lower delay
Phasor windowing	4 to 200+	Delay is $\frac{1}{2}$ the estimate/filter window
Phasor processing	0.2 to 3	PMUs now use DSP/Co-processors, very fast
Data serializing	.001 to 30	Be careful of sub-rate links in comm chain
Modem processing	1 to 80	Analog modems not often used now
Transmission distance	.1 to 8	Actually a small contribution unless very long
Communication processing	1 to 15	This varies widely, should be controllable
PDC/controller input processing	.1 to 10	Could be longer, should be short by design

Table 2: Summary of delay elements

The latency issue will become more important with real-time controls. This work is drawn from measurements starting with early phasor measurement systems using RS232 serial through current work with routed IP communication. More thorough analysis of the effects of elements and ways to control them is needed for critical real-time controls. Further in-depth analysis is left to future studies.

REFERENCES

- [1] A.G. Phadke, J.S. Thorp and M.G. Adamiak (1983) A New Measurement Technique for Tracking Voltage Phasors, Local System Frequency, and Rate of Change of Frequency, IEEE Trans PAS 102, Issue 5, pp 1025—1038
- [2] E.M.Martinez, J.L.Marquez, “Application of Phasor Measurement Units in Adaptive Protection of Interconnected Systems”, Comision Federal de Electricidad (CFE), Mexico.
- [3] IEEE Guide C37.244-2013 “IEEE Guide for Phasor Data Concentrator Requirements for Power System Protection, Control, and Monitoring”.
- [4] IEEE Standard C37.118.2-2011 “IEEE Standard for Synchrophasor Data Transfer for Power Systems”.
- [5] IEEE Standard C37.118.1-2011 “IEEE Standard for Synchrophasor Measurements for Power Systems”.