

Applications of Synchronised Power Quality Measurements

Bill Dickerson, P.Eng.

Introduction

Synchronised phasor measurements (synchrophasors), made using Phasor Measurement Units or PMUs, have become widely accepted for wide-area measurements in the electric power grid. They are important because they allow comparison of data made at different locations as though the measurements were all made by a single instrument, with all channels sampled at the same times. This is made possible by synchronisation of the PMUs using a common time distribution system, typically GPS but also possibly including methods such as IEEE-1588 Precise Time Protocol.

Having synchronised data is important, because it allows measurements such as phase angle variations across a wide area. This white paper will show that such synchronisation is also desirable for power quality data, and in fact it is necessary for measurements where data are being compared or used for calculations, and the data originate in different locations.

The existing standard IEC 61000-4-30 is commonly used when making power quality measurements, harmonics in particular. This standard was written primarily for determination of conformance with legal requirements for power quality, and as such it addresses measurements made at a single location. It does describe means to time tag measurements, to aggregate sequences of measurements, and in fact it imposes specific requirements for the measurement method. We will see here that these requirements are not well suited for harmonic measurements made at diverse locations, when an application intends to perform any sort of wide area determination of power quality performance.

Similar observations can be made for other power quality measurements, such as flicker. Flicker measurements are specified in IEC 61000-4-15, and while this standard suggests a block diagram for a flicker meter, it does not impose normative requirements for methods. In fact, any method meeting the performance requirements may be used. This provides additional flexibility, but the fundamental issue still remains: if flicker is to be compared across a wide area, then synchronised measurements are required.

Synchronised measurements vs. time tagging

'Synchronised measurements' are: *measurements made at substantially identical moments in time in different locations or devices*. For instance, the PMU (according to IEEE Standard C37.118.1) generates a phasor estimate exactly at the top of the second (i.e., estimating the parameters of the signal for times xx.00000 seconds), and every T_s thereafter. $T_s = 1/F_s$, where F_s is the reporting rate, typically equal to the system nominal frequency, or half of it (i.e. 25 or 50/s for a 50 Hz system, and 30 or 60/s for a 60 Hz system). This is done using data collected over an interval controlled by a clock synchronised to a recognised standard time source, typically UTC (or local time derived from UTC). Since all PMUs are synchronised to the same time source, they will all make measurements at the same exact points in time, within a narrow tolerance. This tolerance is typically better than one microsecond.

When making measurements this way, for example at 50 Hz and 50 reports per second, the PMU provides an accurate, synchronised estimate of the fundamental magnitude and phase angle for times of xx.0000, xx.0200, xx.0400... xx.9800, for all seconds xx. This is true for all PMUs, so configured, at all points in the network. Thus, the data may be collected at a central location, and meaningful comparison and analysis can be performed. In essence, the group of PMUs making these measurements may be viewed as a single instrument, making all of its measurements at substantially identical moments in time.

Power quality measurements may be made in this way as well, but this requires an instrument designed to work this way. A normal power quality analyser, even if 'synchronised' to UTC, does not do this. In fact, if it faithfully implements IEC 61000-4-7, it cannot do this. The requirements for synchronisation in IEC 61000-4-30 allow an error of up to one cycle – 20 ms at 50 Hz and 16.667 ms at 60 Hz – because each measurement begins at the zero crossing of the fundamental component at that location. As PMU measurements clearly show, the fundamental phase angle varies across the grid. Measurements of harmonic phase angle, for example, made with this level of synchronisation performance would be useless for comparison between multiple PQ analysers.

Many devices make measurements using what are in effect free-running methods, and then apply a time tag to the result. This approach is probably adequate when comparing aggregated power quality data between a utility and customer for regulatory or contract compliance, but it is not sufficient for the applications discussed here.

Introducing the notion of precisely synchronised power quality measurements makes possible the capture and analysis of the dynamic behaviour of PQ disturbances at much faster rates than is now practical, from 10 to 50 or 60 measurements per second, or even faster.

IEC 61000-4-7

This standard has been around for many years, and in fact was previously known as IEC 1000-4-7, before IEC renumbered their standards. Earlier versions of this standard merely set performance requirements, and did not attempt to standardise on a particular measurement method. This approach allows the greatest flexibility and provides the opportunity for new techniques to be developed that advance the state of the art.

Later versions of this standard, however, focus on obtaining comparable measurements using different instruments by specifying that harmonic measurements shall use a 200 ms nominal data window, with uniform weighting (rectangular window). This window is required to comprise either 10 (50 Hz) or 12 (60 Hz) complete cycles of the input signal frequency, so in reality it will only be exactly 200 ms long when the power frequency is at nominal. Further it requires that the first measurement window of a sequence shall begin at the first zero crossing of the fundamental following the top of the second; as we have seen with synchrophasors, windows so defined do not coincide at different locations.

The purpose of the 10- or 12-cycle window with rectangular weighting is to minimise the error that results in FFT analysis from the Gibbs phenomenon, when the number of cycles in the measurement interval is not an integer. This error causes a spreading of the signal spectrum, with energy appearing in many bins of the resulting spectrum. This complicates analysis of the results.

Implemented perfectly, this method can work well. However, under dynamic conditions this method often fails to follow the signal frequency perfectly (if indeed there is any such thing – what is the ‘frequency’ when it is changing?). The result is that spectral spreading still occurs. Furthermore, harmonics in the power system are not always exact multiples of the fundamental – they may result from mechanisms (such as sub-synchronous motors) that create ‘near-harmonics’: those that rotate in phase relative to an ‘ideal’ harmonic. Other methods exist to reduce the problem of signal spreading, and as we will see these methods are better suited to synchronised power quality analysis.

The notion of synchronisation was introduced in IEC 61000-4-30 many years ago with the aim to make comparison between results measured with two different instruments possible. This was a serious issue in those days. In essence, two instruments were supposed to make their measurements over the same cycles, and since the data were to be aggregated over 3 seconds and 10 minutes, synchronisation errors of 20 ms (one line cycle) were negligible. It was a major improvement at the time.

The requirements of IEC 61000-4-7 and -4-30 fundamentally preclude making truly synchronised power quality measurements (even if device synchronisation is significantly better than one line cycle). As experience with PMUs has shown, the frequency varies at different points in the grid. Further, even when the frequency is stable enough that it might be considered constant (that is, for a normally quite short period of time), the phase angle is different at different locations. Measurements synchronised to the zero crossings and using an integer number of cycles will inevitably be made over different intervals at different locations.

The bottom line is that 61000-4-7 is perhaps acceptable for devices whose function is to determine compliance with regulatory requirements for power quality at a single point of interconnect. Its usefulness is limited, likely non-existent, for measurements which are to be used to understand system performance across a large area, where separate PQ analysers are needed to make the measurements.

Modifications to 61000-4-7 and -4-30 for synchronous applications

We have implemented a method in our Model 1133A Power Sentinel PMU/PQ analyser that avoids the problems caused by having a non-integer number of cycles in the measurement window, and allows us to use truly synchronised measurement methods. In fact, we use the exact same data samples that we use for the PMU measurements. These samples will be taken at the same moment in time, across the measurement area and even across the world, in each and every PMU/PQA device (within a very small error, of ~200 ns or less).

The Model 1133A uses overlapping, 100 ms data windows (each containing 1024 samples), using Hann weighting of the data window. The Hann window is a 'raised cosine' window:

$$W = 0.5 - 0.5 \cos(2\pi \cdot x);$$

where x varies from zero to one from the beginning to the end of the window. This window has the desirable characteristic that when subsequent measurement results are aggregated, each data sample has unity weighting in the resulting sum, except for the 50 ms at the beginning and end of the first and last window. The initial and final 'tails' have the same shape as the original window function: a half cosine, for the Hann window.

Dr. Andrew Roscoe calls this process 'tessellation.' This term is taken from the building of mosaics, and in computer graphics when 'tiling' images to create the illusion of 3D. Here, it involves building up a longer measurement by adding together results of a series of shorter measurements. This process may be applied repeatedly, for instance, to aggregate data over windows of 3 seconds and then 10 minutes. Windows used to perform tessellation require certain properties. Such windows maintain the desirable characteristic of tessellation over longer and longer intervals, when the data are properly aggregated.

The rectangular or uniform window, and a cascade (convolution) of two rectangular windows – the triangular or Bartlett window, also the P-class filter in IEEE Std C37.118.1 for synchrophasors – are other examples of window functions having this property. Windows that share this property have a Fourier expansion in which all of the even-order frequency terms have zero magnitude, except for the dc term. Even-order terms reinforce in such an aggregation, where odd-order terms cancel; so the only even-order term that should be non-zero is the dc or 0th order term. The Hann window has this property; so do the Hamming and Bartlett, but the Blackman does not – it includes a 2nd order component that creates a ripple of the weighting of the points when aggregated.

Another way of saying this is that for each point of the window function $W(n)$ for $n = 0 \dots N/2$, if $W(n) + W(n+N/2) = K$ with K a constant (ideally 1), then the window has this 'tessellation property.' For a window function to be used with the FFT, N is generally a power of 2; but that is not required for this observation to be true. Many windows don't have the tessellation property; but there are several that do.

In the Model 1133A, we provide once per second aggregated results of 20 windows, starting with the one *beginning* at the previous top of second and ending with the one *centred* at the current top of second. (These results may be further aggregated in applications software.) Aggregation is performed by root-mean-square averaging of the moduli of each harmonic in the 20 data sets. Absolute (i.e. relative to UTC top of second, as for synchrophasors) phase angle is determined as the argument (arc-tangent) of the data for the window centred at the current top of second. A total of 600 numbers are provided each second: absolute phase and magnitude for each of 50 harmonics, for all six channels. (Relative phase between voltage and current can easily be found by subtracting the phase of the harmonic voltage and current. Absolute phase cannot be found using the opposite, as with a standard PQ analyser.)

The Hann window (like most window functions except the rectangular or uniform window) tapers to zero (or for some windows, a small number) at the ends of the window. This is what minimises the signal-spreading problem caused by a non-integer number of fundamental cycles in the measurement window. The spreading results in effect from the transient caused by the 'jump' which would appear when the data in the window repeats. The FFT in effect presumes that the data in the window is a representative sample (period) of a periodic sequence that starts at negative infinity and continues to positive infinity. If the points at the beginning and end of the window do not 'connect' seamlessly, a step is generated, and this step has a frequency spectrum that extends from dc to infinity. These additional frequency components cause the spreading of the spectrum as it appears in the output of a Fourier spectral analysis.

The trade-off of using weighting with a window function is that the main 'lobe' of the resulting spectrum is broadened. This is a consequence of all weighting filters, and in fact the greater the suppression of the signal spreading due to the Gibbs phenomenon provided by a given window function, the more the main lobe of the resulting spectrum is broadened. Simply adding together the total energy in the resulting broadened spectrum can compensate for this. This is three bins in the case of the Hann window. Or, the exact transfer function for the window function at off-nominal frequencies may be determined and used for a more precise correction. However, the accuracy requirements for harmonics imposed by the standards (5% of harmonic magnitude) do not require this level of correction, and the simpler method suffices. The standards are likely adequate in this particular regard, since in the power system, harmonic magnitudes are quite variable on a second to second basis.

For the Hann window, a signal at the centre of a frequency bin (e.g. a 50 Hz signal for a 100 or 200 ms window) will be spread into three bins: the expected one plus the two adjacent ones. The commonly used Blackman window, which has even more rejection of signal spreading, spreads the spectrum into five bins. Other, less-common filters, such as the Nuttall 4-term (7 bins), provide even more out of band rejection at the expense of even broader signal spreading.

For our purposes, we determined that the Hann window was the optimum one for this application, because it provides sufficient rejection of spreading effects due to the Gibbs phenomenon, and provides the further benefit of uniform weighting of input samples during aggregation; the Blackman and Nuttall windows do not.

Note also that measurement of interharmonics is significantly improved by this approach, since these signals are not harmonically related to the fundamental and the ‘integer number of fundamental cycles’ method will not in general provide reduction of spectral spreading for the non-harmonically-related interharmonics. This effect is quite severe for the uniform window. Weighting with a window function works equally well for all signals, regardless of their relationship to the fundamental frequency or the window length.

With this modification to the requirements of IEC 61000-4, we can provide harmonics measurements that have the same feature as PMU measurements: they are made for windows centred exactly at the top of second and each 50 ms thereafter. This will be true, within the synchronisation performance limits of the Model 1133A, for all devices located anywhere in the world. Further, the harmonic measurements are compliant with IEC 61000-4 in all performance regards.

Flicker (IEC 61000-4-15); Sags/Swells/Interruptions (IEC 61000-4-30)

Everything said here about the theory of the measurements, and noted below regarding applications, also applies to flicker and for that matter, to sags, swells, and interruptions. These measurements should also be performed using synchronised data, and the measurement algorithms should preserve that synchronisation. We have followed this approach in our Model 1133A, and suggest that other vendors should do the same to promote interoperability.

Flicker is measured using a complex model to estimate the human annoyance caused by amplitude variations of the power signal. This model, specified in 61000-4-15, includes infinite-impulse-response (IIR) filters that have long decay times; nevertheless, we have found that the instantaneous flicker perception data can be meaningfully compared between instruments and locations using synchronised data samples. Despite the ‘obfuscation’ introduced by the relatively complex signal-processing model, significant improvements result from synchronisation.

Sags, swells, and interruptions, like 61000-4-7 harmonics, also use zero crossings to start the measurement. The definition of these impairments does not as readily allow identical measurement windows across a network; but using synchronised samples does allow precise determination of the time of occurrence of sags, swells, and interruptions.

Applications

Propagation of power quality disturbances

There is a desire to understand how power quality disturbances propagate throughout the power grid, both locally and over wide areas. This is basically a *transimpedance* measurement: i.e., the ratio of a voltage measured at a remote location, caused by a current injected at a given point:

$$Z_T = \frac{V_{meas}}{I_{inj}}$$

To make this measurement accurately, both the injected current and the resulting voltage (which are complex quantities, i.e. having magnitude and phase) must be measured over the same interval, using the same time scale. This requires synchronised measurements. This measurement can be made for any of the disturbances described above: harmonics, flicker, and sags/swells/interruptions. (Note that the Arbiter 1133A provides measurements of flicker on both current and voltage channels – the standard applies only to voltage measurements; but with measurement of flicker on induced (load) currents, this same method may be applied to flicker as well, enabling transimpedance measurement.)

Many practical considerations are likely to enter into the problem when making these measurements. Specifically, performing an ANOVA (analysis of variance) can, with enough independent variables and sufficient data, separate the various PQ contributions from different sources in a grid. It remains to be seen how many different measurement points, and how long an interval of observation, will be required. Nevertheless, with synchronised measurements there exists the possibility of separating the effects.

Another potential contribution of synchronised measurements is that auto- and cross-correlation methods may be used if the data are measured using complex numbers (i.e., including phase angle – not simple magnitudes). Complex measurements require synchronisation; non-synchronised measurements do not provide useful phase information. Cross-correlation techniques may prove more powerful than traditional statistical ANOVA methods.

Estimation of line impedance

Estimation of line impedance is a useful part of understanding how PQ disturbances propagate. Line impedance may be estimated at the fundamental frequency using PMU measurements, but what about harmonic frequencies? Clearly the impedance will be quite different at these higher frequencies than at the fundamental, if for no other reason than that in most cases the line impedance is dominated by the inductance of the wire transmission lines, and the reactance of this inductance increases with frequency. However there are other effects which keep this relationship from being a simple linear one, such as the contributions of underground cable inter-wire capacitance, bulk capacitor banks, and other VAR compensation and voltage stability control devices, to name a few.

Line impedance can also be calculated by making measurements at geographically separated points. However, this is only possible if the complex values (voltage and current, also possibly power flows) are made using a consistent time reference, so that the resulting phase differences are meaningful. At harmonic frequencies, this requirement is even more critical than for the fundamental: a time error resulting in one degree phase error at the fundamental (measurable but perhaps tolerable) will cause 50 degrees error at the 50th harmonic – rendering the measurements all but useless. 200 ns synchronisation uncertainty will produce an additional phase angle uncertainty of 0.18 degree at 2.5 kHz, the 50th harmonic of 50 Hz. This level of performance is most likely required if useful measurements are to be made at these harmonic frequencies. (This requirement also applies to other devices in the measurement chain, such as current and voltage transducers. This is not a trivial requirement, either; though there is reason to believe that it can be satisfied through careful selection and possibly calibration of the transducers.)

Clearly, for line impedance estimation at harmonic frequencies, measurements must be made with very carefully controlled synchronisation between the various analysers. Free-running PQ analysers (or those synchronised only to the level required by IEC 61000-4-30) will provide results that are for all purposes random (i.e. useless) when trying to perform line impedance estimation using measurements from different locations, or even different instruments at the same location.

Summary

This white paper has discussed why synchronised measurements, as now practised widely in Phasor Measurement Units, are also important and useful for power quality analysis. Existing standards are somewhat problematic in this regard, especially for harmonics where, followed faithfully as to specified method, they preclude use of the appropriate techniques. Further, for all measurements, the IEC 61000-4 standards suggest that very loose time synchronisation is acceptable, and many devices presumably do not perform much better than this requirement. As we show, the requirement for accurate synchronised measurements at geographically separated locations, for applications such as those discussed here, is approximately five orders of magnitude more demanding than the standards suggest (i.e., 200 ns vs. 20 ms).

This paper has described methods developed at Arbiter Systems, Inc. to perform synchronised power quality measurements using methods similar to those used in PMUs. These methods necessarily differ from accepted international standards, but the variations are justified for these applications and do not preclude compliance with the performance requirements of the standards (only the details of the measurement methods are different).

Acknowledgements

Many thanks to Jean-Pierre Braun of METAS in Switzerland, to Gert Rietveld and Helko van den Brom, both of VSL in the Netherlands, and to Andrew Roscoe of the University of Strathclyde in Scotland, for their review and suggestions, which have helped to improve this paper significantly. Their contributions are gratefully acknowledged.