Synchrophasors:

A Primer and Practical Applications

Jonathan Sykes — **SRP**

Kristian Koellner — **SRP**

Dr. William Premerlani — **GE Global Research**

Dr. Bogdan Kasztenny — **GE Multilin**

Mark Adamiak — **GE Multilin**

Presented at the 33rd Annual
Western Protective Relay Conference
Spokane, Washington October 17-19, 2006

Synchrophasors:
A Primer and Practical Applications

Jonathan Sykes — **SRP**

Kristian Koellner — **SRP**

Dr. William Premerlani — **GE Global Research**

Dr. Bogdan Kasztenny — **GE Multilin**

Mark Adamiak — **GE Multilin**

Abstract
Although the concept and definition of Synchrophasors dates back to 1980, the combination of 2nd generation IED platforms and power system needs has brought the technology into high-visibility in the electric power industry. As synchrophasor technology has matured, nuances of the measurement of a synchronized phasor have been identified and the details of "how" a phasor is defined, synchronized to absolute time, reported, and communicated have subsequently been re-codified in the recently revised IEEE standard: Synchrophasors for Power Systems – C37.118. This paper reviews the concept of the synchronized phasor in light of the IEEE C37.118 standard. Specifically, details of an "un-corrected" Fourier based phasor during off-nominal frequency conditions are presented as well as techniques for correction to meet the Synchrophasor standard. The need for IED input transformer and filter characterization / correction is presented. Simulations of various system transients and their phasor response are presented – specifically, the response of synchrophasor calculation to dynamic system conditions. Finally, this paper reviews application of synchrophasors on the SRP system today as well as plans and needs in the future.

Introduction
As the electric power grid continues to expand and as transmission lines are pushed to their operating limits, the dynamic operation of the power system has become more of a concern and has become more difficult to accurately model. In addition, the ability to effect real-time system control is developing into a need in order to prevent wide scale cascading outages. For decades, control centers have estimated the "state" of the power system (the positive sequence voltage and angle at each network node) from measurements of the power flows through the power grid. It is very desirable to be able to "measure" the system state directly and/or augment existing estimators with additional information.

Alternating Current (AC) quantities have been analyzed for over 100 years using a construct developed by Charles Proteus Steinmetz in 1893[1] known as a phasor. In the power system, phasors were used for analyzing AC quantities assuming a constant frequency. A relatively new variant of this technique that synchronizes the calculation of a phasor to absolute time has been developed [2], which is known as synchronized phasor measurement or synchrophasors. In order to uniformly create and disseminate these synchronized measurements, several aspects of phasor creation had to be codified [3]. The following text spells out the definitions and requirements that have been established for the creation of synchronized phasor measurements.

Synchrophasor Definition
An AC waveform can be mathematically represented by the equation:

$$x(t) = X_m \cos(\omega t + \phi) \qquad \text{Eqn. 1}$$

where: X_m = magnitude of the sinusoidal waveform

$\omega = 2 * \pi * f$ where f is the instantaneous frequency

ϕ = Angular starting point for the waveform

Note that the synchrophasor is referenced to the cosine function. In a phasor notation, this waveform is typically represented as:

$$\overline{\mathbf{X}} = X_m \angle \phi$$

Since in the synchrophasor definition, correlation with the equivalent RMS quantity is desired, a scale factor of $1/\sqrt{2}$ must be applied to the magnitude which results in the phasor representation as:

$$\overline{\mathbf{X}} = \frac{X_m}{\sqrt{2}} \angle \phi$$

Adding in the absolute time mark, a synchrophasor is defined as the magnitude and angle of a cosine signal as referenced to an absolute point in time as shown in figure 1.

In figure 1, time strobes are shown as UTC Time Reference 1 and UTC Time Reference 2. At the instant that UTC Time Reference 1 occurs, there is an angle that is shown as "+θ" and, assuming a steady-state sinusoid (i.e. – constant frequency), there is a magnitude of the waveform of X_1. Similarly, at UTC Time Reference 2, an angle, with respect to the cosine wave, of "-θ" is measured along with a magnitude or X_2. The range of the measured angle is required to be reported in the range of $\pm \pi$. It should be emphasized that the synchrophasor standard focuses on steady-state signals, that is, a signal where the frequency of the waveform is constant over the period of measurement.

In the real world, the power system *seldom* operates at exactly the nominal frequency. As such, the calculation of the phase angle, θ, needs to take into account the frequency of the system at the time of measurement. For example, if the nominal frequency of operating at 59.5Hz on a 60Hz system, the period of the waveform is 16.694ms instead of 16.666ms – a difference of 0.167%.

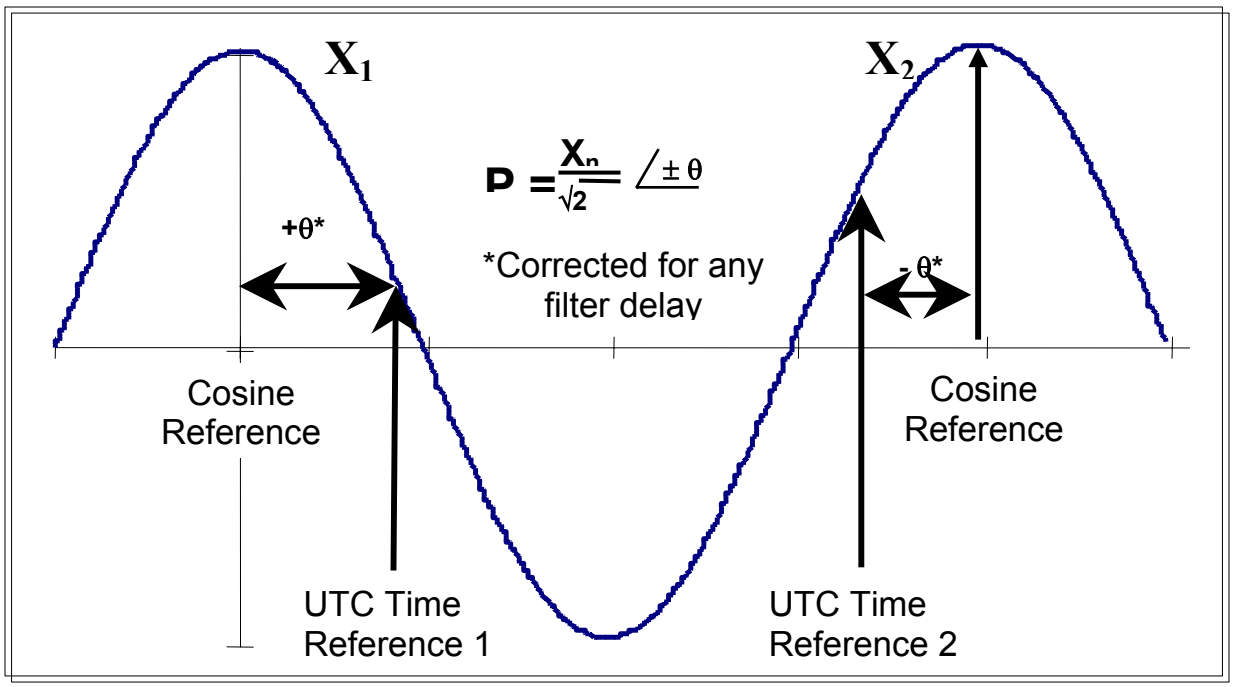

Figure 1
Synchrophasor Definition

The captured phasors are to be time tagged based on the time of the UTC Time Reference. The Time Stamp is an 8-byte message consisting a 4 byte "Second Of Century – SOC", a 3-byte Fraction of Second and a 1-byte Time

Quality indicator. The SOC time tag counts the number of seconds that have occurred since January 1, 1970 as an unsigned 32-bit Integer. With 32 bits, the SOC counter is good for 136 years or until the year 2106. With 3-bytes for the Fraction Of Second, one second can be broken down into 16, 777,216 counts or about 59.6 nsec/count. If such resolution is not desired, the C37.118 standard allows for a user-definable base over which the count will wrap (e.g. – a base of 1,000,000 would tag a phasor to the nearest microsecond). Finally, the Time Quality byte contains information about the status and relative accuracy of the source clock as well as indication of pending leap seconds and the direction (plus or minus). Note that leap seconds (plus or minus) are not included in the 4-byte Second Of Century count.

Synchronized Phasor Reporting

The IEEE C37.118 revision of the IEEE 1344 Synchrophasor standard mandates several reporting rates and reporting intervals of synchrophasor reporting. Specifically, the proposed required reporting rates are shown in Table 1 below.

Table 1
Synchrophasor Reporting Rates

System Frequency:	50 Hz		60 Hz				
Reporting Rates:	10	25	10	12	15	20	30

A given reporting rate must evenly divide a one second interval into the specified number of sub-intervals. This is illustrated in figure 2 where the reporting rate is selected as 60 phasors per second (beyond the maximum required value, which is allowed by the standard). The first reporting interval is to be at the Top of Second that is noted as reporting interval "0" in the figure. The Fraction of Second for this reporting interval must be equal to zero. The next reporting interval in the figure, labeled T_0, must be reported1/60 of a second after Top of Second – with the Fraction of Second reporting 279,620 counts on a base of 16,777,216.

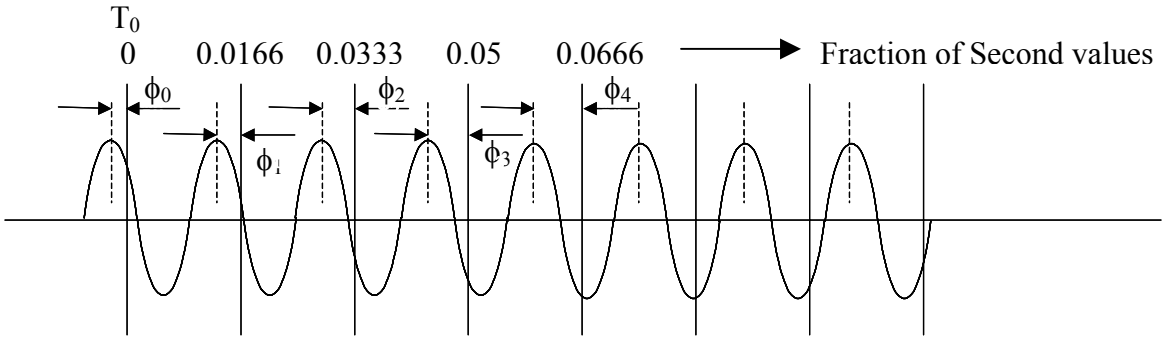

Figure 2
Synchronized Reporting Intervals

Performance Criteria

The measurement of a synchrophasor must maintain phase and magnitude accuracy over a range of operating conditions. Accuracy for the synchrophasor is measured by a value termed the Total Vector Error or TVE. TVE is defined as the square root of the difference squared between the real and imaginary parts of the theoretical actual phasor and the estimated phasor – ratioed to the magnitude of the theoretical phasor and presented in percent (equation 2).

$$\varepsilon = (\sqrt{[((X_r(n) - X_r)^2 + (X_i(n) - X_i)^2) / (X_r^2 + X_i^2)]}) * 100 \qquad \text{Eqn. 2}$$

where: X_r and X_i represent the theoretical exact synchrophasor

and: $X_r(n)$ and $X_i(n)$ represent the estimated synchrophasor

In the most demanding level of operation (Level 1), the synchrophasor standard specifies that a Phasor Measurement Unit (PMU) must maintain less that a 1% TVE under conditions of ±5 Hz off-nominal frequency, 10% Total Harmonic Distortion, and 10% out-of-band influence signal distortion. The next section examines the issues that result from implementation of the phasor measurement using the classical Discrete Fourier Transform.

Off-nominal frequency effect in a classical phasor estimator

For a classical one-cycle, N samples-per-cycle algorithm, an rms phasor is estimated by the following centered computation:

$$\hat{\mathbf{X}} = \frac{\sqrt{2}}{N} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} x\left[\Delta t\left(k+1/2\right)\right] \cdot e^{-j(k+1/2)\frac{2\pi}{N}}$$

$$\hat{\mathbf{X}} = \text{one - cycle phasor estimate} \qquad \text{Eqn. 3}$$

$$\Delta t = \frac{1}{N \cdot f_{nominal}}$$

$$x\left[\Delta t\left(k+1/2\right)\right] = \text{current or voltage sample taken at t} = \Delta t\left(k+1/2\right)$$

It is assumed that N is even. A 1/2 sample offset is used in both the sampling and the complex exponential to achieve exact centering around the on-time mark.

Suppose that the samples are from a single frequency component, not necessarily at the nominal frequency, with a phase angle measured with respect to t=0:

$$x\left[\Delta t\left(k+1/2\right)\right] = \sqrt{2}\,Real\left[\overline{\mathbf{X}} \cdot e^{j(k+1/2)\frac{2\pi}{N}\cdot\frac{f}{f_{nominal}}}\right]$$

$$\overline{\mathbf{X}} = \text{actual, "true", phasor value of the sequence of samples} \qquad \text{Eqn. 4}$$

$$f_{nominal} = \text{nominal frequency base for fixed rate sampling}$$

$$f = \text{actual frequency of the sequence of samples}$$

By substituting equations 4 into equations 1, and simplifying, it can be shown that the one-cycle phasor estimate is related to the true value as follows:

$$\hat{\mathbf{X}} = A \cdot \overline{\mathbf{X}} + B \cdot \overline{\mathbf{X}}^*$$

$$\overline{\mathbf{X}}^* = \text{complex conjugate of } \overline{\mathbf{X}}$$

$$A = \frac{\sin\left[\pi \cdot \left(\frac{f}{f_{nominal}} - 1\right)\right]}{N \cdot \sin\left[\frac{\pi}{N} \cdot \left(\frac{f}{f_{nominal}} - 1\right)\right]} \qquad \text{Eqn. 5}$$

$$B = \frac{\sin\left[\pi \cdot \left(\frac{f}{f_{nominal}} - 1\right)\right]}{N \cdot \sin\left[\frac{2\pi}{N} + \frac{\pi}{N} \cdot \left(\frac{f}{f_{nominal}} - 1\right)\right]}$$

The expressions for A and B may seem daunting, but it is possible to draw some conclusions from them. First, it can be seen that as the actual frequency approaches the nominal frequency, A approaches 1 and B approaches zero, so that the phasor estimate is exactly equal to the true phasor value:

$$\text{when } f = f_{nominal} \text{ then } A = 1, B = 0 \qquad \text{Eqn. 6}$$
$$\text{and } \hat{\mathbf{X}} = \overline{\mathbf{X}}$$

In other words, when the actual frequency is equal to the nominal, the centered window used in equations 3 produces a phasor estimate that is free from any phase or gain error.

For off-nominal frequency, equations 5 indicate there is both phase and magnitude error. As the actual frequency moves away from the nominal, A drops away from 1, and B moves away from zero, with a positive value for a frequency increase, and a negative value for a frequency decrease, introducing a distortion.

To get a better idea of what the nature of the distortion is, it is useful to recast equations 5 in terms of separate real and imaginary components:

$$\hat{\mathbf{X}} = A \cdot \overline{\mathbf{X}} + B \cdot \overline{\mathbf{X}}^*$$
$$Real\left[\hat{\mathbf{X}}\right] = \left[A + B\right] \cdot Real\left[\overline{\mathbf{X}}\right] \qquad \text{Eqn. 7}$$
$$Imag\left[\hat{\mathbf{X}}\right] = \left[A - B\right] \cdot Imag\left[\overline{\mathbf{X}}\right]$$

Equations 7 represent an ellipse. One way to see it is to recast the equations into a more recognizable form:

$$\frac{Real\left[\hat{\mathbf{X}}\right]}{\left[A + B\right]} = Real\left[\overline{\mathbf{X}}\right]$$
$$\frac{Imag\left[\hat{\mathbf{X}}\right]}{\left[A - B\right]} = Imag\left[\overline{\mathbf{X}}\right] \qquad \text{Eqn. 8}$$
$$\left[\frac{Real\left[\hat{\mathbf{X}}\right]}{\left[A + B\right]}\right]^2 + \left[\frac{Imag\left[\hat{\mathbf{X}}\right]}{\left[A - B\right]}\right]^2 = \left[Real\left[\overline{\mathbf{X}}\right]\right]^2 + \left[Imag\left[\overline{\mathbf{X}}\right]\right]^2 = \left|\overline{\mathbf{X}}\right|^2$$

Thus it can be seen that the locus of all phasor estimates, for a given phasor amplitude and variable phase angle, is an ellipse with the major and minor axes aligned with the axes of the complex plain. The real and imaginary intercepts of the ellipse are given by:

$$RealIntercept = \left[A + B\right] \cdot \left|\overline{\mathbf{X}}\right|$$
$$ImagIntercept = \left[A - B\right] \cdot \left|\overline{\mathbf{X}}\right| \qquad \text{Eqn. 9}$$

As can be seen from equations 3 and 7, the eccentricity of the ellipse worsens as the frequency deviation grows. It can also be seen for small positive frequency deviations that the major axis of the ellipse is aligned with real axis. For small negative frequency deviations the major axis of the ellipse is aligned with the imaginary axis. Further examination of equations 5 and 9 also reveals that the magnitude of the phasor estimate "shrinks", because A is always less than 1 for off-nominal operation. Figure 3 shows the outline of the rotating phasor for a 65Hz input signal.

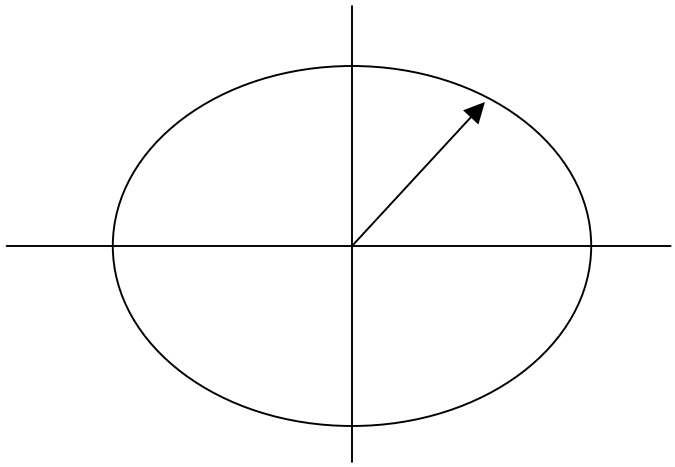

Figure 3
Classic Fourier Elliptical Response at 65Hz Input

A four parameter phasor model

One way of representing raw current and voltage signals that from which we are trying to compute phasors is in the form of the real part of Taylor's expansion according to (10).

$$x(t) \approx \sqrt{2} \cdot Real\left(\left(\overline{\mathbf{X}} + \dot{\overline{\mathbf{X}}} \cdot t\right) \cdot e^{j2\pi \cdot f \cdot t}\right)$$
$$x(t) = \text{instantaneous current or voltage}$$
$$t = \text{time}$$
$$\overline{\mathbf{X}} = \text{constant portion of the phasor} \qquad \text{Eqn. (10)}$$
$$\dot{\overline{\mathbf{X}}} = \text{first derivative of the phasor}$$
$$f = \text{nominal power system frequency}$$

We call this a four-parameter model, because the specification of the constant portion of the phasor and its first derivative with respect to time requires a total of four real parameters to specify the model. In this model, the phasor is meant to be a stationary value. If the actual power system frequency is equal to the nominal value and the power system is in steady state, $\overline{\mathbf{X}}$ is a stationary value, and $\dot{\overline{\mathbf{X}}}$ is identically equal to zero. The goal is to accurately determine $\overline{\mathbf{X}}$ in the face of non-zero values of $\dot{\overline{\mathbf{X}}}$.

It is also possible to represent a time varying phasor with a Taylor's expansion in polar coordinates rather than Cartesian coordinates. While this would be a better fit for the off-nominal frequency situation, the mathematics would become more complicated. For a one-cycle window, there would not be much difference between an

expansion in polar coordinates and an expansion in Cartesian coordinates.

Equation (10) is applied over a limited time window in which the expansion is approximately true, on the order of a few cycles, in which the frequency and other parameters of the representation can be considered constant. Of course, the same model can be applied in a piecewise fashion over all cycles, by selecting the appropriate parameters for each cycle.

It is convenient to express (10) in terms of phasors and their complex conjugates to eliminate the *Real* operator as follows:

$$x(t) \approx \frac{\sqrt{2}}{2}\left(\left(\overline{\mathbf{X}} + \dot{\overline{\mathbf{X}}} \cdot t\right) \cdot e^{j2\pi \cdot f \cdot t}\right) + \frac{\sqrt{2}}{2}\left(\left(\overline{\mathbf{X}}^* + \dot{\overline{\mathbf{X}}}^* \cdot t\right) \cdot e^{-j2\pi \cdot f \cdot t}\right) \qquad \text{Eqn. (11)}$$

The derivative term in (10) or (11) can be used to approximate either slightly off-nominal frequency or power swings. In the case of slightly off-nominal frequency operation, there is a phasor derivative that is perpendicular to the phasor. In the case of a power swing, there is a phasor derivative that is parallel to the phasor.

Before proceeding to a four-parameter solution, it is useful to examine the error introduced by a classical computation when currents and voltages are given by (11). Suppose a phasor is computed from samples of current and voltage with a variation of the classical algorithm, using the one-cycle, N-sample, centered "Boxcar" algorithm specified by (12).

$$\overline{\mathbf{Y}} = \frac{\sqrt{2}}{N} \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} x[n] \cdot e^{-j(n+1/2)\frac{2\pi}{N}} \qquad \text{Eqn. (12)}$$

$$x[n] = x\left(\frac{n+1/2}{N \cdot f}\right)$$

We prefer the use of window that is centered on the time reference for the phasor, because some types of errors exactly cancel that way. N is assumed to be even. The combination of even N and the centered window leads to a shift in the sampling time equal to ½ of the sampling period.

Substituting (11) into (12), it can be shown that the computed one-cycle phasor produced by (12) for the signal model given by (11) is given exactly by (13).

$$\overline{\mathbf{Y}} = \overline{\mathbf{X}} + j \cdot \frac{\dot{\overline{\mathbf{X}}}^*}{2N \cdot f \cdot \sin\left(\frac{2\pi}{N}\right)} \qquad \text{Eqn. (13)}$$

The detailed steps in deriving (13) are not given here because of space limitations. The derivation is not difficult. The principal challenge is reducing sums of powers of complex exponential expressions to simpler forms

Under steady-state conditions with the actual power system frequency exactly matching the nominal value, the derivative term vanishes, and (13) indicates that the classical algorithm recovers the correct value, $\overline{\mathbf{X}}$. However, if there is an off-nominal frequency condition, or a power swing, or other transient condition, there is a non-vanishing first derivative of the phasor, and there is a backward rotating error phasor given by the second term in (13). If we could somehow determine the actual value of the derivative, we could compensate for the error.

It can be shown that a positive sequence phasor computed from per-phase phasors computed by equation (12) is related to positive and negative sequence quantities as follows:

$$\overline{\mathbf{Y}}_{pos} = \overline{\mathbf{X}}_{pos} + j \cdot \frac{\dot{\overline{\mathbf{X}}}^{*}_{neg}}{2N \cdot f \cdot \sin\left(\frac{2\pi}{N}\right)} \qquad \text{Eqn. (14)}$$

This means that if the negative sequence component is zero, and if the positive sequence phasor moves in a straight-line trajectory in the complex plane, the classical algorithm will produce the correct value for the positive sequence phasor. For the case of slightly off-nominal frequency operation or gradual power swings, the trajectory can be approximated as a gradually changing straight-line trajectory, so that we can conclude that a classical algorithm will produce an approximately correct estimate of positive sequence quantities, provided that the negative sequence component is zero. If there is a negative sequence component, (14) indicates that off-nominal frequency operation will produce an error in positive sequence quantities computed using the classical algorithm.

Before proceeding to developing an algorithm for correcting the error in the classical algorithm that is generated by a changing phasor, we will need an estimate of the time derivative of a phasor. One approach is to examine the sequence of phasors computed by the classical algorithm. Examination of (13) reveals that the error in a sequence of computed phasors is the same for each phasor in the sequence, provided the derivative of the sequence is approximately constant. Therefore, the derivative can be estimated simply by taking the differences of pairs of computed phasors in the sequence.

For example, suppose that a phasor value is computed once for each power system cycle, at the nominal, fixed, power system frequency. Furthermore, suppose that the phasor derivative over a few cycles is approximately constant. Then, the phasor derivative at cycle M can be estimated from the classically computed phasors at cycle M and cycle M-1 using (15). Note that it is convenient to compute the ratio of the derivative to the nominal power system frequency.

$$\frac{\dot{\overline{\mathbf{X}}}_{(M)}}{f} \approx \overline{\mathbf{Y}}_{(M)} - \overline{\mathbf{Y}}_{(M-1)} \qquad \text{Eqn. (15)}$$

Equation (15) is exact for a phasor trajectory that is a straight line in the complex plane. It is approximately true for off-nominal frequency operation when the classically computed phasors trace an elliptical trajectory in the complex plane.

By substituting (15) into (13) and rearranging, it is possible to derive (16), which is a simple formula for the corrected estimates of phasors, starting with the classical algorithm and a "Boxcar" window.

$$\overline{\mathbf{X}}_{(M)} \approx \overline{\mathbf{Y}}_{(M)} - j \cdot \frac{\left(\overline{\mathbf{Y}}^{*}_{(M)} - \overline{\mathbf{Y}}^{*}_{(M-1)}\right)}{2N \sin\left(\frac{2\pi}{N}\right)} \qquad \text{Eqn. (16)}$$

Equation (16) corrects classically computed phasors for the error generated by the derivative of the phasors, provided the phasors are computed using (12). It is also possible to derive similar formulas for other variations of the classical phasor computation by first computing the error introduced by the derivative of the phasor.

Equation (16) is not difficult to implement in practice. The constant denominator of the second term can be pre-computed, and the division can be replaced by a multiplication by the reciprocal of a constant. The effect of combining the multiplication by j with the conjugate operator can be achieved by suitable sign changes and the swapping of real and imaginary components.

The accuracy of (16) will depend on how well the actual conditions are described by the assumed Taylor's expansion. For slight frequency deviations and for power swings, a first order Taylor's expansion is a reasonable approximation.

It is instructive to examine the numerical results of two simulations, figures A-D. In each case both the classical one-cycle algorithm and the four parameter model were used to compute single-phase phasors for various phasor trajectories. Nominal power system frequency is 60 Hz, sampling is 64 samples per cycle. In each case, a centered window is used. Phasors are plotted for every sample and are rotated so they are stationary for a steady 60 Hz waveform.

The response to a cosine ramp is shown in figures 4 - 6. The real part of the phasor swings as a cosine function from -1.0 to +1.0 over 10 cycles. The imaginary part stays at zero. An example of a situation that would approximate this case is the current on a transmission line for which the voltages at opposite end of the line swing in opposite directions. The power transmitted by the line falls to zero and reverses. The real and imaginary components for three stages in the calculation are shown in figures 4 and 5. The actual phasor is given in dashed black lines. The phasor computed by the classical algorithm is shown in blue. The result produced by the four-parameter algorithm is shown in green. It is difficult to distinguish between the three plots in Figure 4, because they fall close to each other. It is easier to see the results in Figure 5. The distances between the computed phasors and the true values are plotted in Figure 6. The classical algorithm suffers a peak total error of 2.5% while the four-parameter algorithm is better than 1%. At around 100 milliseconds, when the ramp rate is greatest, the error of the four-parameter algorithm is close to zero.

An off nominal frequency of 58 Hz is shown in Figure 7. The performance of the four-parameter algorithm is much better than the classical algorithm, less than 0.4% as compared to more than 1.8%.

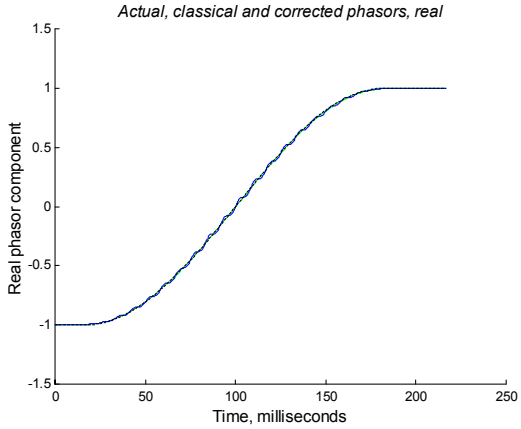

Figure 4
Real Response to a Cosine Ramp

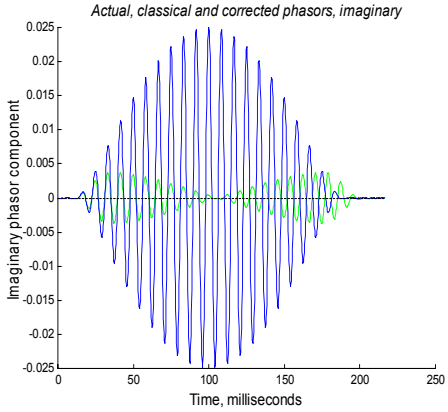

Figure 5
Imaginary Response to a Cosine Ramp

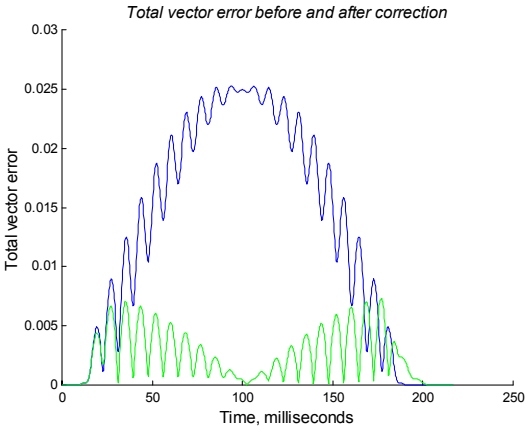

Figure 6
Vector Error During a Cosine Ramp

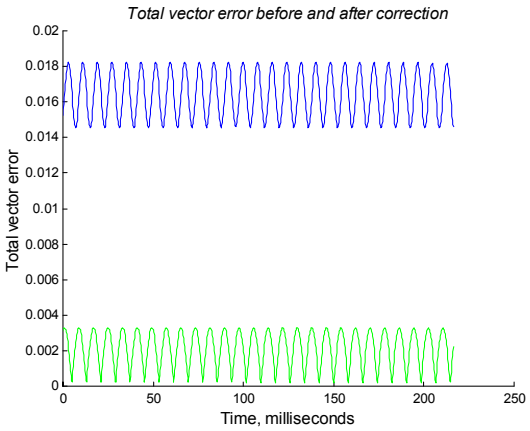

Figure 7
58 Hz Vector Errors –
Before (Blue) and After (Grn) Correction

Implementation of synchrophasors on P&C IED platforms

It is self evident that wide penetration of PMUs facilitating both faster gathering of field data in preparation for advanced applications, and redundancy of measurements required for the future critical applications of synchrophasors, can be naturally achieved by integrating PMU functions with protection and control platforms. Successful integration of SOE and DFR capabilities with protective relays is a historical lesson to follow when considering cost-efficient and universal deployment of PMUs.

Modern protection platforms are capable of supporting synchrophasor measurements, local recording and reporting. This relates to internal architectures, time synchronization, metering accuracy, communication capabilities, and processing power required to comply with the C37.118 requirements.

However, microprocessor-based protection relays have been designed so far without regard to the notion of absolute time. Time stamping for SOE and DFR recording is probably the only instance of reference to an absolute time in protective relaying. Sampling and synchronization, even in critical and high performance systems such as the line current differential protection, is typically achieved without reference to the absolute time. This is a prudent protection approach as it limits exposure of mission critical protection functions to availability and misbehavior of other devices.

As a rule protection relays sample asynchronously with respect to the absolute time, but in sych with power system frequency. The latter is to keep the digitally implemented measurements accurate should the power frequency depart from its nominal value.

This section provides some insight on implementation of synchrophasors on a typical microprocessor-based relay. It presents some solutions, and highlights certain aspects that need to be understood and evaluated by a protection engineer to make sure the extra functionality put on a relay does not jeopardize the core protection task of the device.

Design principles when implementing synchrophasors on existing platforms

It is prudent to follow these design principles when implementing synchrophasor measurements on existing or new protection platforms:

1. The underlying sampling process of the relay shall not be altered. Sampling and data collection potentially affects all other functions of the relay. To minimize the risk, this area shall not be modified. Sampling in synchronism with the absolute time is not only unnecessary; it actually yields a substandard solution from the point of view of metering accuracy.
2. The synchrophasor calculations shall be added in parallel to the existing protection, control and metering functions to minimize the risk of affecting these critical functions.
3. Hardware modifications shall be minimized for the reason of stability of the design.
4. Calculations shall be organized in a way that the extra processing power is optimally distributed and can be accommodated by existing platforms with appropriate security margin, even under fault conditions and other periods of increased activity.

The key design areas for implementation are: timing accuracy; sampling and correlating input signals with the absolute time, algorithms for accurate measurement of the phasors, data storage, recording and streaming. The following subsections address these fundamental implementation issues.

Timing accuracy

The Total Vector Error concerning synchrophasors has three major components: magnitude error, angle error as related to the input signals, and angle error as related to the measurement of the absolute time.

It is enlightening to think of time as a quantity that needs to be “measured” by a given device based on a standard physical input, such as the 1 pulse per second (1 pps) marker embedded in the standard IRIG-B input. Assuming 1% TVE target, and budgeting accordingly for the three sources of error, leaves up to 5-8 microseconds for the total timing error.

Not only does a given device needs to sych with the 1pps signal but the device must internally maintain a very precise notion of time so that each of the synchrophasor reference points (synchrophasor interrupts) occurring within the period of the full second is maintained with an error not larger than few microseconds.

Figure 8 below illustrates this process. In one particular implementation a precise phase lock loop is run with the objective to null out a positional error between the 1pps signal and the last synchrophasor interrupt that ought to occur exactly at the top of the second. This phase lock loop compensates for the natural drift of the internal IED oscillator, and a finite resolution of the latter. For example, a given oscillator could have an error of say 25 parts per million. This means it could drift up to 25 microseconds over a period of 1,000,000 microseconds (1 second). Such values would prevent successful implementation. Moreover, the error can change with temperature and between different samples of the oscillator. The drift, however, is easily measurable with the aid of the 1pps signal. When measured, validated, and averaged, the drift of the oscillator is an input to the phase lock loop making the internal time keeping mechanism extremely accurate.

Another issue is the required resolution of the internal crystal. Assume 60 synchrophasors are produced per second. If so, the synchrophasor interrupts are to be generated every 1,666.66(6) microseconds. When this number is rounded to a practical oscillating frequency, an error would accumulate making the last synchrophasor interrupts in a given second inaccurate. In addition, assume the oscillator is too slow by 12 microseconds in each second (example). To compensate for the drift each synchrophasor interrupt will have to be adjusted by 12/60 = 0.2 microsecond, while a practical resolution of the oscillator can be in the range of a quarter of a microsecond. The 0.25 – 0.20 = 0.05 microsecond error repeated 60 times would yield 3 microseconds of error over 1 second - eating away from the tight error budget.

A dithering algorithm is applied to achieve high accuracy of timing the synchrophasor interrupts. An internal variable is used to count the time with nanosecond accuracy, while the interrupts are generated with a 0.25 microsecond resolution. The device keeps track of the error accumulated due to the finite resolution of the oscillator. Once the error reaches half the oscillation period, the synchrophasor interrupt is moved by one oscillation period. In this way the error is kept below half the period of the oscillator, and never accumulates.

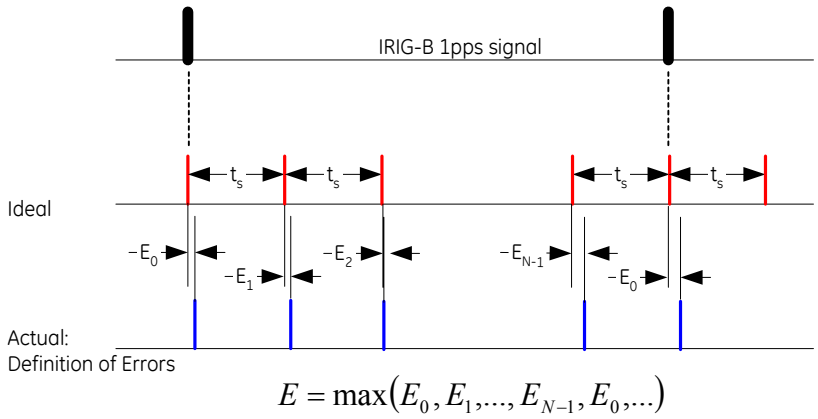

Figure 8. Defining synchrophasor interrupts and timing errors.

Sampling for protection and synchrophasors

Protective relays typically do not sample synchronously with respect to the absolute time. Instead they sample asynchronously and often apply frequency tracking or compensation so that the measurements retain accuracy even if the system frequency departs from the nominal value. It is a common misconception that calculating synchrophasors requires sampling in accordance with an absolute clock.

Another misconception is that staying in synchronism with the system frequency (for accuracy) and staying in synchronism with the absolute time (for phase reference) are contradicting targets, and require convoluted solutions such as measuring the magnitude and angle using different algorithms. The former is about adjusting the length of the data window so that it covers pre-selected multiplies of power cycles; the latter is about positioning of this data window so that the measurement complies with the C37.118 angle convention. Both can be controlled independently with no major obstacles.

Moreover, the samples, referenced to absolute time, can be taken at any time instant. Figure 9 presents a solution in which the samples are collected asynchronously with respect to absolute time. The platform applies frequency

tracking to keep the number of samples constant in the actual period of the waveform as the period changes. When the synchrophasor interrupt is asserted, the device locks the sample index and collects half its data window from the samples that follow the interrupt and half – from the samples preceding the interrupt. In this way, without altering the sampling process, the device gets a data window that is placed very closely with respect to the ideal position.

Note that in this approach:

- The length of the data window is already correct and adequate as the sampling period is controlled by the frequency tracking mechanism;
- The position of the window is within half of the sampling period from the required position as per the synchrophasor convention.

The device calculates the center of the window by averaging the time stamps of the samples within the window. This averaging is done using any time reference, not necessarily the absolute time reference. In one implementation a free running microsecond counter is used to calculate the position of the center of the data window. The same free running counter is used to capture the time of the synchrophasor interrupt asserted based on the true time. Even though the free running microsecond counter is not a true time, the time difference between the synchrophasor interrupt (point when the center of the window should be), and the calculated center of the window (point when the data window actually is) is meaningful and can be used for compensation.

Again, following the window selection procedure illustrated in Figure 9 one places the window to within few degrees to the synchrophasor interrupt. The inherent displacement is precisely measured and is used for very precise compensation of the calculated phasor (a plain shift by 2-3 degrees).

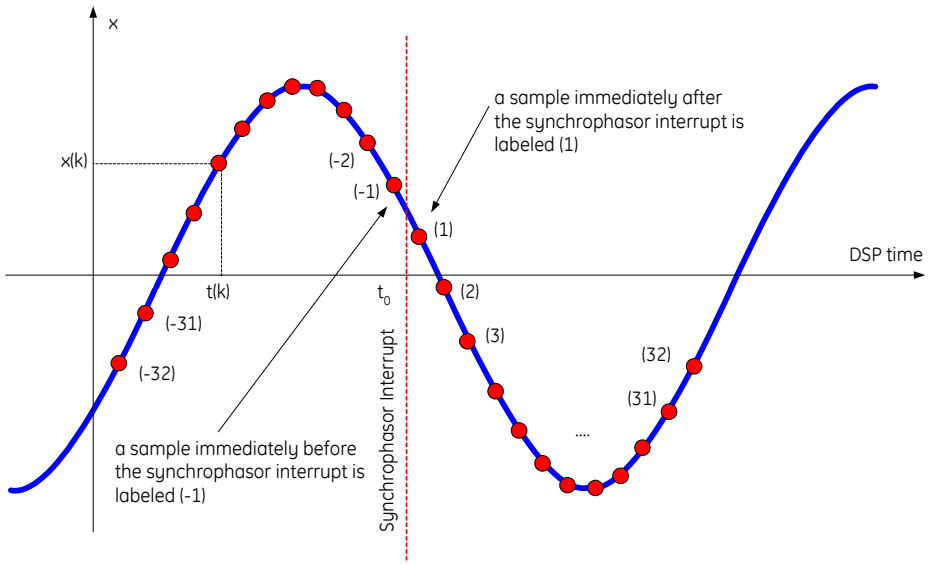

Figure 9. Data window based on asynchronously taken samples.

The method of correlating asynchronously taken samples with the absolute time depicted in Figure 9, can be better explained by looking at Figure 10. In this architecture the CPU synchronizes to the 1pps signal and executes the phase lock loop that generates precise synchrophasor interrupts. These interrupts are captured by a DSP using a "local DSP time" in the form of a free running counter. The interrupt triggers calculations for the synchrophasor instant and allows the DSP to obtain the notion of time, and produce the phasor precisely aligned with the time mark as driven by the interrupt.

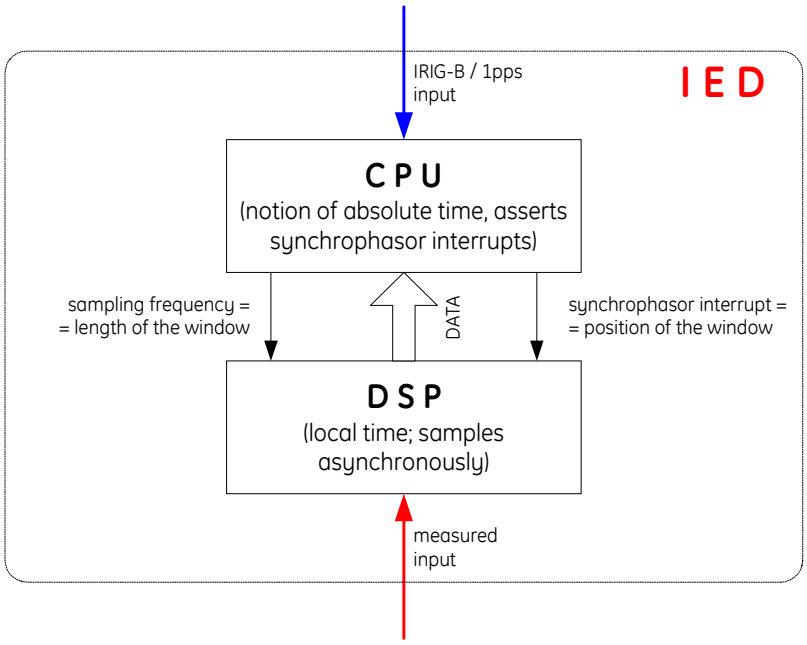

Figure 10. CPU & DSP architecture for synchrophasor implementation.

Post-processing and extra filtering

As depicted in Figure 11, the device uses "best-placed" windows for synchrophasor measurement without altering the sampling process (windows X). It measures the small shift between the required and actual positions of such windows and compensates for the difference by a simple phasor rotation. This yields synchronized full-cycle Fourier windows (windows Y).

The X and Y windows are produced at nominal system frequency regardless of the recording or reporting rates set for the PMU function. The pair of Y windows (the present and past windows) is used to implement the four-parameter signal estimator previously described. As a result, a new, more accurate estimate of the phasor is calculated at the rate of nominal system frequency (windows P in Figure 11). The P-values are calculated assuming the phasor may change in time, and as such are extension of the C37.118 synchrophasor standard, aimed at future dynamic applications of synchrophasors.

In order to control the balance between speed and accuracy of the measurement, the device further implements a user controllable post-filtering: a number of P-measurements can be combined into the filtered synchrophasor output, S, effectively extending the estimation window. The post-filtering is not a straight average, but takes into account the value and rotation speed of each of the used P-values.

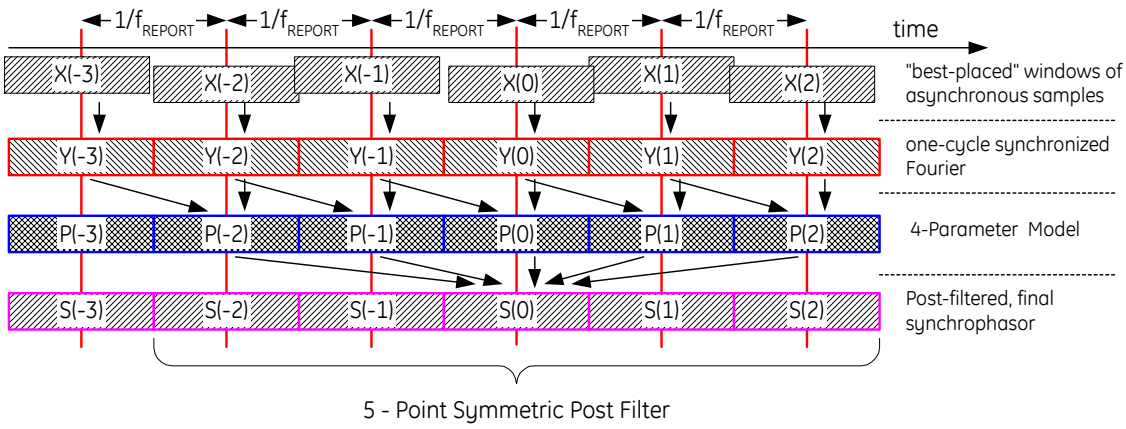

Figure 11. Processing of best-placed raw data windows into synchrophasor values.

Compensating for analog errors

Synchrophasor implementation calls for accuracy above a typical protection accuracy or metering accuracy as provided on protective relays. When implemented on a protection platform, synchrophasors may need correcting for errors of input transformers and anti-aliasing filters.

Figure 12 below shows a correcting function for the current inputs: the correction depends on both the magnitude and frequency of the signal. In particular at very low signal levels and lower frequencies the excitation current of the input transformers starts causing some angular errors, and the device applies higher correction for the measured angle for the current inputs.

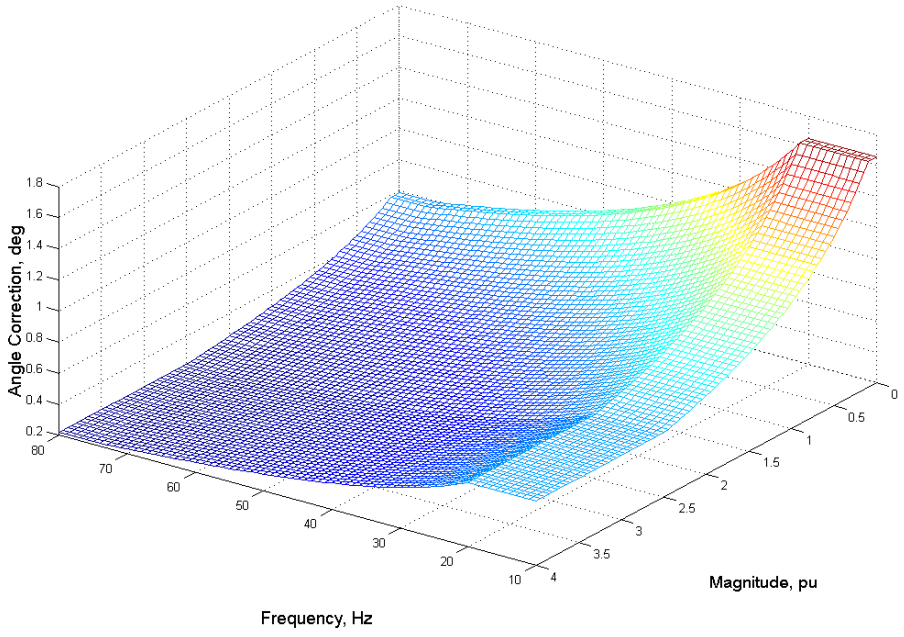

Figure 12. Correction of current input transformers.

Figure 13 below shows the correction applied to the voltage inputs. The required angle shift to keep the measurements accurate is smaller, and again depends on the magnitude and frequency of a given voltage input.

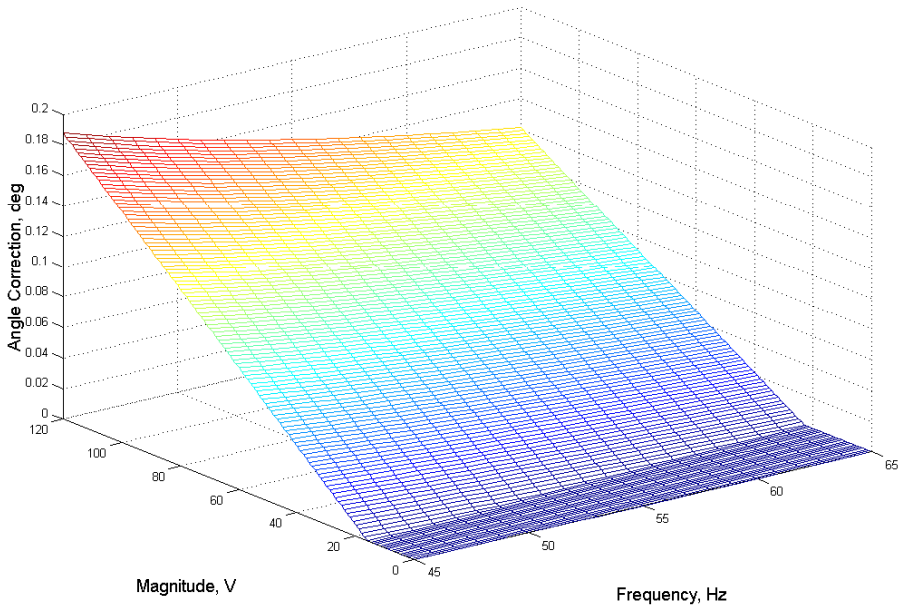

Figure 13. Correction of voltage input transformers.

Analog filters, necessary in any digital measuring system, introduce a phase shift too which need to be compensated in the end result. When the analog filter is set relatively high, the phase shift for the frequency band around the nominal is very linear, and can be easily compensated. Figure 14 below shows the measured (red dots) and applied (blue line) correcting angles accounting for the impact of analog filters.

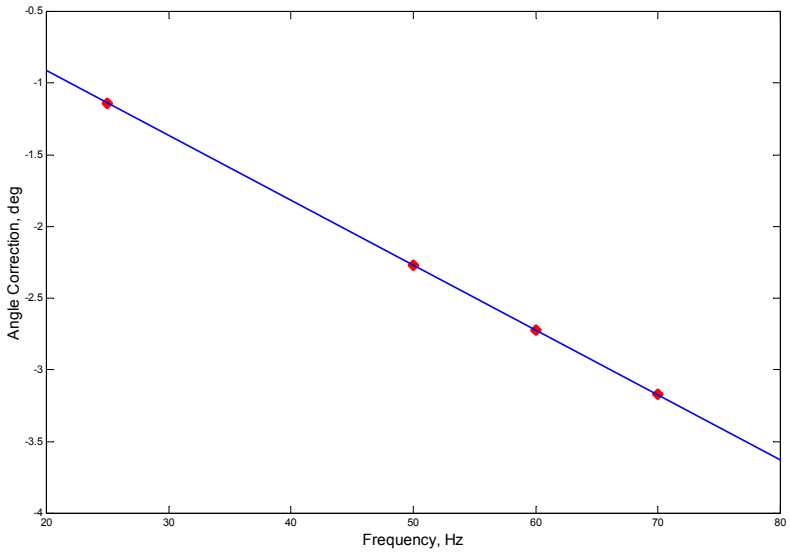

Figure 14. Correction for the Analog Filter.

The few implementation details outlined above are meant to direct attention to the way synchrophasors are implemented on protective relays, and the potential impact on the existing mission-critical functionality. In the outlined implementation, minor hardware changes were required to provide a synchrophasor interrupt from the central processing unit having the notion of the absolute time (IRIG-B input) to the digital signal processor responsible for the majority of the calculations but having no direct relationship with the absolute time. All the other

aspects of the synchrophasor implementation have been accommodated in software, in subroutines completely detached from the key protection functions. This minimizes the risk and allows claiming a very secure implementation.

The processing power required to provide for the PMU function, even when reporting at the rate of 60 phasors a second, is at the level of one zone of distance protection, thus very moderate.

Accuracy

Having implemented the process described above, the following summarizes the steady state performance as tested on the IED hardware:

- TVE for voltages, frequency range 45-70Hz < 0.30%
- TVE for currents, frequency range 45-70Hz < 0.40%
- TVE at 10% of THD, nominal frequency < 0.45%

Figure 15 presents results of the interfering frequency test when reporting at 60 times per second, and using an optional 7-point post-filtering algorithm.

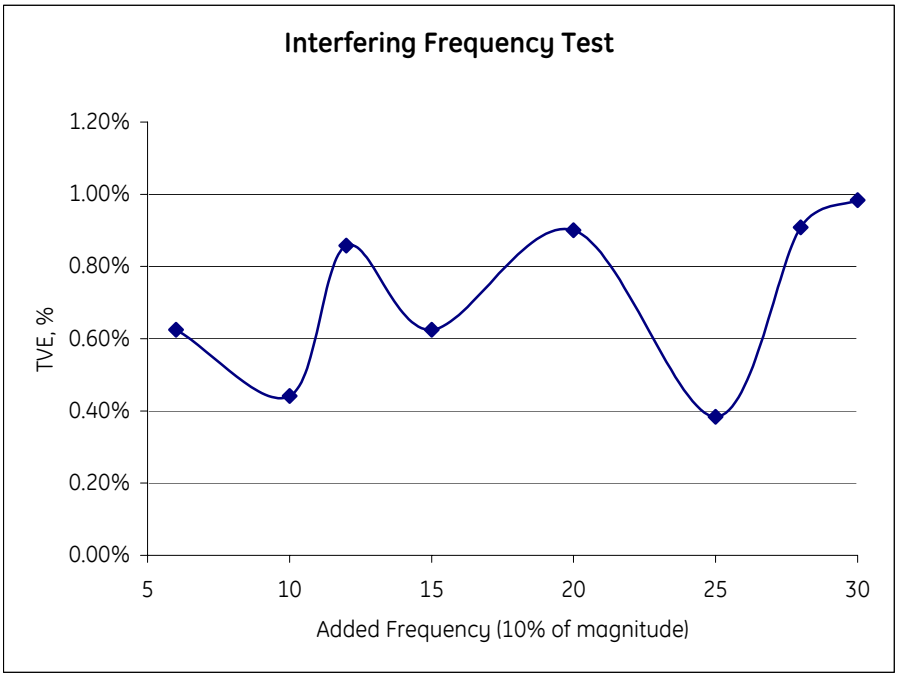

Figure 15. TVE under interfering frequency tests (reporting at 60/second, 7-point post-filter applied).

The described implementation details, and the test results demonstrate that when carefully engineered, modern P&C platforms allow for both secure and accurate implementation of synchrophasor measurement, recording, and reporting. When integrated with protection platforms the PMU functionality is provided universally with wide coverage of the metering points, at a fraction of the cost of legacy stand-alone PMU solutions.

Business Case and Roles for Implementing a Synchrophasor System

In the beginning, building a business case at SRP (Salt River Project) for the implementation for Synchrophasors was a challenge - even though SRP is a company that focuses on new-technology. Technologies that could immediately benefit the bottom line or address a definite need are more easily implemented and herein lies the predicament: how to implement a technology that does not have an immediate positive impact on the budget and does not have a definite need.

A "pull" for Synchrophasors has come from the Western Electric Coordinating Council (WECC) - a Regional Reliability Organization (RRO) in the western USA, Canada, and Mexico. WECC has a program that requires the installation of Phasor Measurement Units (PMUs) at strategic locations throughout the WECC territory. Because of the geographic enormity of the area, the PMUs are implemented to monitor WECC-wide disturbances and generally don't have enough visibility for local disturbances. These units are installed on power corridors or large generation centers. SRP has PMUs installed at three sites per the WECC requirements. The data from the PMUs plays an important role in the analysis of problems and verification of models today, however PMUs clearly are tools that can be used to "detect" and "react" to an event and ultimately avoid system outages.

In early 2005, SRP realized the future potential of the synchrophasor and created a cross-departmental Synchrophasor Team. This team has been investigating all the issues including application and implementation of Synchrophasors at SRP and has been monitoring the efforts of other RROs and the industry. Initial funding focused on testing of the technology as well as investigation of available applications and monitoring of new developments.

As experience with the technology increased, a short-term goal of the team was to established which was to implement a working system. The first application of synchrophasors occurred in support of a Black Start exercise in April of 2005 (discussed in the application section). Additionally, goals were established to bring together data from various manufacturers, archive the data in a central location, and bring more devices on line. The team plans to address the exchange of data with the WECC PMU system and develop "real" applications.

Implementation of the synchrophasors involves many different groups both internal and external to SRP. The monitoring devices are relays and that means the installation of the devices will be the responsibility of the System Protection department and the Relay Shops. The Communication Engineering department will provide the necessary path for the data to get from the various remote sites to the centralized Phasor Data Concentrator (PDC). The Control Engineering and Computer Applications groups at SRP will provide the interface between the synchrophasor system and the Energy Management System (EMS). Computer Applications will also write the algorithms to include the data in the State Estimator and the operator consoles once the various applications are realized.

The operators were the first group at SRP to utilize synchrophasor data during the Black Start exercise. When the applications and implementation matures the Transmission and Distribution Planning groups intend to use the data to verify models and determine reactive margins. All groups within SRP and the RRO will be able to use the data to perform post event performance analysis including analysis of line voltage and reactive power support.

Synchrophasor System Configuration and Architecture

The components of a synchrophasor system include the Phasor Measurement Unit (PMU), an absolute time source (typically a GPS based clock), a communication network, a Phasor Data Concentrator (PDC) with storage and the user applications (figure z1). The first piece, the PMU, can take many forms and the PMUs on a given system may be of many different vintages. Older PMUs are generally stand-alone devices, designed for that dedicated purpose. They may communicate over leased-line modem connections or other low-bandwidth mediums. Those units designed and built before the IEEE Std C37.118-2005 "Standard for Synchrophasors for Power Systems", may support the older IEEE P1344-1995 format or a proprietary, vendor-specific format. Older PMU devices were a significant investment because they required a dedicated infrastructure to support the synchrophasor data retrieval.

Today, as a result of the C37.118 standard, many different vendors can support this functionality in a wide range of products. In a post C37.118 world, these products stand a much higher likelihood of being integrated together. Today's PMUs can now be imbedded within microprocessor-based protective relays. This offers tremendous cost savings because protective relays are already installed in key locations on the grid. All that needs to be done is to

add a "hook" into the relay to retrieve the data. Compared with justifying, designing, and installing a stand-alone device, this is a relatively easy task. This allows for rapid penetration of PMUs on the electric grid and offers data gathering possibilities that have not been feasible in the past. Table 2 below compares the "Pros" and "Cons" of relay-based PMUs with stand-alone PMUs:

Table 2
Pros and Cons of PMU Systems

Relay-based		**Stand-alone**	
Pros	Cons	Pros	Cons
❑ Already in place at key locations on the grid ❑ Low cost implementation ❑ C37.118 support enables vendor inter-operability	❑ Adding PMU capability may require setting changes ❑ May require relay firmware/HW upgrades ❑ CT and PT accuracy questions	❑ Dedicated functionality ❑ Has no impact on relaying system ❑ Can utilize metering accuracy CTs and PTs	❑ More Costly ❑ May be more difficult to justify funding

As indicated above, relay-based PMUs have a significant cost advantage compared to a stand-alone PMU. The primary drawback to a relay-based PMU is the fact that setting changes and/or firmware / HW upgrades may be required. Anyone who has ever issued relay-setting changes or performed a firmware upgrade on a microprocessor-based protective relay can attest to the fact that it is not always a simple task. Performing a firmware upgrade or a hardware component change-out to enable PMU capability may require a new relay checkout and/or signal injection on the protected component. Accomplishing this will often require engineers and technicians at both ends of the line and could involve a somewhat larger time and manpower investment compared to a stand-alone PMU. In as much as the relay-based PMU connects to trip circuits, there is a slight risk of inadvertent tripping activity. For this reason, the evolution of a synchrophasor system must be planned carefully so as to minimize the number of times that the relay-based PMU must be maintained.

Synchrophasor System Deployment

Given a PMU, the question arises as to "where" to locate it. The deployment of the individual PMUs has been given some thought – specifically, in response to providing "complete" visibility of the "state" of the power system. Studies to date indicate that locating PMU's throughout about 1/3 of all stations in a system can provide complete system visibility. This is achieved by not only measuring the local positive sequence voltage but also by measuring the positive sequence currents on all the exiting lines and estimating the remote substation bus voltage. The Eastern Interconnect Phasor Project – EIPP – has focused on these issues and has created a taskforce to develop placement criteria. In general, any utility, given a range of constraints, will choose from one of the following placement criteria:

"N"

This is the simplest deployment strategy out there. Simply put, a number of units "N" is chosen for deployment and those units are distributed in a virtually random fashion. "N" could be 10, 20, 50, or 100 units. This approach is very low-tech but can be useful when just getting started. This approach is particularly useful when you need to establish some momentum in the early stages of your program and show progress to management. For example: "In order to get the ball rolling, we are going to install ten PMUs on our 230kV and 500kV systems before the end of the year."

"Geographic"

The "Geographic" strategy is just slightly more advanced than "N". It takes into account the fact that utility systems are usually distributed over a wide geographic area. Without giving consideration to the actual electrical system connections, one can easily state that we want "X" PMUs in each geographic region of the system – North, South, East, West, etc. This improves upon "N" by at least recognizing that you may need to space your PMUs across your system in an orderly fashion. For example: "We need three PMUs on every major transmission corridor with over 500 MW of flow."

"Whatever makes it easier"
This strategy can advance your PMU deployment dramatically because you are leveraging some element of the installation that is now easier to accomplish. For example, you can activate the PMU capability in every relay of a certain firmware level at every station or you can add PMUs at every station that already has Ethernet communications. By reducing the barriers to deployment, you will ultimately be able to deploy more units. For example: "We will initially only install PMUs at stations with Ethernet capability because that will minimize our communication installation costs."

"Event based"
This approach looks at what it would take to adequately capture or visualize events of a certain nature. Examples of events are the tripping of a generating unit or a high-voltage line, or monitoring of a congested transmission corridor. The PMU sampling rate will then have to be justified in terms of your desire to capture those events that are important.

"Statistical analysis"
This is the most complete strategy that looks to dynamically measure the system state. This approach takes into account the system topology and looks at what samples would be required to directly measure or infer the total system state. This mode of placement is typically driven by a system operator that has a strategy to migrate to the "state measurement" concept. This strategy needs to take into account the overall system availability requirements and to subsequently plan for redundant measurements in order to meet the stated requirement.

"New construction driven"
Moving forward, this is the natural deployment strategy. At this point the value of the PMU installation has become well understood and well justified and it is now cost-effective to include the PMU deployment as part of your design and construction standards. PMU technology can be specified in certain relaying packages or substations of a certain voltage level. For example: "PMU capability shall be built into every 69kV, 230kV, and 500kV line relaying package."

The progression of PMU deployment philosophies is shown in the figure 16 below:

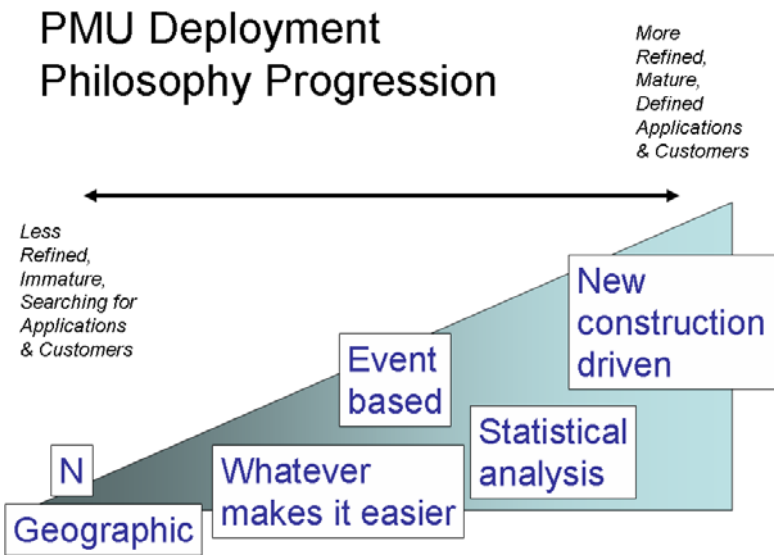

Figure 16
PMU Deployment Philosophy

SRP's approach has been a mixture of all of the above techniques, although we are quickly moving to a new construction driven (standards) approach. This approach has been driven by our belief that the technology will be needed "down the road" and the low cost-adder of relay-based PMUs.

At SRP we currently have a mix of older, stand-alone PMUs at key receiving stations along with newer, relay-based PMUs at other substations. The older PMUs communicate over a dedicated communication channel while the newer units transmit synchrophasor data via TCP/IP or UDP/IP protocols over Ethernet. The older PMUs were installed primarily for WECC disturbance monitoring reasons and are associated with a Phasor Data Concentrator (PDC) at a neighboring utility. The newer units transmit data directly to an SRP PDC. Most of SRP's recent efforts have been focused on expanding our installed base of relay-based PMUs. It has been found that by enabling the PMU capability of certain relays already installed on our 500kV and 230kV transmission systems, adequate coverage can be obtained. The figure 17 below shows where relay-based PMUs are already in-service on SRP's transmission system along with future planned locations.

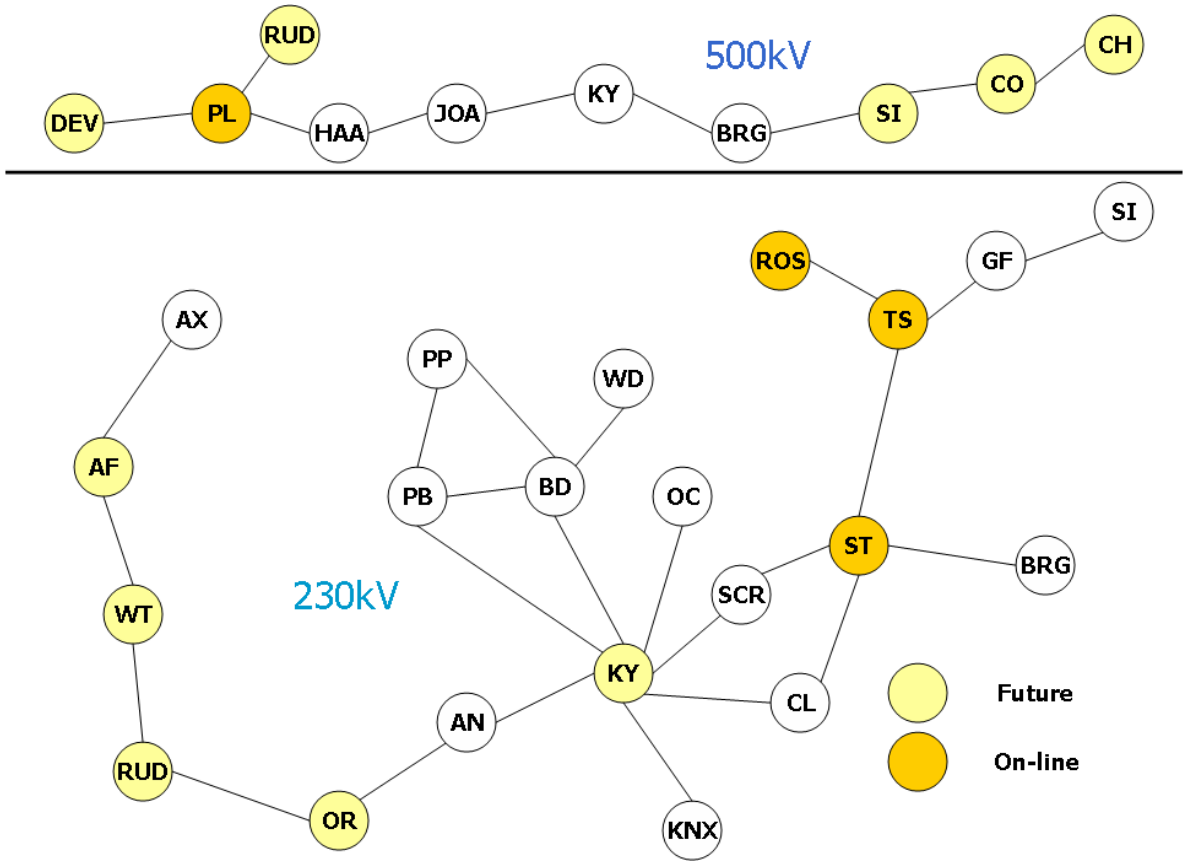

Figure 17
Present & Planned PMU Locations on the SRP System

Communication Bandwidth Requirements

Bandwidth is driven by the user's appetite for data, data type, and rate of delivery. For example, choosing a phasor reporting rate of 60 phasors/sec for 1 voltage, 5 currents, 5 Watt measurements, 5 Var measurements, frequency, and rate of change of frequency – all reporting as floating point values – will require a bandwidth of 64,000 bps. On the other hand, a reporting rate of 12 phasors/second for 1 voltage, 5 currents, and frequency – reported in 16 bit integer format – can be accommodated over a 4800 bps channel.

In making the data requirement/channel bandwidth choices, both present and future requirements should be considered. In particular, the possibility of future closed loop control should be considered in as much as the speed of the control is a function of the available bandwidth.

The next architectural choice is that of a physical communication channel and the related protocol stack. This choice is often driven what is available from a given manufacturer, however, some choices include:

- Ethernet
- Serial (RS-232, RS-485, IEEE C37.94, Direct Fiber)
- IP Multicast
- Protocol Choice: UDP vs. TCP

The clear trend in this field is the enterprise-wide deployment of a high-speed network-based communications infrastructure. Most utilities now have the ability to send synchrophasor data from the field to the office in a standard format over a corporate Wide Area Network (WAN) to a central Phasor Data Concentrator (PDC). This capability has made widespread data gathering virtually "plug and play" and no longer requires dedicated, maintenance-intensive communication circuits for synchrophasor data. If a direct Ethernet output is unavailable on a particular relay, a separate serial-to-Ethernet converter can be utilized to immediately convert the serial output to Ethernet. This has been successfully done at SRP and for now is the standard way to retrieve data from a particular class of relays. Other relays can provide a direct Ethernet output which simplifies the data gathering even more.

On the choice of TCP vs. UDP, then choice has to be made based on streaming rate and data reliability requirements. If only 1 to 5 phasor sets per second are being sent, reliable delivery becomes critical. The use of TCP is recommended as it will "guarantee" deliver of this data. If, however, higher streaming rates are chosen, trying to "guarantee" data delivery with TCP becomes problematic as channel bandwidth to support retransmission becomes an issue. As such, the use of UDP is recommended for this application. It should be noted that at the higher data rates, there is a good amount of "information" redundancy built into the data stream and periodic loss of a packet is typically not an issue.

Phasor Data Concentrator (PDC) Implementation Options

On the PDC side of things, the options fall into two main categories – an integrated hardware/software solution or a software solution that runs on standard "server" hardware. In addition, there is a super-set of the PDC known as the "Super PDC" which is designed to integrate data from a few hundred PMUs. PDC application SRP has had some experience with both types of systems and both have their advantages and disadvantages. Issues to evaluate with either solution include:

- Hardware hardening
- Access security
- Communication capabilities
- Software upgradeability
- Storage options (including data compression)
- Configurability options / functional capability

SRP is currently leaning more towards a software-based PDC approach due to the dynamic state of this field at the moment. Getting tied into a specific hardware solution can limit your system's evolution. Eventually a hybrid approach may prove best. Hardware-based PDC solutions seem to be better suited for receiving the synchrophasor data streams from the individual PMUs while software-based PDC solutions are better at visualization, data archival and "hooks" into other systems.

PDC Storage and Data Access

A critical feature of any PDC is synchrophasor record storage and provisions for data access. Record storage is vital for post-event analysis purposes and having "standard" data access for visualization and other applications will enable "interoperability" of these processes. The IntelliGrid recommended interface [4] for this purpose is OLE for Process Control (OPC) and though not in the original IntelliGrid recommendation, clearly, the Structured Query Language – SQL would be another recommended "standard" interface. Figure 18 illustrates a PDC architecture that implements these recommendations.

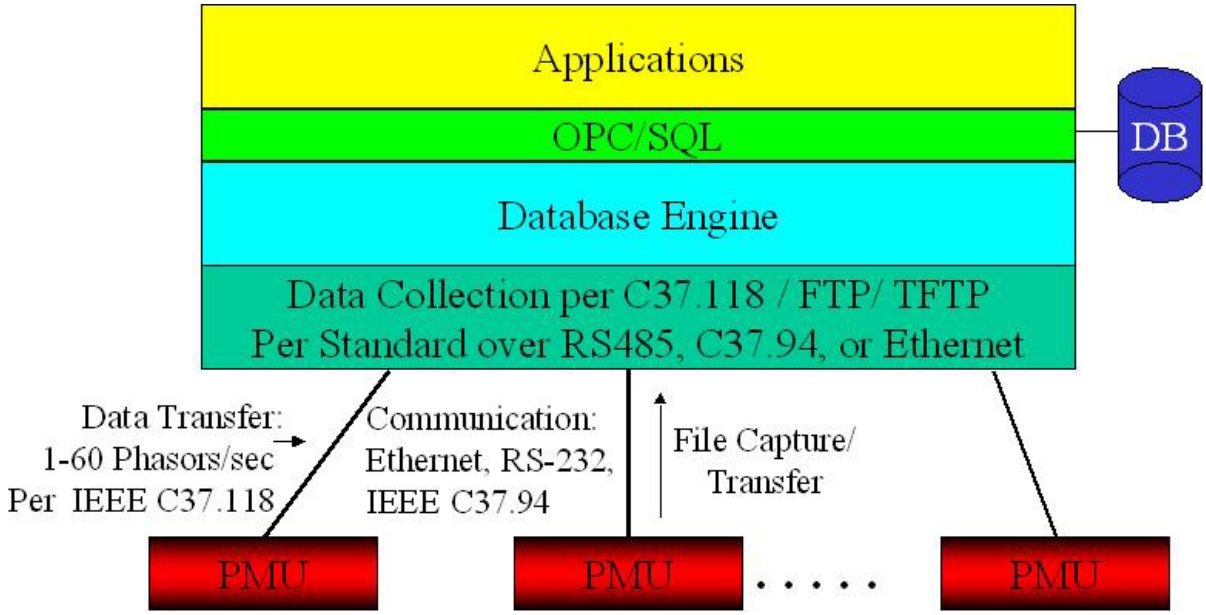

Figure 18
Proposed Phasor Data Concentrator Architecture

Some sort of data compression is desirable due to the vast quantities of data that can be created by PMUs reporting at up to 60 records per second. Data compression, however, should not make it unbearable for users to locate and extract events of interest. Standardized data extraction for visualization and other applications is very desirable if not required and it is also important to be able to take raw data and easily manipulate it using common applications such as Microsoft Excel or a simple Comtrade viewer. Terabyte data storage systems are available on today's server platforms and routine tasks such as moving data from short term to long-term storage are completely automated. Software packages are also available for vital tasks such as monitoring storage volume, monitoring the health of PDC applications, and monitoring communications availability. Table 3 shows some examples of these applications.

Table 3
Software Packages

Monitoring the health of applications	Monitoring communications availability	Running applications as services	Monitoring storage volume
Attention! Software	*Pingplotter Pro*	*Firedaemon*	*SequoiaView*

Other PDC topology considerations involve distributed, master-slave PDC approaches versus a centralized PDC. A distributed PDC approach – where a PDC is installed in the field, at a receiving station for example – is ideal for sites where many individual PMUs are located. The distributed PDC (likely a hardware-based PDC) can collect data from multiple PMUs on-site over an Ethernet (or serial) data connection and can then aggregate the data and send it back to a master PDC via Ethernet. A centralized PDC approach is architecturally desirable for gathering data from multiple, widely dispersed stations that may only have a single PMU in each of them. Ultimately, the PMU implementation strategy and physical topology of a particular system will dictate the approach.

Synchrophasor System Operation

Once a synchrophasor system has been designed and deployed, it must now be operated. The distributed nature of a synchrophasor system presents operation challenges in that complete functionality depend on many diverse components, such as time synchronization quality, communications channel availability, PMU hardware, PDC software, etc – where the failure of any one component may result in the loss of data.

The first step is becoming aware of such a problem. For this reason it is important to build into your system alarms and annunciation for all the different failure modes described above. If a system is successfully deployed it will have customers, and those customers will come to depend on the applications supported by the PMU synchrophasor data. Without the data the applications will fail and the customers (likely other departments at your same utility) will have problems operating or analyzing the electric system.

Some common problems that we have seen during the course of our system operation are shown in table 4.

Table 4
Common Synchrophasor System Problems

Problem	Frequency
❑ IRIG-B signal failure o With no IRIG-B signal, there is no synchrophasor data. IRIG-B clock problems, antenna problems, and even interference can cause the loss of IRIG-B signal.	Fairly common, attributed to interference or clock-PMU incompatibility
❑ Communication channel failure o Ethernet is a fairly robust communications medium. We do see communication channel failures occasionally but they generally heal themselves.	Not particularly common or persistent
❑ PDC software lockup o PDC software, like any software, can occasionally lock up and may require an occasional re-boot.	Somewhat common, attributed to the relative "infancy" of the software
❑ Inter-vendor interoperability issues o Different vendors' equipment does not always like "to play together."	Somewhat common, attributed to the relative "infancy' of synchrophasor technology and the just-released revision of the Synchrophasor standard
❑ Different interpretations of the C37.118 standards o Some elements of the standard are not defined as stringently as others	Somewhat common, attributed to different interpretations of the same standard
❑ Inter-utility data sharing issues o Cyber-security developments such as firewalls are a hindrance to sharing data from utility to utility.	Very common, attributed to the need to maintain a secure network

Ultimately these problems will need to be resolved before synchrophasor applications can become robust enough to enter the mainstream utility environment. Some resolutions depend on vendor involvement and cooperation; others are issues unique to each utility.

Synchrophasor Applications

Black Start

The first application of synchrophasors at SRP was to support the Black Start exercise. SRP is required to provide a Black Start capability upon wide spread loss of the Bulk Electric System - BES. The SRP solution is to start a generator at one of the dams on the Salt River. This will allow some load to be served and also the combustion turbines about 50 miles away to be started. The process is to start generation then add load then start generation then add load until the BES could be restored. The exercise was a success and SRP was able to start the combustion turbines.

Of course the BES cannot be turned off just to perform a black start test and consequently, there was a great effort to isolate the black start area without affecting any customers. The combination of hydro-generation and load in the black start area created a slightly unstable power system. Synchrophasors were subsequently used to monitor the stability of the black start area. The synchrophasor angles were also used to help synchronize the black start area with the BES when the test was finished.

Figure19 is a simplified one-line diagram showing the Black Start area. Hydro generation in the north (top of diagram) is connected to the combustion turbines in the south by a 230kV transmission system. It is SRP's opinion that once the combustion turbines are started that more load and more generation can be brought on line.

69kV Islanding

In another application, SRP will implement synchrophasors within various 69kV electrical islands within its system. The 69kV islands are connected to a 230kV receiving station that is in turn interconnected with many other 230kV-receiving stations. The 69kV islands may not be interconnected with other 69kV islands. SRP's distribution system is a radial 12kV grounded wye that is fed by a 69kV sub-transmission system. At times the electrical angle across the 12kV feeder circuits of neighboring islands can be extreme. The map shown in figure 20 is a graphical representation of two 69kV islands within SRP's service territory. The blue area is one island that is not interconnected with the brown area to the south. The red lines indicate that connections between the substations exceed a safe voltage angle and the feeders should not be paralleled.

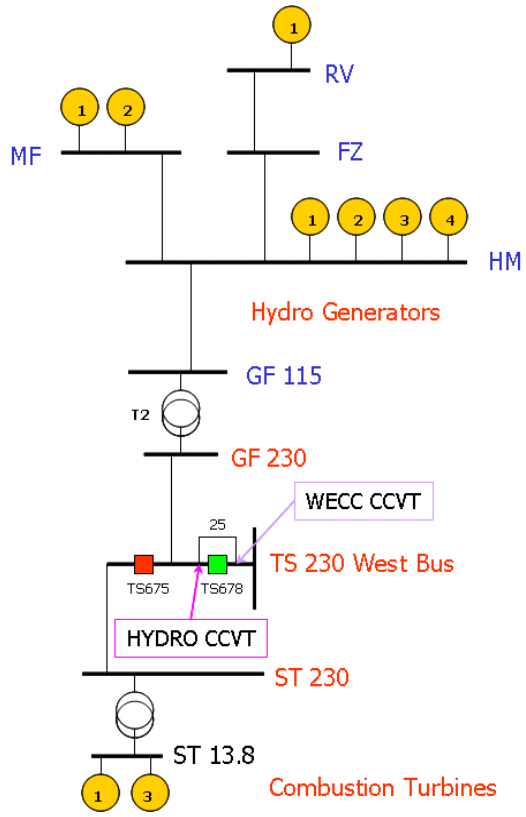

Figure 19
Black Start System Configuration

The applications of synchrophasors at the various substations will allow the distribution operators to determine if it is safe to parallel islands. Synchrophasor data can also be used to confirm the existing analysis used to produce these study results.

State Estimation Verification / State Measurement

SRP also plans to incorporate the synchrophasors into the State Estimator that runs on the EMS. In the short term the data available may only allow gross verifications to be made to the State Estimator. However, as more data is brought into the EMS, the state estimator should be able to be speeded up and made more accurate.

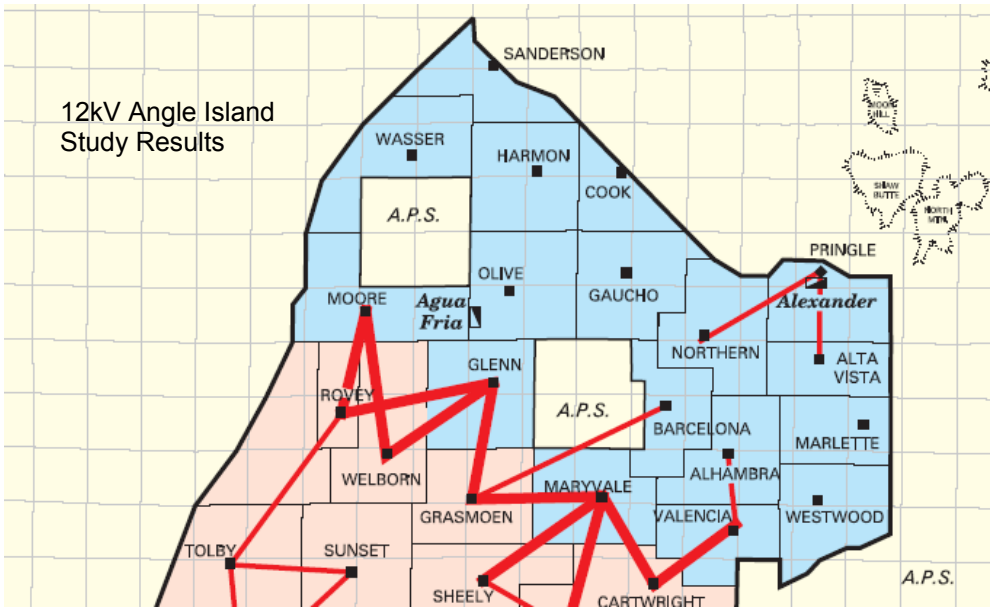

Figure 20
69kV Island - Excessive Angle Visualization

Future Applications

There are any many potential applications that will be realized as more data is made available in more locations. For example, operational visibility of reactive margins, voltage support, and power transfer capabilities. System Protection and Planning groups will be able to verify line impedance and may be able to use the data to determine voltage collapse margins.

Bibliography

1. Complex Quantities and their use in Electrical Engineering; Charles Proteus Steinmetz; Proceedings of the International Electrical Congress, Chicago, IL; AIEE Proceedings, 1893; pp 33-74.
2. A New Measurement Technique for Tracking Voltage Phasor, Local System Frequency, and Rate of Change of Frequency; A. Phadke, J. Thorp, M. Adamiak; IEEE Trans. vol. PAS-102 no. 5, May 1983, pp 1025-1038
3. IEEE Standard for Synchrophasors for Power Systems; C37.118 – 2005.
4. IntelliGrid Architecture (formerly the Integrated Energy and Communication System Architecture; www.e2i.org.

Biographies

Dr. William Premerlani is employed at GE's Global Research Center in Schenectady, New York and a holder of over 20 patents. His current research interest is in phasor measurement, advanced algorithms for diagnostics and protection, with applications to motor diagnostics and power system relaying. Other recent work included object - oriented technology, with application to achieving interoperability among communicating intelligent devices in power systems. He is a co-author of the popular textbooks: "**Object Oriented Modeling and Design"** and **"Object-Oriented Modeling and Design for Database Applications".**

Dr. Bogdan Kasztenny received his M.Sc. and Ph.D. degrees from the Wroclaw University of Technology (WUT), Poland. He joined the Department of Electrical Engineering of WUT after his graduation. Later he was with the Southern Illinois and Texas A&M Universities. Currently, Dr. Kasztenny works for GE Multilin as the Chief Application Engineer and a Project Leader for the Universal Relay family of protective relays. Bogdan is a Senior Member of IEEE, has published more than 100 papers and holds several patents on protection and control.

Mark Adamiak received his Bachelor of Science and Master of Engineering degrees from Cornell University in Electrical Engineering and an MS-EE degree from the Polytechnic Institute of New York. Mark started his career with American Electric Power (AEP) in the System Protection and Control section where his assignments included R&D in Digital Protection, relay and fault analysis, Power Line Carrier and Fault Recorders. In 1990, Mark joined General Electric where his activities have ranged from development, product planning, and system integration. In addition, Mr. Adamiak has been actively involved in developing the framework for next generation relay communications and was the Principle Investigator on the Integrated Energy and Communication System Architecture (IECSA) project. Mark is an IEEE Fellow, past Chairman of the IEEE Relay Communications Sub Committee, and a member of the US team on the IEC 61850 Working Group on Substation Communication.

Jonathan Sykes received an Electrical Engineering Degree from the University of Arizona. He is also a registered professional engineer in the State of Arizona. Jonathan received the 2004 IEEE PES Phoenix Chapter Award for Outstanding Engineer. He has been with Salt River Project for 24 years and started working in the planning and generation departments but has more than 20 years experience in System Protection. Jonathan and is currently responsible for EHV relaying and Remedial Action Schemes at SRP and has been involved in automation and integration throughout the SRP system. He is presently a Senior Principal Engineer with SRP in System Protection and a member of NERC, IEEE, and WECC Committees.

Kris Koellner is a registered Professional Engineer in the state of Arizona. Kris graduated from Arizona State University with a B.S.E degree in Electrical Engineering. He has been with Salt River Project (SRP) since 1994 and has worked in the System Protection department since January, 2004. Kris has also worked in SRP's Power Quality and Distribution Planning groups.