PNNL-25643

U.S. DEPARTMENT OF ENERGY

Prepared for the U.S. Department of Energy
Under Contract DE-AC05-76RL01830

Measurement of Phasor-like Signals

H Kirkham
A Riepnieks

June 2016

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401, fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service,
U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161
ph: (800) 553-6847, fax: (703) 605-6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/ordering.htm

This document was printed on recycled paper.
(8/00)

PNNL-25643

Measurement of Phasor-like Signals

H Kirkham
A Riepnieks

June 2016

Prepared for
U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Preface

This report is a summary of work done principally in 2015 and 2016. The work began, at a low level of funding and activity, a year or so earlier, in an attempt to understand what *exactly* was being measured by a device known as a phasor measurement unit. This was important because I was a member of the working group of experts who were writing the improved Standard for PMU (phasor measurement unit) performance.

The signal going into a PMU was clearly not a phasor, despite the name of the measuring instrument, for a phasor has certain characteristics not possessed by this input signal. In particular, a phasor is a signal whose domain is infinite in time—its amplitude, frequency and phase never change. Yet the PMU was asked to measure something called rate of change of frequency. By the definition of a phasor, that parameter had to be zero.

For that matter, what was the meaning of "frequency" when the frequency had a rate of change? As a new member of the working group, I sought the views of the other experts. I was reminded of the notion of Putnam, who claimed not to be able to tell an elm tree from a beech tree. In a thought experiment in (Putnam H. , 1981) he says

> My concept of an elm tree is exactly the same as my concept of a beech tree (I blush to confess) . . . you and I both defer to *experts who can tell elms from beeches*.

The italics are mine. Unlike Putnam, I found I could not defer to the experts. My PMU experts might have been able to tell elms from beeches, but they could not tell me how to distinguish changing phase from changing frequency. As L. Frank Baum taught us, sometimes the wizard is just a person.

Finding the solution to this dilemma began with an intellectual journey to discover what it meant to make a measurement. What are we doing when we do the things that we do to obtain a result, to perform a calibration, and so on? At this time, after the last couple of years, I find that the answers to this kind of question are clear enough. They are answers that have been hinted at for over a hundred years, by the likes of Carey Foster and Kelvin, but never elucidated. They are answers that draw on the work of Shannon, and go further. The answers are part philosophy, part epistemology, part semantics, and part mathematics.

But at the end of the journey a clear and understandable statement is possible. When we make a measurement of this kind, we are *finding the coefficients of an equation*.

This idea, the notion of measurement being about solving an equation, was at first greeted with strident opposition from several members of the PMU community. When that had subsided somewhat, there was some misunderstanding over which equation was being solved. Was it the equation of a phase-modulated cosine wave?

No it wasn't. But stipulating what the equation is, is simplicity itself. The equation is whatever model we have chosen to represent the signal for the purpose of our measurement. It is the equation that includes the various parameters we ask of the measurement device. In the case of the PMU, those are the amplitude, the frequency and the phase. (Some would include the rate of change of frequency, too.) There are not many choices about how these can be combined.

It seems that the last year or so has been a time of significant advance in the fundamentals of the science of measurement. My colleague Artis Riepnieks and I have published several papers, versions of most of which are included as part of this document.

As this is written, one manufacturer is implementing one of the methods that come out of this work. It is not a licensed effort, for our work product is at the level of basic knowledge, the sort of thing that should be in a textbook. Doubtless, there will be others who follow.

For my own part, it is pleasing to think that I may have initiated a change in the way metrologists think about some kinds of measurements. The advances will no doubt continue, as new methods are tried, and other new ideas developed.

Harold Kirkham
Richland, WA, June 2016

Executive Summary

We describe a new method of characterizing a power system signal, the problem of the phasor measurement unit. The work is founded on several insights. First, insights about the measurement:

(1) Measurement in general uses the effects of the real world to find the values of parameters of a conceptual model.
(2) Measurement of the kind made in a PMU is essentially the act of solving an equation. That equation, the model, is the equation of a phasor.
(3) The real world is distinct from the conceptual model. Nature and our knowledge of nature are not the same thing.
(4) It follows that the result of a measurement is the value of a parameter in the conceptual world, it is *not the value of anything in the real world*.

The engineering community is limited by linguistic processes. An accurate description of what we are measuring is needed so we can write standards specifying it. This report shows that we should cut back on the labelling process, and embrace the idea of a mathematical model. Then, a word like “frequency” can be seen as just a label for a particular parameter in a particular equation.

Next, insights about the implementation

(5) Since the form of the equation is fixed by the physics, regarding the measurement as what mathematicians call a *fitting problem* is appropriate. In fact, any measurement is a matter of fitting a selected model to the observed (sampled) data.
(6) Implemented as a fitting problem, each parameter is treated as a “primitive” quantity. Frequency is not found as the derivative of phase, for example.
(7) Measurement in general should take advantage of the information contained in the residuals left over from the fit. This is information about the signal that is (by definition) not contained in the result of the measurement. Until this present work, that information has not been acknowledged.

* * *

With “clean” signals, the method described here (sometimes called the “new PMU” method) gives values that are limited in accuracy only by the capability of the computer that does the calculations.

The new PMU method differs from existing PMUs in that the measurements are independent, no matter the window width used. The method is also are able to use far fewer sample values than a commercial PMU. For example, four values of the signal, sampled over a period of just 3 ms (a fifth of a cycle), is sufficient to give precise results, including the rate-of change of frequency with a magnitude as low as a few μHz/s.

To our knowledge, no other system comes close to this level of performance.

The “new PMU method” has been simulated as part of a software structure that generates a test signal with known and controllable parameters, and that can add (separately and controllably) phase noise, amplitude noise and frequency noise.

The ability to control the added noise makes it clear that the “perfect” PMU performance we obtained is possible because of the perfection of the input signal. That signal imperfections add performance limitations is recognized by those calibrating PMUs in the laboratory. However, until our method was developed the effect of the noise was just an unknown in the measurement process.

The report shows that in the real world, even the slightest added noise—certainly less than exists all the time in the power system—makes the measurement of rate of change of frequency in a short observation window impossible. That is not because of any implementation issue, but is a fundamental aspect of the measurement.

A new capability to assess—in real time—the *quality* of a particular measurement is an outcome of the method we use to make the measurement, but it can be applied to the measurement results given by any PMU, furnishing valuable additional information to the user in real time.

The work opens the door to a host of measurements beyond the PMU. New algorithms can be developed. New insight is given into the process of calibration.

Many new (and perhaps radical) ideas are developed and demonstrated in this report.

Philosophical background for some of the development is given.

Acknowledgements

The authors would like to thank Zak Campbell of AEP and Ray Hayes of ESTA International, late of AEP, for the oscillography data from the 345 kV system. The data enabled us to demonstrate very clearly that the conventional PMU is actually solving an equation.

We thank Prof Mihaela Albu of “Politehnica” University of Bucharest, Romania for many valuable and insightful discussions.

We thank Dr Alex McEachern of PSL for sample values data from a distribution system. This data stream allowed us to explore the Allan variance method for a real signal.

The authors also acknowledge the support of Phil Overholt at the US Department of Energy. His continued sense of the value of the effort as it developed was of importance in maintaining momentum.

Table of Contents

List of Figures

1.0 Background and Introduction

1.1 Foundations

Foundational studies of measurement seem to have been few and far between, and they are not particularly well known. John Simpson of the U.S. National Bureau of Standards offered the view that as a result the literature on measurements was of an *ad hoc* nature, and widely scattered (Simpson, 1981).

The present work is about as foundational as it gets for metrology. This report elucidates a new foundation of one aspect of metrology: the characterization of power system waveforms. The work has much broader implications, however.

Doubtless some of the ideas presented here can be quite generally applied, as that is the nature of foundational work. Any measurement that involves the characterization of a signal is likely to be affected. That means the measurement of any voltage or current, for example, from the 1.5 volt hearing-aid battery to the 120-volt stuff coming out of the wall-outlet, even to the 1-MV surge used in testing the withstand capability of an insulation system.

While we will focus here on the phasor measurement unit, the foundational aspects of the work are necessarily broadly based. The investigation of the physical world is ordinarily the domain of the engineer and the scientist. The clarification of the logical and methodological processes of that investigation is more the domain of the philosopher. This report includes elements of both domains. As far as possible, we will restrict such matters to the few pages of this first Section, and concentrate on engineering afterwards. But it is important to take the philosophical matters seriously. In some ways, they reflect the most fundamental outcome of the work.

1.2 Non-stationary Parameters

The work on the measurement of phasor-like signals was triggered by an investigation into the meaning of the word frequency when frequency was changing, a problem faced in the PMU. One of the present authors (Kirkham) was a member of the working group writing a new standard for the performance of the PMU, and wondered how one could assess the accuracy of the PMU when the thing being measured (frequency) was not defined if it was changing. Consider the signal in Figure 1, which represents a sine-wave with increasing “frequency.”

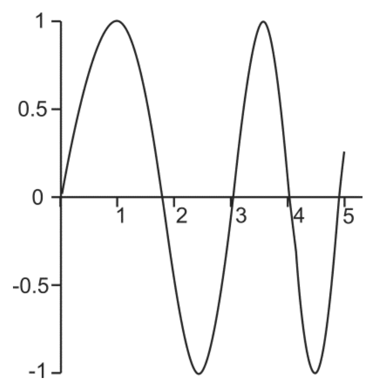

Figure 1 Signal with increasing "frequency"

Several questions demand an answer. For a wave as shown in the figure, even a simple method of measuring frequency based on measuring the time between zero-crossings would not work. The intervals are clearly decreasing half-cycle by half-cycle. How then is the measurement to be made?
More importantly, what would be the meaning of the result of the measurement? Measurement is a process that, once completed, has a *result* that is unchanging afterwards. Should the result be the frequency (seemingly low) at the start of the interval shown, or the last value at the end of the period, or some sort of average?

Whichever option is selected, it seems that the result of the measurement would depend on the way the measurement is made. That is not something one expects of a measurement. If one measures a length with some appropriate means such as a ruler or a tape-measure, one does not expect a different result. It is a problem caused by a lack of definition.

The matter seemed serious, for papers describing the calibration of PMUs at NIST included graphs such as Figure 2.

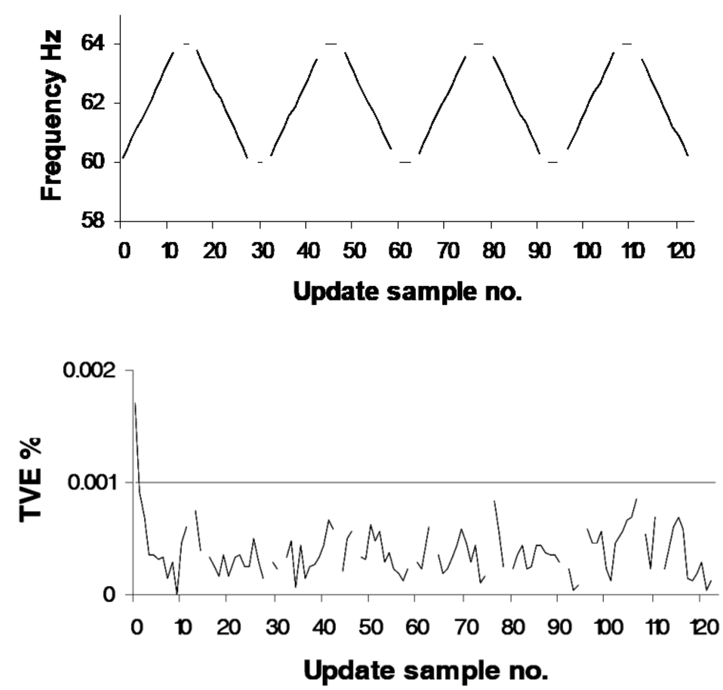

Figure 2 PMU calibration at NIST

The graphs, taken from (Stenbakken & Zhou, 2007), have obvious gaps. The paper says we should "Note that the frequency and errors at the transitions are not reported." We noted that. We did not think it exactly fair to the ultimate user of the PMU, nor necessary to the tester.

One reason the gaps are there is that the authors did not have a definition of the frequency during the "transitions." The real-world PMU will not have the option of simply "not reporting" a result. Yet, lacking a calibration, we are left guessing what might happen at a "transition" in the real world. The real world is, after all, not a laboratory where signals hold steady while you measure them.

Colleagues on the working group, experts all, were unable to help. In terms of Putnam's linguistic division of labor (Putnam, 1973), they were *supposed* to have the answer.[1] As measurement people about to specify tests for the accuracy of commercial products, our group *needed* to have an answer.

1.3 The Meaning of Measurement: Semantics

The matter quickly went beyond the narrow purview of just considering the *definition* of frequency. It became a matter of finding a *meaning* for the word.

In a report to DOE in 2015 (Kirkham H. , 2015), it was observed that the "information" content of messages that had been set aside by Shannon in his Theory of Communication (Shannon, 1948) might be exactly applicable to metrology. Shannon had written

> The fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities.

His next sentence was: "These semantic aspects of communication are irrelevant to the engineering problem." From this beginning, Shannon went on to create the most important theory of communications.

It is strange, therefore, that Shannon's foundational work on communications came to be known as Information Theory. Shannon himself was troubled by that (Shannon, 1956); understandably since he had dismissed the semantic aspects of the communications as not relevant to the problem he was addressing. We do not criticize Shannon; by restricting the scope of his work, he was able to develop the important results that we are familiar with, based on notions of probability and entropy.

The DOE report stressed that the "message" coming from a measuring instrument had "meaning." That was the semantic connection to the "physical or conceptual entity" that Shannon chose to set aside. The *result* of a measurement was indeed *information*.

It may be that Shannon struggled with separating the "physical or conceptual entities." Indeed, his use of exactly that phrase suggests that he did.

The separation between the physical and the conceptual was emphasized by the physicist and philosopher Rudolf Carnap (Carnap, Philospohical Foundations of Physics, 1966). To underscore the separation, he introduced a notation consisting of a small circle "○" to represent the physical operation of joining. Using this notation, the additive rule for length becomes

$$L(a \circ b) = L(a) + L(b) \tag{1-1}$$

[1] Philosopher Hilary Putman argued that in linguistics it was not necessary for everyone to know how to distinguish gold from other materials, but there had to be a "subclass" of speakers who knew how. One might call them experts. The ordinary speaker would rely on them for the meaning of the word. The PMU group did not have the linguistic expertise to distinguish, for example, a changing frequency from a changing phase, or to explain what the "frequency" was at a transition. It should also be noted that "frequency" is a word that implies something constant for all time, and that therefore even the label on the ordinate of the top graph of Figure 2 is not defined.

Here we note that the equals sign must be treated with caution. It can be taken to mean "has the same value as," but it does not mean "is the same thing as." That is, the equation could be expressed "The length of *a* joined to *b* is given by the length of *a* added to the length of *b*." This notation avoids the mistake (that Carnap said is not uncommon) of writing

$$L(a + b) = L(a) + L(b) \qquad (1\text{-}2)$$

The reason for making the distinction is that the "+" sign is a mathematical operation indicating addition, and while one can add two numbers, one cannot "add" objects. Another way of looking at the distinction is to note that the two objects *a* and *b* are observable, they are physical things. One cannot "add" physical things. On the other hand, their lengths are conceptual or theoretical things, and they are treated by mathematics.

Put another way, the left side represents nature, the right side our knowledge of nature.

The work to be presented here, derived from a search for the meaning of "frequency" when frequency was changing, is based on the notion that the process of measurement is a bridge between the left side of Equation (1-1) and the right side. A bridge between nature and knowledge. The distinction is a useful one to retain, as we shall see.

1.4 Calibration – a background

Calibration is an aspect of metrology that allows us to illustrate much of what has been introduced so far. In metrology terms, calibration is part of a process that allows measurements to be traced to standards. That tracing process consists of a number of steps that relate a standard (for example at a national metrology institute) to a system in the laboratory. The sequence of steps unavoidably adds uncertainty at each calibration, so understanding (and documenting) errors and uncertainties is an essential part of the process.

At each step in the sequence, calibration consists of two parts. First, a comparison is made between what is called the *declared value* of an attribute of a calibrating artifact, such as a reference standard, and the declared value of an attribute of the unit under test. The outcome of this step is that an *error* is measured. Second, taking account of the relative uncertainties, a correction factor (or a series of such factors) is calculated that should be applied to the declared value of the instrument during a regular measurement to yield a corrected result. The correction will have the effect of reducing the measurement error, making the system being calibrated as accurate as it may be.

There are four ways to do the comparison of the first step that are slightly different in principle. The choice depends on the nature of the unit under test.

1.4.1 Reference measuring system

A straightforward configuration can be used if a property of the unit under test provides the stimulus. For example, the unit under test could be a signal generator, or a voltage source. The declared value of the unit under test is its nominal value (if it is something fixed) or the indicated output over some range of values. The calibrating artifact provides the sensor. The calibrating artifact's declared attribute value is displayed or otherwise shown, as in Figure 3. In the figures that follow, we indicate the instrument with the lower uncertainty by an increased display resolution.

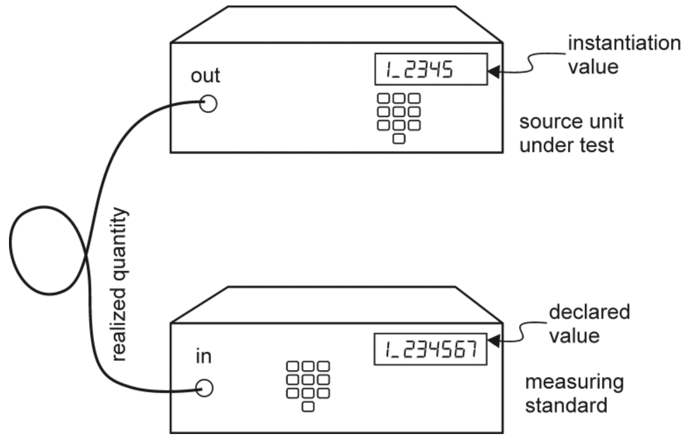

Figure 3 Calibrating a source

1.4.2 Reference source

In a reversal of the previous configuration, the calibrating artifact provides the stimulus. The calibrating artifact's declared value is its nominal or indicated value. The unit under test is the sensor. The sensor responds to the stimulus and drives a display. The displayed reading is the declared attribute value, as in Figure 4.

The connection between the instruments is the same, so the method is fundamentally the same as the previous one, but the uncertainties are allocated differently. We will return to the topic of uncertainties later.

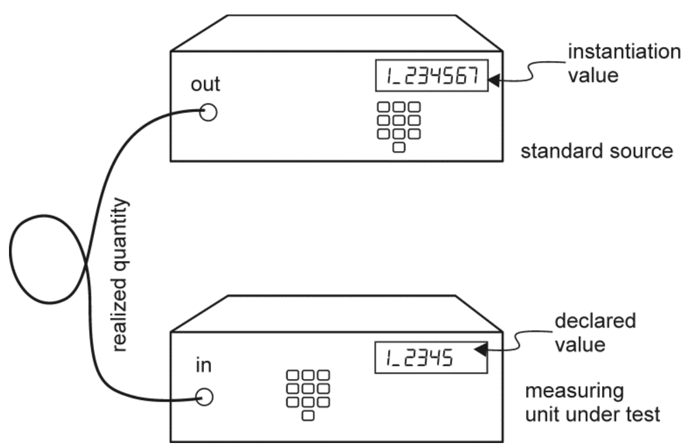

Figure 4 Calibrating a measuring device

The test system we have developed for simulating the PMU and evaluating its performance is essentially of this kind. In Figure 4, a signal is generated that is quite precisely known and it is applied to the measurement device under test. In our system, described in the next section, we never actually generate the signal. Instead, we generate the sample values that would result from a perfect A/D sampling process, and we submit these to the measuring system.

For the sake of completeness, we continue the story of calibration a little longer.

1.4.3 No stable source

In the third configuration, the stimulus is supplied by a source external to both the calibrating artifact and the unit under test. Each responds to the stimulus and drives a display. The displayed readings are the declared attribute values of the calibrating and calibrated unit, as in Figure 5.

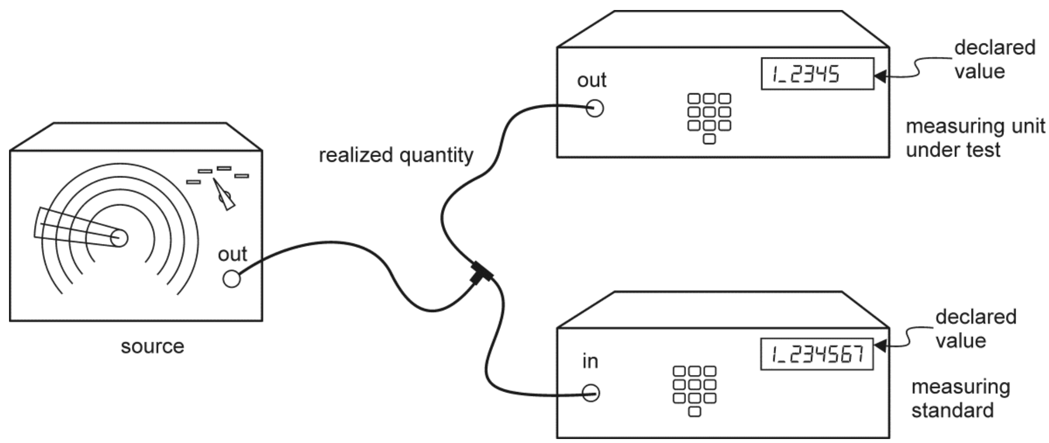

Figure 5 Calibration using (relatively) low-quality source

This process is used when the source is not particularly stable or repeatable: the unit under test and the measuring standard should agree, even if the stimulus value changes as the calibration proceeds. In the Figure, the source is shown as an old-style signal generator, with a scale and pointer to display the frequency. (One may imagine the tubes inside causing frequency drift as they warm up!) Applied to the measurement of the electric power system, one may imagine that the signal source is the power system, and the measuring standard is one of the “gold” or “platinum” PMU that people are trying to develop.

In other words, this kind of calibration is more representative of applied metrology than it is of pure metrology. Here the signal does not “hold still” while it is measured. Unfortunately, the PMU standard has failed to acknowledge this, and it deals with the changeability problem with gaps in the record (such as we saw in Figure 2), based on statements such as “An adequate settling time shall be allowed for each test signal change to prevent parameter change transient effects from distorting the measurement” and “The error calculation shall exclude measurements during the first two sample periods before and after a change in the test ROCOF. (Extracts taken from IEEE C37.118.1, 2011, p 17 and p 19.)

We addressed the issue of blanks in the record in a paper (see Section 7.6) using the simulated PMU that we developed at PNNL. We propose that this PMU method be further developed so that it operates in real time. If it can be made to operate at such a speed (and we have no reason to think it cannot), it is fair to think of it as the best that can be done. Since it is based on a least-squares estimator, its results are the best possible, in the least square error sense. It would be the “ultimate” platinum PMU.

1.4.4 Networks

In all the calibration methods so far described, the error of the unknown is disclosed by a simple comparison of two numbers. The numbers are known: they are the declared value of two instruments. It is assumed that the uncertainties of the instruments are also known.

Systems are also used which may be regarded as an adaptation of the method of Figure 5. The comparison is done in the signals involved by means of a network of some kind, and adjustments are made *in the network* to make the difference zero. There is no declared value comparison to produce an error number *per se*, although automated systems are capable of performing the required calculations do display the end result. (Think self-balancing bridge.)

In essence, the quantity of interest is still the difference between the signals involved, but the comparison is not done numerically it is done physically, and the difference is disclosed by the network values.

The most familiar version of such a calibration is probably the Wheatstone bridge. The basic arrangement is shown in Figure 6.

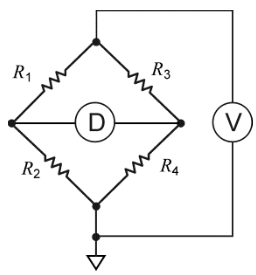

Figure 6 Basic arrangement known as Wheatstone Bridge

In this configuration, no matter the voltage V, the fraction of that voltage that appears at the junction between R_1 and R_2 is given by $R_2/(R_1 + R_2)$. Similarly, the voltage that appears at the junction of R_3 and R_4 is given, as a fraction of the total voltage, by $R_4/(R_3 + R_4)$. When there is no difference between these voltages, either the voltage V has become zero or the bridge (said to be "balanced") is in a state in which

$$\frac{R_2}{R_1 + R_2} = \frac{R_4}{R_3 + R_4}$$

whence

$$R_2(R_3 + R_4) = R_4(R_1 + R_2)$$

$$R_2R_3 + R_2R_4 = R_4R_1 + R_4R_2$$

$$\frac{R_2R_3}{R_4} + R_2 = R_1 + R_2$$

and so finally

$$\frac{R_2R_3}{R_4} = R_1$$

If three of the resistors are known, the unknown can be found in terms of their values. The voltage applied does not come into the equation. Further, the voltage across the detector is required only to be adjusted to zero, so even an accurate voltmeter is not required.

A number of such networks have been developed over the years, suitable for evaluating ac components, or high voltage components, and for components at the extremes of high and low values. For example, the Kelvin Double Bridge, an adaptation of the Wheatstone, was designed to eliminate the effects of contact resistance in the measurement of very low ohmic values.

Other methods of this general type are the current comparator and the transformer ratio arm bridge. These are devices that are essentially passive, and the comparison is done in the magnetic field of a shielded transformer core. With care, uncertainties of a few parts per million are achievable. These (and other bridge-type networks) are generally analyzed on the assumption that the components are linear and the waveforms sinusoidal.

The point of all this is to note that something takes place in the calibrations of this type that has no counterpart in the figures that precede Figure 6. That difference is that some manipulations *are* taking place on the left side of Carnap's equation. There is combining of information in a way that achieves some specific objective on the right side of the equation: typically the need to know only a zero of voltage.

We shall not be considering this kind of calibration in this report. We are concerned here with characterizing signals that are not necessarily sinusoidal, and are measured by sampling. It is interesting to speculate how the world of networks might be combined with such measurements in the future.

1.5 Historical Support

The idea that measurement is a bridge between the physical and the conceptual is not widely acknowledged, even by metrologists, but it is of great significance. Measurement is an act that uses the physical world to find the values of parameters in the conceptual world. It seems safe to say it has always been so. In the world of electrical measurements, the "equals sign" of the measurement used to be manifested when the needle on a moving-coil meter stopped moving: then the forces tending to increase the reading (usually those of a current derived from some network in the physical world) were balanced by those tending to decrease it (usually a spring force). At that moment of time, the conceptual quantity could be obtained as the reading on the meter.

Nowadays, the reading on the meter is replaced by a digital display, and the calculations that lie behind the reading (usually called the *declared value* by metrologists) can be quite complex. It is crucial for the metrologists to communicate his or her needs to the digital design team. In the PMU, one has to tell the team what "frequency" means if one expects consistency.

There is an echo here of the problems faced by Maxwell, after he had been installed as the Cavendish Professor of Experimental Physics at Cambridge University. In a report (Maxwell, 1877) to the University, he wrote

> It has been felt that experimental investigations were carried on at a disadvantage in Cambridge because the apparatus had to be constructed in London. The experimenter had only occasional opportunities of seeing the instrument maker, and was perhaps not fully acquainted with the resources of the workshop, so that his instructions were imperfectly understood by the workman . . .

This two-way communication need is still with us. The difference is that these days, our “instrument makers” are software writers. And even if it is clear to the software people building the equipment that they are solving an equation, the message seems not to have reached the team writing the standard, and was certainly not evident to the user.

And yet it is crucially important. It was observed that (Heisenberg, 1959)

> . . . since the measuring device has been constructed by the observer, we have to remember that what we observe is not nature in itself, but nature exposed to our method of questioning. Our scientific work in physics consists in asking questions about nature in the language that we possess and trying to get an answer from experiment by the means that are at our disposal.

As engineers, our “method of questioning” is the measuring instrument. But we have not generally been fully aware of some of the assumptions behind the design and calibration of the instrument.

What we have discussed so far may be considered “pure” metrology. The metrologist interested in having measurement results traced back to national or international standards, and the metrologist in the national metrology laboratory, work to make the thing being measured match the mathematical model. That is, these metrologists ensure a match of the form of the left side and the right side of Equation 1-2. We see this in GUM (Metrology, 2008). GUM states (Sec. 3.4.1) that “implicit in this Guide is the assumption that a measurement can be modelled mathematically to the degree imposed by the required accuracy of the measurement.” The assumption that the model is accurate means it matches the realized quantity. That match is not guaranteed in “applied” metrology.

It may seem a trivial example, but if we use a voltmeter capable of measuring only steady dc voltage, we will be told the average value of whatever waveform we submit for measurement, even if it is an alternating quantity. The result of the measurement might be close to zero, even if there is a large ac voltage present. A similar situation was reported in (Kirkham H. , 2015) when a dc power supply was observed by an instrument incapable of reporting the ripple. In other words, it is perfectly possible to use the “means at our disposal” to ask the wrong question.

We shall see that once we understand that the instrument is providing answers to our questioning that it frames as an equation, we can see that the equipment can even tell us whether the question is well-framed. That does seem to be a new concept for metrology.

We begin the next Section by examining how it is that a digital measurement system spans the space between the physical and conceptual, in a version of the measurement framework first described in last year’s Report to DOE.

2.0 Measurement Framework

The 2015 report to DOE on the topic of measurements (Kirkham H. , 2015) introduced a framework for measurement, with the intent of showing how the various parts of the process of measurement interacted. The comments of Carey Foster (Carey Foster, 1881) discussed in that report were influential in stressing the development of the measurand, for example, but neither he nor Kelvin, who famously spoke of the importance of mathematics (Kelvin, 1889), ever made the connection between the measurand and the use of an equation.

The present work takes that notion a step further.

2.1 Calibration in the Measurement Framework

Consider the measurement system shown in Figure 6.

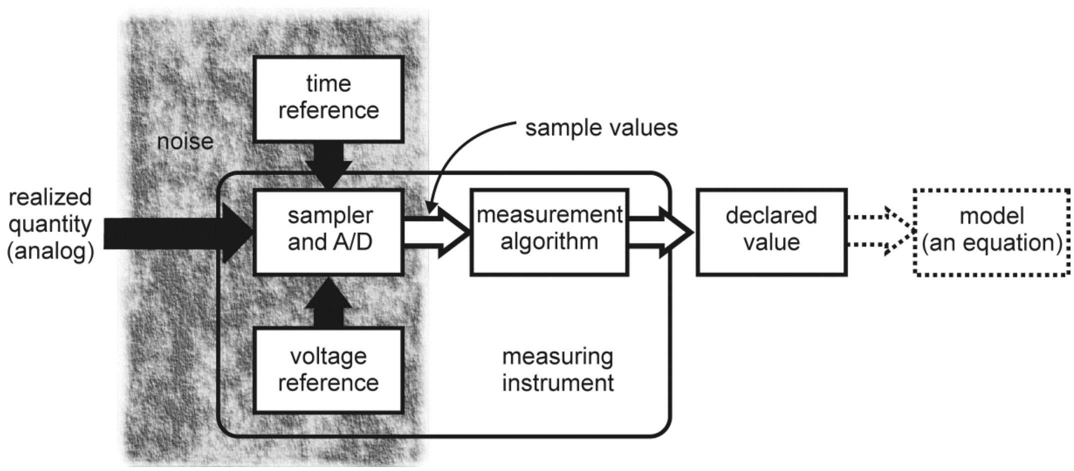

Figure 7 Digital measurement system

The realized quantity, the thing made real for the purposes of measurement, is shown going into an A/D converter. The framework diagram in the 2015 Report shows this entity being responsive to the model shown here on the right of the diagram. What this version of the diagram emphasizes is the separation of the physical entities on the left side and the conceptual ones on the right. The physical things are subject to the effects of noise, shown here as a sort of hairy carpet under that part of the block diagram.

As applied to the PMU, the measuring system in Figure 6 gets its time reference typically from GPS, and has a voltage reference that is typically a band-gap diode. (Both these references are subject to noise.) This noisy sampled signal is then given as a data stream to some sort of algorithm whose job it is to measure the signal amplitude, frequency, and phase. (Phase in the PMU is relative to a well-defined conceptual signal that can be created in the measurement system based on the time from the time reference.)

While this particular diagram may never have been drawn exactly this way before, it will not seem strange to a metrologist. Its counterpart in a signal generating system can be drawn similarly, as in Figure 7. In the figure, the process of generating the signal begins with the selection of the appropriate equation, and the selection of a set of values for the coefficients. In Figure 7 these are called an *instantiation.* The name is used here partly because the word (borrowed from computer technology) suggests something temporary, and partly because this set of values requires a name of its own, rather than to be known as another "reference." It is, after all, a different kind of reference, for it is adjusted by the experimenter to suit the specific requirements of the moment The voltage reference, on the other hand, is as constant as can be for all time.

The equation is known to the generating algorithm, and given the instantiation numbers, it creates a stream of sample values that goes to the D/A converter. Using the time and voltage references, the converter creates the realized quantity.

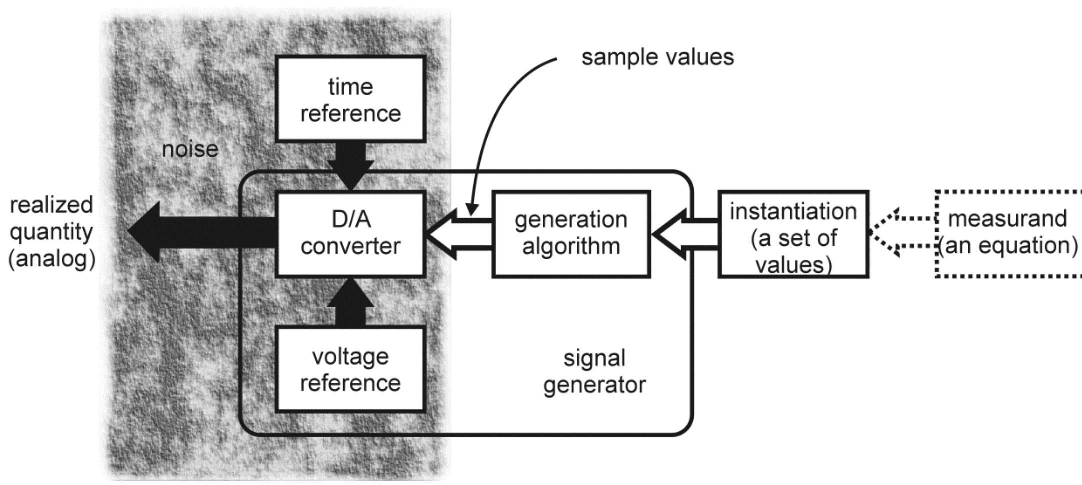

Figure 8 Digital signal generation

The measurement diagram and the signal source diagram can be combined to represent the setup for a calibration, as in Figure 8.

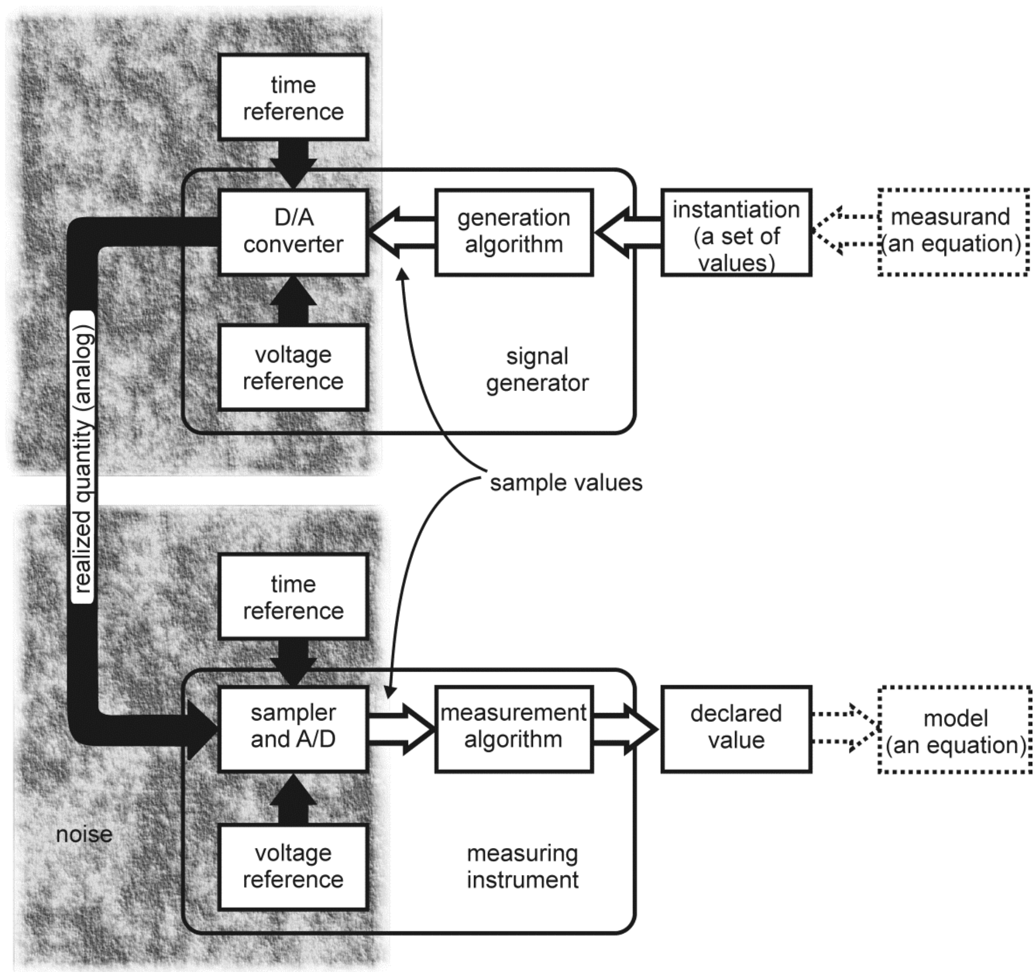

Figure 9 Calibration block diagram

We should point out, before going further, that the left side of Figure 8 is the part of the world that Shannon considered, a noisy channel carrying information that may have meaning. Our concern is not with the information-carrying capacity of the channel between the sender and the receiver, but with the part of the system that lies outside that realm. We are concerned with the meaning of the message, to put it in Shannon's terms.

What the figure shows is essentially two versions of Carnap's equation. For some general variable *X*, and where the + sign substitutes for a general mathematical operation, *w*hat is revealed is this:

$$\overset{\xleftarrow{\text{signal generation}}}{\underset{\xrightarrow[\text{measurement}]{}}{X(a \circ b) = X(a) + X(b)}} \qquad (2\text{-}1)$$

The process of signal generation moves information from the conceptual world to the physical, and the process of measurement moves information from the physical world to the conceptual. The process of calibration involves both steps.

2.2 Error and Uncertainty

In a calibration, the difference is noted between the instantiation values and the declared values. The equation used for the source model and the equation used for the measurand should be the same: then in a perfect world, the difference between the instantiation values and the declared values would be zero. If there is a difference, it is the difference between two quantities that are known, and that are both firmly on the conceptual side of the diagram. The methods of mathematics apply. This quantity is called an *error.*

In the process of *signal generation* noise and signal distortion may be added to the realized quantity, so that it is a less-than-perfect realization of the instantiation values in the measurand equation. Similarly, in the general process of *measurement* (including the measurement part of a calibration), some other noise and signal distortion may also be added to the realized quantity.

The differences between the realized value (known only through measurement) and either the declared value or the instantiation value are termed *uncertainties.* As the label indicates, they cannot be known exactly. An uncertainty is associated with the process of generating the signal in a calibration, and another uncertainty is associated with the process of making the measurement of a realized quantity.

These notions are shown in Figure 9.

It is usually said that an error is the difference between the true value of something being measured and the measured value. However, this statement usually comes with the caveat that the true value cannot be known.

Here is a new thought for metrologists to ponder: In a calibration, the instantiation value can be thought of as the true value, and it *is* known.

Even if the improbable happens and the realized quantity in a calibration is a perfect realization of the instantiation, that fact cannot be *known*, because the two things are on opposite sides of Carnap's equation. The same is true for a measurement. The declared value may, at some time, be exactly the same as the true value. But that is something that can never be known, because they are on opposite sides of Carnap's equation.

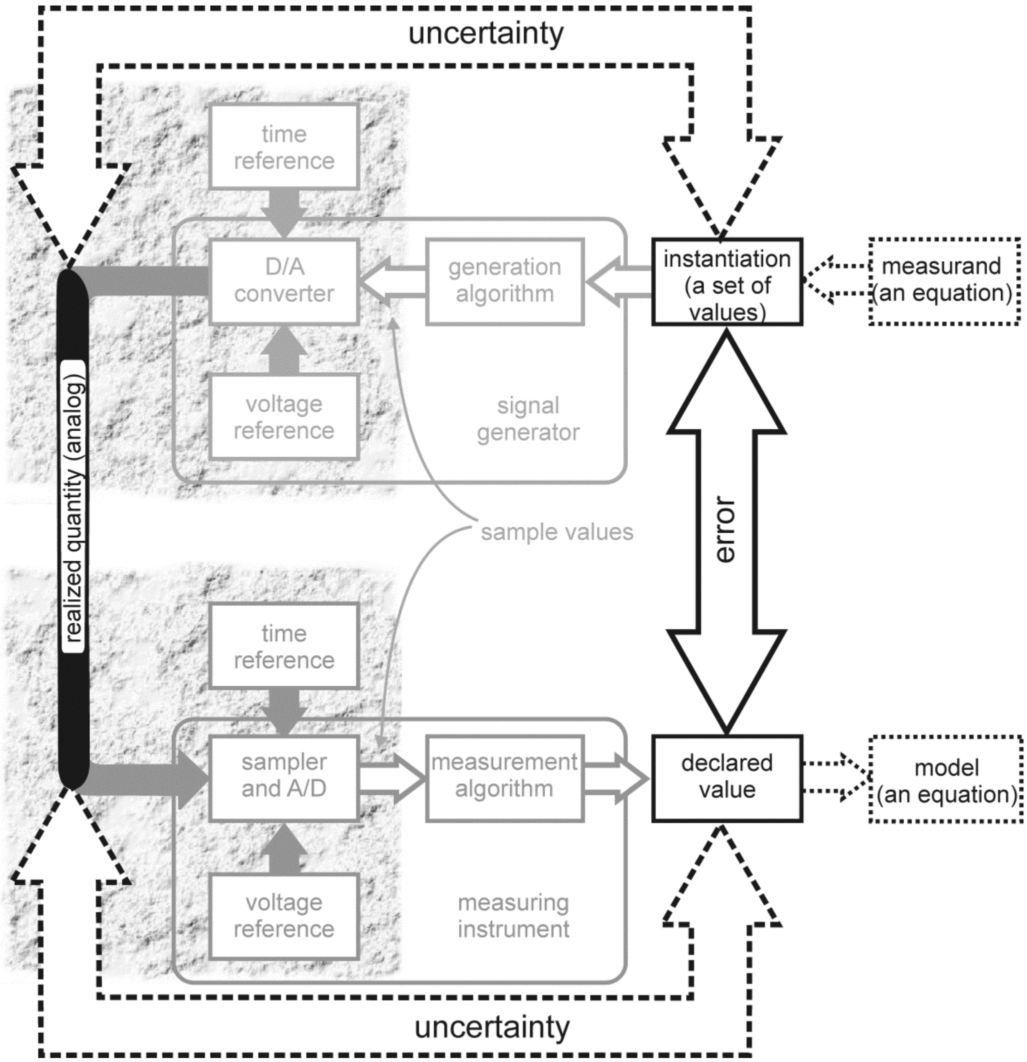

Figure 10 Calibration Framework showing Error and Uncertainty

In a measurement (as contrasted with a calibration) the uncertainty of the measuring device must be added to any uncertainty that may be associated with the source to the realized quantity (for example by an instrument transformer).

In a calibration, one uncertainty is expected to be much less than the other, so that any error can be ascribed to the unit with the higher uncertainty. By convention, the ratio (called the test uncertainty ratio) is required to be four or larger (NCSL International 174 Writing Committee (ASC Z540), 2006) if the calibration is to result in an acceptance decision.[2]

---

[2] The TUR of 4 provides a sufficient margin if the errors are considered random and are treated by combining them according to a root-mean-square calculation. Suppose the calibration equipment has an uncertainty of 1, and the measurement system an uncertainty of four. The rms value of the combined uncertainty is $\sqrt{1^2 + 4^2} = 4.12$. The larger uncertainty dominates the total, so that the smaller can be considered negligible. Some metrologists prefer a larger ratio.

Figure 8 showed noise, and we have treated that as a source of uncertainty. In fact, there are other things that can contribute to uncertainty (such as drift of electronic component values) and these should be taken into account in the estimating.

But the figure shows something remarkable. Error and uncertainty may be aspects of a single system, but they are clearly distinguishable. There is no excuse for confusing them.

Yet as recently as 2006, there was so much confusion that in their introduction to GUM (Metrology, 2008) for a report of the Australian Department of Trade and Industry (Kirkup & Frenkel, 2006) wrote

> At present, students often encounter texts that are mutually inconsistent in several aspects. For example, some texts use the terms *error* and *uncertainty* interchangeably, whilst others assign them distinctly different meanings. Such inconsistency is liable to confuse students, who are consequently unsure about how to interpret and communicate the results of their measurements.

The distinction between error and uncertainty is drawn clearly in Figure 10. There should be no further excuse for confusion.

2.3 Model Quality

In addition, and not shown in the figure, there is an unrecognized assumption in measurement (stated as an "implicit assumption" in GUM 3.4.1), that the realized quantity is an exact embodiment of some measurand that is precisely known. If that were the case, the measurement process would reveal what is sometimes called the true value. But it is never the case. There is always some inexactly known thing that the 2015 Report to DOE called *semantic coloration.*

Consider the case of the power system. In the words of (Clement & Johnson, 1960) "It has been observed that roughly 95 per cent of the electrical energy consumed in the United States comes from sinusoidal sources." That justifies the study of systems that are characterized by such signals. But the truth is that *zero* percent of the energy is generated that way. There is *always* distortion of one kind or another. Always.

Figure 10 shows a physical system which is creating a real signal that has some "true value" and some semantic coloration. The two are combined in the physical system. This combining is the sort of thing that Carnap indicated by his little circle symbol – it is a physical operation, not a mathematical one. Perhaps the signal is operated on nonlinearly, for example, or perhaps power supply ripple adds to the signal at the same frequency.

This *combined* signal is what we are measuring, and is the thing from which we hope to find the "true value." If the "pollution" were of the characteristic of noise, we might stand a chance. Since it is semantic coloration, it combines with the value we would rather have and produces a signal that has a different semantic value. There may be no way that these things can be separated once they are combined. That is what semantic coloration means.

But of course, we try to separate the pollution from the signal. We recognize that the process of measurement furnishes a value of a parameter of a *model* that we have constructed to represent the physical world. The model is an idealized and simplified thing, and it must be so if it its use is to be tractable.

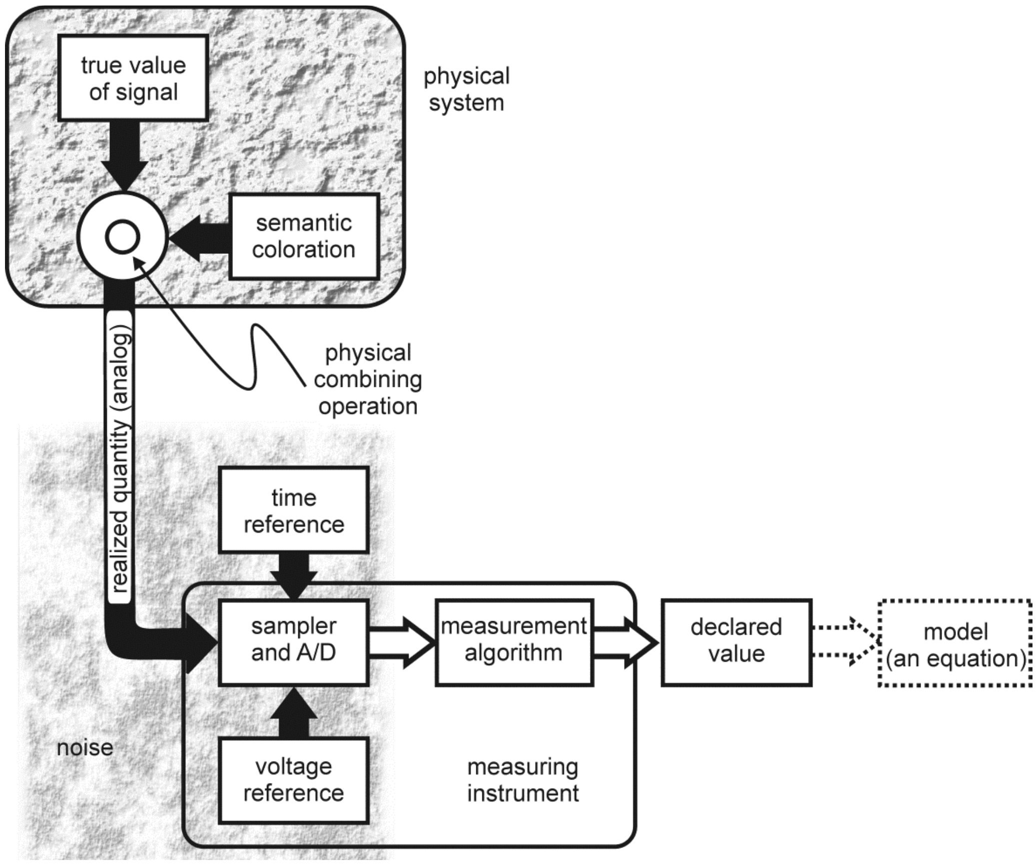

Figure 11 Measurement in Framework view, with Semantic Coloration

If there is no large amount of power carried by the harmonics on the power system, we may safely ignore their effect in our calculations. If we make a measurement that filters out the harmonics from the signal, we might stand a chance of getting closer to the underlying "true" value. The measurement algorithm in Figure 10 might then include filtering.

As (Rutman J. , 1977) observed:

> In an ideal world, the correspondence between the physical reality and the model should be perfect. In fact, the real world is so complex that many details are ignored in the model and therefore, the correspondence is incomplete and imprecise. On the other hand, some properties that have no direct meaningful counterparts in the real world have to be introduced in the model in order to make a tractable model (eg : stationarity properties of random processes). In any case, the model must be well defined from the mathematical point of view.

If we have a model of the power system that assumes operation at only the fundamental frequency, and we select out the harmonics in the selection of the model or in the signal processing of measurement, we may get a tractable measurement process, but we get a result that is not such a good representation of the real world.

It is sometimes important not to forget that.

2.4 Simulation of Calibration

In our work we considered only digital measurements and digital signal generation. What that means is that the instantiation combined with the generation algorithm can be simulated in software, and do not require any specialized hardware of the kind that might be used in a real calibration. Further, controlled noise can be added.

Further, the data stream that would be fed into the measuring instrument algorithm can be simulated, and the entire noisy arrangement of D/A and A/D converters can be bypassed. Figure 12 shows what remains.

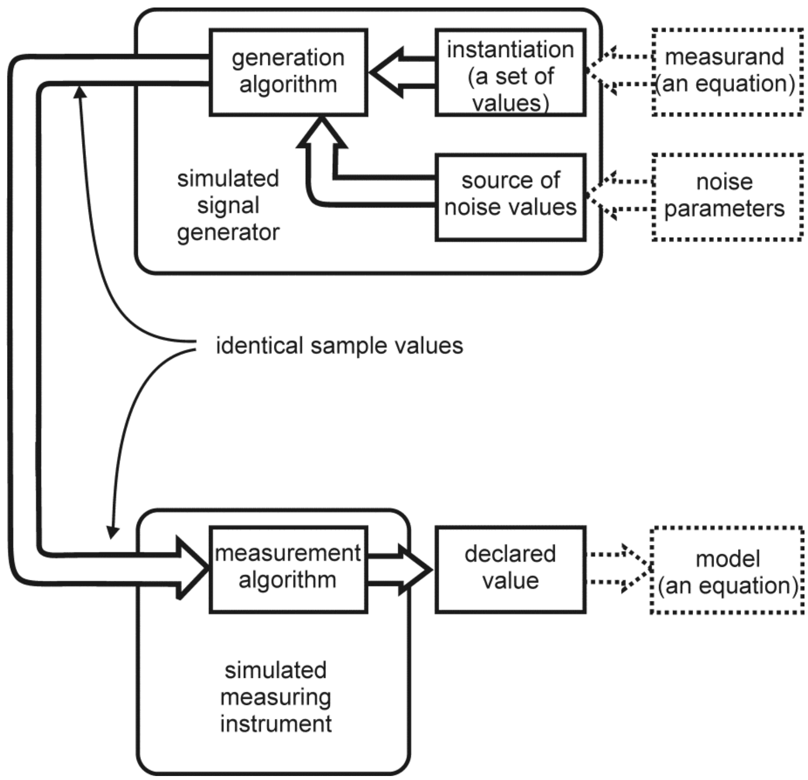

Figure 12 Simulated calibration system.

In this simulated calibration system, a perfect signal can be submitted for measurement to an algorithm. Measurement algorithms can be compared and modifications to algorithms can then be evaluated. The theoretical performance limit can be established.

In a metrology laboratory, the measurement algorithm is usually considered to be the part of the system being evaluated. For a PMU, for example, accurate timing is provided by a link to UTC (usually via GPS) and the quality of that timing is a matter of an economic trade-off. Similar remarks can be made about the quality of the A/D converters in the measurement equipment. We assume that the manufacturers are providing the digital representation of the analog input signal with what they judge to be adequate performance capability.

It is therefore fair to say that a system built along the lines of Figure 12 could be used to evaluate the performance of a PMU's algorithms.

2.5 Implementation

The "new PMU" method achieves its measurement by fitting the PMU equation to the samples of the signal. The form of the equation is fixed by the physics, so this is not an equation-fitting method so much as a parameter fitting. Some preliminary results in the Report to the DOE in 2015 showed promise: this Report gives much more complete results.

It is seemingly new to metrology to regard measurement as what mathematicians call a fitting problem, but it seems quite appropriate. The groundwork is laid by the notion of a mathematical model on one side of Carnap's equation and the physical world on the other. Some investigators have had much of this vision (Carnap, obviously, for one) without quite completing the connection to fitting. Rutman, in a review of the field of oscillator specifications (Rutman J. , 1977) made the following observation:

> At first, it is important to emphasize on the fact that the two following facets are often confused:
>
> - the real world, with its physical devices, measurement apparatus, experimental results derived from reading meters, counters, dial settings and so on.
> - the mathematical model, with the means and rules for operating with the symbols introduced in it.

In an implementation using the method of Figure 12 the signal samples that stand in for the sample values from the realized quantity must have noise and distortion added if a real-world signal is to be modeled. Therefore, we considered how to add noise, and how to give it various well-controlled properties.

The addition of random numbers to the sample values is not an effective way to represent the sort of noise of interest in measurements. Such noise may be characteristic of the problem solved by Shannon, but our experiments showed that no amount of such noise, no matter how it might be filtered, was observed as phase noise or amplitude noise, for example.

What must happen, of course, is that the measurand (model, equation) must be modified to allow for noise. If the signal is modeled as a sinusoid, represented in the time-domain by

$$x(t) = X_{\mathrm{m}} \cos(\omega t + \varphi). \tag{2-2}$$

where the amplitude X_{m}, the frequency ω and the phase φ are ideal and stationary, we must modify the equation to allow the values to change:

$$x(t) = [X_{\mathrm{m}} + \epsilon(t)]\cos(\omega t + \varphi + \Phi(t)). \tag{2-3}$$

where $\epsilon(t)$ is a random amplitude noise and $\Phi(t)$ is a random phase noise. This implementation is an embodiment of Figure 12 with the addition of noise, as needed. The top part of the figure, corresponding to the generation of the signal, was implemented in MATLAB or Excel, and the bottom part, corresponding to the measurement process, was done in MATLAB using their solver. Output files were Excel compatible.

The sequence of operations was as follows.

1. Decide on the case(s) to be run, what signal model, how many samples per nominal cycle, what changes in parameters, what noises to add. Or choose a file of oscillography results.

2. Set these numbers into the software to create the "PMU input signal."
3. Run the solution software against the input signal to create output files. These contained the results of a string of measurements. For some experiments, this software would run overnight.
4. Examine the output data, and graph the interesting results.

If noise was added to the signal, the results had a statistical nature, and a proper understanding of the results required multiple "measurements." Sometimes a meaningful statistical analysis required several hundred cases be run. The software might then be operated overnight.

Figure 13 shows the three options for the software operation.

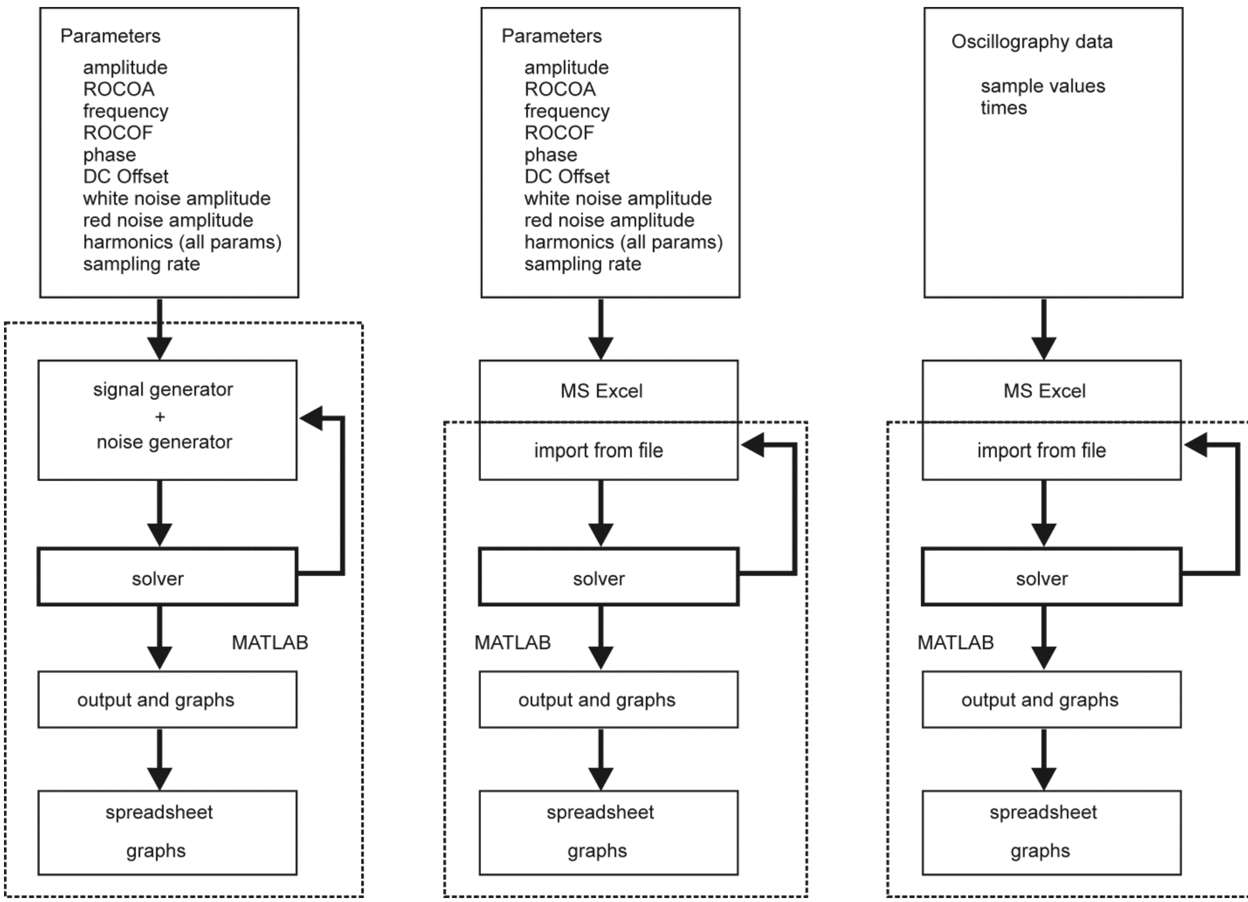

Figure 13 Block diagram of software operations

2.5.1 Input data is generated in MATLAB environment.

The capability to specify phasor or "phasor-like" parameters allowed us to create steady sine waves or steady ramping of frequency and/or amplitude. Later versions of the code allowed the addition of harmonic distortion with unlimited harmonic count. For the harmonics, amplitudes and phase shifts were separately controllable. A way to generate normal distribution white noise and red noise (Brownian noise) was implemented.

Creating input data in MATLAB allows for faster code execution and is more convenient than data import from an Excel file, which was how we began the effort. That is described next.

2.5.2 Input data is imported from *.xlsx* file.

MATLAB supports importing data from files, and the majority of our synthetic data was created in Excel software. The data was created as a large array of sine values. Spreadsheet data generation allows not only for steady state signal generation but also multiple parameter value changes in a single data stream, things like phase jumps, amplitude jumps, ROCOF value changes. A spreadsheet was also used to generate harmonics and noise (filtered and raw white noise). By generating data this way, the MATLAB script took more time to run, and changes in the data are not easy to make. On the other hand, unique changes could be made. For some experiments done this was the only way to generate the desired input signal.

Figure 14 shows how a realistic signal with harmonics and noise can be built up in Excel.-wave data points and Δt

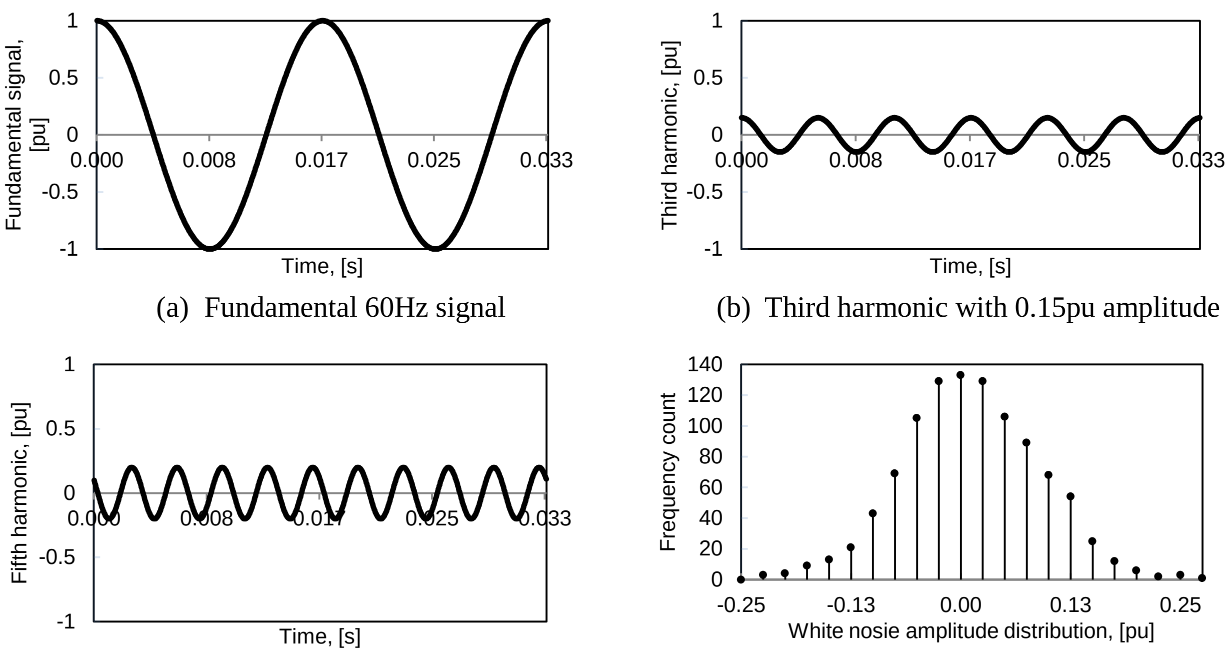

(a) Fundamental 60Hz signal (b) Third harmonic with 0.15pu amplitude

(c) Fifth harmonic with 0.2pu amplitude (d) White noise distribution

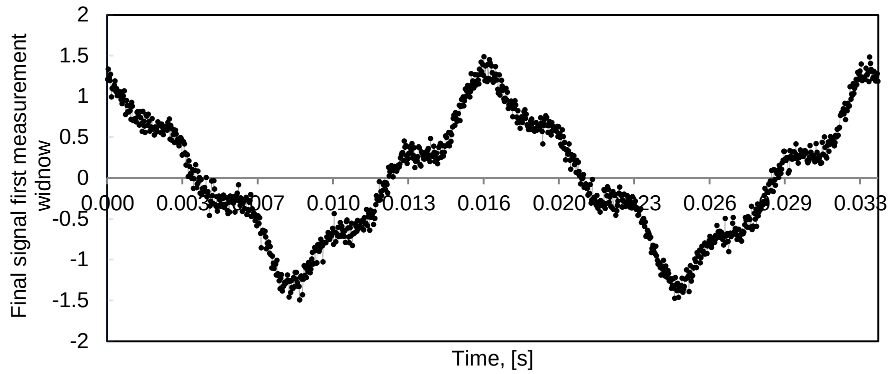

(e) Final signal imported in MATLAB

Figure 14 Signal construction in Excel

Data import from a file was, of course, the only way to import real-world oscillography data into MATLAB. We used *.xlsx* or *.csv* file extension to import the data. A conversion from COMTRADE to an Excel-compatible format was done for us by Ray Hayes of ESTA International.

Figure 15 shows a short section of a data stream (oscillography) containing a fault. The relay doing the data collection was sampling at 64 samples per nominal cycle.

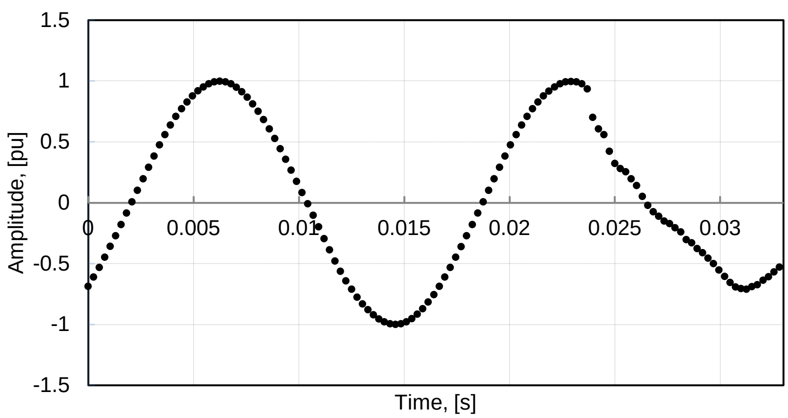

Figure 15 Real data import

2.5.3 Method – Nonlinear robust least squares fitting algorithm

To obtain the coefficient estimates, the method minimizes the overall residuals, specifically the summed square error (SSE). The residual for the i_{th} data point r_i is defined as the difference between the observed response value y_i and the fitted response value Y_i, and is identified as the error associated with the data

$$r_i = y_i - Y_i \,. \qquad (2\text{-}4)$$

The summed square of residuals is given by

$$S = \sum_{i=1}^{n} r_i^2 = \sum_{i=1}^{n} (y_i - Y_i)^2. \quad (2\text{-}5)$$

The cosine signal is nonlinear, so the equation is difficult to fit because the fitting coefficients cannot be found by simple matrix techniques. Therefore we used an iterative technique.

We assume that errors follow normal distribution and very large values are sparse. Since the method was developed to deal with real data, where outliers can occur, we used a robust least-squares regression with *bisquare weights* in order to minimize their influence. (Squaring an extreme value gives very large errors.) The further the data point is from the fitted line, the less weight it gets. Data points that are outside the estimated random chance region get zero weight (The MathWorks Inc., 2015).

For most cases, the bisquare weight method is preferred over Least Absolute Residuals (LAR) because it simultaneously seeks to find a curve that fits the bulk of the data using the usual least-squares approach, and it minimizes the effect of outliers.

The algorithm follows this procedure:

I. Start with previously chosen initial coefficients (start values).

II. Construct the curve. The fitted response value *Y* is given by

$$Y = f(X, b) \quad (2\text{-}6)$$

and involves the calculation of the Jacobian of $f(X, b)$, which is defined as a matrix of partial derivatives taken with respect to the coefficients (The MathWorks Inc., 2015).

III. Fit the model by weighted least squares.

IV. Compute the adjusted residuals and standardize them. The adjusted residuals are given by

$$r_{adj} = \frac{r_i}{\sqrt{1 - h_i}} \quad (2\text{-}7)$$

where h_i is leverage that adjust the residuals by reducing the weight.

The adjusted standardized residuals are

$$u = \frac{r_{adj}}{K_S} \quad (2\text{-}8)$$

where *K* is a tuning constant *4.685*, and *s* is the robust variance *MAD/0.6745* (*MAD* -Median Absolute Deviation) (The MathWorks Inc., 2015).

V. Calculate the robust weights as a function of *u*. The bisquare weights are given by

$$w_i = \begin{cases} (1 - (u_i)^2)^2 & |u_i| < 1 \\ 0 & |u_i| \geq 1 \end{cases}. \quad (2\text{-}9)$$

The final weight is the product of the regression weight and the robust weight.

VI. Finish if the fit has converged. If not, perform another iteration of the fitting procedure.

Adjust the coefficients and determine whether the fit improves. The direction and magnitude of the adjustment depend on the trusted region. This is the MATLAB default algorithm and is used because coefficient constraints can be specified. It can solve difficult nonlinear problems more efficiently than the other algorithms. Iterate the process by returning to step 2 until the fit reaches the specified convergence criteria (The MathWorks Inc., 2015).

2.5.4 Output data

Output data can be any variable created or imported into MATLAB code. We chose to use all phasor and phasor-like quantities as outputs together with GoF metric, calculated from residual RMS values. Usually we output arrays of results so that they can be exported to MS Excel for graphical representation. (We prefer their appearance to the MATLAB graph.)

2.5.5 Code sample for real-world data estimation

This is simple one measurement code. For different experiments additional loops and conditions were used that implemented more complex code.

```
clear
clc
%% Read data from file;
%64 samples per cycle
y = xlsread('name_of_file.xlsx',1,'AB100:AB227'); %Import data from file
T = xlsread('name_of_file.xlsx',1,'T100:T227');
%% Measurement process
[xData, yData] = prepareCurveData( T, y );      %preparing data for curve fitting
% Set up fittype and options.
ft = fittype( 'a*cos(2*pi*c*x+2*pi*(d/2)*x*x+e)', 'independent', 'x', 'dependent', 'y' );
%Fitting equation
opts = fitoptions( 'Method', 'NonlinearLeastSquares' ); %Set method
opts.DiffMaxChange = 0.0001;                         %Max step change
opts.Display = 'Off';                                %Disable display option
opts.Lower = [0 -50 55 -5 -3.14159265358];           %Lower trust region boundaries
opts.MaxFunEvals = 1000;                             %Maximum evaluations allowed
opts.MaxIter = 1000;                                 %Maximum iterations
opts.Robust = 'Bisquare';                            %Select bisquare robust fitting
opts.StartPoint = [1 0 60 0 0];                      %Start values
opts.TolFun = 1e-8;                                  %Termination tolerance for the function
opts.TolX = 1e-8;                                    %Termination tolerance for x
opts.Upper = [1.5 50 65 5 3.14159265359];            %Maximum trust region values
%Call for MATLAB solver for curve fitting with selected options and outputs
[fitresult, gof, fitinfo] = fit( xData, yData, ft, opts );
%OUTPUT
RMS = 20*log10(1/gof.rmse); %Calculated GoF values
f = fitresult.c;            %Frequency values
A = fitresult.a;            %Amplitude values
ph = fitresult.e;           %Phase values
C_w = fitresult.d;          %ROCOF values
%% write data to file
filename = 'name_of_file.xlsx';
xlswrite(filename,A',1,'A1')
xlswrite(filename,f',1,'B1')

%%FILTERING
%Import the signal from file
test=xlsread('name_of_file.xlsx',1,'A1:A30720');
```

```
Fs = 30720;
%Sampling Frequency (samples per second)
%%Butterworth Lowpass filter designed using FDESIGN.LOWPASS
fpass = 3;          %Passband Frequency
Fstop = 100;        %Stopband Frequency
Apass = 1;          %Passband Ripple (dB)
Astop = 6;          %Stopband Attenuation (dB)
match = 'stopband';
%Band to match exactly
%%Construct an FDESIGN object and call its BUTTER method.
h  = fdesign.lowpass(Fpass, Fstop, Apass, Astop, Fs);
Hd = design(h, 'butter', 'MatchExactly', match);

ttt = filter(Hd,test);   %Filtering signal
```

2.6 PMU Theoretical Limits of Performance

It should come as no surprise to find that the setup of Figure 12 reveals that the method proposed in (Kirkham & Dagle, 2014) works *perfectly*. That is to say, when the measurement algorithm is a fitting method, adjusting the values of the "PMU equation," the solution is found with errors so small they are at the level of the accuracy of the computer.

For all *practical purposes, the method produces error-free results.*

Complete perfection would be an overstatement; we must add a little qualification. When the signal being measured is a pure sinusoid, the method produces error-free results. When the signal is a sinusoid with non-zero rate of change of frequency, the method produces error-free results. When the amplitude is allowed to change (this goes beyond the PMU's purview) the results are still error-free. It is only when the signal includes components that are not described by the "PMU equation" that there are errors.

2.6.1 The PMU Equation

The PMU equation, assuming the amplitude is not allowed to change, is

$$x(t) = X_{\mathrm{m}} cos\left\{\left(\omega_{\mathrm{ALF}} + \frac{C_\omega}{2} t\right) t + \varphi\right\} \quad (2\text{-}10)$$

where X_m is the signal amplitude, ω_{ALF} is the apparent local frequency (the word "local" applying to time), C_ω is the rate of change of frequency and φ the phase. The term ω_{ALF} includes terms in changing phase, and apart from the C_ω term, the equation retains the general form of the well-known phasor equation

$$x(t) = X_{\mathrm{m}} \cos(\omega t + \varphi) \quad (2\text{-}2)$$

Note that C_ω (which deals with the acceleration of the power system) is omitted in this version of the equation. The matter is dealt with in two of the papers that follow (Abstract 3.5, Paper 4.5: Abstract 3.7, Paper 4.7).

In a calibration such as shown in Figure 12, the measurement algorithm can be designed in many ways. The "new PMU" method treats the problem as a fitting problem in mathematics, as we will see below. The commercial PMU does not. But no matter, the results of the measurements are expected to be the same. If we ask for a value for the power system frequency, for example, we do not expect to get an answer that depends on how the measurement was made.

The point is that the "measurement" of the signal parameters *can* be done by curve fitting. A somewhat unexpected finding is that in theory, given the method of curve fitting, far fewer sample-value points are needed to make the measurement than are typically found in PMUs. In fact, if three parameters are to be found, the total number of sample values needed is three. With just three points, and three parameter values to be found, the fitted curve is exact: the fitting method is, of course trying to minimize the residuals, the differences between the numbers generated by the equation and the sample values furnished as data. Three points can be fitted *exactly* by an equation (of the proper form) with three adjustable parameters.

The difference between the number of samples taken of the signal and the number of parameters to be found is known as the *number of residual degrees of freedom* (Cuthbert & Wood, 1971). If there are more points and they are "perfect" in terms of timing and amplitude, the fit can still be perfect. It is only when the additional sample values mean that the signal changes from its pure cosine form (including some defined changing parameters) that the fit becomes less than perfect.

2.6.2 Changing Window Width

We have shown that with a very few samples and a very short observation window, the PMU equation can be solved with good accuracy by fitting. That, at least, is the case with a "perfect" noise-free signal. It is interesting to see how the method works when the signal is a real-world one.

To explore the limits of real-world performance, we examined oscillography data from AEP. The oscillography records are one second in duration, and are triggered when a fault is detected. Since the fault is cleared long before the end of the oscillography record, the data present the opportunity to examine measurement performance under relative normal and relatively distorted conditions. Figure 16 through Figure 19 shows the measured results for the frequency.

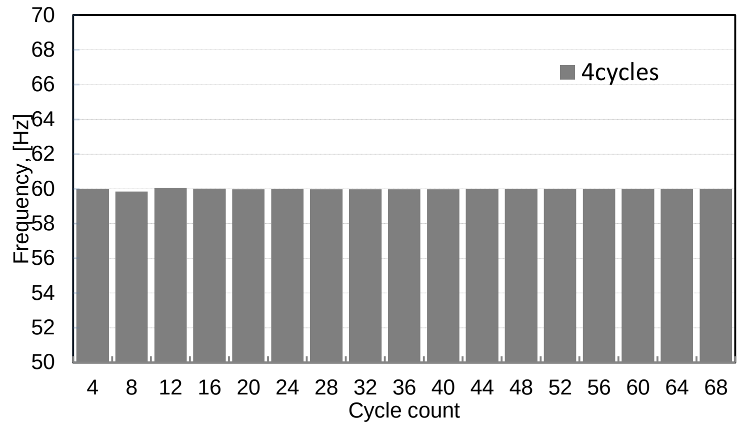

Figure 16 Oscillography results, window at four cycles

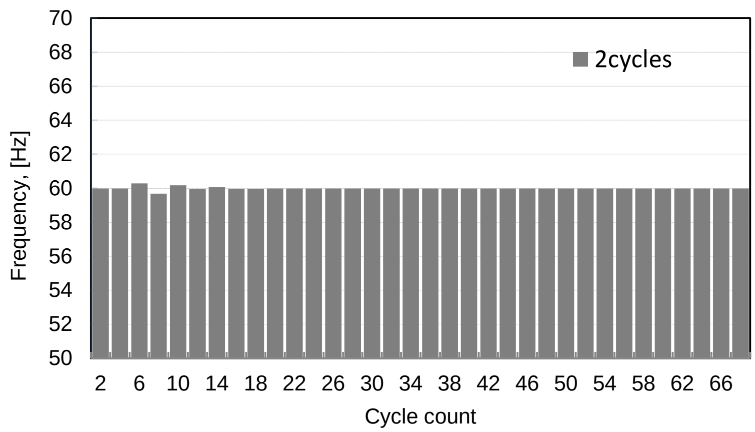

Figure 17 Oscillography results, window at two cycles

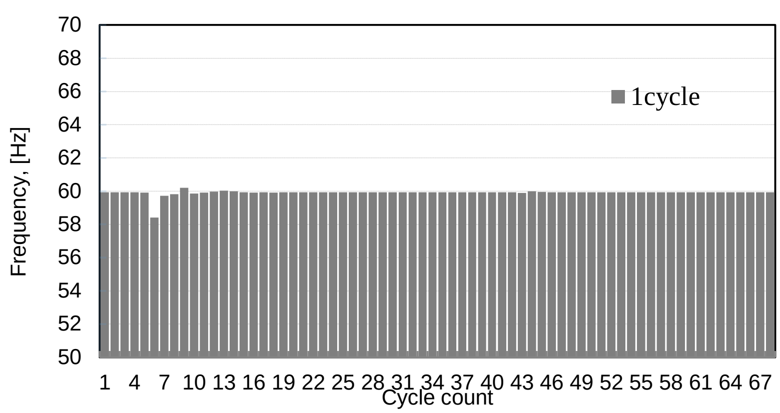

Figure 18 Oscillography results, window at one cycle

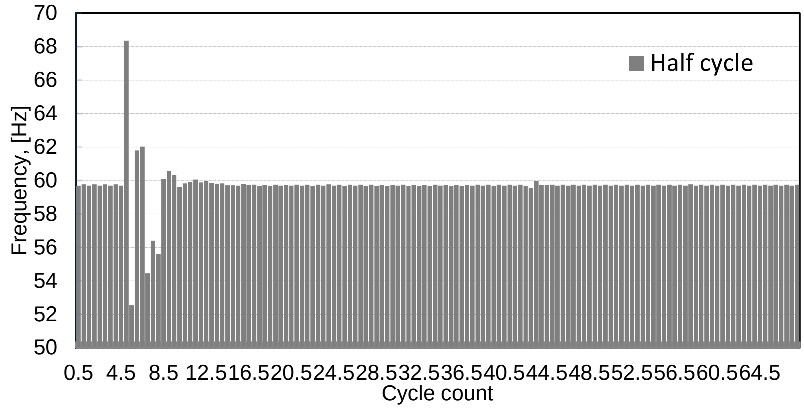

Figure 19 Oscillography results, window at half cycle

The graphs show a clear trend toward more “detail.” With a four-cycle window, the frequency seems hardly to be disturbed by the fault. With the window down to half a cycle, the frequency is quite erratic. *However, there is no useful information in all this detail*. In the region of greatest apparent frequency excursion, the waveform is as shown in Figure 20.

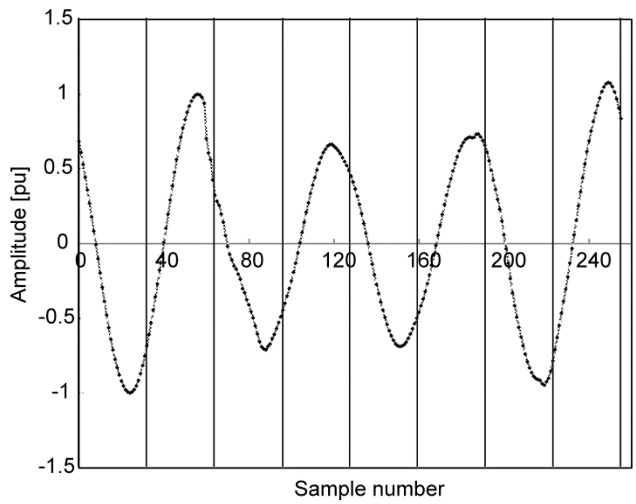

Figure 20 two cycles of waveform with fault

The vertical lines across the graph mark the boundaries of the half-cycle measurements. For the results in Figure 16 through Figure 19, the “new PMU” was set up exactly like a PMU would be: the model amplitude was assumed constant during the window. Since it is abundantly clear that the real amplitude does in fact change across the two cycles shown here, it follows that the Goodness of Fit is not particularly “good.”

A frequency of 68 Hz, followed a half-cycle later by a frequency of 52 Hz is something not physically possible. Yet that is what is indicated in the graph, and is the result of a measurement known to give the best estimate in the least squares sense.

It is crucial at this point to remember what is being done in the act of measurement. As we established, the measurement finds the value of the coefficients of the model used to represent the measurand. In the case of the PMU, that model is the PMU equation

$$x(t) = X_{\mathrm{m}} \cos(\omega t + \varphi) \tag{2-2}$$

In other words, we are asking the measurement system to answer the question: assuming that the signal is a perfect cosine wave, and the sampled data are accurate, what is the amplitude, the frequency and the phase? We should at this point recall the words of Heisenberg: We have to remember that what we observe is not nature in itself, but *nature exposed to our method of questioning*.

Our method of questioning makes no sense. The signal is obviously not a perfect cosine wave. It does not repeat periodically. It does not have a fixed amplitude. And yet we demand of the PMU to tell us the answer as if all those things were true.

Put another way: it is perfectly possible to measure the PMU parameters in just a half cycle of time.[3] But the usefulness of the result is likely to be less than the value obtained with a four-cycle window. It is as if we expect that the original signal is somehow cobbled together from little pieces of independent cosine waves. It is not even as if knowing the harmonics of the fundamental would help: the idea of harmonics is something based on the notion of periodicity. This is obviously not a situation in which periodicity is applicable. The problem is basic: *the phasor equation is the wrong model.*

2.6.3 Independent Measurements

There has been a push from some in the PMU community to obtain results ever more rapidly. We have shown above that that is complete nonsense. If the results are "smooth," the data may be meaningful but there is no more information in the repeated statement of results. If the results are not smooth (as in the half-cycle results here, there is *less* useful information in them, because the wrong question is being asked.

It is possible, perhaps even likely, that that the push comes from a misunderstanding of the concept of independent measurements. If the window of observation is advanced for each measurement by an amount less than its width, the measurement results are not independent. Each will contain some information from the previous measurement, or perhaps more.

For example, we have heard people demand a reporting rate of 120 reports per second. That is a half-cycle rate, as in Figure 19. However, it the PMU making the measurement is a conventional P-class PMU, the window width is two cycles. Therefore, the result of each measurement contains information from the seven previous half cycles. This is far from an independent measurement. How such a thing should be handled in a data set is not something that the customary statistics could handle.

Such a thing is technically possible to do, but its meaning would be far from clear. We give an example in Figure 21. It is assumed that the reporting time is the center of the window, and therefore the report is nominally one cycle after the timestamp on the report. Note that the scale has been changed from the scales used in Figure 16 through Figure 19.

---

[3] It should not be thought this half-cycle time represents any particular limit. The limit is set by the noise on the signal, and we (the PMU community) have no knowledge of the noise on power systems. Our study stopped at half a cycle for convenience. Our point was demonstrated.

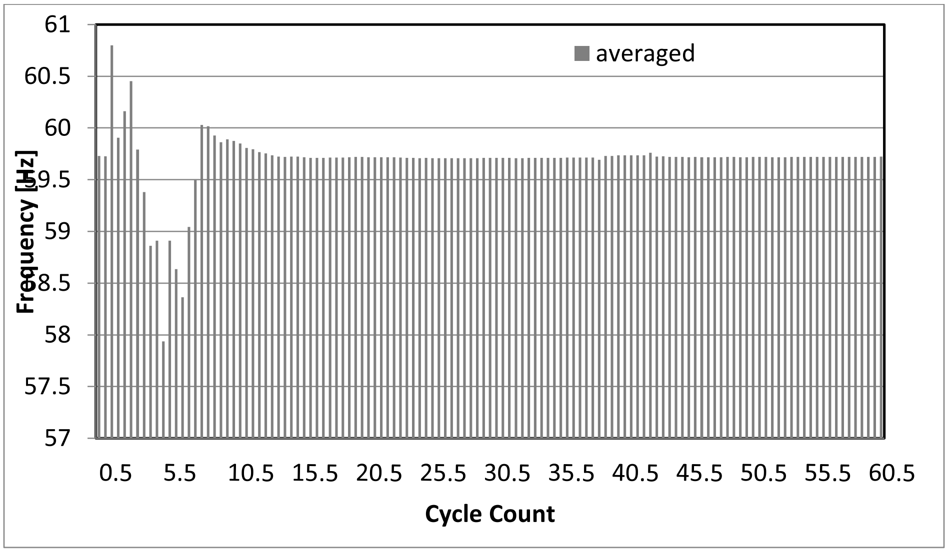

Figure 21   Two cycle averages of half-cycle results

A large “spike” in frequency is evident shortly after the graph begins. This is the report that is the first to contain the apparent 68 Hz report from the half-cycle results. Because of the averaging effect, its magnitude is reduced from the original 68 Hz.

The next report contains the 52 Hz report of the half-cycle results, and this reduces the effect of both the high- and the low-frequency anomalies.

The reader may be interested to note that the result of this averaging does not produce the same result as was obtained by the two-cycle window measurement. This data is plotted in Figure 22 for comparison.

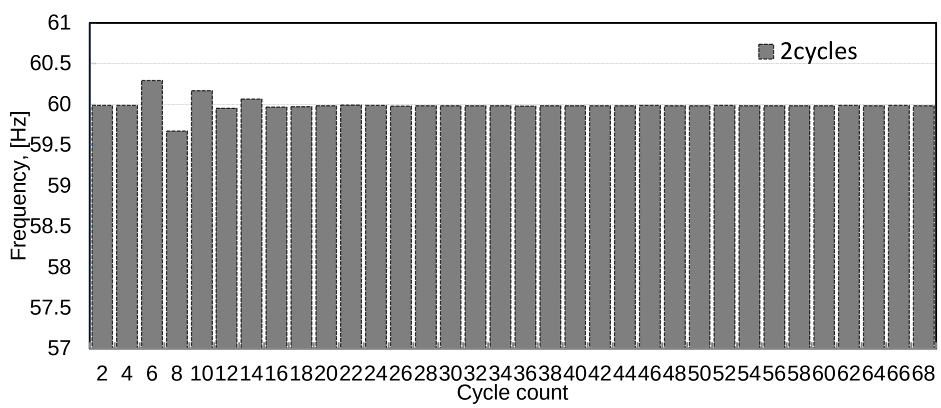

Figure 22 Two-cycle data, re-scaled

Not only does there appear to be more “detail” in the two-cycle-averaged half-cycle graph, the numbers are not the same as the two-cycle results. The variance is much greater.

It is probably fair to say that *all* that is being obtained by the process is an increase in the variance. The *meaning* of the results is not clear at all. Nor is it obvious how filtering the averaged results of the independent measurements would produce understandable data.

An important conclusion from our work, supported by this analysis of a real signal, is that the push for ever more rapid measurements should be resisted. Those asking for the fast measurements should be asked what it is they expect to find. If the truth is explained, they may realize they are unlikely to find what they hope for.

3.0 The “van der Pol Problem”

3.1 Background

The term “phase” is nowadays reserved for φ, the time invariant part of the argument of the cosine in Equation (3-1).

$$x(t) = X_\mathrm{m} \cos(\omega t + \varphi) \tag{3-1}$$

It was not always so, and up to about fifty or sixty years ago, the word phase might have meant the complete argument. An exploration of the terms was undertaken by the great B. van der Pol. Brought out of the Netherlands toward the end of WWII, he gave a paper in 1945 on frequency modulation in which he explored the topic (van der Pol B. , 1946). In this paper, van der Pol shows that over the years many investigators preferred to refer to the entire argument of the cosine term, that is ($\omega t + \varphi$), as the phase. He argued that considering this pair of terms as the phase allowed the phase of signals of different frequency to be described mathematically.

If we use the “total phase” notion, and denote this phase as Ψ, we can rewrite Equation (1) as

$$x(t) = X_\mathrm{m} \cos \Psi \tag{3-2}$$

This is a more general representation than (3-1) because there is nothing about the equation that requires the total phase to be a linear function of time, as (3-1) does. That is worth remembering, since the real world of electric power signals is not likely so well-behaved. It is interesting to note that both (3-1) and (3-2) are used in the IEEE PMU standard (IEEE C37.118.1, 2011).

3.2 Instantaneous Frequency

In a review of the term “instantaneous frequency,” (Boashash B. , 1992) comments on the 1945 paper. He observes that van der Pol approached the problem of formulating a definition for the instantaneous frequency by analyzing an expression that is fundamentally the same as (3-1):

$$y(t) = A\, cos(\omega t + \varphi) \tag{3-3}$$

where A is amplitude, ω is frequency of the oscillation, φ is a phase constant, and the argument of the cosine function, namely $(\omega t + \varphi)$, is the phase. van der Pol went through a sequence of defining the various modulation methods, changing each coefficient in turn. For example, amplitude modulation can be written by modifying (3-3) so that the amplitude, A, varies as a function of t, as in (3-4):

$$A(t) = a_0[1 + mg(t)] \tag{3-4}$$

where $g(t)$ is the (audio) modulating signal, and m a coefficient he identified as the modulation depth. Similarly, he defined phase modulation by

$$\varphi(t) = \varphi_0[1 + mg(t)] \tag{3-5}$$

van der Pol then turned his attention to frequency modulation. Having used essentially the same expression to modify the amplitude and the phase terms in the phasor equation, he noted that it would be erroneous simply to substitute ω in (3-3) by

$$\omega(t) = \omega_0[1 + mg(t)] \tag{3-6}$$

because it “would lead to a physical absurdity.” Boashash points out that by substituting (3-6) into (3-3), the resultant phase does not yield (3-5). van der Pol reasoned that expression (3-3) for harmonic oscillations must be rewritten in the form:

$$y(t) = A\,cos\left(\int_0^t \omega dt + \theta\right) \tag{3-7}$$

where the whole argument of the cosine function is the phase $\psi(t)$.

3.2.1 Physical Absurdity?

We did not regard the arguments an discussions of van der Pol and Boashash as particularly enlightening. For several years, in our work to make PMU measurements based on the notion that the measurand should be a mathematical specification (Kirkham & Dagle, 2014), (Kirkham H. , 2015), (Kirkham H. , 2016) (Riepnieks & Kirkham, 2016), (Kirkham H. , 2016), an implementation of (3-6) has been used. The purpose has been to create a test signal for a PMU implementation. The “physical absurdity” has been observed, but it took more thought to understand it.

It is worthy of further explanation and comment.

The PMU is required to measure four things. Three are the three “fixed” parameters of a phasor (the amplitude, the frequency and the phase). The fourth is the rate of change of frequency. (When work on making the first-ever PMU began, it must have seemed reasonable to assume that the frequency might vary, and that presumably its rate of change would be something of interest.)

If we start with the phasor equation, and add only the change that allows the frequency to change, we define what might be called the “PMU equation”:

$$x(t) = X_{\mathrm{m}} cos\left\{\left(\omega + \frac{C_\omega}{2}t\right)t + \varphi\right\}, \tag{3-9}$$

where X_{m} is amplitude, ω the frequency, C_ω the rate of change of frequency (ROCOF) and φ is the phase. Note that with no rate-of-change variable, the equation is that of a phasor.

However, the equation has the form of Equation (3-6), which was described by van der Pol as leading to a physical absurdity. The various papers already cited show that the results of measurements based on this equation *do not suffer any physical absurdity*.

Yet van der Pol was correct.

Understanding this matter is what we have called the "van der Pol problem."

3.2.2 Solving the van der Pol problem

As we know, the PMU measures its input signals for a duration known as a *window*. By analyzing the information in the window, it outputs values for the three phasor parameters plus ROCOF. In the testing of the fitting-method PMU, a signal is generated synthetically and fed directly to the PMU algorithm. In a real PMU, the input signal is analog, and the PMU converts to digital by the usual A/D converters. The PMU of interest in our work has not used any A/D converters, relying instead on this synthetic signal.

The PMU examines the incoming data stream by looking at sections of it that correspond to measurement windows in time. It is not a real-time method, however, and the synthetic signal is pre-calculated and stored for use. Once a window of data has been solved, the results are output and the next window is examined.

The problem lies in the creation of the synthetic signal. Suppose we use a spreadsheet to create a signal based on (3-6) or 3-(9), and suppose we begin with $C_\omega = 0$. After some time, say t_k, we give a value to C_ω. From $t = 0$ to $t = t_k$ the spreadsheet values follow in sequence and describe a cosine wave. At t_k the sequence of numbers looks very much like those of a cosine wave, but the rate of change of frequency is subtly changing the phase. (Recall that phase is measured with respect to a fixed-frequency reference.) So far, so good.

Suppose that at some later time t_m we decide to reset the C_ω to zero, so that altogether we have created a sequence that should produce a ramped step in frequency. The spreadsheet has no problem generating the numbers, because it is given equation (3-9). However, at the *second* change in C_ω the phase no longer has the value that it started with at $t = 0$.

There is therefore a phase jump in the sequence of numbers, because the spreadsheet generator is still solving (3-9) with the original value of φ to output its number sequence. The PMU may very well measure the correct phase relative to the reference within its window, but our synthetic signal generator is out of touch with the constraints on a real signal. One measurement result may well reflect a phase jump.

If instead of simply looking at (3-9) with constant values for the coefficients, it was instead solving for a new frequency and phase at each step, the problem would not exist. That is in essence what the real world would do, and that is what (3-7) describes. A signal generator operating so as to simulate the real world in this way would bypass the van der Pol problem.

It is worth noting that the problem is evident only at the *second* (or later) change in ROCOF, if the first change is the one at which ROCOF becomes non-zero. While the frequency is constant at the reference value, the relative phase does not change, and there is no phase jump when the step in ROCOF takes place. For this reason, tests on the fitting-PMU have tended to be of this nature.

3.2.3 The van der Pol problem at large

We cannot end this discussion without noting that others have failed to observe and account for the problem. A paper was brought to our attention (as a possible help when the computer of the editor of the IEEE Transactions told us that we had too few references, itself a dismally unintelligent comment). We will not cite the paper in our bibliography, because we do not wish to give it another citation. (Citations are assumed good, by default.) The paper, entitled “A Fast Recursive Algorithm for the Estimation of Frequency, Amplitude, and Phase of Noisy Sinusoid” by P.K. Dash and Shazia Hasan, appears in the Transactions on Industrial Electronics, Vol 58, No 8, October 2011. It includes the graph shown marked up in Figure 18.

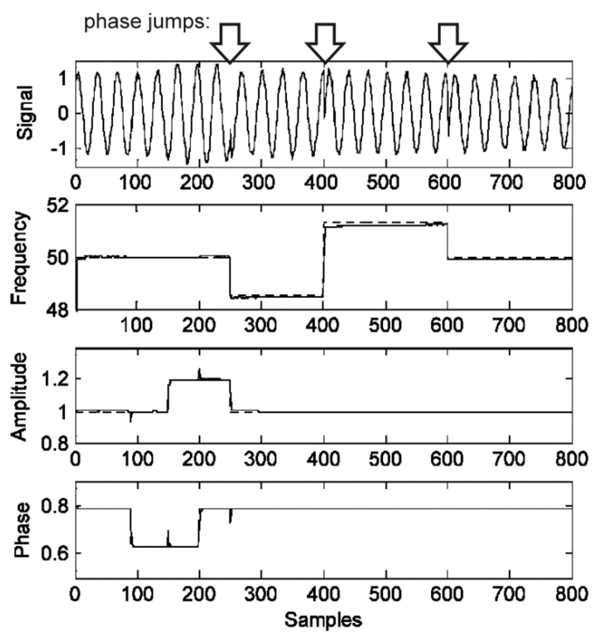

Figure 23 Fig 4 from P.K. Dash paper

We added the fat arrows at the top marking the phase jumps that are evident in the generated wave labelled “signal” at the top. They correspond to changes in the frequency parameter. This means that the waveform is not likely to be representative of anything in the real world. But that is not the problem.

More importantly, the supposedly measured phase (at the bottom of the graph) shows no evidence of being aware of the phase jumps. For example, the large and very clear phase jumps at samples 400 and 600 are completely absent from the output on the curve at bottom.

We conclude that the measurement algorithm that is the subject of the paper is actually non-functional. The claims made for its excellence in the Conclusions of the paper are not supported by the evidence presented.

We will discuss with the editor in question the possibility of having this paper tagged by an “expression of concern” in IEEExplore, along with several others by the same first author that exhibit similar shortcomings. We have never come across such a thing in IEEE, but it is a protective feature of the medical research field.[4]

A suitable expression of concern might be something like this:

> The Editors have received information that the paper entitled “A Fast Recursive Algorithm for the Estimation of Frequency, Amplitude, and Phase of Noisy Sinusoid,” published by this Journal [citation], draws conclusions that are not supported by the graphs presented. The Editors of the Transactions bring this problem to the attention of readers and suggest that this information should be taken into account in making reference to this paper, and in judging its content.

We considered asking for a retraction. However, while there is no doubt that the paper contains nonsense, some of the responsibility for the fact that it got published lies with the IEEE. A retraction implies (we think) that the authors were deliberately trying to deceive, and that is not a claim that we make. A retraction might be appropriate for a situation in which plagiarism had been committed, or deliberate changes made to data so as to affect the outcome of a study. This does not seem to be the case. The publisher (IEEE) should not be allowed to cause harm to the authors because they (IEEE) published something that did not make sense.

There are two lessons. First, professors really should pay attention to what their students produce. The paper by Ayrton and Haycraft (Ayrton & Haycraft, 1895) is well crafted and describes a well-designed experiment. The paper by Dash and Hasan clouds the issue with pages of unintelligible mathematics, and presents results that undermine its own conclusions. Second, it is incumbent on us all is to be better reviewers, and to stop this sort of erroneous paper being ever made public. A good review would have identified the problem and rejected the paper.

[4] A series of flowcharts of associated processes in the medical field may be found at the Website of the Committee on Publication Ethics at http://publicationethics.org/files/Full%20set%20of%20flowcharts_0.pdf (accessed 5/23/2016)

4.0 Future Work

4.1 Noise and the Allan Variance

In the 2015 Report to the DOE the Allan Variance was explained. That is a scheme of evaluating the results of a series of measurements using a two-sample variance, one that looks only at two samples at a time. In particular, the method assumes that the signal parameters are constant while the measurement is re-made with a different width of the sampling window in the measurement system. Very often the variance of the readings would change in a particular way, as shown in Figure 27.

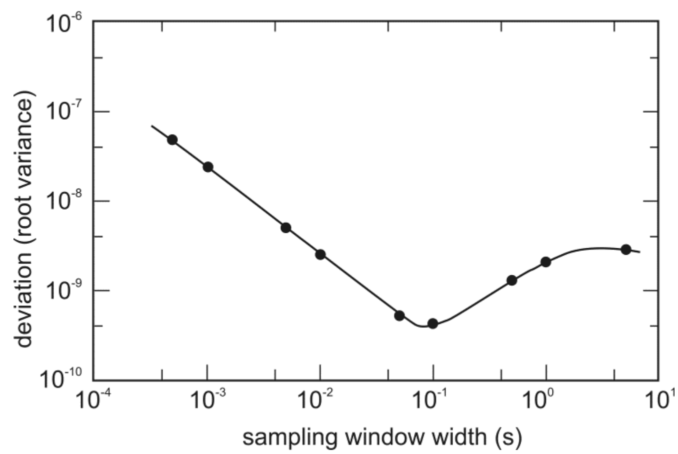

Figure 24 Deviation as a function of window width

What is generally observed is that if the measurement duration is made short, the effect of noise on the signal predominates. If the measurement window is longer, other effects (called semantic coloration in the 2015 report) start to become evident in the result. The lower limit of the variance curve establishes what may be thought of s the noise floor of the measurement.

The graph shape is something that depends on the noise in the system. For window width one may use the detector time-constant as an analogy from an analog measurement system: the longer the time constant, the more sensitive thee measurement. (However, one should also beware the results of too long time constant, as discussed in the 2015 report.)

In an analog system, the signal is typically not windowed, it is always connected to the instrument. The reading on the display is filtered by the time-constant of the detector and the response-speed of the indicating instrument. It has long been known that if the noise was random, the average value was expected to be zero, so designs for analog instruments quite often made the detector time constant as long as possible, consistent with not boring the user with a too-long response time.[5]

[5] Users of these analog measurement devices sometimes became adept at interpreting the movements of the needle. Vibration might indicate ripple on a power supply, for example. With the advent of digital measurements, things

We have already noted that the act of measurement of PMU parameters takes a finite amount of time. It may be inferred that the signal being measured must not change during the Allan Variance measurement. It is, of course, assumed constant for the duration of a window, so the declared value will apply. If it is not constant for longer, the notion of finding a statistical distribution for repeated measurements does not work. The notion of changing the window duration also does not work, as the signal parameters will be different if the signal changes.

These observations lead to the conclusion that the real power system is unlikely to be useable as a way to establish an Allan Variance curve. Indeed, in most experiments to do that, the equipment is set up very carefully indeed to ensure stability, because very often thousands of measurements are needed at the same parameter settings to allow noise averaging. For these measurements, the signal must be stable and the noise assumed stationary.

Our synthetic signal source is capable of allowing the addition of controlled noise, and altogether can satisfy the requirements on the signal and the noise for doing an Allan Variance study. What is not known, however, is the level of noise that the power system includes. Until that is better understood, no useful Allan Variance study can be done.

Preliminary work, however, has revealed that the optimum window is not at all likely to be the same for the various parameters being measured. ROCOF, for example, almost certainly could not be measured by a P-class PMU in any real power system.[6]

4.2 Sampling Rate

The effect of sampling rate has not been studied, as far as we know, with respect to PMU performance in a noisy environment. As with the Allan Variance, such a study would require stability on the part of the signal being measured and stationarity of the noise being added.

It is our opinion that a two-sample variance similar to the Allan Variance would be a worthwhile study. We have named the variance after its primary inventor, the second author of this report, Artis Riepnieks.

The Riepnieks Variance will form part of the study for Artis's PhD from Riga Technical University, in Latvia, his home country.

changed. The features that corresponded in the digital instrument were the sampling window, and the refresh rate on the display. But it is scarcely possible to discern the reading on a digital display that is changing rapidly.

[6] It should be noted, however, that the commercial PMU does not give its results as independent measurements. Indeed, the declared values are heavily filtered, and it may be that some useful indication of ROCOF is possible, though it can scarcely be attributed to any one particular sampling window.

4.3 Double exponent

4.3.1 Background

The report so far has concentrated on the characterization of the alternating signals measured by the PMU. However, once it is accepted that the measurand should be an equation, and because the equation can be solved by the fitting method, it is clear that other signals can be subject to similar treatment. We consider the case of high voltage testing.

High voltage testing is governed by IEEE Standard 4-2013, [IEEE Standard for High-Voltage Testing Techniques, IEEE, May 2013] one of the oldest IEEE standards, and one that is periodically updated to reflect advances in knowledge and technology.

High voltage testing is often a more obviously statistical kind of measurement than other measurements: one may be interested in the gap that will withstand 50% of the high voltage pulses applied, for example. That information would establish the middle of a statistical distribution (assumed normal) of flashover probabilities, and once the variance is known, the voltage for any given flashover likelihood can be calculated for the gap in question.

We set aside the measurement of direct high voltage and of alternating high voltage, and concentrate on impulse measurement. The standard informs us that the impulse waveshape is represented by a double exponent. The equation is

$$V(t) = Ve^{-t/\alpha} - e^{-t/\beta} \tag{4-1}$$

The time constant α establishes the rise of the impulse. For lightning, numbers on the order of a few μs are used. The time constant β controls the "tail" of the pulse, and for lightning is a few tens of μs. When the time constants α and β are 1 and 100 μs, the curve shown in Figure 28 is generated. Switching impulses are much slower, typically by a factor of about 1000.

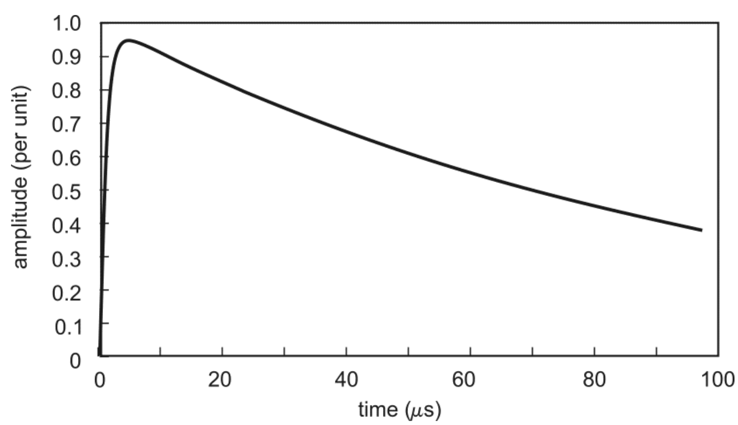

Figure 25 double exponent impulse, time constants 1.2 and 50 μs

We have shown that the curve fitting method is capable of fitting the double exponent waveform of the figure. An example of a fit (with noise added to the impulse) is given in Figure 29. The residuals are, in essence, a copy of the noise. The fit to the "clean" signal is almost at the level of machine accuracy.

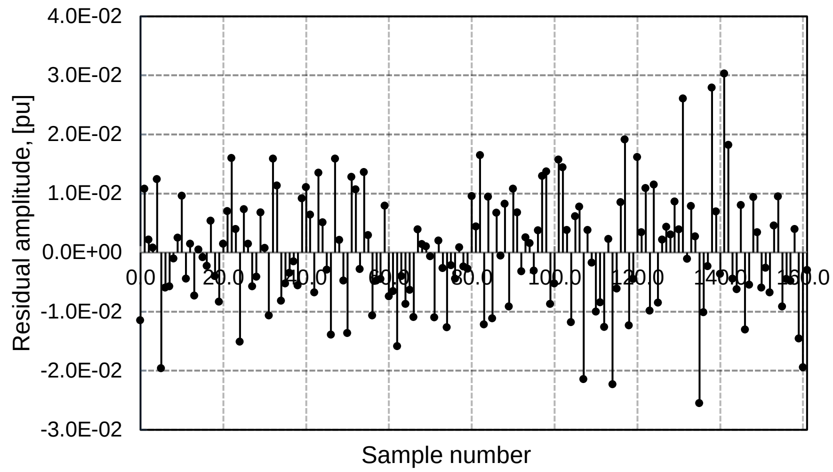

Figure 26 Residuals of fitted double exponent

In practice, impulse generators do not create such "perfect" impulses as shown in Figure 28, in particular at the beginning. The region below an amplitude of 30% of the peak is often far from the curve generated by the double exponent. The impulse is therefore characterized as shown in Figure 30.

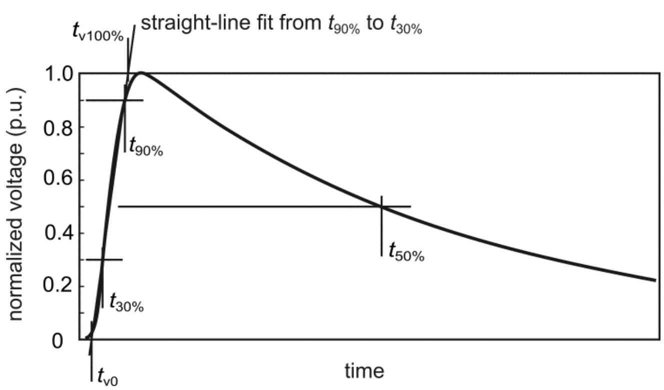

Figure 27 Impulse waveform as given in IEEE Std 4

Because of the modified shape of the wave, for the purposes of testing, the impulse is not *defined* by the two exponents. What is needed is the peak voltage, the *front time* and the *virtual steepness*. The way these terms are defined suggests the use of a fitting method such as ours, although the definitions surely arose long before curve-fitting could be done by computer.

According to the standard, a best-fit straight line is drawn for signal between 30% and 90% of the peak value. The front time of a lightning impulse is defined as the 1.67 times[7] the duration for which the signal is between 30% and 90% of the peak value on the test voltage curve.

At the time of writing, we have tested only the ability of the method to fit the exponents. In Figure 31, a synthetic noisy signal is shown.

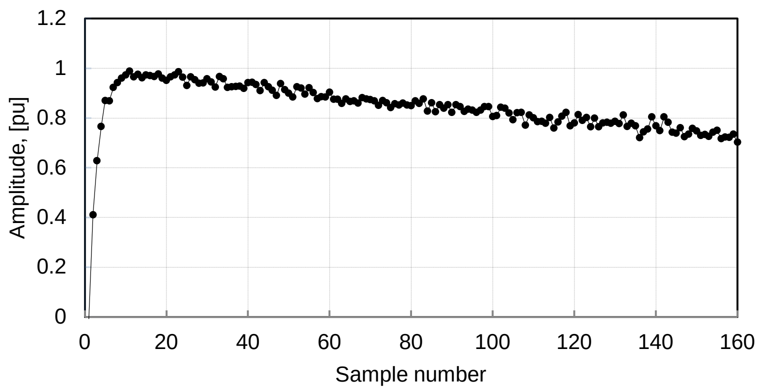

Figure 28 Synthetic double exponent with 5% Gaussian noise on sample values

The fit obtained was quite acceptable, with residuals many orders of magnitude down on the signal, as shown in Figure 32.

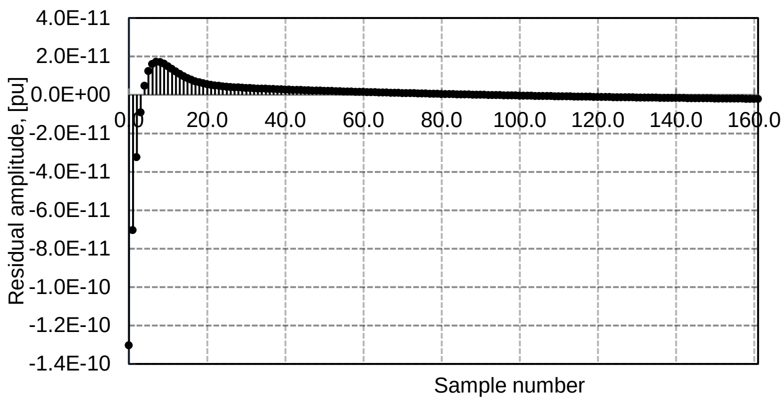

Figure 29 Result of fitting noise double exponent: the residuals

The next steps in the development would be to restrict the fitted data to the region above 30% of peak. However, peak itself does not have a simple definition, and may have to be found by curve-fitting. We are presently studying the matter.

[7] The number accounts exactly for the 100% of the line of which 60% is defined.

5.0 Summary Remarks

The most important conclusion to emerge from this work is this. When we make a measurement, we are finding out something about a model of the thing we are measuring. We are using the physical world to do it, but it is about the *model* that we are being informed. That is a strange thought, indeed.

The story was told in the previous DOE report of this project about a power supply whose output measured 400.00 V on a 5-digit voltmeter. According to the customary rules, the statement of the result of the measurement should have been something like “400 volts, with a 95% probability that the true value lies between 399.98 and 400.02 volts.” That sort of statement would be conventional, and yet would have been inadequate to describe the situation. The power supply was oscillating, with about 100 V pk-pk oscillation superposed on the dc. The model used by the voltmeter was not the right model to characterize the signal.

A more complete statement would have included the observation, easily made in the measuring instrument, that the signal and the model did not match. Such an observation comes from regarding measurement as a mathematical fitting problem. Measurement, however it is implemented, *is* a way of fitting a model to the physical world. That thought takes some getting used to. Measurement is a process that selects a model and adjusts its parameters to fit the thing being measured, the realized quantity.

For the characterization of a phasor, most implementations find rms quantities and use frequency-domain methods (such as Fourier transforms). But we have shown that a direct mathematical fitting method is an excellent way of making the measurement. Applied to power system signals, in windows of size from half a cycle to many cycles, the amplitude, the frequency and the phase (with respect to the IEEE standard reference wave) can be simultaneously measured.

For the PMU, the method relies on the use of an equation that in this report we have labelled the “PMU equation.” The equation is the equation of a phasor

It has been shown that, although the commercial PMU is designed with no regard to that equation, it does in fact furnish the same values. That should not be a surprise. At least in the steady state, a signal has only one value for each of the three parameters, and any measuring device should give the value.

The fitting method that we have employed uses a least-squares estimator, so the results can be said to be the best possible, in the least-squares sense. If some advantage could be demonstrated, the method could easily be adapted to provide “center-weighted” results. At present, based on even weighting of the residuals, all the sample values carry equal weight.

If the ROCOF parameter is included in the equation, it can be evaluated as a “primitive” quantity, not a derivative of anything else. The fitting method allows perfect estimation of the value of ROCOF—provided there is no noise on the signal. We have shown that even a very small amount of noise will swamp the ROCOF contribution to the signal over a narrow window.

The results of the measurement made by our method in this report are independent measurements, unlike the measurements made by commercial PMUs. They are obtained by what is called a rectangular window. The independence means that they can be treated by conventional statistical means—something that cannot be done to results that are not independent.

The use of fitting to an equation in this way is an advance in measurement theory, and it has led to other theoretical gains, some interesting offshoots in other aspects of measurement theory. If we set aside the progress made in instrumentation, we are left with useful results in measurements.

After fitting, next on the list is the fact that once the fit is complete, the residuals can be examined. With a real signal, no fit is perfect, and the residuals contain useful information that is, by its very definition, not part of the result of the measurement. A logarithmically compressed value of the reciprocal of the root-means-square residuals has been shown to be a useful metric for the quality of the match between the signal and the model. We label the metric the Goodness of Fit.

Signals from a 345-kV system whose data we have obtained typically show a Goodness of Fit from our method of about 35 dB, indicating a mismatch between the signal and the phasor model of a little over 1%. (A commercial PMU measuring the same signal has shown a GoF metric a few dB worse.)

When the power system is faulted, and the waveform is distorted, the GoF becomes much worse, dropping by 10 or 20 dB. That implies a much worse match between the model and the signal, and casts into doubt the value of the result of the measurement.

The oscillating 400-V power supply has a goodness of fit metric of 24 dB, certainly low enough to cause concern.

The GoF metric is something that we argue should be attached to the statement of the result of the measurement—indeed, to the result of *any* measurement. It is all very well to know that the 400-V supply is producing 400 V, but that statement of results conceals some important information.

It would be interesting to see how the calculation of the phase sequence values should be modified to account for such information as GoF.

Measurement theory benefitted a while ago from efforts by metrologists to adapt the method of the Allan Variance to measurements other than time. The method demands stability, because measurements must be repeated hundreds or even thousands of times. This is the domain of “pure” metrology, with its controlled conditions. To allow such an evaluation of our fitting method, we created a system whereby we could add noise to the sample-values used to represent a sampled signal. The noise could be added to each of the parameters of the PMU equation, and could be filtered to generate white noise, pink noise or red noise. The application of signal of this kind to a PMU (any PMU) gives the underlying stability needed for an evaluation using the method of Allan Variance. Our early work on that indicates that the optimum window is not the same for the various parameters of the equation.

A new two-sample variance that we have named the Riepnieks Variance has been defined. Whereas the Allan Variance plot is has the window width as independent variable, the Riepnieks Variance has the sampling rate. Our early work with this suggests that the parameter should also be useful in gaining an understanding of the effects of noise on the result of the measurements.

Appendices

6.0 Papers: Introduction and Abstracts

This section of the report lays out most of what the new method has established. It does so by reproducing the papers that have been published (or have been submitted) describing the work.

These papers cover a good deal of the new ground opened up by the fundamental work reported above, but the coverage is not complete because the work is still continuing. Mr Riepnieks will be returning to his home town (Riga, in Latvia) at the start of June of this year, and there he will continue this effort into his PhD studies at the Riga Technical University.

The papers, and information about where they were published or submitted, are briefly summarized next.

6.1 NASPI Goodness of Fit (ABSTRACT)

While some presentations had been made in meetings of the Instrumentation and Measurements Committee of the IEEE Power and Energy Society over the last year or so, a NASPI meeting in March 2016 was the first time the product of this work was formally aired. In response to an Abstract submitted in advance, it was agreed that the work on Goodness of Fit would be presented to the PSRVTT, the Performance, Standards, Requirements Verification Task Team of NASPI. The presentation uses data obtained from the real-world power system.

Making a measurement is the act of using the physical effects of the real world (such as the magnetic field of a current) to find values for the parameters of a mathematical model. Sometimes, that model is simple: a direct current is describable by $i(t)$ = const. for example. Four parameters are measured by a PMU: the amplitude, the phase, the frequency, and the ROCOF. These four parameters can be combined into a mathematical model that resembles the equation of a phasor. That equation is the model.

In the world of digital instrumentation, the measurement results are derived from a sequence of A/D samples of the signal. Once the measurement results have been obtained, the values can be put back into the equation of the model, ideally reconstructing the original signal. Residuals can be calculated: differences between the values at the time of each A/D sample as predicted by the model and as found in the signal. For the direct current, for example, the residuals will be small if the ripple is small. If there is a lot of ripple, the residuals will not be small.

We have developed for the PMU a way to use the residuals to evaluate a metric we call “Goodness of Fit.” We present results calculated for PMU data from the “real world,” information obtained during off-nominal power system conditions (such as a Line-Ground fault). We argue that the GoF is a new metric that could become useful to the PMU user/application. Unlike the measurement uncertainty value, this metric indicates how well the model matches the signal that was measured. During a rapid change in the power system, for example, the phasor equation may be a poor fit to the signal. The metric thus indicates the degree of confidence that can reasonably be placed in the result.

6.2 The Measurand: the Problem of Frequency

Inasmuch as the NASPI presentation was within the PMU community, this will be the first “external” paper to see the light of day. It is to be presented at a meeting of the IEEE Instrumentation and Measurements Society in May 2016.

The PMU equation is given, and the notion of having a mathematical measurand is presented. The idea of instantaneous frequency is floated.

The conceptual entity that metrologists term the measurand is a model selected to represent the physical entity being measured. In a world of digital measurements, it should be defined first mathematically, and only then put into words. Human linguistic processes lack the precision required

when all we do is use labels. In this paper, reactive power and frequency are used as examples.

The act of measurement finds the values of the coefficients of the model. In other words, it solves an equation.

In a digital instrument, information about the quality of the fit between the physical entity being measured and the conceptual model is often available. In essence the instrument can comment on the selection of the model. This comment should be reported as part of the statement of the result of the measurement, along with the declared value and the uncertainty.

6.3 Dealing with non-stationary signals: Definitions, Considerations and Practical Implications

This paper will be the second paper to go beyond the PMU community. It is planned to be presented at the IEEE PES General meeting in Boston in July 2016. The paper introduces mechanical damping factor as an example of a parameter that is described generally only as a term in an equation. That serves to establish the validity of the technique for defining frequency. The notion of finding the parameters of a phase modulated calibration signal is ruled out. Residuals are examined, and the effects of phase jumps shown.

The paper addresses the question of how to deal with non-stationary power signals. The first part of the solution is at a fundamental level: the recognition that the thing being measured is known by some kind of label in a model. The label is attached to a some parameter in an equation, and is often identifiable by its position in the equation. The paper presents measurement as the act of solving the equation to find the value of the parameter. In other words, the equation is what metrologists term the measurand, and the measurement equipment must be designed around it. To measure a time-varying signal, in a world of digital measurements, one of the first questions that must be addressed is the relationship between the sampling window of the measurement system and the rate at which the signal is varying. A goodness of fit metric is identified. Several changing-frequency cases are examined.

6.4 Error Correction: a proposal for a Standard

This is a paper with coauthors from outside PNNL. One is Dr Eddy So, a metrologist at the Canadian National Research Council, the other is Mr James McBride, owner of JMX Services, a company that does high-voltage calibrations. The paper is short, but it manages to introduce the measurement framework and the idea that measurement is solving an equation. The paper will be presented at a conference of metrologists, the Conference on Precision Electromagnetic Measurements, in Ottawa, Canada, in July 2016. This is a meeting for experts from the various national metrology laboratories. While Dr So has had papers at this conference before, this is only the second time for us.

Some of the errors in transducers such as instrument transformers can be corrected as part of the digital processing for the measurement. The instrument transformer can be characterized in such a way that allows the Transducer Electronic Data Sheet of IEEE Std 1451 to transfer the information to the measurement system. A modification would allow the measurement system to perform a high-quality self-calibration whenever a transducer was replaced. That levies requirements on the characterization accuracy of the instrument transformer.

6.5 Introduction to Goodness of Fit for PMU Parameter Estimation

This paper is in the review cycle at IEEE. It was written specifically to be a Transactions paper, with archival value.

An earlier version of the paper was given an "administrative reject" by a computer at IEEE, which decided we did not have enough citations. Hard to believe, they form letter informing us of this said that the reviewers needed citations, apparently so they would know what we were talking about. We have objected to this stupidity.[1] We suspect that the real reason is the academics for whom publish or perish is a reality will benefit from increased citation counts.

In this paper, the Goodness of Fit is expressed logarithmically. Several examples of GoF for real-world signals (distribution as well as transmission) are given. The case is made that GoF is a metric that should be part of the statement of the result of any measurement.

It is posited that the process of measuring the various parameters that characterize a signal is equivalent to a fitting problem in mathematics. The equation being fit can be written based on the "physics" of the signal. The Fourier transform or rms calculations in a phasor measurement unit furnish the values of the coefficients. Regardless of exactly how the measurement is made, a metric we define and call the Goodness of Fit allows the measuring system to comment on the match between the signal it is observing and the model. The metric is based on the residuals, the differences between the signal itself and the value calculated from the result of measurement. Results from real-word phasor measurement units and real world signals illustrate that the equation of the PMU is well solved during steady conditions. We examine the effect of a fault in the transmission system on the Goodness of Fit metric for a PMU. We also apply the metric to results from a microPMU in the distribution system.

[1] It is not that we do not think there should be many citations in a paper. A citation may be needed to support a contentious or difficult point, or to show the point of departure for new work. For original and foundational work such as this, that sort of citation is not to be found. Shannon's great papers on what became Information Theory had only ten citations. We had matched that number.

6.6 Rate of Change of Frequency measurement

This paper has been submitted to an IEEE conference in Riga, Latvia, the home of the second author of this report. The paper begins by looking at an omission in the testing of PMUs, and concludes by making a strong statement about the non-measurability of the parameter called rate-of-change-of-frequency. It also stresses the need to understand the noise on the power system.

The measurement of amplitude, frequency, rate-of-change of frequency, and phase of an alternating waveform is the purpose of the Phasor Measurement Unit, PMU. Performance requirements, specified by standard, are tested with constant values of each of these parameters, using a synthetic waveform with values that are precisely known. However, device performance requirements are not defined during transitions from one set of values to another. We investigated measuring across a transition. Our investigation revealed something interesting about ROCOF, the rate of change of frequency. We conclude that until power system noise is better understood, the attempt to measure real-world ROCOF during a short PMU measurement window should be abandoned, but measurements during calibration transitions might still be possible and need not be excluded from the standard.

6.7 Students' Simple Method for Determining the Parameters of an AC Signal

This paper, like the previous one, is aimed at a conference in Riga.

The title is a tribute to the great William Edward Ayrton. His 1894 paper "Students' Simple Apparatus for Determining the Mechanical Equivalent of Heat" set aside details of the then common instrumentality to show how, using direct-reading electrical methods, a quantity identified as the mechanical equivalent of heat could be measured in about ten minutes. That time contrasted with the years that scientists had spent (along with considerable effort, time—and money) making this measurement, with no greater accuracy.

Our goal, like Ayrtyon's, is to give the student of measurement the clear and concrete idea. These days that comes from implementations in MATLAB, a convenience that avoids the need of wires and transformers, just as the implementation of Ayrton and Haycraft avoided the need for what a reviewer described as "a lesson in calibration, or in the principle of the tangent galvanometer."

Although, like Ayrton, we have developed a new method of measurement, and although we think it to be capable of giving better results than anything presently available, it has yet to be shown that it can work in real time, and many measurement people are not yet "comfortable" with fitting as a measurement method. We do not use this paper to advocate for the method.

Instead, the ideas that we communicate are that

1. the performance of a measurement system is limited by noise in the system and on the signal being measured

2. a flexible measurement system (such as this) can be used to explore the limits of performance

3. the meaning of the measurement result is not necessarily obvious.

The paper sets aside details of instrumentality to reveal the nature of the problem addressed by measurement. Its title is based on the title of a 1894 paper by Prof. W.E. Ayrton and his student H.C. Haycraft. They described a new and simplified method of measurement to improve the teaching of their underlying topic, and that is the goal of this paper. In the work described here, the measurand is taken to be an equation representing an alternating signal, and the declared values of the measurement are estimates of the parameters of the equation. It is shown that the parameters of the ac signal can be found by curve-fitting. Lessons can be drawn about the role of noise in measurement and about the very meaning of the result.

7.0 Papers

7.1 NASPI Goodness of Fit (ABSTRACT and presentation)

H. Kirkham, Harold.Kirkham@pnnl.gov (509) 371-6721

A.Riepnieks, Pacific Northwest National Laboratory and Riga Technical University, Riga, Latvia, artis.riepnieks@pnnl.gov (509) 375-4508

R. J. Murphy, Macrodyne, jay.murphy@macrodyneusa.com (518) 281 8101

R. M. Hayes, ESTA International, ray.hayes@estainternational.com, (740) 438 9624

Goodness of Fit

Making a measurement is the act of using the physical effects of the real world (such as the magnetic field of a current) to find values for the parameters of a mathematical model. Sometimes, that model is simple: a direct current is describable by $i(t)$ = const. for example. Four parameters are measured by a PMU: the amplitude, the phase, the frequency, and the ROCOF. These four parameters can be combined into a mathematical model that resembles the equation of a phasor. That equation is the model.

In the world of digital instrumentation, the measurement results are derived from a sequence of A/D samples of the signal. Once the measurement results have been obtained, the values can be put back into the equation of the model, ideally reconstructing the original signal. Residuals can be calculated: differences between the values at the time of each A/D sample as predicted by the model and as found in the signal. For the direct current, for example, the residuals will be small if the ripple is small. If there is a lot of ripple, the residuals will not be small.

We have developed for the PMU a way to use the residuals to evaluate a metric we call “Goodness of Fit.” We present results calculated for PMU data from the “real world,” information obtained during off-nominal power system conditions (such as a Line-Ground fault). We argue that the GoF is a new metric that could become useful to the PMU user/application. Unlike the measurement uncertainty value, this metric indicates how well the model matches the signal that was measured. During a rapid change in the power system, for example, the phasor equation may be a poor fit to the signal. The metric thus indicates the degree of confidence that can reasonably be placed in the result.

Goodness of Fit

NASPI
March 2016

Harold Kirkham, Artis Riepnieks, PNNL
Jay Murphy, Macrodyne
Ray Hayes, ESTA International

Pacific Northwest NATIONAL LABORATORY

Measurement (1)

Given an alternating voltage, we could

- Measure amplitude
- Measure frequency
- Measure phase relative to some other signal

Or, if we wish, we could measure all together

That is the same as solving the equation

$$x(t) = A\cos(\omega t + \varphi) \qquad (1)$$

Pacific Northwest NATIONAL LABORATORY

Measurement (2)

Measurement of this kind is thus what mathematicians call a "fitting problem."

Only in this case, the form of the equation is fixed by the physics.

As a fitting problem,

- Need multiple samples
- Min # samples = (# degrees of freedom in equation) +1
- (1) can be solved with 3 samples if no noise

Modify the equation to include ROCOF, need 4 samples

Pacific Northwest
NATIONAL LABORATORY

Measurement (3)

- The equation is a model: Equation (1) was a cosine model
- The *results* of the measurements are the model coefficients
- To solve as a fitting problem, one minimizes the residuals

But the measurement does not have to be made as a fitting problem to take advantage of the method*ology*

- One can find residuals, however the measurement was done
- One can look at how big the residuals are

- That is a very informative thing to do!

Pacific Northwest
NATIONAL LABORATORY

AEP* 345-kV data fits the equation

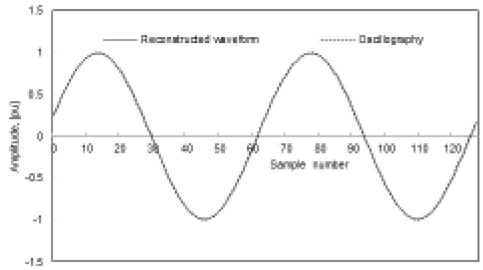

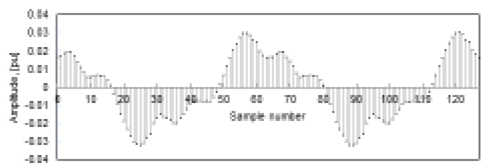

*Many thanks to Zak Campbell (AEP) for the data sets

Pacific Northwest
NATIONAL LABORATORY

AEP data

- The residuals are few percent of the signal
- There is a fundamental-frequency component, indicating phase or timing error somewhere

But residuals are small, indicating

the model is a good representation of the signal

- Not always the case: during fault, for example, or phase jump
- Define Goodness of Fit:

$$\mathrm{GoF} = 20\log\frac{A}{\sqrt{\frac{1}{(N-m)}\sum_{k=1}^{N}(u_k - v_k)^2}}$$

Pacific Northwest
NATIONAL LABORATORY

Fault response (voltage)

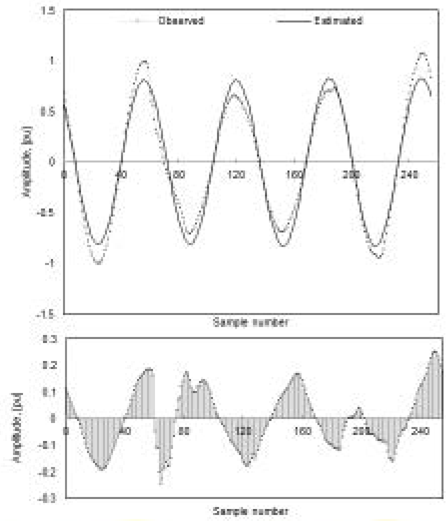

Fault response: Current & GoF

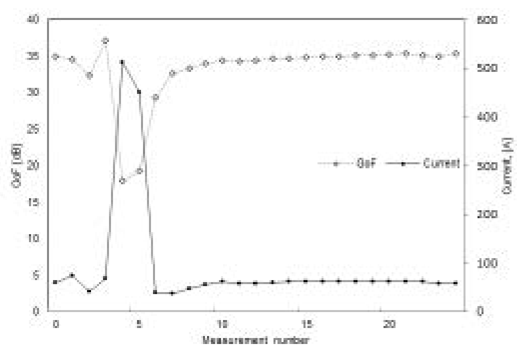

Pacific Northwest
NATIONAL LABORATORY

Final Remarks

"Since the measuring device has been constructed by the observer, we have to remember that what we observe is not nature in itself, but nature exposed to our method of questioning"

The PMU answers this question:
If this signal were a cosine wave, what would the amplitude, frequency and phase be?

But the signal may not be a cosine wave . . . And GoF will tell you that without delay

W. Heisenberg, Physics and Philosophy: The Revolution in Modern Science, London: George Allen and Unwin, 1959.

Pacific Northwest NATIONAL LABORATORY

Glaucon: And the arts of measuring and numbering and weighing come to the rescue of the human understanding—there is the beauty of them—and the apparent greater or less, or more or heavier, no longer have the mastery over us, but give way before calculation and measure and weight?

Socrates: Most true.

Plato *Republic* Book X, 360 BCE

Pacific Northwest NATIONAL LABORATORY

7.2 The Measurand: the Problem of Frequency

7.2.1 INTRODUCTION

The theme of this paper is the evolving nature of the thing that metrologists call the measurand. Perhaps in this digital world, the evolution has finally stopped. But the current significance of the measurand is not yet well-enough known.

7.2.1.1 Progress in Measurement Technology

Measurand is the word used to describe the thing being measured. An early comment on the need for a good definition was made by George Carey Foster in his address [1] as the incoming President of the Society of Telegraph Engineers and Electricians. He noted that

> Before methods of measurement can be devised, it is evident that clear conceptions must be formed of the things to be measured. Such conceptions usually grow up by degrees in many minds from indistinct beginnings, until, in some one mind, they take definite shape and receive the precise expression which makes it possible for them to become the subject of mathematical reasoning.

This *conception* of the thing to be measured we will here call the measurand. The actual thing being measured is better called the realized value. GUM [2] notes (page 49) that "Ideally, the quantity realized for measurement would be fully consistent with the definition of the measurand. Often, however, such a quantity cannot be realized and the measurement is performed on a quantity that is an approximation of the measurand." That clarification allows us to use the word measurand to mean just the definition.

These days, it is more likely that the precise expression would be created by a standards working group in the IEEE rather than in some one mind. But note that Carey Foster saw the need for mathematical reasoning. (He also recognized that measurement was "application-driven" and spoke of that in the same address.)

The time was the 1880s, and instruments we would call direct-reading were making their first appearance. William Ayrton and John Perry, in particular, had major impact on the world of measurements with their instruments. It is no exaggeration to say that Ayrton separated the world of electrical engineering from the world of physics when he presented a paper [3] that showed an electrical way to determine the "Mechanical Equivalent of Heat" in November 1894.

The paper showed how electrical measurements could give a value for what today we would call the latent heat of water in a matter of tens of minutes. The world of physics, having spent years making such a measurement (with no greater precision), was not well pleased. As electrical engineers and metrologists, however, we can thank Ayrton that we are not still using the tangent galvanometer.

For a maker of direct-reading instruments, the realized quantity was readily selected (current or voltage) and the measurand was somewhat ill-defined. However, the result of the measurement was readily available. Though Kelvin had been the mentor of Ayrton (and an admirer), he did not favor direct-reading instruments, even though he had said, a few years earlier [4] that

> . . . when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a

> meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts advanced to the stage of science. . .

Kelvin saw that measurement was the key to progress. His measurements work was truly insightful. One thinks particularly of the double bridge, designed to overcome problems of connection resistance and at the same time rely only on a steady zero-point on the indicating instrument.

Difficulties with measurements can be traced back even earlier. In the 1870s, Cambridge University installed James Clerk Maxwell as Cavendish Professor of Experimental Physics. In a report [5] to the University in 1877, he wrote

> "It has been felt that experimental investigations were carried on at a disadvantage in Cambridge because the apparatus had to be constructed in London. The experimenter had only occasional opportunities of seeing the instrument maker, and was perhaps not fully acquainted with the resources of the workshop, so that his instructions were imperfectly understood by the workman . . ."

This two-way communication gap is similar to the still unsolved problem we have today. The difference is that these days, our "instrument makers" are software writers. We are at a time after what has been called the "digital revolution in measurement" has taken place. The equivalent of frequency response and dynamic range are now all determined as part of a digital measuring system, unconstrained by details of bearing friction and needle-width. We need to tell our instrument makers quite precisely what we want to measure.

7.2.1.2 Significance of Frequency Measurement

Measurement is data compression, but it is not just arbitrary compression. Measurement is the connection between the physical world and the conceptual. Measurement results have very particular meanings. It is these meanings that have to be defined in the measurand.

The present author was a participant in a standards working group for the performance requirements of an instrument called the phasor measurement unit (PMU). Troubled by the way frequency was being handled, he was invited to consider the possibility of a definition for "frequency" that would be suitable to the problem at hand. After all, only when something is well-defined can a calibration be done to assess how well it is being measured.

Frequency is one of those words that everybody understands. Yet when asked for a definition, most of us start to wave our hands. Frequency is not exactly the number of times something repeats in a second, because it might not repeat in a second, and because there might not be a second in which to make the measurement.

Strictly, the phasor representation forces the frequency to be constant over all time, and for the PMU we needed a new definition for when it (whatever "it" is) is changing.

7.2.2 DIGITAL MEASUREMENTS

7.2.2.1 High-Level View of Digital Measurement

Digital measurements force closer consideration of the process of measurement. They may have started as digital version of analog measurements, but they are in many ways more capable, and these new capabilities deserve careful evaluation.

Because of digital measurements, it is evident as it never was before that the process of measurement begins with the selection of a conceptual model to represent the realized quantity. That is what the measurand really is: a model of the thing being measured.

The measurand is the bridge from the physical to the conceptual. On one side, the conceptual model furnishes data to an application. On the other side it matches the real quantity being measured. A couple of electrical measurement examples will illustrate.

Suppose the problem to solve is the measurement of a value of direct current. The thing to be measured could be described by a simple equation $i(t) = I_c$. That says that the current is constant over all time. The quantity is converted to a voltage (perhaps by means of a resistor, whose value thus contributes to the measurement uncertainty) and sampled in an A/D converter. A sequence of numbers is then furnished to some process that selects a few numbers from the sequence, does some kind of averaging (in case the numbers are not quite all the same) and gives as a result the thing called the declared value.

The measurement system is solving the equation. That seems to be a new thought. Carey Foster talked about mathematical reasoning, but never quite made the connection to solving an equation. Kelvin came close when he talked about expressing what you were speaking about in numbers. It is hard to see that a tangent galvanometer, even when used by an expert, is solving an equation. That, nevertheless, is what is going on.

The bridge between the physical and the conceptual is the model, the equation, the measurand. The process of measurement is the process of finding the parameters of the equation [6].

Alternating quantities can be handled with a similar process. Perhaps the rms value is needed: the sequence of (digitized) values is used to solve an equation like this:

$$v_{\mathrm{rms}} = \sqrt{\frac{1}{T_2 - T_1} \int_{T_1}^{T_2} [v(t)]^2 dt} \qquad (1)$$

The difference between the times that define the integral is the period of the incoming wave, which should therefore be known for the calculation to proceed accurately. However, the result is not terribly sensitive to errors in the time, so the calculation is feasible.

7.2.2.2 More on AC Measurements

Since power delivery as alternating voltage and current began, it has been known that some loads have the characteristic that they draw current from the supply that is not aligned in time with the voltage. The current does not correspond to power delivered.

Since delivered power is usually the quantity being metered and billed, it is to the advantage of the power supplier to minimize the non-power current. The term *power factor* was created [7] to represent the ratio of the true to the apparent power. At the time, 1891, it was observed that

> If the currents and pressures were simple sine functions, then the power-factor in that case would be the cosine of the angle of lag of primary current behind the primary terminal potential difference.

And electrical engineering students ever since have known that power is Volts times Amps times Cosine-ϕ.

Something called the apparent power can be written $S = V{\times}I$, and the power factor is P/S. A handy Pythagorean triangle relates the apparent power S to the real power P by means of an imagined quantity called reactive power:

$$S^2 = P^2 + Q^2 \tag{2}$$

The problem is, the voltages and currents in the power system are not always perfect sine-waves. When there are harmonics in the system, they change the numbers.

Constantin Budeanu proposed to deal with them harmonic by harmonic [8], and for years this method was favored. However, there were other ways to handle the problem. Some modern instruments even allow the user to choose.

The argument has sometimes become heated. Proponents of one method describe the other as, for example, "deeply erroneous" [9]. But there is no experiment that can say one way of handling harmonics is right and another wrong. Calorimetry can confirm the delivery of energy, but no singular experiment can be designed for reactive power (however defined) since it is just a fiction. Nelson [10] identified at least ten different methods in use.

"Wrong" in this context seems to mean simply that the Pythagorean relationship above is not satisfied. But there is no reason that Pythagoras is applicable. The quantities involved are scalar, and not associated with any particular "direction." It may be hard to accept, but any of the ten methods is "right" as defined.

The use of mathematics as the measurand allows the digital instrument maker to understand the problem. The software can make exactly the required measurement. But describing two different things with the same name is a problem with our linguistic process. The name is no more than a label, and labels fail to take advantage of the wonderful subtleties of language.

7.2.2.3 Frequency

That same linguistic deficiency is at the root of the problem of measuring changing "frequency." In the world of electrical power, it has been taught for generations that the speed of the generators is so constant that you can run a clock from the power system, and keep good time. Given that, power engineers think of the quantities in the power system as being described by phasor equations such as:

$$x(t) = X_{\mathrm{m}} \cos(\omega t + \varphi) \tag{3}$$

Here $x(t)$ is the instantaneous value of the function, X_{m} is the maximum value, ω is the angular frequency and φ is the phase.

The equation is a phasor, meaning the parameters are stationary. It is easy to measure the frequency of such a thing: just count the cycles. No matter that it takes a long time, the answer is believable because the technology is so straightforward.

Results showing that the frequency was far from constant in the short term, and was also a function of location, came as a surprise to many power engineers. Figure 1 shows the frequency at four locations after a large generator had been abruptly removed from the power system.

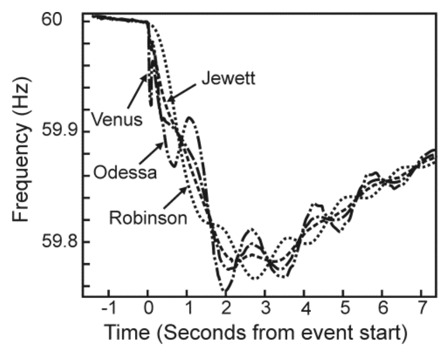

Figure 1 Observed frequency following generator loss

It is clear that the thing called frequency is not constant. Though the speed variation is not large, it is evident that Equation (3) is at best an approximation. The difference came about because of a new way of measuring the thing called frequency. It was based on different technology [11], one that allowed the frequency to be measured very rapidly.

The question really should be this: if the value is changing, can it be called frequency?

Suppose we re-write the phasor equation (3) and allow everything except the amplitude to vary with a constant rate of change.[2] We get

$$x(t) = X_{\mathrm{m}} \cos\left\{\left(\omega' + \frac{C_{\omega}}{2} t\right) t + \left(\varphi' + \frac{C_{\varphi}}{2} t\right)\right\} \tag{4}$$

where the C parameters are the rates of change. We have replaced the symbols ω for frequency and φ for phase by ω' and φ' since ω and φ are customarily stationary, and we wish to overcome that prejudice. We are preparing for a short sequence of data to be evaluated to give these terms, and they will not generally be the same from one sequence to the next.

The term in $C_{\varphi} t$ cannot be distinguished from, and is therefore moved to, like term ω'.

[2] This is not a requirement that the power system behave this way, only that we are going to model it this way in the short term.

$$x(t) = X_{\mathrm{m}}\cos\left\{\left(\omega' + \frac{C_\varphi}{2} + \frac{C_\omega}{2}t\right)t + \varphi'\right\} \tag{5}$$

The equation now has two terms inside the cosine argument that have a linear dependence on *t.* They are together occupying the place that frequency occupies in the phasor equation. If it were not for the fact that the word was already defined, we could call their sum the frequency. Perhaps we should call them the apparent frequency. The point is that when you measure something that has the characteristics of frequency and is fixed across the measurement window, you are measuring these two terms, and assigning them one name, apparent frequency, or perhaps just frequency.

The important thing is that the mathematics defines the thing measured: the name is just a label.

If you measure something that you allow to change across the window, you are including the term in t^2, and you could call the time-varying collection the instantaneous frequency. The point is not whether the equation *always* defines apparent frequency, or instantaneous frequency, but we could assert that it defines the measurand of an instrument well enough for it to be unambiguously evaluated, ie, measured. One can ask no more. Maxwell would be pleased.

That was the situation with reactive power. The label is not definitive. There is more than one mathematical definition claiming to be reactive power. Provided the definition is clear and unambiguous mathematically, a measurement can be made. Any such definition must be admissible, but the mathematics should be published.

Only then should a name in a written or spoken language be considered. The linguistic labeling process simply lacks the precision of mathematics.

7.2.2.4 Phase and the PMU

Equation (3) contains a term in what is commonly referred to as phase. In fact, the word phase did not always have the meaning shown in (3). In 1945, the eminent B. van der Pol presented a paper [12] that discussed the significance of two meanings of phase. The alternative to the present meaning is that the complete argument of the cosine is the phase. That interpretation of the word was favored by van der Pol, but it was an argument that he lost. It has the advantage that with that meaning it is possible to define the relative phase of two signals that are not at the same frequency. Otherwise, such a thing is not defined.

While it might be thought straightforward to estimate the relative phase of two signals whose frequencies are close, what are we to expect of a measurement for the situation shown in Figure 2?

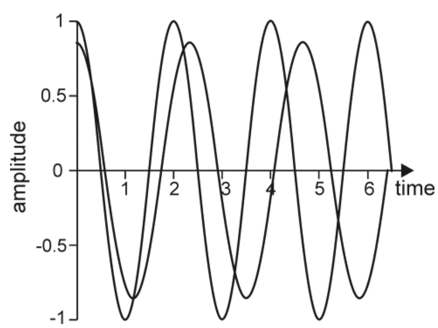

Figure 2. Two sine waves

The signals shown in Figure 2 are not at the same frequency, and so their relative phase (in the usual sense of the word) is a function of the time. However, if the total phase of the signals is known, the difference can be found. In essence, that is what goes on in the PMU.

The total phase is something that originates at time zero and accumulates from there on. In the PMU, a reference wave is created that is a true phasor. For this wave, every cycle is the same as every other, and so one can set $t = 0$ at the top of any cycle of the wave without losing generality. If the larger of the two waves in Figure 2 is the reference, it is possible to set $t = 0$ at time 2, 4 or 6, for example. Up to the rate at which the reporting rate equals the power frequency, the PMU is required to report the phase value at the times where $t = 0$, so the problem becomes one of solving Equation (5) for the value of φ'.

It is important to note that the result is temporally local. An application for the result of the measurement may be interested in the total phase accumulated since some much earlier time. Midnight as a time origin has the feature that the total phase would be a direct indication of the clock error, for example. Such a quantity can be calculated, but is not in the measurand defined for the PMU.

7.2.3 MEASUREMENT QUALITY

Designers of digital instruments may have known they were solving equations, but somehow the notion has not registered much in the world of metrologists. The notion offers new capabilities for digital measurements. A particularly interesting example is the calculation of the quality of the “match” between the measurand (model) and the input data stream.

Once the equation is solved, however it is solved, for many measurements the solution values can be inserted into the measurand, and residuals calculated between the sampled sequence and the model predictions across a given measurement window.

For a cosine wave such as Equation (5), it is possible to normalize the residuals according to the measured amplitude. In work with phasor measurement system, we have found that a perfect synthetic input signal can be measured with residuals that are many orders of magnitude smaller than the signal—limited only by the computer and the accuracy of the time signals. That means that across the measurement window, the declared values match the signal extremely closely.

With real-world signals that have visible distortion (estimated to be about 3%, but not actually measured), we have found the residuals are larger, about a hundred times smaller than the signal. That still indicates a reasonable match between the model and the data.

With a phase jump inserted into the signal, the residuals for the window containing the jump are much larger. The large value for the residuals indicates that there is a mismatch between the model and the signal. The residuals clearly show that the measurement results are not to be trusted for that one window.

The technique would have value in many uses. A power supply design producing a constant output voltage might be tested by measuring with a multi-digit voltmeter, but that voltmeter would mask any ripple in the output. A voltmeter inadvertently set to measure dc instead of ac, and then plugged into a power outlet, would indicate a very small voltage. However, it could also tell the user that the model (the equation $i(t) = I_c$) was a really bad one!

7.2.4 THE FUTURE

Recognition that the process of measurement is the process of solving an equation, and that the equation is the measurand of the system, changes the way one thinks about making measurements.

One hears about the need for more rapid measurements of some quantities, for example, frequency. At the moment, it is possible to get a value of the apparent frequency of a power system with a resolution of about a mHz, and a measurement window just a few cycles long. (That is a feat that would require about 17 s if the measurement were made by counting zero crossings.)

What is "instantaneous frequency"? That is a topic on which many papers have been written (see reviews [13, 14]). How quickly can a frequency measurement be made? That is a question that gets a different answer today than when only zero-crossings were used. Efforts to speed up the process while maintaining accuracy were reported as long ago as 1946 [15], but even the more recent PMU efforts have not explored the theoretical limits.

If the problem is to solve an equation such as (3), it should be seen that there are three parameters to estimate: the amplitude, the frequency and the phase. It follows that at least three samples must be taken of the input wave. What is needed beyond that will depend on the noise and distortion on the signal.

In my laboratory, we have treated the measurement problem as a fitting problem. It is not the usual fitting problem, because the form of the equation is given by the physics. With this approach, we have measured the three parameters plus the rate of change of amplitude and the rate of change of frequency with five samples spread over half a cycle. With a clean signal, the result of the measurement gives residuals ten or more orders of magnitude below the signal. We are now proposing to explore the impact of noise of various kinds, a matter that will depend on the estimator used.

7.2.5 CONCLUSION

We refine the meaning of the thing known as a measurand. Loosely, it has been said to be the thing being measured; it is more properly a definition thereof. We argue that the measurand is indeed a description; it is a model of the thing being measured.

We further propose that this model should first be defined mathematically, before a linguistic version is stated. Human linguistic processes that simply give the measurand a brief label lack the ability to say unambiguously what mathematics can easily elucidate. The measurand is an equation to be solved by the instrument.

Once a solution has been found, many measurands have the property that the sequence of input samples can be compared to the calculated result of the measurement, and residuals found. Those residuals are a comment we have called the Goodness of Fit, a judgment of the match between the model and the input signal. This metric should become a useful part of the expression of the result of a measurement, along with the declared value and a statement of the uncertainty.

7.2.6 REFERENCES

[1] G. Carey Foster, Inaugural Address, 1881, Journal of the Society of Telegraph Engineers and of Electricians, Vol 10, No 35, p 4 *et seq*

[2] Metrology, W. G. 1. o. J. C. f. G. i., 2008. GUM: Evaluation of measurement data--Guide to the expression of uncertainty in measurement, Paris: BIPM.

[3] Students' simple apparatus for determining the mechanical equivalent of heat, W.E. Ayrton and H.C. Haycraft , Proceedings of the Physical Society of London, Nov 23, 1894, pp 295 - 309

[4] W. L. Kelvin, "Electrical Units of Measurement (May 3, 1883)," in *Popular Lectures and Addresses*, London, Macmillan, 1889, pp. 79-136.

[5] J.C. Maxwell, Report on the Cavendish Laboratory for 1876 (from the Cambridge University Reporter, ca May 1877) On page 10 of A history of the *Cavendish laboratory*, 1871-1910: https://www.google.com/search?q=Report+on+the+Cavendish+Laboratory+for+1876+&ie=utf-8&oe=utf-8&aq=t&rls=org.mozilla:en-US:official&client=firefox-a

[6] H. Kirkham, "A Conceptual Framework for Measurement, Report PNNL-24071," Pacific Northwest National Laboratory, Richland, WA, 2015. Available from https://certs.lbl.gov/publications/conceptual-framework-measurement

[7] On some effects of alternating-current flow in circuits having capacity and self-induction, J.A. Fleming, Journal IEE, Vol 20 No 94, May 7, 1891, pp 362 - 408)

[8] C. Budeanu, Puisances reacives et fictives," Instytut Romain de l'Energie, Bucharest, 1927.

[9] L. Czarnecki, "Budeanu and Fryze: Two frameworks for interpreting power properties of circuits with nonsinusoidal voltages and currents," *Electrical Engineering,* pp. 359-367, 1997.

[10] Nelson, T, "Definitions for calculations of VA, VAh, VAR, and VARh for Poly-Phase Electricity Meters, NEMA C12.24 TR-2011," NEMA, 2011.

[11] A.G. Phadke, J.S. Thorp and M.G. Adamiak (1983) A New Measurement Technique for Tracking Voltage Phasors, Local System Frequency, and Rate of Change of Frequency, IEEE Trans PAS 102, Issue 5, pp 1025—1038

[12] B. van der Pol, "The Fundamental Principles of Frequency Modulation," *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering,* pp. 153-158, 1946.

[13] L. Cohen, Time-Frequency Distributions—a Review, *Proceedings of IEEE,* vol. 77, no.7, pp. 941-981, July 1989.

[14] B. Boashash, "Estimating and Interpreting the Instantaneous Frequency of a Signal--Part 1: Fundamentals," *Proceedings of IEEE* , vol. 80, no. 4, pp. 520-538, April 1992.

[15] H.B. Law, "An instrument for short-period frequency comparisons of great accuracy" , Journal of the Institute of Electrical Engineering – Part III: Radio and Communication Engineering, 1947, Vol 94 Issue 27, pp 38—41

7.3 Dealing with non-stationary signals: Definitions, Considerations and Practical Implications

7.3.1 Introduction

Consider the equation below:

$$a\frac{d^2\theta}{dt^2} + b\frac{d\theta}{dt} + c\theta = Gi \tag{1}$$

This is the equation of motion of a galvanometer, an instrument of much interest in measurements in historical times. In the equation, θ is the angle of deflection in radians, a is the constant of inertia, *b* is the damping constant, due to air friction and the elastic hysteresis of the suspension. *G* is the displacement constant, the product of the field strength in the air gap, and the area of the coil. That, at least, is the explanation given by Golding [1]. Equations of this form are familiar to us all.

Harris [2] gives a slightly different version. He writes

$$P\frac{d^2\theta}{dt^2} + K\frac{d\theta}{dt} + U\theta = \frac{GE}{R} - \frac{G^2}{R}\frac{d\theta}{dt} \tag{2}$$

where *P* is the moment of inertia, *K* is the mechanical damping coefficient, *U* is the suspension stiffness, *E* is the applied voltage and *R* the resistance. (Harris gives *G* the name “motor constant.”) Equation (2) can be rearranged to give

$$P\frac{d^2\theta}{dt^2} + \left(K + \frac{G^2}{R}\right)\frac{d\theta}{dt} + U\theta = Gi \tag{3}$$

so that the comparison with (1) is very clear. In particular, note that if *b* is the damping constant, then so is the expression $\left(K + \frac{G^2}{R}\right)$.

The point is this. The label we attach to the parameter is one that is determined by its location in the equation. In the equation of motion above, the coefficient of the first order term is the damping. We propose that this is the solution to the problem of the measurand for non-stationary signals.

7.3.2 Defining the Measurand

“Measurand” is a term that should be carefully considered. It is a description of the thing being measured. While it is a word sometimes used to signify the physical thing being measured, it is used here in the strictly conceptual sense of a description.

In the world of digital measurements, the measurement is made as a calculation of some sort, solving the equation that is the embodiment of the measurand for a digital instrument.

There are several examples of the measurand as an equation in Table I.

The equation for dc may seem trivial, but as we shall see (in Section V), it is actually something useful.

Note that there are two measurand equations shown for power factor. These two are actually but two of many. Nelson [3] identifies ten different ways the calculation is done in digital instruments. A problem with that is that the result of the calculation is not always the same if the waveform being measured contains harmonics, an example of something called “semantic coloration” in [4]. The difference between semantic coloration and noise is that semantic coloration changes the value (meaning) of the measurement result, whereas the average effect of random noise does not. Only in the absence of this coloration (ie, pure sine-waves) do all the definitions of power factor give the same result.

The fact that ten different equations giving ten different results are given but one label (power factor) is an indication of problems with our linguistic labeling process. The identification of this problem as a linguistic one can lead to a solution for dealing with non-stationary signals.

Table I Mathematical forms of Measurand

Name	Measurand	Note
Direct current	$i(t) = I_c$	Declared value may be computed as an average over some window
Phasor	$x(t) = X_m \cos\{\omega t + \varphi\}$	Find parameters X_m, ω, φ. These values apply for all time.
rms	$v_{rms} = \sqrt{\frac{1}{T_2 - T_1}\int_{T_1}^{T_2} [v(t)]^2 dt}$	Solve for v_{rms}. The period (T_2 – T_1) must be known
Power Factor	$\mathrm{PF} = \frac{\mathrm{P}}{\mathrm{V} \times \mathrm{I}}$	Solve for PF. Solutions of these equations are the same only for sinusoids
	$\mathrm{PF} = \frac{\mathrm{P}}{\sqrt{\mathrm{P}^2 + (\sum \mathrm{V_h I_h} \sin\theta_h)^2}}$	
PMU Ramp Test	$x(t) = X_m \cos\left\{\left(\omega_{app} + \frac{C_\omega}{2} t\right) t + \varphi_{app}\right\}$	Find parameters X_m, ω_{app}, C_ω, φ_{app}
PMU Modulation Test		

7.3.3 The Phasor Measurement Unit

Consider now the familiar equation of the phasor:

$$x(t) = X_m \cos(\omega t + \varphi) \tag{4}$$

where $x(t)$ is the value of the phasor at time t, X_m is the maximum value, ω is the angular frequency and φ is the phase at time t = 0. (The parameter φ is not explicitly a function of the time. However, in the phasor measurement unit the value is given with respect to a reference that may not be at the same frequency as the signal being measured.)

Equation (4) might represent a voltage or a current, and can be solved in several ways, by a PMU given an input stream of samples. Frequency can be found using Fourier transform methods, for example. The relevant IEEE standard [5] mentions finite difference methods based on adjacent measurements of phase. In our laboratory, we have implemented a method based on a curve fit. We use a modified least-squares estimator, and obtain excellent results.

Can a curve-fit be a measurement? Why not? Progress in digital measurements has meant that measurement systems are no longer simply imitations of their analog predecessors. The fact that our system uses an estimator in its solution is evidence that there is no functional difference between an estimate and a measurement.

In fact, the PMU is not asked to solve Equation (4). That is a phasor equation, and a phasor is a stationary equation; it has constant parameters for all time. That is very useful for teaching purposes, but not a good representation of the real power system. Of particular interest (and challenge) is the fact that the frequency is rarely constant.

If the frequency is not constant, the value of ω will not be constant, and φ will not be constant because the phase is measured relative to a constant-frequency reference. We can represent these changes by letting both of these parameters have a constant rate of change. That is the next simplest assumption after assuming a constant value. If we do that, we can rewrite (4) as

$$x(t) = X_{\mathrm{m}} \cos\left\{\left(\omega' + \frac{C_\omega}{2}t\right)t + \left(\varphi' + \frac{C_\varphi}{2}t\right)\right\} \tag{5}$$

where C_ω is the rate of change of frequency (ROCOF), and C_φ is the rate of change of phase. The parameters ω and φ are now marked with a $'$ to indicate that they no longer obey the stationarity requirement of the phasor. They are fixed only for the duration of the measurement.

We can generate some new labels for these terms. The term in $C_\varphi t$ cannot be distinguished and estimated separately from ω', therefore is moved in the equation. This has the effect of leaving φ' as a constant:

$$x(t) = X_{\mathrm{m}} \cos\left\{\left(\omega' + \frac{C_\varphi}{2} + \frac{C_\omega}{2}t\right)t + \varphi'\right\} \tag{6}$$

Combining terms, we obtain

$$x(t) = X_{\mathrm{m}} cos\left\{\left(\omega_{\mathrm{ALF}} + \frac{C_\omega}{2}t\right)t + \varphi'\right\} \tag{7}$$

where ω_{ALF} combines C_φ and ω', and can be labeled Apparent Local Frequency (ALF). We could note that since the term occupies the same place in (7) that frequency occupies in (4), it could also be called frequency. But that would be a failure to distinguish the term from the term used in the forever-constant phasor. ALF is a combination of two things, just as the damping term in (3) combined mechanical friction and electromagnetic damping.

The time being limited by the measurement window makes the term "local" imply temporal locality. We added the term "apparent" because the calculation of the cosine argument also includes a term in t^2. If the PMU is attached to a generator whose speed is changing, the $C_\omega t^2$ term will be non-zero, and the total "frequency" term will not be ALF. The total value might be called the instantaneous frequency. We will look at changing frequency later. Before we do that, we will demonstrate a way to see the effect of non-stationarity on the measurement.

7.3.4 Epistemology and Semantics

We consider a signal that is far from stationary. Suppose there is a step change in the phase. An example is shown in Fig 1. (Shown is a two-cycle stretch of the input.)

The signal amplitude has not changed, the frequency has not changed, only the phase has changed, and is discontinuous. There are two related questions: what do you (as user) *want to know*, and what will your instrument *tell you*.

The answer to the second question is that the instrument will fit the model that the PMU has as measurand to the incoming data stream. There will be values for the amplitude, frequency, phase and ROCOF. The values could then be compared to the input.

In evaluating the comparison, we found the situation improved if we changed the solution method to allow for the amplitude to change during the measurement window. (That is not required of a PMU.) The result is then as shown in Fig 2.

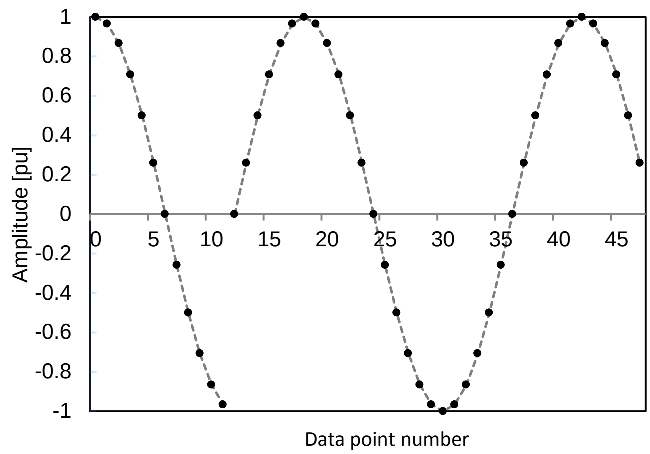

Fig.1. Input signal for single measurement window with 90^0 phase jump

The solution is saying, in effect, if you want to represent this waveform by that model, this is the best that can be done. A PMU working by another solution method would likely give a different solution, but the solutions would have one thing in common: they would not be very good.

The differences between the input signal and the model result are called the residuals, and they tell the message of Fig 2 in a different way, as in Fig. 3.

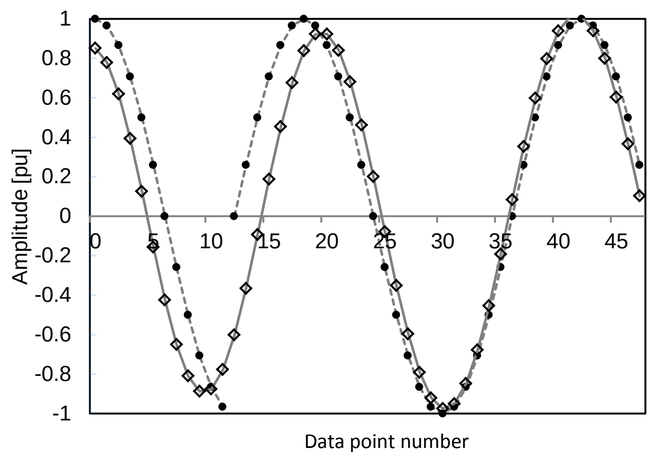

Fig.2. Phasor-like estimate of signal with 90^0 phase jump

The residuals can be normalized by dividing by the rms value of the signal in the same measurement window and finding the means square. The reciprocal of this number can be used as an indication of the goodness of fit between the data stream and the measurand. When the signal of Fig 1 was inserted into an otherwise "clean" data stream, the goodness of fit number showed a sudden decrease, as in Fig. 4

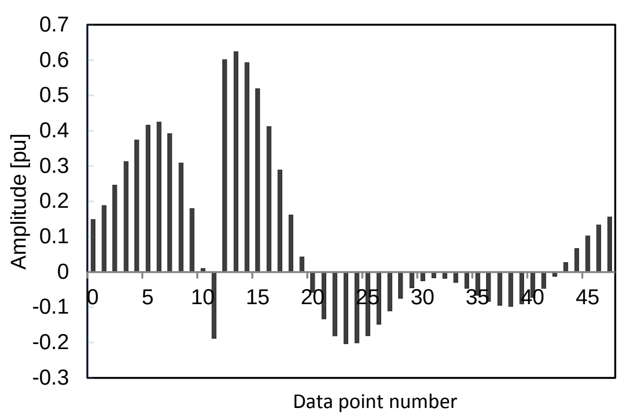

Fig.3. Residuals of situation with 90^0 phase jump

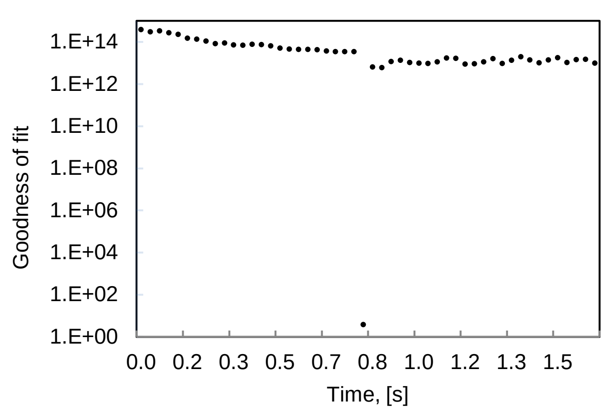

Fig.4. Goodness of Fit for measurement series with 90^0 phase jump

(The slow decline and small “wiggles” in the goodness of fit numbers are artifacts of our solution method. They result from the tolerance of an iteration, and depend on initial conditions.)

About 0.8 s into the measurement sequence, the model and the world do not make a good match. That is the information conveyed by the goodness of fit metric. In essence, the measurement system is saying “Here is how well the model you have for the signal matches the actual signal.”

Werner Heisenberg famously said “. . . since the measuring device has been constructed by the observer, we have to remember that what we observe is not nature in itself, but nature exposed to our method of questioning” [6]. The question we have asked by connecting a PMU does not include any mention of a phase jump. The answer will therefore not include mention of a phase jump, and we will know nothing of such an event when we see the PMU output.

The study of what is *known* is epistemology. Semantics is the study of *meaning.* Measurement is surely the process by which something in the physical world is given meaning, and becomes something known. That is why we said earlier that a kind of noise process that changed the value of the result of a measurement could be called semantic coloration, to distinguish it from random noise.

The measurement process works by finding the value of the parameters of an equation, the measurand. The measurand is, in other words, a conceptual model of the physical world, or at least a little part of it. It determines the question you ask of nature, and therefore, the kind of answer you will get.

7.3.5 Changing values: Frequency

We have selected frequency to illustrate the problem of a changing parameter because there is no universal agreement on the meaning of terms associated with changing frequency. A very readable discussion is given by Boashash [7], who reviews the treatments given to the issue since the introduction to broadcasting of frequency modulation in the 1930s.

Whatever the label, with a perfectly “clean” signal, a PMU can make a frequency measurement of great accuracy in a couple of cycles of the power system. The resolution can be as low as 1 mHz on such a measurement, equivalent to 17-s of zero-crossing counts.

7.3.5.1 Ramping frequency

But what if the frequency is ramping? First, we need to consider how much the frequency changes. During the few seconds immediately following loss of a generator from the interconnected power system, the frequency might fall at a rate between 10 mHz/s and 100 mHz/s, depending on where it was, and the amount of lost generation. If the change in frequency is not dealt with in the definition of the measurand, the goodness of fit will be reduced. However, the practical implication is minor, because noise on the signal will tend to mask the ROCOF signal.

We verified the effect of noise experimentally for a ROCOF of 10 mHz/s. With no noise on the signal, if the model fails to include ROCOF, the goodness of fit would drop from about 10^{12} to about 10^{9}. That change is no cause for concern: a number such as 10^{9} still means that the model is a very good match to the signal.

But in practice, the number is not likely to be achieved because of quantization effects in the A/D converters, and noise in the signal. When the signal contains 0.2% noise (the maximum mentioned in the standard), the residuals are larger and the goodness of fit, even with ROCOF included, is reduced to about 2.5×10^3.

This value of goodness of fit means that the effect of noise is to mask any problem caused by failure to include ROCOF in the measurand. In other words, with this level of changing frequency, there is no problem that need be dealt with. The signal is close enough to stationary that the result is the same.

Suppose the ROCOF is much larger. Load rejection by a generator could result in overspeed, and a positive value of ROCOF that might be as much as 3 Hz/s, say. In this situation, with the same level of noise as before, if the model accounts for ROCOF, the goodness of fit is still 2.5×10^3. This means that the noise on the signal is limiting the goodness of fit.

If the measurand does not include ROCOF, the goodness of fit is reduced to about 900. The goodness of fit is no longer limited by the noise, it is low because the model (without ROCOF) is a poor match for the signal.

The difference between those numbers, 900 and 2500, means that when the ROCOF is large enough to get the signal out of the background noise, the non-stationary signal is adequately dealt with by including the change in the measurand equation. In other words, one deals with the non-stationary nature of the signal by explicitly acknowledging its change.

Figures 5 and 6 present the information graphically. In Fig 5, the goodness of fit for the two cases where the ROCOF is large enough to be noticed are compared. The variations in the bands are the result of the randomness of the added noise.

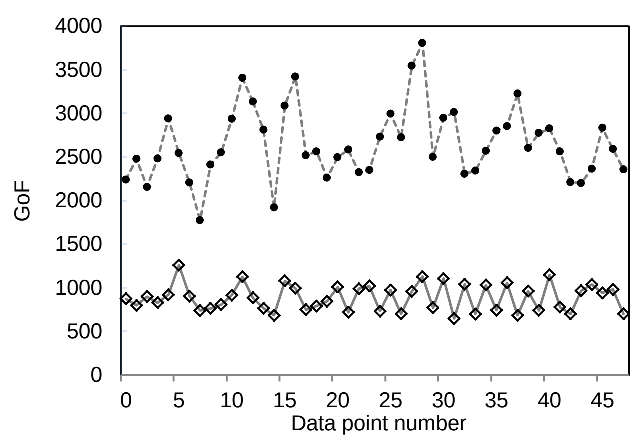

Fig.5. Goodness of fit for ROCOF= 3 Hz/s, included in measurand (top) and not included in measurand (bottom), for 50 measurements

Figure 6 shows the residuals over a two-cycle period for the case where the ROCOF is 3 Hz/s. It is clear that the residuals contain information, though further study would be needed to understand the message.

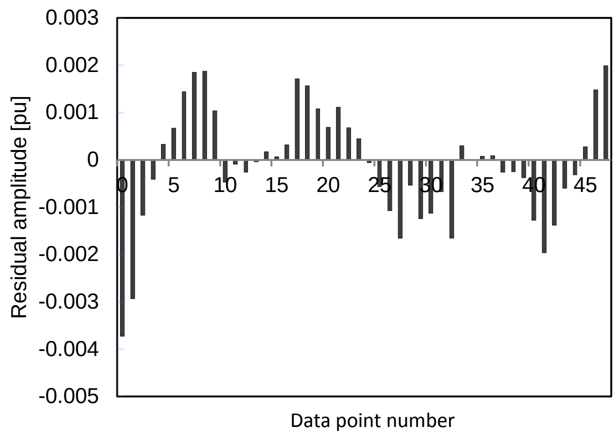

Fig 6. Residuals over two cycles (48 samples) with ROCOF = 3 Hz/s but not included in the measurand.

7.3.5.2 Phase modulation

We have been asked on more than one occasion about the PMU measuring parameters such as the modulation frequency during a compliance test. The test consists of submitting (say) a phase-modulated signal input to the PMU: the 60-Hz system frequency is modulated at a rate on the order of 1 Hz. The generating equation is of the form

$$v(t) = X_{\mathrm{m}}\cos(\omega_0 + k_{\mathrm{a}}\cos\omega_{\mathrm{mod}}t) \qquad (9)$$

where ω_0 is the power system frequency, k_{a} a modulation amplitude factor related to modulation index, ω_{mod} is the modulation frequency. Since the PMU makes a measurement in just a few ms, it should be clear that it cannot, from that measurement, estimate the value of a modulating signal with a frequency on the order of a Hz. Nor, it should be pointed out, is the PMU *asked* to make such a measurement. The parameters are not in its measurand. That is not to say that the parameters could not be deduced from multiple measurements in sequence. The digital measurement system making a fast measurement of a changing signal is, in essence, sampling a waveform.

Within the window of measurement, the PMU measurand includes a term in ROCOF and the changing frequency problem is solved the only way a PMU can solve it: it treats it as a ramp.

If the sample windows are long compared to the modulation period, the model would have to be changed to account for the changes in the character of the signal. That is beyond the scope of the present paper.

7.3.5.3 Conclusion

A solution to some of the problems of measuring a changing signal is to ascertain, from the measuring system, the quality of the match between the measurand and the signal being measured.

The goodness of fit parameter can indicate the presence of noise, and changes in the signal that affect the match. It is a metric that can be attached to any PMU, and other measurements beside.

A signal that has a small rate of change may be well-enough approximated by a stationary measurand. A signal with a large rate of change that is unaccounted for in the measurand may make itself evident in the residuals.

Examination of the residuals can suggest a change in the measurand—for example, the addition of a new parameter to the equation being solved. We have observed, for example, that some A/D converters seem to have a small dc offset, and the quality of fit can be improved by correcting for that by including a dc term in the equation.

The question of how to measure time-varying signals seems solvable. The question that has to be addressed is how to interpret the results of the measurement. We suggest that mathematics provides the rules for that.

7.3.6 REFERENCES

[1] E. Golding, Electrical Measurements and Measuring Instruments, London: Pitman and Sons, 1938.

[2] F. Harris, Electrical Measurements, New York, NY: Wiley and Sons, 1962.

[3] Nelson, T, "Definitions for calculations of VA, VAh, VAR, and VARh for Poly-Phase Electricity Meters, NEMA C12.24 TR-2011," NEMA, 2011.

[4] H. Kirkham, "A Conceptual Framework for Measurement, Report PNNL-24071," Pacific Northwest National Laboratory, Richland, WA, 2015.

[5] *IEEE Standard for Synchrophasor Measurements for Power Systems,* pp. 1-61, Dec, 2011.

[6] W. Heisenberg, Physics and Philosophy: The Revolution in Modern Science, London: George Allen and Unwin, 1959.

[7] B. Boashash, "Estimating and interpreting the instantaneous frequency of a signal - Part1: Fundamentals," *IEEE,* vol. 80, no. 4, pp. 520-538, 1992.

7.4 Error Correction: a proposal for a Standard

7.4.1 Introduction

Measurement connects the physical world to a conceptual model: an equation. In a moving coil instrument, the "equals sign" is manifested by the balance of the torque generated by the current in the coil against the torque of a return spring. In a digital instrument, the equation solving is more obvious. The equation being solved is the measurand. It was shown in [1] that the notion that measurement is equation-solving has many useful consequences. The simplifications and approximations of linguistic labeling are done away with: a concrete definition is readily found for measurands that have been problematical in the past—frequency, for example, when frequency is changing. The residuals can be used to calculate a metric that indicates the quality of the measurement.[2]

In this paper, we use the notion that the measurand is a model, an equation, to examine high voltage measurement.

7.4.2 Measurement Framework

The measurement framework, drawn as a block diagram, establishes the relationships between parts of the measurement process well known to metrologists. It can also be drawn to show the transducers used, as in Figure 1. In the figure, the solid arrows represent physical links, the open arrows conceptual ones. Figure 1 shows a measurement system such as a phasor measurement unit (PMU), in which the measuring instrument is made as accurate as required, and calibrated from its input terminals. In use, the measurement result will inevitably contain the artifacts added by the transduction system, typically using instrument transformers (IT).

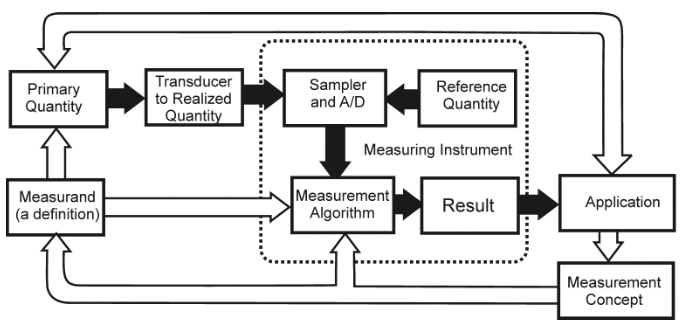

Fig. 1.Measurement framework, with transducer included

Applied to high voltage measurement, an isolation transformer, a CCVT or some active system based on field measurement may be involved. The point is that the primary quantity (the thing of concern to the application) is separated from the measuring instrument. The realized quantity is supposedly a scaled copy of the original, but the scaling will be imperfect, to a greater or lesser degree.

7.4.3 Changing the Equation

The artifacts of the transducer that cause differences from perfect scaling are generally known as the cause of Type B uncertainties. In particular, the scale factor could be off-nominal, and there will be some frequency-response effects. These uncertainties are generally characterized during calibration tests.
For high voltage and current measurements, instrument transformers (IT) are commonly used, with specified "accuracy class." Both IEC and IEEE specify the relationship between ratio error and phase error graphically. The IEC standard [3] shows the boundary of the permitted region by rectangular boxes on a plane of ratio error vs phase error for each IT accuracy class without restrictions on the range of their load power factor. The IEEE standard [4] uses parallelogram boxes to describe their accuracy class, with the load power factor being limited within the range of 0.6 to 1.0. There are proposals to change the IEEE standard to the "square box" accuracy classes.

It is proposed that at least some of these artifacts can be compensated for by changing the solution method in the measurement algorithm. Calibration data for the transducer can be added to the information available to the algorithm, and a solution found for the measurand that more closely represents the primary quantity. That could put the result of the measurement well *inside* the uncertainty box.

Consider the example of the PMU. The PMU produces four parameters in its result. It assumes the signal is sinusoidal, and gives the amplitude, frequency, phase (with respect to a well-defined reference) and the rate of change of frequency.

While devices are made that demonstrate extremely accurate results during lab testing, no allowance is made for the errors that may exist in the IT.

The requirements levied [5] on the PMU limit the uncertainty of the amplitude and the phase results, which are referred to as the *synchrophasor*. The standard combines the allowed error on these parameters in a vector fashion that is tantamount to an upper limit of about the same magnitude as that of the IT. In other words, a compliant PMU and a compliant IT may produce a non-compliant result.

Correcting for the errors would allow improved system analyses. The idea is not new. Passive correction was applied as long ago as 1912 [6] to reduce amplitude errors. In 1991, an active electronic system was demonstrated as a range extender [7] with low errors over a wide frequency range. An on-line system was simulated [8] in 2000. A system of correcting for errors that were temperature dependent was simulated in 2008 [9]. One of us (McBride) has been correcting for PT frequency-response errors off-line for some while, with excellent results. But none of this indicates routine use.

7.4.4 IV. Standardization

It may as well be assumed that measurements are done by digital equipment. The capability and flexibility of measurement systems greatly exceeds that of analog predecessors. The technology has been shown to exist to inform the measurement algorithm of the parameters of the transducer, and have the effect of the errors reduced.

The process of informing the algorithm could be standardized, with the information transferred by means of the Transducer Electronic Data Sheets (TEDS) defined in [10].

Accomplishing this would mean that the characteristics of the IT would have to be described for the TEDS, a greater effort than merely certifying compliance. But it is a procedure that can be accomplished on a Type-test basis.

The process can go further. A modified TEDS could also store a recording of the output of the A/D converters in a measurement system used during that characterization. Applied to the measuring system, that recording should reproduce a result that represents the primary quantity. Thus, a modified PMU that could read from a file instead of from its internal A/D converters could perform a self-check on demand.

Recalibration frequency depends on the stability of the IT: transformers are usually long-term stable, but CCVTs often exhibit drift. If the transducer is changed or recalibrated, a new recording should be available in the TEDS. The measuring system can use it to perform a self-calibration, because on playback, the same primary quantity results should be reproduced as before.

That notion levies a requirement on the characterization of the IT for the TEDS. The calibration system (including the A/D converters) that created the recording should have a test uncertainty ratio of at least four [11], so the self-calibration can be considered equally good. If that is done, a threshold for an out-of-tolerance result can be set.

7.4.5 Conclusion

A system view of measurement shows that the technology exists to correct for some transducer errors in real time. The standards exist to support the method. Automatic system self-calibration following transducer replacement is possible.

If these steps were taken, measurement accuracy and device user experience would be improved, and greater confidence in the measurement system would be established.

The authors are considering proposing standard changes to the Instrumentation and Measurement Society.

7.4.6 References

[1] H. Kirkham, “A Conceptual Framework for Measurement, Report PNNL-24071,” Pacific Northwest National Laboratory, Richland, WA, 2015. Available from https://certs.lbl.gov/publications/conceptual-framework-measurement

[2] Kirkham, H. and Riepnieks, A. “Dealing with Non-Stationary Signals: Definitions, Considerations and Practical Implications” to be presented at the IEEE PES General Meeting 2016, in Boston MA.

[3] IEC TC 38 Instrument transformers, IEC Std 61689-2007.

[4] IEEE Std C57.13 IEEE Standard Requirements for Instrument Transformers, 2008

[5] IEEE Standard for Synchrophasor Measurements for Power Systems, pp. 1-61, Dec, 2011.

[6] Craighead, J.R. “Potential Transformer Testing,” Trans AIEE, 1912, Vol XXXI, No 2, pp 1627-1633

[7] Miljanic, P. N.; Angelo, D.J. “High precision range extenders for an AC calibrator up to 600 V and 100 A,” IEEE Trans I&M, 1991, Vol 40, No 2, pp 393-395

[8] T.K Hamrita, B.S. Heck, A.P.S. Meliopoulos, On-line correction of errors introduced by instrument transformers in transmission-level steady-state waveform measurements, IEEE Trans Power Delivery, 2000, Vol 15, No 4.pp 1116-1120.

[9] Sun, J., Wen, J., Lau, L, and Li, X "Steady-state error analysis and digital correction for capacitor voltage transformers" International Conference on Electrical Machines and Systems, ICEMS 2008, pp 769-773, Pp 769 – 773
[10] IEEE Standard for a Smart transducer Interface for Sensors and Actuators – Common Functions, Communication Protocols, and Transducer Electronic Data Sheets (TEDS) Formats, IEEE Std 1451.0, Sept 2007.
[11] ANSI/NCSL Z540.3-2006 (R2013), "Requirements for the Calibration of Measuring and Test Equipment".

7.5 Introduction to Goodness of Fit for PMU Parameter Estimation

7.5.1 Introduction

For voltages that are supposed to be alternating periodically, the signal is usually described by a phasor equation:

$$x(t) = Acos(\omega \cdot t + \varphi) \tag{1}$$

where A is the amplitude, ω is the angular frequency and φ is the phase.

Some of the most critical measurements in power system are assumed to be represented by such phasors, and power system measurements have long been based on this representation. Following the observation of slow phase changes in the power system [1] by phasor measurement units (PMUs) were developed and are now widely applied [2]. The phasor is an ideal construct: its parameters are time-invariant, ie, characterized by stationary values. However, it is a characteristic of the power system that the voltage or current amplitudes, frequency and phase angles are rarely constant for long. These power system quantities can be referred to as “phasor-like”, but actually there is no clearly set definition of the term [3]. Nonetheless, a phasor measurement unit (PMU) is expected to measure the real signal.

The PMU allows for the changeable nature of the signal in two ways. First, it measures over a short-duration window. During the window, things it measures are reported as constant: the *result* of a signal measurement is always a fixed thing, even if the signal is changing. Because the frequency might be changing, the PMU is also required to measure the rate of change of frequency, or ROCOF. This, too, is assumed constant for the duration of the measurement window.

The things the PMU measures, therefore, are the three coefficients of Equation (1), plus the rate of change of the frequency. Some preliminary results were given in [3] that showed these parameters being found simultaneously by regarding the problem as a fitting problem in mathematics. Further work in our Laboratory since then has confirmed the value of that approach to measurement.

However, this paper is not concerned with how the measurement of the PMU signal is implemented. Our focus in this paper is on one particular outcome that can be applied to any PMU. Our topic is goodness of fit, as applied to the results of a phasor measurement unit.[3]

7.5.2 Goodness of Fit

7.5.2.1 Definition

We showed in [4] that when an experimental measuring system similar to a PMU is used to solve (1), there are residuals. These are artifacts of the real world, the result of noise and distortion on the signal, essentially the part of the signal that is not explained by the results of the measurement. Residuals exist after any measurement is made, and our purpose here is to show how the residuals from a measurement can be used

[3] The term goodness of fit is connected with phasor measurement units only once in IEEEXplore, and that application is quite different. The interest of that paper is a matter of fitting state estimation results to an entire power system model. (Choi & Meliopoulos, 2016)

to increase the value of the result. That goes beyond just phasor measurement: the measurement of a supposedly constant voltage is given below as an example.

In [4], we defined a Goodness of Fit metric that ranged over many decades. At the suggestion of Ray Hayes of ESTA International, we have now adopted a compressed (logarithmic) version, in which the numbers are more manageable. We propose the use of a logarithmic description because of the large dynamic range needed, and the use of the reciprocal because the number would be bigger when the fit was better.

We define the Goodness of Fit metric (GoF) as the reciprocal value of the fit standard error [5] [6], normalized. The calculation depends on the number of degrees of freedom of the equation. Expressed in decibels:

$$\mathrm{GoF} = 20\log\frac{A}{\sqrt{\frac{1}{(N-m)}\sum_{k=1}^{N}(u_k - v_k)^2}} \tag{2}$$

where N is the number of samples, m is the number of parameters being estimated in the equation (one more than the number of degrees of freedom), A is the signal amplitude, u_k is the signal sample value and v_k is the estimated sample value. The parameter $(N - m)$ is called the *residual degrees of freedom* in [5].

7.5.2.2 Measuring dc

Consider the example in [7] of a dc/dc power supply designed to furnish a steady 400 volts. On a multi-digit voltmeter, that was what it seemed to do. Unbeknownst to its designer, it was actually oscillating. Figure 1 shows the output voltage.

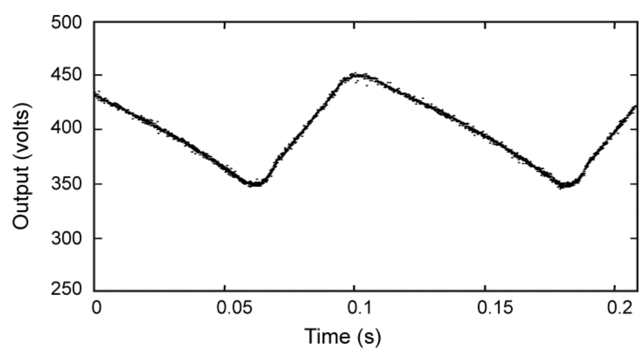

Fig.1. Input voltage of unstable power supply

The peak signal amplitude in this case is 450 V. The oscillating part of the output has an almost triangular waveform between 350 V and 450 V, ie, it has an amplitude of 50 V. The rms residual value is thus 50/√3, or about 29 V. We can use that number to represent the residuals, on the assumption that the *model* is of a constant voltage. The GoF is therefore 20 log(450/29), about 24 dB. Had the “ripple” been just one volt rms, the GoF would have been about 58 dB.

GoF is thus seen as a comment on the *match* between the selected model (in this case, dc) and the signal, rather than an estimate of the accuracy of the measurement. In the case of the dc/dc converter, the voltage shown in Fig 1 may be declared by a measuring system with a long integration time as 400.0 ±0.1 V with a confidence level of 95%. But it is still not a good measurement, because the model of direct voltage is a poor fit to the signal. For the power supply designer, it would have been useful to know that the fit was poor, had the measuring instrument been capable of indicating it.

Actually, the instrument had all the information needed to indicate the GoF, and the calculation would not have consumed much in the way of computing resources.

We will see next how the PMU model compares, measuring an alternating signal.

7.5.2.3 Measuring ac: background

We were able to obtain data from (commercial) devices monitoring a 345-kV EHV transmission system. The data consisted of two kinds: one was a PMU, set to a 30-per second reporting rate in accordance with the PMU standard [8]. The other was what is called oscillography, recording sample values at 64 samples per cycle. The results we show below are from these devices.

Our procedure to verify that the PMU was solving (1) was to take the PMU measurement results for a given report-time and insert the parameter values into equation (1) along with a version of the time that corresponded to the times of the sampled values as recorded by the relay. The extraction of that time information requires some explanation.

7.5.2.4 Timing definition

The PMU standard requires that the measurement of the signal phase as reported by the PMU is the phase relative to a reference signal that is standardized with reference to UTC. Most PMUs rely on the widespread availability of a timing signal from navigation aids such as GPS, and it is assumed in the standard that time is known with good precision—certainly good enough for this purpose.

The reference wave is defined as a cosine wave that reaches a positive maximum at the time of the UTC second tick, and has a frequency that is exactly the nominal value of the power system.[4] There are exactly 60 (or 50) of these peaks in a second, and for reporting rates up to 30 (or 25) per second, every report from the PMU will be time-tagged at one of them. The information used in the measurement comes from the power system signal in the time before and after the peak of the reference wave. For a P-class PMU, there will be about one cycle before the time tag, and one after.

For the reference wave, the definition means that time = zero corresponds to one of the positive peaks. Since the wave is (by definition) a perfect cosine wave, each cycle is the same as the next. That means that a time of zero can be assigned to *any* wave of the reference at its positive peak. Those times are known precisely, now and forever.

All sample-times are just offsets from those defined times. The time of any second tick can be called zero without changing the result of the measurement, and the other sample times re-labeled accordingly.

---

[4] Standards can make statements like that. There are exactly 2.54 cm in one inch, because that is what the standard says defines the inch. No matter that the real world of measurement always involves some uncertainty. While the uncertainty of timing on a PMU oscillator may be a microsecond or two, the reference wave reaches its peak at the second tick.

Therefore, in the recorded data files, any time with an integer number of clock-seconds (UTC) can be called zero, and so on. The PMU equation can thus be evaluated at the times used in the relay to sample data: the process should regenerate the data that the PMU would have internally generated, if it had been sampling at the same rate as the relay.

Therefore, for each sample value of the relay, a corresponding estimate value can be reconstructed (using equation (1)), and the residual equation (2) can be evaluated.

7.5.2.5 Phase definition

It is true that the relative phase of two signals that are not at the same frequency is not mathematically defined. Consider the situation where two signals that are close in frequency are compared. If their frequencies are constant, they will appear to "beat" at a constant rate, going in and out of phase periodically. The phase is not defined because it is a time varying function that will depend on the details of the frequency relationship between the two signals.

For the PMU, the reference signal is nicely defined, but the power system apparent frequency wanders about as loads change and generators correct whatever drifts are detected. Yet we demand a value for something we label "phase." How is the PMU to deal with this situation?

It is not a new question. An exploration of the term "phase" was undertaken by the great B. van der Pol, who gave a paper [9] on the topic in London in 1945. In this paper, van der Pol shows that many investigators preferred to refer to the entire argument of the cosine term, that is ($\omega t + \varphi$), as the phase. We could call it the "total phase." He argued that considering this pair of terms as the phase allowed the phase of signals of different frequency to be described mathematically. He explained his reasoning as follows:

> This definition has, among others, the advantage of enabling one to speak of a phase difference of two oscillations of different frequencies. This phase difference is then simply a linear function of the time, just as one phase by itself is already such a function of the time.

In essence, van der Pol lost the argument. In the convention of power engineers, the cosine argument of equation (1) is evaluated at time $t = 0$, and the result is called the phase.

However, it must be pointed out that the information needed to fix the frequency-dependent part of the total phase is available in the PMU, so that the relative phase of two signals may be *evaluated* for any *given time*. In terms of sampled values, all the information needed to make the calculation is available.

As an aside, we mention that the temporal locality (ie, the zero of time is nearby) means that the phase reported by the PMU is always the principal value of the phase, the value closest to zero. If the zero of time were fixed at some point that was not local (say, midnight yesterday), the relative phase of the power system and the reference could become quite large. A time error of a second on the power system is equivalent to a relative phase of 60 cycles or 21 600 degrees. The PMU would report that as zero, nevertheless. The PMU could be designed to record accumulated phase since a given reference, but it is not. The result is the need for some users to "unwrap" the accumulated phase value by considering the whole history of the phase since their chosen reference time.

7.5.2.6 Measuring ac: results

Figure 2 shows two cycles of data from part of a high voltage transmission system. One curve is based on the values from the relay, the other on the values reconstructed from the PMU measurement results. Figure 3 shows the residuals from these two data sets.

The dashed line in Fig 2 is based on oscillography. The solid line is based on the PMU reconstruction. The lines are very close together, but there is in fact a small separation. The recordings last one second, and start before a fault on the system. The figure shows results taken from times near the end of the one-second record from the relay, because at that time the power system has settled down to normal operation.

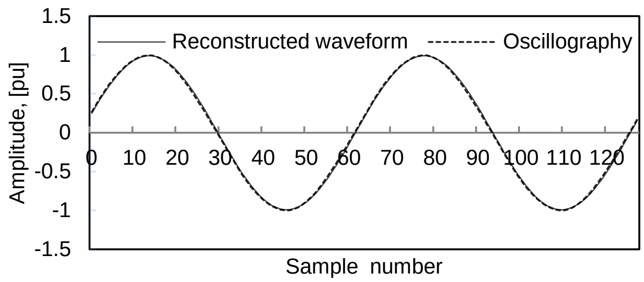

Fig.2. PMU and relay data compared, 345 kV system (1).

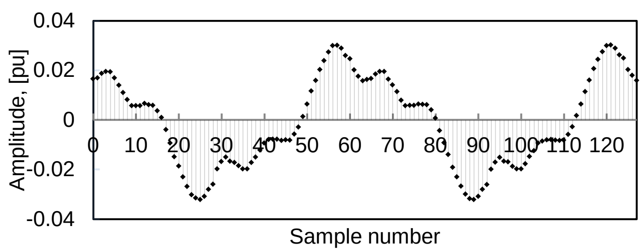

Fig.3. Residuals for the two signals in Figure 2.

The residuals in Figure 3 make the difference clearer. The residuals show a considerable harmonic content, but there is a large fundamental component. It appears to be about 90 degrees out of phase with the main signal. There seems to be a constant phase angle difference between the signal and the measured result of almost one degree. We have no insight into the cause. We speculate it may be a timing error, but possibly it is the result of instrumentation transformer error.

Note that the residuals are about 0.03 pu (peak), or about 3% of the size of the signal.

Phase error notwithstanding, we conclude from the close match between the relay samples and the reconstructed wave that the PMU is doing a fine job of solving equation (1) even though it was not designed with that equation in mind.

We see similarly good fits for other data-set pairs from other times and places. Figures 4 and 5 show records from records immediately *before* a relay sampled a fault. This is information from a different location and a different time, and therefore based on a different set of equipment. As before, the commercial PMU is giving results whose fit is very good. The residuals here are just slightly larger in magnitude.

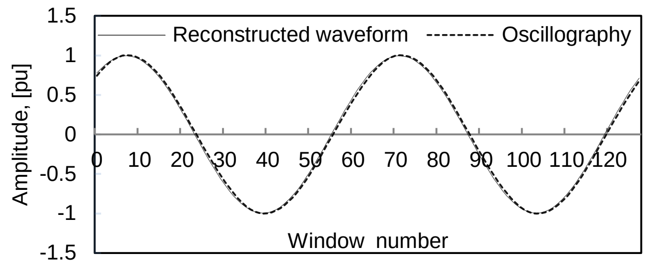

Fig.4 PMU and relay data compared, 345 kV system (2).

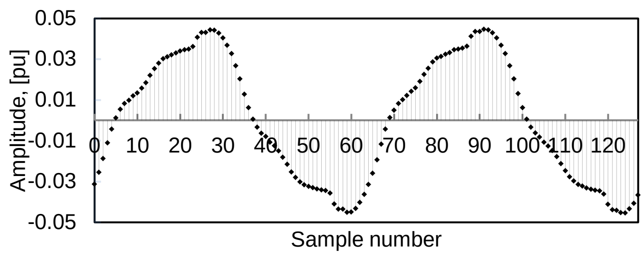

Fig.5. Residuals for the two signals in Figure 4.

The PMU, *any* PMU, can describe the world only as a cosine wave. When a cosine wave is not a good picture of the world, the PMU cannot change the things it reports. Therefore, when there is a fault, the PMU will report the signal as if it were a cosine wave, even if it clearly (to the human eye) is not. The result is a lower value for GoF.

The Goodness of fit metric was calculated for all the results through a one-second fault recording. Figure 6 shows the GoF across the recording that gives Figure 2 and 3, along with the current values reported by the PMU. At the time corresponding to Figs 2 and 3, the GoF is about 35 dB. Figure 7 corresponds to Figures 4 and 5, and again adds the PMU current magnitude result. The GoF is between 30 and 35 dB except when there is a fault.

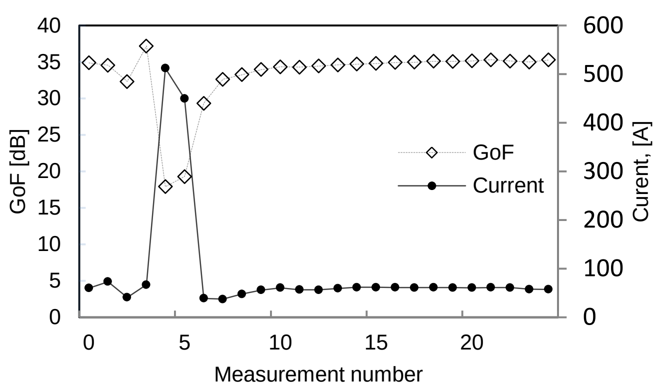

Fig.6. Measured current and Goodness of Fit during fault (1).

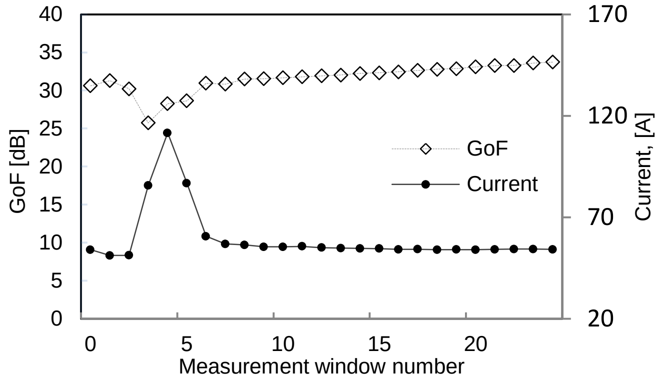

Fig.7. Measured current and Goodness of Fit during fault (2).

It is evident in Figures 6 and 7 that the GoF drops to about 20 or 25 during the fault (indicated by the current spikes), a decrease of about 10 dB. Note that the GoF value is calculated at the same time as the other values that come from the PMU. It is not necessary to wait until a comparison of numbers in a sequence show that there *was* a fault. The GoF provides information that the wave is no longer sinusoidal in shape.

That is new information, and it is available in real time. If there were an application using the PMU data in some sort of control scheme, that would be good information to have. It may even prove to be of value to a human operator to know that some particular PMU results are not to be trusted.

7.5.2.7 Distribution system data

"Anonymous" but real point-on-wave data for a medium voltage distribution system was provided by the ARPA-E microphasor project [10]. Unfortunately, the PMU measurement results corresponding to these sample values were not available. Therefore, we calculated our own results, using our fitting-solution method of measurement. (It is a version of the work reported in (Kirkham, H; Dagle. J., 2014).) The input data to our "PMU" consisted of the analog/digital converter output signal without any filtering or other manipulations. The sampling frequency is 30 720 samples per second (512 samples per cycle). We are confident that the results are not significantly different from those that would be obtained from the micro-PMU.

From the input data we constructed two-cycle measurements windows and ran our estimation algorithm. Some results for frequency are shown in Fig. 8.

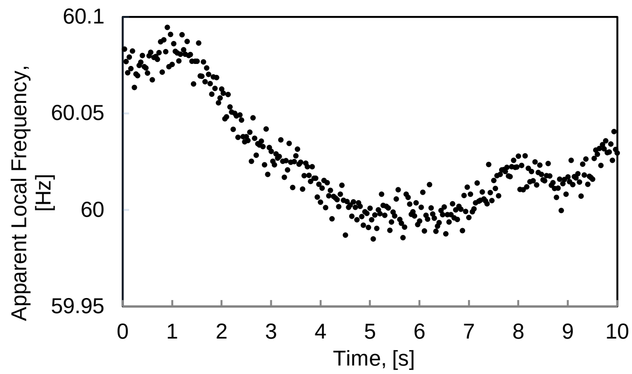

Fig.8. Apparent Local Frequency estimate for 10 second real data.

The spread of numbers around an average is ten or twenty mHz. Bearing in mind that there is no frequency-domain filtering in our processing, and that these are the results of *independent* measurements of two cycles width, we were reasonably satisfied with the results.

We found that the residuals were slightly smaller (40 dB down instead of 30) than those for the transmission system signals. This can be seen in Fig. 9, which shows how GoF varied during the same period. We do not have information to say that the improvement is due to our method of solving equation (1), since we do not have the matching PMU measurement results.

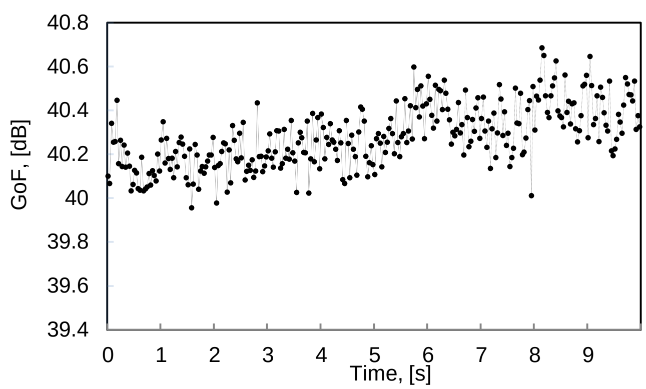

Fig.9. GoF values for fitted data.

Examination of the waveform revealed some distortion, as shown in Fig. 10.

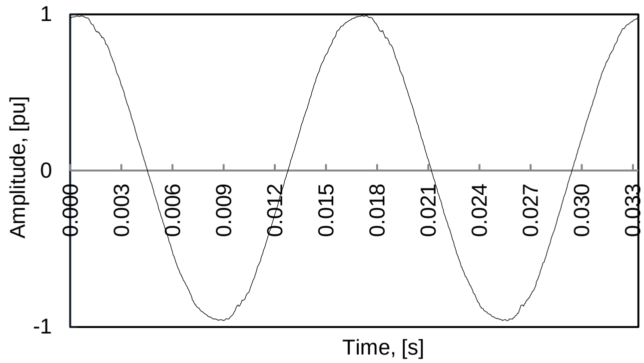

Fig.10. Waveform for distribution network real-world signal.

We subjected the sample stream to FFT analysis. It can be observed from Figure 11 that the signal contains 3rd, 5th and 11th harmonics along with noise. These harmonics could account for at least 2% of total signal magnitude.

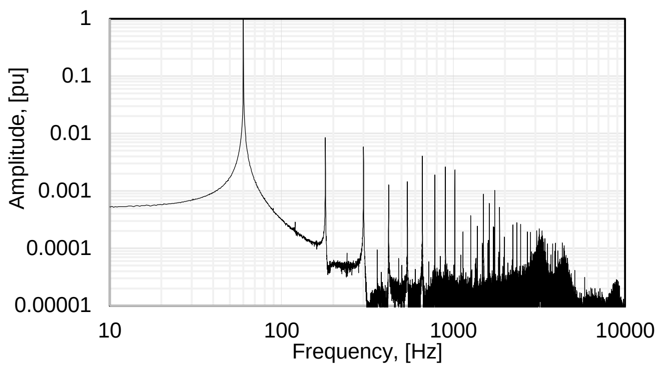

Fig.11. Spectrum for distribution network real-world signal.

7.5.2.8 Fault response

In situations that involve faults and switching operations there might be significant distortions and phase jumps in the signal. We saw in Figure 6 and Figure 7 some excursions on the Goodness of Fit. We thought it would be instructive to present the signal waveforms that correspond to fits of this kind.

We illustrate matter in Figure 12, which shows the oscillography obtained during the fault (corresponding to Fig.6). This data is for the line voltage on phase C, the two readings that correspond to the times when the fault current is above 300 A. There is a slight break in the “estimated” curve at sample number 128, when the reconstruction is based on the second of the two PMU results and not the first.

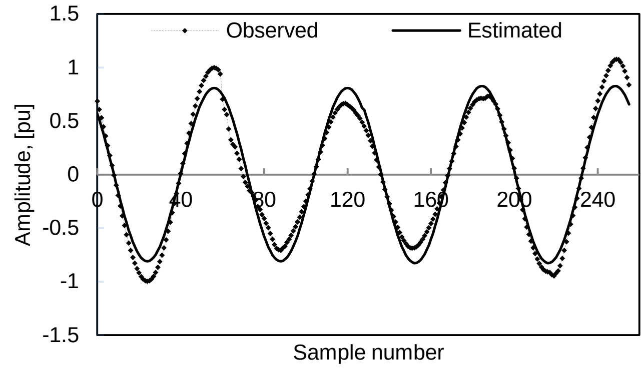

Fig.12. Line voltage on transmission system during fault.

A P-class PMU makes a measurement of the input signal using an observation window that is two cycles wide. Considering just the amplitude of the peaks of the signal, it is evident from the figure that while the fault is applied, there are no two adjacent cycles with the same peak amplitude. Yet the equation that the PMU solves, equation (1), forces the PMU to assign a constant value to the amplitude. As a result of that constraint, the PMU estimate is sometimes larger than the signal, sometimes smaller. The residuals for this time in the record tell the story: they are ten times larger than any we have presented so far from the real world, and they have a large in-phase component, as in Figure 13.

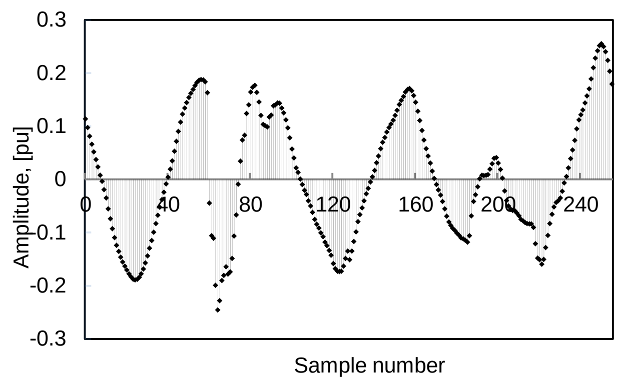

Fig.13. Residuals during fault.

7.5.3 Setting a Threshold

We wondered early on whether a threshold can be established to distinguish an acceptable GoF number from an unacceptable one. We have so far a rather limited number of examples to consider, but it does seem that it may be possible.

For the unfaulted signals we have examined, the GoF is rarely above 35 dB, however, and for the faulted examples it may be only 30 dB as the fault is starting or ending. There does not seem to be much “wiggle room.” However, an adaptive scheme that set a threshold say 6 dB down from the average of twenty of thirty consistently “good values” might reasonably be interpreted to be signaling something has changed if the GoF dropped abruptly in that manner.

Of course, a sudden drop below a fixed threshold of, say, 20 dB would also be indicative of a problem. Expressed in the manner of an r-squared value, it would indicate that fully ten percent of the signal was not explained by the PMU model.

It is curious that the GoF for the distribution system results seems to be slightly better, even though the waveform is quite obviously distorted. We are presently investigating possible reasons.

7.5.4 Discussion

We have heard users discuss the inability of a PMU to give meaningful values of parameters like the current during a fault. There is more than a hint in these discussions that the problem is one of filtering. Perhaps that lies behind the push for ever-faster reporting rates from PMUs.

The problem is not one of filtering; it is considerably more fundamental. The PMU measurement gives its results on the assumption that the signal it is measuring is a cosine-wave. If the signal is not a cosine wave, we should be cautious in interpreting the results of the measurement.

Werner Heisenberg famously said ". . . since the measuring device has been constructed by the observer, we have to remember that what we observe is not nature in itself, but nature exposed to our method of questioning"[11]. In other words, if you ask what kind of fixed-amplitude cosine wave the signal is, that is going to be the nature of the response.

We can interpret the fault results as follows. A low Goodness of Fit number informs the user that the model (the measurand, the description of the thing to be measured) and the sample stream are not a good match. It is not a matter of assigning blame: the simple fact is that the data are not particularly well represented by the model. . In Heisenberg's terms, we have selected an inappropriate method of questioning.

It is interesting to speculate on why the notion of a goodness of fit has not been applied to measurement before now. We think that the reason may be in the "different" world of what might be called "pure" metrology compared to "applied." Pure metrology is defined, in large part, by documents such as GUM, the Guide to the Expression of Uncertainty in Measurements [12]. This document guides metrology around the world. In particular, it makes clear that the presentation of the result of a measurement should include not only the declared value (the reading of the instrument), but also a statement of the uncertainty. The uncertainty statement should consist of two numbers that essentially define the location and dispersion of a distribution of other possible results. One might write, for example, that a given value was measured as 120 volts ±1% (2σ). That would mean that if you made 100 measurements of this quantity using this instrument, you could expect close to 95 of them to be within 1% of 120 V.

There is an assumption behind that way of stating the result, and it is written in GUM. GUM observes (Sec. 3.4.1) that "implicit in this Guide is the assumption that a measurement can be modelled mathematically to the degree imposed by the required accuracy of the measurement." In other words, the assumption is made that the model is accurate.

That is generally true in the metrology laboratory. For the engineers and scientists in that laboratory, measurement is the process of transforming a set of repeated observations into a single number (or a set of numbers), and dealing with extraneous influence quantities that are inexactly known. That is not the world of the PMU. The PMU gets one brief look at its world, and has to accept that as its realized quantity. There will be no repeated measurements.

The assumption that the model is accurate, the underlying assumption behind "pure" metrology, is what is tested by the goodness of fit metric. If the signal is distorted, the PMU can still report only values for amplitude, frequency and phase. But it can (and we think it should) include a comment on the fit of the model. That information can help the user. That is likely to be of increasing importance if the user is automated equipment.

One of the things many PMUs do is to calculate a positive sequence value for the things it measures. It seems like a fair question to ask whether the phase sequence values should also be tagged with the GoF metric. If the calculation of the sequence values is based on an assumption of sinusoidal shape (and that seems likely, as it is written that way in the textbooks), what will the result mean if the GoF is poor?

The matter of GoF goes beyond the PMU. Any signal/measurand pair could suffer a mismatch, and any digital instrument could make the GoF evaluation. The multi-digit voltmeter used to indicate the "precision" of the 400-volt power supply would have been much more informative if it had furnished the GoF. We have seen that the metric applies to any PMU. The user could adapt appropriately.

There is information in the residuals that could be subject to further analysis, either in real time or after the fact. It seems that "ordinary" residuals (Fig 3) do not resemble fault residuals (Fig 13) in magnitude or overall "shape." We suggest that there is interesting work to be done studying PMU residuals.

We have seen that *real* power system data is very noisy, and that leads to general worsening of Goodness of Fit values. We are now investigating the effect of noise on the fit, particularly to see if and when the ROCOF estimation can be so masked by noise that its values cannot be believed.

7.5.5 Conclusion

The Goodness of Fit parameter, developed from an idea in [7] and an offshoot of the estimator described in [3], has the potential to be very useful with real PMUs and real signals. It indicates in real time the degree of match between the signal (changing with the power system), and the measurand (fixed by the design of the PMU

1) About the Goodness of Fit Metric

- A Goodness of Fit level can be calculated by *any* PMU. The calculation is straightforward, and does not depend on the measurement method;
- The Goodness of Fit indicates that near-ideal results can be obtained with an ideal signal;
- The Goodness of Fit has such a wide dynamic range, it is probably best to use decibels for expressing the normalized reciprocal rms value of the residuals.

2) About the noise

- The true nature of real world power system noise characteristics is still not fully understood. In view of the influence of noise on the performance of the PMU, a program of characterizing power system noise might be justified.

We are optimistic that the GoF method will show good results on more real-world data and real-world PMUs than we have yet had access to. We plan to continue testing the method using PMU and oscillography data.

The method is not limited to PMUs, however. It could be implemented as part of any digital measurement whose measurand equation can be elucidated.

Altogether, GoF seems to be a promising technique for a large class of digital measurements.

7.5.6 Appendix on the PMU Equation

To see how the PMU equation relates to the real (changing) world, let us give each parameter of (1) a rate of change, with the exception of amplitude. Let us suppose that the rate of change is also a value to be measured, and therefore assume it to be constant across the window. We obtain:

$$x(t) = Acos\left\{\left(\omega' + \frac{C_\omega t}{2}\right)t + \varphi' + \frac{C_\varphi t}{2}\right\} \tag{2}$$

where A is amplitude, C_φ is the rate of change of phase, and C_ω is the rate of change of frequency (ROCOF). A mark has been added to the original variables to indicate that they no longer obey the stationarity requirement of the phasor, though they are stationary for the duration of the measurement window.

The term in $C_\varphi t$ cannot be distinguished and estimated separately from ω′, therefore it is moved, leaving φ′ as a constant. Equation (2) is then re-written:

$$x(t) = Acos\left\{\left(\omega' + \frac{C_\varphi}{2} + \frac{C_\omega t}{2}\right)t + \varphi'\right\}. \tag{3}$$

Equation (3) is a model of the power system with changing parameters. It is seemingly identical in form to equation (1). However, the coefficient of the *t* term in the cosine argument (the coefficient normally labeled “frequency”) now consists of three components. The first is the original frequency of the system, here labeled ω'. The second is due to the changing phase, and that may be due only to the fact that the power system frequency is not nominal, whereas the reference frequency is (by definition) at the nominal power frequency value. If the system is not only at an off-nominal frequency but also accelerating or decelerating, an additional term appears in the frequency coefficient, ${C_\omega t}/{2}$. As we saw earlier, C_ω is the ROCOF, so the effect on the frequency term can be seen to be to make it quadratic in nature.

The value of the quadratic element is very small compared to the other elements, because both the signal and the time are very small, the time being limited to the width of the observation window and centered on zero. We can estimate some values.

A large ROCOF, such as would occur when a large generator is dropped from the system, might be 30 mHz/s, that is 3×10-2 s-2. The third term in the cosine argument is thus that number times the square of half the window width. For a class P PMU with a two-cycle window at 60 Hz the last sample in the window contributes a value of 2.61×10-5. The corresponding nominal frequency contribution is 60 times half the window width, or 6.3 × 100, a number 5 orders of magnitude larger. Any reasonable value of ROCOF will therefore make only a tiny contribution to the quality of the fit. This accounts for the difficulty experienced measuring it in the usual noisy environment. It allows the approximation to the form of (1) to be acceptable.

Apart from that problem, equation (1) is indeed solved by the PMU [7]. The commercial PMU is not designed with equation solving in mind, but in fact the PMU does a very good job of solving it.

7.5.7 Acknowledgments

We thank Alex McEachern of Power Standards Laboratory for sharing with us the real-world sine-wave data recorded with a Power Standards Lab ARPA-E micro-synchrophasor instrument. We also thank Zak Campbell of AEP for furnishing oscillography data from the AEP power system. We also thank Raymond M. Hayes from ESTA International for many valuable discussions.

7.5.8 References

[1] M. Adamiak, "PAC World," Autumn 2007. [Online]. Available: https://www.pacw.org/en/issue/autumn_2007_issue/synchrophasors/synchrophasorsin_the_beginning.html. [Accessed 15 July 2016].

[2] J. De La Ree, V.Centeno, J.S. Thorp, A.G. Phadke, "Synchronized phasor measurement applications in power system," IEEE Transactions on smart grid vol 1, No.1, pp. 20-27, 2010.

[3] H. Kirkham, J. Dagle, "Synchronous phasor-like measurements," IEEE PES Inovative Smart Grid Technologies Conference (ISGT), Washington, DC, Feb, 2014.

[4] H. Kirkham and A. Riepnieks, "Dealing with Non-Stationary Signals: Definitions, Considerations and Practical Implications," in Proceedings IEEE Power and Energy Society General Meeting, Boston, MA, 2016.

[5] D. Cuthbert, F.S. Wood, Fitting Equations to Data, New York: John Wiley & Sons, Inc., 1971.

[6] The MathWorks Inc., "Documentation," 2015. [Online]. Available: http://www.mathworks.com/help/matlab/.

[7] H. Kirkham, "A Conceptual Framework for Measurement, Report PNNL-24071," Pacific Northwest National Laboratory, Richland, WA, 2015.

[8] IEEE Standard for Synchrophasor Measurements for Power Systems, C37-118.1-2011, pp. 1-61, Dec, 2011.

[9] B. van der Pol, "The Fundamental Principles of Frequency Modulation," Journal of the Institute of Electrical Engineers, vol. 93, no. 23, pp. 153-158, 1946.

[10] University of California Berkeley, "Micro-Synchrophasoprs for Distribution Systems," Advanced Research Projects Agency - Energy, 2015. [Online]. Available: http://arpa-e.energy.gov/?q=slick-sheet-project/measuring-phase-angle-change-power-lines.

[11] W. Heisenberg, Physics and Philosophy: The Revolution in Modern Science, London: George Allen and Unwin, 1959.

[12] Joint Committee for Guides in Metrology, WG 1, "GUM: Guide to the expression of uncertainty in measurement," 2008.

[13] S. Choi and A. Meliopoulos, "Effective Real-time Operation and Protection Scheme of Microgrids Using Distributed Dynamic State Estimation," IEEE Transactions on Power Delivery, 2016.

7.6 Rate of Change of Frequency measurement

7.6.1 Introduction

The phasor measurement unit measures the amplitude, frequency, phase and rate-of-change of frequency of an alternating signal from a power system. For a PMU to be “compliant,” each of these parameters must be measured with an error no greater than that specified in the relevant standard, currently C37.118.1 as amended by C37.118.1a, [1][2]. For example, the frequency error must not exceed 5 mHz over a range ±2 Hz from the nominal power frequency.

The accuracy of a PMU is verified during what is called in the standard a compliance test. A signal source of low uncertainty is used to generate a high-quality signal, the values of the relevant parameters being known. Broadly speaking, the test is a steady-state test, with the various parameters being set to known values and held steady for a while.

In the words of the standard:

> The reference condition specified for each test is the value of the quantity being tested when not being varied. Only the parameters specified for each requirement shall be varied as the effects shall be considered independent.

The parameters that define the phasor are not something that can be measured at a single point in time.

A frequency measurement by a zero-crossing method will take a long time to establish good resolution. Measurements other than those based on zero-crossings are possible, but all require more than one sample value [3].

Therefore, however the measurement is done in detail, the measuring device is obliged to examine a stretch of the signal that occupies some finite amount of time. This time may be referred to as the observation window, the measurement interval, or the sample period. The measured parameter may be said to be the value at its center.

For a window during which all the phasor parameters are constant, all is well. However, it seems there is no way to state the “correct” result for a measurement whose window spans a transition in a parameter value. Part of the window is characterized by a phasor model with one set of values, the other part by a model with different values.

Because of this, the standard has avoided levying performance requirements during transitions. In addition to specifying that the quantity being tested will not be varying, it also says that the measurements made either side of a transition will not be required to meet requirements on accuracy:

> The error calculation shall exclude measurements during the first two sample periods before and after a change in the test ROCOF. Sample periods are the reporting interval, 1/Fs, of the given test. For example, if the reporting rate Fs = 30 fps, then measurements reported during a period of 67 ms before and after a transition shall be discarded.

The purpose of avoiding the region around a transition may be to make some allowance for the filtering that is often used in measurements. This would have the effect of inserting some information from one measurement window into the result of the next, and could affect the declared value inappropriately. PMU reports are, indeed, not independent measurements because of this.

Nevertheless, the effect of that omission is somewhat unexpected. Figure 1 shows a figure from a paper describing a "dynamic phasor measurement unit test system" [4]. It is an illustration of the "Linear voltage frequency pattern" used for a test. The reader is told to "Note that the frequency and errors at the transitions are not reported."

We noted the blank spaces. We did not think them either fair to the PMU user or necessary to the PMU tester. We thought we could show a method whereby a justifiable "true value" can be derived for transitions such as the ones shown.

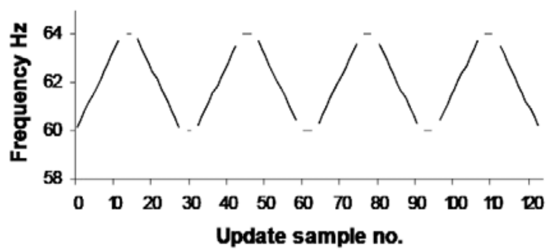

Figure 1. Copy of Figure 10 from Stenbakken-Zhou paper

Therefore, we tried to measure the PMU parameters across the transition region, expecting to find some sort of gradual swing from one value to the other. What we found instead was a curious aspect of measuring ROCOF, the rate of change of frequency.

7.6.2 Fitting as a Measurement Method

Several papers and reports have been written on the use of the fitting method as a mechanism of measurement [5] [6] [7]. The method combines several principles. First, it is recognized that the act of measurement is one that uses an effect in the physical world (the force exerted by a current in a magnetic field, for example) to obtain a value for a parameter in the conceptual world (the magnitude of a current). Second, the conceptual entity is viewed as a model of the real world. Third, the thing that metrologists call a measurand (a description of the thing to be measured) is that model. Fourth, it is best written in mathematical terms, rather than linguistic.

Measurement is then seen as solving an equation. Whether or not the PMU is designed with the idea of solving an equation, it does so very nicely. Figure 2 shows two curves representing the voltage on part of the real 345-kV system. One curve is the "PMU equation" with the values plugged in from the results of a phasor measurement (ie, the amplitude, the frequency and the phase) at the sample times of a recording device. This is identified as the "reconstructed waveform". The other curve is the sample values recorded, identified as oscillography.

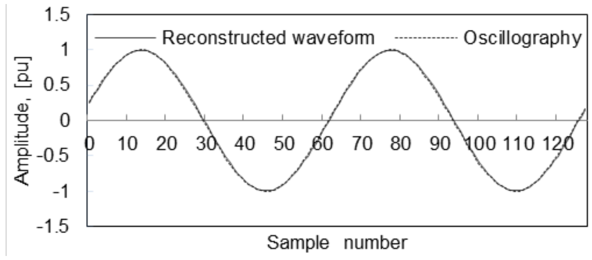

Figure 2. Match between PMU solution and sample values

The match is excellent for this measurement window. What this means is that the PMU equation is actually being solved by the PMU (the mathematical model is a close match to observed reality), even though the designer likely had no awareness of the possibility. A moment's reflection will remind us, however, that amplitude is just amplitude, and however we measure it, we should get the same answer. The same goes for the other parameters of the phasor, too, until we start to allow things to change.

The PMU equation (for our present purpose) is

$$x(t) = X_{\mathrm{m}} cos\left\{\left(\omega + \frac{C_\omega}{2} t\right) t + \varphi\right\}, \tag{1}$$

where X_{m} is amplitude, ω the frequency, C_ω the rate of change of frequency (ROCOF) and φ is the phase. Note that with no rate-of-change variable, the equation is that of a phasor.

We have implemented a synthetic "PMU" in MATLAB, measuring the PMU parameters by solving the fitting problem. (The entire method is in the digital domain: we do not create an analog signal.) With noise-free signals (such as used in calibration) the fit is excellent. That means that the model—equation (1)—is a good match for the signal.

In [7] we defined the Goodness of Fit metric (GoF) as the reciprocal value of the fit standard error, normalized. The calculation depends on the number of degrees of freedom of the equation. We now give the metric in dB to make the numbers more manageable. Expressed in decibels:

$$\mathrm{GoF} = 20 \log \frac{A}{\sqrt{\frac{1}{(N-m)} \sum_{k=1}^{N} (u_k - v_k)^2}} \tag{2}$$

where *N* is the number of samples, *m* is the number of parameters being estimated in the equation (one more than the number of degrees of freedom), *A* is the signal amplitude, u_k is the signal sample value and v_k is the estimated sample value. The parameter $(N - m)$ is called the residual degrees of freedom in [8]. When $N = m$, the metric blows up, but if the solver is doing its job, the fit is perfect.

A typical PMU has more than enough points to find the three or four parameters required of the PMU. There are devices available with as few as 24 samples per nominal cycle, and as many as 512 [9]. In the presence of noise and distortion, the fit will not be perfect, even if the results of the measurement are perfectly accurate.

The perfect-fit situation does not arise when the PMU is asked to characterize a non-phasor signal. This obviously creates residuals, and is the general case in the real world.

We thought it should be possible to use our "PMU" to find the values of the PMU parameters that minimize the residuals of the measurement across the transition region in Figure 1. The results of this measurement would at least be a plausible set of values for the transition regions of the signal. We wondered how bad the GoF could get during the transition and what the measured values would be.

It is worth noting that our implementation does not use any filtering. Each measurement window in our PMU is a rectangular window whose duration is set before the measurement. If we want a measurement that occupies just one cycle, only the data from that cycle will be used. For the present test, we set the length to two cycles, as that is the fastest typically used for a P-Class PMU [2].

7.6.3 Transitions

In essence, we "scrolled" the transition through our two-cycle window sample by sample. The results of our measurements are illustrated in Figure 3 and Figure 4, which shows the GoF and the value of ROCOF. The transition is from ROCOF = 0 to ROCOF = 3 Hertz per second. The measurement window consists of 2 cycles with 24 samples per (nominal) cycle. Note that the GoF graph does not include two measurements (on both sides) when the transition is not inside the window (0 samples into window) and all the way in the window (48 samples in the window), because then the GoF values are hundreds of dB, and the large vertical axis range would hide the GoF trend across the transition.

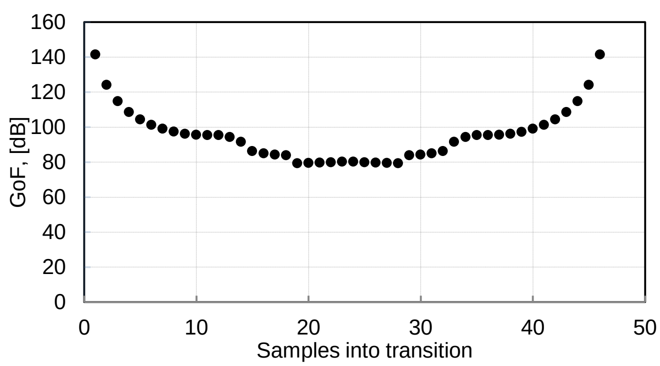

Figure 3 Results of GoF as transition moves across measurement window

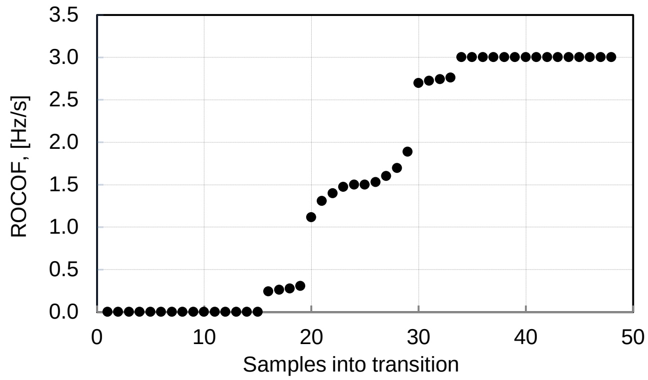

Figure 4 Results of ROCOF as transition moves across measurement window

The Goodness of Fit barely dips below 80 dB as the “corner” of the transition moves across the middle of the window. A value of 80 dB means that the fit, though not as good as when the corner is not inside the window, is still very good indeed. The residuals even at the “worst” fit point, are on the order of 0.01%, as shown in Figure 5. In the real world, that sort of performance is not achievable. The GoF for the measurement shown in Figure 2 is about 44 dB, for example.

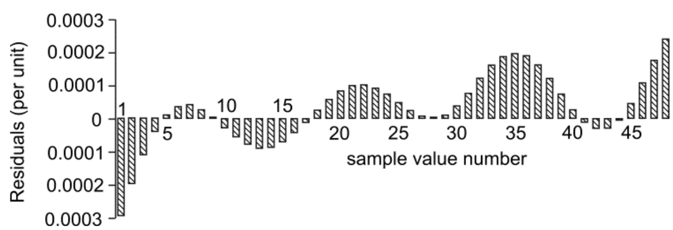

Figure 5 Residuals corresponding to window 20 in Figure 3

The reported ROCOF seems to have some very strange characteristics, with sudden changes in value. At the beginning of the transit, the reported ROCOF remains at zero longer than expected, and the final value (that is, 3) is reported before the transit is completed. The result of the measurement is odd and unconvincing, yet the GoF remains resolutely good and ROCOF graph looks at least plausible. What is going on?

7.6.4 Real transitions

Even the casual observer can see that Figures 3 and 4 are symmetrical. We hypothesized that because clean synthetic data was used, the “PMU” was able to minimize residuals, and the result is “accurate” by that definition. We decided to test the method on a real-world “dirty” signal.

We have already mentioned 345-kV system oscillography data. These are recordings that are triggered by a large system transient event, typically a fault. Note that the kind of transition created occupies a time period much larger than 2 system cycles (our chosen and also the PMUs measurement window), and likely contains not only a ROCOF change, but also amplitude changes and phase jumps. Further, real-world oscillography contains noise and harmonic distortion.

We scrolled some “fault” data through our measurement window. We used the “native” sampling rate (64 samples per second) of the oscillography. The GoF values reported are shown in Figure 6.

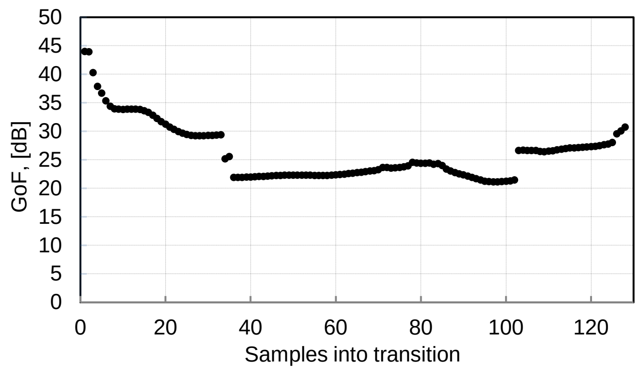

Figure 6 Results of GoF as real system event moves across measurement window

The minimum GoF for this transition is 21 dB. That is much worse than for our clean signal with a ROCOF step, yet the graphs have broadly similar shape. The real-world graph represents more than one change, and they are spread over a longer time. (It is worth mentioning that steady state GoF values for this oscillography are about 44 dB. This means that data contains a considerable amount of noise and distortion.)

The ROCOF values reported during this sequence are shown in Figure 7.

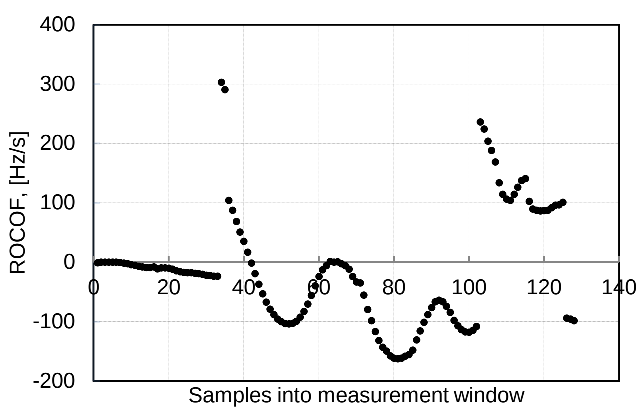

Figure 7 Reported ROCOF values during fault transition through measurement window

What Figure 3 and Figure 7 taken together show is that the value reported as ROCOF has only minor influence on the goodness of fit. While the values with our clean signal are plausible, they don’t “feel” right. The values from the real world are clearly nonsensical, and arise because our method of fitting allows the ROCOF parameter to achieve unrealistic values if that is what it takes to give a good fit. That is a problem with the method of fitting, but it highlights a deeper problem with the relative size of the “ROCOF signal.”

For us the next question was stimulated by the idea that the ROCOF signal, because of its role in equation (1), is very small: Can ROCOF even be measured in the real world? PMUs that meet the performance standards under test evidently struggle to measure ROCOF with real-world signals.

7.6.5 ROCOF Role in the Model

To see the effect of a ROCOF error, we can just consider the magnitude of the elements that make up the argument of the cosine in the "PMU equation" (1).

Let us use the value we used before, 3 Hz/s. Assume the ROCOF remains at this value until the end of the window. At the zero-time sample, the first term in the cosine argument is $60 \times 2\pi$. With a two-cycle window and a constant ROCOF all the way across, the frequency would be measured as 60.05 (60 plus half a thirtieth of 3). One might imagine that value to apply at the center of the window, if only it were possible to know a frequency at a single sample time [3]. The first term in the cosine argument is now $60.05 \times 2\pi$. The change due to a non-zero ROCOF is 50 mHz or 0.08%.

7.6.6 Discussion

Consider first how much the frequency changes in the real world. During the few seconds immediately following loss of a generator from the interconnected power system, the frequency might fall at a rate between 10 mHz/s and 100 mHz/s, depending on where it was, and the amount of lost generation [7].

What our calculation showed is that with ROCOF = 3 Hz/s, the cosine argument is within a tenth of a percent of the ROCOF = 0 value. ROCOF is just not making much change to the fit.

We think that statement is equivalent to saying that a ROCOF of 3 Hz/s is not measurable within the two-cycle PMU window, because of the distortion introduced by the transition.

The difficulty of ROCOF measurement is not caused by the need to differentiate one signal to obtain a value for another. Differentiation is known to be a noise-sensitive operation, but our solution method does not rely on differentiation of a frequency signal to obtain its rate of change. All the parameters in the equation are regarded as what Carnap termed "primitive" as opposed to derived [10].

The effect is caused by a rather moderate mismatch between the signal and the model used by the PMU. In terms of our solution method, we imagine a flat minimum in the multi-dimensional error function. In terms of other measurement methods, it would give the appearance of trying to measure a very small signal in the presence of some very large ones. In essence ROCOF is comparable to low level noise signal.

In short, the signal cannot be measured in this window except in the complete absence of disturbing factors.

We have seen that our solution method is able to measure very small values of ROCOF (say, 3 mHz/s or even 3 μHz/s) with excellent accuracy (say, ppb) if the signal is ideal. But with even the moderate mismatch of this transit, the ROCOF value becomes very approximate.

In the noise-free world of a calibration signal, our "PMU" can make the measurement of ROCOF. So could a commercial PMU. But in the noisy environment of the real world, noise masks the ROCOF.

We had hoped that we could find some plausible values for a ROCOF measurement transitioning across a window. It is not yet clear that we were successful.

7.6.7 Conclusions and Future Work

It has been shown that parameter transitions do not necessarily result in completely unmeasurable parameters. With clean signals, a plausible value is obtainable with the transition anywhere in the window. We note that further work has to be done before performance standards could be levied, and that seems to require a better understanding of the meaning of the measurement result.

The work shows, unfortunately, that even a poor ROCOF measurement may not be detectable in the GoF metric. Its contribution is just too small.

In the real world, with a power system that normally is not quite in steady state, the measurement of ROCOF also seems to be impractical within such a small measurement window.

We should not blame the PMU for its inadequacies. The fact is, the ROCOF signal is so small that it is swamped by real-world effects. Until we have a better understanding of the actual noise and distortion on the power system, we will be able to make little headway in making the measurement. We suggest that a program of noise study be started, with the ultimate goal of setting some reasonable performance requirements for the PMUs and defining them even through parameter transitions.

7.6.8 Acknowledgments

We thank Zak Campbell of AEP for furnishing oscillography data from the AEP power system. We also thank Raymond M. Hayes from ESTA International for many valuable discussions.

7.6.9 References

[1] IEEE C37.118.1, "IEEE Standard for Synchrophasor Measurements for Power Systems," IEEE, 2011.

[2] C37.118.1a, "IEEE Standard for Synchrophasor Measurements for Power Systems. Amendment 1: Modification of Selected Performance Requirements," IEEE, 2014.

[3] J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress," Proceedings of the IEEE,, vol. 66, no. 9, pp. 1048-1075, 1978.

[4] G. Stenbakken and M. Zhou, "Dynamic Phasor Measurement Unit Test System," in Proceedings IEEE PES General Meeting, Tampa, FL, 2007.

[5] H. Kirkham and J. Dagle, "Synchronous phasor-like measurements," in IEEE Conference on Innovative Smart Grid Technologies, Washington, DC, 2014.

[6] H. Kirkham, "A Conceptual Framework for Measurement, Report PNNL-24071," Pacific Northwest National Laboratory, Richland, WA, 2015.
[7] H. Kirkham and A. Riepnieks, "Dealing with Non-Stationary Signals: Definitions, Considerations and Practical Implications," in IEEE PES GM, Boston, MA, 2016.
[8] D. Cuthbert and F. Wood, Fitting Equations to Data, New York: John Wiley & Sons, Inc, 1971.
[9] von Meier, A; Culler, D; McEachern,A, and Arghandeh, R., "Micro-synchrophasors for distribution systems," in IEEE PES Innovative Smart Grid Technologies Conference, Washington, DC, 2014.
[10] R. Carnap, Philosophical Foundations of Physics, New York, NY: Basic Books, 1966.

7.7 Students' Simple Method for Determining the Parameters of an AC Signal

7.7.1 Introduction

The title of this paper is a tribute to the great William Edward Ayrton. On 23 November 1894, his paper "Students' Simple Apparatus for Determining the Mechanical Equivalent of Heat" was presented. The paper [1] was co-authored by his student Hermann Clark Haycraft. It showed how, using electrical methods, a quantity identified as the mechanical equivalent of heat could be measured in about ten minutes. That time contrasted with the years that scientists had spent (along with considerable effort, time—and money) making this measurement, with no greater accuracy.

Though the use of direct reading instruments, as described in the paper, was deprecated by some of the eminent physicists of the day [2], such methods did predominate in the end. The Ayrton-Haycraft paper marked the separation of modern electrical engineering from the antediluvian science then predominant in metrology.

The present paper is also written by a senior researcher and a student, and it is their hope that it marks a similarly profound moment in the science of measurement. The digital revolution in measurements [3] is an ongoing effort. Future developments will take the field further from merely imitating the kind of analog measurements Ayrton might have made, to a time of truly digital measurements whose results are firmly understood as the parameters of mathematical models.

7.7.2 Measurement of a phasor

The observer expects to get the same result no matter how a measurement is done. That should be a statement with complete generality. What takes place when the measurement is being made is that information carried by the forces and fields of the physical world are transferred to the conceptual world of numbers and mathematics. Keeping these worlds separate in our minds is crucial [4][5], and yet difficult. When the measurement is complete, we cannot help but think we have learned something about the physical world. In fact, what we have learned is the value of one or more parameters in the conceptual world of whatever mathematical model we are using to represent the quantity being measured. Our understanding of real world processes may still be incorrect.

In the measurement of alternating signals such as those that characterize the power delivery system, one way that the "usual" model becomes incorrect occurs if the frequency is changing. The signal is then no longer periodic. It is well known and accepted that periodic electrical signals can be described using a phasor equation and measured by a device known as a Phasor Measurement Unit, or PMU. The parameters that PMUs are asked to find are the amplitude X_{m}, the frequency ω and the phase φ, and something called the rate of change of frequency, or ROCOF. Apart from ROCOF, the measured parameters can be regarded as the variables on the right side of a continuous-time-domain equation such as (1)

$$x(t) = X_{\mathrm{m}} \cos(\omega t + \varphi). \tag{1}$$

If we allow for a frequency changing at a fixed rate (the simplest assumption), the model must be modified, and the equation becomes:

$$x(t) = X_{\mathrm{m}} \cos\left(\left[\omega + \frac{C_\omega}{2} t\right] t + \varphi\right). \tag{2}$$

where C_ω is the rate of change of frequency. The equation no longer represents a sinewave, but we still speak of measuring a phasor, which is a representation of a sinusoidal entity.

7.7.3 Implementation

The authors set themselves to demonstrate that measurement depends on a conceptual model, and once the model is established, the method of finding the parameters does not change the values of those parameters. The "simple method" of our title achieves its results by making measurements as the solution of a mathematical *fitting problem*. We assert that Equation (1) is the model, and we fit the model parameters to samples of the signal in the discrete-time domain.

For our demonstration, there is thus a two-part problem. First, we have to create a set of values that represent the signal we wish to characterize, and second, we have to make the measurements. The form of equation (1) is fixed by the physics, so this is not an equation-fitting method so much as a parameter fitting. Some preliminary results [6] showed promise. This paper gives details of the method and some results.

7.7.3.1 Fitting Solution

It is seemingly new to metrology to regard measurement as what mathematicians call a fitting problem, and to make the measurement by curve-fitting, but it seems quite appropriate. We have shown in [7] [8] and [9] that when the measured parameter values declared by a PMU are inserted into equation (1) along with the times, the sample-values can be reconstructed. These devices are therefore solving the equation. In the commercial PMU each parameter is measured separately: fitting just measures them all together. The measurement then becomes one of adjusting the parameters of equation (1) and fitting the result to samples of the waveform whose parameters are needed.

In a real PMU, the samples arrive from an A/D converter at a rate set by the designer, triggered by a clock linked to UTC. For our purposes, not requiring a real-time solution, we used files of pre-recorded data, and wrote a MATLAB software script to perform the parameter fitting. We were able to select files of data from the real world (samples known as *oscillography* in honor of a very old technology for recording such signals), or files of our "synthetic" data created by MATLAB or an Excel spreadsheet.

Since our purpose was to explore the method, we required flexibility. The software allowed for the variation of all the parameters in equation (1), including the amplitude, something that the PMU is not required to do. It allowed for changes in the sample rate, and allowed for the addition of noise and distortion to the signal. Further, in an endeavor to explore the limits of performance of the system, we allowed for changes in the window width.

Fitting is not a new idea (except as applied to measurements), so we selected and used an already available method. Our choice was to use a least squares minimization technique that searches for the minimum of summed square residuals. The residual r_i is defined as

$$r_i = y_i - Y_i \,. \tag{3}$$

where y_i is the observed value and Y_i, is the fitted estimate of i_{th} sample. The difference then is identified as the residual associated with the data.

The summed square of residuals is given by

$$S = \sum_{i=1}^{n} (y_i - Y_i)^2. \tag{4}$$

Linear fitting coefficients for equation (1) cannot be found by simple matrix techniques because the cosine function is nonlinear. Therefore we used an iterative approach.

We also assumed that the residuals are normally distributed. That might not always be a good assumption since the method would have to deal with real data, and outliers do sometimes occur in real data. Therefore, we used a robust least-squares regression with *bisquare weights* in order to minimize outlier influence (squaring an extreme value gives very large errors). In essence, it means that the further the data point is from fitted line the less weight it gets. Data points that are outside estimated random chance region receives zero weight [10].

The least squares algorithm follows MATLAB specified procedure and can be found in their documentation [10]. We will not go in the details of the fitting algorithm itself because that can be easily found online.

1) Start with previously chosen initial coefficients (start values).

2) Construct the curve. The fitted curve Y is given by

$$Y = f(X, b) \tag{5}$$

   and includes the Jacobian of $f(X, b)$, which is a matrix of partial derivatives taken with respect to the coefficients..

3) Fit the model by weighted least squares.

4) Adjust the residuals and standardize them.

$$r_{adj} = \frac{r_i}{\sqrt{1 - h_i}} \tag{6}$$

   where h_i is leverage for reducing the residual weights.

   The adjusted standardized residuals are

$$u = \frac{r_{adj}}{Ks} \tag{7}$$

where K is a tuning constant and s is the "robust variance" [10].

5) Calculate the robust weights as a function of u. The bisquare weights then are defined as

$$w_i = \begin{cases} (1-(u_i)^2)^2 & |u_i| < 1 \\ 0 & |u_i| \geq 1 \end{cases}. \quad (8)$$

The final weight is the product of the regression weight and the robust weight.

6) Finish, if the fit has converged. If not, perform another iteration of the fitting procedure.

7) Adjust the coefficients and determine whether the fit improves. The direction and magnitude of the adjustment depend on the trusted region. This is the MATLAB default algorithm and is used because coefficient constraints can be specified. Iterate the process by returning to step 3 until the fit reaches the specified convergence criteria [10].

7.7.3.2 Noise Addition

In [6] the kind of "engineering noise" considered by Shannon [11] is discussed. Shannon's work showed that messages could be encoded in such a way that they could be transmitted without error even via a noisy channel. One might imagine that this noise was the sort of "hash" you might hear on a radio channel, or something responsible for changing the occasional zero to a one on a digital link. In [12] the authors consider (very briefly) noise as a way of changing the meaning of the message. In [6] that is termed "semantic coloration" to distinguish the process from what [12] calls "engineering noise." These terms can be clarified by considering equation (1).

If the signal that (1) represents is considered as a stream of digital values $x(t)$, and the digital values are corrupted by random changes in the values, that is the sort of noise that Shannon considered. The effect of such noise is a problem solved by the methods developed after Shannon.

However, if the noise is added to the terms on the right side of the equation, we are assuming that some sort of process in the physical world is *changing the way we should model the signal*. The matter is considered with respect to oscillators by modifying (1) to produce the following equation:

$$x(t) = [X_\mathrm{m} + \epsilon(t)]\cos(\omega t + \varphi + \Phi(t)). \quad (9)$$

where $\epsilon(t)$ is a random amplitude noise and $\Phi(t)$ is a random phase noise. (This is similar to the notation that Jacques Rutman used in [13]. The interested reader is urged to study that paper for further information on the stationarity assumption of these random processes.)

The addition of noise in our signal generation process can therefore be done controllably and separately to the amplitude and the phase. We could also modify the frequency, recalling that any change in the frequency will also change the phase. It is therefore interesting to see what the measurement result shows about such a situation.

7.7.3.3 Sequence of Operation

The sequence of operations was as follows. While the IEEE standard [14] divides PMUs into two classes, essentially based on their response speed, which depends on the window width and filter characteristics, we allowed for narrow windows of observation (less than a cycle) as well as multi-cycle windows, but added no filtering. The parameters were fitted using only data within each window, that is, we used a rectangular window. This was the simplest possible window, and it enabled us to make truly independent measurements. Spectral leakage generally rules out the window for methods based on the Fourier transform, but our method is a time-domain method, not a frequency-domain one.

Our process was:

1. Define the case or cases to be run: what signal model, how many samples per nominal cycle, what, where in time and how often were changes in parameters to be made, what noises or harmonics to add.
2. Set these numbers into the software to create the "PMU input signal."
3. Run the solution software against the input signal to create the output in a file or files. These contained the results of a string of measurements. We also stored the measurement residuals and the solver parameters.
4. Examine the output data, and generate graphs of the more interesting results.

If noise was added to the signal, the results had a statistical nature, and a proper understanding of the results required multiple "measurements." Sometimes a meaningful statistical analysis required several hundred or even a thousand cases be run.

7.7.4 Input data

7.7.4.1 Excel

It is a straightforward matter to create a spreadsheet that contains all the values of hypothetical samples of a signal, with all the parameters in Equation (1) or (2) known. These values can then be used as input to the fitting program.

It is possible to change a parameter part way through a sequence, producing a step change in a parameter value. Such a step input is often used as a test of filter response, but in our implementation we have used no filtering. It is therefore interesting to observe a step-response in the output of the fitting solution. If the input step-change occurs at the start of a window, the full step is reported immediately, otherwise some sort of intermediate value is reported for one window.

It is possible to implement more than one step in the signal, though the method exposes a problem when a second step is applied to the rate of change of frequency. Since the phase of the signal is a path-dependent integral of the whole argument of the cosine, the spreadsheet phase value will not ordinarily be correct at the second step. We have come to call this the "van der Pol problem" after the Dutch engineer who wrote about it with respect to the mathematics of representing frequency modulation [15]. If the phase is not corrected, there will be a phase jump in the signal. (That is something that can occur in the electric power system, but it is beyond the scope of this paper to explore. As the text books say “it is left to the student to show . . . ”)

7.7.4.2 MATLAB

The capability to specify phasor or “phasor-like”[16] parameters allowed us to create steady sine waves or waves with steady ramping of frequency and/or amplitude. Later versions of the code allowed the addition of harmonic distortion with unlimited harmonic count. For the harmonics, amplitudes and phase shifts were separately controllable. A way to generate normal distribution white noise and red noise (Brownian noise) was implemented.

Creating input data in MATLAB allows for faster code execution and more convenience compared to data import from an Excel file, which was how we began the effort. However, data import from a file was the only way to import real-world oscillography data into our MATLAB system. We used files with *.xlsx* or *.csv* extensions to import the data. A conversion from COMTRADE to an Excel-compatible format was done for us by Ray Hayes of ESTA International.

MATLAB supports importing data from files, and all our real-world data and much of our synthetic data was imported from Excel files. The artificial data was created as a large array of sine-wave data points and Δt values. Spreadsheet data generation allows not only for steady state signal generation but also multiple parameter value changes in a single data stream, things like phase jumps, and amplitude jumps. A spreadsheet was also used in harmonics and noise (filtered and raw white noise) generation. Generating data this way, the MATLAB script took more time to run, and changes in the data were not easy to make. On the other hand, unique changes could be made. For some experiments this was the only way to generate the desired input signal.

Random noise could be added to any of the parameters to create more realistic versions of the signal. The noise is characterized by its amplitude distribution (an example is shown in Figure 1) and a filter characteristic. “White noise” can be filtered in MATLAB to produce amplitudes that fall of in various accepted ways to produce pink noise or red noise.

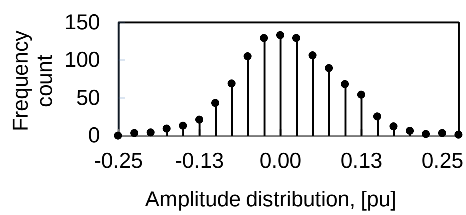

Figure 1. White noise amplitude distribution

Figure 2 shows how a realistic signal with harmonics and noise can be built up in Excel. In the example, a small amount of random noise is added to the amplitude signal (as in Equation (9)). Fifth and third harmonics are also added.

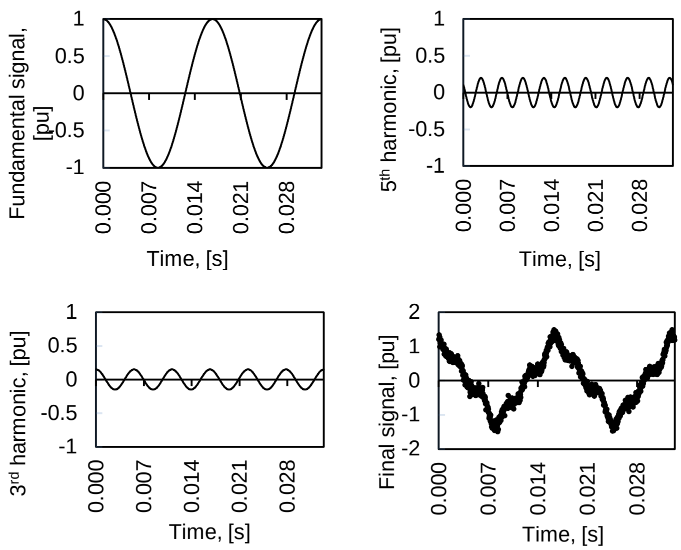

Figure 2. Signal construction in Excel.

The “pollution” in the signal of Figure 1 is a mixture of harmonics (which are obliged to track the fundamental frequency when it changes) and random noise. These are at least stationary processes, with constant parameters. The real power system is not always well-characterized by such signals.

Figure 3 shows a short section of a data stream (oscillography) from a 345-kV power system in the US. The signal shows evidence of something quite non-stationary: a fault. The relay doing the data collection was sampling at 64 samples per nominal cycle. We will discuss the results of measuring this below.

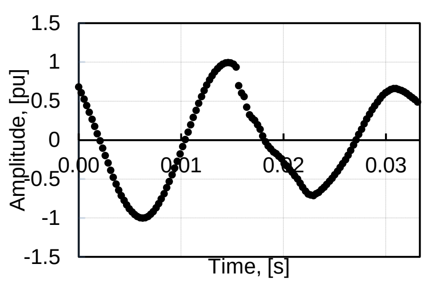

Figure 3. Real-world data during system fault.

7.7.5 Output data: Results

Output data can be any variable created or imported into the MATLAB code. We chose to use all phasor and phasor-like quantities as outputs together with GoF metric, calculated from residual RMS values. Usually we output arrays of results so that they can be exported to MS Excel for graphical representation. (This is a personal preference. MATLAB will allow graphs, but our choice was the interface and appearance provided by Excel.)

7.7.5.1 Synthetic data

For example, we have simulated a signal with a smooth increase in frequency, using clean (noise-free) signals. Our measurement system was configured to represent a PMU with a measurement window two cycles long. The PMU is asked to report a changing frequency, and our fitting solution gave a value of 0.199999999996 Hz s^{-1}. The correct value was 0.2 Hz s^{-1}. In other words, the "error" appears in the twelfth significant figure. For a clean signal such as this, that level of error in the result shown in Figure 4 is typical.

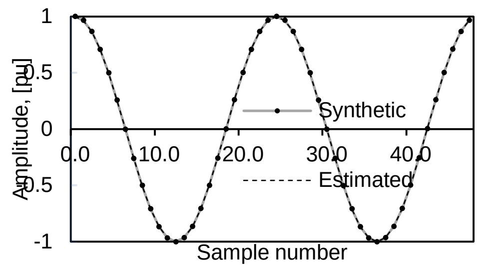

Figure 4. Measuring a wave of changing frequency

However good it is, it is unrealistic to expect such performance in the field. We show elsewhere [7] that in the noisy environment of the power system, the rate of change of frequency that will typically be seen in the power system (the value above is representative) is still immeasurably small. The ROCOF signal is, simply put, drowned out by noise.

To have any hope of understanding the limits of performance of a device such as a PMU, a knowledge of the noise is essential. At this point, there is (as far as we know) no information available on power system noise that would be of value here. For the remainder of this paper, we will omit consideration of ROCOF.

In a power system, a phase jump may occur because of switching operations. We simulated that situation by changing the phase in the middle of a two-cycle measurement window, as shown in Figure 5.

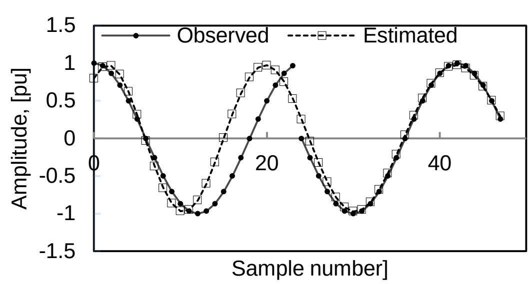

Fig.5. Input signal with 90^0 phase jump.

The fitting solution is also shown in the Figure. The signal is *missing* a 90-degree section of the cosine wave. The frequency that gave the best fit was higher than the nominal frequency. The result (about 87 Hz, as seen in Fig 4) is not meaningful *considered as a frequency*: but the question really is, what is the meaning of the word frequency for the waveform of Fig 5? There is no simple answer to that question.

Figure 6 demonstrates the ability of the method to give abrupt changes in the declared values. What is reported as the phase (right axis) jumps from zero to a negative number before it goes positive, and the frequency has a single high value that is clearly anomalous.

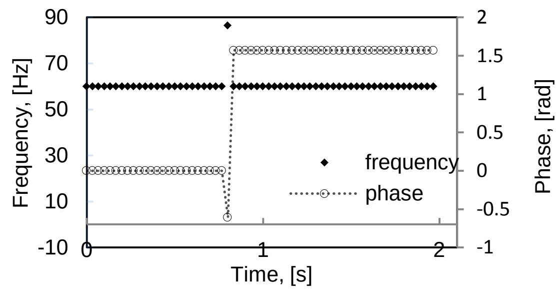

Fig.6 Method output to 90^0 phase jump introduced in input signal.

7.7.5.2 Oscillography data

Fig 7 shows the oscillography sample values of Fig. 3 and the values estimated by our fitting method, calculated for each sample instant.

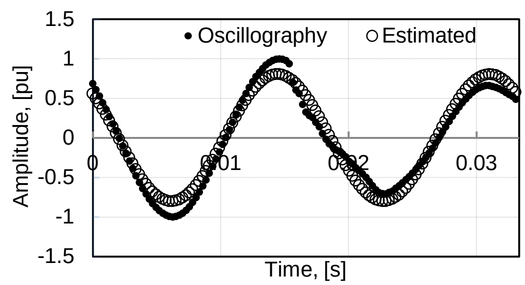

Figure 7. Estimated curve for real-world fault data

We have seen that with a clean signal and a fully representative model, the residuals are very small. That is obviously not the case in this example. Put another way, the mathematical model does not do a great job of representing physical reality. The PMU is obliged to use this model, however, even if it is a poor representation of the signal. Werner Heisenberg famously said [17]

> . . . since the measuring device has been constructed by the observer, we have to remember that what we observe is not nature in itself, but nature exposed to our method of questioning.

In the case of the PMU, and possibly all digital measurements, the “method of questioning” is fixed by the equation we are solving. The PMU did not ask about changing amplitude. The observer is therefore cautioned to remember that the declared value of the measurement is the answer to a very specific question, but the question may not have been the best one to ask.

In Table 1 we show what a PMU reported from the same signal, and what our fitting method estimated over the fault period (as shown in Figure 5), along with the next two 2-cycle measurement windows. We think that the differences are because the fitted values are truly independent and no filtering is used, and the PMU values show some effects of filtering.

TABLE I
PMU reported value comparison to estimated values

PMU			Estimator		
Amplitude [pu]	Frequency [Hz]	Phase [°]	Amplitude [pu]	Frequency [Hz]	Phase [°]
0.809	59.961	45.264	0.802	60.291	46.159
0.826	59.889	43.081	0.835	59.673	45.738
1.017	60.069	43.05	1.029	60.169	44.482

7.7.5.3 Code Sample

Below is an example of the code for a single measurement. For different experiments additional loops and conditions were used that implemented more complex code. Given here is the core code for calling for the MATLAB solver and stating settings, as well as an example of filter setup.

```
clear
clc
%% Read data from file;
%64 samples per cycle
y = xlsread('name_of_file.xlsx',1,'AB100:AB227'); %Import data from
file
T = xlsread('name_of_file.xlsx',1,'T100:T227');
%% Measurement process
[xData, yData] = prepareCurveData( T, y );      %preparing data for
curve fitting
% Set up fittype and options.
ft = fittype( 'a*cos(2*pi*c*x+2*pi*(d/2)*x*x+e)', 'independent', 'x',
'dependent', 'y' );     %Fitting equation
opts = fitoptions( 'Method', 'NonlinearLeastSquares' ); %Set method
opts.DiffMaxChange = 0.0001;                            %Max step change
opts.Display = 'Off';                                   %Disable display
option
opts.Lower = [0 -50 55 -5 -3.14159265358];              %Lower trust
region boundaries
opts.MaxFunEvals = 1000;                                %Maximum
evaluations allowed
opts.MaxIter = 1000;                                    %Maximum
iterations
opts.Robust = 'Bisquare';                               %Select bisquare
robust fitting
opts.StartPoint = [1 0 60 0 0];                         %Start values
opts.TolFun = 1e-8;                                     %Termination
tolerance for the function
opts.TolX = 1e-8;                                       %Termination
tolerance for x
opts.Upper = [1.5 50 65 5 3.14159265359];               %Maximum trust
region values
%Call for MATLAB solver for curve fitting with selected options and
outputs
```

```
[fitresult, gof, fitinfo] = fit( xData, yData, ft, opts );
%OUTPUT
RMS = 20*log10(1/gof.rmse); %Calculated GoF values
f = fitresult.c;            %Frequency values
A = fitresult.a;            %Amplitude values
ph = fitresult.e;           %Phase values
C_w = fitresult.d;          %ROCOF values
%% write data to file
filename = 'name_of_file.xlsx';
xlswrite(filename,A',1,'A1')
xlswrite(filename,f',1,'B1')

%%FILTERING
%Import the signal from file
test=xlsread('name_of_file.xlsx',1,'A1:A30720');
Fs = 30720;
%Sampling Frequency (samples per second)
%%Butterworth Lowpass filter designed using FDESIGN.LOWPASS
fpass = 3;          %Passband Frequency
Fstop = 100;        %Stopband Frequency
Apass = 1;          %Passband Ripple (dB)
Astop = 6;          %Stopband Attenuation (dB)
match = 'stopband';
%Band to match exactly
%%Construct an FDESIGN object and call its BUTTER method.
h  = fdesign.lowpass(Fpass, Fstop, Apass, Astop, Fs);
Hd = design(h, 'butter', 'MatchExactly', match);

ttt = filter(Hd,test);   %Filtering signal
```

7.7.6 Discussion

The 1894 paper by Ayrton and Haycraft was quite controversial. That may have been the intent of the authors. The paper compared the result of a ten-minute experiment by some students with the results of years of elaborate work by established scientists. Objections were raised about the use of direct-reading instruments, and about the avoidance of calibrations and "corrections." A few days later, *The Electrical Review* pointed out [2] that

> To our mind, the points which are here objected to are the very ones which make it of greatest educational value for experimental work. The object of the apparatus is to give young students a clear, *concrete* idea, of the mechanical equivalent of heat . . .

We have a somewhat similar object with this paper. Our goal, like Ayrtyon's, is to give the student the clear and concrete idea. These days that comes from implementations in MATLAB, a convenience that avoids the need of wires and transformers, just as the implementation of Ayrton and Haycraft avoided the need for "a lesson in calibration, or in the principle of the tangent galvanometer."

Unlike Ayrton, however, we are not proposing a new method of measurement. Although we have developed a new method of measurement, and although we think it to be capable of giving better results than anything presently available, it has yet to be shown that it can work in real time, and many measurement people are not yet "comfortable" with fitting as a measurement method.

The clear and concrete ideas that we communicate are that

1. the performance of a measurement system is limited by noise in the system and on the signal being measured
2. a flexible measurement system (such as this) can be used to explore the limits of performance
3. the meaning of the measurement result is not necessarily obvious.

Measurement is a process that uses information in the physical world to produce information in the conceptual world. But this world is bounded by human preconceptions that the observer may not even be aware of. These preconceptions lie behind the very design of the instrument. Remember Heisenberg!

Forgotten conventions are exposed by the fitting method, too. The usual definition of frequency, for example, is the derivative of phase. For power engineers, that does not work for the PMU with the usual definitions, because "phase" to a power engineer is the φ in equations (1) and (2). In the PMU this phase is with respect to a hypothetical reference at exactly the nominal power frequency. A signal with a constant phase is thus at exactly the power frequency, and yet the derivative of "phase" is zero. For the derivative relationship to hold, the phase has to be the whole of the cosine argument in equations (1) or (2). Power engineers have mostly forgotten that their conventional use of the word "phase" means just the phase of the stationary phasor.

7.7.7 Conclusion

The "experiment" of making a measurement by curve-fitting brings home to the student that the act of measuring is one that can be done in various ways, but the end result should not depend on the method selected. Most importantly, it teaches that measurement is the act of using signals from the real world to find parameters of a model. That model is almost always a simplification.

As Rutman [5] observed:

> In an ideal world, the correspondence between the physical reality and the model should be perfect. In fact, the real world is so complex that many details are ignored in the model and therefore, the correspondence is incomplete and imprecise. On the other hand, some properties that have no direct meaningful counterparts in the real world have to be introduced in the model in order to make a tractable model (eg : stationarity properties of random processes). In any case, the model must be well defined from the mathematical point of view.

And so we can see that having stationary values for the phasor parameters is a practical requirement that the physical world need not accede to. Nevertheless, any PMU finds the values of equation (1), and reports them as the (stationary) results within its observation window. It is up to the observer to understand the significance of those parameters and of the models they are part of.

Measurement continues to move in the direction established by Ayrton. He showed that understanding was improved by setting aside some of the details of the instrumentality. As the digital revolution in measurements progresses, that setting aside will be increasingly straightforward. We are left with the conceptual problem that has been there all along: what does the result of the measurement actually *mean?*

7.7.8 Acknowledgments

We thank Zak Campbell of AEP for furnishing oscillography data from the AEP power system. We also thank Ray Hayes of ESTA International, and late of AEP, for many valuable discussions.

The first author of this paper would like to say a special “thank you” to Jerry Stenbakken, late of NIST, who used the phrase “digital revolution in measurements” in a conversation, and thereby set much in motion.

7.7.9 References

[1] W. Ayrton and H. Haycraft, "Students' Simple Apparatus for Determnining the Mechanical Equivalent of Heat," *Proceedings of the Physical Society,* vol. 13, pp. 295-309, January 1895.

[2] anonymous, "Determination v Experiment," *The Electrical Review,* p. 9, Jamuary 1895.

[3] H. Kirkham, "The Digital Revolution in Measurements," in *IEEE Conf on Innovative Smart Grid Technologies (ISGT),* Washington, DC, 2103.

[4] R. Carnap, Philosophical Foundations of Physics, New York: Basic Books, 1966.

[5] J. Rutman, "Oscillator Specifications: A Review of Classical and New Ideas," in *31 st Annual Symposium on Frequency Control*, 1977.

[6] H. Kirkham, "A Conceptual Framework for Measurement, Report PNNL-24071," Pacific Northwest National Laboratory, Richland, WA, 2015.

[7] A. Riepnieks and H. Kirkham, "Phasor Measurement as a Fitting Problem," *IEEE Transactions on the Smart Grid*, 2016.

[8] H. Kirkham and A. Riepnieks, "Dealing with Non-Stationary Signals: Definitions, Considerations and Practical Implications," in *Proceedings IEEE Power and Energy Society General Meeting*, Boston, MA, 2016.

[9] A. Riepnieks and H. Kirkham, "An Introduction to Goodness of Fit for PMU Parameter Estimation," *Transactions on the Smart Grid*, 2016.

[10] The MathWorks Inc., "Documentation," 2015. [Online]. Available: http://www.mathworks.com/help/matlab/.

[11] C. Shannon, "A Mathematical Theory of Communication," *The Bell System Technical Journal,* vol. 27, pp. 379-423, 623-656, July, October 1948.

[12] C. Shannon and W. Weaver, The Mathematical Theory of Communication, Illini Books, 1963.

[13] J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress," *Proceeding of the IEEE,* vol. 66, no. 9, pp. 1048-1075, 1978.

[14] *IEEE Standard for Synchrophasor Measurements for Power Systems, C37-118.1-2011,* pp. 1-61, Dec, 2011.

[15] B. van der Pol, "The Fundamental Principles of Frequency Modulation," *Journal of the Institute of*

Electrical Engineers, vol. 93, no. 23, pp. 153-158, 1946.

[16] H. Kirkham, J. Dagle, "Synchronous phasor-like measurements," IEEE PES Inovative Smart Grid Technologies Conference (ISGT), Washington, DC, Feb, 2014.

[17] W. Heisenberg, Physics and Philosophy: The Revolution in Modern Science, London: George Allen and Unwin, 1959.

8.0 Bibliography

anonymous. (1895, Jamuary). Determination v Experiment. The Electrical Review, p. 9.

Ayrton, W., & Haycraft, H. (1895, January). Students' Simple Apparatus for Determnining the Mechanical Equivalent of Heat. Proceedings of the Physical Society, 13, 295-309. doi:10.1080/14786449508620699

Boashash, B. (1992, April). Estimating and Interpreting the Instantaneous Frequency of a Signal--Part 1: Fundamentals. Proceedings of IEEE , 80(4), 520-538.

Boashash, B. (1992, April). Estimating and Interpreting the Instantaneous Frequency of a Signal--Part 2: Algorithms and Applications. Proceedings of IEEE, 80(4), 540-568.

Carey Foster, G. (1881). Inaugural address of the President. Society of Telegraph Engineers and Electricians.

Carnap, R. (1966). Philospohical Foundations of Physics. New York, NY: Basic Books.

Carson, J., & Fry, T. (1937). Variable frequency electric circuit theory with application to the theory of frequency-modulation. Bell System Technical Journal, 16(4), 513-540. doi:10.1002/j.1538-7305.1937.tb00766.x

Choi, S., & Meliopoulos, A. (2016). Effective Real-time Operation and Protection Scheme of Microgrids Using Distributed Dynamic State Estimation. IEEE Transactions on Power Delivery. doi:10.1109/TPWRD.2016.2580638,

Clement, P., & Johnson, W. (1960). Electrical Engineerng Science. New York, NY: McGraw-Hill.

Cuthbert, D., & Wood, F. (1971). Fitting Equations to Data. New York: John Wiley & Sons, Inc.

Eriksson, K., Sasksvik, O., Holmgren, T., & Mitsch, K. (2014). Celilo Pacific DC Intertie Upgrade presentation. CIGRE meeting (pp. 1-8). Paris: CIGRE.

Fleming, J. (1891, May 7). On some effects of alternating-current flow in circuits having capacity and self-induction. Journal IEE, 20(94), 362 - 408.

Fleming, J. (1892, May). Experimental Researches on Alternating Current Transformers. Journal IEE, XXI, 594-686.

Golding, E. (1938). Electrical Measurements and Measuring Instruments. London: Pitman and Sons.

Harris, F. (1962). Electrical Measurements. New York, NY: Wiley and Sons.

Heisenberg, W. (1959). Physics and Philosophy: The Revolution in Modern Science. London: George Allen and Unwin. Retrieved 10 28, 2015, from https://archive.org/stream/PhysicsPhilosophy/Heisenberg-PhysicsPhilosophy_djvu.txt

IEEE Standard for Synchrophasor Measurements for Power Systems, C37-118.1-2011. (2011, Dec). IEEE Std.C37.118.1-2011 (Revision of IEEE Std.C37.118-2005). pp. 1-61.

IEEE Standard for Synchrophasor Measurements for Power Systems. Amendment 1: Modification of Selected Performance Requirements. C37.118.1a. (2014).

De La Ree, J., Centeno, V., Thorp, J.S., Phadke, A.G. (2010). Synchronized phasor measurement applications in power system. IEEE Transactions on smart grid vol 1, No.1, 20-27.

Joint Committee for Guides in Metrology, WG 1. (2008). GUM: Guide to the expression of uncertainty in measurement.

Kelvin, L. (1889). Electrical Units of Measurement (May 3, 1883). In Kelvin, Popular Lectures and Addresses (pp. 79-136). London: Macmillan. Retrieved from https://archive.org/stream/popularlecturesa01kelvuoft#page/n7/mode/2up accessed May 18, Kirkham, H. (2103). The Digital Revolution in Measurements. IEEE Conf on Innovative Smart Grid Technologies (ISGT). Washington, DC: IEEE. doi:10.1109/ISGT.2013.6497875

Kirkham, H; Dagle. J. (Feb, 2014). Synchronous phasor-like measurements. Washington, DC: IEEE PES Inovative Smart Grid Technologies Conference (ISGT). doi:10.1109/ISGT.2014.6816510

Kirkham, H. (2015). A Conceptual Framework for Measurement, Report PNNL-24071. Richland, WA: Pacific Northwest National Laboratory. Available from https://certs.lbl.gov/publications/conceptual-framework-measurement

Kirkham, H. (2016). The Measurand: the problem of Frequency. Proceedings of the International Instrumentation and Measurements Conference, I&M Society. Taipei: IEEE.

Kirkham, H., & Riepnieks, A. (2016). Dealing with Non-Stationary Signals: Definitions, Considerations and Practical Implications. IEEE PES GM. Boston, MA.

Kirkup, L., & Frenkel, R. (2006). An Introduction to Uncertainty in Measurement. Cambridge, UK: Cambridge University Press.

Maxwell, J. (1877). Report on the Cavendish Laboratory for 1876. Cambridge University Reporter (p. 10). Cambridge: Cambridge University Press.

Metrology, W. G. (2008). GUM: Evaluation of measurement data--Guide to the expression of uncertainty in measurement. Paris: BIPM.

NCSL International 174 Writing Committee (ASC Z540). (2006). ANSIINCSL Z540.3-2006: Requirements for the Calibration of Measuring and Test Equipment. American National Standards Institute.

Nelson, T. (2011). Definitions for calculations of VA, VAh, VAR, and VARh for Poly-Phase Electricity Meters, NEMA C12.24 TR-2011. NEMA.

Putnam. (1973). Meaning and Reference. Journal of Philosophy, 699-711.

Putnam, H. (1981). Reason, Truth and History. Cambridge: Cambridge University Press.

Riepnieks, A., & Kirkham, H. (2016). An Introduction to Goodness of Fit for PMU Parameter Estimation. *Submitted to* Transactions on the Smart Grid.

Rutman, J. (1977). Oscillator Specifications: A Review of Classical and New Ideas. 31 st Annual Symposium on Frequency Control (pp. 291-301). IEEE. doi:10.1109/FREQ.1977.200163

Rutman, J. (1978). Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress. Proceedings of the IEEE,, 66(9), 1048-1075.

Rutman, J., & Walls, F. (1991). Characterization of frequency stability in precision frequency sources. Proceedings of the IEEE, 79(7), 952-960. doi:10.1109/5.84972

Shannon, C. (1948, July, October). A Mathematical Theory of Communication. The Bell System Technical Journal, 27, 379–423, 623–656. Retrieved from http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf

Shannon, C. (1956). The Bandwagon. IRE Transactions on Information Theory, 2(1), 3. doi:10.1109/TIT.1956.1056774

Shannon, C., & Weaver, W. (1963). The Mathematical Theory of Communication. Illini Books.

Simpson, J. (1981). Foundations of Metrology. Journal of Research of the National Burea of Standards, 86(3), 281-292.

Stenbakken, G. & Zhou. (2007). Dynamic Phasor Measurement Unit Test System. Proceedings IEEE PES General Meeting. Tampa, FL.

The MathWorks Inc. (2015). Documentation. Retrieved from MATLAB Help: http://www.mathworks.com/help/matlab/

University of California Berkeley. (2015). Micro-Synchrophasoprs for Distribution Systems. (Advanced Research Projects Agency - Energy) Retrieved from http://arpa-e.energy.gov/?q=slick-sheet-project/measuring-phase-angle-change-power-lines

van der Pol, B. (1946). The Fundamental Principles of Frequency Modulation. Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering, 153-158. doi:10.1049/ji-3-2.1946.0024

von Meier, A;Culler, D; McEachern,A Arghandeh, R. (2014). Micro-synchrophasors for distribution systems. IEEE PES Innovative Smart Grid Technologies Conference. Washington, DC.