

SANDIA REPORT

SAND2014-19137

Unlimited Release

Printed October 2014

Sampling and Filtering in Photovoltaic System Performance Monitoring

Anton Driesse, Joshua S. Stein, Daniel Riley, Craig Carmignani

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2014-19137
Unlimited Release
Printed October 2014

Sampling and Filtering in Photovoltaic System Performance Monitoring

Anton Driesse
PV Performance Labs
Canada / Germany
www.pvperformancelabs.com

Joshua S. Stein, Daniel Riley, Craig Carmignani
Photovoltaic and Distributed Systems Integration
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS1033

Abstract

The performance of photovoltaic systems must be monitored accurately to ensure profitable long-term operation. The most important signals to be measured—irradiance and temperature, as well as power, current and voltage on both DC and AC sides of the system—contain rapid fluctuations that are not observable by typical monitoring systems. Nevertheless these fluctuations can affect the accuracy of the data that are stored. This report closely examines the main signals in one operating PV system, which were recorded at 2000 samples per second. It analyzes the characteristics and causes of the rapid fluctuations that are found, such as line-frequency harmonics, perturbations from anti-islanding detection, MPPT searching action and others. The operation of PV monitoring systems is then simulated using a wide range of sampling intervals, archive intervals and filtering options to assess how these factors influence data accuracy. Finally several potential sources of error are discussed with real-world examples.

ACKNOWLEDGMENTS

This work was funded by the U.S. Department of Energy SunShot Program.

CONTENTS

1. Introduction.....	11
2. Background.....	12
2.1 The Measurement Chain.....	13
2.1.1 Physical to Electrical Conversion.....	14
2.1.2 Analog Processing.....	14
2.1.3 Analog to Digital Conversion.....	14
2.1.4 Digital Processing.....	15
2.2 Measurement Chain Parameters.....	15
3. Methodology.....	16
3.1 Test Cases.....	17
3.2 Evaluation.....	19
4. Hardware and Software Description.....	21
4.1 PV System.....	21
4.2 Data Acquisition.....	22
4.3 Calculation Platform.....	23
5. Signal Characteristics.....	25
5.1 Irradiance.....	26
5.2 Temperature.....	31
5.3 DC Current, Voltage and Power under Stable Conditions.....	35
5.3.1 DC Current.....	35
5.3.2 DC Voltage.....	38
5.3.3 DC Power.....	41
5.4 AC Current, Voltage and Power under Stable Conditions.....	44
5.4.1 AC Voltage.....	45
5.4.2 AC Current.....	48
5.4.3 AC Power.....	51
5.5 DC and AC signals with Rapid Irradiance Changes.....	54
5.6 Summary of Signal Characteristics.....	58
6. Simulated Measurement Chains.....	59
6.1 Simulated Measurement of Mean AC Power.....	59
7. Real-World Observations.....	68
7.1 Line Frequency Noise.....	68

7.2	Campbell Scientific Line Frequency Noise Rejection.....	69
7.3	Anti-islanding Perturbations	71
8.	Conclusions.....	75
9.	References.....	76
10.	Distribution	78

FIGURES

Figure 1 Performance data showing inconsistent trends in irradiance and power.....	13
Figure 2 Frequency response of the filter types that were used.....	18
Figure 3 Step response of the filter types that were used	19
Figure 4 The PV system used for the study	21
Figure 5 Three irradiance sensors used for the study	22
Figure 6 Stable irradiance signals from three sensors	26
Figure 7 Simultaneous disturbance of three irradiance signals	27
Figure 8 Average spectrum for the irradiance signal from the thermopile.....	27
Figure 9 Average spectrum for the irradiance signal from the photodiode	28
Figure 10 Average spectrum for the irradiance signal from the reference cell	28
Figure 11 Irradiance measurements with line frequency ripple removed.....	29
Figure 12 Irradiance measurements with moving clouds	30
Figure 13 Irradiance measurements with moving clouds (close-up).....	30
Figure 14 Midday ambient temperature signal	31
Figure 15 Midday module temperature signal.....	31
Figure 16 Step change in the ambient temperature signal (close-up).....	32
Figure 17 Step change in the module temperature signal (close-up).....	32
Figure 18 Average spectrum of the ambient temperature signal	33
Figure 19 Average spectrum of the module temperature signal	33
Figure 20 DC current signal under stable conditions.....	35
Figure 21 Average spectrum of the DC current signal under stable conditions	36
Figure 22 DC current signal with 60 Hz harmonics removed	37
Figure 23 DC current signal with 60 Hz and 35.9 Hz harmonics removed (close-up).....	37
Figure 24 DC current signal with and without high-frequency content	38
Figure 25 DC voltage signal under stable conditions	39
Figure 26 Average spectrum of the DC voltage signal under stable conditions.....	39
Figure 27 DC voltage signal with 60 Hz and 35.9 Hz harmonics removed (close-up)	40
Figure 28 DC voltage signal with and without high-frequency content.....	40
Figure 29 DC power signal under stable conditions	41
Figure 30 Average spectrum of the DC power signal under stable conditions	41
Figure 31 DC power signal with 60 Hz and 20 Hz harmonics removed (close-up).....	42
Figure 32 DC power signal with and without high-frequency content.....	43
Figure 33 Average spectrum of the instantaneous AC voltage signal	45
Figure 34 RMS AC voltage signal under stable conditions.....	46
Figure 35 Average spectrum of the RMS AC voltage signal under stable conditions	46

Figure 36 RMS AC voltage signal with 60 Hz and 20 Hz harmonics removed (close-up).....	47
Figure 37 RMS AC voltage signal with and without high-frequency content	47
Figure 38 Average spectrum of the instantaneous AC current signal	48
Figure 39 RMS AC current signal under stable conditions	49
Figure 40 Average spectrum of the RMS AC current under stable conditions	49
Figure 41 RMS AC current signal under stable conditions (close-up).....	50
Figure 42 RMS AC current signal with and without high-frequency content.....	50
Figure 43 Average spectrum of the instantaneous AC power signal.....	51
Figure 44 Mean AC power signal under stable conditions.....	52
Figure 45 Average spectrum of the mean AC power signal under stable conditions.....	52
Figure 46 Mean AC power signal under stable conditions (close-up).....	53
Figure 47 Mean AC power signal with and without high-frequency content.....	53
Figure 48 DC current signal with rapidly changing irradiance.....	54
Figure 49 DC voltage signal with rapidly changing irradiance	55
Figure 50 DC power signal with rapidly changing irradiance.....	55
Figure 51 Mean AC power signal with rapidly changing irradiance.....	56
Figure 52 RMS AC current signal with rapidly changing irradiance	56
Figure 53 RMS AC voltage signal with rapidly changing irradiance.....	57
Figure 54 Graphical results of one simulated measurement chain	60
Figure 55 $P_{ac,mean}$, mean error	61
Figure 56 $P_{ac,mean}$, RMS error	62
Figure 57 $P_{ac,mean}$ with 1 Hz analog low-pass filter, mean error.....	63
Figure 58 $P_{ac,mean}$ with 1 Hz analog low-pass filter, RMS error	63
Figure 59 $P_{ac,mean}$ with 0.1 Hz analog low-pass filter, mean error.....	64
Figure 60 $P_{ac,mean}$ with 0.1 Hz analog low-pass filter, RMS error.....	64
Figure 61 $P_{ac,mean}$ with analog low-pass filter adapted to sampling rate, mean error	65
Figure 62 $P_{ac,mean}$ with analog low-pass filter adapted to sampling rate, RMS error	65
Figure 63 Slowly varying offset in 1-second measurements caused by line frequency drift	68
Figure 64 Illustration of two AC power line noise rejection methods.....	69
Figure 65 Frequency response of two power line noise rejection methods.....	70
Figure 66 Slowly varying offset in 10-second measurements due to anti-islanding ripple.....	71
Figure 67 Aliasing of anti-islanding ripple with one-second measurements.....	72
Figure 68 Aliasing causing slowly varying offset in average of 1-second measurements	72
Figure 69 Aliasing causing periodic groups of peaks in current signal.....	73
Figure 70 Aliasing with anti-islanding test causing frequent isolated peaks.....	73

TABLES

Table 1. Measurement chains for different signals in the test system	24
Table 2. Summary of signal observations.....	58
Table 3. Numerical results of one measurement chain.....	60

NOMENCLATURE

AC	alternating current
dB	decibel
DC	direct current
DOE	Department of Energy
FFT	fast Fourier transform
MPPT	maximum power point tracking
NMR	normal mode rejection
POA	plane-of-array
PV	photovoltaic
SNL	Sandia National Laboratories

1. INTRODUCTION

The number of PV systems feeding electricity into the grid is growing rapidly. At the same time, awareness is growing of the benefits of carefully monitoring the operation of these systems. This is in part brought on by the increasing system sizes, which involve larger investments and larger revenue streams. Keeping close tabs on system operation is one way to mitigate financial risks.

The capabilities of the available monitoring equipment, and also the self-monitoring capabilities of many major components such as inverters, have been continually enhanced to meet this need. The net result is that PV systems are now capable of generating large volumes of operational data, and it can become a challenge to transmit, store and/or process this data, and extract the desired intelligence in a timely and effective manner.

This study has three components. First, it takes a detailed look at the most important signals that are monitored in PV systems using data collected at 2000 samples per second to find out what measurement challenges each signal holds. It then explores how different ways of processing and sampling these signals can affect the accuracy of the data that are archived in logs and databases. And finally, it presents several real-world examples of measurement problems that can be understood and managed or even avoided with the understanding gained in the first two sections.

2. BACKGROUND

Much of what has been written about PV monitoring focuses on the analysis and interpretation of operational data that has been collected. This is not surprising because the value of the monitoring effort lies there. However the prerequisite measurement processes cannot be taken for granted. It is very possible to come to the wrong conclusions about a PV system's status based on inaccurate data.

Accuracy involves much more than a number on a spec sheet. If a data logger can take a single voltage measurement with an accuracy of $\pm 1\%$, what does that mean for the accuracy of hourly values in the log files? To answer this and similar questions, we need to examine the process by which measurements are converted to log entries. This typically involves taking more than one measurement and calculating an average, but there are many other relevant details.

Examining the measurement process alone is not enough: we must go further in both directions. On the measurement side we must consider what periodic measurements can or cannot tell us about the continuously varying signal; and on the analysis side we must understand how the log entries will be interpreted. For example, there have been a number of papers that discuss how rapidly irradiance conditions can vary [1], and how taking average irradiance values over different time intervals can lead to different estimates of energy yield, in particular due to non-linear effects such as inverter power limiting. [2,3] On the front end we have the question of how frequently irradiance measurements should be taken in order to produce an accurate average value for a given time interval; and on the back end we have the question of how short or long that time interval should be to produce an accurate system simulation.

The difficulty with the back end is that we don't always know in advance how the data will be used or understood. The usual interpretation of a value associated with a time interval is that the value represents an average or a sum over the interval. However for some purposes it could be better to have periodic instantaneous values, or moving averages spanning multiple intervals, or maximum or minimum values over the interval. PV monitoring systems could—at least theoretically—generate multiple values from the same measurements to cover all anticipated uses, but this is not commonly done yet.

Another detail that is subject to interpretation is the value of the time stamp associated with measurements. A time stamp could indicate the beginning of the averaging interval, the middle, or the end. Using the wrong interpretation can lead to wrong conclusions, particularly if the time intervals are long. Fortunately this is one of the few things that can usually be adjusted later in the data logs by uniformly shifting all time stamp values. (Shifting the measurement intervals themselves, for example to be centered on the half hour rather than on the full hour, is not so easily done.) The most important ingredient with respect to time stamps is clear documentation about the conventions that are used.

When one is presented with performance data from a PV system, it is common to find some problems, or some things that do not make sense. A recently published report [4] presents Figure 1 as an example of data anomalies, because an increase in irradiance from one time step to another does not always lead to an increase in output power. Sometimes there may be a logical

explanation for such anomalies, but if not enough additional information is available, the only recourse maybe to remove inconsistent data from the data set so that they do not distort the analysis. However there is no guarantee that all anomalies can be reliably detected and removed. Although there is likewise no guarantee that all anomalies can be prevented, we believe that PV monitoring system design and operation can be improved to significantly reduce their number.




Figure source: [4] page 20.

Figure 1 Performance data showing inconsistent trends in irradiance and power

If the documented function of a specific PV monitoring system is to record the average values of each input signal over specified time intervals, then we could evaluate its accuracy by separately measuring the true average values over those intervals using a second measurement system of known accuracy, and comparing with the archived values. This would produce an objective indication of overall accuracy including contributions from analog and digital circuits, hardware and software capabilities, etc. Such knowledge could be useful to the user of that specific system, but to better understand its strengths and weaknesses we also need to know what is going on inside. After all, many monitoring systems have options and settings that can be adjusted, and these can perhaps influence the overall accuracy.

2.1 The Measurement Chain

PV monitoring system operation can be decomposed into a sequence of basic steps. Some monitoring systems may offer a little more or less functionality of course, but we are interested in differences arising in the same core operations. For example, all systems do analog to digital conversion, but they may do so at different rates. Whether this analog to digital conversion takes place at the data logger or inside a smart combiner box is less important.

The complete sequence of steps leading from a signal to a value in a data log or archive is referred to as a “measurement chain”. In the following section we describe the steps comprising a measurement chain and identify the key parameters in each step that may vary between different monitoring systems.

2.1.1 Physical to Electrical Conversion

Many of the signals in PV monitoring systems are electrical in nature at the source, allowing a voltage or current to be measured directly or via an electrical transducer. Other signals require what would normally be referred to as a “sensor” to convert some other physical quantity to an electrical one. The main sensor types are the irradiance sensors, such as pyranometers and reference cells, and temperature sensors, such as thermocouples and temperature-dependent resistors (RTDs).

2.1.2 Analog Processing

Electrical signals, whether they are produced by a sensor or obtained directly from an electrical system, frequently need some processing before they can be digitized. The main types of processing are:

- **Amplification.** A signal that is too large or small requires an amplifier with the appropriate gain. An amplifier providing isolation may also be needed to ensure safe operation.
- **Line-frequency filtering.** Signals are often tainted with line-frequency noise (60 Hz in North America) and harmonics that must be removed.
- **Low-pass filtering.** A signal may contain legitimate but unwanted high-frequency content. To capture high-frequency content a high sampling rate is required. (Theoretically twice the highest frequency, and in practice even higher.) When lower sampling rates are used, the high-frequency content from the original signal reappears in the sampled signal as low frequencies that are impossible to distinguish from the original low-frequency content, an effect called aliasing. This can be prevented by filtering out high-frequency content above half of the sampling rate *before* sampling. A low-pass filter used for this purpose is referred to as an anti-aliasing filter.
- **RMS conversion.** For AC current and voltage signals the rapid oscillations may not be of interest, but a simple low-pass filter would only produce an average value of zero. The root-mean-square (RMS) value can be calculated instead, to obtain a magnitude that is similar to DC current and voltage.

2.1.3 Analog to Digital Conversion

The A/D conversion step is the transition from the continuous analog world to the discrete digital world: voltage levels are converted to numbers at regular time intervals. The main parameters of interest here are the magnitude resolution, specified by a number of binary bits, and the time resolution or sampling rate. Analog amplification should ensure that the signal range fully utilizes the available range of discrete digital values, and analog filtering should take into account the sampling rate to prevent aliasing.

There are other useful features, such as the ability to integrate a signal over a specified time interval rather than take a (nearly) instantaneous reading. Integration over the duration of a line frequency cycle (16.7 ms. for 60 Hz) can be used to filter out line frequency noise, for example. When multiple signals are converted, some products can do this simultaneously, but most process them sequentially.

2.1.4 Digital Processing

After the analog signal is reduced to a series of bits the processing possibilities are almost endless, at least in theory. In practice many data loggers support only a few basic operations:

- **Unit conversion.** The binary number from the A/D conversion is transformed to a number in engineering units—units that correspond to the physical quantity being measured. This might be a simple multiplication or a more complex transformation such as for thermocouples or thermistors whose voltage/resistance does not vary linearly with temperature.
- **Signal combination.** Sometimes two sensor signals need to be combined to calculate a physical quantity. Examples are, using reference cell current and temperature to calculate irradiance, or using current and voltage to calculate power.
- **RMS conversion.** If a root-mean-square (RMS) value is needed, and not produced on the analog side, then it must be calculated on the digital side.
- **Data reduction.** Commonly not all measurements are stored in the data logs, so an essential processing step is to take sequence of measurements and calculate a smaller number of values for archiving. This is sometimes referred to as decimation. Different decimation rates may be used (ratio of measurements in to values archived), and also different calculations are possible.
- **Time stamping.** Finally, a real-time clock must be consulted to put an appropriate time stamp on each archive value.

2.2 Measurement Chain Parameters

It should be clear from the foregoing descriptions that a large variety of different measurement chains is possible. In a single monitoring system, there may be several variations to accommodate different signal types, but even for a single signal type, a range of different parameters can be used. In this study we will focus on the following parameters:

- the type of analog filter applied before conversion to digital form
- the design or cut-off frequency of the analog filter
- the sampling rate (or interval) of the analog to digital conversion
- the type of data reduction calculation used to determine which values are stored, which is implemented as a digital filter
- the design or cut-off frequency of the digital filter
- the archive rate (or interval) that determines how many values are stored

3. METHODOLOGY

In this study we examine the effect of different measurement chain parameters on overall measurement accuracy. We do this not by building different measurement systems, but by simulating them. Suitable signals are obtained from a normal operating PV system with a grid-connected inverter using a high-performance data acquisition system capable of acquiring measurements at 2000 samples per second simultaneously on four channels. Any signal content above 900 Hz is filtered out before the analog-to-digital conversion. We assume that frequency content above this threshold is either irrelevant or non-existent, and treat the digital signals we obtain as accurate representations of the “true” signals. (More details about the hardware are found in the following section.)

The signals we have selected are:

- Irradiance in the plane of the array (POA), as measured by three different types of sensors
 - thermopile pyranometer
 - silicon pyranometer
 - PV reference cell
- Ambient temperature measured with a thermocouple
- Module temperature measured with a thermocouple
- Instantaneous array output voltage
- Instantaneous array output current
- Instantaneous array output power (calculated offline from voltage and current)
- Instantaneous inverter output voltage
- Instantaneous inverter output current
- Instantaneous inverter output power (calculated offline from voltage and current)

In general we examine each signal type individually. The exceptions to this are the power signals, which are derived from current and voltage measurements.

The first part of the analysis is to determine what types of variations in signal level are apparent at different time scales. This is done by zooming in on portions of the signals in the time domain, and also by examining the spectrum and identifying components of the signal in the frequency domain.

In the second part we simulate the operation of different measurement chains, to determine what values would be measured and stored in each case. Analog filters are implemented as digital calculations, and analog-to-digital conversion simply consists of taking evenly-spaced samples from the set of high-frequency measurements. Data reduction options are implemented by another combination of digital filter and selection of evenly-spaced samples. The objective of the second part is to assess a wide range of measurement chain parameters and explore how they correlate with overall measurement accuracy.

In the third and final part, we look at some individual scenarios and illustrate how inaccuracies come about.

3.1 Test Cases

The measurement chain test cases are composed by varying the following parameters:

1. the type of analog filter applied before conversion to digital form
2. the design or cut-off frequency of the analog filter
3. the sampling rate (or interval) of the analog to digital conversion
4. the type of data reduction calculation used to determine which values are stored, which is implemented as a digital filter
5. the design or cut-off frequency of the digital filter
6. the archive rate (or interval) that determines how many values are stored

There are a few real-world options that at first glance do not seem to fit into this scheme. The first is the ability of the A/D converter to integrate a signal in order to eliminate line-frequency noise. This operation has the same effect as placing a moving average filter before a non-integrating A/D converter, and it can therefore be represented as two distinct steps in the simulated measurement chain.

The RMS calculation for AC current and voltage is also a special case. In real-world PV monitoring systems this is normally done by specialized hardware. For the simulations the RMS values had to be calculated, but we did not explore different parameters that might be used there.

Another real-world option is the use of a decimation filter that combines the functions of digital filtering and sample-rate reduction in a computationally efficient manner, for example a cascaded integrator comb (CIC) filter. [7] The effect of a CIC filter on the signal can be simulated in two separate steps as well: first applying a moving average filter, and second, selecting every R^{th} sample, where R is the ratio of sampling rate to archive rate. Higher-order versions of the CIC filter have the same effect as applying the moving average filter multiple times to the same signal.

The ranges of values or options for each parameter that we investigate are as follows:

- **Analog filter type** is one of:
 - none
 - single-pole low-pass filter
 - two-pole low-pass filter
 - four-pole low-pass filter
 - single pass moving average (60Hz only)
- **Analog filter design frequency**
 - 60 Hz (moving average)
 - 1 Hz (low-pass)
 - 0.1 Hz (low pass)
 - $1 / \text{sampling interval}$ (low-pass)
 - $1 / \text{sampling interval} \times 2$ (low-pass)
- **Sampling interval** ranges from 0.1 to 900 seconds (15 minutes)
- **Digital filter type** is one of:

- none (take the most recent sample)
- single pass moving average (rectangular window; first-order CIC)
- dual pass moving average (triangular window; second-order CIC)
- **Digital filter design frequency**
 - 1 / archive interval
 - 1 / archive interval x 2
- **Archive interval** ranges from 1 to 1800 seconds (30 minutes)

The design frequency for the filters is defined as the cut-off frequency for the low-pass filters, and as the location of the first notch, or zero, in the frequency response of the moving average filters. The width of the moving average window is therefore equal to the inverse of the design frequency.

The frequency response of each filter type is shown in Figure 2. The two types of filters have very different characteristics, the moving average type having very distinct notches at multiples of the design frequency, and the low-pass filters showing no preference for specific frequencies but rolling off in a gradual manner with frequency. A comparison of a moving average filter and a low-pass filter with the same design frequency is not meaningful in general as the design frequency has to be adjusted to achieve specific filtering goals. For example, if a low-pass filter were used to suppress line-frequency noise, its design frequency would have to be much lower.




Figure 2 Frequency response of the filter types that were used

The frequency response characteristics are important when dealing with continuing oscillations in a signal, whether they be sharply defined such as line frequency or broadly distributed in the spectrum such as thermal noise. Changes in irradiance are generally not periodic (discounting the diurnal and annual cycles) but are rather more like randomly occurring step changes with varying slopes, magnitudes and durations, and this drives the characteristics of many of the other signals as well. The spectrum of such a signal can be calculated over different intervals, but it is forever changing.

Another important characteristic of the filters, therefore, is the step response. Figure 3 shows the positive and negative step response of the same set of filters, with the design frequencies adapted to a realistic time scale. For the moving average filter, additional passes that improve noise suppression also slow the response to rapid step-like changes. The higher-order low-pass filters do not respond quite so sluggishly, but have the possible disadvantage of shooting beyond the actual signal level. With complex signals, therefore, both frequency domain and time domain characteristics of the filters are important.




Figure 3 Step response of the filter types that were used

3.2 Evaluation

To evaluate the accuracy of each measurement chain, an optimal measurement chain must also be defined. To define this optimal chain, we must ask the following question: given a certain target archive interval, what value should ideally be stored for each interval? The answer to this question depends on what the stored values will subsequently be used for, and since there may be multiple uses, there may be multiple answers. It is also possible that all the ways of using the data are not known in advance.

The most important signals in PV monitoring are those that represent power: irradiance, DC power and AC power. Key system performance indicators such as yield, efficiency and performance ratio are derived from these quantities. Integrated over time, these power signals measure energy, therefore the ideal value associated with each archive interval is the *average value over the interval*, because only that value, multiplied by the duration of the interval, gives the correct energy value.

For DC current and voltage, the ideal value is not as easily defined. To begin, the average power over an interval does not in general equal the product of the average voltage and the average current. Thus, for consistency with DC power values, some voltage or current value other than the average value over the interval might be more appropriate. On the other hand, in some systems *only* the DC current or *only* the DC voltage may be measured, so ideal archive values should be defined independent of each other and of the DC power.

There are other potentially interesting values besides the mean that could be extracted from an archive interval, such as the *median* or the *mode*, or even the most recent value, however we know of no examples where these descriptive statistics are used in downstream analysis of PV systems. It is perhaps easier to imagine a downstream use for the *maximum* and *minimum* DC voltage and DC current over an interval, and this naturally leads to the idea of storing multiple values for an interval rather than a single one. Together, the mean, standard deviation and two extremes for an interval, for example, provide much more complete and potentially useful information about a signal than any one of them alone. And this idea is of course applicable to any of the signals being measured, not just DC voltage and current. There is a trade-off involved though: extra storage is required for the descriptive statistics. If this extra storage is available it could just as well be used to shorten the archive intervals and thereby improve the resolution of actual signal variations.

Despite the foregoing observations, we come to the conclusion that if a single value is to be stored per interval, it is preferable to store the average value for all the signals under consideration. The principal reason for this is that PV system performance is frequently compared to simulation program output, and such programs generally make use of time interval average values from weather files, and produce time interval average outputs. It should also be noted that using the mean for some signals and another metric for other signals could easily lead to confusion and errors in practice, and in the absence of a compelling reason to do otherwise, consistency is preferred. Lastly it is worth mentioning that older formal guidelines for PV monitoring systems, such as EUR IEC 61724:1998 [5] and IEC 16338 EN [6], provide clear recommendations to calculate and store mean values or sums for all signals.

Thus, for the purpose of this study, the ideal measurement chain produces the true average value of the signal over each archive interval. The less-than-ideal measurement chains attempt to produce the same value by taking occasional samples and processing those in different ways. The difference between the two possible archive values is the *error* resulting from that particular measurement chain. (Note that this definition does not work for instantaneous AC voltage and current, since they have mean values of zero. We have to work with the RMS values for these two signals.)

A single simulation takes a three-hour segment of an input signal, calculates the error for each archived value, and at the end retains only the mean error and the RMS error. After running simulations for a large number of measurement chains, it is possible to see how sampling and archive rates relate to the magnitude and distribution of the errors, and also how the different filtering options change those relationships.

4. HARDWARE AND SOFTWARE DESCRIPTION

4.1 PV System

In order to evaluate several possible measurement chains which may be used with PV systems, Sandia has measured irradiance, temperature, array current, array voltage, inverter output current and inverter output voltage on an operating PV system.

The PV system which was measured consisted of sixteen series-connected MS145GG-02 copper indium gallium diselenide (CIGS) modules from the MiaSolé company. The modules were mounted south-facing with a latitude tilt of 35 degrees from horizontal. The PV array output was inverted to AC power by a Fronius IG Plus 3.0 inverter with a nominal 208 V_{AC} output. Four modules were instrumented with self-adhesive type T thermocouples located on the module back.

Reference irradiance instruments were located just above the PV array, oriented in the same plane as the PV modules. The irradiance instruments consisted of: one Eppley PSP first-class broadband pyranometer, one Licor LI200 silicon photodiode pyranometer, and one EETS RC1 split silicon reference cell. One type-T thermocouple probe in an aspirated radiation shield served as an ambient air temperature sensor, located approximately 6 meters from the array.



Figure 4 The PV system used for the study



Figure 5 Three irradiance sensors used for the study

4.2 Data Acquisition

The data acquisition system was built around a National Instruments Compact RIO (cRIO) model 9076 with a single NI 9239 input module. The cRIO sampled the signals at a 2 kHz rate and stored the data (in binary) on a standard USB flash drive.

Most signals, however, could not be measured directly by the 9239 module's $\pm 10\text{V}$ input, and some signal conversion was required. The signal conversion for each measured signal is detailed in the following list.

- DC voltage – The DC voltage was converted to a 0-10 V signal by an Ohio Semitronics transducer, model VT7-010D powered by a 60 Hz, 120 V_{AC} source. The VT7 output was read by the 9239.
- DC current – The DC current was converted first by a 10A/100mV shunt from Empro. Then converted to a 0-10 V signal by an Ohio Semitronics transducer, model VT7-016D. The VT7 output was read by the 9239.
- AC voltage – The AC voltage was converted to a $\pm 10\text{ V}$ signal by an Ohio Semitronics transducer, model VT7-008D-11. The VT7 output was read by the 9239.
- AC current – The AC current was converted to a small voltage by a 20A/50mV shunt from Empro, then converted to a $\pm 10\text{ V}$ signal by an Ohio Semitronics transducer, model VT7-015D-11. The VT7 output was read by the 9239.
- Pyranometer irradiance, G_{pyr} – The Eppley PSP produces a full-sun signal of approximately 11-15 mV. In order to increase the signal level into the range of the 9239, the pyranometer voltage was amplified by an Analog Devices 7B40 isolating amplifier before being read by the 9239.

- Photodiode irradiance, G_{diode} – The Licor photodiode produces a full-sun signal of approximately 10-15 mV, and thus also required an Analog Devices 7B40 isolating amplifier prior to being read by the 9239.
- Reference cell irradiance, G_{ref} – The EETS RC1 reference cell produces a full-sun signal of approximately 50 mV, and also used a 7B40 amplifier prior to being read by the 9239.
- Reference cell voltage – The open circuit voltage of the EETS RC1 reference cell, when illuminated, is approximately 0.6 volts, and required no amplification. It was measured directly by the 9239.
- Module temperature – The type-T thermocouple used for module temperature measurement requires cold-junction compensation and amplification by a separate module. The OMEGA Engineering amplifier, model TXDIN 6120, was used to read the module temperature and produce a 0-10V output. The TXDIN output was read by the 9239. The TXDIN was calibrated with an OMEGA CL27 thermocouple calibrator rather than using the linear input/output expected from its settings.
- Ambient temperature – An OMEGA TXDIN 6120 with similar settings to the module temperature TXDIN was used. The ambient temperature TXDIN was calibrated with the CL27 thermocouple calibrator.

The essential measurement characteristics for each step in the measurement chain of this system are summarized in Table 1.

4.3 Calculation Platform

The acquired data were processed using MatLab and its signal processing toolbox. The latter was used primarily to calculate the filter parameters for the digital equivalents of the low-pass filters according to Butterworth designs, but also in preliminary work in exploring the signal and filter characteristics.

The volume of data (7,200,000 values per signal per hour) presented some challenges, and eventually led to the practical limit of 3 hours for the measurement chain simulations in section 6. With further optimization, however, full days could realistically be simulated.

Table 1. Measurement chains for different signals in the test system

Measurement step	Irradiance	Irradiance	Irradiance	Temperature
physical to electrical conversion	thermopile	photodiode	reference cell	thermocouple
analog processing	amplification: 100x	amplification: 100x	none	non-linear amplification
	low-pass filter: 10 kHz, one-pole low-pass filter: 10 kHz, four-pole	low-pass filter: 10 kHz, one-pole low-pass filter: 10 kHz, four-pole		low-pass filter: response time, 95% < 500ms
	low-pass filter: 900 Hz	low-pass filter: 900 Hz	low-pass filter: 900 Hz	low-pass filter: 900 Hz
analog to digital conversion	2000 samples/second	2000 samples/second	2000 samples/second	2000 samples/second
digital processing	conversion: volts to W/m ²	conversion: volts to W/m ²	conversion: volts to W/m ²	conversion: volts to C
	store value with timestamp	store value with timestamp	store value with timestamp	store value with timestamp
	DC voltage	DC current	AC voltage	AC current
physical to electrical conversion	none	10A/100mV shunt	none	20A/50mV shunt
analog processing	amplification: 0.017x	amplification: 100x	amplification: 0.025x	amplification: 200x
	low-pass filter: 10 kHz	low-pass filter: 10 kHz	low-pass filter: 10 kHz	low-pass filter: 10 kHz
	low-pass filter: 900 Hz	low-pass filter: 900 Hz	low-pass filter: 900 Hz	low-pass filter: 900 Hz
analog to digital conversion	2000 samples/second	2000 samples/second	2000 samples/second	2000 samples/second
digital processing	conversion: volts to volts	conversion: volts to amps	conversion: volts to volts	conversion: volts to amps
	store value with timestamp	store value with timestamp	store value with timestamp	store value with timestamp

5. SIGNAL CHARACTERISTICS

This section looks at small segments of each signal to explore what kinds of variations are present, and if possible, explain the origins of those variations. This is helpful when deciding later whether to ignore, suppress, or capture those variations. Unfortunately it is not always possible to say whether an observed signal variation originated in the signal source or in some element of the measurement chain. This is especially true of the line frequency and its harmonics, which are omnipresent.

For some diagrams, specific high-frequency components have been filtered out to better illustrate what low-frequency variations they mask. A rectangular window (or moving average) filter was used for this, applied once in the usual direction and once in reverse, thereby avoiding any time/phase shifts. This technique is mainly suitable for offline processing.

The spectrum of each signal segment is also calculated and shown. The magnitude scale is logarithmic and shown in decibels (dB) where 20 dB corresponds to factor of 10. The magnitudes are normalized to largest peak (usually the DC value) to give a sense of proportion. For example, if a sharp peak occurs at 60 Hz with a magnitude of -40 dB, that corresponds to 60Hz ripple with an amplitude that is 1% of the signal's DC value. (2% peak to peak.) With broad peaks or humps, or a harmonic series of peaks, the magnitudes are not as easily interpreted. The temperature spectra are not normalized to the DC value, but rather the DC component is removed. The magnitude is therefore absolute, but also presented on a dB scale. (0 dB = 1.0 C, -20 dB = 0.1 C, etc.)

The starting point for the examination of characteristics of each signal is a 5-minute segment of data where the operating conditions are nearly constant. As only four high-speed acquisition channels were available, these data segments could not all be collected at the same time. Instead, irradiance signals, temperature, and electrical signals were collected during three consecutive periods.

5.1 Irradiance

Irradiance is the most fundamental measurements for PV system monitoring. Three types of sensors are commonly used: thermopile pyranometers, photodiode-based pyranometers, and reference cells. Signals were obtained simultaneously from one unit of each type.




Figure 6 Stable irradiance signals from three sensors

Figure 6 shows that the three irradiance signals are very close to each other in the range $1005 \text{ W/m}^2 \pm 1\%$, but what is more interesting in this context is that they evidently have different levels of noise: most for the photodiode sensor, and least for the reference cell. Another interesting feature is the infrequent occurrence of larger impulses, which sometimes even appear simultaneous.




Figure 7 Simultaneous disturbance of three irradiance signals

In Figure 7 the three irradiance signals are magnified considerably, showing not only simultaneous impulses but also a brief ringing afterwards. Clearly this is not caused by a change in irradiance since thermopile could not respond that quickly. Some kind of static discharge in the data acquisition equipment seems possible, but a definite explanation cannot be given.

To investigate the unequal level of noise on the three signals, we calculate the spectrum for each signal, focusing on the frequency range 0–200 Hz.




Figure 8 Average spectrum for the irradiance signal from the thermopile




Figure 9 Average spectrum for the irradiance signal from the photodiode




Figure 10 Average spectrum for the irradiance signal from the reference cell

There are both similarities and differences in the above three figures. The narrow peaks at multiples of 40 and 60 Hz are similar to those found in the AC voltage spectrum, which is a strong indicator of their origin. Their amplitudes are all very small at less than 0.1% of the signal.

A more interesting feature is the profile of the remaining noise, which is different for each sensor. The thermopile cannot respond to fast irradiance changes because of its thermal mass, so above a few Hz there is only white noise. The reference cell shows a rapid drop in high-frequency content down to 40 Hz. In theory it could respond quickly, but its large area has an averaging effect. The photodiode, with its small area and ability to respond quickly has the

longest tail, indicating irradiance fluctuations above 100 Hz are likely present in very small quantities. The base white noise level for the thermopile and photodiode are higher because their signals are smaller to start. They also have higher impedances than the reference cell shunt resistor.




Figure 11 Irradiance measurements with line frequency ripple removed

Figure 11 shows the same three irradiance signals after the line frequency components and harmonics were removed. The pyranometer signal has very little remaining high-frequency content, while the other two sensors still show rapid fluctuations in irradiance that the pyranometer could not capture.

The final two graphs in this section show measurements from a different time interval where the irradiance fluctuates rapidly due to moving clouds.




Figure 12 Irradiance measurements with moving clouds




Figure 13 Irradiance measurements with moving clouds (close-up)

While it is not so clear in Figure 12, the delayed response of the pyranometer is quite apparent in Figure 13. The peak reading near the midpoint arrives 1-2 seconds after the other two sensors record it, and the peak magnitude is also not quite reached.

5.2 Temperature

Module temperature and ambient temperature are two other key parameters for PV performance monitoring. Usually temperatures do not fluctuate as rapidly as many other parameters because the thermal mass of materials limits the rates of change. In addition, the two thermocouples used here are connected to amplifiers that have specified response times of 500 ms, so even rapid changes in the thermocouple signals, including noise and line-frequency ripple, should be attenuated significantly. Nevertheless, our data acquisition system picked up some extremely rapid fluctuations.




Figure 14 Midday ambient temperature signal




Figure 15 Midday module temperature signal

The above two graphs show the same 5-minute time segment for each temperature signal in the middle of a very clear day. The range over which each signal fluctuates is approximately the same about 1.5 °C and presumably driven by air movement since irradiance is nearly constant.

Two remarkable features of these graphs are periodic narrow pulses—the largest one about 1 °C in magnitude—and a uniformly thick line, perhaps indicating some high-frequency content.




Figure 16 Step change in the ambient temperature signal (close-up)




Figure 17 Step change in the module temperature signal (close-up)

Close-ups of the two temperature signals over a 10-second period (Figure 16 and Figure 17) confirm that the pulses occur precisely at the same time and last exactly one second. The

presumed high-frequency content is still not clearly seen, but unlike the noise seen on the irradiance signals, its amplitude is extremely consistent. Closer examination of the time plots (not illustrated) shows that the high-frequency content looks very much like a sine wave, and that the pulses come at regular intervals of approximately 65 seconds. However, some smaller glitches are also visible between the larger pulses and not all pulses are observed on both channels (e.g., 240 s).




Figure 18 Average spectrum of the ambient temperature signal




Figure 19 Average spectrum of the module temperature signal

The spectra of the two signals (Figure 18 and Figure 19) confirm that there is indeed a clean sine wave near 126 Hz on each temperature signal. The frequencies are slightly different for the two signals and could be artifacts of inductive coupling of power or signal within the amplifiers, used

to isolate the input stage from the power supply and output. As the peak-to-peak amplitude of this ripple is less than 0.1 C this is certainly negligible for performance calculations.

The rest of the spectrum is largely the result of the pulses, which produce a long harmonic series of very gradually decreasing peaks in the frequency domain. Due to the facts that the pulses mostly occur simultaneously on the two independent input channels, and that the level change occurs from one sample to the next (0.0005 s) this phenomenon must originate somewhere in the main data acquisition unit. For PV monitoring these signal pulses would normally be negligible as well. Occasional erroneous measurements are usually absorbed in longer-term averages, and this would therefore appear to be a negligible source of error as well. However, as its origin is not understood, its potential to cause further problems cannot be assessed.

5.3 DC Current, Voltage and Power under Stable Conditions

The following sections examine the electrical signals during a 5-minute interval in the middle of a different sunny day. With the external conditions (nearly) constant, the observed signal variations are principally caused by inverter operation.

5.3.1 DC Current




Figure 20 DC current signal under stable conditions

Figure 20 shows the “raw” DC current signal. From this it is immediately obvious that there is considerable fluctuation despite the relatively stable mid-day conditions. The measured values range from 6.0 to 6.2 A, a range of $\pm 1.6\%$.




Figure 21 Average spectrum of the DC current signal under stable conditions

To identify periodic components within this signal, we examine the spectral composition of the signal segment and find quite a number of prominent peaks. (Figure 21)

The highest peak occurs at 120 Hz, which is the second harmonic of the line frequency and also the frequency at which the instantaneous AC power fluctuates at the output of the single-phase inverter. The large DC bus capacitors inside the inverter are there to decouple this power flow, but there appear to be some residual fluctuations here on the DC side. These are not purely sinusoidal though, since various harmonics of 60 Hz are visible as well.

To visualize the contribution of these 60 Hz harmonics, we remove them from the signal by applying a moving average filter that has a period equal to $1/60^{\text{th}}$ of a second (once forward and once in reverse to avoid time shifts). The resulting filtered signal is shown in Figure 22.




Figure 22 DC current signal with 60 Hz harmonics removed

This filtered signal appears much more focused than the original, but the largest fluctuations still remain. Looking back at the spectrum we see that the second largest peak occurs at approximately 35.9 Hz, and there is a long series of peaks at 2.4 Hz and multiples thereof. We therefore zoom into the signal to look whether an explanation could be found there.




Figure 23 DC current signal with 60 Hz and 35.9 Hz harmonics removed (close-up)

Figure 23 shows a one-second segment of the filtered signal and the approximately 36 peaks of the 35.9 Hz signal are readily apparent here. There are also rounded impulse-like perturbations approximately 0.4 seconds apart, which explain the peak at 2.4 Hz (period of 0.416 s) and the associated train of harmonics. But why are they there?

The explanation comes from the fact that the same periodic impulse-like perturbations are seen very strongly on the AC current (Figure 41) and AC power (Figure 46). In fact, they occur exactly once every 25 line cycles and are a byproduct of the inverter's active islanding detection system, which probes the grid at regular intervals to ensure that it is still up and connected. (A similar effect was observed every 30 cycles in another system with a different inverter.) The brief perturbation in output current and power destabilizes the DC operating point briefly, and the 35.9 Hz oscillations are probably evidence of under-damped control of the operating point intended to get it back to where it should be as quickly as possible.




Figure 24 DC current signal with and without high-frequency content

When we now remove the 35.9 Hz, and the 2.4 Hz components (once again using a simple moving average filter of the appropriate width), the resulting signal is very clean. (Red trace in Figure 24) The fluctuations that remain during this stable period are clearly the searching action of the MPPT algorithm. A step change in DC current occurs every 3.33 seconds, and the step size is about 0.33%.

5.3.2 DC Voltage

The DC voltage signal is analyzed for the same stable 5-minute period.




Figure 25 DC voltage signal under stable conditions

At first glance the voltage signal in Figure 25 appears to be very similar to the current signal, and indeed it is. The values range from about 313 to 324 V, which is $318.5 \pm 1.7\%$, and the higher frequency variations appear to have similar magnitudes.

What is different, however, is that whenever the current goes up, the voltage goes down. The two are of course constrained by the I-V characteristics of the array, which has a negative slope.




Figure 26 Average spectrum of the DC voltage signal under stable conditions

The spectrum of the DC voltage (Figure 26) shows that the relative amplitudes of all frequency components are virtually identical to those of the DC current.




Figure 27 DC voltage signal with 60 Hz and 35.9 Hz harmonics removed (close-up)

A comparison of the close-ups in Figure 27 and Figure 23 confirms whenever voltage goes up, current goes down.




Figure 28 DC voltage signal with and without high-frequency content

As with the current, removing the high-frequency components from the DC voltage again makes the step changes caused by the MPPT operation are clearly visible. In contrast to the DC current, however, the level of these voltage steps does not drift during individual steps, demonstrating that the MPPT search logic is prescribing the voltage level rather than a current level. Each step appears to be about 1 V.

5.3.3 DC Power

The DC power signal is created by multiplying individual voltage and current measurements.




Figure 29 DC power signal under stable conditions

A closer look at the scale of Figure 29 shows that the power range in this interval is about ± 10 W ($\pm 0.5\%$), and on shorter intervals it is about ± 2 W ($\pm 0.1\%$)—much narrower than the ranges of either voltage or current in the same interval. This is also consistent with the movement of the operating point on the IV curve. At the MPP dP/dV is zero, and it remains small for small variations around MPP.




Figure 30 Average spectrum of the DC power signal under stable conditions

The spectrum of the DC power signal (Figure 30) shows some similarities as well as some differences compared to the DC current and voltage spectra. Multiplying two sinusoidal signals creates two new sinusoidal signals with frequencies that are the sum and difference of the original two frequencies, so it is not surprising to see a few different peaks.

The large DC components in each signal amplify the original harmonics of the other signal, but the resulting magnitude at each frequency also depends on the relative phases, which the IV curve causes to be largely out of phase. Nevertheless, this explains why it is still possible to find a small peak at 35.9 Hz.




Figure 31 DC power signal with 60 Hz and 20 Hz harmonics removed (close-up)

As Figure 31 shows, removing the 60 Hz and harmonics leaves only very small fluctuations of $\pm 0.02\%$, and with further filtering of the 20 Hz multiples we are again left with only the anti-islanding perturbation visible at this time scale. The amplitude of this, it must be stressed, is very small.




Figure 32 DC power signal with and without high-frequency content

As with the voltage and current signals, once all these higher frequency components are removed, a much smoother DC power signal emerges. (Red line in Figure 32.) In contrast to the current and voltage signals, however, there is barely any evidence of the step changes caused by the maximum power point tracking.

5.4 AC Current, Voltage and Power under Stable Conditions

AC signals have instantaneous values that oscillate at the line frequency, but for PV monitoring we are primarily interested in the RMS values. Since we sampled the signals at 2,000 times per second we actually recorded the instantaneous values, and therefore must convert these to RMS.

RMS values are only meaningful over a specified time interval: it is over this interval that the square root of the mean of the squared signal is calculated. With a periodic signal, the ideal averaging period is the repetition period because it produces a constant RMS value. This averaging operation has the same frequency response characteristics as shown in Figure 2 so all the harmonics of the fundamental period would be suppressed. However the real-world the AC signal is not the same from cycle to cycle, and other frequency components are normally also present. This leads to ripple and non-periodic fluctuations in the calculated RMS value.

Analog RMS transducers usually do not calculate a true moving average, but use an internal low-pass filter instead. In order to achieve sufficient suppression of the AC components of the squared signal, the cut-off frequency must be quite a bit lower than the line frequency, and as a result, the 95% response to step changes is typically around 100-200 ms. In digital RMS transducers, true averaging can be implemented, but the averaging period may be extended to multiple cycles to reduce residual ripple. In this study the RMS values are calculated with a moving average period of one AC cycle to get a better understanding of the signal content and also to avoid making optimistic assumptions about the cleanliness of the signals.

5.4.1 AC Voltage

The AC voltage where the inverter is connected is nominally a 208 V_{RMS} sine wave, and in practice it should not deviate much from that. Figure 33 shows the average spectrum for the 5-minute period, with the expected peak at 60 Hz and a variety of other smaller peaks representing distortion of the pure sine waveform. Distortion observed here is not necessarily caused by this inverter; it could also be caused by other equipment nearby.




Figure 33 Average spectrum of the instantaneous AC voltage signal

The next graph (Figure 34) shows the calculated RMS voltage, which is on average about 2.5% above the nominal value. It also exhibits some gradual changes as well as more rapid oscillations giving a combined range of about ± 0.5 V (± 0.2 %). Since the PV system power output is known to be virtually constant during this period, the slower fluctuations must be caused by other equipment on the grid.




Figure 34 RMS AC voltage signal under stable conditions




Figure 35 Average spectrum of the RMS AC voltage signal under stable conditions

In the spectrum of the RMS AC voltage signal (Figure 35) the harmonics of the fundamental frequency are minimal, but the interaction between some of the intermediate frequencies in the RMS calculation has left a series of peaks 20 or 40 Hz apart.




Figure 36 RMS AC voltage signal with 60 Hz and 20 Hz harmonics removed (close-up)

Figure 36 demonstrates how removing these harmonics brings the signal closer to the somewhat longer-term average value that is of interest.




Figure 37 RMS AC voltage signal with and without high-frequency content

In the final figure for the RMS AC voltage, the contrast between the signal with and without the high-frequency content is illustrated. (Figure 37)

5.4.2 AC Current

The AC current is considerably more interesting, and perhaps challenging, than the AC voltage. While the RMS value of the voltage usually stays within a narrow band, the RMS current varies widely to transfer the available energy to the grid. The inverter only has a minimal and indirect influence on the AC voltage, but it is in full control of the current it pushes into the grid.




Figure 38 Average spectrum of the instantaneous AC current signal

The AC current is ideally a sine wave, but in reality it is an approximation whose characteristics are determined by the inverter circuit topology, modulation and control methods, and filtering. The spectral content of that signal can therefore be complex. Efficiency is usually the most important design goal as long as minimum power quality requirements are met, so it is not surprising to see that the noise floor in this signal is higher than in most of the other signals. Clearly visible above the noise floor is the series of 2.4 Hz harmonics that is also seen on the DC side. But they are much more prominent here on the AC side, as it is here that the active anti-islanding test is carried out.




Figure 39 RMS AC current signal under stable conditions

After calculating the RMS AC current (Figure 39) it becomes apparent that despite the filtering inherent in that calculation there are still very large fluctuations of ± 1 A or $\pm 11\%$. The spectrum of the RMS AC current (Figure 40) shows the effect of that filtering around the line frequency harmonics, but the 2.4 Hz harmonics remain prominent.




Figure 40 Average spectrum of the RMS AC current under stable conditions




Figure 41 RMS AC current signal under stable conditions (close-up)

Examining a 1-sec portion of this signal in Figure 41 we see that these anti-islanding perturbations are principally responsible for the large signal swings. Removing the 2.4 Hz harmonics from the signal leads to enormous improvements in signal stability, as Figure 42 demonstrates.




Figure 42 RMS AC current signal with and without high-frequency content

5.4.3 AC Power

As mentioned earlier, we calculate AC power from voltage and current. More specifically, we multiply the instantaneous values of the AC voltage and current signals to obtain the instantaneous power, which is a value that nominally oscillates from 0 to some peak value at twice the line frequency. Power factor therefore does not have to be taken into account as it would when multiplying the RMS values.




Figure 43 Average spectrum of the instantaneous AC power signal

In the spectrum of this power signal in Figure 43 the oscillations at 120 Hz dominate, but the important value is the DC value that represents the mean power generated over this interval. For monitoring purposes it is the mean value that should be recorded, and we calculate this over one line cycle to be consistent with the RMS calculations.

The AC power signal spectrum shows the same strong 2.4 Hz harmonics as the AC current. This is necessarily so because the AC voltage is not easily perturbed by these impulses, unless an islanding condition occurs! This situation is rather different from the DC side, where the current and voltage perturbations were out of phase, and the power exhibited only very small 2.4 Hz harmonics.




Figure 44 Mean AC power signal under stable conditions

The mean AC power signal in Figure 44 therefore fluctuates by $\pm 11\%$ just like the RMS AC current. In fact, the RMS AC current and mean AC power spectra in Figure 40 and Figure 45 respectively are barely distinguishable from one another.




Figure 45 Average spectrum of the mean AC power signal under stable conditions




Figure 46 Mean AC power signal under stable conditions (close-up)

When these sharp perturbations (Figure 46) are removed from the signal, the true mean becomes apparent (Figure 47).




Figure 47 Mean AC power signal with and without high-frequency content

5.5 DC and AC signals with Rapid Irradiance Changes

In this section we briefly illustrate how the different electrical signals vary in response to irradiance changes. The time span is 5 minutes like in the previous section, but because of the large swings in irradiance the vertical scales are larger. All signals are shown with and without high-frequency content, giving another perspective on the relative magnitudes. All signals were acquired simultaneously.




Figure 48 DC current signal with rapidly changing irradiance

Although we could not measure irradiance during this period, the DC current normally follows irradiance changes quite closely, and Figure 48 demonstrates the irradiance changes were large and fast.




Figure 49 DC voltage signal with rapidly changing irradiance

In contrast to the DC current, the DC voltage (Figure 49) adjusts comparatively slowly. In fact, the MPPT algorithm appears to impose an upper limit on the rate of change equal to 3 steps per 10 seconds, or 0.3 V/s. There are just a few occasions where the inverter fails to maintain its target voltage, such as during the rapid rise in irradiance near the 140 s mark and the subsequent drop after the 180 s mark. The otherwise regular staircase steps are distorted here.




Figure 50 DC power signal with rapidly changing irradiance

The DC power (Figure 50) has no trouble following the rapid changes, but it is hard to tell visually how close it actually stayed to the maximum power point. The long linear rise in DC voltage during the nearly constant conditions in the first 50 seconds suggests that the MMPT had gotten a bit behind in following the true V_{mpp} .




Figure 51 Mean AC power signal with rapidly changing irradiance

In Figure 51 the AC power follows the same general profile as the DC power, but is of course a little lower due to the losses in the inverter. The high-frequency content is really evident here, and is seen to decrease with the power level. The noise caused by the anti-islanding test dominates in these two signals, and since the periodic changes in magnitude are a by-product of attempts to induce phase or frequency shifts, it seems reasonable that the magnitude of the perturbations would be more or less proportional to the average value of the signal.




Figure 52 RMS AC current signal with rapidly changing irradiance

As the AC voltage is nominally constant, the profile of the AC current signal (Figure 52) is virtually indistinguishable from that of the AC power, noise included.

Despite being nominally constant, the AC voltage at the connection point is measurably influenced by the injected power. (Figure 53) But a voltage change of 0.7% caused by PV output fluctuation is fairly insignificant considering that grid voltage tolerances are typically $\pm 5\%$ or more from their nominal values.




Figure 53 RMS AC voltage signal with rapidly changing irradiance

5.6 Summary of Signal Characteristics

The table below recaps some of the observations made on the various signals. It should be stressed that not all the features that are noted would necessarily be significant sources of error. In some cases the effects are too small, in other cases it depends on the characteristics of the monitoring equipment used to measure and archive the signals, which is the subject explored in the next section.

It is important to note that the observations are made on a *single* PV system consisting of *specific* components in *one* location. The ranges given and the features noted should therefore not be interpreted as being representative of all, or even a subset of all systems. However they can give some ideas of what to look for when checking the accuracy of measurements, either at commissioning, or when a suspicion arises that something doesn't quite add up.

Table 2. Summary of signal observations

Signal	Range of high frequency content	Main high-frequency features
Irradiance, thermopile	$\pm 0.4\%$	- uniform white noise - residual line frequency harmonics - occasional extremely brief spikes
Irradiance, photodiode	$\pm 0.3\%$	- white noise slowly diminishing with frequency - residual line frequency harmonics - occasional extremely brief spikes
Irradiance, ref. cell	$\pm 0.6\%$	- white noise rapidly diminishing with frequency - residual line frequency harmonics - occasional extremely brief spikes
Ambient temperature	n/a	- ripple from transducer - recurring one-second pulses of varying amplitude
Module temperature	n/a	- ripple from transducer - recurring one-second pulses of varying amplitude
DC current	$\pm 1.6\%$	- line-frequency harmonics - anti-islanding test perturbations - maximum power point tracking step changes
DC voltage	$\pm 1.7\%$	- line-frequency harmonics - anti-islanding test perturbations - maximum power point tracking step changes
DC power	$\pm 0.1\%$	- line-frequency harmonics
RMS AC voltage*	$\pm 0.2\%$	- 20/40/60Hz harmonics
RMS AC current*	$\pm 11.0\%$	- anti-islanding test perturbations - higher general noise level than other signals
Mean AC power*	$\pm 11.0\%$	- anti-islanding test perturbations - higher general noise level than other signals

* calculated over one powerline cycle

6. SIMULATED MEASUREMENT CHAINS

Looking closely at the acquired and derived signals has given insight into various higher frequency components that may be present there, particularly how they appear in the time and frequency domains. The filters used to separate the signals from the noise were not the kind that can be easily implemented in practice, but they do illustrate the potential for appropriately chosen filters to reduce signal fluctuations and thereby *potentially* also reduce errors in the archived values.

6.1 Simulated Measurement of Mean AC Power

The first signal to be considered is mean AC power, that is, the AC power calculated by multiplying instantaneous values of AC current and voltage, and applying a one-cycle moving average. As an introduction to this section, we will look at the results in one measurement chain in more detail. The parameters of this measurement chain are as follows:

- analog filter type: single pass moving average
- analog filter design frequency: 60 Hz
- sampling interval: 300 s (5 minutes)
- digital filter type: single pass moving average
- digital filter design frequency: 1/archive interval
- archive interval: 1800 s (30 minutes)

In other words, this is a logger that has an integrating A/D converter to reject line frequency noise and take measurements every five minutes. Every half hour it stores the average value of the last 6 measurements. (The intervals are long for the purpose of the illustration.)

The results of a simulation of this one measurement chain are shown in Figure 54. The “true” power signal can be seen in the background and the black dots identify the measurements that the logger took at five-minute intervals. The half-hour means of the five-minute measurements are shown as red staircase plot, which can be compared to the actual half-hour means of the signal in the green staircase plot. Finally the mean values for the entire three-hour period are shown as horizontal lines.




Figure 54 Graphical results of one simulated measurement chain

The three-hour simulation is evaluated based on the difference between the simulated archive values and the actual half-hour means, also referred to as the correct archive values. In the graph these errors are seen as the vertical distances between the red and green lines.

For this one simulation the numeric results are listed in Table 3, which also shows the two accuracy indicators that summarize the simulation. One is the mean error, which is the distance between the long red and green horizontal lines. The other is the RMS value of the six individual errors, which is a common way to quantify the range of the errors. Both of these accuracy indicators can also be expressed as a percentage of the signal mean, which is how they are presented in the remainder of this section.

Table 3. Numerical results of one measurement chain

Time	Correct archive values	Simulated archive values	Error	
	(W)	(W)	(W)	(%)
11:30	1,843	1,843	0	0.0%
12:00	1,852	1,971	119	7.5%
12:30	1,927	2,052	125	7.9%
13:00	1,317	1,190	-127	-8.0%
13:30	1,233	1,261	28	1.8%
14:00	1,343	1,468	126	7.9%
Mean	1,586	1,631	45	2.8%
RMS			102	6.4%

When this simulation is repeated for a wide range of sampling and archive intervals, the mean percentage error and RMS percentage error can be plotted as shown in Figure 55 and Figure 56 respectively. Note that both horizontal and vertical scales are logarithmic in order to better illustrate the relationships over a wide range.

At first it might seem odd that the mean error in Figure 55 depends only on the sample interval, and not on the archive interval. That is true in this simulation set because each archive value is a mean value of a certain number of samples, and it doesn't matter how you group those samples to arrive at the overall mean. (Above the sample interval of 60 s there are a few connecting lines that do not coincide because certain sample intervals do not divide evenly into certain archive intervals and were therefore not simulated.)




Figure 55 Pac_{mean}, mean error

The feature that is most remarkable here is that as sampling rates grow smaller from 10 s down to 0.5 s, the mean error does not go down, but remains stable with even a slight upward trend. This effect is also seen in the RMS errors in Figure 56. Although the mean error remains less than 0.1% in this range, the RMS error is well over 1% in some cases.




Figure 56 Pac_{mean} , RMS error

Since the dominant source of high-frequency noise on the mean AC power was identified to be the anti-islanding test perturbations, it makes sense to see what happens when the 2.4 Hz harmonics are removed by an analog filter. Indeed, when a 2.4 Hz moving average filter is used both mean and RMS errors go down significantly. (Not shown.)

A more practical analog filter, a 2-pole low-pass filter with design frequency of 1.0 Hz has a virtually identical effect on the errors, and this result is shown in Figure 57 and Figure 58. The mean error for the entire affected sampling interval range from 0.5 s to 10 s is much improved and has become a continuation of the trend from longer sampling intervals. However the delay in the step response of this filter (see Figure 3) leads to somewhat larger errors at the highest sampling rates.




Figure 57 Pac_{mean} with 1 Hz analog low-pass filter, mean error




Figure 58 Pac_{mean} with 1 Hz analog low-pass filter, RMS error

Since the 1 Hz filter led to an optimal sampling interval of 0.5 s—where the minimum mean and standard deviations occur—it is reasonable to wonder whether the design frequency of this filter can be adjusted to minimize the error for a chosen sampling interval. Indeed Figure 59 and Figure 60 show that the use of a 0.1 Hz filter frequency minimizes the errors for a 5 s interval.




Figure 59 Pac_{mean} with 0.1 Hz analog low-pass filter, mean error




Figure 60 Pac_{mean} with 0.1 Hz analog low-pass filter, RMS error

In a final test we adjust the analog filter design frequency to twice the sampling rate for each simulation and see in Figure 61 and Figure 62 that this can significantly reduce the errors at any sampling rate. Both mean and RMS errors are reduced by roughly an order of magnitude. This means that this low-pass filter is also effective in dealing with those large, sudden swings of irradiance that are normally the main source of error at very long sampling intervals.




Figure 61 Pac_{mean} with analog low-pass filter adapted to sampling rate, mean error




Figure 62 Pac_{mean} with analog low-pass filter adapted to sampling rate, RMS error

The principle at work here is quite simple. From a time-domain perspective, the filter is pre-calculating quasi-average values so that the samples that are taken are no longer instantaneous values. It no longer matters whether a peak value falls between two samples since the filter output takes them into account.

From a frequency-domain perspective the filter is acting like an anti-aliasing filter. Having a filter with a cut-off frequency that is half of the sampling frequency, as above, corresponds exactly to the theoretical filtering requirement. The two-pole filter is certainly does not provide the perfect cut-off characteristics, but it nevertheless has a very positive effect on the archived values.

A low-pass filter like this could certainly be included in PV monitoring system designs where high sampling rates are not achievable, and it is also something that could be retrofitted quite easily as well—provided a suitable product can be located. Nevertheless, some further study would be warranted to explore the influence a broader range of filter types and design parameters using additional signals.

6.2 Discussion of Additional Simulations

Despite being able to condense the results of many individual measurement chain simulations in graphs such as those in the previous section, the breadth of the parameter range makes it impractical to present all our simulation results at the same level of detail. This section therefore presents the certain high-level observations based on those simulations.

First of all, there is very little to say about signals that vary slowly. Module temperature is a case in point where just taking one-minute samples and archiving them—without any analog or digital filtering—resulted in a mean error of only about 0.05 C and an RMS error of around 0.5 C. This result was obtained despite several large fluctuations between 30 C and 55 C during the three-hour period that was evaluated.

The effect of line-frequency noise on the mean and RMS error could be seen at sample intervals shorter than 1 s for signals where that was the predominant type of noise, such as DC current. And the corresponding improvement due to the addition the cycle-averaging line-frequency filter was also clearly visible in this range. However at longer sampling intervals the errors from not following rapid irradiance changes overshadow the former effect.

Among the analog low-pass filters, the two-pole version was observed in several situations to be more effective at reducing errors than either the single-pole or the four-pole version. The likely explanation for this is found in their step responses (Figure 3), where the two-pole version offers the best compromise between fast response and small overshoot. This trade-off could be fine-tuned by varying the internal filter parameters rather than using the values of the Butterworth design.

On the digital side five filtering options were examined, but most of the results are not very meaningful. Using no digital filter at all, in other words just archiving the most recent sample, is equivalent to making the sampling interval equal to the archive interval. This is one case that is already included in simulations involving other digital filtering options. The remaining four options are the single and dual-pass moving average (equivalent to the first and second-order CIC filter) applied over either a single or a double archive interval. Since the evaluation of the simulation chain errors is based on a comparison with the true average signal values over single archive intervals the single-pass moving average has a natural advantage. Indeed, using either a dual-pass or a double-interval average always produced very large errors.

This result does not mean that the dual-pass and double-interval averaging filters should be discarded, but rather that they should be evaluated against a more appropriate objective. Essentially we need to know more about how the data will be used and analyzed down stream in order to make an optimal selection for a digital filter or data reduction calculation. All three of these filter alternatives smooth the signal beyond what was expected, and this could be exactly what is required in certain applications.

7. REAL-WORLD OBSERVATIONS

7.1 Line Frequency Noise

It is common practice to suppress line frequency noise when taking almost any kind of measurement. But what can happen when, despite best efforts, a signal still contains some line frequency component, especially when measurements are taken many seconds or even minutes apart?

When the line frequency is *exactly* 60 Hz, and the measurement interval is an *exact* multiple of one second (or actually a multiple of $1/60^{\text{th}}$ second), then theory predicts that the 60 Hz will produce an alias frequency at 0 Hz or DC! The value of that DC component will depend on the phase relationship between the 60Hz signal and the sampling intervals. This is easy to visualize: if the first measurement is taken at a peak in the 60Hz waveform, so will the measurements that follow it, and all measurements will have an offset equal to the amplitude of the ripple.

In the real world, neither the grid frequency nor the sampling rate will correspond exactly to their nominal values. Moreover, the grid frequency fluctuates slightly over time (and perhaps the sampling interval as well), so the aliased version of that not quite uniform ripple becomes an unpredictably fluctuating measurement offset. Figure 63 shows a portion of the signal discussed in section 5.3.1, with samples taken every second. Around the 40 s mark the samples are taken in the troughs of the ripple, and around the 70 s mark they are taken at the peaks. Over the long term this produces a slowly varying measurement error.




Figure 63 Slowly varying offset in 1-second measurements caused by line frequency drift

Given the pervasiveness of line frequency electromagnetic fields around cables and components, it is important not only to implement countermeasures in PV monitoring systems, but also verify that these measures are adequate, and that line frequency noise is not compromising the accuracy of the measurements.

7.2 Campbell Scientific Line Frequency Noise Rejection

The Campbell CR800 and CR1000 (and perhaps other models as well) are able to perform an integration over the duration of one AC cycle in order to filter out 60 Hz or any multiples thereof. However this is possible only on the millivolt input ranges; on the ± 2.5 V and ± 5.0 V input ranges another technique is used. Instead of the full integration, two short samples are taken $\frac{1}{2}$ cycle apart, and it is readily apparent from the illustration in Figure 64 that the average of these two samples is independent of the amplitude of the AC waveform, and therefore effectively filters out 60Hz.




Figure source: [8]

Figure 64 Illustration of two AC power line noise rejection methods

Looking at the frequency response of the two methods (Figure 65), we see that these two methods produce nearly the same attenuation for frequencies up to 60 Hz, but beyond this frequency they differ considerably. Full cycle integration completely blocks both odd and even harmonics of 60 Hz and also mildly attenuates all frequencies between. The $\frac{1}{2}$ cycle technique on the other hand blocks only odd harmonics, passes even harmonics untouched, and provides very little attenuation for large bands of the remaining high frequencies.




Figure 65 Frequency response of two power line noise rejection methods

In many situations the $\frac{1}{2}$ cycle technique may provide sufficient power line noise rejection, but we have seen that even harmonics are present on various signals as well. These signals would usually be obtained through isolation amplifiers, therefore it would be important to use a product that also includes an appropriate level of filtering (which many do).

7.3 Anti-islanding Perturbations

During the signal analysis discussion (section 5) we identified several other periodic fluctuations that can result in measurement errors. In our example, the pattern resulting from anti-islanding operation had a fundamental frequency of 2.4 Hz, which makes 12 complete cycles in 5 seconds. Therefore if the measurement interval is an exact multiple of 5 seconds (or actually a multiple of $5/12^{\text{th}}$ second) then this periodic variation can also produce a DC offset. Or in practice, a slowly varying offset.

Figure 66 shows how measurement taken every 10 seconds on the AC RMS power signal (with line frequency and harmonics removed) do indeed exhibit errors that vary slowly over time, reaching as high as $\pm 6\%$! The alignment of the 10-second samples was manually adjusted in this example so that the errors would be large during the chosen five-minute period, but a different sample alignment would just move the largest errors to another period. Clearly if one were to use one-minute, or even five-minute averages from these measurements to calculate inverter efficiency, there would be some unrealistic values.




Figure 66 Slowly varying offset in 10-second measurements due to anti-islanding ripple

Sampling a signal once per second transforms all high-frequency content into alias signals in the range 0 to 0.5 Hz (the Nyquist frequency). A set of evenly-spaced harmonics, like those observed in Figure 45 will look quite different after they are all aliased into that range. In this case the 2.4 Hz fundamental frequency produces an alias frequency of 0.4 Hz, but second harmonic at 4.8 Hz produces a lower alias at 0.2 Hz. In fact, with these exact frequencies and ratios all the higher harmonics would all alias to either 0.0, 0.2 and 0.4 Hz. With real-world signals the aliased spectrum would look a bit more blurred, but in the time domain the 5-second repetition period of the new 0.2 Hz fundamental frequency is very visible. (Figure 67)




Figure 67 Aliasing of anti-islanding ripple with one-second measurements

Note, however, that some of the higher harmonics also alias to 0 Hz, which means there is the possibility of a DC offset again. And indeed looking at the whole 5-minute period in Figure 68, it is clear that samples taken at 1-second intervals exhibit also a slowly varying measurement error. There is a definite improvement compared to the 10-second samples, but still the possibility of a $\pm 1.5\%$ error is still significant.




Figure 68 Aliasing causing slowly varying offset in average of 1-second measurements

Anti-islanding can be more difficult to cope with than line-frequency noise. IEEE 1547 requires that the inverter “detect the island and cease to energize the Area EPS within two seconds of the formation of an island”, but does not prescribe a method. This means that the magnitude and frequency of perturbations (if any) can vary from system to system, making it hard to design a

universal filter. And with perturbations potentially occurring as much as two seconds apart, they can become hard to distinguish from signal fluctuations caused by changing irradiance.

In a different PV system using an older three-phase Xantrex PV20-208 inverter we observed that the anti-islanding detection perturbed the DC current and DC voltage at 0.5 seconds intervals. Because the monitoring hardware for this system uses sampling intervals of 1.1235 s for the current signals, and 1.2280 s for the voltage signals, the aliasing effect produced patterns recurring at different intervals as shown in Figure 69 and Figure 70. What's more, the perturbations are not symmetrical in this system: DC current briefly goes down, while DC voltage briefly shoots up.




Figure 69 Aliasing causing periodic groups of peaks in current signal




Figure 70 Aliasing with anti-islanding test causing frequent isolated peaks

The solution to this problem lies in designing filters not based on the signal content, but based on the measurement and archiving intervals. These should be known quantities in the monitoring systems. As was shown in section 6, a low-pass filter with a cut-off frequency of half the sampling rate was very effective at reducing both mean and RMS errors of the archived values.. This does not come as a surprise, as this is just the general anti-aliasing filter requirement applicable to any sampling of any signal.

The underlying assumption with many data loggers is that fast signal variations are both small and irrelevant, and they therefore they do not contain this type of anti-aliasing filter. We have demonstrated that for PV monitoring systems these assumptions do not necessarily hold; and when they do not, filters are indeed beneficial.

8. CONCLUSIONS

In this study we have investigated how the interplay between complex signals and the way in which they are measured can lead to errors in the data delivered by PV monitoring systems.

To assess the signal complexity we acquired data of the most important signal types from a grid-connected PV system and associated weather instruments using a very high sampling rate (2,000 samples per second). In this high-resolution view of the signals we were able to observe various rapid fluctuations and we found explanations for the origins of many of them. The most prominent fluctuations were found to be caused by the active anti-islanding system, which caused periodic changes of AC current and power at intervals equal to 25 line cycles. The effect of this was also seen on the DC voltage and current signals.

The signal characteristics we report were observed over relatively short time intervals on a single PV system, which means that they certainly cannot be considered as representative of all systems. But the fact that they were all found on the first system we examined, suggests that signals in other systems are likely also affected, although maybe not for the same reasons and not in the same proportions. Further evidence needs to be gathered from other operating PV systems before more general statements can be made.

To explore how measurement errors can arise in PV monitoring systems, we simulated their operation using a wide range of sampling intervals and archive intervals, and using several different filtering options. We saw how the anti-islanding system perturbations dominated the measurement errors in AC power over a specific range of sampling rates, and found that a simple two-pole low-pass filter preceding the analog-to-digital conversion could be tuned to reduce those measurement errors. Furthermore, we showed that the low-pass filter could be tuned to reduce the measurement error at any sampling rate. As PV monitoring systems often have sampling rates that are too low to capture rapid fluctuations in irradiance and power, the addition of a low-pass filter presents itself as a possible solution to obtaining more accurate average values for many signals.

Because there is a strong link between various parameters of a monitoring system and the quality of the archived values, is very important to know what those parameters are. Documentation for PV monitoring systems should always include specifications of filtering methods, sampling rates, summary calculations, archive rates and time stamp conventions, and this metadata should always accompany the data files that are produced.

The present study focused on individual signals and the accuracy of the archive values that were produced. Future work should explore the accuracy of the relationships *between* signals, which are the key to understanding and assessing system performance. The true relationships between signals can be distorted by excessively long averaging periods, by taking sequential measurements rather than simultaneous ones, or by unequal response rates somewhere in the measurement system. We expect that a better understanding of all these effects will lead us to more accurate PV system performance assessments.

9. REFERENCES

1. Yordanov, G.H.; Midtgård, O.-M.; Saetre, T.O.; Nielsen, H.K.; Norum, L.E., *Over-irradiance (Cloud Enhancement) Events at High Latitudes*, IEEE Journal of Photovoltaics, vol.3, no.1, pp.271,277, Jan. 2013.
2. Burger, B and Ruther, R. Inverter sizing of grid-connected photovoltaic systems in the light of local solar resource distribution characteristics and temperature. Solar Energy, 80 (2006), pp. 32–45.
3. Luoma, J, Kleissl, J., Murray, K., *Optimal inverter sizing considering cloud enhancement*, Solar Energy, 86 (2012), pp. 421–429.
4. Mokri, J. and Cunningham, J., *PV System Performance Assessment*, SunSpec Alliance, www.sunspec.org, Version 2.0, June 2004.
5. IEC, *IEC 61724 Std. Photovoltaic system performance monitoring-Guidelines for measurement, data exchange and analysis*, Int. Electro-Tech. Comm., 1998.
6. Blaesser G. and Munro D., *Guidelines for the Assessment of Photovoltaic Plants Document A Photovoltaic System Monitoring*, Joint Research Centre - Institute for Systems Engineering and Informatics, EUR 16338 EN, 1995.
7. Hogenauer, E.B., *An Economical Class of Digital Filters for Decimation and Interpolation*, IEEE Transactions on Acoustics, Speech, and Signal Processing, Vol. 29, No. 2, April 1981.
8. Campbell Scientific, *CR1000 Measurement and Control System – Operator’s Manual*, Campbell Scientific Inc., Revision 10/10, 2010.

10. DISTRIBUTION

- 1 MS0899 Technical Library 9536 (electronic copy)
- 2 This report available electronically on Sandia National Laboratories Photovoltaic web page: pv.sandia.gov

