

Maximizing Inverter-Based Generation Performance for PV and Storage Generation Facilities

Brian Waldron
Schweitzer Engineering Laboratories, Inc.

Presented at the
2026 PAC World Africa Conference
Johannesburg, South Africa
April 20–23, 2026

Originally presented at the
60th Annual Minnesota Power Systems Conference, November 2024

Maximizing Inverter-Based Generation Performance for PV and Storage Generation Facilities

Brian Waldron, *Schweitzer Engineering Laboratories, Inc.*

Abstract—This paper examines three topics related to performance for inverter-based generation: selecting system components and designing communication systems, analyzing and identifying individual component limitations within the system, and ascertaining the type of performance that can be expected from generation facilities with well-designed system components compared to typical equipment and methods.

Many engineers at inverter generation facilities do not coordinate the communications infrastructure, selection of the power plant controller, site metering resource, selection of protocols, and configuration of inverter response behavior together early in the design process with the intention of meeting specific performance criteria. Many interconnection agreements and regulating bodies have well-established criteria for generation sites responding to a set-point change or reaching a desired set-point change within a certain time period. This paper covers the process of identifying performance metrics, then selecting the best methodologies necessary to meet those metrics with a focus on selecting the correct infrastructure to implement those methodologies.

Often site engineers only discover limitations in their design choices when commissioning the site while trying to achieve those performance requirements. Addressing these limitations can lead to increased site commissioning time, the potential need to replace existing equipment or change the settings of devices that are already commissioned. Alternatively, one part of the system may be limiting the performance of the rest of the system. For example, the deadband for the real and reactive power measurements on the revenue meter may be set too high to be reported in the one-second DNP3 class poll. This setting prevents the controller from receiving the necessary feedback to make correct control decisions. This paper focuses on the process used to find the limitations in the system and provides recommendations on how some of these limitations can be accommodated through settings changes.

Provided a system has the necessary infrastructure to support data update rates faster than 500 milliseconds (ms) (a common limitation for systems with standard supervisory control and data acquisition [SCADA] protocols and communications infrastructure), how fast can set-point targets be achieved? Is sub-100 ms possible? This paper looks at modeled data to identify the type of performance that could potentially be achieved with additional communications equipment and a system designed for performance.

This discussion prepares power system engineers to effectively troubleshoot their inverter generation facility performance concerns, identify common resolutions and strategies to improve performance, and balance the improvements in system performance against the cost considerations of the additional equipment.

I. INTRODUCTION

A power plant controller manages the set points for individual generating resources to create a collective output that achieves a desired target for the generation facility. This paper focuses on the performance of inverter-based generation. Inverters are commonly used in photovoltaic (PV) and storage generation facilities. It explores various factors that affect the overall response of the generation facility. The generation facility contains several components in addition to the generators; collectively, these generators and components will be called a system.

In a power plant controller, the performance of the system is typically measured by the time that elapses between when a set point is received by the power plant controller and when that value is achieved at the measurement point. That time is composed of several different system components, which depending upon the site, can significantly impact that overall time. If the target goal is to change a voltage set point in 3 seconds (s), but it takes 2 s to get the information from the meter to the controller and another 2 s to get the data from the controller to the inverter, it does not matter how quickly the controller calculates new set points, that 3 s target cannot be achieved since the overall communications transmission time is 4 s. This is an easy position to end up in. The meter may be polled every 500 milliseconds (ms) for changed values, but there is an incorrectly configured deadband that requires the monitored value to change by a large amount, such that the present value only gets reported via an integrity poll every 2 s. This is a case where it may appear the controller should be getting the data much faster, but due to a settings issue a significantly longer delay is introduced. The inverters may be configured on a low-bandwidth RS-485 communications network where the 30 inverters share one communications line. Each poll takes 50 ms to respond. The time it takes to get a present status from each inverter at 50 ms that adds up to is 1.50 s. If controls need to be sent to each inverter, the polling rate takes longer overall. Thus, the average time to interweave controls and polls across all the inverters may easily exceed 2 s. There are changes to the settings that can take place, and perhaps some rewiring that can occur, to bring the system performance within the target specification in this example. But without understanding the time each part of the system takes, it is difficult to troubleshoot or make effective changes to the system to improve the overall response time.

Determining the overall system response time is a practice that should occur in power plant control. Once that practice is complete, the end users must allocate a portion of that response time to each part of the system and then make technological decisions to support that time budget. Once onsite, it is important to verify that each part of the system matches the time budget that was allocated. This paper explores the potential design process and key considerations for the essential components that make up the control system. By accounting for each component in the design process, unexpected surprises can be prevented once on site while reviewing the performance of the generation facility.

II. IDENTIFYING SYSTEM PERFORMANCE REQUIREMENTS

Many generation facilities have a variety of requirements from different stakeholders that must be met. Regulatory requirements from independent system operators (ISOs) [1] may dictate IEEE 2800 compliance, or the utility may impose different response characteristics. Understanding these requirements before developing a time budget for the control system is critical. Some regulatory bodies [1] make a distinction between completing a set-point change versus starting a response. There are many instances where there are requirements that the system starts changing in one or two seconds; this requirement is quite different from achieving the total system response. Identifying the target goal for the total response time or the start of the response is important for the system design. However, determining these requirements prior to selecting equipment and design is not always done. To ensure there are no surprises when the site is deployed and commissioned, it is important to identify these requirements before the selection of equipment and design.

III. WHAT IMPACTS PERFORMANCE?

The components that impact system performance the most include the metering source, the time it takes to move data from the meter to the controller, the controller processing interval, the time it takes to move data from the controller to the inverter, and the time it takes for the inverter to implement the set-point change. The time associated with moving the data between devices is significantly impacted by three aspects: the communication medium (fiber or copper), the transport mechanism (serial or Ethernet), and the protocol (DNP3, Modbus, C37.118, etc.). The paper describes each of these system components and their nuances in later sections. Fig. 1 shows the order and relationship between system components.




Fig. 1. System Components

A. Metering

The meter's sampling rate is typically not a concern in power plant control applications. The controller benefits from a sample rate of 0.2 Hz to 10 Hz. However, when considering the metering source, it has a tendency to overlap with the communications component. While the meter may sample the

power system at a significantly faster rate, it may only update the root-mean-square (rms) quantities to the communications protocols between 1 Hz and 10 Hz. The meter may respond to data requests every 50 ms. But if it is only capable of providing new information to the communications interfaces every 500 ms then the meter is always adding at least 500 ms delay to the system response. This delay is something that typically cannot be managed by the end user, but communications protocol may have some impact over the time delay. Though it cannot be managed by the end user, it is an important aspect of understanding the overall time budget. Meter documentation should be able to provide this information.

B. Communications Infrastructure From Meter to Controller

The communications infrastructure from the meter to the controller is an important aspect to take into consideration. Because the information from the communications infrastructure is typically fed into the closed-loop control logic in the controller, the controller decides whether more or less watts or VARs are needed in the system. The rate at which the controller can calculate new set points cannot be faster than the time it takes for the meter to sample the power system and then deliver those data to the controller. Out of each individual portion of the system, this communications infrastructure typically has the largest impact on overall performance. The choices in this category often have financial implications for the project. Copper is frequently cheaper and easier to work with and often handles the speed necessary for power plant control. However, if there is significant electrical noise in the environment then data polls may get corrupted and no communications will be possible at all. Using fiber eliminates electrical noise but is a more expensive material and has higher installation costs. The decision to select copper or fiber should be driven by other site considerations. The impact on performance is negligible as long as there is no data corruption in the copper due to electrical noise.

Ethernet may have a reputation for being faster than serial communications, it typically allows for large amounts of data to be passed through; however, power plant controllers typically only need real power, reactive power, voltage, and frequency. Given there is only a small amount of data that can be efficiently transferred, using serial communications may be a suitable choice. It is important to take into consideration other factors when choosing between serial communications and Ethernet. Generally, speaking, Ethernet is the standard solution, but serial communications may be a valid solution for the power plant controller if there is a sufficient baud rate above 38400 bps. Most systems are able to take advantage of the other benefits of Ethernet that are not discussed in this paper. It should simply be noted that serial communications with sufficiently high baud rates is capable of delivering acceptable performance for power plant control applications.

The communications protocol will potentially have a large impact. As discussed in Section III.A., by using protocols, such as DNP3, Modbus, or Manufacturing Message Specification (MMS), the meter may only update the value passed to the communication protocol on a certain time interval. So, even if

the meter is capable of responding at 50 ms via Modbus, the data value may not change for several iterations. The Synchrophasor Protocol IEEE C37.118 was designed to operate at higher speeds. Many devices which implement this protocol offer 60 Hz phasor data or potentially even up to 240 Hz [2]. This can provide updates to the power plant controller between approximately 4 ms and 16 ms. The responsive capabilities of the inverters may be exceeded when using data rates greater than 5 Hz in 200 ms intervals. And because the Synchrophasor Protocol and the associated standard require the data values to update at these intervals [2], there is not the same concern as there is with DNP3 or Modbus in terms of update frequency [3] in the meter-to-the-communications protocol. The primary downside of synchrophasors is the cost of implementing them. Synchrophasors require precise time alignment [2], which typically requires a GPS clock that delivers a better than 1 us accuracy signal to the meter and power plant controller. The generation facility may or may not have this infrastructure already in place. A dedicated Bayonet Neill-Concelman (BNC) cable may need to be run to the meter to deliver the IRIG-B signal. Alternatively, Precision Time Protocol (PTP) may be used to deliver the time accuracy needed through the Ethernet network. Both Synchrophasors and PTP require an Ethernet network with sufficient bandwidth to deliver the data reliably. This may require additional communications infrastructure or additional engineering to utilize these options. But they do provide the highest data delivery capability to the portion of the system that typically tends to have the largest limitations.

In summary, for the best performance possibilities and to eliminate any possible bottlenecks in this category, selecting fiber, Ethernet, and synchrophasors with a data rate of 10 Hz (10 msg/sec) or greater delivers the best performance possible [3]. This decision typically allows the total response times for the system to be well under 1.5 s. Response times under 300 ms are possible with this infrastructure. With this communications infrastructure, the next likely bottleneck is the inverter response time itself. But this is the most expensive option in terms of hardware, engineering design, and settings configuration.

A more balanced approach, which is significantly cheaper, is to use either copper or fiber wires (depending upon site conditions), Ethernet, and Modbus. This setup does not require a GPS clock to coordinate time between devices, instead, a very simple Ethernet network can be deployed with minimal engineering. Modbus is a good protocol to use because it ensures that the present value is returned. There are no deadband settings to configure in the meter or any other protocol settings that may be configured incorrectly, leading to an assumed data update rate of a given interval, but in reality, also a much slower interval due to misconfigured or unknown settings in the protocol implementation. The simplicity of using a Modbus protocol helps ensure greater performance. For systems which deploy this approach, there is often total system response times in the 2 to 5 s range with the initial response time in under 1.5 s [3]. There are many commissioned systems that meet regulatory requirements with this infrastructure approach.

The counterargument to using Modbus compared to DNP3, MMS, or many other supervisory control and data acquisition (SCADA) related protocols would be bandwidth consumption. Even if the value has not changed, those data will be returned, consuming more bandwidth. This would be a valid concern if the power plant's control applications required large amounts of data. At a bare minimum, most power plant control applications need updates for voltage, real power, reactive power, and frequency. Assuming these are 32-bit values, that is a total of 128 bits. The potential bandwidth savings over a localized communications network are not worth the potential settings misconfiguration that may occur with DNP3, MMS, or other SCADA protocols that use deadbands or exception reporting.

C. Controller Execution Rate

The execution rate of the controller is the least likely to be a limiting factor in the performance of the system. Nonetheless, it is a good parameter to understand. The controller's execution rate is not the rate at which new set points are calculated. It is the rate at which the controller processes logic and samples data from its communications protocols. For example, if the controller execution rate is 200 ms but the desired rate to calculate new set points is 300 ms, the actual rate will be adjusted to 400 ms because 300 ms is not evenly divisible by 200 ms. This allows it to execute the set-point calculation logic in the following execution scan. As your rate to calculate new set points becomes smaller, making sure the execution rate of the controller is a value that is divisible is an important factor.

D. Communication Infrastructure From Controller to the Inverter

Compared to the communications infrastructure between the meter and the controller, the communications infrastructure between the controller and the inverters do not have as big of a performance impact with one notable exception. The reason this portion of the system does not have as large of an impact is because this portion of the system is primarily used to distribute control set points to individual inverters. The metering feedback from individual inverters typically does not have as significant of an impact on the closed-loop control decisions. When the controller calculates a new set point for the inverter, it is usually distributed on demand. Since it is a short message, many communications protocols that are implemented recognize it as a priority message and will temporarily pause polling requests to the inverter to send the new set point. Aside from one exception, controls are distributed from the controller to the inverter in less than 100 ms regardless of communications protocol. Some protocols are capable of distributing controls faster than others. However, many inverters do not offer a wide selection of communications options. Many only support Modbus, and thus, a detailed conversation comparing benefits of the different protocols for control distribution is not discussed in this paper because inverters usually do not provide those options.

The most important consideration in the communications infrastructure from the controller to the inverter is to ensure that all inverters have their own unique communications session. If inverters have to share a communications session where the controller needs to wait to send controls or request information for an inverter before sending those data to a different inverter, this slows down the performance of the system. Sharing a communications session between inverters is the exception in which the communications infrastructure significantly impacts system performance. This type of setup is usually referred to as a multidrop communication line or party line. It is most common in serial-based communications systems that use EIA-485. For inverters that support an Ethernet communications option this is typically not a concern. While Ethernet communications to inverters seems to be the trend in new system deployments, it is an important aspect to take into consideration in the design of the system. If the inverter only supports serial communications, it is highly recommended that the system design use dedicated communications channels from the controller to the inverters. This can increase the cost of the communications network by running individual connections from the controller to the inverters, but the cost in material and installation is worth the performance benefits. Depending on the quality of the multidrop communication this can increase the distribution of set points from less than 100 ms up to seconds. This issue scales with more inverters added on a multidrop line. If 30 inverters share a line and each inverter takes 100 ms to distribute a single round of set points from the controller, that would be 3 s. The controller cannot calculate new set points until the distribution of set points is complete. This example would require the controller to be limited to calculating new set points greater than 3 s. This could easily become the largest bottleneck in the system that no settings could adjust.

The metering data from the individual inverters is sometimes used in control algorithms but typically does not transmit feedback directly into closed-loop control calculations, so it does not need to update as quickly as the primary metering connection. Faster updates are always beneficial, but most systems benefit from metering updates from the inverter in 3 s or less. If the communications infrastructure can support it, a good target would be between 500 ms to 1000 ms (1 to 2 Hz).

E. Time for Inverters to Implement Set Points

This paper does not discuss the typical performance of time it takes to implement a set point that an inverter has received. However, various site deployments experiences, settings and software revisions, and new hardware show a variety of response times can be achieved, which in some circumstances have limited the rate at which new set points can be calculated for the system by the controller. For example, if an inverter takes 2 s to start responding to a set-point change then the controller should not calculate new set points faster than the rate at which the inverter is able to respond to the set points. It is

important to understand what the response characteristics of an inverter are. Because the response characteristics can impact the rate at which new set points can be calculated it is important to understand the response characteristics. Examining these characteristics from three perspectives is helpful. The best way to identify these characteristics is to place the controller into an open-loop mode where a single set-point change is issued to the inverter and the response of individual inverters can be examined.

First, how long does it take for the inverter to start changing its output after a new set point has been received? This paper does not examine or discuss typical inverter behavior, such as how long an inverter takes to start responding to a set-point change. However, site deployment experience has shown in occasional circumstances a delay that was longer than the rest of the system bottle necks.

Second, how long does it take for the inverter to achieve the set-point change?

Third, how accurately does the inverter settle at the set point? This is identified by the overshoot of the target set point, oscillations around the target set point, or even ramp response from 0 to 80 percent of the set-point change, then a significantly slower ramp from 80 to 100 percent of the set-point change. If the inverter's response consistently displays issues settling in at the target set point, this may impact the rate at which new set points for the site are calculated. Or it may require tuning or other adjustments of the closed-loop control for the overall site.

While addressing these types of behavior would ideally be done in the inverter to create a smoother response, that is not always possible, and the power plant controller may need to adjust instead of relying on a smoother response from the inverter. While the exact solution needs to be considered for each individual site, the most typical adjustments are to increase the time in between how often new set points are calculated by the controller to allow the inverter more time to settle before asking the controller to make more set-point changes, or to have more conservative tuning parameters if using proportional integral derivative (PID) closed-loop control.

IV. HOW TO PICK AN EVALUATION PERIOD AND EXPECTATIONS FOR SYSTEM RESPONSE

So far, this paper has discussed a wide variety of factors that affect the ability for the controller to execute set points in a timely manner. These factors restrict what rate set points can be calculated. But a user armed with the information discussed needs to be able to identify what the evaluation period should be. The evaluation period is the rate at which new set points are calculated at. When looking at the five different parts of the control system, usually the evaluation period can be set to a slightly larger value than the largest time. For example, if we had the following characteristics, as shown in Table I, we would pick the following evaluation period.

TABLE I
EXAMPLE EVALUATION PERIODS FOR DIFFERENT SYSTEM CHARACTERISTICS

	System 1 (ms)	System 2 (ms)	System 3 (ms)
Meter	500	16	500
Meter-to-Controller Communications	500	16	3,000
Controller Task Interval	100	8	50
Controller-to-Inverter Communications	100	50	100
Inverter Response Time	200	100	100
Evaluation Period	600	150	3,200
System Response Expectations	1,200 to 6,000	300 to 1,500	6,400 to 32,000

As shown in Table I, the evaluation period is set to slightly above the largest bottleneck in the system. The evaluation period will not be the same as the overall response time of a set-point change. Depending on the steady-state losses and tuning of the system, the overall system response will likely be between 2 and 10 evaluation periods. The best way to calculate your evaluation period is to identify the desired overall response time and divide that response time by 2 and 10. This is because the controller will likely take multiple steps to achieve the target set point. This process will take multiple evaluation periods, thus causing the overall system response to be a multiple of these evaluation periods. Tuning the closed-loop control will significantly influence how many evaluation periods are needed to achieve the desired system response.

That will give you a range of evaluation periods that will be necessary to meet that performance time. Then, compare that calculated evaluation period against the time for each part of the system response. If the desired evaluation period is smaller than any of the system component times, then there is an incompatibility between system capabilities and performance expectations. In this circumstance, either settings or design changes need to be made to ensure the calculated evaluation period is larger than any of the individual system components. Alternatively, take the largest system component, make that the evaluation period, and have the expectation that system performance will be 2 to 10 times the evaluation period. Based upon all site experience, 2 to 10 times are the extreme cases. Based upon the average site experience, a typical response is usually within 3 to 5 times the evaluation period.

V. WHAT HAPPENS WHEN EXPECTATIONS ARE MISALIGNED WITH SYSTEM CAPABILITIES

When a system is not meeting the target performance there are two primary aspects that can be modified to adjust the performance of the control system: tuning parameters and the evaluation period of the controller. The best parameter to modify is the evaluation period. This causes the controller to calculate new set points for assets more frequently. But at a certain point the evaluation period will become faster than the limitation of the system. When this occurs, there are generally

two characteristics that occur in the overall system response: either there is an overshoot of the target set point every time a set-point change occurs, or the system oscillates constantly or several times, slowly dampening to the target set point. When this occurs, there are two options: either calculate new system set points at a longer interval, or identify the bottleneck in the system components, as previously discussed, and reduce that bottleneck to be below the evaluation period. When set points are calculated faster than a system's bottleneck, the power plant controller calculates a set point, and sees that the response value has not changed from the last time a set point was calculated. The power plant controller asks for additional power to bring the site response to the set point. The controller cannot tell if the system has losses that need more power or if the data have not had a chance to update yet. To the controller, these two situations look identical. As a result, the controller increases the set point. This is standard behavior for a PID-based controller. In this case, the set point has increased twice as large before the inverter has had a chance to respond or before the response data has had a chance to make it back to the controller. This double (or more, depending upon settings) accumulation of error in the set point causes overshoot or oscillations in the response. Changing the tuning parameters can reduce the magnitude of overshoot or oscillations but, fundamentally, the only solution is to either increase the time of the evaluation period or change the system to get updated response information back to the controller before a new set point is calculated. There are examples of what this behavior looks like in the following sections.

In Fig. 2 there is a set point and a response for a single inverter, which both oscillate. This is a case where the inverter response does not change or decrease when the set point for the inverter increases. In this case, the inverter does not respond to several set points. As the controller accumulates more errors and requests for a larger set point, the inverter suddenly reacts, causing an overcorrection. Then as the controller attempts to bring the site closer to the target set point, the inverter lags behind the controller's set points once again. However, when looking for smaller set-point changes the inverter takes longer to respond to the set-point changes than the controller's evaluation period. To resolve this behavior, either the inverter settings need to be adjusted or the evaluation period of the controller can be set to a longer interval to allow the inverter to fully respond to the set point, preventing the controller from building additional error.




Fig. 2. Delayed Inverter Response Causes Oscillation

Fig. 3 highlights a case in which the system settles at the target set point but there is a significant overshoot, which slows down the evaluation period and eliminates the overshoot.




Fig. 3. Delayed Inverter Response Causes Overshoot

In both Fig. 2 and Fig. 3 the tuning could reduce the overshoot or the oscillations, but it would also have the effect of significantly slowing down system response as well.

VI. WHAT ARE REASONABLE EXPECTATIONS FOR SYSTEM PERFORMANCE?

This paper discusses many different aspects of a generation system that affect the performance of a set-point change. However, what are reasonable expectations for system performance? This paper shows that system performance can vary significantly. But it generally demonstrates that more planning prior to construction and installation of the control, metering, and communications equipment can deliver better results. Here are some estimates to create high-level expectations based upon design effort:

For a system with:

- Ethernet network and DNP3 between the meter and controller
- Serial communications and Modbus between the controller and inverters with independent communications channels
- Little to no consideration of detailed settings or parameters

The system will likely be able to achieve set-point changes between 3 s and 10 s [3]. The evaluation period is likely to be between 1 s and 2 s.

For a system with:

- Ethernet network and Modbus between the meter and controller
- Ethernet network and Modbus between the controller and inverters with independent communications channels
- Some consideration to settings and parameters

The system will likely be able to achieve set-point changes between 1 s and 5 s [3]. The evaluation period is likely to be between 750 ms to 1.5 s.

For a system with:

- High-bandwidth Ethernet network and synchrophasors with a 30 Hz update rate between the meter and controller

- Ethernet network and Modbus between the controller and inverters with independent communications channels
- Detailed consideration to settings and parameters

The system will likely be able to achieve set-point changes between 250 ms and 2 s [3]. The evaluation period is likely to be between 50 ms and 200 ms. In these systems, the bottleneck will be the inverter's response characteristics.

VII. CONCLUSION

This paper has gone into detail about what requirements are needed for a system to perform well while maximizing inverter-based generation performance for PV and storage generation facilities. By identifying these performance criteria, the user will be better able to avoid complications when commissioning a site. The six main components that impact system performance are provided in detail. The paper then showcases the variety of factors that influence the ability of the controller to affect set points in a timely manner. This information leads the user to be able to identify what evaluation period will meet stakeholders' desired performance time.

In addition to discussing the correct steps to take to align the system capabilities with stakeholder expectations, the paper also goes into what happens when expectations are not aligned. This misalignment can cause system bottlenecks, and the paper provides suggested tuning parameters and the evaluations of the controller to achieve ideal set-point changes. The paper concludes with estimates to create high-level expectations based upon designing a control system considering the topics discussed, which will allow the commissioning of the power plant controller to be more successful and meet system expectations.

VIII. REFERENCES

- [1] AEMO and Western Power, "Generator Performance Guideline," September 2019. Available: westernpower.com.au/media/3457/generator-performance-guideline-20190918.pdf.
- [2] IEEE Std C37.118.1-2011, *IEEE Standard for Synchrophasor Measurements for Power Systems*.
- [3] D. Anichkov and B. Waldron, "Case Study: Phasor-Based Control for Managing Inverter Generation," proceedings of the 57th Annual MISYCON, St. Paul, MN, November 2021.

IX. BIOGRAPHY

Brian Waldron is a senior automation engineer with Schweitzer Engineering Laboratories, Inc. (SEL). He has over a decade of experience in designing and troubleshooting automation systems, communications networks, and distributed energy resource (DER) control systems. He has authored several technical papers, application guides, and teaching presentations focusing on integrating automation products. Brian graduated from Gonzaga University with a BS degree in electrical engineering.