Quantifying Shading's Economic Impact

By Paul Grana and Paul Gibbs

Charts and illustrations courtesy Folsom Labs

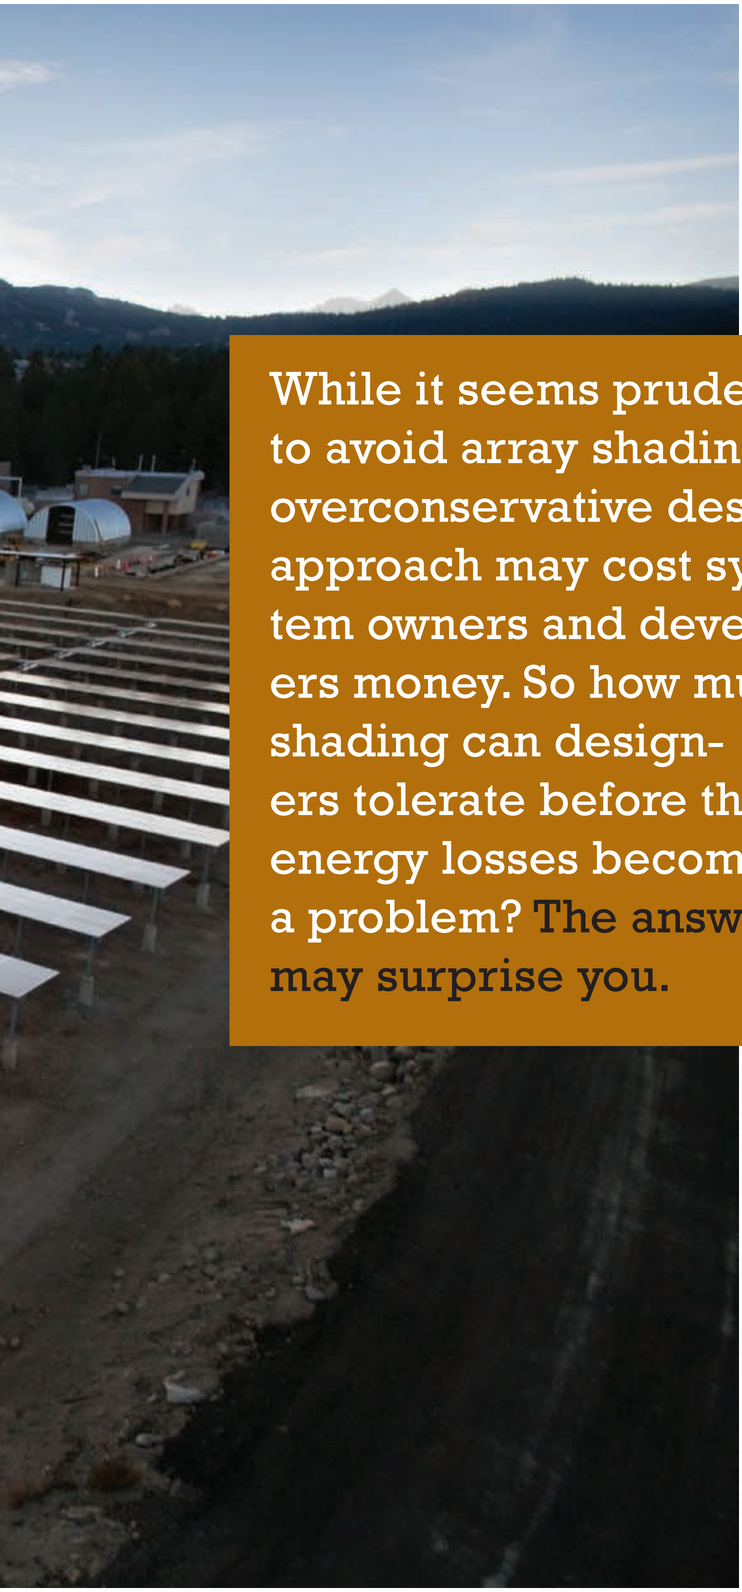

Courtesy SPG Solar

While it seems prudent to avoid array shading, an overconservative design approach may cost system owners and developers money. So how much shading can designers tolerate before the energy losses become a problem? The answer may surprise you.

Project developers and engineering firms typically approach array shading according to rules of thumb that are based on low tolerance for shade or that seek to avoid it altogether. For example, the system designer might eliminate modules if they are shaded during a specific window of time, or if they are located within a certain proximity to an obstruction. These conservative design approaches are largely based on assumptions that have lost their relevance.

First, traditional design approaches for dealing with shading developed at a time when modules were the most expensive components of a PV system. Therefore, it made sense for system designers to prioritize production efficiency. Second, for many years no software tools on the market were capable of calculating the actual energy production and mismatch effects of shaded modules. As a result, designers could not determine which modules to include or exclude based on an evaluation of economic performance at the system level.

Today, module prices make up a smaller percentage of total project costs, having fallen by approximately 60% over the last 4 years, and system designers have access to new software tools that can calculate the production of shaded modules. In light of these changes, it is worth reevaluating traditional design approaches to array shading, and considering instead a cost-benefit approach that looks at component costs in relation to potential revenue.

Here we explore the system-level effects of shade to better understand optimal design approaches to array shading. We first consider common shade types and traditional design approaches for dealing with the system-level effects of shade. We then discuss energy losses associated with shading and consider the results of detailed shading analyses performed using simulation software tools. Finally, we present the results of a cost-benefit case study, identifying the optimal amount of shade tolerance for a space-constrained PV system based on specific shade profiles.

Shade Types and Frequency

Shade impacts PV systems of all sizes, and a wide variety of obstructions can create shade. While many system designers think of shading as a problem confined to residential systems, obstructions are often present on and around commercial rooftops, as well as within and around ground-mounted arrays. To understand shading losses, we can classify shade into three categories: self-shading, near-object shading and far-horizon shading.

Self-shading. The most common example of self-shading is row-to-row shading, which results when a tilted row of modules shades an adjacent row of modules. Because adjacent rows of tilted modules are typically located close to one another, row-to-row shading can affect the system throughout the entire year. If the system designer does not anticipate and account for self-shading, it can have a large impact on system production.

Near-object shading. This type of shading results when objects directly shade the array, causing shadows to move across it. In some cases, the obstructions responsible for near-object shading, such as rooftop units, vents or parapet walls, are relatively close to the array (within 10 feet). However, obstructions located in the middle distance (up to 100 feet), such as trees, utility poles, water towers or nearby buildings, can also cause near-object shading. Obstructions that are farther away typically affect the array only during certain times of day, whereas shade from nearby objects is more persistent.

Far-horizon shading. This type of shading results from obstructions on the far-horizon line and impacts the entire array. Mountain ranges or city skylines are typical examples of far-horizon shading. Because these obstructions are so far away, they are generally assumed to shade all modules at once whenever the sun is below the horizon line. As a result, there is no specific design or engineering response to far-horizon shading. However, designers do need to calculate its impact on total system performance.

TRADITIONAL DESIGN RESPONSES

PV system designers and developers typically take a minimum-tolerance approach to array shading. One of the most common design standards is to remove any modules that are shaded between 10am and 2pm on the winter solstice. Alternately, designers might adhere to setback rules based on a distance multiple. For example, depending on the site latitude, they might set modules back from an obstruction at a distance of two or three times the object's height.

System designers specifically apply these design standards in response to self-shading and near-object shading. In the case of row-to-row shading, designers typically calculate the precise minimum interrow spacing based on the module size, tilt angle and site latitude. For near-object shading, designers usually translate objects into an array exclusion area based on object size and site latitude.

Except where fire or building codes require array exclusion areas that are deeper than shading setbacks, as might be the case with skylights on a commercial rooftop, traditional design responses to near-object shading mean that each object reduces the system's peak power capacity. These design standards do not necessarily result in optimal system design in terms of economic performance.

System-Level Effects of Shade

To optimize PV system designs in terms of shade tolerance, it is important to understand the effects of shade at the system level. System design factors, such as module construction and PV source-circuit performance characteristics, determine these effects, as do the specific components of sunlight.

System design factors. Crystalline silicon (c-Si) modules are typically composed of 60 or 72 series-connected solar cells, and six to 20 c-Si PV modules are generally connected in series to form a PV source circuit. If a single cell is shaded so that its maximum current is less than the

CONTINUED ON PAGE 38

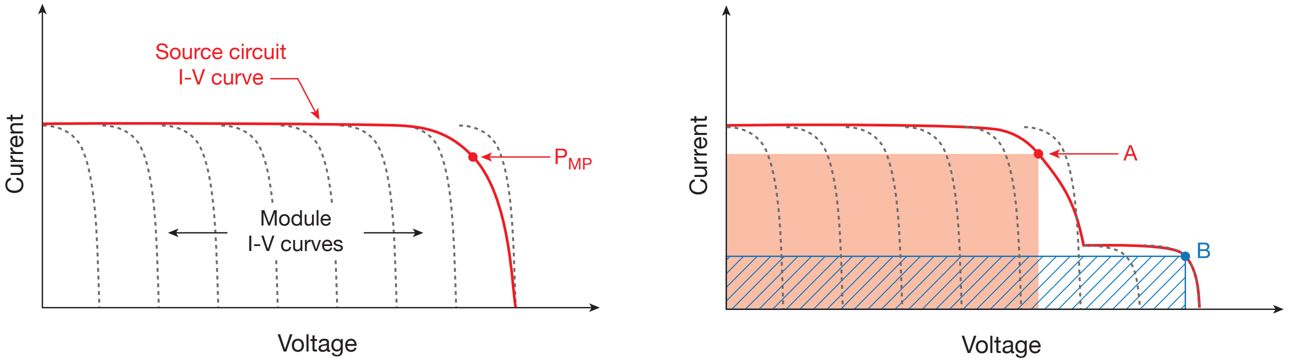

Figures 1a and 1b Figure 1a (left) illustrates how the I-V curve for a single PV source circuit is a composite of individual module I-V curves. In Figure 1b (right), two of these modules are shaded and receive diffuse irradiance only, which causes their current to drop proportionally. This results in two possible operating points that are locally optimal: At Point A, the system bypasses the shaded modules to keep the unshaded ones operating at full current; at Point B, the system reduces the source-circuit current to match that of the shaded modules.

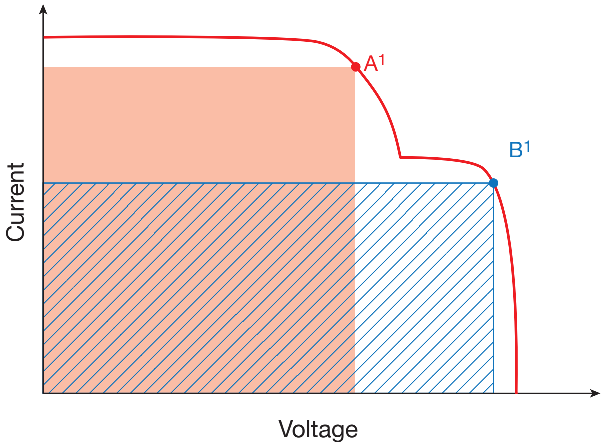

Figure 2 When a system connects multiple PV source circuits in parallel, the effects of shading may lead to different potential operating points, as represented by Points A^1 and B^1.

maximum current of the PV source circuit, that solar cell consumes power. To prevent this, modules contain bypass diodes that remove shaded cells from the source circuit under specific conditions. While bypass diodes primarily serve to improve product safety, they also mitigate shade impacts.

In the National Renewable Energy Laboratory (NREL) conference paper "Partially Shaded Operation of a Grid-Tied PV System" (see Resources), Chris Deline explains how bypass diodes improve module performance under shaded conditions: "The bypass diode allows current from non-shaded parts of the module to pass by the shaded part." He continues: "When a bypass diode begins conducting, the module voltage will drop by an amount corresponding to the sum of cell voltages protected by the bypass diode plus the diode forward voltage, but current from surrounding unshaded groups of cells continues around the group of shaded cells."

A shaded module does not produce zero energy because some amount of diffuse irradiance always strikes its surface.

Today, most modules have three bypass diodes. When one bypass diode activates, that effectively removes one-third of the cells from the circuit. This explains why a small amount of shading can have a disproportionate impact on module performance. According to Deline: "Shading half of one cell negates all the power produced by the 18 cells in that bypass diode group. Therefore, the reduction in power from shading half of one cell is equivalent to removing a cell active area 36 times the shadow's actual size."

Of course, reducing the voltage of a single module by one-third is preferable by far to restricting the current of an entire PV source circuit. Recall that series-connected solar cells and PV modules must all operate at the same current. On one hand, without any bypass diodes in the modules, hard shade on a single PV cell could shut down an entire PV source circuit. On the other, with three bypass diodes per module, hard shade on a single cell of a 12-module source circuit reduces the string voltage by less than 3%.

Components of sunlight. Sunlight is primarily composed of direct and diffuse irradiance. *Direct irradiance* is the beam of sunlight that comes directly from the sun, in a straight line. *Diffuse irradiance* is the scattered sunlight that arrives equally from all directions.

When a module is shaded, it primarily loses direct irradiance. Depending on how much of the sunlight is obstructed, the module still absorbs most, if not all, of the diffuse irradiance. Generally speaking, the effective diffuse irradiance accounts for 10%–40% of the total incident radiation. A shaded module does not produce zero energy because some amount of diffuse irradiance always strikes its surface.

SHADE EFFECTS

When an obstruction shades an array, a shadow moves across the array, based on the sun's angle in the sky and the obstruction's size. When this shadow hits modules in the array, it results in both irradiance and mismatch losses at the system level.

Irradiance. The irradiance effect of shade is the sunlight lost before it hits the module surface, compared to the irradiance that would have reached the modules without the obstruction. Since current and power in a PV source are directly proportional to irradiance, reduced irradiance due to shade results in direct energy losses.

Mismatch. The mismatch effect of shade is the difference between how each module could have performed at its individual maximum power point and its actual performance based on system constraints. Module-to-module mismatch due to shade results in indirect energy losses. Mismatch losses from shade are highly nonlinear and system dependent, based on component selection and system design specifics.

A PV source circuit impacted by shade has two potential operating points, as shown in Figures 1a and 1b (p. 36). On one hand, the source-circuit current could drop to match the restricted current of the shaded cells, as illustrated by Point B in Figure 1b. In this case, unshaded modules are running below their potential maximum output power and the resulting lost energy is counted as mismatch

CONTINUED ON PAGE 40

Production Modeling Software Tools

Since energy production is site specific and varies based on system configuration and rate structure, designers need software tools that can model plant performance and economic returns. Several products available on the market can calculate the energy yield of systems related to shading losses.

HelioScope: Developed by Folsom Labs, HelioScope is a cloud-based PV system design and performance-modeling program. In the software, users lay out a system based on the physical location and electrical connections of the PV modules. HelioScope then calculates the operating characteristics of each module individually and uses that data to calculate system mismatch effects based on each module's electrical behavior and circuit connections. This enables the program to calculate the irradiance and mismatch effects of shading based on 3D modeling of the obstructions in SketchUp. For row-to-row shading, HelioScope factors in both direct and diffuse effects, with added electrical effects for cell string-level performance. (We used HelioScope for most of the analyses in this article.)

PV Designer: Solmetric developed PV Designer as a companion product to its SunEye site evaluation tool. System designers can incorporate SunEye shade measurements taken at specific roof locations. The software then creates a map of the average irradiance across the array and uses this to model system performance. PV Designer does not calculate the mismatch effects of shade.

PVsyst: The eponymous software developed by PVsyst is the industry standard for modeling PV power plant production. PVsyst uses an "infinite sheds" approach to calculate row-to-row shading losses, modeling a generalized distance between rows of modules and their corresponding irradiance losses. The software can calculate shading effects from obstructions with a custom 3D near-object shading design tool. The software also approximates the string effects of shade obstructions through user-defined electrical effects.

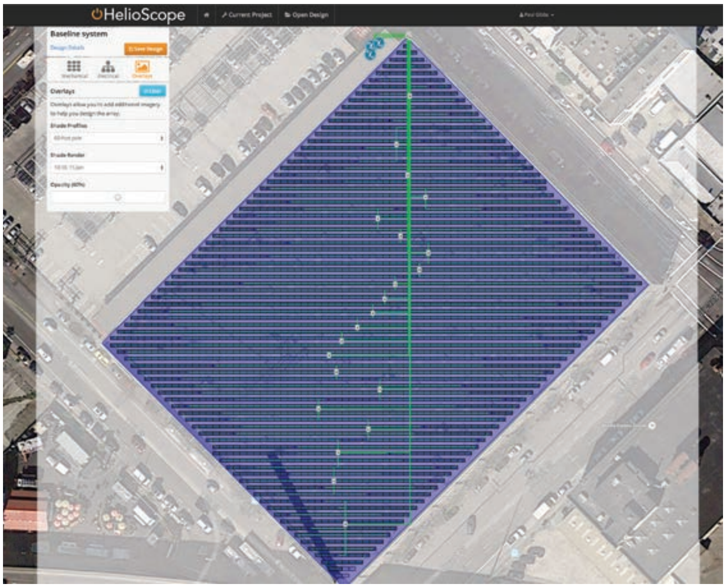

Production model We completed the near-object shading analyses in this article using HelioScope, which allows users to find a site on Google Earth and import its 3D layout into SketchUp. HelioScope then analyzes annual shade effects—such as the impact of a 60-foot pole, shown here—at both the module and the system level based on a 3D model.

PVWatts: NREL developed PVWatts, which estimates array production based on a weather file and a series of user-defined deratings. The software does not calculate system effects, instead requiring the user to input monthly loss factors due to obstruction or horizon shading.

Simuwatt: Developed by concept3D, Simuwatt is a 3D design application available for the Apple iPad. The application renders a system in 3D, including modules and obstructions, and calculates the shading that will hit each module. For performance calculations, the software utilizes the NREL System Advisor Model (SAM). ●

loss. On the other hand, bypass diodes could activate and remove the shaded cells from the circuit, as illustrated by Point A in Figure 1b (p. 36). In this case, the shaded modules are operating at a voltage that is above their unique maximum power point, which is another source of mismatch loss.

In a larger system with multiple source circuits in parallel, the shaded source circuit must produce a voltage similar to the other strings. This scenario requires modules on a shaded source circuit to run above their maximum power point to compensate for the voltage lost from shaded modules. This lost energy is again counted as mismatch loss. Figure 2 (p. 38) provides an example of mismatch losses in a larger system.

Calculating Obstruction Shade Losses

To estimate the shaded production of a PV array and determine the system effects of shading, system designers must use software to calculate the operating characteristics of every PV module and model performance based on how all of these modules are connected in series and parallel.

For the following analysis, we used HelioScope to calculate the system effects of shade. HelioScope can import shade patterns from SketchUp, a free 3D modeling platform, and use them to estimate system-level energy production. These estimates account for irradiance

CONTINUED ON PAGE 42

Shading Losses Associated with Various Obstructions

Scenario (net height)	Irradiance loss	Mismatch loss	Total shading loss
Pole (20 feet)	0.4%	1.4%	1.8%
Pole (40 feet)	0.8%	2.2%	3.0%
Pole (60 feet)	1.1%	2.9%	4.0%
Small tree (25 feet)	1.7%	2.8%	4.5%
Large tree (50 feet)	2.7%	4.2%	6.9%
Nearby building (70 feet)	2.4%	2.1%	4.5%
Tree line (40 feet)	3.1%	3.3%	6.4%

Table 1 This table details the irradiance, mismatch and total shading losses associated with common obstructions. We determined these values using software capable of calculating module-level operating characteristics in response to geolocated 3D shade patterns and modeling the resulting performance of the PV system as a whole.

and mismatch losses based on module-level irradiance and power calculations.

Our analysis assumes a 300 kW fixed-tilt, ground-mounted PV system in Southern California. We modeled obstruction-shade losses associated with different types of objects. To generate conservative results, we modeled near-object shading associated with obstructions located to the south of the array.

Baseline losses in a shaded array. To start our analysis, we set the row-to-row spacing and fully populated the available array area with modules. To analyze the full system impacts of obstruction shading, we left all of the shaded modules in the system. We then imported shade patterns for a variety of obstructions: 20-, 40- and 60-foot poles; 25- and 50-foot trees; a nearby building; and a 40-foot tree line. We then calculated the total system losses associated with each of these obstructions compared to an unshaded array.

Table 1 details the results of these shading loss simulations. Note that mismatch losses comprise a larger percentage of the total shading losses for skinny objects such as power poles. Meanwhile, irradiance losses are more significant for wider objects. These effects are intuitive, given that wider objects

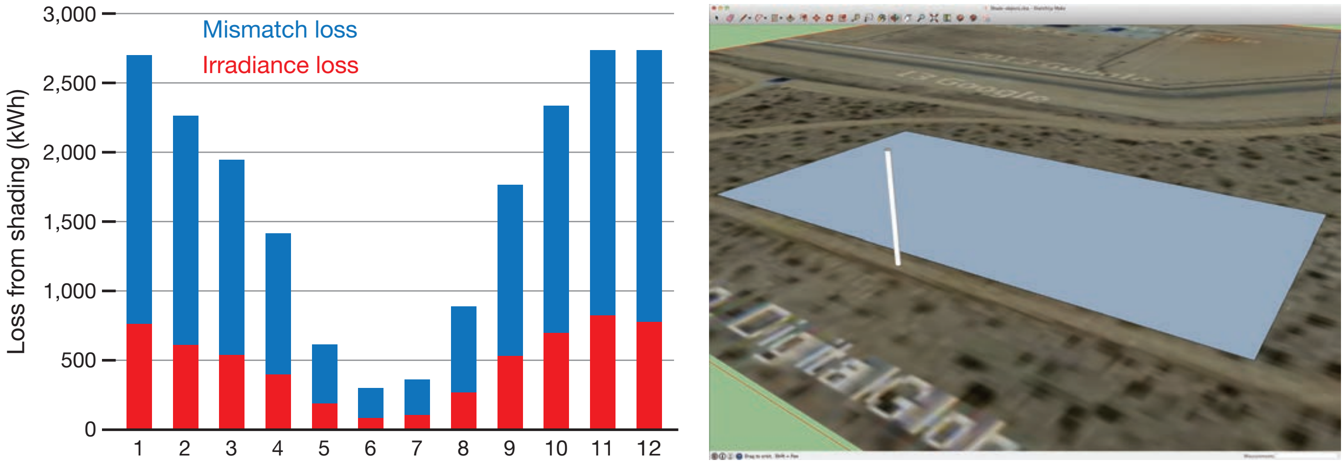

Figures 3a and 3b Figure 3a (left) shows the monthly energy losses resulting from a 40-foot pole located directly south of a 300 kW PV array in Southern California. Figure 3b (right) shows the 3D obstruction model we used to generate the shade pattern for the 40-foot pole.

block more direct sunlight. However, you may find the overall results less intuitive, as the system level losses associated with near-object shading are relatively modest given that we did not attempt to mitigate the shade effects.

To get a better idea of what is going on, we need to drill down into the details. For instance, Figure 3a shows the monthly energy losses associated with a 40-foot pole located directly south of the PV array. The annual system-level losses

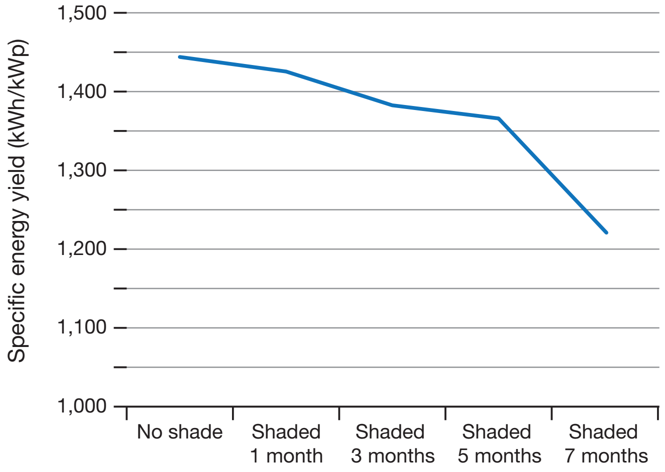

Figure 4 This figure details the specific yield associated with different groups of modules. We have grouped the modules according to the number of months in which each one experiences near-object shading from a 40-foot pole on the south edge of the array.

> The optimal design response to near-object shading varies for different shade profiles.

associated with a 40-foot pole are estimated at 3%, as reported in Table 1 (p. 42). However, Figure 3a clarifies that these losses vary significantly by season. The energy losses approach 2,700 kWh (>7%) in the winter when the sun is low on the horizon and drop below 500 kWh (<1%) in the summer when the sun is directly overhead.

Figures 3a and 3b help explain why predicted annual shading losses are modest. System designers typically evaluate shade patterns by looking specifically at winter months. While this is an informative exercise, it is also inherently conservative. PV systems in the Northern Hemisphere generate the most energy in June and July when the sun is overhead and shading losses are lowest.

We get similar results if we repeat this monthly energy loss analysis for other types of obstructions responsible for near-object shading. However, the system-level effects are generally less than expected. This analysis helps us understand the losses associated with disregarding shade effects entirely. But what if we take a more nuanced approach based on partial tolerance to seasonal shading?

Partial shade tolerance. Designers have a range of options for dealing with near-object shading from obstructions. For example, rather than simply viewing modules in an array as either shaded or unshaded, we could break the array into groups based on how often each module is shaded.

In this scenario, the unshaded modules serve as the reference group, since they have the highest production efficiency or specific yield as measured in kilowatt-hours per kilowatt. Using the production efficiency of the unshaded modules as the basis of comparison, we can then look at the relative specific yield for other groups of modules, such as those that are shaded in 1 month (December), 3 months (November–January), 5 months (October–February) and so on.

Figure 4 shows an example of this type of production efficiency analysis, detailing the shade effects associated with a 40-foot pole located directly south of the array. Note that the shading losses are once again relatively modest. This is especially true for modules that are shaded in December only, as the production efficiency of the module group shaded in 1 month is just 1.3% less than that of the unshaded group. The production efficiency losses are 4.2% for the module group shaded in 3 months and 5.3% for the group shaded in 5 months.

These data can provide system designers with valuable insights and allow them to take a more informed approach to shade tolerance. For example, if the designer includes modules shaded in 1 month in the final design, the capacity of the PV system increases by 6% compared to including unshaded modules only. Further, if the designer includes modules that are shaded in 3 months, the total system capacity increases another 7%.

When we ran the same analysis for tree shade, we found very similar production efficiency values for the groups of modules shaded in 1 or 3 months. However, the production efficiency curves for tree shade versus pole shade diverge significantly for groups of modules shaded in 5 or more months. Because tree shading is wider than pole shading, it reduces the production efficiency of affected modules more severely.

These results suggest that the optimal design response to near-object shading varies for different shade profiles. But how do we determine the best project-specific design approach?

THE ECONOMICS OF SHADE

To optimize a design response to near-object shading, it is important to look at the relative economic performance of various design options. On one hand, the designer could plan a system that maximizes the production efficiency of the array by eliminating any shaded modules from the design. On the other, the designer could increase system capacity based on different levels of shade tolerance.

CONTINUED ON PAGE 46

Reliability Implications of Shading

Shaded modules run at reduced power levels. A shaded module still receives sunlight, but because of system mismatch effects it is likely not operating at its maximum power point. Further, if the sunlight hitting a module does not convert to electricity, the laws of thermodynamics suggest that it converts to heat. This additional heating could negatively affect module reliability.

Similarly, module shading can cause bypass diodes to activate. When a bypass diode conducts current, its temperature increases. Both of these effects—increased temperature and diode activation—could also negatively affect module reliability.

Even though shading appears to put additional stresses on PV modules, it is difficult to assign a precise number to the reliability implications of these effects. If you have questions about the viability of a partially shaded design, we recommend contacting your module supplier's product applications engineer. It is also important to read and follow the manufacturer's instructions to ensure that you are using the modules in accordance with the warranty terms. ●

We conducted a cost-benefit analysis to evaluate the economic performance of different design responses to the pole and tree shading scenarios. We premised our results on two important assumptions. First, we assumed that the available array area is space constrained, which limits our design options for dealing with shade. We cannot simply move modules that we know will be shaded to an unshaded location; we can either incorporate the shaded modules into the system or remove them. Second, we assumed that the PV-generated energy does not exceed the building load or the terms of the off-take agreement. Obviously, there is no need to increase PV array capacity by locating modules in partially shaded areas if you do not have a good use for that energy.

Benefits of higher array capacity. A larger PV system has multiple financial benefits. First, even with a constant profit margin, a larger system yields more money for the system owner and developer because it has more modules and produces more energy. More important, a larger system also improves the array's cost structure and profitability by spreading out the project's fixed costs. The overhead to develop and build a project includes a lot of scale-independent tasks, such as permitting and interconnection agreements, construction administration and general overhead. By making a system larger—say, by 10%–20%—designers can spread fixed costs over the increased system capacity, typically reducing these fixed costs by 10%–20% on a per-watt basis.

To analyze the optimal system capacity for two partially shaded arrays—one shaded by a pole and the other by a tree—we used a lifetime system cost model. This model incorporates cost assumptions representative of a commercial-scale PV system: fixed costs of $0.50 per watt, marginal costs of $1.50 per watt and an energy value of $0.15 per kilowatt-hour. As before, we broke down the arrays into module groups according to how often the obstruction shades the modules. We then treated each group of modules as an independent design option and looked at its cost versus production benefits. This approach allowed us to analyze the profit contribution for each group of modules. Figure 5 summarizes the results of this profitability analysis.

The data in Figure 5 provide concrete evidence that different shade types lead to different design optimization decisions. It appears to be profitable to add module groups shaded by the pole into the final design, even those shaded in 5 or 7 months of the year. With regard to tree shade, however, it is clearly not profitable to include modules that are shaded in more than 5 months of the year.

Shade Mitigation Design Strategies

Once the designer knows which module groups improve the economic performance of a PV system, the designer can use the same type of cost-benefit analysis to evaluate other responses to the effects of shading. When considering different shade mitigation

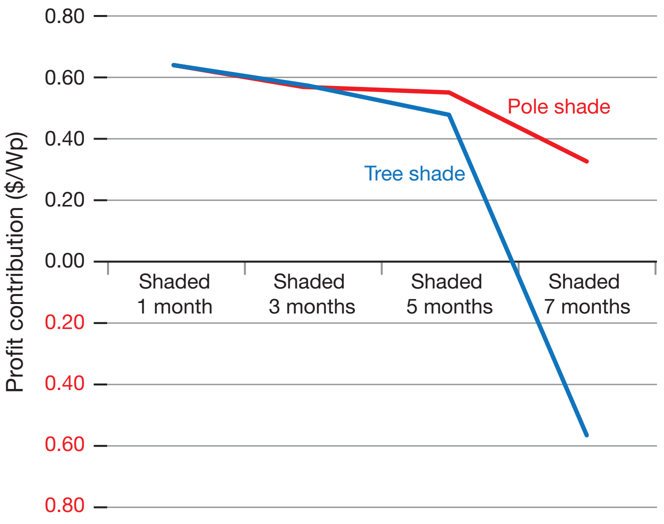

Figure 5 This figure plots the relative profit contribution of different module groups based on shade type and frequency. While the module groups subject to shading by a pole (the red line) generally remain profitable, persistent shading by a tree (the blue line) can erode project profitability.

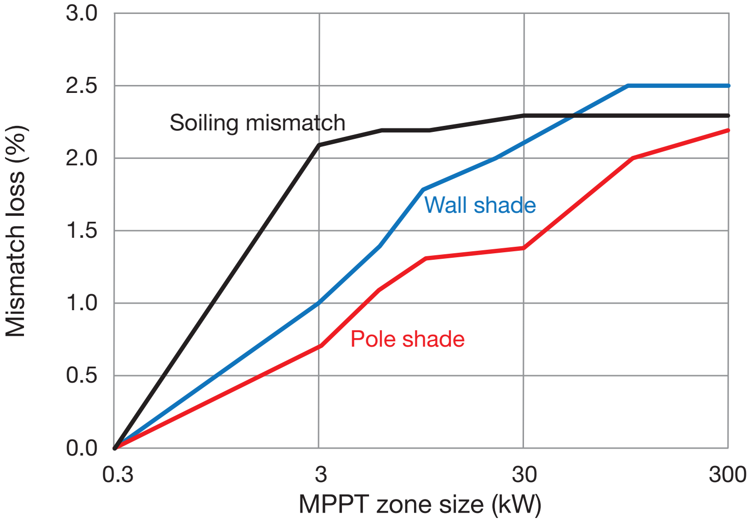

Figure 6 This figure plots the relative mismatch losses associated with different shade types based on MPPT zone size. Note that the benefits of increased MPPT granularity depend on the shape of the obstruction and its shade profile, which drives the series and parallel mismatch effects at the system level.

options, it is helpful to keep the two-part nature of shade losses—irradiance and mismatch—in mind.

Since irradiance losses result from sunlight that never reaches the modules, the only possible design response is to eliminate the obstruction or move the modules. There is really nothing that the system designer can do to replace the lost irradiance.

Mismatch losses vary depending on how you configure the system, which means you can take many possible approaches. For example, you could configure the PV source circuits to minimize shade effects. You could increase the granularity of the MPPT zone. Or you could improve system shade tolerance by using thin-film instead of c-Si modules.

But what shade mitigation options are most cost-effective? Does it make sense to increase the granularity of the MPPT by using module-level power electronics or 3-phase string inverters? Is it more cost-effective to use a single central inverter and to wire the PV source circuits in a way that minimizes shade effects? What happens when the system uses thin-film modules?

These are all scenarios that the system designer can model and evaluate using software. Performing this type of cost-benefit

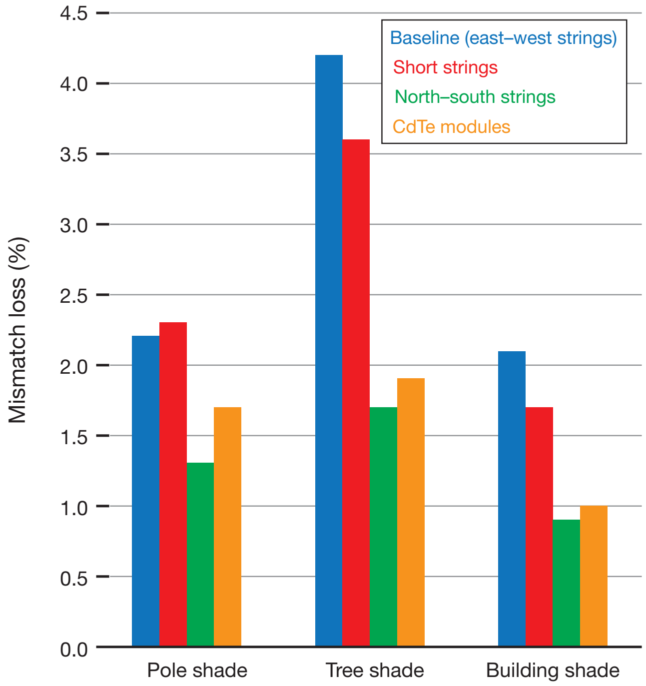

Figure 7 This figure details the mismatch losses associated with different source-circuit configurations based on shade type. Note that orienting the PV source circuits in the same direction as the shade reduces the mismatch losses by about 50% compared to the baseline case.

analysis provides designers with a quantitative means of evaluating shade mitigation design strategies, as shown in the following examples.

Optimal MPPT zone size. Today, system designers have more hardware options than ever for mitigating mismatch losses. Rather than specify a single 300 kW central inverter, a designer could specify 15 20 kW string inverters or eliminate mismatch losses altogether by specifying 1,200 250 W microinverters or dc-to-dc converters. Some products in development offer bypass-diode–level MPPT, which will optimize the production of individual cell strings or groups and allow for tighter module spacing and higher ground-cover ratios.

To determine the most profitable design option, the system designer needs to quantify the increased energy production associated with more-granular MPPT zones, and then model financial returns over the life of the project based on the cost structure of different design options. As shown in Figure 6 (p. 47), the optimal design response to mismatch losses depends in part on the source of the shading. For example, if a pole is the source of the shade, the designer can reduce mismatch losses approximately 50% by specifying 10 kW MPPT zones rather than a 300 kW central inverter. However, the only way to meaningfully reduce soiling mismatch losses is to utilize substring-level MPPT.

Optimal source-circuit configuration. It often makes sense to include partially shaded modules in a final design regardless of whether the system includes string inverters or module-level power electronics to mitigate mismatch losses. However, if the system includes a central inverter, which typically puts many source circuits in parallel on a single MPPT input, the designer may want to evaluate the potential benefits of different source-circuit configurations.

For example, you might be able to mitigate mismatch losses by specifying shorter source circuits. Alternately, you could wire the source circuits so that they are physically oriented in the same direction as the shade, rather than perpendicular to it. The designer might also be able to mitigate mismatch losses by using cadmium telluride (CdTe) thin-film modules.

As shown in Figure 7, when we modeled these different design options, we found that source-circuit wiring is generally the most effective of these shade mitigation design options, regardless of shade type. To the extent that the PV source-circuit wiring groups shaded modules together, system-level mismatch losses become less of a problem. Changing the module type has a similar effect, in part due to the fact that CdTe modules have a lower fill factor than c-Si modules. CdTe modules also result in shorter PV source circuits. However, shorter strings are actually a liability in conjunction with pole shade.

Put It in the Shade?

While designers should generally minimize array shading, our analysis indicates that a measured shade tolerance can optimize a PV system's profitability. Ultimately, the optimal design will vary depending on the obstruction type and its unique shade profile. However, traditional design approaches to near-object shading, especially those that base decisions on the extent of shading on the winter solstice, may be overconservative. Designers should instead evaluate shade tolerance and shade mitigation strategies to determine which responses improve the economic performance of the PV asset over the life of the system. ⏚

» CONTACT

Paul Grana / Folsom Labs / San Francisco, CA /
paul.grana@folsomlabs.com / folsomlabs.com

Paul Gibbs / Folsom Labs / San Francisco, CA /
paul.gibbs@folsomlabs.com / folsomlabs.com

RESOURCES

Deline, Chris, "Partially Shaded Operation of a Grid-Tied PV System," NREL conference paper (NREL/CP-520-4601), June 2009, nrel.gov/docs/fy09osti/46001.pdf