43rd European Photovoltaic Solar Energy Conference and Exhibition

This presentation was selected by the Sc. Committee of the EU PVSEC 2026 for submission of a full paper to one of the EU PVSEC's collaborating peer-reviewed journals.

UNIVERSAL ANALYTICAL MODEL FOR INTER-ROW SHADING LOSS ESTIMATION IN PV PLANTS

Maitheli Nikam, Karel De Brabandere, Julien Deckx, Gofran Chowdhury
3E
Kalkkaai 6, 1000 Brussels, Belgium — karel.debrabandere@3e.eu

ABSTRACT: Inter-row shading in high-density PV plants causes two distinct losses: a linear irradiance deficit and a non-linear electrical mismatch loss. The mismatch term carries most of the uncertainty and normally requires explicit current-voltage curve composition, which is far too slow to evaluate continuously over an operating fleet. This paper presents a closed-form model that reproduces the mismatch loss without any iteration. An array is described by three integers: the number of vertical partitions in parallel, the number of series partitions within each of them, and the number of cells per series partition. Two floor functions locate the shade edge inside that structure, and the loss follows algebraically for both common inverter wiring strategies. Evaluated against an explicit current-voltage composition over eleven configurations, the model is two and a half to thirty-eight times closer to that reference than an area-proportional assumption, and it evaluates a plant-year at one-minute resolution in milliseconds rather than hours. It is further applied to two operating plants in Germany, where it resolves the output of individual module rows. An implementation has been prepared for the open-source pvlib library.
Keywords: Shading, Modelling, Performance, Monitoring, PV Array

1 INTRODUCTION

As land becomes scarce, developers favour higher capacity on the same footprint, so rows are spaced more tightly and inter-row shading becomes a routine operating condition rather than an edge case. This matters well beyond plant design. Asset managers and O&M providers judge a plant by its Energy Performance Index (EPI), the ratio of measured output to the output a performance model says the plant should have produced. Every expected-loss term inside that model propagates directly into the indicator used to decide whether to send a technician. If the expected shading loss is too low, real production is flagged as underperformance and crews are dispatched to chase shadows. If it is too high, a genuine fault hides behind shading that was over-predicted.

Total shading loss separates into two components. The irradiance deficit is linear in the shaded area and is cheap to evaluate. The electrical mismatch loss is not: a partially shaded cell limits the current of every cell in series with it, bypass diodes switch in and out as the shadow advances, and the power-voltage characteristic develops several local maxima between which the global maximum power point jumps. The magnitude of this term depends on module technology, on module orientation, on how many module rows are stacked on a table, and on how those rows are connected to the inverter's maximum power point trackers (MPPTs).

Capturing this behaviour rigorously means composing current-voltage (I-V) curves cell by cell, which is accurate but iterative and far too slow to run for every string of a portfolio at every timestep. Simplified models are fast but generally smooth away precisely the bypass-diode behaviour that dominates the loss. Neither trade-off suits operational monitoring, which needs a per-string, per-timestep number over a whole fleet.

This paper presents a closed-form model that contains no iteration and no transcendental solve, yet reproduces the staircase behaviour that the bypass diodes impose. Section 2 states the shared irradiance and geometry framework and the I-V reference model. Section 3 develops the partition model. Section 4 reports accuracy, computational cost and an evaluation against two operating plants. Section 5 positions the work against existing methods and Section 6 states its limitations, before Section 7 concludes.

2 MODELLING FRAMEWORK

2.1 Irradiance decomposition and shading geometry

Both the reference model and the proposed model share the same front end. Global horizontal irradiance, from satellite or on-site sensors, is decomposed into its horizontal components using the model of Skartveit et al. [1] and transposed to the plane of array with the Hay and Davies model [2], giving a beam component, a circumsolar diffuse component, an isotropic sky diffuse component and a ground-reflected component.

Under row-to-row shading only the beam and circumsolar components produce a sharp shadow boundary and hence spatially non-uniform irradiance on the array; the isotropic sky and ground-reflected components are treated as spatially uniform over the active surface. The electrical mismatch model therefore acts on the beam and circumsolar components only. The geometric shaded fraction is obtained from standard two-dimensional infinite-shed projections using solar position and the plant layout parameters (tilt, azimuth and ground coverage ratio), as implemented in pvlib-python [3].

2.2 Reference model: explicit I-V composition

To establish a ground truth for the mismatch loss, an explicit I-V composition was implemented in pvlib-python. Single-diode parameters are evaluated separately for shaded and unshaded cells, with reverse-bias breakdown represented by the Bishop extension [4] so that shaded cells dissipate rather than simply stop conducting. Cell curves are combined in series by summing voltages, with bypass diodes acting as ideal switches that clip negative voltages; series blocks are combined along a row, and rows are combined in parallel by summing currents. The maximum of the resulting array power-voltage curve is the operating point.

Fig. 1 shows the result for four stacked rows of landscape modules with three bypass diodes per module. The left panel gives the array power-voltage curve at several optical shaded fractions. Shading creates a second

10.4229/EUPVSEC2026/4BO.6.3
020243-001

local maximum; as the shadow rises, the global maximum power point transfers from one local maximum to the other each time a bypass diode activates or deactivates. The right panel traces the resulting global maximum against shaded fraction and shows the characteristic staircase, with the markers on each curve of the left panel carried across. It is this staircase, not the optical deficit, that a fast model has to reproduce.

3 THE PARTITION MODEL

3.1 Three integers

Inter-row shadows advance upward from the bottom of a table, so the only structure that matters is how the array is divided along the vertical axis. The model reduces any array to three integers (Fig. 2):

N_par, the number of vertical shading partitions in parallel. This is the number of modules stacked vertically on a table, doubled for portrait modules with half-cut cells because the two half-strings are electrically independent.

N_ser, the number of vertical shading partitions in series within each parallel partition. For landscape crystalline-silicon modules the three bypass-diode blocks are stacked along the vertical axis and N_ser is 3; in most other arrangements it is 1.

N_cells, the number of cells per series partition along the vertical axis: typically 1 for thin-film cell strips, 2 for landscape crystalline silicon and 10 or 12 for portrait modules.

Common configurations map onto this triple directly. A table of n landscape rows is [n, 3, 2]; n rows of portrait full-cell modules is [n, 1, 12]; n rows of portrait half-cut modules is [2n, 1, 12]; thin-film strips in portrait are [n, 1, 1]. Beyond these three numbers the loss model needs no further information about the hardware.

3.2 Closed-form loss function

Let f_sh be the optical shaded fraction of the array, between 0 and 1, and let

$$N_{\text{tot}} = N_{\text{par}} \times N_{\text{ser}} \quad (1)$$

be the total number of vertical partitions. Two floor functions count how many partitions are completely shaded, across the whole array and per parallel partition respectively:

$$N_{\text{tot,sh}} = \lfloor f_{\text{sh}}\, N_{\text{tot}} \rfloor \quad (2)$$

$$N_{\text{par,sh}} = \lfloor f_{\text{sh}}\, N_{\text{par}} \rfloor \quad (3)$$

Their difference locates the fully shaded series partitions inside the one parallel partition that is part-way through its transition:

$$N_{\text{ser,sh}} = N_{\text{tot,sh}} - N_{\text{par,sh}}\, N_{\text{ser}} \quad (4)$$

What is left over is the shade state of the single partition in transition,

$$f_{\text{cell}} = f_{\text{sh}}\, N_{\text{tot}} - N_{\text{tot,sh}} \quad (5)$$

Because the cells of a parallel partition are in series, the partially shaded cell limits the current of every series partition in that partition which is not already bypassed. It is shaded over f_cell of the height of one cell, and there are N_cells cells along the vertical axis of a series partition, so the current-mismatch term is

$$L_{\text{cur}} = f_{\text{cell}}\, N_{\text{cells}}\, (N_{\text{ser}} - N_{\text{ser,sh}}) \quad (6)$$

Equation (6) is where the model departs from a purely combinatorial partition count: the sub-partition ramp is derived from the cell count rather than fitted, so the loss rises continuously from zero as the shadow enters a partition instead of switching the partition off in one step.

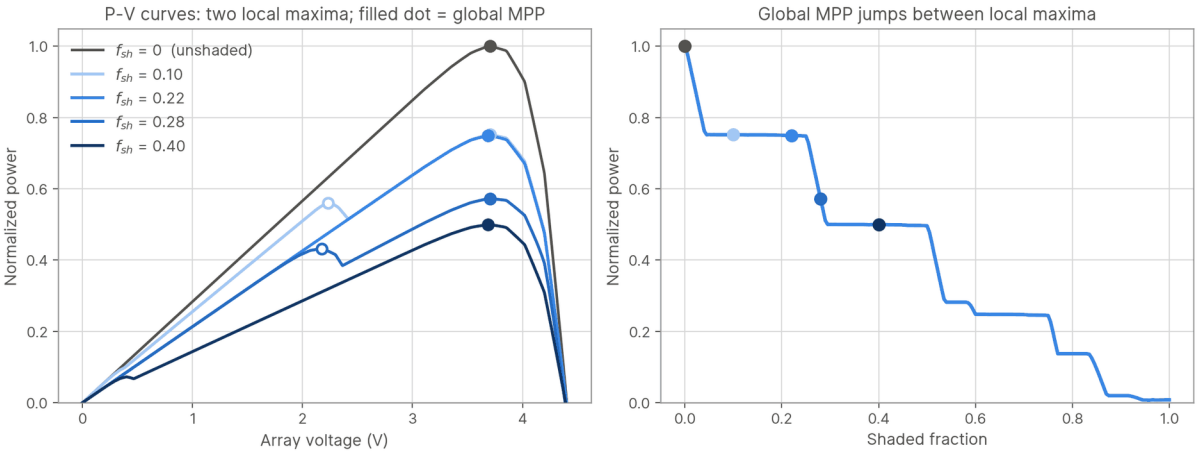

Figure 1: Array power-voltage curves at several optical shaded fractions (left; filled marker = global maximum power point, open marker = the competing local maximum), and the resulting global maximum against shaded fraction (right). Four rows of landscape modules, three bypass-diode blocks per module.

10.4229/EUPVSEC2026/4BO.6.3
020243-002

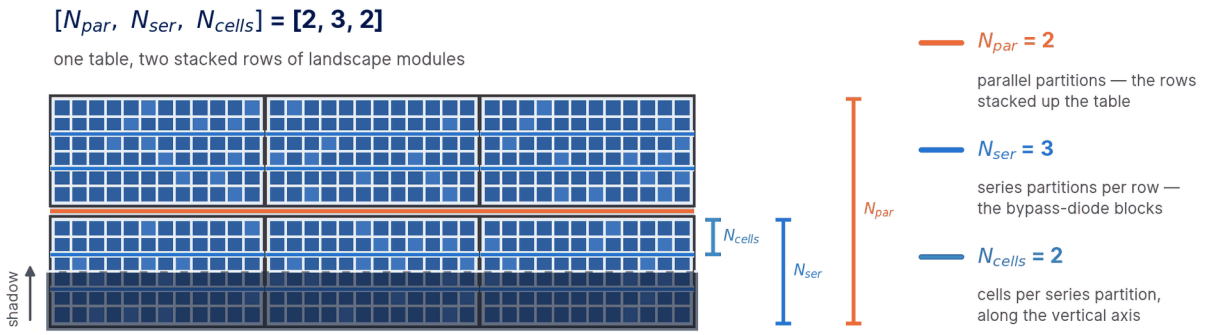

Figure 2: The three integers, shown on a table carrying two stacked rows of landscape modules, configuration [2, 3, 2]. The shadow advances upward from the bottom edge.

3.3 Inverter wiring

The wiring between module rows and MPPTs changes which expression applies, but not the method. When each row has its own MPPT (scenario A, "same row per MPPT") there is no voltage mismatch between rows and the array loss is the average over all partitions:

$$L(f_{sh}) = [\, N_{tot,sh} + \min(1, L_{cur}) \,] / N_{tot} \qquad (7)$$

When several rows share one MPPT (scenario B, "all rows") their voltages are forced equal, the array can no longer sit at the sum of the individual maxima, and a voltage-mismatch term appears which counts the fully shaded series partitions dragging down the rows that are not yet shaded:

$$L_{vol} = (N_{par} - N_{par,sh})\, N_{ser,sh} + \min(N_{par} - N_{par,sh}, L_{cur}) \qquad (8)$$

$$L(f_{sh}) = [\, N_{par,sh} + \min(1, L_{vol} / N_{ser}) \,] / N_{par} \qquad (9)$$

Both expressions return the fraction of beam-driven DC power lost, normalised to 1. Evaluating them requires two floor functions, a product and two minima; there is no iteration and no transcendental equation anywhere. For N_ser = 1 the two scenarios are provably identical, which is the expected result: with a single series partition there is nothing to bypass, so the wiring cannot matter.

3.4 Behaviour across topologies

Fig. 3 evaluates equations (1) to (9) over the configurations that occur most often in the field. Two trends are visible. More parallel partitions give a finer staircase and less loss at low shaded fractions, because the shadow is spread over more independent units. More cells per series partition makes the first step more severe, because one shaded cell row is enough to current-limit the whole partition. Every panel is the same expression with three integers changed.

4 RESULTS

4.1 Accuracy against the I-V reference

The model was evaluated against the explicit I-V composition of Section 2.2 by sweeping the optical shaded fraction over its full range and comparing the predicted normalised loss. Because the reference runs with a finite diffuse fraction (15%) and power is not exactly linear in irradiance, the reference loss was normalised onto [0, 1] by its own endpoints, so that loss shapes are compared rather than the irradiance-to-power conversion. The comparison is made against an area-proportional (linear optical) assumption, which is what a model without any mismatch term would give.

10.4229/EUPVSEC2026/4BO.6.3
020243-003

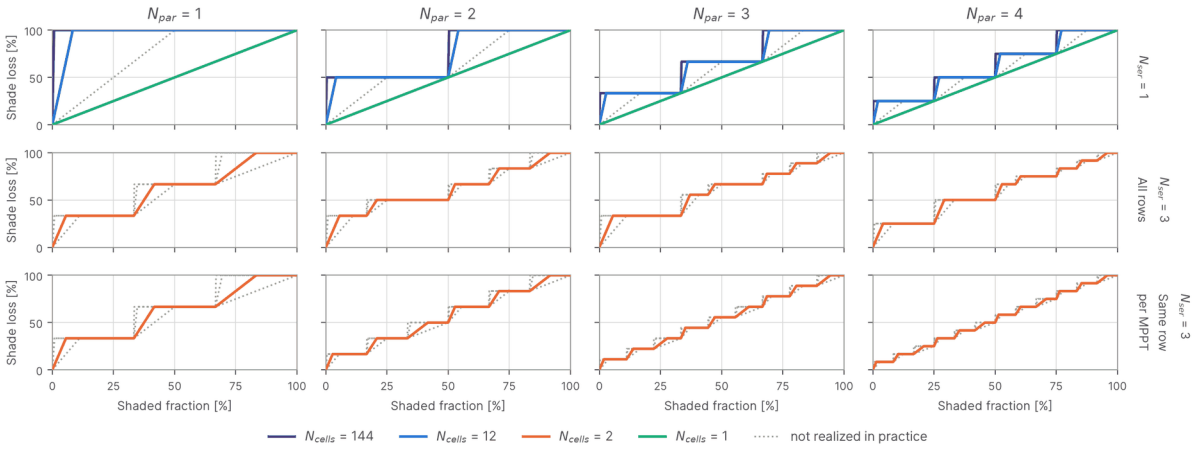

Figure 3: Shade loss against optical shaded fraction across the configuration space. Columns vary the number of stacked rows; bands vary the number of series partitions and the inverter wiring. Solid: realised in practice; dotted grey: combinations that do not occur. "Same row per MPPT" is scenario A, "all rows" scenario B.

Fig. 4 shows one configuration in full. The closed form places the steps at the right shaded fractions and the plateaus at the right heights, in both wiring scenarios, with a mean absolute deviation of 0.017 and 0.014 respectively. The visible residual is concentrated at the step corners, where the model switches a series partition between current-limited and bypassed abruptly while the composed curve crosses over more gradually.

Table I extends the comparison to eleven configurations covering portrait and landscape orientation, full and half-cut cells, one, two and four stacked rows, and both wiring scenarios; for the single row the two coincide, there being only one row to wire. The model stays between 2.5 and 38 times closer to the I-V reference than the area-proportional assumption in every case. The margin is largest for portrait and half-cut configurations, where a single shaded cell row removes a whole partition and an area-proportional argument cannot represent that at all, and smallest for landscape, which is simply where the area-proportional assumption was least wrong to begin with. Four full-cell rows and two half-cut rows reduce to the same triple [4, 1, 12] and accordingly give identical figures, which is the universality claim in miniature. The last column reproduces direct_martinez [5], the closest comparable model in pvlib today, measured the same way; Section 5 discusses why it behaves as it does.

4.2 Computational cost

The second half of the value proposition is speed. Because the loss is a handful of arithmetic operations with no iteration, it vectorises over time and over strings, so a whole fleet is one array operation rather than a loop. Its cost is also independent of plant size: three integers describe the array however large it is. Table II reports the measured cost of a plant-year for one array configuration, on a single desktop CPU core, best of five runs.

Configuration	Sc.	Model	Linear	Martinez
2 rows, portrait	A	0.014	0.213	0.074
2 rows, portrait	B	0.007	0.220	0.076
4 rows, portrait	A	0.007	0.107	0.085
4 rows, portrait	B	0.003	0.113	0.082
1 row, landscape	A, B	0.069	0.173	0.053
2 rows, landscape	A	0.035	0.086	0.093
2 rows, landscape	B	0.035	0.140	0.079
4 rows, landscape	A	0.017	0.043	0.128
4 rows, landscape	B	0.014	0.091	0.095
2 rows, portrait half-cut	A	0.007	0.107	0.085
2 rows, portrait half-cut	B	0.003	0.113	0.082

Table I: Mean absolute deviation from the explicit I-V composition over the full sweep of optical shaded fraction, for the closed form (Model), for an area-proportional assumption (Linear) and for direct_martinez [5]; Sc. is the wiring scenario. Each is normalised on [0, 1] by its own full-shade value, as described in Section 4.1.

10.4229/EUPVSEC2026/4BO.6.3
020243-004

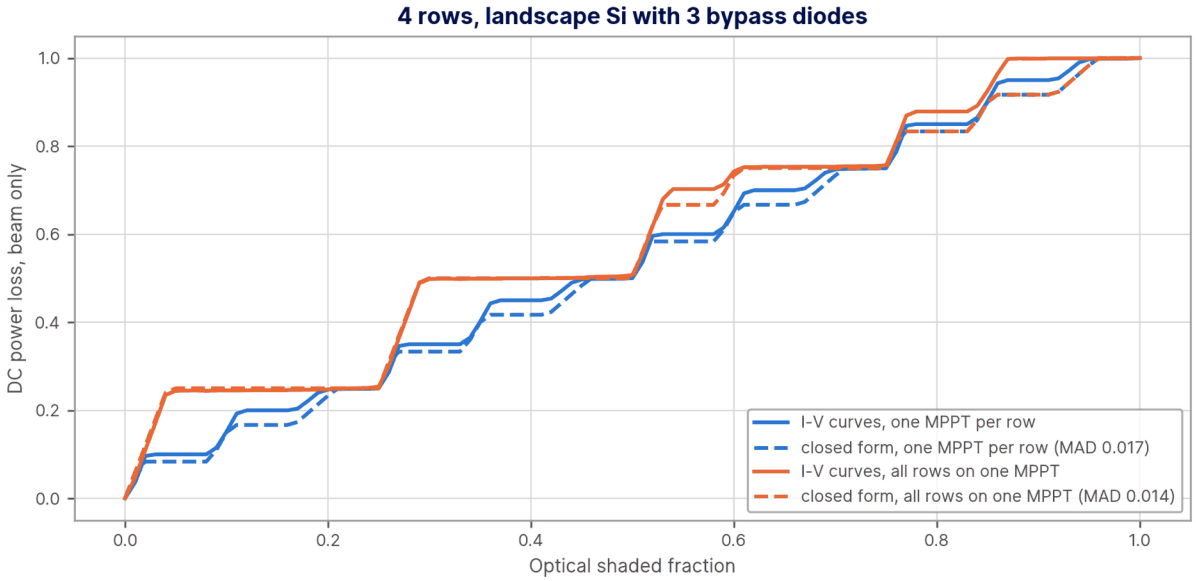

Figure 4: Closed-form model (dashed) against explicit I-V composition (solid) for four rows of landscape modules with three bypass-diode blocks, in both wiring scenarios. MAD denotes the mean absolute deviation over the sweep.

A plant-year of shading loss	Closed form	I-V composition
1-minute steps (525 600)	14 ms	6.6 h
15-minute steps (35 040)	0.2 ms	27 min

Table II: Time to evaluate the shading loss of one array configuration over a full year, vectorised closed form against the explicit I-V composition used as reference.

This comparison should be read for its order of magnitude and not as a fixed factor. Both I-V compositions available to us evaluate one timestep at a time in Python, and a properly vectorised composition would close part of the gap. The point that survives that caveat is operational: hours per plant-year is a study run once at design time, while milliseconds per plant-year is a term that can be recomputed continuously for every string in a portfolio.

4.3 Evaluation on operating plants

The model was applied to two operating plants in Germany, both with two stacked rows of portrait modules: Haßloch, with half-cut cells, configuration [4, 1, 12]; and Barth, with full-cell modules, configuration [2, 1, 12]. For each, modelled DC power was compared with the measured DC power of every inverter input on a shading-affected winter day, each normalised by its own DC capacity (Fig. 5). The simulation was run separately for the lower and the upper module row.

At Haßloch the lower row loses power in the morning and again in mid-afternoon while the upper row does not, and the measured inverter inputs separate in the same way. Barth is the clearer case: its 450 measured inputs fall into three distinct bands, which correspond to the three ways its strings are distributed over the two rows — two strings on the upper row, two on the lower, or one on each. The model reproduces all three bands from the same expression with only the row index changed. This row-level resolution is what allows an underperformance to be attributed to a row rather than to the plant as a whole.

Two honest qualifications apply. The ground coverage ratio of the Barth plant had to be corrected from the value on record before the comparison was meaningful, which is itself a reminder that configuration data is a common error source in this kind of work. And although the field data constrains the overall loss well, it does not cleanly separate the two wiring scenarios: the difference between them is small at these latitudes and clear-sky winter days are rare. The wiring distinction is therefore presented here as a capability of the model rather than as a field-validated result.

5 RELATION TO EXISTING METHODS

The partition concept itself is not new and the novelty claimed here is deliberately narrow. Oliosi et al. [6] presented the PVsyst partition model at this conference in 2021, including discrete partitions, the resulting staircase loss and its dependence on how strings are grouped per MPPT input. Deline et al. [7] published the canonical fast analytical treatment of inter-row shading for large arrays, implemented in SAM, with a computation time independent of array size. Saint-Drenan and Barbier [8] analysed and modelled the same effect from operating data.

Against that background this work contributes two things. The first is continuity within the partition. The PVsyst treatment deactivates a whole partition rectangle as soon as its bottom cell row is shaded and then restores low-shading behaviour with a separately fitted fraction for the electrical effect; equation (6) instead derives that sub-partition ramp from the cell count, so no additional fitted quantity is required. The second is that the result is a formula rather than a procedure. The PVsyst model works from three-dimensional geometry and interpolated shading tables; equations (1) to (9) are a single algebraic expression in three integers, each of which has a direct physical meaning, which is what makes the model portable into a library and cheap enough to evaluate per string.

10.4229/EUPVSEC2026/4BO.6.3
020243-005

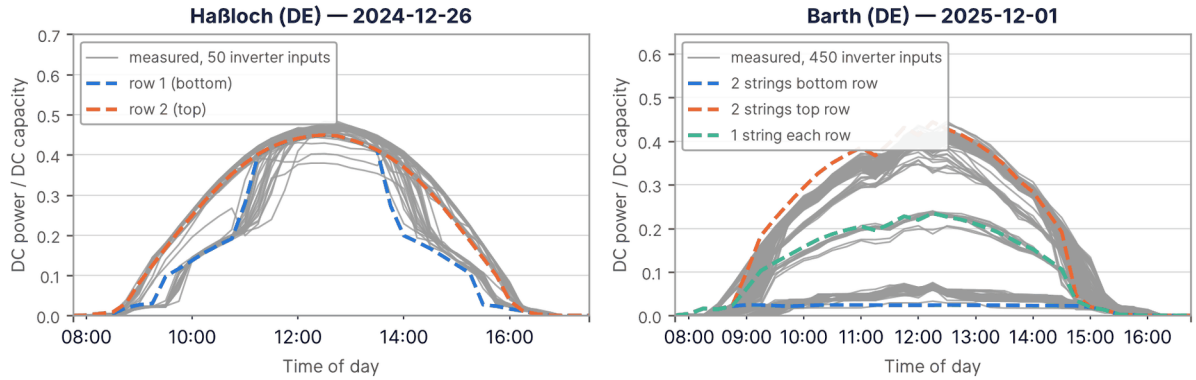

Figure 5: Modelled against measured DC power on a shading-affected winter day at two operating plants. Grey: measured DC power of each inverter input, normalised by its own DC capacity. Dashed: the model evaluated separately for each way the strings of an inverter input are distributed over the two module rows.

The closest comparable function in the pvlib library today is direct_martinez, an experimental model of Martínez-Moreno et al. [5], reproduced in the last column of Table I. On the four-row landscape array of Fig. 4 it deviates from the I-V reference by a mean of 0.128 against 0.017 for the model presented here. The difference is structural rather than a matter of tuning. Its only topology input is the number of blocks along the shaded axis, so the number of cells per block is invisible to it: portrait silicon and thin-film cell strips with the same block count receive the same prediction, where the present model separates them by 0.458. It carries no parameter for the MPPT wiring, so the distinction of Section 3.3 cannot be expressed. And because the shaded block count enters through a ceiling function, the predicted loss steps discontinuously to a finite value the instant any shade touches the array, where equation (6) ramps continuously from zero.

A further modelling difference is worth stating: the circumsolar component is grouped with the beam here, as it shares the same shadow boundary, where the PVsyst treatment handles it as uniform diffuse.

6 LIMITATIONS

The loss factor is applied to the beam and circumsolar components only. A partition that is bypassed or current-limited in reality also forfeits its diffuse-driven contribution, and a factor acting on the beam alone cannot express that. This is the dominant part of the residual in Section 4.1: for the weakest case in Table I, a single landscape row, the mean absolute deviation is 0.069 at a 15% diffuse fraction but falls to 0.018 at 1%. What remains at low diffuse fractions is the sharpness of the step transitions. The same design decision — one scalar factor on the beam — is what makes the model closed-form and fast, so this residual and the speed have a common origin.

The normalisation used in Section 4.1 is a choice, made so that loss shapes could be compared independently of the irradiance-to-power conversion; a different convention shifts the absolute numbers in Table I, though not the ordering. Finally, the model assumes the infinite-shed geometry of its front end, so it inherits that model's treatment of finite row ends and non-uniform terrain.

7 CONCLUSIONS

Inter-row electrical mismatch loss can be written in closed form. Two floor functions, a product and two minima, evaluated on three integers describing the vertical electrical structure of the array, reproduce the staircase that bypass diodes impose, for both common inverter wiring strategies and without any iteration.

Against an explicit I-V composition the model is 2.5 to 38 times closer than an area-proportional assumption across eleven configurations, and it evaluates a plant-year at one-minute resolution in milliseconds where the composition needs hours. Applied to two operating plants it reproduces the separation between module rows that the measured inverter inputs show, which is what makes row-level attribution of underperformance possible.

Operationally this turns the expected shading loss from a design-time study into a term that can be recomputed continuously inside an Energy Performance Index, per string and per timestep, across a portfolio. An implementation has been prepared for the open-source pvlib library, with unit tests and a worked example, and is offered for community review and independent validation.

8 REFERENCES

[1] A. Skartveit, J. A. Olseth, M. E. Tuft, Solar Energy 63 (1998) 173.
[2] J. E. Hay, J. A. Davies, Proceedings First Canadian Solar Radiation Data Workshop, Toronto (1980).
[3] K. Anderson, C. Hansen, W. Holmgren, A. Jensen, M. Mikofski, A. Driesse, Journal of Open Source Software 8 (2023) 5994.
[4] J. W. Bishop, Solar Cells 25 (1988) 73.
[5] F. Martínez-Moreno, J. Muñoz, E. Lorenzo, Solar Energy Materials and Solar Cells 94 (2010) 2298.
[6] M. Oliosi, B. Wittmer, A. Mermoud, Proceedings 38th European Photovoltaic Solar Energy Conference (2021).
[7] C. Deline, A. Dobos, S. Janzou, J. Meydbray, M. Donovan, Solar Energy 96 (2013) 274.
[8] Y. M. Saint-Drenan, T. Barbier, Solar Energy 184 (2019) 127.

10.4229/EUPVSEC2026/4BO.6.3
020243-006