Photodiode Capacitance Is Not One Number

Why reading the zero-bias figure from a datasheet can cost you a factor of four in bandwidth

Jezdan Gomez
Opto Diode

Abstract.
The same large-area silicon photodiode can legitimately specify 1.0 nF of capacitance at zero bias and 260 pF at 12 V reverse bias (a ratio of nearly four) because junction capacitance tracks the depletion width, which grows as the square root of applied voltage. Engineers who take the first capacitance figure they see and compute an *RC* bandwidth are therefore routinely wrong by 4×, and the error runs the dangerous way: at low bias the detector is *slower* than calculated, not faster. This note develops the depletion physics behind the bias dependence, shows that the simple $t_r = 2.2RC$ model reproduces a published 30 ns rise-time specification to within 5%, tabulates capacitance, rise time, and bandwidth versus bias, and quantifies what happens when a short optical pulse meets a detector whose *RC* time is too long: the waveform is integrated, and the peak reading can fall to less than half the true value.

1. One device, many capacitances

A photodiode datasheet does not state *the* capacitance of the device; it states the capacitance under particular conditions (a bias voltage and a test frequency), and the value moves strongly with both. For a representative 100 mm² silicon photodiode, published typical values are 1.0 nF at 0 V and 260 pF at 12 V reverse bias: a factor of 3.8 between two rows of the same table. Since every first-order speed estimate scales directly with capacitance, an engineer who designs against the wrong row inherits that full factor in rise time and bandwidth.

The direction of the error deserves emphasis. Photodiodes are frequently operated photovoltaically (zero bias) for the sake of low dark current and low *1/f* noise. An engineer who computes bandwidth from a capacitance specified at 10 or 12 V and then runs the part unbiased has built a receiver four times slower than the calculation promised. Reducing bias *increases* capacitance; the mistake never fails in the safe direction.

2. Where the bias dependence comes from

A reverse-biased p–n junction is a parallel-plate capacitor whose dielectric is the depleted silicon itself. For active area A and depletion width W,

$$C_j = \frac{\varepsilon_0 \varepsilon_{\mathrm{Si}} A}{W} \tag{1}$$

with $\varepsilon_{Si} \approx 11.7$. The plate spacing W is not fixed by geometry; it is set by electrostatics. For a one-sided abrupt junction with light-side doping density N,

$$W = \sqrt{\frac{2\varepsilon_0 \varepsilon_{Si}\left(V_{bi} + V_R\right)}{qN}} \tag{2}$$

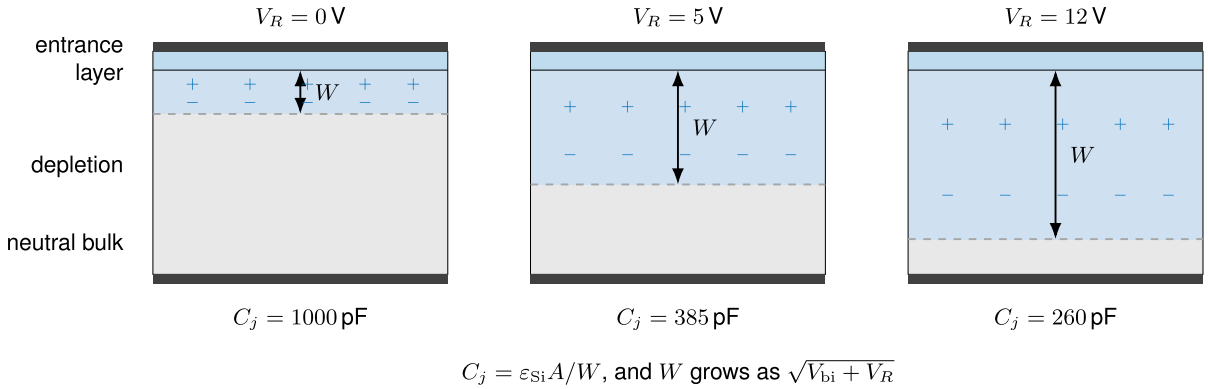

Figure 1. Three cross-sections of the same 100 mm² junction at 0, 5 and 12 V reverse bias, with the depletion width drawn to scale as $1/C_j$ and each panel labeled with its bias and computed junction capacitance. The fixed charge density at the edge of the depleted region is identical in all three panels: reverse bias does not add charge, it pushes that boundary outward, which is the physical picture behind treating the junction in Eq. (1) as a parallel-plate capacitor whose spacing grows with bias.

where V_{bi} is the built-in potential (roughly 0.6–0.9 V in silicon) and V_R the applied reverse bias. Combining Eqs. (1) and (2), the capacitance of a partially depleted junction follows

$$C_j\left(V_R\right) = C_j\left(0\right)\sqrt{\frac{V_{bi}}{V_{bi} + V_R}} \tag{3}$$

Two features of Eq. (3) explain the datasheet numbers. First, because V_{bi} is under a volt, the denominator changes fastest near zero: the first volt or two of reverse bias does most of the work, and the curve is steepest exactly where photovoltaic users operate. Second, the published pair of values is quantitatively consistent with the square-root law: an effective V_{bi} of 0.87 V reproduces both 1.0 nF at 0 V and 260 pF at 12 V exactly, and the textbook value of 0.7 V predicts 235 pF at 12 V, within 10% of the published figure. Equivalently, via Eq. (1), the published capacitances imply the depletion region widens from roughly 10 µm at zero bias to roughly 40 µm at 12 V.

2.1 *The knee: full depletion*

Equation (2) cannot hold indefinitely. Once the depletion region reaches the far boundary of the active layer, W is pinned at the layer thickness, and by Eq. (1) the capacitance stops falling. Beyond this knee, additional bias buys essentially no additional speed: the capacitance–voltage curve goes flat. The bias at which the knee occurs depends on the thickness and resistivity of the active layer and therefore varies from product to product; the datasheet capacitance-versus-voltage curve reveals it directly as

the flattening of the characteristic. Bias beyond full depletion is not wasted (it maintains drift field throughout the volume, which benefits response uniformity and high-current linearity), but the capacitance dividend is exhausted.

3. From capacitance to speed

When the response is limited by the circuit time constant rather than by carrier transit or diffusion (the usual case for a large-area diode into a load), the 10–90% rise time and the −3 dB bandwidth follow directly from the single-pole *RC* response:

$$t_r \approx 2.2\left(R_L + R_s\right)C_j \tag{4}$$

$$f_{3\mathrm{dB}} \approx \frac{0.35}{t_r} \tag{5}$$

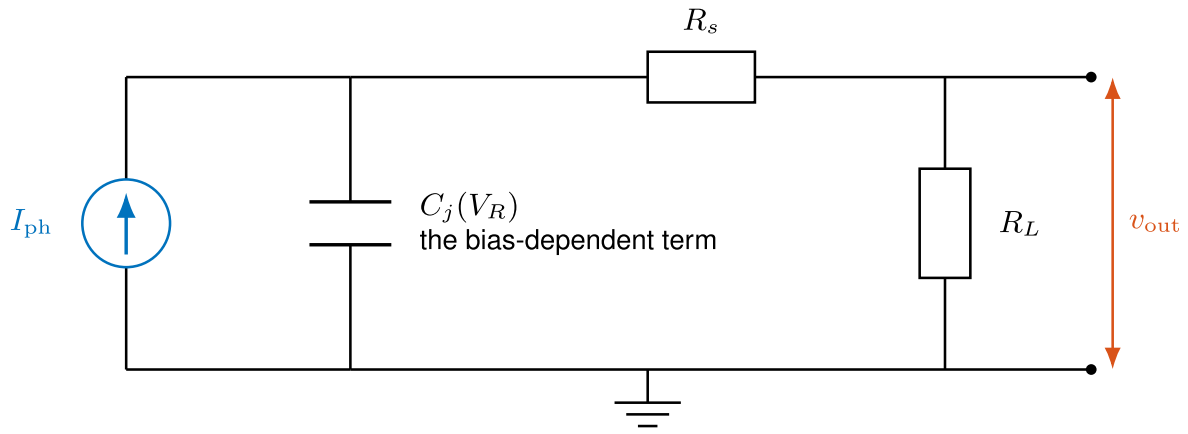

Figure 2. The small-signal circuit behind Eqs. (4) and (5): photocurrent I_{ph} drives the junction capacitance $C_{\mathrm{j}}(V_{\mathrm{R}})$, which charges through the series resistance R_{s} and load R_{L} to produce the output v_{out}. Because C_{j} is the only bias-dependent term in the loop, t_{r} and $f_{3\mathrm{dB}}$ both scale directly with it: whatever factor bias buys in capacitance, it buys in rise time and bandwidth as well.

Both scale linearly with C_{j}, so both inherit the full bias dependence of Eq. (3). Table 1 carries the representative 100 mm² device through Eqs. (3)–(5) into a 50 Ω load, using the effective V_{bi} = 0.87 V that matches both published endpoints; the intermediate rows are interpolated from the square-root law.

Table 1. Capacitance, rise time, and bandwidth versus reverse bias for a 100 mm^2 silicon photodiode into R_L = 50 Ω (R_S neglected). Endpoint capacitances are published typical values; intermediate rows follow Eq. (3).

Bias (V)	C_j (pF)	$t_r = 2.2RC$ (ns)	f_{3dB} (MHz)
0	1000	110	3.2
0.5	797	88	4.0
1	682	75	4.7
2	551	61	5.8
5	385	42	8.3
12	260	28.6	12.2

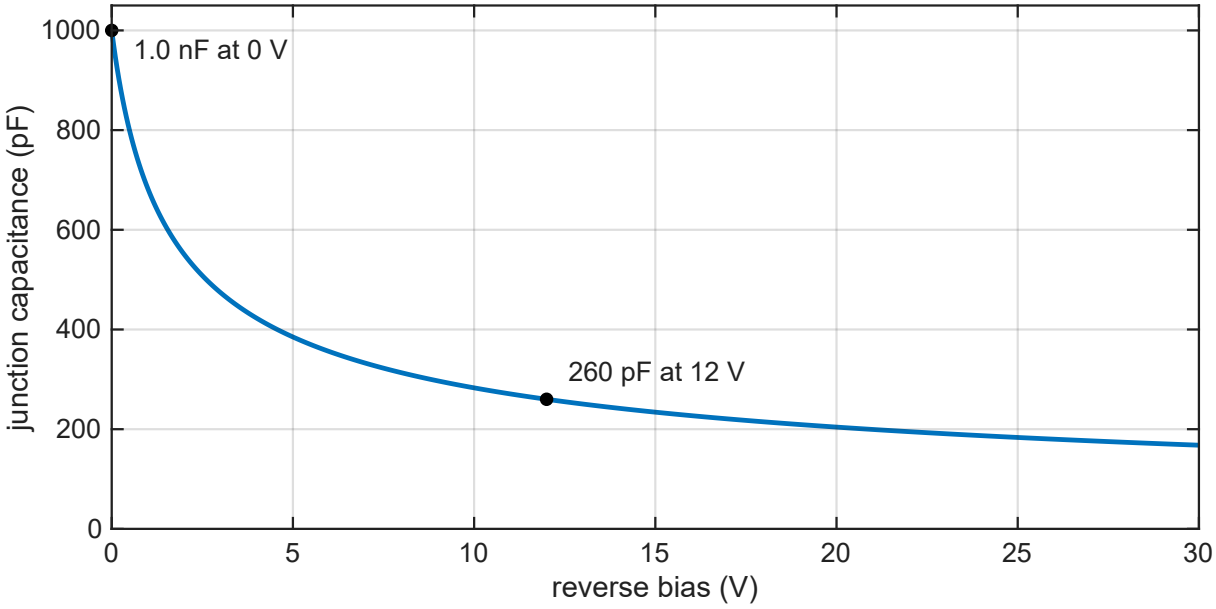

Figure 3. Junction capacitance against reverse bias for a 100 mm^2 silicon photodiode, from Eq. (3) with the curve anchored to the two published specification points (marked). The curve is steepest below about 2 V (precisely where photovoltaic-mode users operate), so a small change in bias there produces a large change in capacitance.

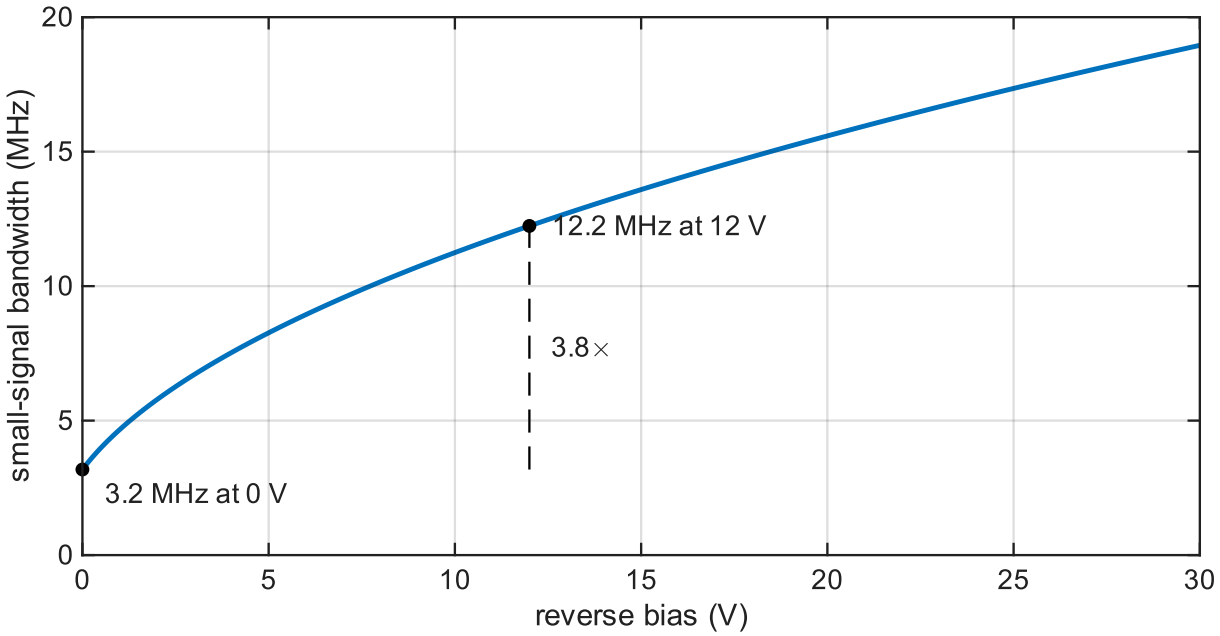

Figure 4. Small-signal bandwidth $f_{3\mathrm{dB}} = 0.35/t_\mathrm{r}$ against reverse bias, for the same device into a 50 Ω load. Bandwidth tracks $1/C$ and therefore rises as the square root of bias: reading the zero-bias capacitance from a datasheet and applying it at 12 V understates the bandwidth by a factor of approximately 4.

The two endpoints of Table 1 are the whole story of the 4× error: 12.2 MHz at 12 V against 3.2 MHz at zero bias, a ratio of 3.8. A design that borrowed the 0 V capacitance for a 12 V calculation (or, more dangerously, the 12 V capacitance for a 0 V circuit) misses by that factor.

Worked check: does the simple model match the datasheet?. The same device publishes a typical rise time of about 30 ns into 50 Ω at 12 V bias. Equation (4) with the published 12 V capacitance gives $t_\mathrm{r} = 2.2 \times 50\ \Omega \times 260\ \mathrm{pF} = 28.6\ \mathrm{ns}$, within 5% of the published figure before series resistance is even included (adding a plausible $R_\mathrm{s} = 5\ \Omega$ brings it to 31.5 ns, bracketing the specification). The lesson: for large-area diodes the lumped RC model is not a rough sketch, it is the specification. Whatever capacitance you insert into it is the speed you will get, so insert the one measured at your operating bias.

4. When the pulse is faster than the detector

The consequence of an overlooked factor of four is mild in a CW measurement and severe in a pulsed one. If the circuit time constant $\tau = (R_\mathrm{L} + R_\mathrm{s})C_\mathrm{j}$ is comparable to or longer than the optical pulse width τ_p, the detector does not resolve the pulse; it integrates it. For a rectangular optical pulse that would produce photocurrent I_0 in a fast circuit, the output only attains

$$I_\mathrm{pk} = I_0\left(1 - e^{-\tau_p\,/\,\tau}\right) \tag{6}$$

before the light switches off and the output decays. The *collected charge* is still correct: the photogenerated carriers all arrive eventually, so an integrating (charge) measurement survives. The

peak reading and the apparent waveform, however, do not. Concretely, for the Table 1 device viewing a 30 ns optical pulse into 50 Ω: at 12 V, τ = 13 ns and the output reaches 90% of the true peak; at zero bias, τ = 50 ns and it reaches only 45%. The same detector and the same pulse produce a factor-of-two disagreement in measured peak power, traceable entirely to which capacitance row governed the operating point. Even a 100 ns pulse attains only 86% of its true peak at zero bias (versus >99.9% at 12 V), an error that looks like a calibration problem and is actually an *RC* problem.

5. The area trade, and one more fine print

Equation (1) contains a second, equally underused lever: capacitance is proportional to area. A 10 mm² device with the same vertical structure as the 100 mm² example carries one tenth the capacitance: 26 pF at 12 V, or a 2.9 ns rise time and 120 MHz into 50 Ω. No bias adjustment can match a 10× area reduction. Choosing the smallest active area that reliably captures the beam is the single most effective speed decision available, ahead of bias, load resistance, or amplifier choice. Large-area photodiodes are intrinsically slow; that is the price of optical collection, not a defect of any particular part.

The measurement frequency matters too. Junction capacitance is specified at a stated test frequency (commonly 100 kHz or 1 MHz), and the two need not agree. Slow trapping responses, conductance in partially depleted material, and package parasitics all make the small-signal capacitance mildly frequency-dependent. When comparing parts across manufacturers, or comparing a measurement to a datasheet, confirm both the bias *and* the test frequency of the quoted value before concluding that anything is out of specification.

Verdict. A photodiode's capacitance is a curve, not a constant: it falls as $1/\sqrt{(V_{bi} + V_R)}$ until the active layer fully depletes, and for a representative 100 mm² device that curve spans 1.0 nF to 260 pF (a 3.8× spread that flows straight through $t_r = 2.2RC$ into bandwidth). Read the capacitance at *your* operating bias, verify that its *RC* time is short compared with the pulses you must resolve, and buy speed first with area, then with bias, remembering that bias helps only up to the full-depletion knee.

Notes

Symbols. C_j junction capacitance; A active area; W depletion width; ε_0 vacuum permittivity; $\varepsilon_{Si} \approx 11.7$ relative permittivity of silicon; q elementary charge; N doping density of the lightly doped side; V_{bi} built-in potential; V_R applied reverse bias; R_L, R_s load and series resistance; t_r 10–90% rise time; f_{3dB} −3 dB bandwidth; $\tau = (R_L+R_s)C_j$ circuit time constant; τ_p optical pulse width; I_0, I_{pk} ideal and attained peak photocurrent.

Basis. Endpoint capacitances (1.0 nF at 0 V, 260 pF at 12 V), series resistance range (0.5–5 Ω), and the ~30 ns rise time are published datasheet-level typical values for a 100 mm^2 silicon photodiode; intermediate-bias rows in Table 1 are interpolated from the one-sided abrupt-junction law of Eq. (3) with an effective $V_{bi} = 0.87$ V fitted to the two published endpoints. Rise-time and bandwidth figures assume a single-pole RC response into 50 Ω and neglect carrier transit and diffusion, which is appropriate for circuit-limited large-area devices but optimistic for small, fast ones. The bias at which a given product fully depletes depends on the thickness and resistivity of its active layer; consult the product's capacitance–voltage curve.

References

1. S. M. Sze and K. K. Ng, *Physics of Semiconductor Devices*, 3rd ed., Wiley, 2007 (abrupt-junction depletion and capacitance).
2. J. G. Graeme, *Photodiode Amplifiers: Op Amp Solutions*, McGraw-Hill, 1996, ISBN 0-07-024247-X (photodiode capacitance in receiver design).
3. B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, 2nd ed., Wiley, 2007 (photodetector response and bandwidth).
4. Manufacturer datasheet, large-area (100 mm^2) silicon photodiode: capacitance vs. bias and rise-time specifications (typical values).