Three diode models

You've almost certainly been told that a diode drops 0.7 V, and the last page showed that it's really an exponential, with no fixed voltage at all. You'll use both descriptions in this course, and the point is to choose between them on purpose, once you know what each one costs. A model is a simplification that comes with an error bar, and choosing one is an engineering decision you'll make several times a day.

This page covers the three models you'll use, from the crudest to the exact one.

Model 1: the ideal switch

Forward: a wire, 0 V. Reverse: an open circuit.

This is the model you run in your head while reading a schematic, to answer "which way does the current go?". There's no voltage or current in it, so the only thing it claims to tell you is the **direction** of the current, and it gets that right.

Model 2: the constant drop

Forward: a 0.7 V battery. Reverse: an open circuit.

This is the one everybody quotes, and it works because of the number you saw on the last page, **77.04 mV per decade**. In the decade table, going from 1 nA to 1 µA multiplies the current by a thousand, while the voltage only goes from 0.419503 V to 0.651480 V, a change of 232 mV.

A device whose current can change by three orders of magnitude while its voltage changes by less than a quarter of a volt looks like a battery if you stand far enough back, and Model 2 is the diode seen from that distance.

Model 3: the exponential

$I = I_0\left(e^{V/nV_T} - 1\right)$*, solved together with the rest of the circuit.*

This one is exact, but you can't solve it by rearranging the equation, so it always comes with a numerical method: either your own three lines of iteration or SPICE's.

What each model costs

Take a single resistor R from 1.8 V into a 1 µm² diode. Model 1 puts no voltage across the diode, so it says the current is $1.8/R$. Model 2 takes 0.7 V off the supply and says $(1.8 - 0.7)/R$. Model 3 is ngspice's operating point, from The operating point, and the two error columns compare the other models with it.

R	Model 1: switch	Model 2: 0.7 V	Model 3: ngspice	Model 1 error	Model 2 error
10 kΩ	180.00 µA	110.00 µA	**90.94642 µA**	+97.9 %	+20.9 %
100 kΩ	18.000 µA	11.000 µA	**10.60153 µA**	+69.8 %	+3.8 %
1 MΩ	1.8000 µA	1.1000 µA	**1.143884 µA**	+57.4 %	−3.8 %

Two of Model 2's errors are under 4 %, but "Model 2 is good" is the wrong conclusion to draw from that. Three other things in this table are worth noticing.

Model 2's error changes sign. It's 3.8 % high at 10 µA and 3.8 % low at 1 µA, because 0.7 V is just where this particular junction happens to sit at a few microamps. Above that current the real voltage is higher than 0.7 V, and below it the real voltage is lower. The constant drop is a tangent point rather than a bound: it's right at one current and wrong on either side of it, so it can't tell you which side of the answer you're on. For a diode running at 100 µA, 0.7 V is a bad guess, and for one at 1 nA it's a terrible one, because the decade table puts that diode at 0.419503 V.

Model 2 is worst at the largest current. At that current the diode's series resistance (979.1 Ω) adds a drop of its own on top of the junction's, which Model 2 has no way to include. At 10 kΩ the model is out by a fifth.

Model 1 is at least 57 % wrong at every resistor value in this circuit, and it's still the right model for reading a schematic, because it was never meant to tell you how much current flows, only which way.

Which one to use

You are asking	Use
which way does current flow; is this branch on or off	**Model 1**
roughly what current, roughly what headroom, on the back of an envelope	**Model 2**, and say "about" out loud
what does this actually do; will it still work at 0 °C; how much does it drift	**Model 3**, in a simulator

> **When to stop using 0.7 V.** If a hand calculation used 0.7 V and its answer matters to better than about 10 %, redo it in the simulator. 0.7 V is a one-significant-figure model, and a two-significant-figure result that came out of it isn't worth defending.

The small-signal model

There's also a fourth model, and it's the most useful of all. You replace the diode with a resistor whose value is $\Delta V / \Delta I$ at the operating point (the slope of V against I, which is the reciprocal of the slope of the usual I–V plot), and that resistor is the **small-signal model**. It's valid only for changes small enough that the curve looks straight over them, and it's how every amplifier in this course is analysed. You've already met one of these resistors: 30.3 kΩ at the 1 MΩ operating point, on The operating point.

What junction area buys

In AD102 a resistor's value turned out to be a matter of geometry: so many squares of doped silicon. A diode is drawn the same way, out of the same kinds of shapes, so it's natural to ask what the shape of a diode buys you.

Lab 02 puts four square junctions (1×1, 2×2, 4×4 and 10×10 µm) at the same 0.500 V, and they carry these currents:

area (µm²)	current at 0.500 V
1	11.09698 nA
4	34.57181 nA
16	118.6544 nA
100	667.9656 nA

If current were proportional to area, the 10×10 junction, with a hundred times the area of the 1×1, would carry a hundred times the current. In fact it carries **60.19** times as much.

To see where the rest went, divide each current by its own area and put the result next to the junction's perimeter-to-area ratio:

area (µm²)	perim (µm)	perim / area	current / area (nA/µm²)
1	4	4.0	11.09698
4	8	2.0	8.6429525

area (µm²)	perim (µm)	perim / area	current / area (nA/µm²)
16	16	1.0	7.415900
100	40	0.4	6.679656

Take **only the first and last rows**, leaving the middle two out, and fit a straight line through them:

$$\frac{I}{A} = 6.188848 + 1.2270344 \times \frac{P}{A} \ \text{nA/µm}^2$$

Then use the line to predict the two rows you left out. At $P/A = 2$ it gives **8.642917**, against a measured 8.6429525, and at $P/A = 1$ it gives **7.415882**, against a measured 7.415900. The line matches those two rows to five significant figures, even though it was never fitted to them.

A straight line like that means the diode behaves as two devices in parallel: a **face**, whose current is proportional to area, and an **edge**, whose current is proportional to perimeter. SkyWater's model card describes the diode the same way, with `js` in amps per unit area and `jsw` in amps per unit length, and these four measurements reconstruct both terms without being told either one.

This gives you a fabrication fact worth remembering: **a small junction leaks more per square micron than a big one**, because perimeter shrinks with the linear dimension while area shrinks with its square. For the same reason, two diodes with the same total area don't behave identically if one is a long thin strip and the other is a square.

You can also turn the 60.19× around. At a *fixed current*, the diode with a hundred times the area sits lower by $77.04 \times \log_{10}(60.19) = 137$ mV, so making a junction a hundred times bigger buys you only **about a seventh of a volt**. Area is a very expensive way to move a diode's voltage, and most analog designers learn this once, by drawing something enormous and getting almost nothing for it.

You'll run this measurement yourself in Lab 02, Part 7, along with three ways of getting a diode's size wrong without any warning.

Check yourself

1. A diode carries 100 µA. Is Model 2's 0.7 V likely to be too high or too low? Use the decade table to justify your answer.
2. Using the fitted line above, what would the current per µm² be for a junction with $P/A = 0.1$?
3. You need to lower a diode's voltage by 77 mV at fixed current. Roughly how much bigger must the junction be, if area alone set the current? Why is the real answer bigger still?

Next: Lab 02: The diode I–V curve.