Parasitic capacitance

If a capacitor is just two conductors near each other, then every wire on a chip is a capacitor, with capacitance to the substrate, to the wires above and below it, and to the wires running alongside it. You didn't draw any of these capacitors, and you can't delete them. They're called **parasitics**, and on a modern chip they are usually the largest capacitance in the design.

Parasitics are why digital chips have a maximum clock rate and why a layout that passes DRC can still be slow, and they're the reason AD104 exists.

Parasitic capacitance in the PDK

The PDK tabulates these capacitances in `/foss/pdks/sky130A/libs.tech/ngspice/parameters/typical.spice`, on lines 2876–2902. The area terms are in F/m² and the sidewall terms are in F/m:

```
+ cp1f = 1.06e-04  cp1fsw = 8.64e-11
+ cl1f = 3.69e-05  cl1fsw = 8.30e-11
+ cm1f = 2.58e-05  cm1fsw = 1.07e-10
+ cm2f = 1.75e-05  cm2fsw = 1.08e-10
+ cm5f = 6.48e-06  cm5fsw = 7.85e-11
+ cm2m1 = 1.28e-04 cm2m1sw = 1.03e-10
```

Here are the same numbers converted to the units a layout engineer works in (1 F/m² = 1e-12 F/µm², 1 F/m = 1e-6 F/µm):

coupling	area	fringe, per edge
poly ⟶ substrate (`cp1f`)	0.106 fF/µm²	0.0864 fF/µm
li ⟶ substrate (`cl1f`)	0.0369 fF/µm²	0.0830 fF/µm
met1 ⟶ substrate (`cm1f`)	0.0258 fF/µm²	0.107 fF/µm
met2 ⟶ substrate (`cm2f`)	0.0175 fF/µm²	0.108 fF/µm
met5 ⟶ substrate (`cm5f`)	0.00648 fF/µm²	0.0785 fF/µm
met2 ⟶ met1 (`cm2m1`)	0.128 fF/µm²	0.103 fF/µm

For comparison, a MIM capacitor is 2.00 fF/µm², and met1 to substrate is 0.0258 fF/µm², which is 78 times less dense. That's the way round you want it, since the capacitor you built on purpose ought to beat the accidental one. Still, 78 times less dense is a long way from zero, and a chip has an enormous amount of wire.

Fringe capacitance of a thin wire

Take a met1 wire of minimum width, which the PDK gives one directory up, in `libs.tech/ngspice/sky130_fd_pr__model__r+c.model.spice`, as `wminm1= 0.14u`, and run it 1000 µm across the chip. Its capacitance to the substrate has two parts:

- **Plate area:** $1000 \times 0.14 = 140$ µm², times 0.0258 fF/µm² = 3.61 fF
- **Fringe:** two edges, 1000 µm each, times 0.107 fF/µm = 214.0 fF

The total is about 217.6 fF, and the parallel-plate part, the one you would have calculated from $\varepsilon A/d$, is only 1.7 % of it.

Almost everyone's first estimate of a wire's capacitance comes from its area, because that's the formula from the previous page, but for a minimum-width wire that estimate is off by a factor of sixty. The reason is geometric: the wire is 0.14 µm wide and sits on the order of a micron above the substrate, so it's more like a string than a plate,

and most of its field leaves through the sides rather than the bottom. `cm1fsw` is the capacitance of that sideways field, as SkyWater measured and tabulated it, and it's the reason extraction tools exist.

As a rule, for anything narrower than it is tall, the sidewall capacitance dominates, so use the `*sw` term rather than the area term.

What a long wire costs

Now put that wire on the output of the smallest inverter, `sky130_fd_sc_hd__inv_1`, which comfortably drives a few femtofarads of gate load. Hang 217.6 fF on it instead, roughly fifty times its normal load, and the rise time goes up by roughly the same factor. That's the mechanism behind the saying that long wires are slow, and it's why place-and-route tools spend most of their effort on wire length rather than gate count.

The same wire also has resistance. 1000 µm of minimum-width met1 is $1000/0.14 = 7143$ squares, and at 0.125 Ω/□ that comes to 893 Ω (see Ohms per square). The resistance and the capacitance are both spread along the wire's whole length, so the wire behaves as a distributed RC line rather than a single node, and the delay through it grows as the square of its length.

In Lab 02, Part 9 (optional) you can draw two wires in SiliWiz and watch their coupling capacitance change as you move them.

VPP capacitors and parasitic extraction

Parasitic capacitance has two more consequences. The first is that it's what makes a VPP capacitor work. Ordinary metal wires couple to their neighbours at about 0.1 fF per micron of edge whether you like it or not, so you can use that coupling on purpose: interleave enough fingers and the accident becomes the component. On Capacitors on a chip, `sky130_fd_pr__cap_vpp_11p5x11p7_l1m1m2m3m4_shieldm5` measured 147.3400 fF, and it's built entirely out of wire.

The second is that every number on this page came from a pre-characterised table, because the real value depends on whatever happens to be near the wire, and a table can't know that. A **parasitic extraction** (PEX) run reads your finished layout, works out these couplings for the shapes you actually drew, and hands back a netlist with hundreds of capacitors your schematic never had. In AD104 you run one and compare the circuit before and after extraction.

At some point every chip designer sees a circuit that simulated perfectly come back from extraction 40 % slower, with nothing in the schematic changed. We wrote this page so that when it happens to you, you'll know why.

Next: Inductors on a chip.