What a picofarad costs

The capacitors in a chip schematic are given in femtofarads, while the ones in your parts drawer are in microfarads. The reason is that on a chip you pay for capacitance in square micrometres of area, and the exchange rate is poor.

In Capacitors on a chip we measured a MIM capacitor rather than quoting a number for it, and it gives about 2 fF per μm². This page works out what that density means for a floor plan.

The exchange rate

you want	MIM plate area	plate is about	in `sky130_fd_sc_hd__inv_1` (3.7536 μm²)
100 fF	49 μm²	7 × 7 μm	13
1 pF	494 μm²	22 × 22 μm	132
10 pF	4 940 μm²	70 × 70 μm	1 316
100 pF	49 400 μm²	222 × 222 μm	13 160
1 nF	494 000 μm²	703 × 703 μm	131 600
1 μF	494.6 mm²	*a plate 2.2 cm on a side*	132 million

These numbers come from the measured 30 × 30 μm device: 1819.782 fF over 900 μm² is 2.0220 fF/μm², so 1 pF needs $1000/2.0220 = 494.6$ μm². The last column gives the same area as a count of `sky130_fd_sc_hd__inv_1` inverter cells, each 3.7536 μm².

Compare the last row with the 100 nF ceramic capacitor you would solder next to a microcontroller without thinking about it. A typical small SKY130 die is about 3 mm × 3 mm, or 9 mm², and the 1 μF capacitor needs 494.6 mm², which is the area of fifty-five entire chips for one component. On-chip capacitors are small because a 1 μF capacitor physically can't fit on a die, and no amount of cleverness closes a gap that size.

Switching to the densest option doesn't get you there either. A MOS capacitor at 7.88 fF/μm² brings 1 μF down to 126.9 mm², which is still fourteen whole dies, and its value changes by 3.57× with the voltage across it.

What sizes are practical

capacitance	typical use
1–50 fF	sampling capacitors in an ADC, amplifier compensation, charge redistribution
50–500 fF	filter capacitors, matched arrays, bandgap compensation
1–20 pF	serious on-die decoupling, a big integrator, an off-chip driver's load
> 50 pF	you are budgeting a large fraction of the die for one component

These limits explain a lot about how analog IC design looks. A discrete designer builds a 1 kHz low-pass filter from 16 kΩ and 10 nF, but that capacitor would take 4.94 mm² on a die, which is more than half of a typical small die. So the on-chip designer doesn't build that filter. Instead they build a switched-capacitor filter, or use a very

large resistance made from a starved transistor, or push the corner frequency up by three decades and filter the rest off-chip. For more on this, see What to build instead of an inductor.

Density and plate size

Divide each measured MIM capacitor by its drawn area (1, 100 and 900 µm²):

```
c_mim_1x1 = 2.642250e-15
c_mim_10x10 = 2.065822e-13
c_mim_30x30 = 1.819782e-12
```

You get 2.64, 2.07 and 2.02 fF/µm², so the density falls as the capacitor gets bigger.

The extra capacitance in the small one comes from the field at the plate's edge, which the model adds as `cpmimc = 0.19e-15` farads per micron of perimeter. Perimeter grows with the side length while area grows with its square, so the edge adds 24 % at 1 µm but only 1 % at 30 µm. **Never estimate a small capacitor from area alone.** In Lab 02, Parts 1 and 2, you measure a ladder of plate sizes and separate the area term from the edge term yourself.

Worked example: one plate or a hundred

You need 1 pF, and you have room for a single 23 × 23 µm MIM. Alternatively, you could place one hundred 2.3 × 2.3 µm MIMs in parallel, which uses the same 529 µm² of plate. Which gives more capacitance? Work it out before reading on.

Most people say "the same", but the perimeter term makes them differ. One big plate gives $2.00 \times 22.975^2 + 0.19 \times 2 \times (22.975 + 22.975) = 1055.70 + 17.46 = 1073.2$ fF. Each of the hundred small plates gives $2.00 \times 2.275^2 + 0.19 \times 2 \times (2.275 + 2.275) = 10.35 + 1.729 = 12.08$ fF, and a hundred of them give 1208 fF. Cutting the plate into pieces gives you 12.6 % more capacitance for the same plate area, because together the pieces have a lot more edge than the one big plate.

Designers really do use this, and it's the reason VPP capacitors exist and the `sky130_fd_pr__cap_vpp_*` structures are drawn as interleaved fingers rather than slabs. In practice, a hundred small plates also need a hundred sets of connections and a lot of spacing between them, so nobody takes it to that extreme; the foundry took the idea as far as it goes when it drew the VPP cells.

Costs besides area

A capacitor also has three costs that don't show up in its value in farads.

Bottom-plate parasitic. A MIM plate is a large sheet of metal above the substrate, so it has its own capacitance to ground whether you asked for it or not, and the `.subckt` has a third terminal for exactly this. In a switched-capacitor circuit this parasitic is the main source of error, and much of the design work goes into arranging for it to be driven by something that doesn't care about it.

Series resistance. The MIM `.subckt` has two resistors in series with the capacitance:

```
rs1 a  b1 'r1' tc1 = {tc1rm3} tc2 = {tc2rm3}
rs2 b1 c1 'r2' tc1 = {tc1rvia3} tc2 = {tc2rvia3}
```

`r1` is the sheet resistance of the met3 plate itself, and `r2` is the resistance of the via array. So what you've drawn is really a capacitor in series with a small resistor, and at high frequency the resistor dominates.

Corner spread. The MIM's area capacitance is `2.00e-15` F/µm² at the typical corner, but `1.778e-15` at the low corner and `2.231e-15` at the high one, a spread of about ±11 %. Process corners covers what to do about that.

Reading capacitor values

When an analog paper says “C_1 = 240 fF”, that capacitor is 119 μm² of somebody’s floor plan, and the area was argued over in a design review. Area is the currency of this whole field, and apart from inductors, capacitors are the most expensive thing you can buy with it.

Next: Parasitic capacitance, the capacitors nobody drew in the schematic.