What to build instead of an inductor

Without a usable inductor, you might wonder how anyone builds a filter, an oscillator or a power supply on a chip. The answer is to work out what the inductor was doing in the circuit and get that some other way. There are four standard ways to do it, and the first three are used all the time.

1. Use R and C, at a different frequency

Most of what an inductor does in a discrete design, such as setting a corner frequency or shaping a response, an RC network can also do, with one fewer pole per stage. The catch is that an RC corner depends on R × C, and you already know what each of those costs in area:

$$f_{-3\ \mathrm{dB}} = \frac{1}{2\pi RC}$$

A 1 kHz corner needs $RC = 159$ µs. If you use the 100 kΩ resistor from What a resistor costs in area (314 µm² of poly), the capacitor has to be 1.6 nF, and that capacitor alone is 791 300 µm², or 0.79 mm². That's about 9 % of a 3 mm × 3 mm die, or the area of 210 800 inverters, for a single pole, so nobody builds one.

So on a chip, the designer moves the corner instead. The corners deck in `labs/passives-decks` builds an RC low-pass from one 3550 Ω resistor and one 206.5822 fF MIM capacitor:

```
--- tt : RC corner frequency, Hz ---
f3db                =  2.15410e+08
```

That's 215.410 MHz. With a resistor and a capacitor of a size you can afford, the corner lands five orders of magnitude above the 1 kHz you wanted, so anything slower than a few megahertz is filtered off-chip, or with one of the next two techniques. Lab 03 builds an RC from the resistor and capacitor you designed and measures its time constant in the time domain.

2. Switched capacitors

Switched capacitors are the most important trick in analog IC design, and they exist because of the two facts you have just learned about what resistors and capacitors cost in area.

Take a capacitor C and a two-phase clock at frequency f_{clk}. Connect C alternately to node A and to node B. Each cycle it carries a charge $Q = C\,\Delta V$ from one node to the other, so the average current is $I = f_{clk}\,C\,\Delta V$. That is Ohm's law with

$$R_{\mathrm{eq}} = \frac{1}{f_{clk}\,C}$$

So a capacitor plus a switch acts as a resistor, and its value is set by a clock frequency rather than by a length of poly.

To put numbers on it, take a 100 fF capacitor (about 49 µm², the area of thirteen inverters) and clock it at 1 MHz:

$$R_{\mathrm{eq}} = \frac{1}{10^{6} \times 100 \times 10^{-15}} = 10\ \mathrm{M}\ \Omega$$

So you get 10 MΩ in 49 µm². Building the same 10 MΩ as a 1 µm wide `res_xhigh_po` resistor would take a strip $(10^{7} - 34.2111)/2118.7619 = 4719.7$ µm long, which is 4.72 millimetres (wider than most dies) and 4720 µm² of poly. The switched capacitor uses about ninety-six times less area, and unlike a 4.72 mm strip, it's something you can actually place.

Switched capacitors have two more advantages:

- **The value is tunable.** Changing the clock changes the resistance, so you can retune it without new masks.
- **The time constant becomes a capacitor ratio.** An RC made from a switched capacitor C_1 and an integrating capacitor C_2 has $\tau = R_{eq}C_2 = C_2/(f_{clk}C_1)$, which is a ratio of capacitors times a clock period. A ratio of matched devices is the most accurate thing you can make on a chip (see Matching and Pelgrom's law), and that's why almost every ADC you will meet is built from switched capacitors.

The price is that the circuit is now sampled, so it only works below $f_{clk}/2$, and it folds noise. AD202 covers both of those properly.

3. Gyrators

An inductor obeys $V = L\,dI/dt$ and a capacitor obeys $I = C\,dV/dt$: the same equation with current and voltage swapped. A circuit that swaps current and voltage, called a **gyrator** and built from two transconductors, turns a capacitor at one port into an inductance at the other:

$$L_{\text{eq}} = \frac{C}{g_{m1}g_{m2}}$$

With $g_m = 100$ µS and $C = 1$ pF you get $L_{eq} = 100$ µH, which is ten thousand times the biggest spiral in the PDK, from about 500 µm² of MIM plus two small amplifiers.

The catch is that a gyrator is *active*: it needs supply current, adds noise, has a limited signal swing, and stops working above the transconductors' bandwidth. So what you get is a circuit that behaves like an inductor over a limited range, which is usually enough. Filters built this way, called "active inductor" or "simulated inductor" filters, are standard practice at audio and low RF.

4. Put the inductor off the die

If you really need an inductor (a switching regulator, an RF matching network, a high-current supply), you buy one and put it on the board. Two on-die pads and a package pin cost a few thousand µm² and get you a component with a Q of 40 and a value in microhenries.

Designers use the package itself, too. A bond wire has roughly 1 nH of inductance per millimetre, and RF designers deliberately make that inductance part of the matching network instead of fighting it.

When to use an on-chip spiral

The three SKY130 spirals still have their uses. Above about a gigahertz, where Q climbs past 10, an on-die inductor is the right answer, and often the only one, for jobs like these:

- the **LC tank** of a voltage-controlled oscillator, where Q sets the phase noise;
- **inductive peaking** on a wideband amplifier, where a couple of nanohenries extends the bandwidth by tens of percent;
- **impedance matching** to a 50 Ω antenna.

Those are the jobs the spiral cells are for. It's a narrow niche, though a real and well-paid one, and neither AD102 nor AD103 works in it.

What the four answers have in common

All four answers rest on the same principle:

> **On a chip, absolute values are expensive and ratios are cheap. So build what you need out of ratios, and get the absolute scale from something that is not silicon: a clock, a bias current, or a component on the board.**

A switched-capacitor filter's corner is a capacitor ratio times a clock. A gyrator's inductance is a capacitance over two transconductances, which are themselves set by a bias current. A bandgap reference's output is a resistor ratio times a physical constant.

The principle follows directly from a measurement you'll see on the next two guide pages, which show that on-chip absolute values are poor and on-chip ratios are essentially exact.

Next on the Course path: Lab 03: A time constant in silicon. After that, read Process corners.