How Many Bypass Capacitors Do We Need?

Istvan Novak

Welcome to the first article in the *QuietPower* columns.

In this introductory article I share my thoughts about a generic question that comes up in nearly all discussions among board designers: how many bypass capacitors do we need? As we usually say with my friend Eric, 'It depends'. However, we should be able to give a more specific answer, at least in general terms, by putting our current design constraints into historical perspective.

Take for instance the computer board shown on the photo below.

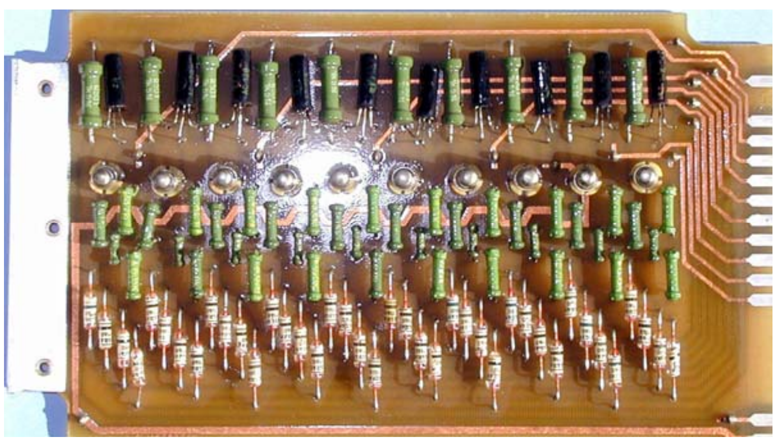

This is an old computer board with Germanium diodes and transistors, performing a simple DTL (diode-transistor-logic) gate function. It was made probably in the sixties or early seventies. I pass this board around in my power distribution design courses by raising the question: how many bypass capacitors can you count on this board? The answer is a surprising zero. Why is it then that today we are scrambling to find room on our boards for the many bypass capacitors we have to use to quiet our electronics? Whether we open a notebook computer or look at a large computer board, we can count hundreds, sometimes thousands of capacitors. To understand the reasons for this striking change that has happened over just a few decades, we have to look at the purpose of bypass capacitors and how bypass capacitors fulfill their roles.

Whether analog or digital, all electronic circuits eventually generate and/or process alternating electrical signals. During their operation, their need for supply current from the power source changes according to the actual function they happen to perform at the moment. Sometimes the circuit needs less current, sometimes more, and this varying demand changes the current in the supply rail. The changing current flows through the wires and printed-circuit board traces/planes between the source and the electronics, and it creates voltage changes:

$\Delta V = L\frac{dI}{dt}$, where ΔV is the voltage change, L is the inductance in the supply path; dI/dt is the rate of change of the current with time.

Small voltage fluctuations will be tolerated by the electronics, but each circuit has a maximum and minimum limit within which the correct operation is guaranteed. To limit the change in supply voltage that otherwise would occur due to the varying current demand, we have to place charge reservoirs between our circuit and the series inductance of the power path. A capacitor is a good charge reservoir, because as opposed to an inductor, the voltage across a capacitor (if we neglect its parasitic elements) will not change abruptly for a sudden surge of current.

To determine what kind and how many capacitors we need, it is more convenient to look at the problem in the frequency domain. A power distribution network (PDN) can be simulated in many different ways, starting from simple lumped equivalent circuits all the way to detailed grid models. For the purposes of our discussions here, a very simple lumped equivalent circuit will do. The graph below shows the simplified lumped equivalent impedance of a PDN. To create a familiar Bode-plot style, we use log-log scale on the axes. Note that for sake of simplicity, the chart shows impedance magnitude only, but as we will see in later articles, the phase is equally important when we use the PDN impedance to calculate the circuit behavior.

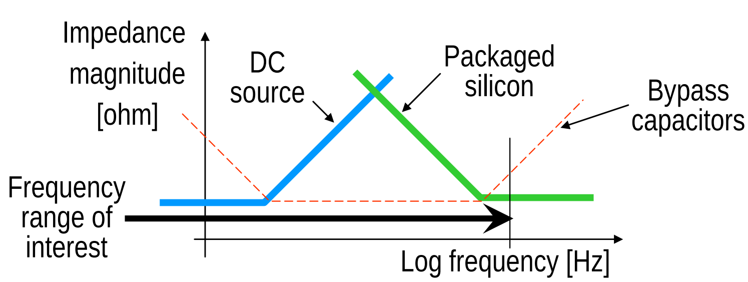

The blue trace on the plot represents the impedance of the power source, which can be a battery or the output of a power converter. At low frequencies its impedance is low, but at higher frequencies the connection inductance of wires or traces or planes will eventually dominate. The green trace illustrates the impedance of the active device. If we take this impedance plot snapshot at the board-package interface, the green trace represents the impedance of the active device together with the package and maybe the PDN capacitors on the package.

We call the active device 'silicon', but it can be any kind of active circuit: on that old computer board the active devices were Germanium transistors. Together the blue and green traces create the triangular-shaped impedance profile with flat tails at low and high frequencies. If we add no bypass capacitor to this system, any noise current hitting repetitively the peak frequency of this triangle would probably create too much noise. The thick black line stretching from left to right above the frequency axis shows what we designate as the frequency range of interest, where the PDN impedance matters. The frequency range of interest is not necessarily a continuous portion of the spectrum, especially for composite circuits with different functions connected to the same supply rail. If the frequency range of interest does not include the frequency range where the impedance peak occurs between the blue and green curves, the circuit will work fine without adding bypass capacitors. If the operation of the circuit excites the impedance peak, we need bypass capacitors to reduce the impedance and thus reduce the noise. The red dashed line on the plot represents the impedance of a set of bypass capacitors that would complement the impedances of the power source and active device, bringing down the overall impedance to low values over a wide frequency band.

So why didn't we need bypass capacitors on the old computer board? Because the Germanium transistors switched so slowly, and the clock frequency was so low, that the transients of switching current did not excite high-impedance portions of the PDN impedance. With today's fast electronics, complex circuit functions and high clock frequencies, most of the time we have to care for a very wide frequency band of PDN impedance and therefore we can not afford having significant impedance peaking between the source and load impedances; therefore we need bypass capacitors. If we fast-forward another few decades, we can easily speculate about having circuit boards again without bypass capacitors. When distributed power sources will be small enough that we can place them very close to the load so that interconnect inductance will be low, at the same time higher amounts of silicon and package capacitance will be available in our active devices, time may come back when our boards will be free again of bypass capacitors.