Decoupling with Integrated Capacitors

Richard Ulrich
University of Arkansas
Fayetteville, Arkansas

Leonard Schaper
Xanodics
Fayetteville, Arkansas

Abstract
Successive generations of ICs demand higher peak current levels and faster current rise times, challenging traditional surface mount decoupling. The considerably lower parasitic inductance of integrated capacitors and the structures associated with them are an enabling technology to prevent decoupling from becoming a system limitation. In addition to better electrical performance, integrated capacitors eliminate the solder joints that must be placed at the hottest part of the board, improving a significant reliability issue, and free up real estate near the chip for other uses.

Introduction
Power distribution is one of the principal functions of electronic packaging, and decoupling is one of the principal aspects of power distribution. Simply stated, the role of power distribution is to supply stable, noise free power, at a constant, specific voltage, to integrated circuits and other components that comprise an electronic system. The trends in electronic systems are for the supply voltages to decrease, required currents to increase, and clock speeds to increase, making it more difficult to distribute noise-free power to all parts of the system. Decoupling capacitors are necessary to achieve stable power distribution, but the use of discrete capacitors in decoupling is becoming less effective due to their parasitic inductance. This opens the door for the use of integrated capacitors, with far less parasitic inductance, for this application.

The ideal power distribution system would look like a battery of constant voltage, regardless of the current draw, connected to the IC by a zero impedance line, as shown in Figure 1. From the viewpoint of the IC, it would see constant voltage no matter how much power it drew or how its current requirements changed with time over the space of a single clock cycle.

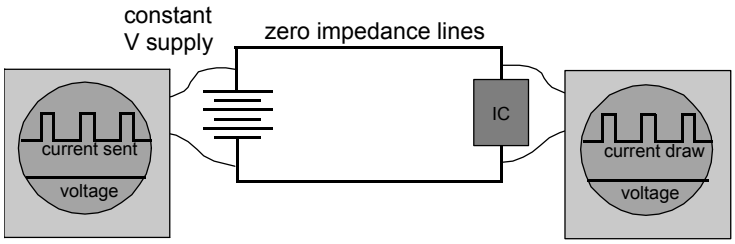

Figure 1 - An Ideal Power Supply and Power Distribution System

In reality, of course, this is not the case. In most systems, one power supply provides power for the entire system, which is distributed from the power supply by a combination of wires, connectors, distribution planes within circuit boards, etc. All of these conductive paths have parasitic inductance, which has no effect at DC, but has a significant effect at the high frequencies typical of IC operation; and parasitic resistance, which affects both DC and AC current. The problem, particularly in the most popular CMOS digital systems, is that as ICs switch many devices between high and low logic levels at each clock cycle, their current demands change rapidly with time. Thus, the power supply is not just supplying a constant current at a particular voltage, it is being asked to supply a highly variable amount of current over a fraction of the clock cycle, which means that the current waveform has many components across a wide range of frequencies, from DC to several GHz. Any impedance present in the real power distribution system will produce a voltage drop as shown in Figure 2. Individual ICs will not see the purely constant voltage they need for proper operation; noise in the voltage supply can cause false logic triggering or insufficient potential to drive signals on or off the chip. The current changes demanded by the IC will not be supplied because of the series inductance.

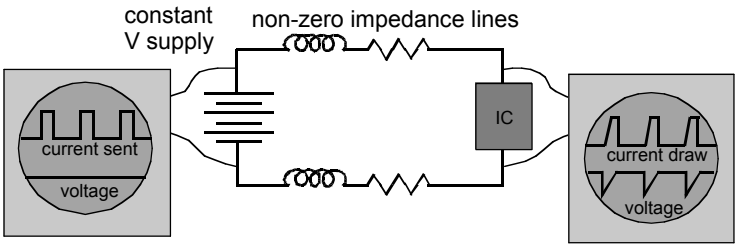

Figure 2 – A Real Power Supply and Power Distribution System without Decoupling Capacitors

The solution to this problem is to put capacitors, called decoupling or bypass capacitors, across the power and ground distribution conductors, physically close to the ICs that are demanding the varying current. These capacitors act as short-term low-impedance reservoirs of charge, and supply current that cannot otherwise be supplied by the power supply because of the low-pass filtering action of the parasitic inductances. They are referred to as "decoupling" because they decouple the power distribution system from the current surges of the IC, or "bypass" because they bypass whatever noise is on the power supply conductors to ground. This is shown in basic form in Figure 3. Viewed as decoupling capacitors, they act as batteries to run the IC for one clock cycle. In between periods of high current demand, the power supply acts as a battery charger to recharge the capacitor. Viewed as bypass capacitors, they are high pass filters that short high-frequency noise generated by the IC and prevent it from getting back into the power distribution system.

Considered still another way, the job of the decoupling capacitor is to make the power distribution system have zero AC impedance when viewed from the IC. Since the parasitic series inductances are relatively small, the IC can "see" all the way to the power supply at low frequencies. That is, low frequency components of varying current can be supplied directly from the power supply and its output capacitors. At high frequencies it is the decoupling capacitor that must provide near-zero impedance. The capacitor should be sized to make its impedance sufficiently low, and ideally zero, over the frequency range of interest to satisfy the voltage stability requirements of the IC. The decoupling capacitor in Figure 3 is shown as a pure component with no parasitics of its own, but it would have a small amount of intrinsic ESR and ESL along with some in the leads and vias between it and the IC.

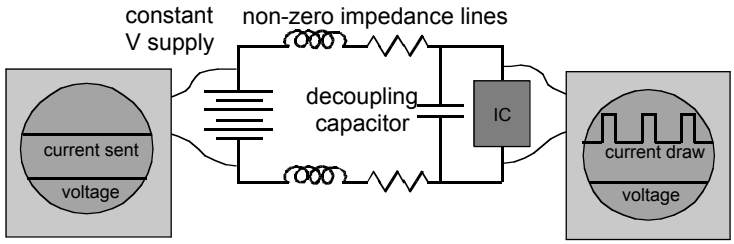

Figure 3 - A Real Power Supply and Power Distribution System with a Decoupling Capacitor

The IC tries to draw a specific amount of charge from the power distribution system in a certain amount of time. Because of the parasitic inductance, the power supply itself is unable to deliver that charge; it all must come from the decoupling capacitor. The capacitor obeys the equation $I\,\Delta t = C\,\Delta V$. That is, pulling a current I out of a capacitor for a time Δt will reduce the voltage on that capacitor by ΔV. The bigger the capacitor, the more charge it can store, and the less voltage drop will be produced by a given current drain. The chain of decoupling capacitors must be able to keep the supply voltage within the tolerance of the IC throughout the period when the IC is drawing current. This establishes the lower limit on capacitor values.[1]

In normal operation the power supply must be able to recharge the capacitor up to the full power supply voltage during one clock cycle, even though the IC is drawing large amounts of current during part of the cycle. If the intervening inductance is too large, this will not be possible. This gives rise to the idea of a hierarchy of decoupling in which each stage closer to the load progressively shapes the current waveforms closer to that required by the chip. This is shown in Figure 4. The power supply simply delivers the average current needed by the whole system. The values of the individual capacitors are a complex function of how the current requirement's frequency components are distributed through the system which is, in turn, a function of the ICs' demands as well as the value and location of the distributed parasitics on the board and in the capacitors.

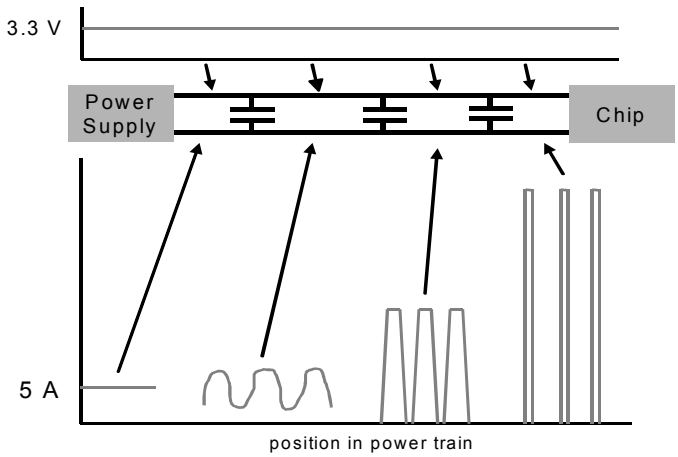

Figure 4 - Hierarchy of Decoupling Capacitors

The picture is further complicated because there are many ICs and many decoupling capacitors in the system, but the principle is the same. The high frequencies and large current demands of the latest microprocessors mean that locating decoupling capacitors immediately adjacent to the IC packages is no longer sufficient; the parasitic inductance of the package itself is enough to block the necessary high frequency currents. Thus we are seeing many decoupling capacitors mounted inside or on the IC packages, and even provided on unused areas of the IC die itself using gate oxide, which is, of course, an integrated capacitor.[2,3]

Also complicating the issue is the fact that not only the clock frequency needs to be decoupled. Because of system power saving requirements, sections of microprocessors and other complex ICs do not operate all the time, but must power up (recharge) instantly, or at least within a couple of clock cycles, on demand to be ready to perform operations. This powering up and down can induce huge delta I demands on the power distribution system at mid-range frequencies well below the clock frequency. All of this charge must be available to the IC in a very short time (little inductance in the loop), which greatly increases the decoupling demand.[4]

Decoupling with Discrete Capacitors

A significant problem with decoupling capacitors, no matter where they are located in a power distribution system, is that capacitors are not ideal devices; they have their own internal effective series resistance (ESR) and inductance (ESL) as was discussed previously. Because of this, no matter how close the capacitors are to the IC, there is still some inductance and resistance that prevents them from decoupling perfectly. For many discrete chip capacitors of the size normally used for decoupling (10 – 100 nF), the ESR can be hundreds of mΩ and the ESL can be several hundred pH. For decoupling, the way to overcome the ESR and ESL limitations is to place multiple capacitors in parallel. All three quantities change favorably; capacitance increases to the sum of the individual components while inductance and resistance decrease. The top curve in Figure 5 shows the impedance vs. frequency for a 100 nF capacitor that also has an ESR of 50 mΩ and an ESL of 300 pH, which are representative values for a ceramic chip capacitor. The bottom curve is for five of these identical capacitors in parallel. The effect is to move the entire curve down while maintaining its overall shape. This increases the frequency range that is below a desired impedance, such as 0.10 Ω.[5–7]

It is not unusual, in a complex, multiboard system, to have hundreds of capacitors used for decoupling, both to provide the necessary charge reservoirs as well as to reduce inductive effects. All of these discretes can occupy a substantial amount of PWB real estate and contribute to cost; the number of solder joints can also be a reliability concern.

As was seen in Figure 5, a parallel combination of capacitors can reduce the impedance compared with one capacitor, but simply paralleling the same value of capacitor does not create the wideband low impedance power distribution system that is necessary. Figure 6 shows the result of a parallel combination of five capacitors of different values, in this case 470 nF, 100 nF, 50 nF, 20 nF, and 10 nF. In real discrete capacitors, corresponding parasitics run approximately in proportion to discrete component value or case size. In this example the ESLs range from 500 to 100 pH, and ESRs from 80 to 20 mΩ. Both the individual capacitor impedances as well as the impedance of the parallel combination are shown. Note that paralleling capacitors of different values rather than all the same value significantly broadens the range of frequencies with low impedance. However, decoupling at frequencies >100 MHz is still difficult because of the intrinsic inductance of standard surface mount devices. This has led to the development of low inductance capacitors (~60 pH), such as the AVX LICA (Low

Inductance Capacitor Array). Multiple devices must still be used in parallel to achieve a power distribution system below 0.1 Ω at frequencies in the GHz range.[8]

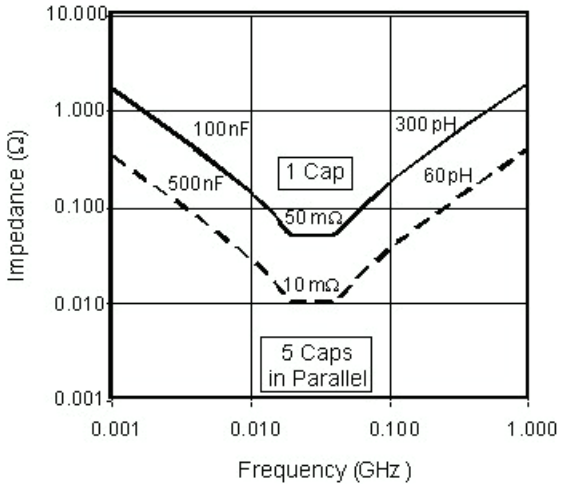

Figure 5 - Comparison of One and Five Identical Decoupling Capacitors in Parallel

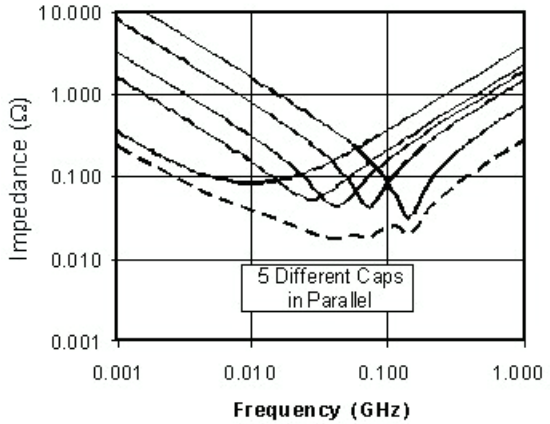

Figure 6 – Effect of Placing Five Different Decoupling Capacitors in Parallel

Lower Inductance in Integrated Capacitors

Why would an embedded capacitor have less inductance? If you want to make an inductor on purpose, you wind a wire into a loop and, the bigger the loop, the more inductance you have. Now think about a surface mount decoupling capacitor mounted on top of the board. For current to travel from the power/ground planes in the board, up to the surface, through the cap, and back down into the board means that it travels in a loop. A simple circular loop 1 mm in diameter would give about 2 nH of inductance, and that's a lot in high-frequency decoupling. But, installing a dielectric directly in between the power and ground planes, as shown in Figure 7, almost completely removes this loop and its associated inductance. The embedded capacitor is planar and in the same plane as what it decouples so that little loop is left, especially if the cap's plates are very close together.

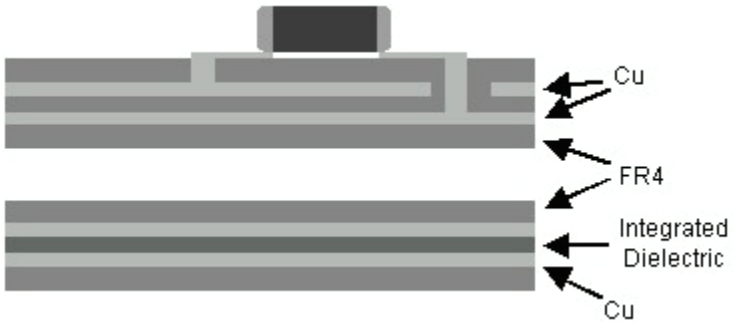

Figure 7 - Vias and Traces, Along with their Inductance, are Absent from Integrated Decoupling Capacitors

On top of that, the inherent inductance of a planar embedded capacitor is less than that of a surface mount because of the way current is directed inside the device. Figure 8 shows how the current in a typical surface mount capacitor travels in the same direction in the plates due to the fact that the leads are on opposite ends of the unit. Currents flowing in the same direction create fields that reinforce each other and amplify the effects of inductance. By contrast, the current through an embedded capacitor is fed in and drawn out of the same edge so the current travels in opposite directions in the plates. Since they travel in opposite directions, their fields mostly cancel out, lowering the inductive effect. Inductance for embedded caps can be even further decreased by connecting the plates along the entire length of their edges, avoiding the current crowding associated with point contacts.

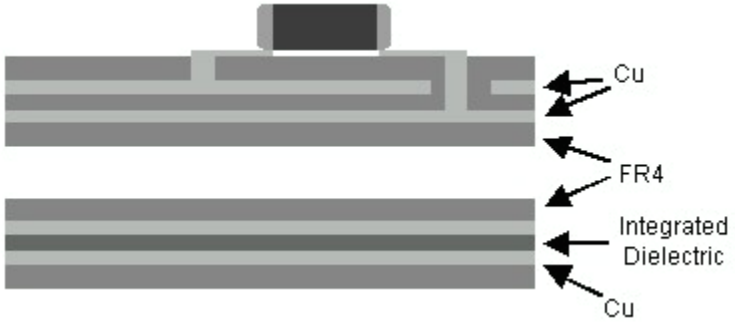

Figure 8 - Inductance is Decreased in Integrated Capacitors Relative to Surface Mount because of Field Cancellation

As mentioned earlier, substantial progress has been made in reducing the inductance of surface mount capacitors. Much of this has come through directing currents inside the device in such a way as to cancel these fields. Arranging the internals to do this requires careful design and manufacture and this is one reason these close-in surface mount decoupling capacitors are many times the cost of the caps used for board-level decoupling, where some inductance is tolerable. These complex surface mount units are actually integrated passive arrays since they contain several capacitors in one unit. Still, the inductance of the loop up to and back from the surface remains, so surface mount decoupling will never be able to achieve the low values that can be had with embedded capacitance. Connecting multiple embedded caps in parallel does not further lower their inductance, so a single large-area integrated capacitor can replace the multiple discrete capacitors mounted in parallel, along with all those pick-and-place operations and solder joints. And, surface space near the chip is freed up for other uses.

Where is Inductance on the Board?

When the parasitic inductance of the capacitors themselves was very high, say 1000's of pH, that inductance dominated the total and it was of paramount importance to decrease it at the cap level. For integrated capacitors and, recently, for some very low inductance surface mount units, the inductance of other board structures dominate, so it is instructive to look at their values in comparison. Table 1 below shows approximate values for some structures common in board-level power trains. The first three items are shown to indicate how low the parasitic inductance of integrated capacitors is by comparison.

Table 1 - Approximate Inductance Values Common in Board-level Power Trains

Structure	Inductance (pH)
1 mm of 10 mil, 1/2 oz. Cu conductor	~700
1 mm dia circle of 10 mil conductor	~3000
bond wire	1000 - 5000
10 mil diameter via, 30 mils long	~300
Solder bumps	10 - 100
Pad/Trace/Via	1000 - 5000
SM Discrete Caps	10's - 100's
Integrated parallel plate	<10, hard to measure

Capacitance Density

If the embedded specific capacitance is low and a large area of the board is used to boost the total available capacitance to the required value, care must be taken in the design to ensure that the entire board-sized capacitor can be seen by the IC it serves. At very high clock rates, the time of flight may be too short for the highest-frequency components to reach very far from the chip. For a given design, let's say that the highest frequency to be considered is six times the clock frequency of 1 GHz. The distance of flight through the board or through an embedded capacitor at 6 GHz is less than two inches. To make matters worse, the actual time that the chip draws current is well under the clock cycle time, so whatever capacitance is intended to decouple these frequencies must be even closer to the IC. It may be necessary to utilize an integrated dielectric with a higher

specific capacitance to put enough decoupling within reach. This is a complicated issue, since propagation rate and capacitance density depend on k, and since the inductance depends on the plate separation, but it's always better to put as much embedded capacitance as close to the chip as possible for both surface mount and embedded solutions.

Technology Availability

Having made the case for embedded decoupling, what is available to use right now? The ideal case would be a single dielectric material that could be placed between the power and ground planes of the entire board, replacing the dozens of surface mount units. An extremely simple approach is to just reduce the thickness of the insulator between the power and ground planes, thereby forming a parallel plate capacitor the size of the board. A two-mil layer of FR-4 between power and ground planes would provide a capacitance of 0.078 nF/cm^2, a very small value. A 1 ft^2 board would have only 0.072 μF from this approach, a fraction of the total required for decoupling everything on a typical computer motherboard. A modification of this involves adding very fine particles of barium titanate or other high dielectric constant materials to the polymer that can increase its specific capacitance by around a factor of ten, to around 5 nF/cm^2 at most. These approaches are commercially available today and provide a limited amount of high frequency decoupling, but not enough to substantially reduce the total amount of capacitance required in discrete devices on the board.

Capacitance densities over 100 nF/cm^2 are on the horizon, and these would provide enough capacitance density to reduce the number of discretes. Ferroelectrics such as barium titanate have very high dielectric constants - thousands - and are soon to be made available as a mil-thick fired coating on one side of Cu foil that can be laminated into the board stack. The first generation of this technology should give around 50 nF/cm^2 with higher values to follow. But keep in mind that the dielectric constant of most ferroelectrics, such as barium titanate, falls rapidly with frequency as you approach GHz levels and this may limit their usefulness. Apart from high k, the other approach to high specific capacitance is to use very thin films of dielectrics. Tantalum, for example, can be sputtered on copper foil and anodized to give less than a micron of Ta_2O_5 and up to 200 nF/cm^2, with excellent high frequency characteristics. Thinner dielectrics always give less inductance, and can provide decoupling at high frequencies in the GHz range where no surface-mount device can. Capacitors we have fabricated, for example, provide 100 nF with less than 4 pH of parasitic inductance and 20 mohms of parasitic resistance.[9]

Device Performance

The actual measured performance of one of these 1-cm^2 devices is shown in Figure 9. The model parameters that gave the best fit to the measured data were 107 nF, 25 mΩ, and 4 pH. Most of the 4 pH of inductance is due to the measurement technique with a single microwave probe; in actual use simulations show that the effective inductance of a multiply connected capacitor is far less. The devices are built on polymer flex, only 50 μm thick, and can be used as discrete surface mounted parts or, potentially, integrated into a circuit board. We are currently doing the work needed to show that capacitors using this technology can be integrated with high yield and reliability.

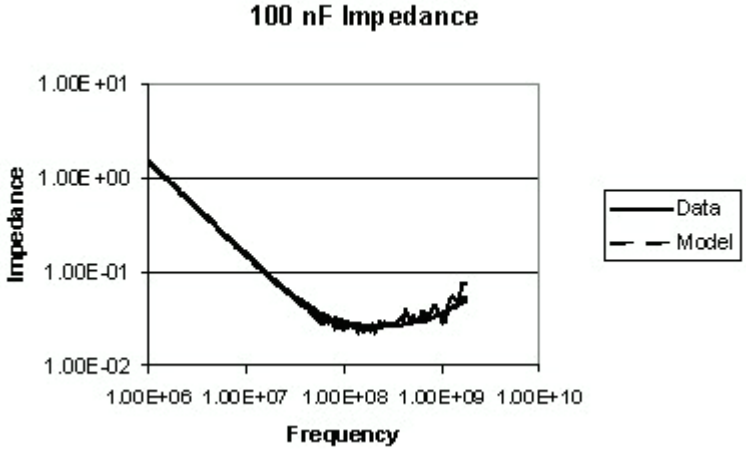

Figure 9 – Impedance of a Thin Film Capacitor on Flex

Conclusions

Decoupling is an ideal initial application for integrated passives. It can solve many of the performance problems associated with the inductance of surface mount strategies, free up board space, and eliminate solder joints. Tolerance issues are not critical as long as a minimum amount of capacitance is delivered. Low specific capacitance materials are available for use now, with higher-performing systems to come.

It's hard to be very specific about how much total capacitance is needed, how the integrated caps should be arranged in the board, and exactly how much inductance is tolerable for each layer of decoupling. That's because these issues are extremely specific to the application and the old rules of thumb for surface mount do not apply.

When decoupling with discretes, much of the capacitance is present only because many capacitors are put in parallel in order to reduce the total inductance. But, due to the inherently low inductance of embedded capacitors, the total capacitance required should be considerably less. The complexity of determining this, particularly with all of the various board, power supply, and IC configurations, probably means that some combination of modeling and experimentation is necessary. The simplest approach would be to just use the same amount that had been used with surface mount but, if that strategy is followed, the buried caps would be very large. Some integrated capacitor technologies could not even provide that much at the scale of the entire board.

So will integrated capacitors be replacing surface mount components for decoupling? As with any new, enabling technology, penetration into the market will be slow, starting with high-performance systems and working gradually down towards commodity products. The point where the majority of decoupling is performed by embedded capacitors is years away, but it will come. The simultaneous issues of supply chain, design tools, and the availability of proven products are being solved and there are ICs around the corner that cannot be decoupled any other way.

References

1. L. Schaper and D. Amey, "Improved Electrical Performance Required for Future MOS Packaging", IEEE Transactions on Components, Hybrids, and Manufacturing Technology, pp. 283-289, September 1983.
2. L. Smith, "Packaging and Power Distribution Design Considerations for a Sun Microsystems Desktop Workstation," Electrical Performance of Electronic Packaging Conference, IEEE, Oct. 1997.
3. L. Smith, "Decoupling Capacitor Calculations for CMOS Circuits," Electrical Performance of Electronic Packaging Conference, IEEE, Nov. 1994.
4. W. Becker et al., "Modeling, Simulation and Measurement of Mid-Frequency Simultaneous Switch Noise in Computer Systems," IEEE Transactions on Components, Packaging and Manufacturing Technology - Part B, Vol. 21, No. 2, May 1998.
5. L. Schaper and G. Morcan, "High Frequency Characteristics of MCM Decoupling Capacitors," Proc. 46th ECTC, pp. 358-364, 1996.
6. A. Murphy and F. Young, "High Frequency Performance of Multilayer Capacitors," IEEE Trans. Microwave Theory Tech., vol. 43, pp. 2007-2015, Sept. 1995.
7. A. Murphy and F. Young, "High Frequency Design and Performance of Tubular Capacitors," IEEE Transactions on Components, Hybrids, and Manufacturing Technology, vol. 16, no. 2, p. 228, March 1993.
8. L. Smith et al, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," IEEE Transactions on Advanced Packaging, Volume: 22 Issue: 3, pp. 284-291, Aug. 1999.
9. D. Nelms, R. Ulrich, L. Schaper, and S. Reeder, "Anodization for forming thin film embedded capacitors in MCM-D and MCM-L substrates," Proceedings of the Electronic Components and Technology Conference: ECTC, pp. 247-251, 1998.

Decoupling with Integrated Capacitors

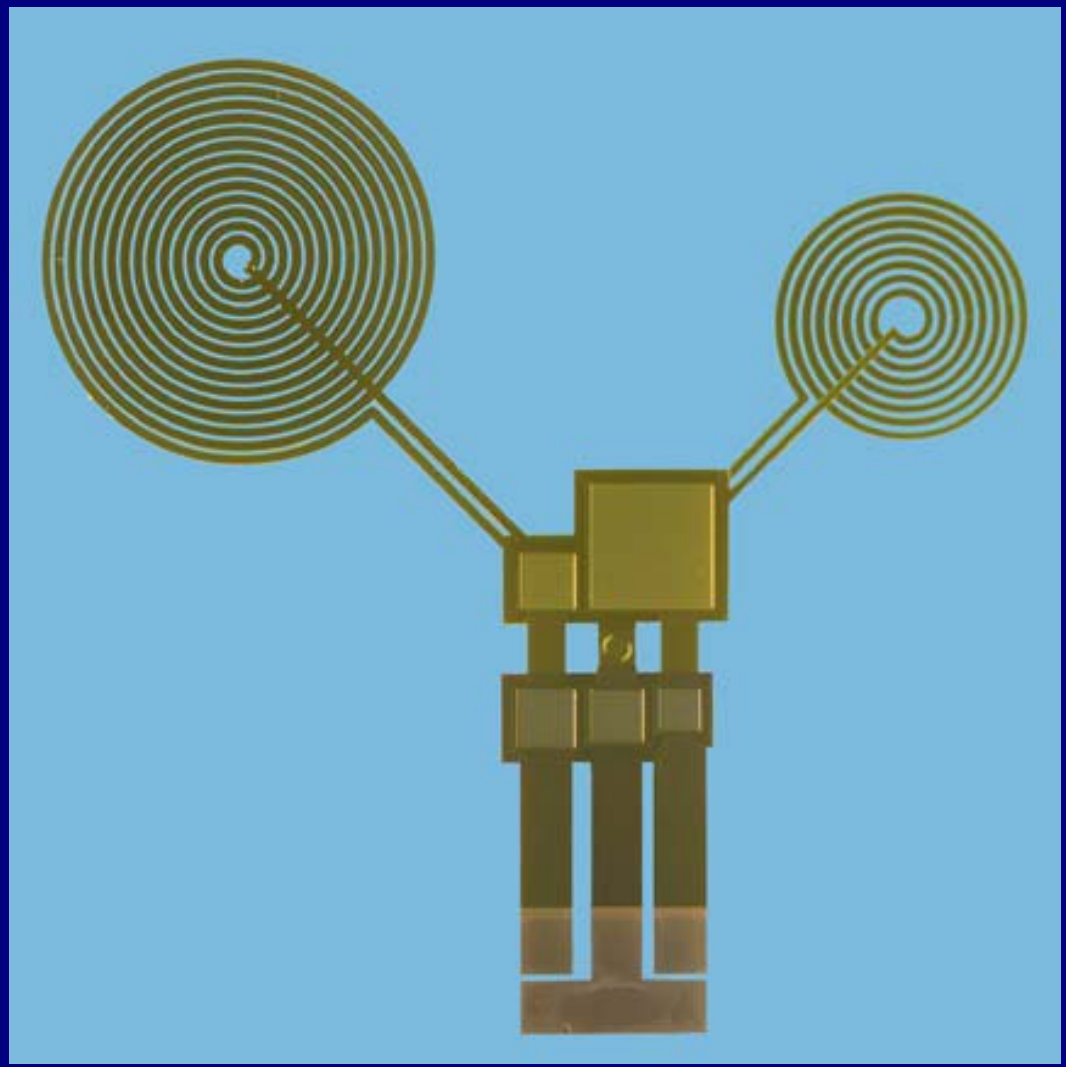

Richard Ulrich
Leonard Schaper

University of Arkansas
Xanodics LLC

IPC Annual Meeting 2003

Minneapolis, MN
Sept. 30, 2003

Overview

What is decoupling?

Using caps to provide low impedance power to chips

This is an excellent application for integrated caps

Their inductance is much lower than discretes

This talk will cover:

1. why and how you decouple with discretes
2. why integrated caps are better for this

Decoupling is an excellent entry-level application for integrated passives

Voltage and Current Draw for a 3.3 V, “5 Amp” IC

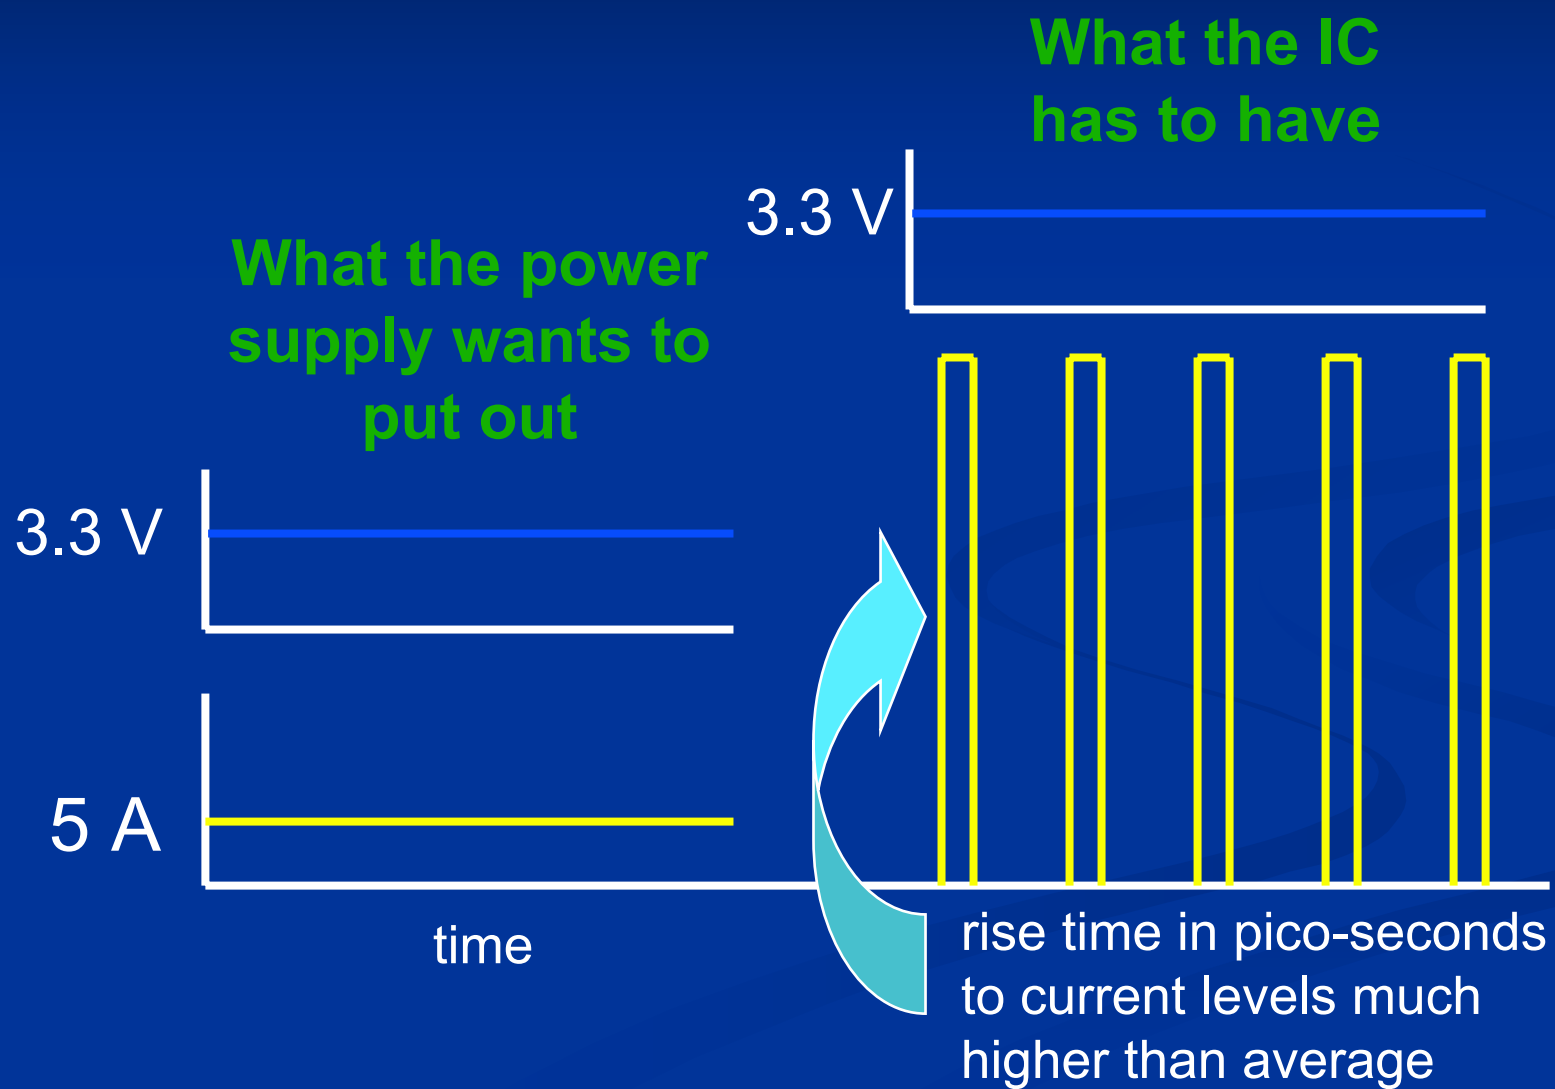

So you have to have very fast current rise times to very high levels, all at constant voltage. What're the enemies?

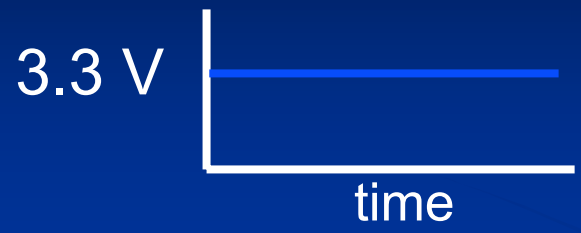

Inductance

$$\Delta V = L \frac{dI}{dt}$$

Resistance

$$\Delta V = I R$$

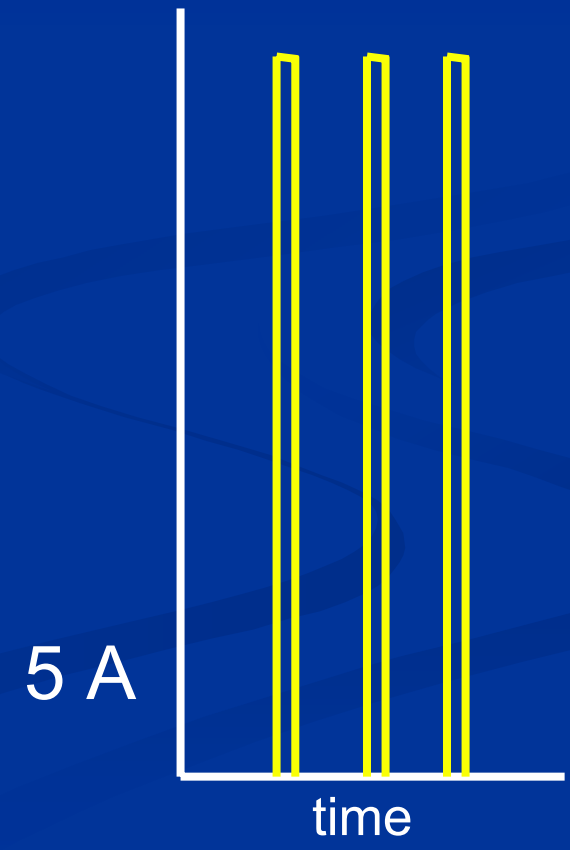

If the power supply could put out the high current spikes and **if** there were no impedance in the power train, then the chip would get the current shaped the way it needs at constant voltage.

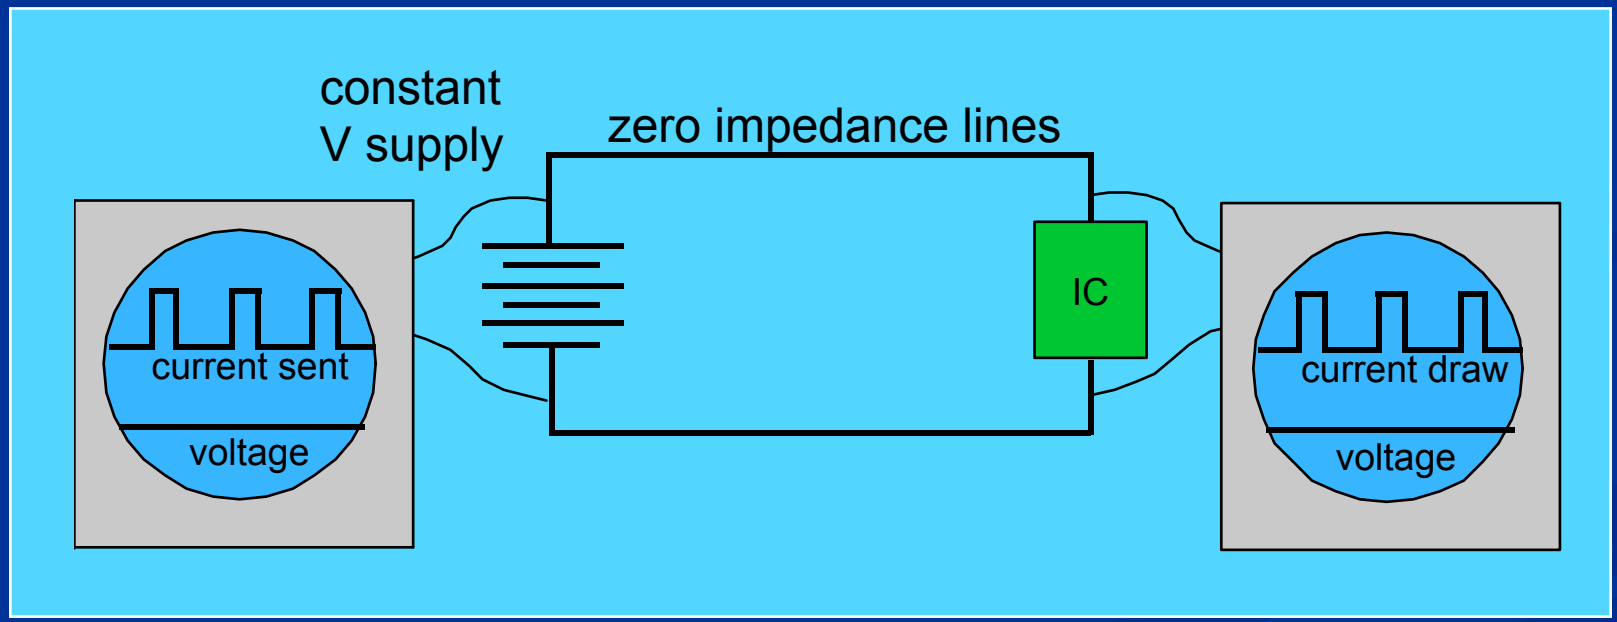

But, the power train has resistance and inductance, leading to distortions of both the desired current and voltage.

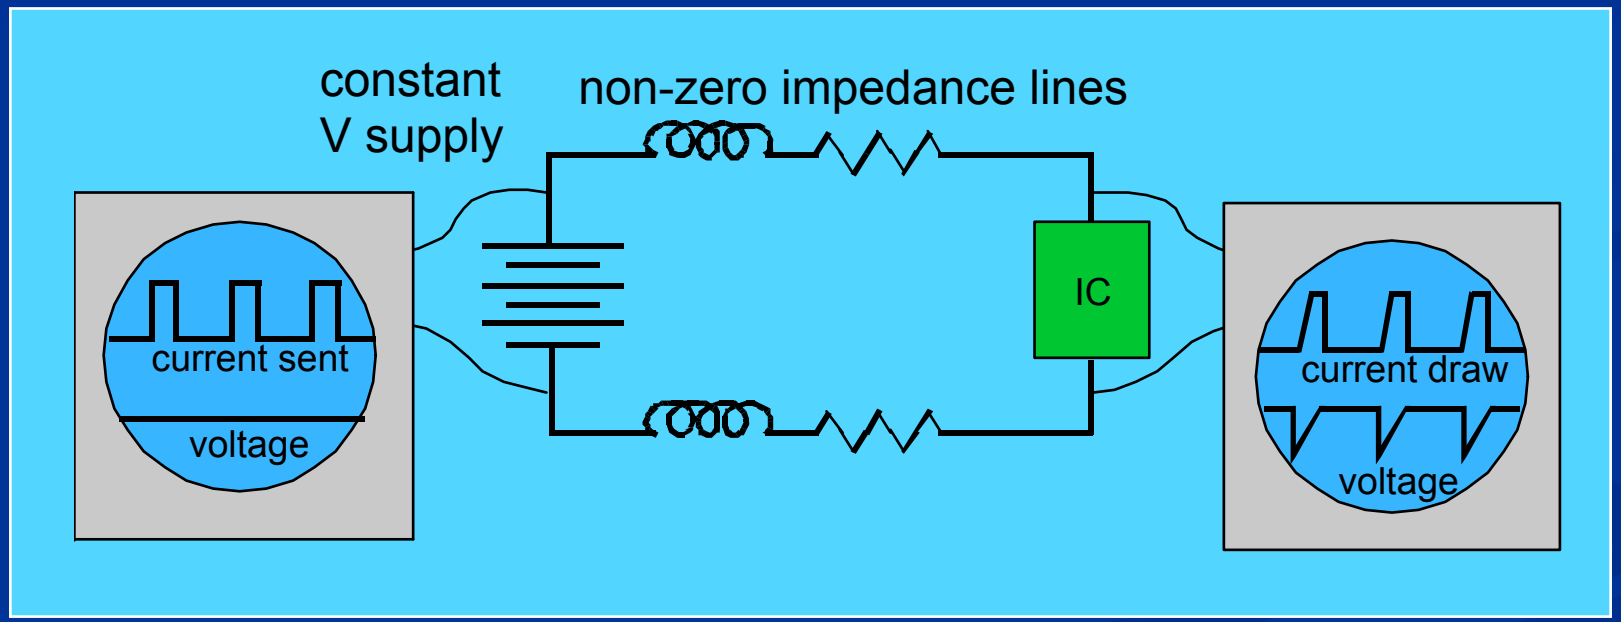

Putting decoupling caps in the power train enables the chip to have the current and voltage forms it needs while, simultaneously, relaxing the requirement for the power supply.

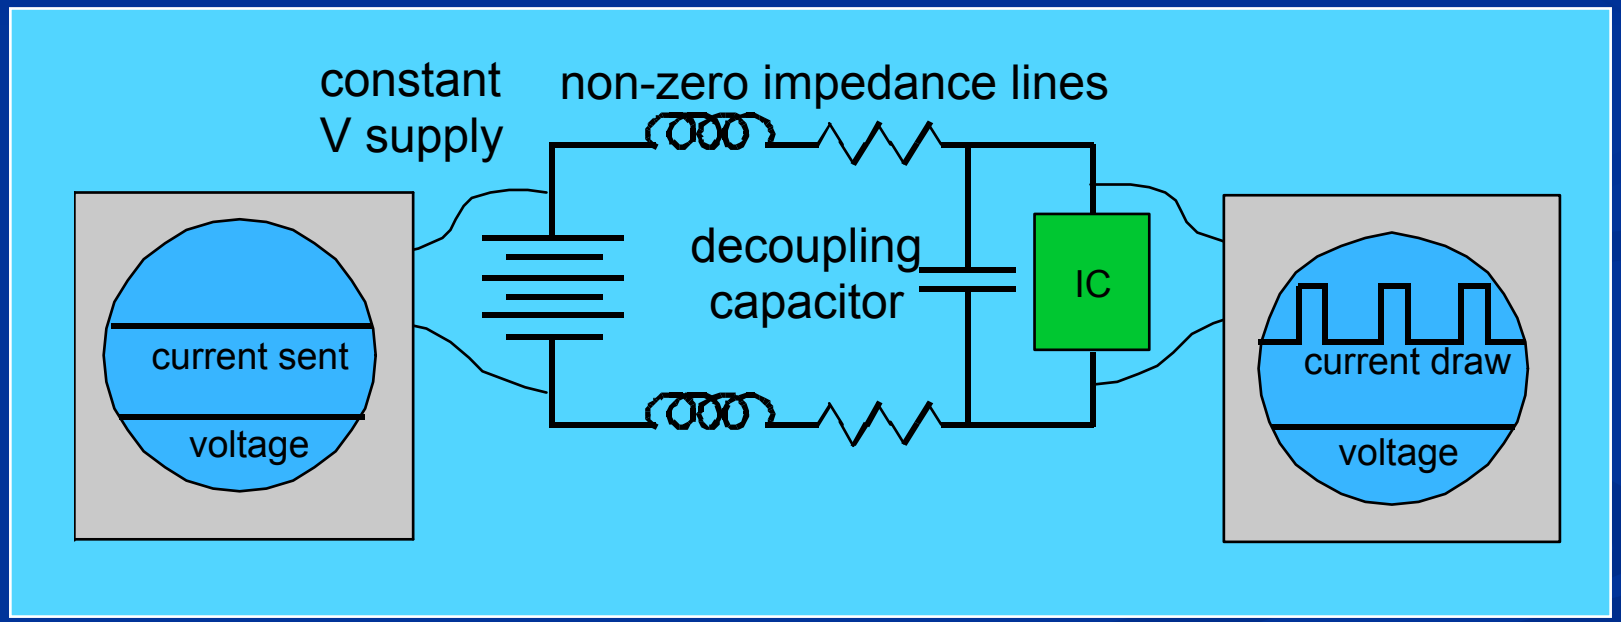

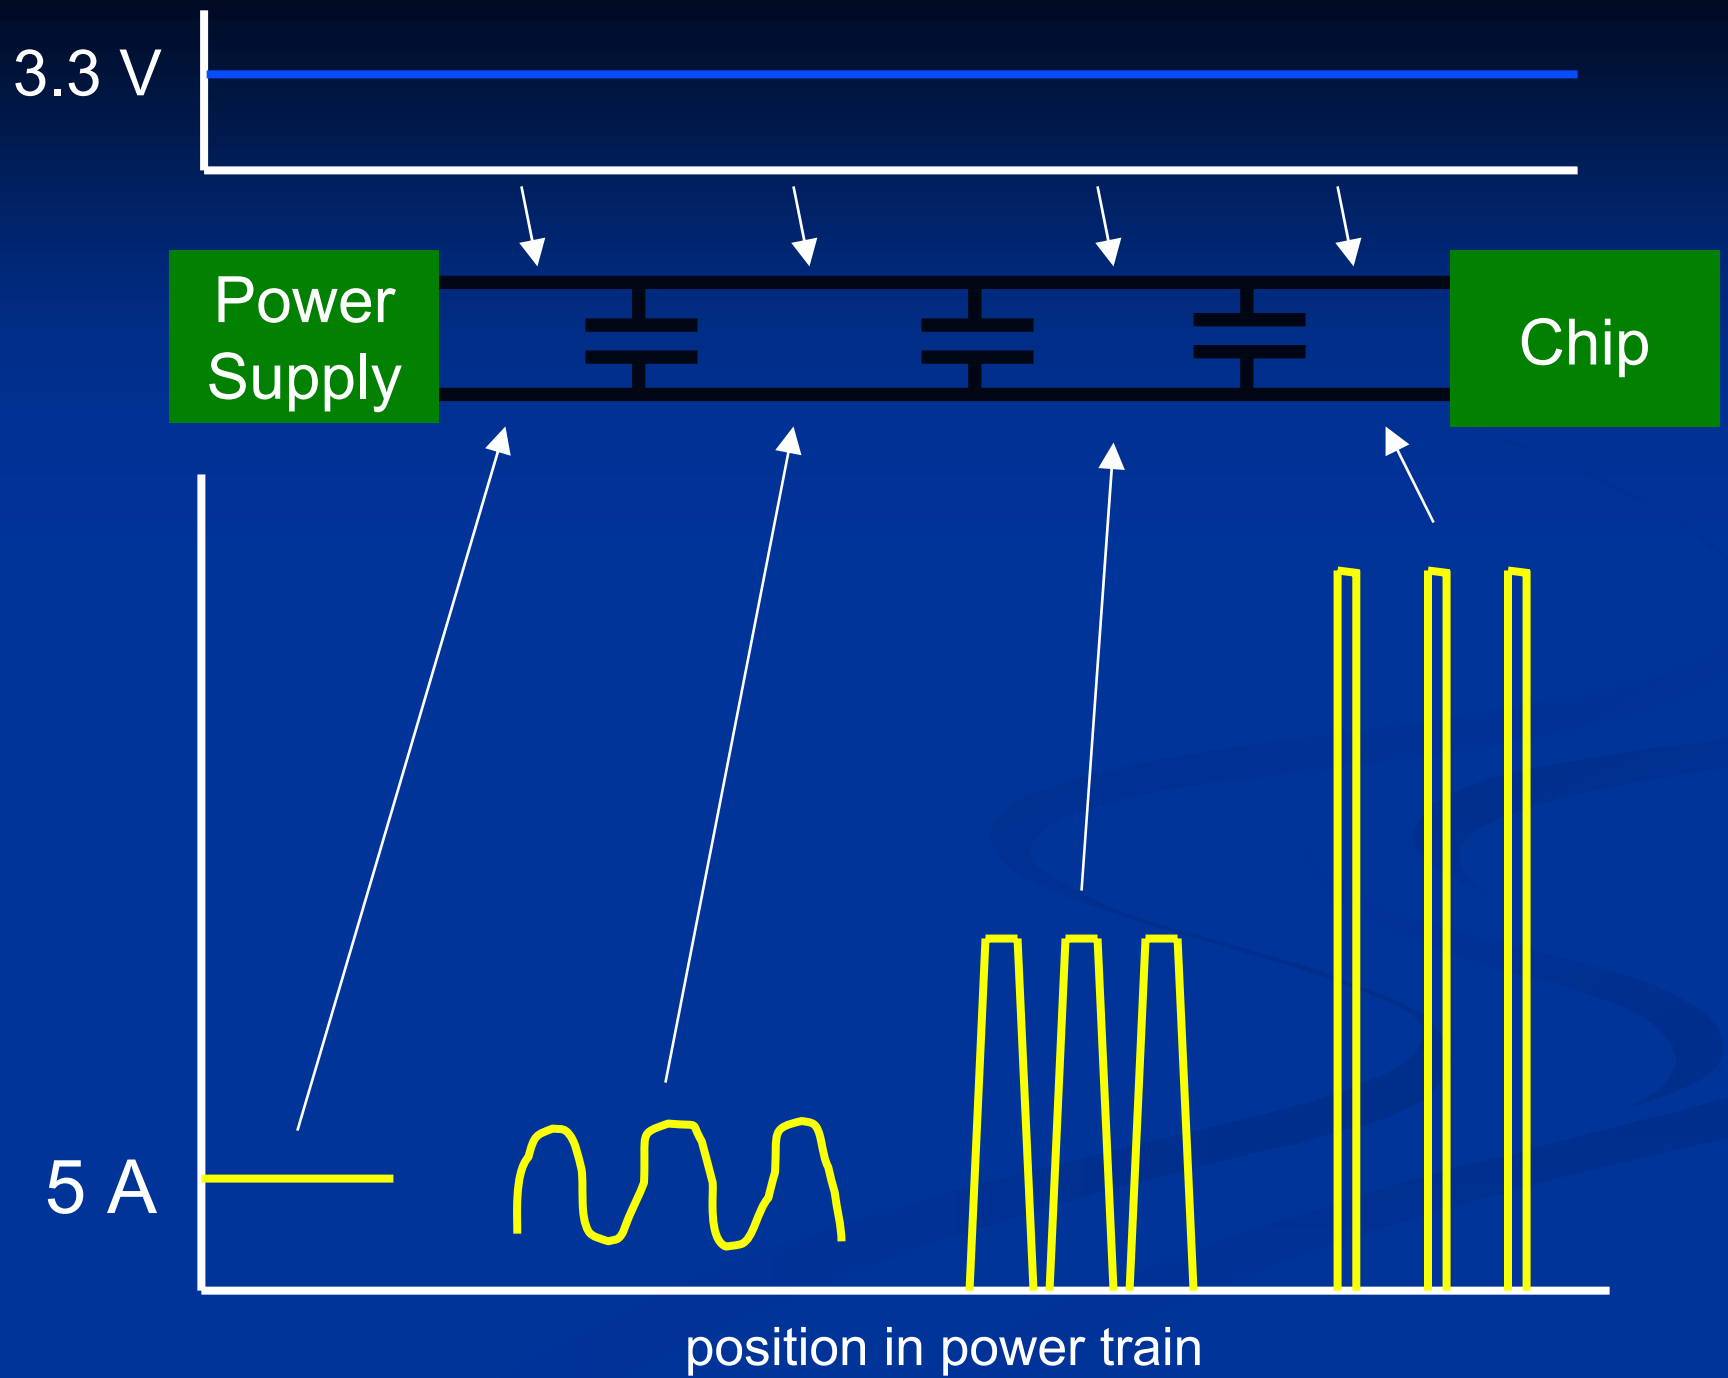
3.3 V
Power Supply
Chip
5 A
position in power train

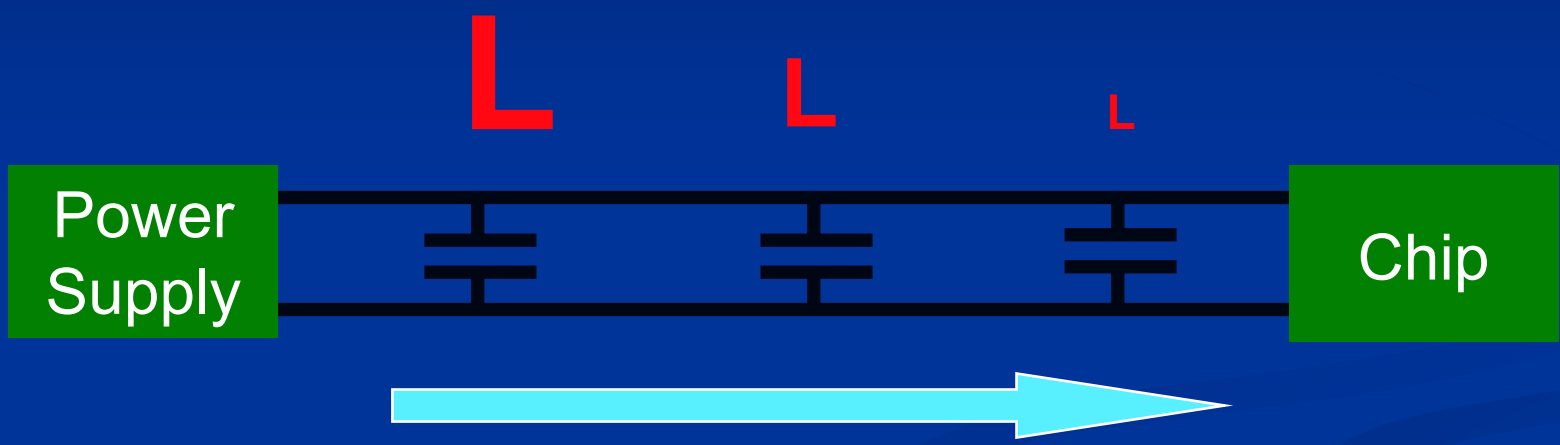

lower parasitic inductance is required in the decoupling caps as you get closer to the chip

Lowering Parasitic Inductance with Surface Mount Discretes

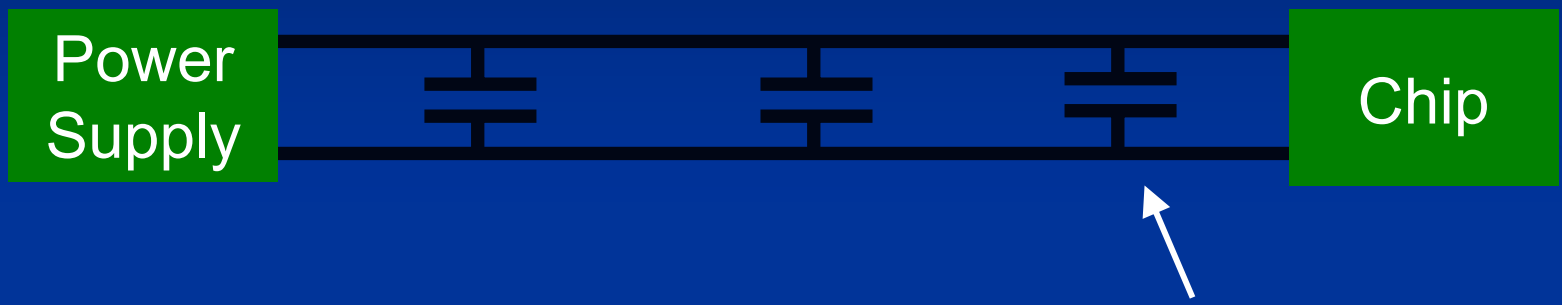

Many small caps are placed in parallel to lower parasitic inductance.

The minimum required capacitance is usually far exceeded

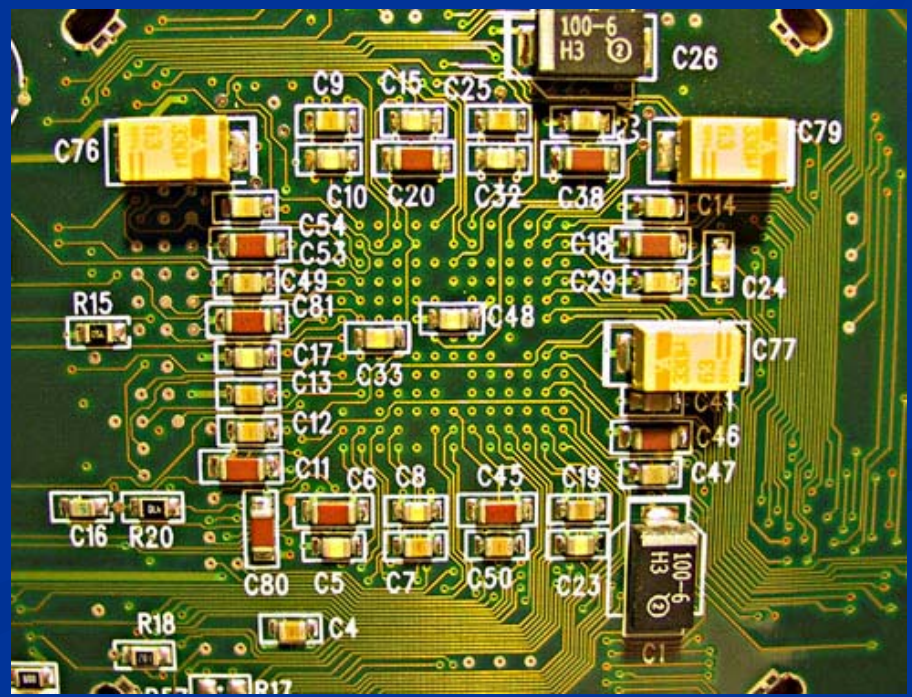

Frequency Behavior of a Single Capacitor

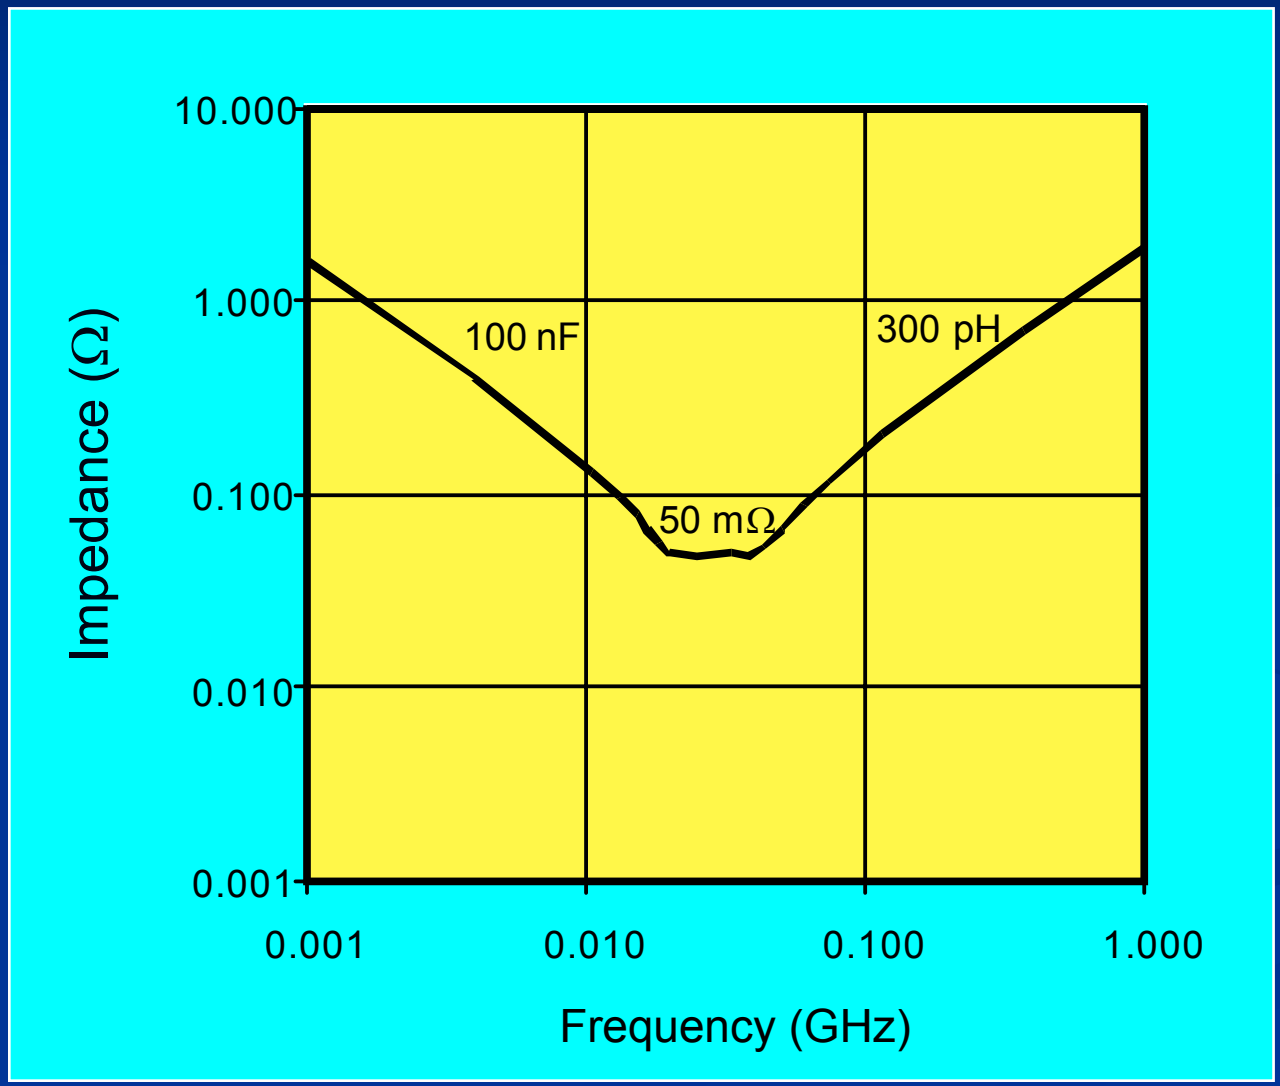

For effective decoupling you want:

1. low impedance at all frequencies
2. high self-resonance frequency

Putting five caps of the same value in parallel moves the curves to lower impedance but does not increase the self-resonance frequency.

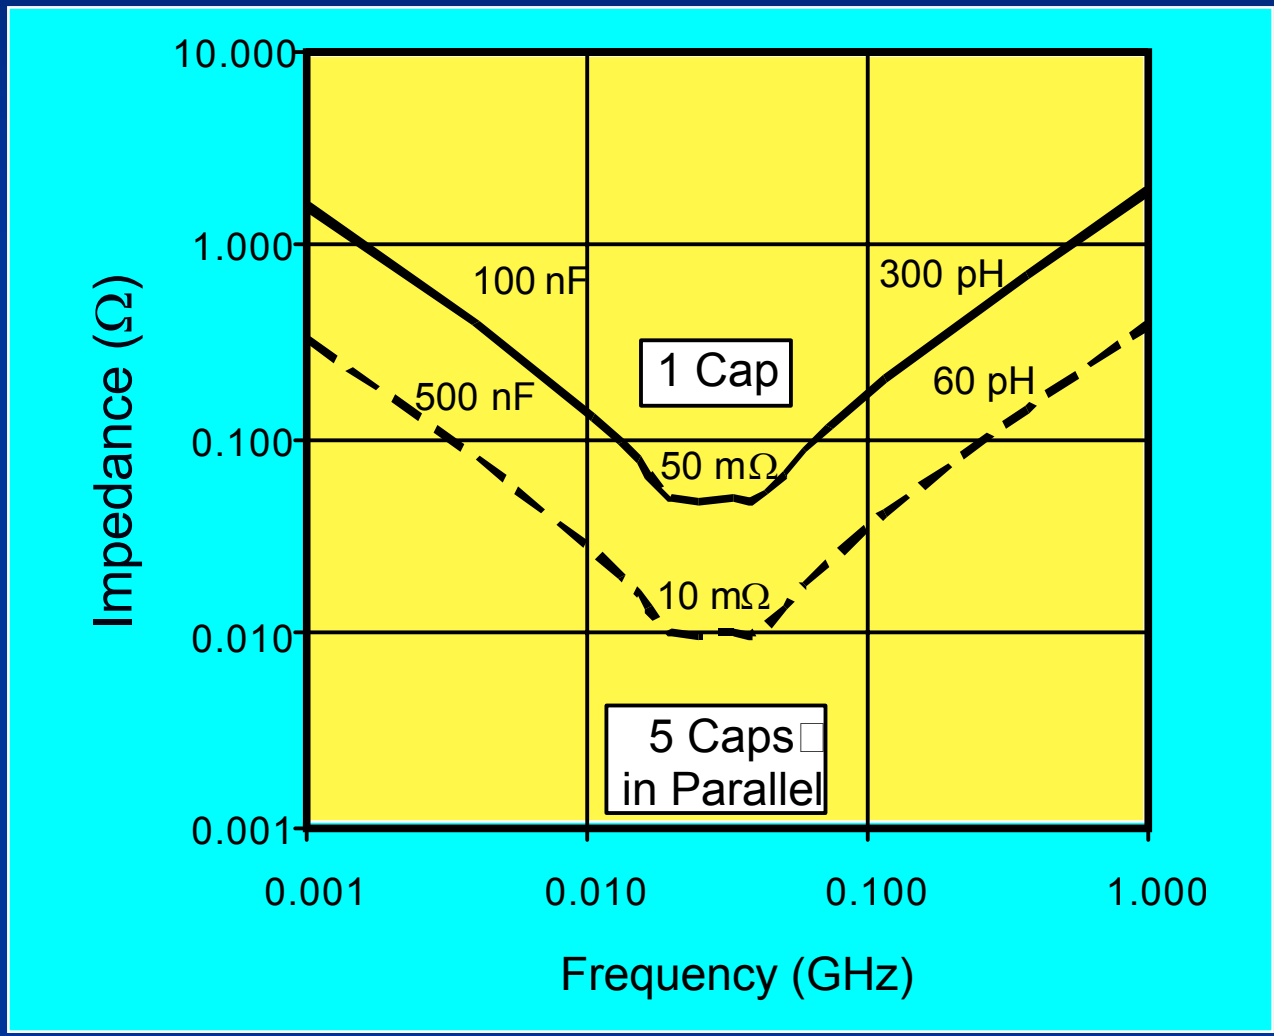

Putting five caps of the different values and case sizes in parallel moves the curves to lower impedance and increases the self-resonance frequency.

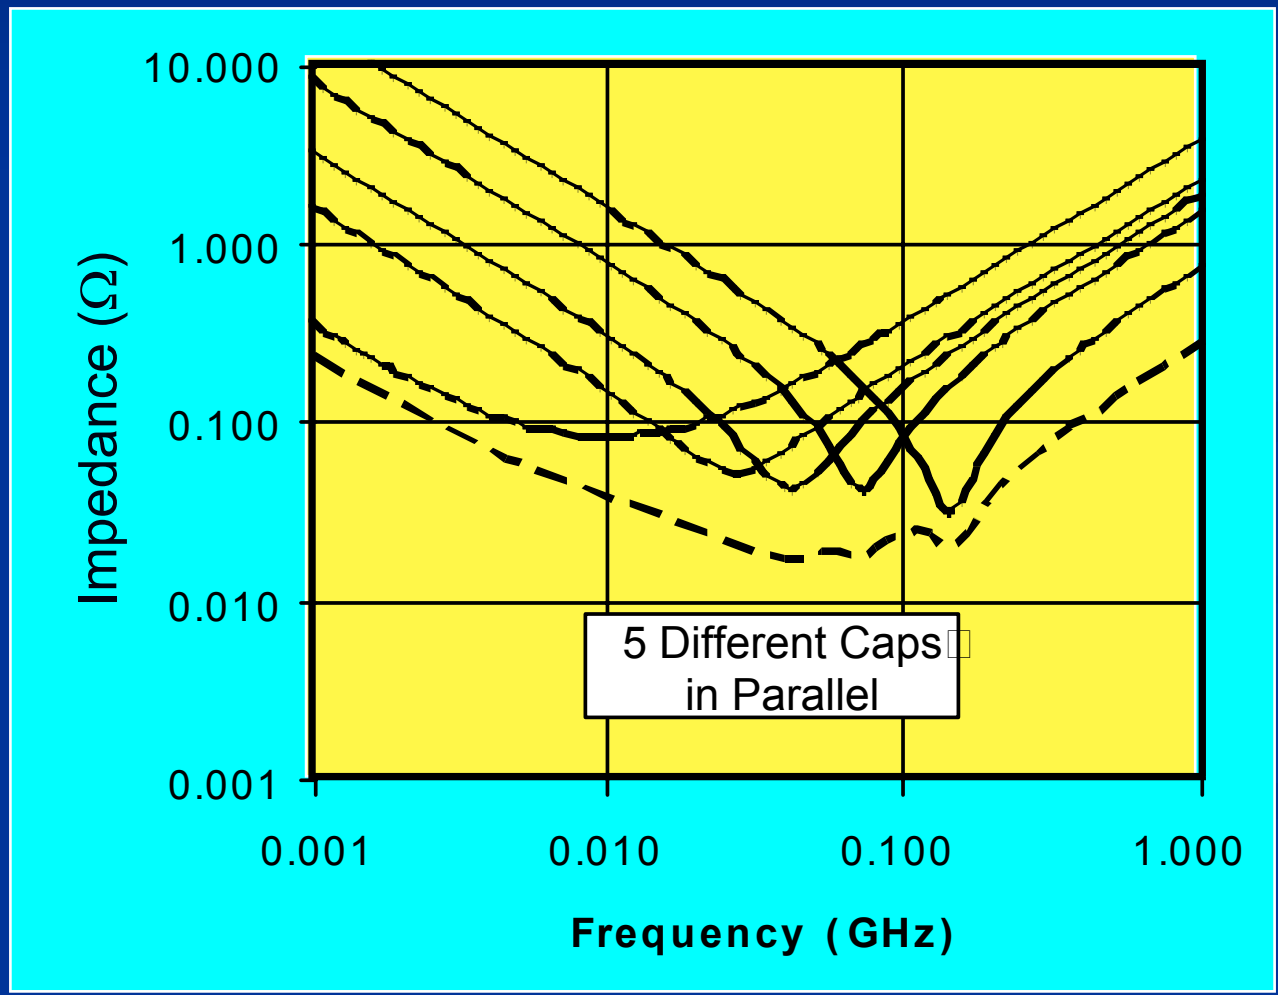

The Situation So Far: Decoupling with Discretes

Multiple surface mount caps are placed in parallel and close to the chip to decrease the overall inductance.

The result is numerous solder joints near the hottest part of the board and valuable real-estate near the chip used up.

Low inductance surface mount caps are being developed, but are expensive and, for coming generations of high-power chips, still too inductive.

Parasitic Inductance in Integrated and Discrete Capacitors

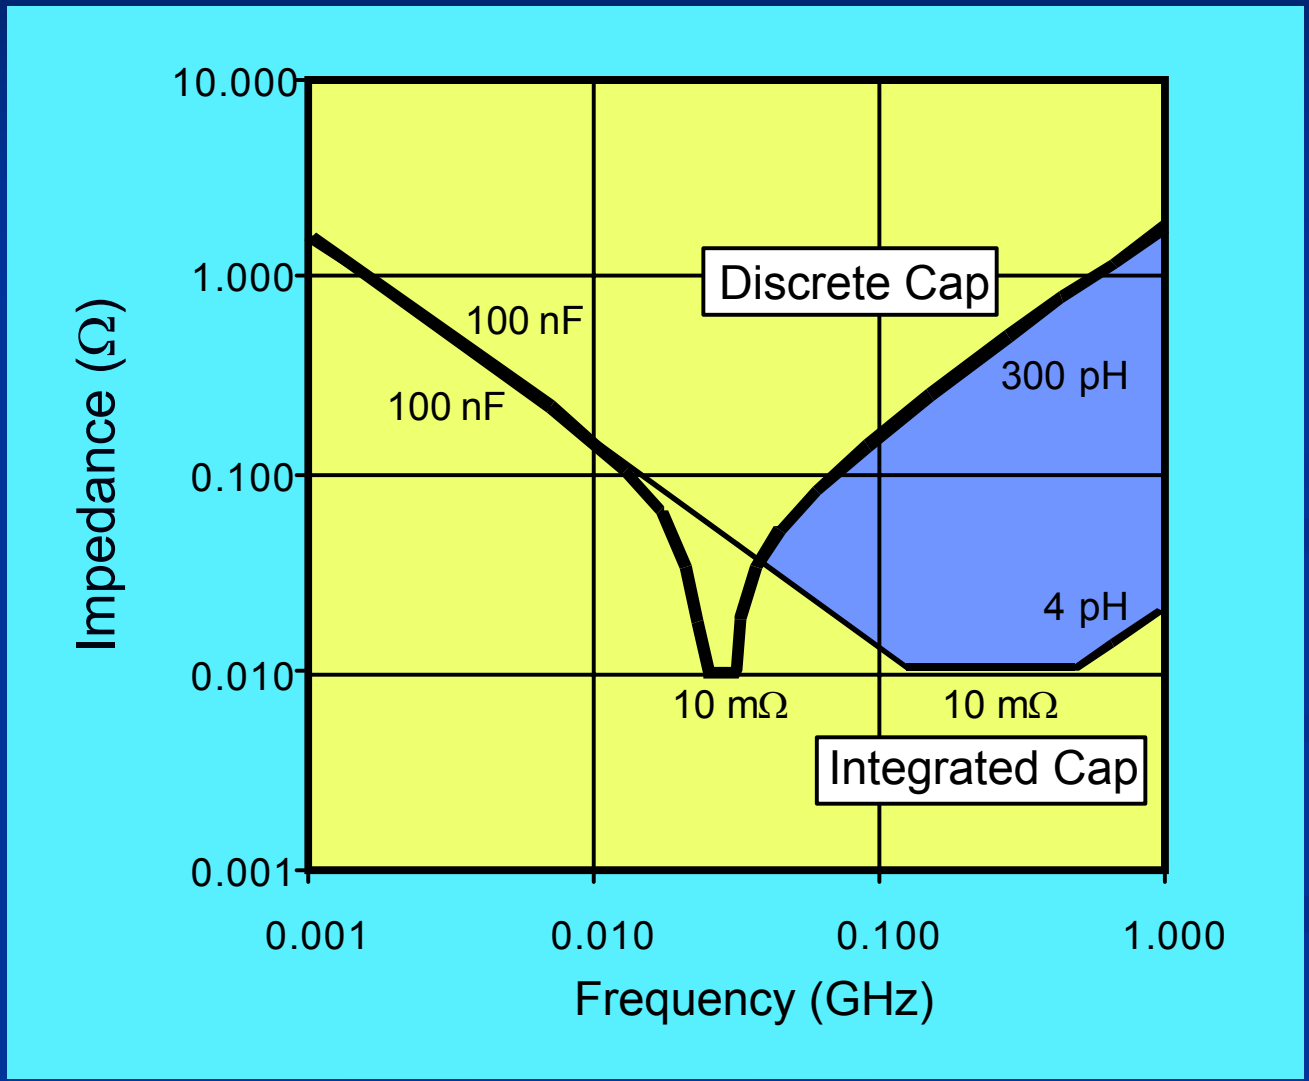

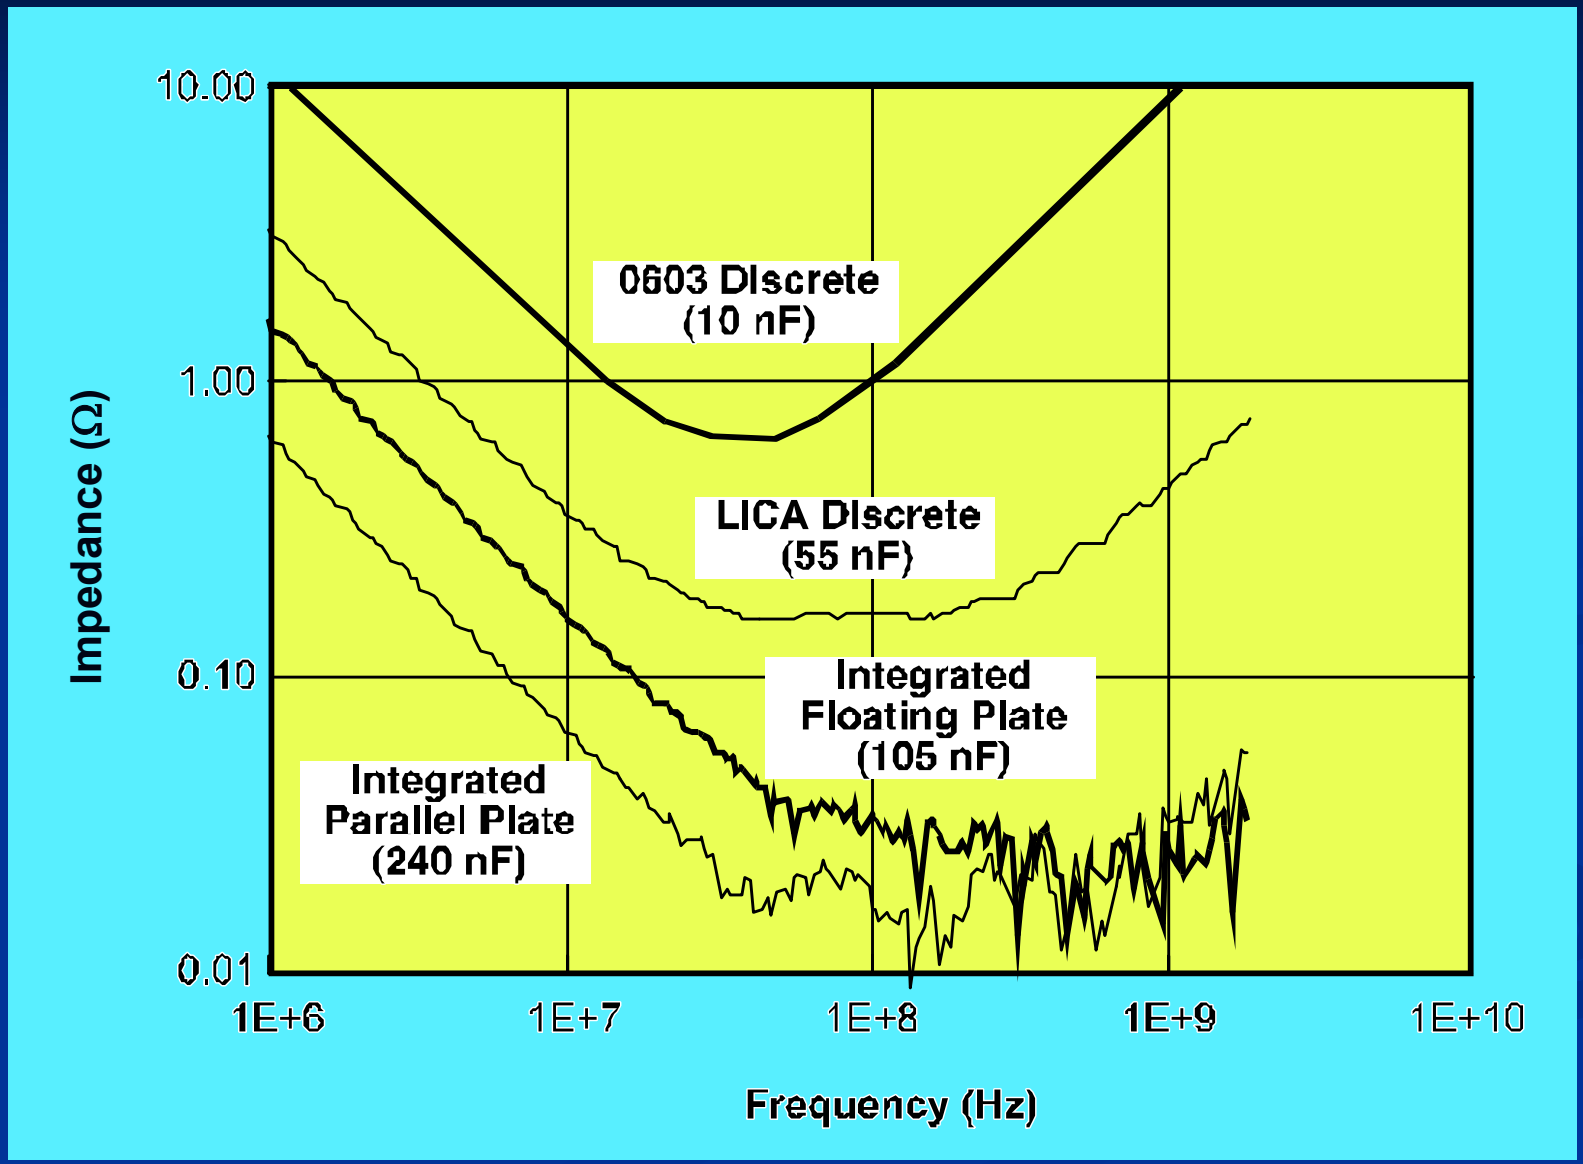
10.00
1.00
0.10
0.01
Impedance (Ω)
0603 Discrete
(10 nF)
LICA Discrete
(55 nF)
Integrated
Floating Plate
(105 nF)
Integrated
Parallel Plate
(240 nF)
1E+6
1E+7
1E+8
1E+9
1E+10
Frequency (Hz)

There are two reasons that integrated capacitors have less parasitic inductance than surface mount:

1. Opposed current directions provide field cancellation

2. Integrated caps can be placed in between the power and ground plane in the board, eliminating inductive vias and traces

Field Cancellation in Integrated Caps

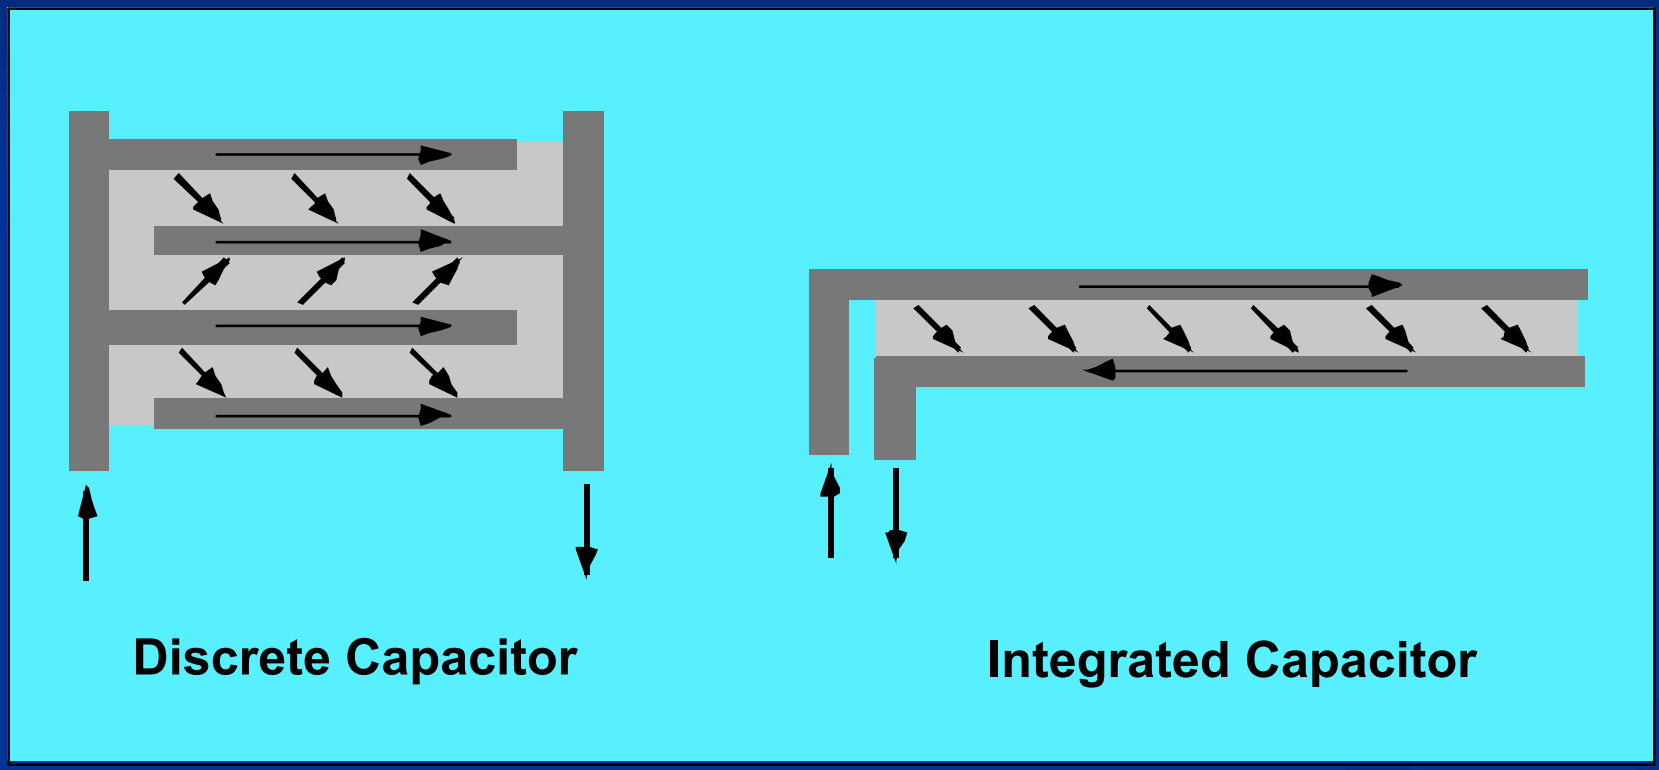

parasitic L of planar, integrated caps w/parallel current:

L in pH/square = 1.26 (dielectric thickness in μm)

Elimination of Inductive Connections

Surface Mount Cap Decoupling

Vias and Traces
are Inductive

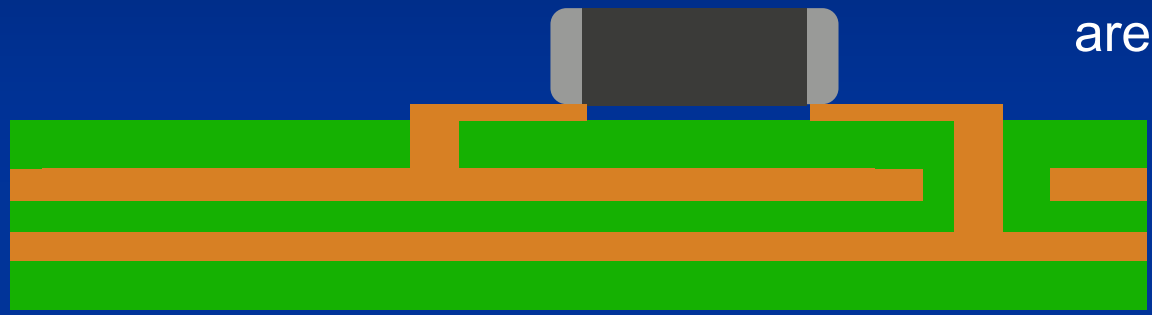

Integrated Cap Decoupling

Integrated
Dielectric

How Much Inductance Is Where?

Just for comparison:

- 1 mm of 10 mil, 1/2 oz. Cu conductor: ~700 pH
- 1 mm dia circle of 10 mil conductor: ~3000 pH
- bond wire: 1000 - 5000 pH

10 mil diameter via, 30 mils long: ~300 pH
Solder bumps: ~10 - 100 pH
Pad/Trace/Via: 1000 - 5000 pH

SM Discrete Caps: 10's - 100's pH
Integrated parallel plate: <10 pH

Below some level of parasitic L, the cap doesn't matter anymore and the limiting factor becomes vias, pads and chip interconnect.

Inductance of Pad/Trace/Via Combinations

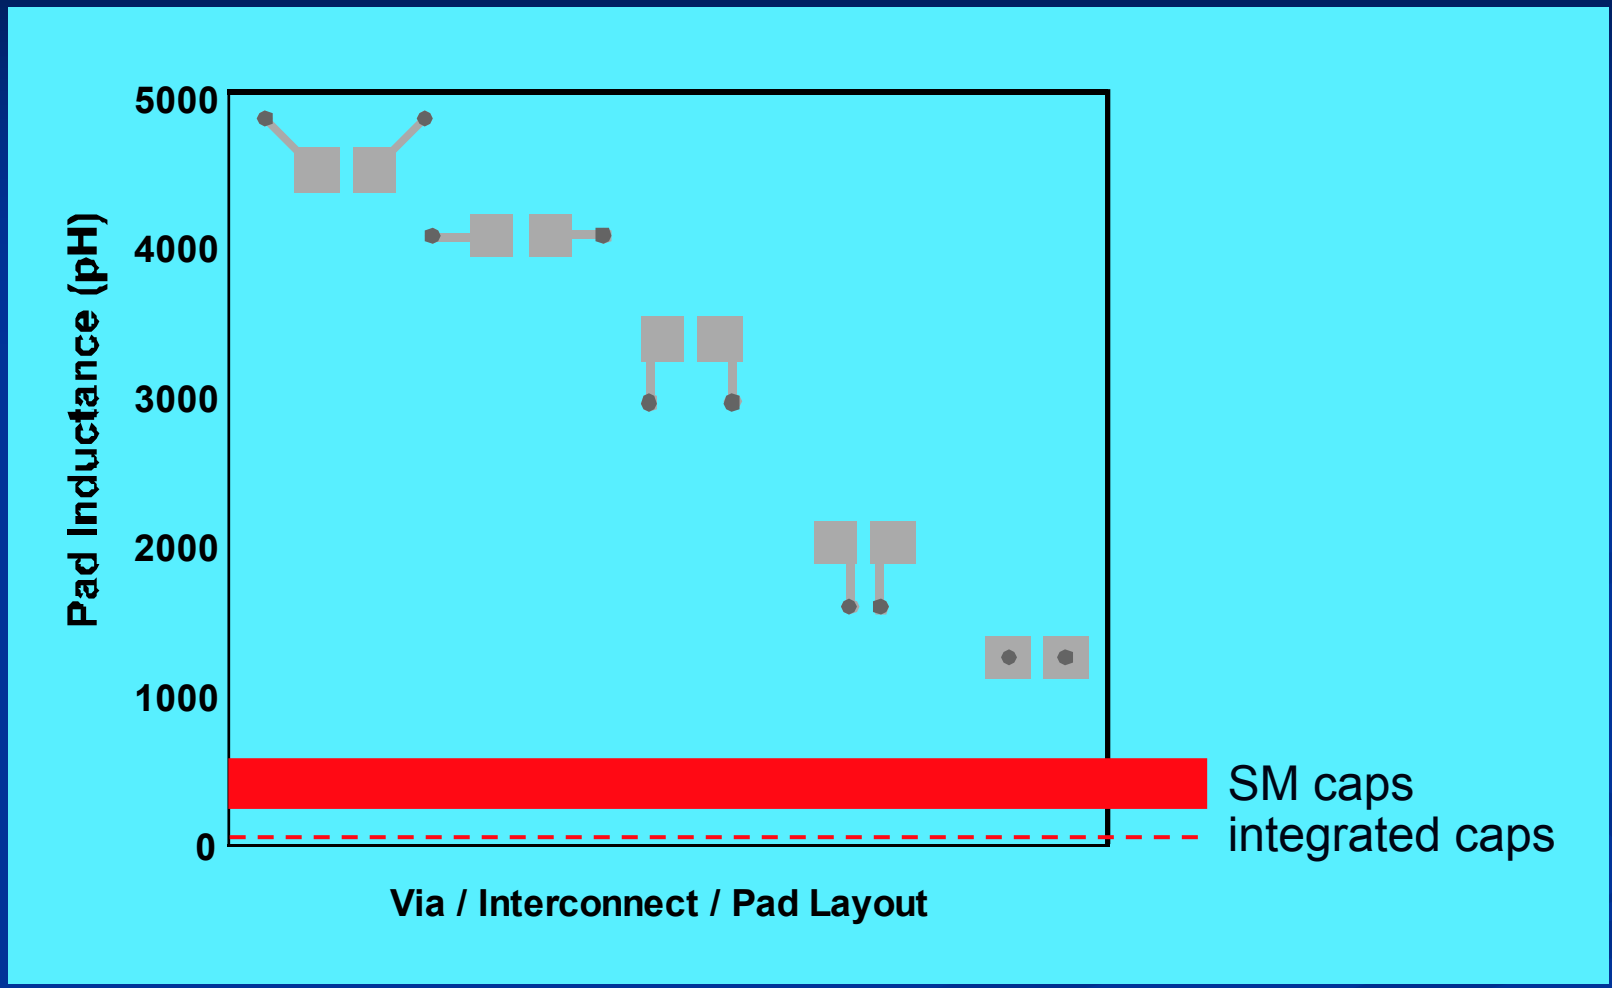

T. Roy, L. Smith "ESR and ESL of Ceramic Capacitor Applied to Decoupling Applications", IEEE Topical Meeting on Electrical Performance of Electrical Packaging, West Point, p. 213, Oct. 1998.

How Much Capacitance Do You Need?

Using surface mount discretes, there is usually much more than necessary because of the large number of caps placed in parallel to reduce parasitic inductance.

Since integrated capacitance is much lower in parasitic L, considerably less C is probably required.

How much less is very application-specific; it's very hard to generalize the answer to this.

How Much Dielectric Constant Do You Need?

Current can only be drawn from a radius that the chip can “see” in one of its clock cycles.

Higher k for an integrated cap would result in slower propagation velocity but more capacitance per unit area. Is it a wash?

Again, inductance matters and probably restricts how fast current can be drawn out.

More k is certainly better, but how much better is also application-specific.

Decoupling with Integrated Capacitors

New ICs demand higher and steeper current pulses
- reducing parasitic inductance is the key
- L of surface mount caps may not be low enough in some cases

Advantages of using integrated caps
- Very low parasitic inductance, does not have to be set in parallel
- Frees valuable real estate near chip
- No solder joints

Requirements - thin films, high capacitance density
- thinnest possible dielectrics - lowest parasitic inductance
- highest possible capacitance density - no good if you can't get to it

Specific requirements are very application-dependent

Integrated Passive Component Technology

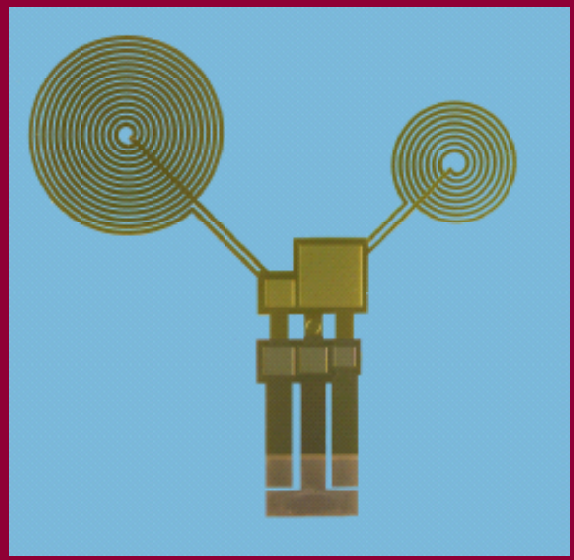

by

Richard Ulrich

Leonard Schaper

Wiley.com
ISBN 0-471-24431-7

Decoupling with Integrated Capacitors

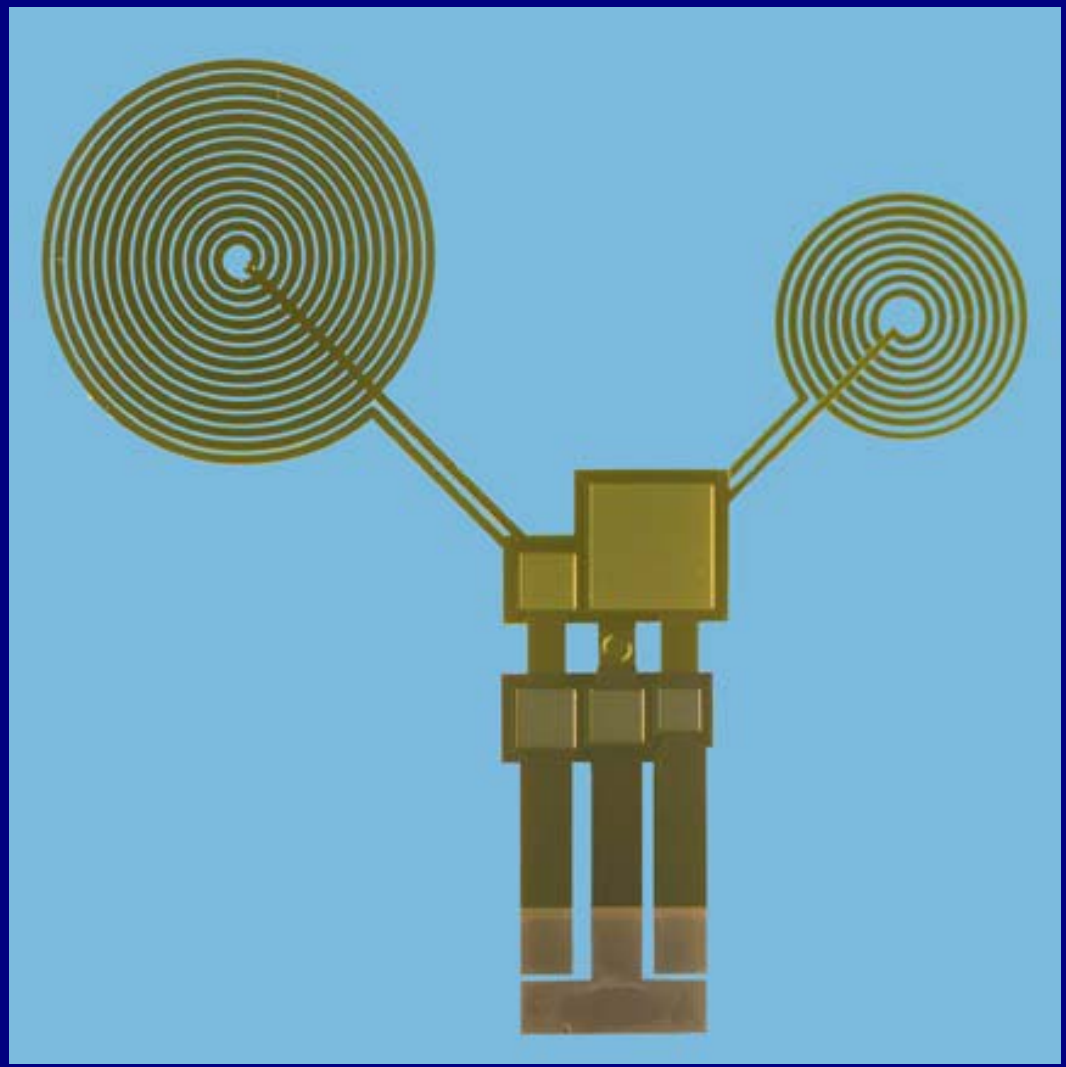

Richard Ulrich
Leonard Schaper

University of Arkansas
Xanodics LLC

IPC Annual Meeting 2003

Minneapolis, MN
Sept. 30, 2003