DesignCon 2005

Decoupling Capacitance Platform for Substrates, Sockets, and Interposers

Josh G. Nickel, Silicon Bandwidth, Inc.
jnickel@siliconbandwidth.com, 408-965-1277

Joseph F. Rosenberger, Silicon Bandwidth, Inc.
joer@siliconbandwidth.com, 408-965-1278

Abstract

As computer system signal speeds continue to rise, the increasing burden on power delivery networks has prompted power integrity to assume a leading focus, along with signal integrity, in system design.

Our solution to power integrity is a novel structural integration of decoupling capacitance between the core power nets and ground to enhance core power delivery. This decoupling capacitance replaces the numerous capacitors suboptimally placed on traditional printed circuit boards (PCBs). Lowered power supply noise and increased core power stability result, permitting greater semiconductor switching frequency while reducing overall system cost.

Studying actual system applications, we compare this technology to a wide range of expensive and largely ineffective decoupling strategies that have been deployed and continue to be proposed, and demonstrate its superiority in both cost and performance.

Author(s) Biography

Josh G. Nickel is a research and development engineer at Silicon Bandwidth, Inc. and is responsible for the development, characterization, and testing of new interconnect technologies for improved signal and power integrity. He received his Ph.D. from the University of Illinois at Urbana-Champaign in 2001. He is an adjunct professor at Santa Clara University where he teaches a signal integrity course. He is the primary author of over 15 technical publications and holds one U.S. patent.

Joseph F. Rosenberger is a Marketing Communications Manager at Silicon Bandwidth, Inc., San Jose, CA. Previously he was with The Panda Project, Boca Raton, FL., and was co-founder and Creative Director of Tarvin Associates, Cincinnati. He has been a Creative Director at numerous Midwest marketing and advertising agencies, and won both Addy and Ad Club of Cincinnati Director's Awards. His previous clients include Procter & Gamble, AT&T Cable, and numerous others. He holds a BS from Xavier University. Co-author and editor of technical publications and articles, he was film critic for Cincinnati Suburban Press. In 2001, he won the Esquire Magazine writers' competition and was a featured presenter at the 2001 Maui Writer's Conference.

Introduction

System data transmission at 10 Gbps and beyond is limited by instability in the core voltages supplying semiconductors. The importance of power management and its impact on the performance of systems is highlighted in the 2003 edition of the International Technology Roadmap for Semiconductors (ITRS) [1] and several microprocessor industry publications [2,3]. The roadmap points out that this problem will become more pronounced with shrinking lithography, increased switching speeds, and greater power requirements.

Traditionally, core voltage “design for stability” has been a compartmentalized endeavor. Inadequate communication among the chip, package, and board designers, as well as a lack of system co-design combine to result in suboptimal performance, including increased bit-error rate (BER) [4], device malfunction, increased distributed noise coupling, and lower semiconductor yield.

Consequently, system designers have resorted to exotic “fixes” and brute-force decoupling schemes at these various system levels. Unfortunately, these “fixes” are inadequate because they are an afterthought rather than a forethought to the system architecture design.

Force-and-Fix Approaches

Maintaining core voltage stability and suppressing switching noise on the power supply are two intimately related problems in today’s systems. From a circuit point of view, these problems can be reduced to simply minimizing the impedance seen by the power delivery system. This goal has come to be focused on the task of effectively providing decoupling capacitance, but it is here that the ineffective “legacy” and “piecemeal” solutions fall short of the task.

The traditional approach to core decoupling has been to endow the PCB, substrate, and die with more decoupling capacitance. This inadequate “brute-force” response has become more pronounced in recent years as system complexity and performance have grown.

Consider typical motherboards, processor cards, or system backplanes with hundreds, even thousands, of discrete decoupling capacitors crowded around the devices at the PCB level. Not only are these discretes unsightly, but on a practical level are extremely wasteful. Every decoupling capacitor requires two pads, two vias, plane connections, some traces, and, worst of all, valuable PCB real estate. The large swaths of area locally required for surface-mount PCB decoupling capacitors prevent other critical components from being located optimally. Inside, signal routing is highly congested and suboptimal. The result: systems suffer unnecessarily from increased bus lengths and remote voltage regulator module (VRM) location, adding latency and loss to data transmission.

Moreover, as simple electrical models demonstrate, decoupling ability and magnitude decreases with distance from the die. Above the resonant frequency introduced by the parasitic inductance from vias, pads, and traces, the capacitor in fact loses its decoupling reactance altogether. This renders ineffective much of the mid-frequency decoupling capacitance on the PCB. These problems become clear in Fig. 1. Power and ground via connections to the topside capacitors suffer high loop inductances and series resistances. Capacitance placed directly under the processor via field, connected between the powers and the grounds, also suffers from via loop inductance.

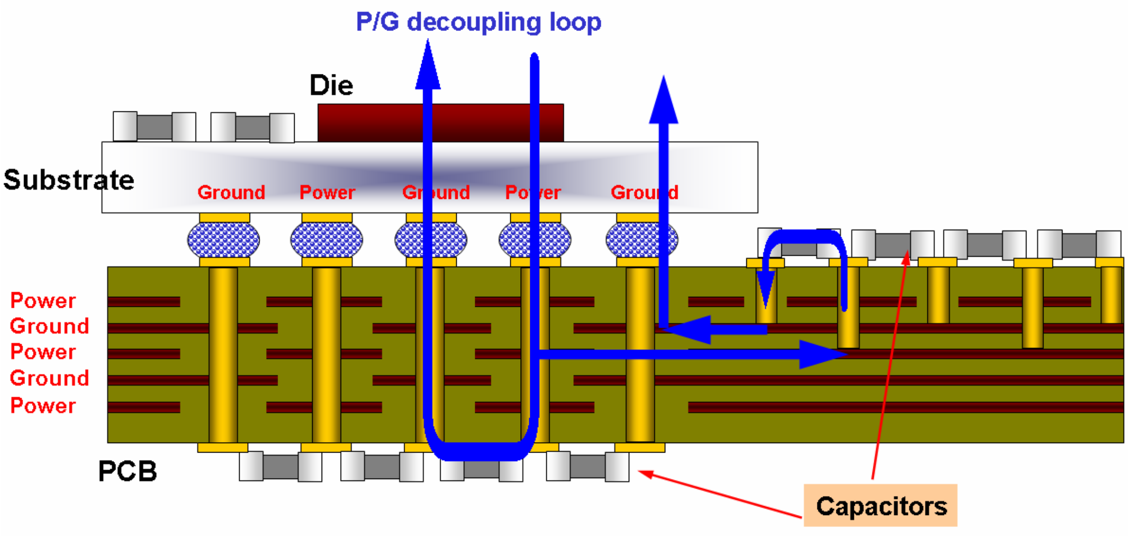

Fig. 1. System decoupling loops for PCB-mounted decoupling capacitors.

To compensate, designers resort to high-density interconnect (HDI) technologies such as microvias, blind vias, and stepped vias, hoping to liberate space, all of which add high cost premiums. Similar solutions aimed at lowering inductance, among them enlarged or double vias, are also prohibitively costly.

Other low-inductance alternatives have been proposed. The more advanced, even exotic, solutions include embedded passives [1,5]. These make use of existing power-ground decoupling planes and utilize the third dimension of the PCB to integrate passives. For most present-day applications, however, extra layers cannot be amortized over the relatively inexpensive cost of surface-mounted discretes. Moreover, densely packed inexpensive discretes still offer orders of magnitude more capacitance in a small form factor.

High-capacitance interdigitated capacitors (IDCs) are designed for low inductance surface-mounting to the substrate or PCB. However, they are expensive components and consume costly substrate real estate. A recent industry trend is to add capacitance directly to the substrate. However, IDCs on the substrate package drive testing costs up owing to the increase in test access points. Moreover, substrates are also singled out by the ITRS as key issues affecting semiconductor operation since increased via density in their cores radically increases cost and lowers transmission, thereby preventing diagnostics and limiting testing.

Thus, any approach to PCB or decoupling that lacks die proximity and increases complexity is de facto suboptimal from the system perspective. Ideally, decoupling should occur as close to the switching semiconductor as possible.

To be sure, on-die capacitance is most ideal and has relieved some of this burden, but at additional cost and at the sacrifice of precious die real estate. New fabrication technologies with increasing complexity continue to provide greater capacitance, but there are also concerns about the increased access resistance of this capacitance.

System Co-design with CapCore

The aforementioned PCB level decoupling “band-aids” have begun or have already reached their limit, as suggested by new power integrity research and new development within the product pre-design cycles. In current designs with clock speeds approaching or surpassing 1 GHz, the noise budget continues to decrease. As signal logic levels fall, the noise budget assumes a larger percentage of the signal, making noise suppression and voltage stability that much more critical. Perhaps the largest cost and performance value of core stability to a system is the ability to clock the processors at a higher frequency.

Unfortunately, the industry suffers from a lack of power grid co-design among the chip, package, and PCB suppliers. This fundamentally design shortcoming is perhaps the root problem plaguing system power integrity. Miscommunication or even a complete lack of communication between the semiconductor design process and the system architects leads to redundancy and inefficiency. Although this problem is well-documented, its solution lags in implementation.

We highlight that this fact was and remains a powerful incentive for a more novel, root-level solution, one that drove the development of this new decoupling platform that integrates discrete or custom capacitance into the voltage core much closer to the die. We will discuss our present prototype development next but emphasize that it is the system-level value that is of paramount interest to the OEMs and end customers.

CapCore Anatomy

In this new and novel configuration, decoupling capacitance is integrated inside the interposer between power and ground pins. The simplest design involves direct solder attach of discrete caps to the power and ground pins.

Core power supply configurations vary from “checkerboarded” to random. The “checkerboard” power-ground grid is typically found in microprocessors where high switching speeds require highly stable power supplies, as it provides the closest possible ground return path and leads to the lowest loop inductance. The checkerboard configuration also readily lends itself to a dense array of packed capacitors.

The first feasible implementation of this solution (trade-named CapCore™) is within an interposer or socket, shown in Fig. 2. It provides substrate solder attach on top and PCB solder attach on the bottom with formed pins. It is designed to be footprint-compatible, adding between 2 and 3 mm in vertical height to the substrate-PCB stack. All pins are mapped 1-to-1 between the substrate and PCB. Figure 3 shows the assembly of this interposer for a general application with the core power-ground grid clustered in the center of the pin array.

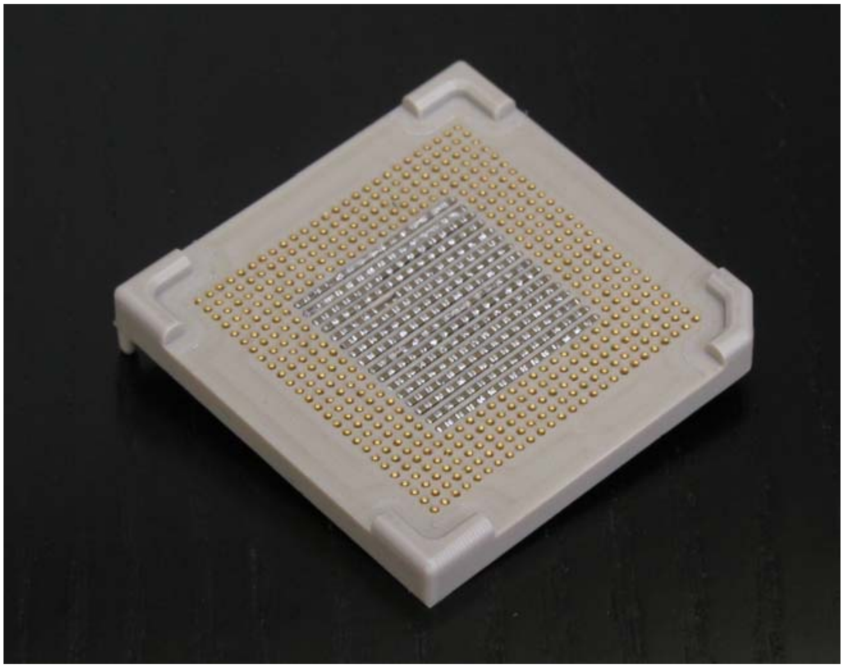

Fig. 2. Photograph of a "checkerboard" configuration of a CapCore prototype on a 1 mm pitch. Core power and ground pins are in the center, surrounded by signal pins.

CapCore at the System Level

Viewed casually, the interposer or socketed CapCore solution is no more than an additional interconnect that adds cost to the system and electrical parasitics to the signals. As an interposer, it adds an additional assembly step and additional vertical height to the PCB-substrate-die stack.

To address the first misconception, we must view the system-wide perspective. Figure 4 shows a reduction in the power-ground decoupling loop to the die. Clearly the decoupling is moved closer to the switching in this "third dimensional" implementation than would otherwise be possible only in the planar axes of the PCB. We will present data in the next section demonstrating the improvement in power delivery and decoupling, but for now it is sufficient to point out that decoupling proximity to the die, as simple electrical models demonstrate, suffers less inductance and is thus far more effective.

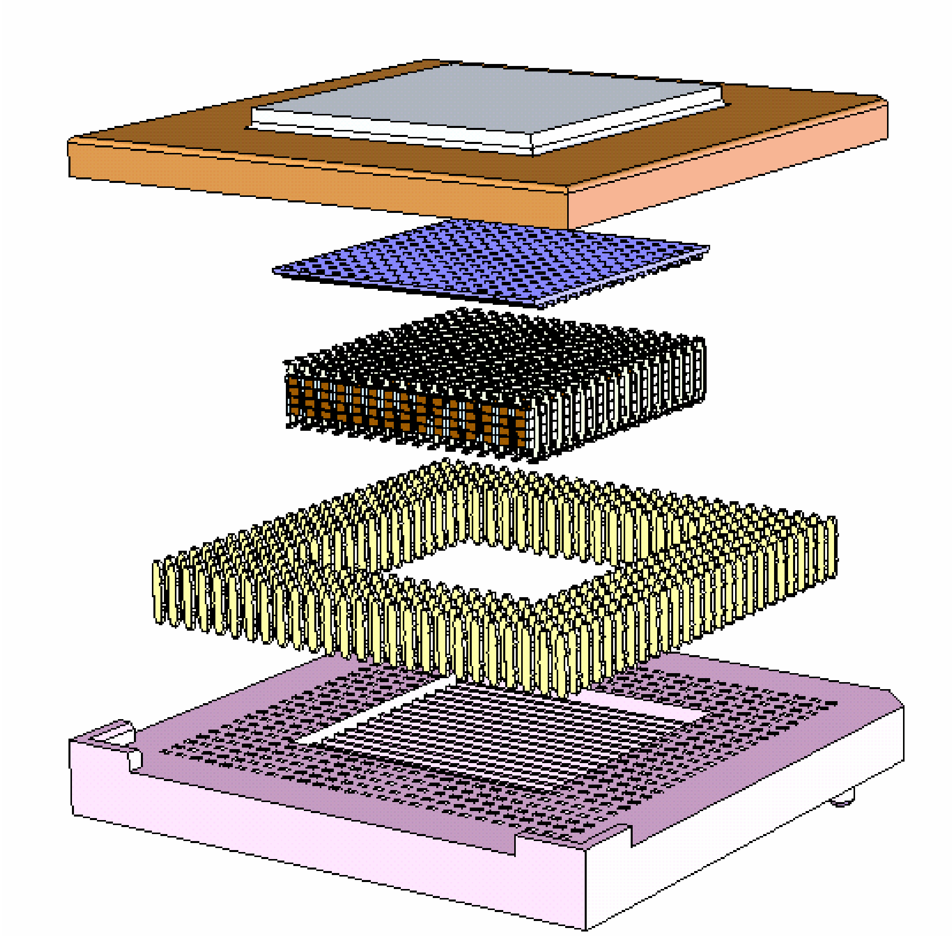

Fig. 3. Exploded view of the interposer assembly. The die/substrate is at top, with a top cover next, then the capacitor core, then signal pins, and finally, the interposer body at bottom.

The relocation of hundreds of PCB decoupling capacitors from the PCB to the CapCore, whether into the interposer, socket, or substrate version, does far more than increase their effectiveness. The newly-liberated PCB routing space, the reduced layer count, the elimination of expensive HDI technology and its exotic via schemes—all work to reduce total applied cost. Two solder joints are removed with each capacitor; assembly and yield costs are reduced. Liberation of PCB space reduces signal congestion and simplifies routing, allowing relaxation of design rules and layout simplification. Our research and customer evaluations indicate that these overall gains reduce system costs.

Now we address the perceived high speed performance degradation through the “signal” pins. Superficially, we should expect degraded performance with the addition of a solder joint, the extra length, and pin parasitics. The additional pin length and the extra ball inductance are immediately sited as problematic.

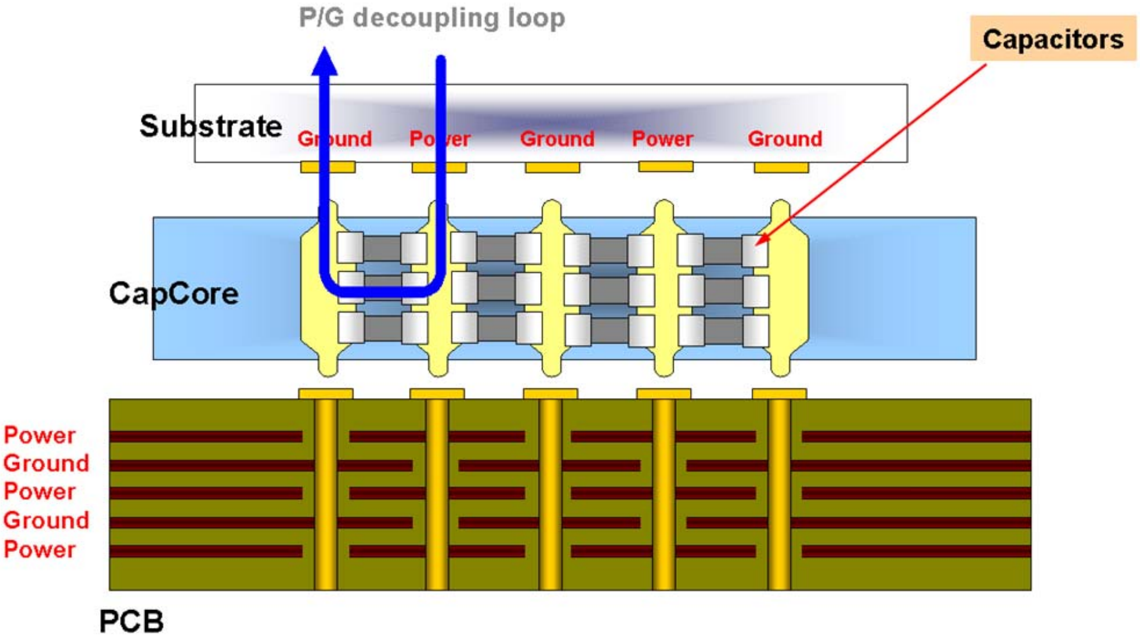

Fig. 4. Decoupling loop with CapCore interposer.

Our solution to this is coaxial signal channels within the interposer. We shall detail application specifics next, but the essential idea is to shield the pins from one another with a grounded main body. Metallization on plastic housings is an established technology and can easily be applied to this interposer body. The pins are then insert-molded or coated with an insulating dielectric, typically plastic. The resulting coaxiality mitigates crosstalk within the pin array by isolating signals, and the controlled impedance is matched to the signal nets.

The additional controlled-impedance length added with the interposer may initially appear troublesome, but we must refocus our perspective and consider the entire system. Elimination of numerous PCB decoupling capacitors eliminates all their supporting pads, vias, and traces. Again, this newly liberated space reduces congestion, simplifies routing, lowers crosstalk, and reduces discontinuities. Studies of present applications indicate that liberated space on the PCB gained from using CapCore readily translates to shortened critical traces, compensating for the approximately 3 mm of additional net length added by the interposer. Moreover, from the perspective of system architecture, shortened bus lengths reduce latency and loss, and allow for closer VRM access.

Thus the main discontinuity introduced with the interposer is the additional solder ball. To estimate its effect, we extracted circuit models for the balls from full-wave simulation Circuit models from full-wave simulation on the interposer reveal that the total loop ball inductance is 77 pH, while the shunt capacitance is 0.04 pF. Two interconnect models, one with the CapCore model and one without, were simulated in Allegro's Allegro PCB SI 630. Linear 50-Ω drivers were applied to the inputs and outputs. A 2.5 Gbps signal was applied to the trace-CapCore-trace interconnect as shown in Fig. 5. The 48-Ω controlled-impedance CapCore was simulated as a 2.7 mm lossy line.

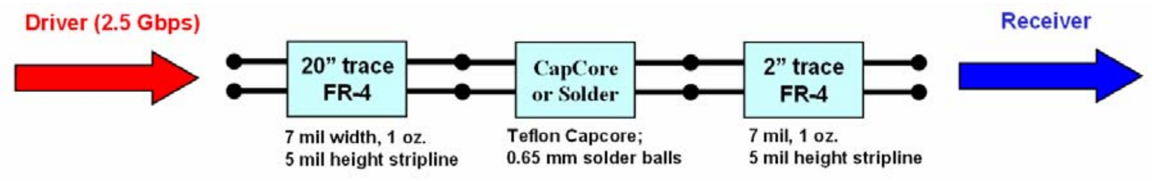

Fig. 5. Interconnect simulation topology for comparison between CapCore and solder balls.

Results are shown for CapCore in Fig. 6, and the single-solder ball (existing solution) in Fig. 7. The CapCore interconnect channel shows an eye opening of 381 mV with 0.55 UI of jitter, while the existing system shows 391 mV eye opening and 0.55 UI of jitter. Clearly this impact on performance due to the extra solder ball is minimal; more importantly, it does not consider all the signal integrity improvements at the PCB level due to liberation of space, reduced congestion, etc. At present we expect few applications to be so sensitive to an extra solder ball; of course an ideal evaluation would involve system-level co-design and analysis with interested customers.

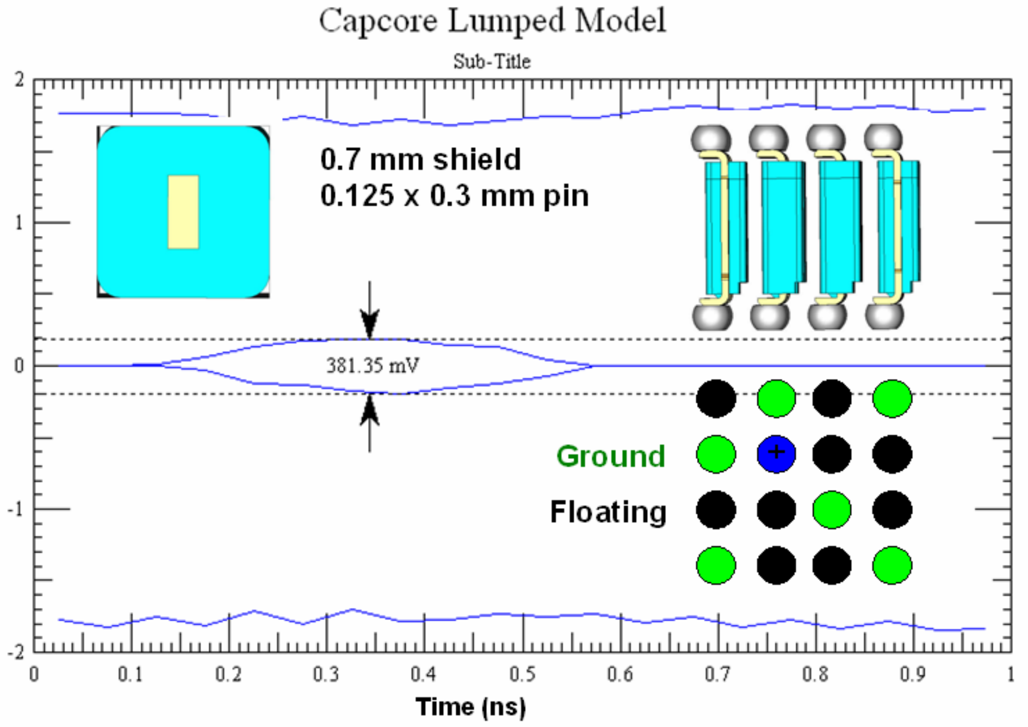

Fig. 6. Eye diagram for 2.5 Gbps trace-CapCore-trace interconnect model. Included is the CapCore controlled-impedance pin cross section and the signaling scheme (darkest pins are ground, while the light pins are floating and the "+" pin is the active line.

Resistance added by the balls and pins is of little concern in most practical boards. Simple DC resistance calculations demonstrate that total signal net resistance is dominated by the traces and vias, with negligible contribution from the ball/pin section of CapCore. Trace widths of 5-10 mil have series resistances on the order of 1 Ω or greater, while the balls or pins, with average cross sections 0.075 mm^2, have resistances on the order of 10 mΩ.

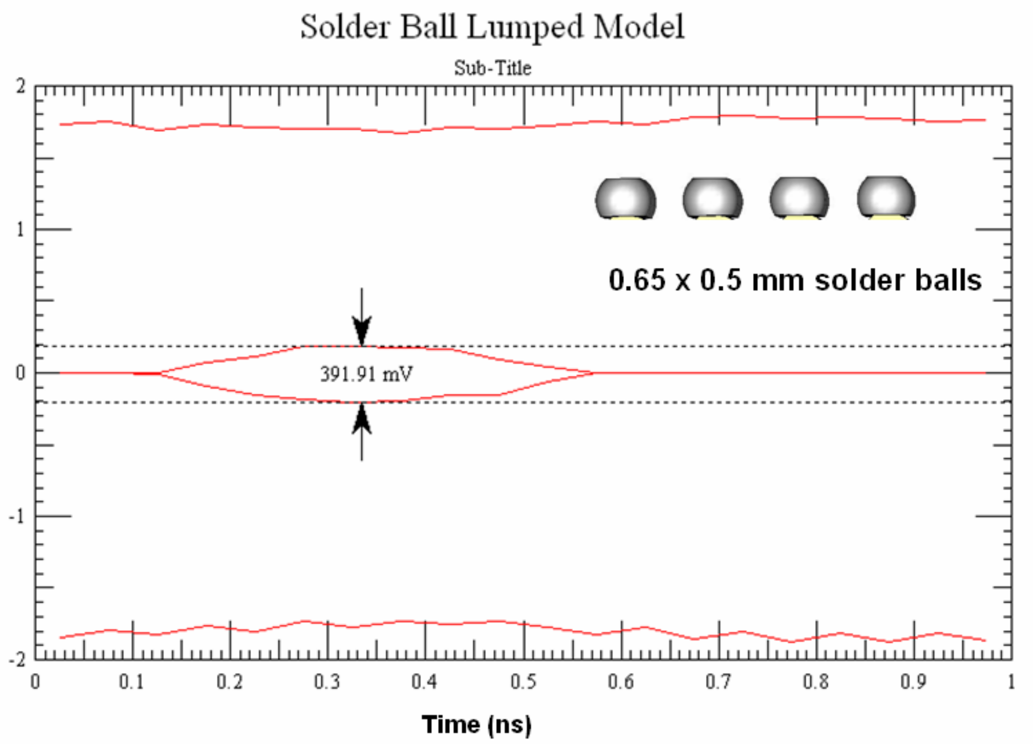

Fig. 7. Eye diagram for 2.5 Gbps trace-ball-trace interconnect model. Signaling scheme is the same as Fig. 6, with solder ball geometry included.

Using Sigrity's Speed2000 package, we analyzed the power impedance of a CapCore circuit with a checkerboard configuration composed of 12 segments of 14 columns of capacitors per segment. Each column had 3 capacitors (ESL = 400 pH, ESR = 0.1 mΩ) with via segment resistance 0.4 mΩ and inductance 200 pH. The power pins were tied together on the substrate and PCB side, as were the ground pins, forming an effective 2-port. An equivalent SPICE model was then created and analyzed. Power impedance is very low in the mid-frequency band, with circuit resonance between 3 and 7 MHz, depending upon capacitor size. The convergence of results above 30 MHz occurs as the resistance and inductance begin to dominate; decoupling for frequencies higher than this must be done on or near the die.

Note that the random configuration applies to layouts in which the power-ground pairs or clusters are distributed around the footprint without any particular order or regular arrangement. We applied another version of this new CapCore interposer to an actual application with this random configuration (22-layer PCB, 10-layer substrate) in which more than 400 mid-frequency decoupling capacitors were required on the PCB for low power impedance and noise suppression. We analyzed the complete system results both with and without an interposer composed of approximately 245 0.22µF 0402 chip capacitors and the identical parasitics of the checkerboard configuration. Results showed a remarkable improvement over the existing solution: nearly 10x lower improvement in the power delivery impedance, as seen from the die, for frequencies up to 400 MHz (see Fig. 8). Moreover, when the 400+ PCB capacitors are removed, the power impedance curve is virtually identical, proving that CapCore replaces the extensive (and expensive) mid-frequency PCB decoupling.

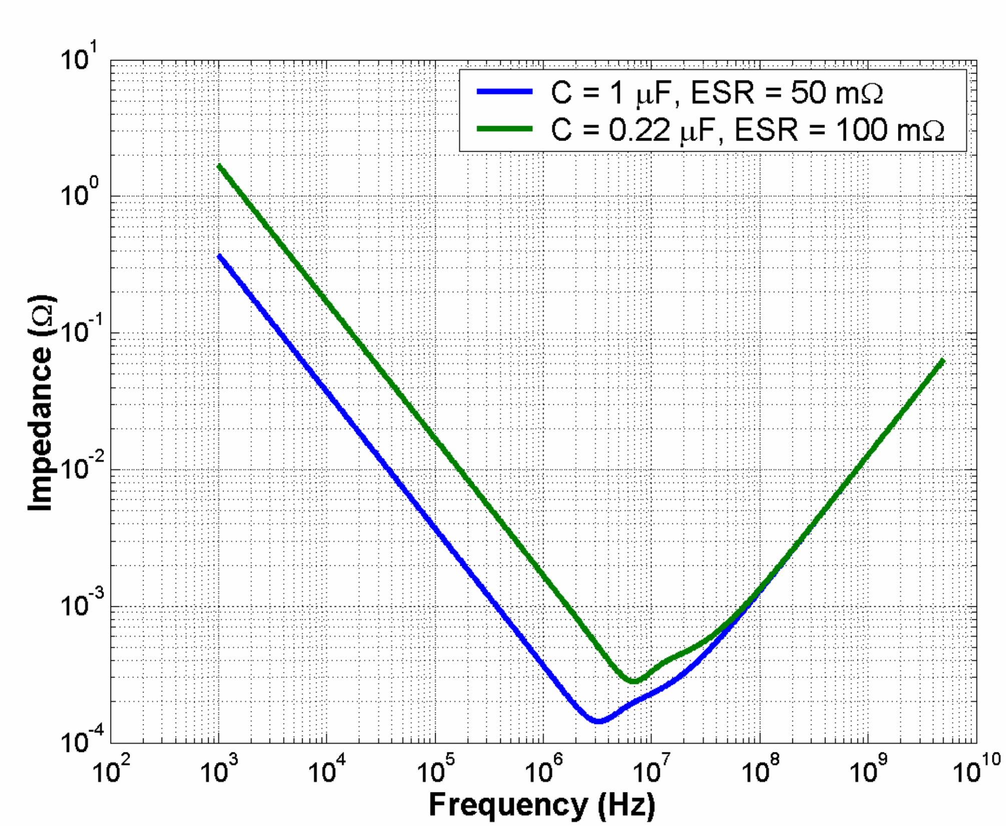

Fig. 5. Power impedance curves for the interposer with all power pins tied together and all ground pins tied together at the top and bottom.

Conclusions

We have presented a new and unique decoupling technology and demonstrated its improvements to system power delivery and suppression of transient noise on the power system. Power impedance results for an actual customer application show that CapCore is a superior alternative to hundreds of PCB decoupling capacitors, providing lower power impedance and stabilizing the core voltage. Proper PCB design with this novel interposer will substantially reduce both PCB complexity and total applied cost. More importantly, it will allow higher microprocessor operating frequencies and better die bin yielding, giving system vendors a key competitive advantage.

References

[1] 2001 Edition of the International Technology Roadmap for Semiconductors, http://public.itrs.net.

[2] A. Waizman, C. Y. Chung, “Package Capacitors Impact on Microprocessor Maximum Operating Frequency,” IEEE Electronic Components and Technology Conference, pp. 118-122, Orlando FL, May 2001

[3] C. Y. Chung, A. Waizman, “Optimum Signal Integrity Through Appropriate Analysis of Signal Return Path and Power Delivery,” IEEE Electronic Components and Technology Conference, pp. 1402-1407, Orlando FL, May 2001.

[4] S. Crane and J. G. Nickel, “It’s the system, stupid!” Connector Specifier *20*, No. 4 (2004).

[5] V. Cyr, I. Novak, N. Biunno, J. Howard, “ARIES: Using Annular-Ring Embedded Resistors to Set Capacitor ESR in Power Distribution Networks,” Electrical Performance of Electronic Packaging Conference, Boston MA, October 2001.

[6] A. C. Cangellaris, “Electromagnetic Modeling Methodologies for Multi-GHz PCB Signaling” Intel Technology Symposium 2001.