Minimizing Ohmic Loss and Supply Voltage Variation Using a Novel Distributed Power Supply Network

M. Budnik and K. Roy
Purdue University
School of Electrical and Computer Engineering
465 Northwestern Avenue
West Lafayette, IN 47907 USA
mbudnik, kaushik@purdue.edu

Abstract

IR and di/dt events may cause ohmic losses and large supply voltage variations due to system parasitics. Today, parallelism in the power delivery path is used to reduce ohmic loss while decoupling capacitance is used to minimize the supply voltage variation. Future integrated circuits, however, will exhibit large enough currents and current transients to mandate additional safeguards. A novel, distributed power delivery and decoupling network is introduced reducing the supply voltage variation magnitude by 67% and the future ohmic loss by 15.9W (compared to today's power delivery and decoupling networks) using conventional processing and packaging techniques in a 130nm technology node.

1. Introduction

As the transistor count and performance expectations of integrated circuits have grown with each technology generation, their power consumption has also grown, approximately doubling every thirty-six months [1]. To minimize the increase in power consumption and maintain dielectric reliability [2], the supply voltage of high performance microprocessors has been scaled down. As the operating voltage decreases and integration density increases, however, more current is required to meet the growing power requirements. Our research addresses two significant problems created by these circumstances, each of which will continue to grow as time (and technology) progress: IR and di/dt events.

In an IR event, ohmic losses occur as current flows through the parasitic resistance of the external voltage regulator module (VRM), the system's board, and the microprocessor's package [1]. Today, with a modest 80A IR event, these ohmic losses can exceed 13W for typical conditions [3, 4]. Recall, the power dissipated in a resistor is proportional to the square of the current. Therefore, the ohmic losses in the power delivery path can be reduced by delivering a microprocessor's input power at a higher voltage and lower current. Analogous to the delivery of power to commercial and residential structures (Fig. 1), we have developed a novel, step-down power conversion architecture for the integrated circuit to deliver the required operating voltage in a distributed fashion (Fig. 2). Each regulation node in our distributed power supply network (DPSN) provides power necessary for a small geographic portion of the die. The regulation nodes we designed "transform" input power from the high input voltage to the required operating voltage.

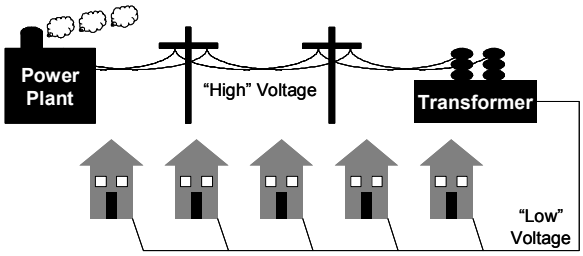

Fig. 1. Example of commercial power delivery

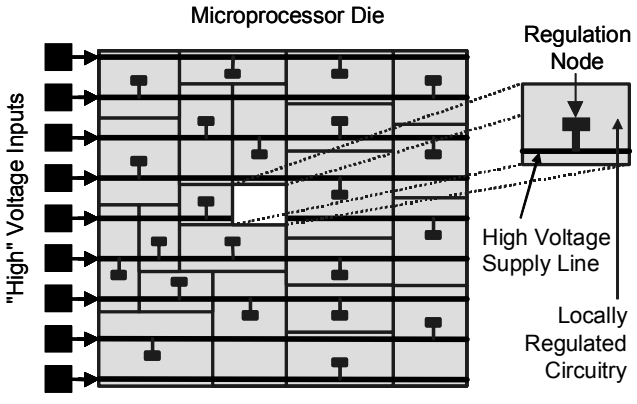

Fig. 2. Distributed Power Supply Network (DPSN)

Our DPSN can also be used to reduce the supply voltage variations caused by changes in the operating current (di/dt events). Fig. 3 shows a conventional power delivery and decoupling network (PDDN). Component values are from [3, 5-7]. The PDDN can be modeled as three decoupled loops, where each loop is a second order system with its own resonant frequency [5, 8]. The first loop is composed of the integrated decoupling capacitance (C_{DIE}), the organic land grid array (OLGA) parasitic resistance and inductance, and the discrete land side capacitors which are mounted within the microprocessor package (C_{LSC}). The second loop is composed of C_{LSC}, the parasitics of the socket and package, and the PC board (PCB) capacitors (C_{PCB}). The final loop is composed of C_{PCB}, the parasitics of the VRM connector and PCB, and the VRM output capacitors (C_{OUT}).

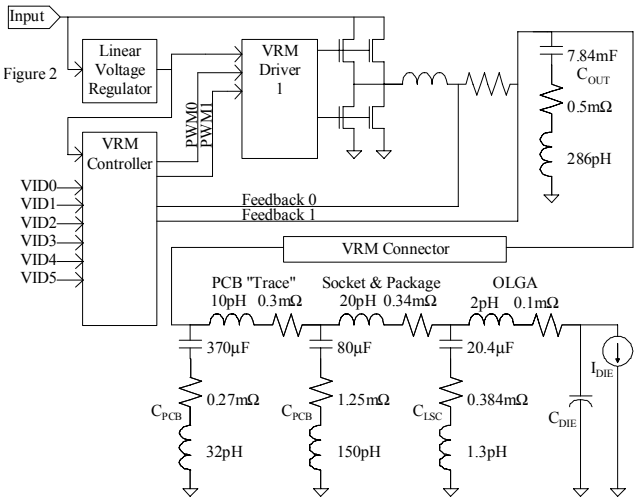

Fig. 3. Example PDDN for a 90nm Microprocessor

If the loops are well dampened and examined in isolation, the resonant frequency for each loop is given by:

$$f_{\mathrm{RESONANT}} = 1/\left(2\pi\sqrt{\mathrm{LC}}\right). \qquad (1)$$

For example, in the first resonant loop, C is given by the parallel value of C_{DIE} and C_{LSC}, and L is the value of L_{OLGA} and the equivalent series inductance (ESL) of C_{LSC} in series. In conventional PDDNs (as in Fig. 3), the inductance and capacitance values increase further away from the microprocessor die. This leads one to expect:

$$f_{\mathrm{RESONANT},3} < f_{\mathrm{RESONANT},2} < f_{\mathrm{RESONANT},1}. \qquad (2)$$

These three resonant frequencies correspond to the three superimposed droops observed in the supply voltage variation in response to a di/dt event (Fig. 4). These supply voltage variations can adversely affect the integrated

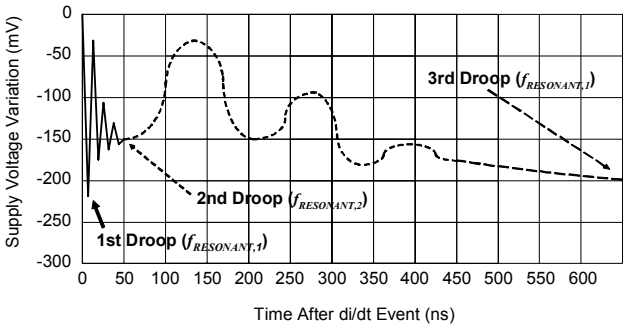

Fig. 4. Example Supply Voltage Variation Droops

circuit's performance [1, 9]. Note, the magnitude of the three droops will vary with the PDDN component values and the magnitude and duration of the di/dt event. As operating currents increase, the magnitude of di/dt events will also increase [10, 11]. This will exacerbate the supply voltage variation. However, the distributed regulation nodes in our DPSN can rapidly respond to di/dt events and minimize the three supply voltage variation droops.

Therefore, by providing distributed voltage regulation in our DPSN, we can achieve significant reduction in the supply voltage variation while minimizing the ohmic losses associated with IR events. In this paper, we present a novel power delivery network which:

- Reduces future IR event ohmic loss by over 15W
- Reduces di/dt event voltage droop magnitude by 67%
- Reduces di/dt event voltage droop over time by 98%
- Uses a standard silicon CMOS fabrication process
- Uses standard microprocessor packaging and package mounted (C_{LSC}) decoupling capacitors

2. Conventional Power Delivery and Decoupling Networks

For conventional microprocessors, supply voltage variation is minimized with external decoupling capacitance (C_{OUT}, C_{PCB}, and C_{LSC}) and integrated decoupling capacitance (C_{DIE}). We subjected the Fig. 3 PDDN (with C_{DIE}=1µF) to four different 100A current transients to quantify the supply voltage variation. The results are shown in Fig. 5 and Table 1. The maximum supply voltage variation magnitude increases as the duration of the 100A di/dt event decreases. Decreasing the duration of the current transient past the time constant of the resonant loop, however, has minimal impact upon the voltage variation [5]. Therefore, we will use 100A/1ns current transients for the comparison of our DPSN to conventional PDDNs.

Table 1 also illustrates the maximum first droop magnitude for several values of C_{DIE}. Also shown is the minimum die area required to implement C_{DIE} in a 130nm technology [12]. The inverse relationship between

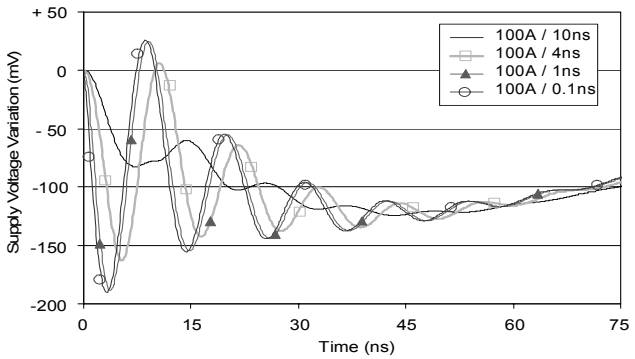

Fig. 5 Supply Voltage Variation in PDDN

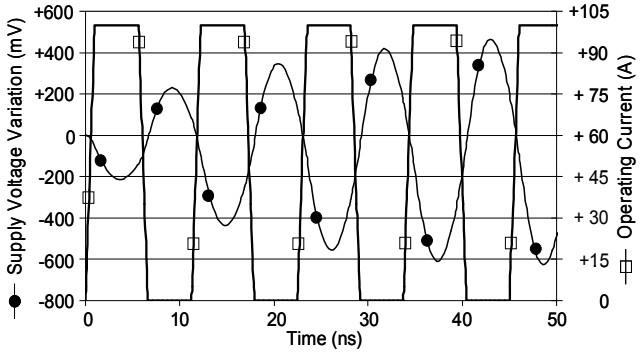

Fig. 6. Worst Case Supply Voltage Variation

C_{DIE} (nF)	Area (cm^2)	Supply Variation (mV)
1000	0.901	186
500	0.45	225
250	0.225	354

Table 1. C_{DIE} vs. Theoretical Supply Voltage Variation

decoupling capacitor die area and supply voltage variation is clear. Note, even with 0.901cm^2 of die area dedicated to C_{DIE}, the supply voltage variation approaches 200mV.

The maximum supply voltage variation magnitude is the key metric used in measuring the supply voltage variation in [1, 5, 7, 10]. It does not, however, define the quantity of the supply voltage variation over time. Therefore, we propose a new metric, VNS (Volt-ns), to quantify the amount of voltage droop over time. The VNS value of the supply voltage variation is found by integrating the negative area of the voltage droops. Once the variation magnitude is less than 10mV, we consider the microprocessor operating voltage to be within specification and the variation to have ended. Table 2 summarizes the maximum supply voltage variation droop magnitude and VNS values for the 90nm microprocessor conventional PDDN with C_{DIE} = 1µF in response to the 100A di/dt events in Fig. 5:

di/dt Event	1st Droop Magnitude	VNS
100A / 10ns	82.5 mV	-9.63
100A / 4ns	163 mV	-9.86
100A / 1ns	186 mV	-9.99
100A / 0.1ns	189 mV	-9.99

Table 2. PDDN Supply Voltage Variation

If repeated operating current transients occur near the first loop resonant frequency, larger supply voltage variations can arise. We subjected the 90nm micro-processor conventional PDDN (Fig. 3), to a periodic 100A/1ns current transient. The worst case supply voltage variation magnitude of 628mV was observed at 93MHz (Fig. 6). This variation is occurring faster than a VRM can respond [3, 4]. Therefore, it will continue as long as the periodic operating current transients persist.

Between di/dt events, ohmic loss occurs in all the elements of a PDDN power delivery path: VRM MOSFET switches, inductors, sense resistors, connector, PCB, microprocessor socket, and the microprocessor package. From [4], one can calculate the ohmic loss in the VRM MOSFET switches, inductors, and sense resistors. For an 80A IR event, we calculated the ohmic loss in each power delivery path component. The results are shown in Table 3. Summing the dissipated power in the conventional PDDN components, we found the typical ohmic loss during an 80A IR event to be 13.6W. (If worst case component values are used in the power delivery path, however, the ohmic loss increases to 21.8W). [13] predicts an increase in microprocessor operating current to 248A by 2016 (delivered at 0.8V). If future conventional PDDNs can achieve 80% efficiency [3] under these conditions, they will still result in 49.5W of ohmic loss.

PDDN Component	Ohmic Loss (W)
VRM MOSFETs	6.01
Inductors	1.60
Sense Resistors	0.64
PCB	3.20
Socket and Package	2.18

Table 3. Typical Ohmic Loss in Conventional PDDN

3. Distributed Power Supply Network

In our proposed distributed power supply network (DPSN), the distributed voltage regulation nodes convert the input power from the "high" VRM output voltage to the desired operating voltage (Fig. 1). The size and placement of the voltage regulator nodes are determined by the expected current draw of the circuits within each node voltage regulator's local area. We implemented a 2:1 switched capacitor voltage regulator [14] appended with a linear voltage regulator (SCLVR) for each of the distributed voltage regulation nodes in our DPSN (Fig. 7). The 2:1 switched capacitor voltage regulator provides the high

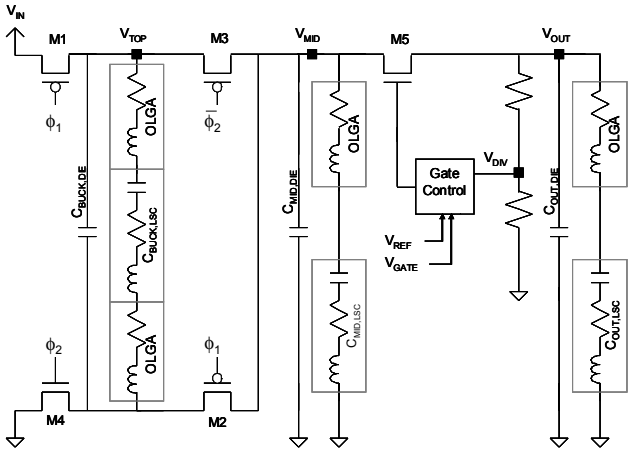

Fig. 7. SCLVR Functional Diagram

efficiency power conversion and the linear voltage regulator filters the switching noise and enables a fast response to di/dt events.

In our SCLVR circuit, a two-phase, non-overlapping clock is used to toggle an array of MOSFET switches, transferring charge from the input to C_{MID}. During the first clock phase, charge is transferred from the input, through C_{BUCK}, to C_{MID}. V_{TOP} approaches V_{IN}, and V_{MID} approaches $V_{IN}/2$. During phase two, the load current draws charge from C_{BUCK} and C_{MID}. This results in the decay of V_{TOP}, V_{MID}, and V_{OUT}. If the output voltage varies from its expected value, this is reflected at V_{DIV}. The gate control block then adjusts the gate voltage of M5 until V_{OUT} reaches its expected value. This allows the SCLVR to respond quickly to a di/dt event.

There are two differences between conventional voltage regulators and our SCLVR [15]. First, we provided a single, external reference voltage to all the distributed regulator nodes on the die. This eliminates the need to replicate a consistent reference voltage for each node. Second, we provided an external "high" gate voltage to the gate control block to increase the linear voltage regulator current drive [16].

In our research, we implemented a simple DPSN of twenty-three SCLVR nodes in a 130nm technology [12]. Each node is capable of sourcing 4.39A. M1-M4 were implemented with 2.5V "thick oxide" transistors. M5 used a 1.5V "thin oxide" MOSFET. Biasing M5 with the external "high" gate voltage, however, may result in dielectric breakdown due to gate leakage [2]. Therefore, we simulated the gate leakage current of M5 with the Taurus device simulator. We determined for a SiO_2 thickness greater than 1.5nm, the dielectric failure rate is less than 100PPM over ten years for worst case conditions (125C and 2.5V supply voltage). [17]

Based upon the sizes of M1-M5 in Fig. 7, the minimum area consumed by M1-M5 was 0.27cm^2 using the 130nm technology aggressive design rules [12]. (We will see this area can be offset by a reduction in C_{DIE}.) The DPSN uses the same external power delivery and decoupling network found in the conventional PDDN (Fig. 3).

The layout for our 130nm test chip is shown in Fig. 8.

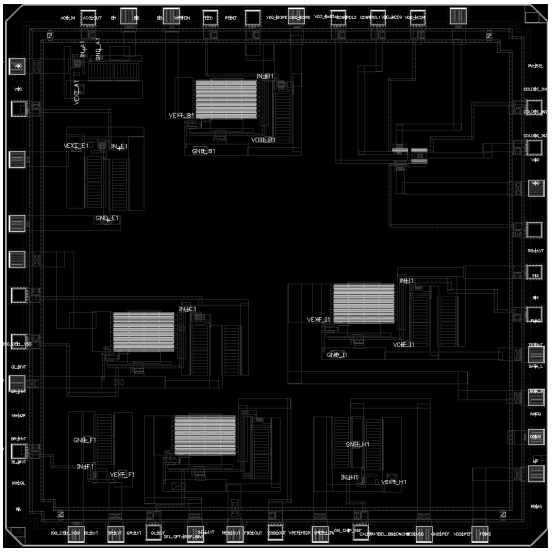

Fig. 8. DPSN Test Chip

4. Reduction of Supply Voltage Variation

We compared the supply voltage variations of the conventional PDDN and our DPSN under various conditions. In the conventional PDDN, we switched the current source from 1A to 101A. For the DPSN, we switched all of the current sources in the twenty-three nodes simultaneously from 43.5mA (1A total) to 4.39A (101A total). Our desired microprocessor operating voltage was 1V. The externally supplied V_{GATE} signal we used was 2V. We used a net integrated decoupling capacitance (C_{DIE}) of 1μF for both networks. Based upon the findings of [5], we used a simple VRM model for all simulations. This model is convenient to use and allows for very accurate modeling of the first and second droops in response to a di/dt event.

From Fig. 7, as V_{MID} approaches 1V, our DPSN SCLVR efficiency increases, but the regulation nodes will deliver less energy following a di/dt event. This reduces the ability of our linear voltage regulator circuit to minimize the supply voltage variation. Fig. 9 illustrates how reducing V_{MID} impairs the DPSN's ability to reduce the supply voltage variation following a 100A di/dt event.

Fig. 9 also shows how the DPSN linear voltage regulator with an input voltage of 1.1V is unable to regulate the operating voltage before the second droop supply voltage variation occurs (from 20-70ns). If the DPSN's nodes were able to provide more charge at a reduced input voltage, the improvement would be more drastic. Recall, the drain current of a transistor is a function of the width of the

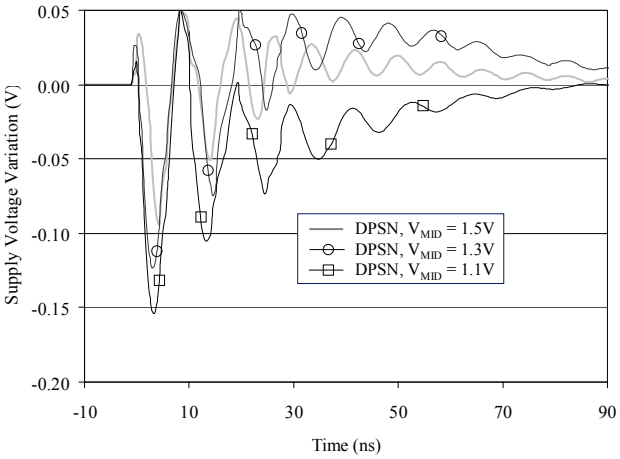

Fig. 9. DPSN Supply Voltage Variation vs. V_{MID}

transistor [16]. Therefore, if M5 in our SCLVR is widened, it will be able to deliver additional charge in the same amount of time. This will improve the ability of our DPSN to reduce the supply voltage variation. We lowered V_{MID} to 1.05V and increased the width of M5 by an order of magnitude. (With V_{MID} = 1.05V, the maximum efficiency of our SCLVR linear voltage regulator would be about 95%.) All other parameters were maintained from the previous simulation. In Fig. 10, we compare the performance of our modified DPSN to a conventional PDDN. Our DPSN reduces the first and second droop supply voltage variation, and has a much improved VNS value (-3.53VNS vs. -9.99VNS). Our simulations show the supply voltage variation would continue to decrease if we further increased the width of M5 in our SCLVRs.

Finally, we examined the performance of our DPSN with V_{IN} = 2.5V. Our switched capacitor circuits operated at 10MHz. The externally supplied gate voltage was 2.5V. All other parameters were from the previous simulations. The supply voltage variation of our DPSN and the PDDN are shown in Fig. 11. Our DPSN has reduced the maximum supply voltage variation magnitude from 186mV to 62mV (a 67% improvement). In addition, our VNS value is now only -0.174VNS (vs. -9.99VNS), a 98% improvement.

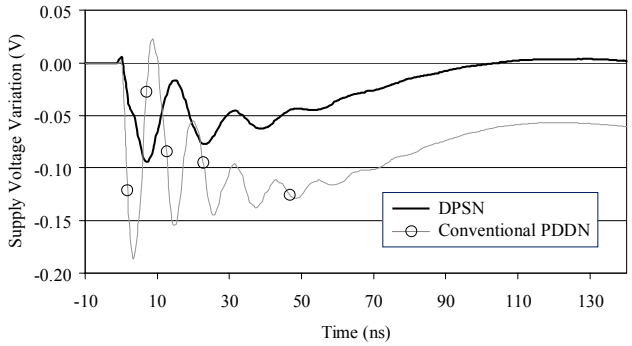

Fig. 10. DPSN with V_{MID} = 1.05V vs. Conventional PDDN

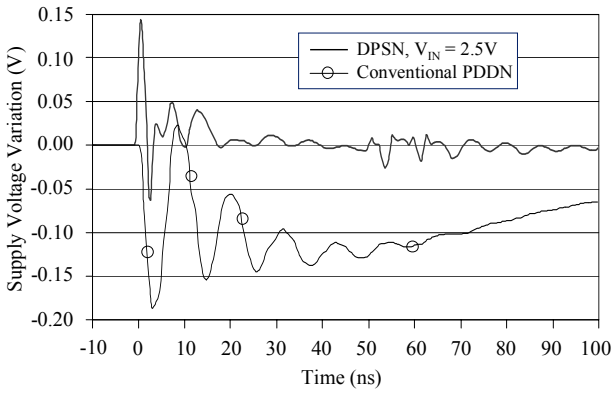

Fig. 11. DPSN vs. Conventional PDDN

The amount of reduction in the supply voltage variation will depend strongly upon the input voltage, the amount of integrated decoupling capacitance, and the size of the transistors in our SCLVR circuit. An appropriate trade-off between the improvement in the supply voltage variation, the DPSN efficiency, and the die area consumed by the DPSN would need to be made for each application space.

5. Reduction in IR Event Ohmic Loss

In 2016, microprocessors are expected to consume 198W of power (248A at 0.8V). Based upon a conventional PDDN efficiency target of 80% [3], the external power delivery path components will dissipate:

$$P_{Loss} = P_D(1-\eta)/\eta \qquad (3)$$

$$P_{Loss,PDDN} = (198\text{W})(1\text{-}0.8)/0.8 = 49.5\text{W}$$

Let us now examine the ohmic loss of our DPSN in 2016. If the DPSN input voltage is 1.7V, the DPSN power conversion efficiency can be nearly 90%. Therefore, from (3), the power lost in the DPSN converting the input power from 1.7V to 0.8V would be:

$$P_{Loss,DPSN} = (198\text{W})(1\text{-}0.9)/0.9 = 22.0\text{W}.$$

Transmitting the microprocessor's input power at 1.7V instead of 0.8V, however, can greatly reduce the ohmic loss in the external power delivery path components. Recall, the power dissipated in a resistive element is proportional to the square of the current. Therefore, from [3, 4], we calculated the DPSN can reduce the ohmic loss in the external power delivery path components from 49.5W to 11.6W. (This 37.9W improvement includes the decreased switching loss in the VRM.) There is a net improvement in the system's ohmic loss of:

$$P_{Savings} = \left(P_{Loss,PDDN} - P_{Loss,External,DPSN}\right) - P_{Loss,DPSN} \qquad (4)$$

$$P_{Savings} = (49.5W - 11.6W) - 22.0W = 15.9W$$

If our DPSN can reduce the microprocessor's integrated capacitance (C_{DIE}), it can result in additional ohmic loss savings. According to our simulations, the DPSN can achieve the same supply voltage variation magnitude as the PDDN (186mV), while reducing C_{DIE} from 1µF to 500nF. [18] shows the leakage power of a 175nF integrated capacitor to be 10W (with 50% of the capacitance contributed by 1.5nm "thin" silicon dioxide capacitors and 50% by 2.5nm "thick" silicon dioxide capacitors). This results in 57.1W of leakage power loss for a conventional microprocessor with 1µF of integrated decoupling capacitance. The smaller, 500nF DPSN microprocessor decoupling capacitance, however, will result in only 28.6W of leakage power. If other elements of the DPSN leakage current do not significantly increase, the improvement in the ohmic loss of a 500nF DPSN microprocessor system would be:

$$P_{Savings} = (49.5 - 11.6W) - 22.0W + (57.1W-28.6W)$$

$$P_{Savings} = 44.4W$$

In addition, if our DPSN can reduce C_{DIE} from 1µF to 500nF, Table 1 indicates the integrated decoupling capacitance die area can be reduced by $0.451cm^2$. This reduction more than offsets the $0.27cm^2$ DPSN power transistors (M1-M5) die area.

6. Conclusion

As microprocessor operating currents increase, the supply voltage variation associated with di/dt events and the ohmic loss due to IR events will continue to increase. We have proposed a novel distributed power supply network to address both problems in future microprocessors using conventional CMOS processing and standard microprocessor packaging. Compared to conventional power delivery and decoupling networks, our proposed distributed power supply network is capable of reducing the maximum supply voltage variation magnitude by 67% and the supply voltage variation over time by 98%. Our network also shows a future reduction in the ohmic loss of 15.9W.

Acknowledgements

The authors gratefully acknowledge the generous support of a MOSIS Educational Program (MEP) Research Grant in the development of this research project.

References

[1] R. Mahajan, *et al.*, "Emerging directions for packaging technologies", pp. 62-75, Intel T., J., Vol. 6, Issue 2, 2002.

[2] J. H. Stathis, "Reliability limits for the gate insulator in CMOS technology", IBM J. of Res. and Dev., pp. 265-286, March, 2002,

[3] Intel Corporation, "Voltage regulator module and enterprise voltage regulator-down 10.1", June 2004.QQQ

[4] Maxim Integrated Products, "MAX8524/MAX8525 2- to 8-phase VRM 10/9.1 PWM controllers with precise current sharing and fast voltage positioning", April, 2004.

[5] Y. L. Li, *et al.*, "Enhancing power distribution system through 3D integrated models, optimized designs, and switching VRM model", 2000 Electronic Components and Technology Conference, pp. 272-277.

[6] Intel Corporation, "Intel xeon processor with 800 MHz system Bus", June, 2004.

[7] H. Braunisch, S. N. Towle, R. D. Emery, C. Hu, G. J. Vandentop, "Electrical performance of bumpless build-up layer packaging", 2002 Electronic Components and Technology Conference, May, 2002.

[8] P. -L. Wong, F. Lee, X. Zhou, J. Chen, "VRM transient study and output filter design for future processors," Proceedings of the 24th Annual Conference of the IEEE, Vol.1, pp. 410-415, August, 1998.

[9] P. Larsson, "Power supply noise in future IC's: a crystal ball reading", Custom Integrated Circuits, 1999, Proceedings of the IEEE 1999, pp. 467-474.

[10] P. Gelsinger, "Discontinuities driven by a billion connected machines", IEEE Design and Test of Computers, pp. 7-15, January, 2000.

[11] B. Pohlman, "Overcoming the barriers to 10GHz processors", Microprocessor Forum Keynote, October, 2001.

[12] IBM Microelectronics Division, "CMOS8RF (CMRF8SF) Design Manual", May, 2005

[13] International Technology Roadmap for Semiconductors, "2004 Overview and Summary", 2004.

[14] Maxim Integrated Products, "Application Note: 718, Properties of the Charge-Pump Voltage Splitter" www.maxim-ic.com/an718.

[15] Maxim Integrated Products, "Application Note: 737, Choosing the Right Power-Supply IC for your Application", www.maxim-ic.com/an737.

[16] Y. Taur, T. H. Ning, *Fundamentals of modern VLSI devices*, Cambridge University Press, Cambridge, 1998.

[17] M. Budnik, K. Roy, "Implications of SiO2 Breakdown in an integrated nanoscale power supply," accepted for publication in the Proceedings of the 2005 International Semiconductor Device Research Symposium (ISDRS), Bethesda, MD, December 7-9, 2005.

[18] H. H. Chen, J. S. Neely, M. F. Wang, G. Co, "On-chip decoupling capacitor optimization for noise and leakage reduction", Proc. 16th Symp. on IC and Systems Design.