Feasibility of Monolithic and 3D-Stacked DC-DC Converters for Microprocessors in 90nm Technology Generation

Gerhard Schrom[1], Peter Hazucha[1], Jae-Hong Hahn[2], Volkan Kursun[3], Donald Gardner[1],
Siva Narendra[1], Tanay Karnik[1], Vivek De[1]

[1] Circuit Research, Intel Labs. [2] Mobile Platforms Group, Intel. [3] EE and CE Dept., University of Rochester, NY.

2111 NE 25th Ave., M/S JF3-334, Hillsboro, OR 97124, USA.
+1(503)712-4349

gerhard.schrom@intel.com

ABSTRACT

Rapidly increasing input current of microprocessors resulted in rising cost and motherboard real estate occupied by decoupling capacitors and power routing. We show by analysis that an on-die switching DC-DC converter is feasible for future microprocessor power delivery. The DC-DC converter can be fabricated in an existing CMOS process (90nm-180nm) with a back-end thin-film inductor module. We show that 85% efficiency and 10% output voltage droop can be achieved for 4:1, 3:1, and 2:1 conversion ratios, area overhead of 5% and no additional on-die decoupling capacitance. A 4:1 conversion results in 3.4x smaller input current and 6.8x smaller external decoupling.

Categories and Subject Descriptors B.7.0 Hardware, Integrated Circuits, General

General Terms Design, Performance, Theory

Keywords 3-D integration, DC-DC converter, integrated magnetics, on-die switching converter, power delivery

1. INTRODUCTION

Maximum current consumption, current density and current transient demands of high performance microprocessors have been increasing by 50% per generation in spite of supply voltage (V_{CC}) scaling (see Figure 1). Reduction of V_{CC} makes the problem of delivering larger currents with high conversion efficiency even more challenging, especially since the maximum acceptable V_{CC} variation is on the order of 10% of the target V_{CC} value [1]. Employing traditional methods [3] to meet V_{CC} variation targets on the microprocessor die in the presence of large current transients requires a prohibitively large amount of on-die decoupling capacitance (decap). Alternately, the motherboard voltage regulator and converter module (VRM) is required to operate at a higher frequency. Expensive solutions need to be employed to minimize impedance (Z_{ext}) of the off-chip supply network carrying high current from the VRM to the die across board, socket and package traces and reduce the parasitic resistance and inductance between the VRM output and the on-die power grid [4] (see Figure 2). Excessive losses in the low-voltage, high-current distribution network are also imposing significant

burdens on system cooling. Increasing input voltage to the VRM [5], and moving the VRM closer to the microprocessor by integrating it either in the package or the die itself alleviates both problems. Introduction of thin-film inductors operating at frequencies above 100MHz [6][7] has opened the possibility of integrating a DC-DC converter on a single silicon chip.
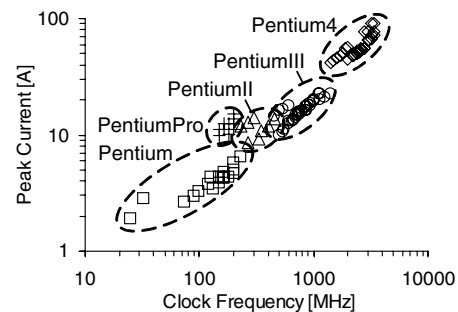

Figure 1. Increasing peak supply current of high-performance microprocessors.
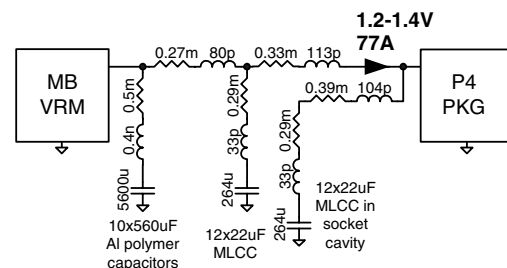

Figure 2. Simplified power delivery network for Intel Pentium4 processor on 90nm process.

1.1 Near-Load Converter Insertion

Inserting a DC-DC converter near the load (see Figure 3) will reduce the VRM current I_{ext} and allows to increase the impedance Z_{ext}, i.e., reduces the decoupling requirement.
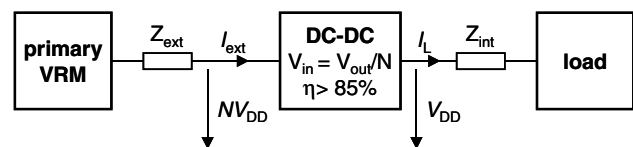

Figure 3. Insertion of a DC-DC converter near the load.

For a given conversion ratio of N:1 and an efficiency η, the current reduction is $I_{ext}/I_L = 1/N\eta$. With a converter-added droop of 5% the reduction of decoupling requirement is $1/(0.5N^2\eta)$. Figure

4 shows the expected improvements for η=85% efficiency. For a conversion ratio of 4:1 the VRM current is 3.4x smaller, and the off-chip decoupling capacitance is 6.8x smaller.
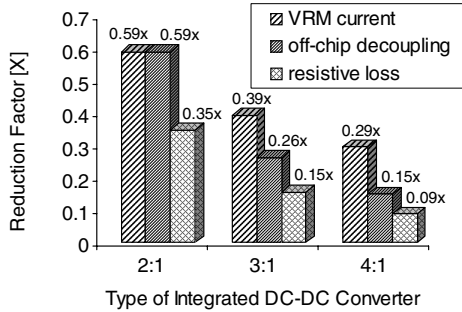

Figure 4. Reduction in VRM current, resistive loss, and off-chip decoupling requirement for integrated DC-DC converters with 85% efficiency.

1.2 DC-DC Converter Integration Schemes

In this paper, we propose and evaluate monolithic (1-chip) and 3D-stacked (2-chip) power delivery schemes utilizing an integrated DC-DC converter and a thin-film inductor technology. In the 1-chip scheme (Figure 5a), the converter is implemented on the microprocessor die packaged using flip-chip technology with Controlled Collapse Chip Connection (C4 bump) between a die and the package. The 1-chip scheme adds complexity to the logic processing technology, but provides the added benefit of reducing C4 bump currents that are limited by reliability considerations. In the 2-chip scheme (Figure 5b), a separate converter chip is "stacked" on top of the microprocessor die using a three-dimensional (3D) "through-hole" assembly technology in order to put the two chips in the closest possible proximity. This allows the process technology to be optimized separately for the converter chip, and does not impact the already scarce interconnect resources on the microprocessor chip.
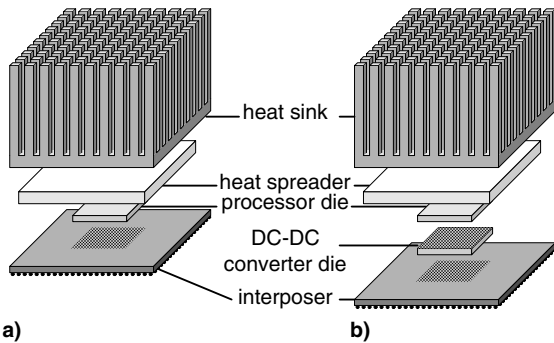

Figure 5. DC-DC converter on the microprocessor die (a) and on a 3D-stacked die with through vias (b).

We propose various converter topologies and concepts based on regular and coupled inductors [8], and compare their effectiveness in terms of efficiency, die area and impact on process complexity for high performance microprocessors in 90nm technology generation. We assumed a supply voltage, V_{CC}, of 1.2V and load current density of 100A/cm^2 for conversion ratios (N) ranging from 2:1 to 4:1. We also evaluate efficiency improvements

achievable by energy-recycling drivers and zero-voltage-switching (ZVS) applied to inductor bridges.

2. INTEGRATED DC-DC CONVERTERS

Integrated DC-DC converters differ from the traditional off-chip VRM in several ways. Typically, the VRM is a synchronous buck converter (Figure 6) that uses slow high-voltage transistors as switches in the bridge, a high-quality inductor L and a large decoupling capacitance C.
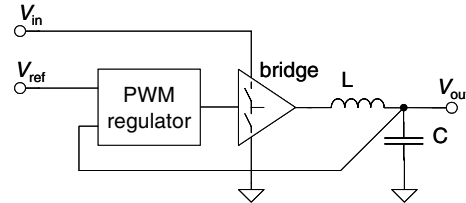

Figure 6. Buck converter topology.

An on-die converter, in contrast, has to meet the same V_{CC} droop requirement with much smaller decap already present on the microprocessor die to minimize die area impact. This requires smaller L and higher bridge frequency along with fast regulators. As a result, switching and conduction losses in the converter become much higher, thus making it difficult to achieve high conversion efficiency. Figure 7 illustrates the fundamental limitation of a buck converter's response to a load current step: The minimum time to accommodate a fast load current change ΔI is limited by the inductance L, i.e., $T_{ind} = L\Delta I/\min(V_{out}, V_{in}-V_{out})$, and the regulator loop delay T_{reg}. During this time, the current change has to be supported by the decoupling capacitor $C > \Delta Q/V_{droop}$, $\Delta Q = \Delta I (T_{reg} + T_{ind}/2)$.
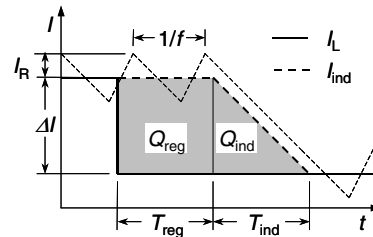

Figure 7. Load current step response in buck converters.

Smaller capacitance C can be achieved for a smaller filter inductor L which leads to a large ripple current. A multi-phase interleaved buck topology is used to cancel out the ripple current and suppress the ripple voltage at the output. Therefore, the I²R losses in the inductors and bridges, rather than the output voltage ripple impose the limit on the minimum size of L and the achievable droop of a buck converter.

2.1 Coupled Inductors

The tradeoff between the ripple current and output impedance is greatly alleviated by coupled-inductor topology (see Figure 8 and Figure 9). Compared to the equivalent buck converter, in circuits using coupled inductors with a coupling factor k, the effective output inductance L_s (leakage inductance), which is responsible for the droop response, is reduced by a factor of $L_s/L = (1-k)/(1+k)$. The total ripple current I_R, however, is still determined by L: $I_R = V_{out}(1-V_{out}/V_{in})/2fL$.
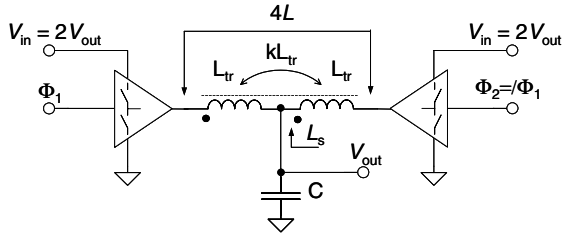
Figure 8. 2-phase 2:1 converter with coupled inductors.

Figure 9 shows a 4-phase converter using coupled inductors. The conversion ratio can be 4:1, 2:1, and 4:3. The circuit in Figure 9 can be generalized for 2^m phases by parallel connection two 2^{m-1}-phase converters with two coupled inductors, allowing V_{out} to be a multiple of $V_{in}/2^m$. Further coupled-inductor topologies are shown in [8].
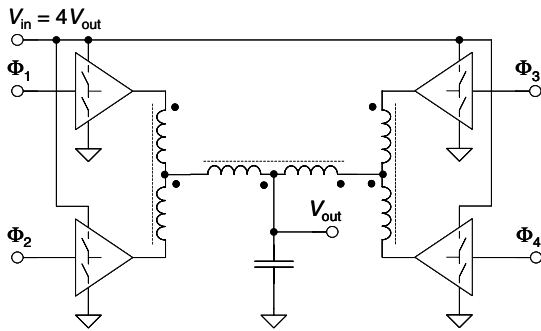

Figure 9. 4:1 converter topology with coupled inductors.

One of the key advantages of coupled-inductor topologies is that for small L_s, the transient voltage droop becomes $V_{droop} \approx \Delta I \sqrt{(L_s/C)} < 0.1 V_{CC}$, even without any regulation. By choosing $k<1$ or by adding a small extra inductance at the output V_{out} can be regulated to some extent at the cost of degraded droop response and output voltage ripple. The trade-off improves, however, as the number of phases is increased.

2.2 High-Voltage Drivers

High-voltage transistors are not readily available in high-performance logic process technologies used for microprocessors or stacked converter chip. Cascode bridge drivers can be used to support input voltages greater than V_{max}, where V_{max} is the reliability-limited highest transistor gate-source voltage allowed by the logic technology. Figure 10 shows a cascode bridge for $V_{in}=2V_{max}$ converters. The NMOS cascode M1 M2 forms a switch connecting the output y to ground when V_{G1} is V_{max}. When the switch is turned off ($V_{G1}=0$), the output voltage can rise to $2V_{max}$ whereas V_{D1} will rise only to $V_{max}-V_T$. The PMOS cascode M3 M4 operates accordingly, controlled by V_{G4} which lies between V_{max} and $2V_{max}$. The timing control ensures non-overlapping operation of the switches. None of the V_{GS} / V_{DS} will exceed V_{max}. The current into the auxiliary rail V_{max} is the supply current difference of the drivers U1 U2, which is typically small,.the V_{max} rail can be supplied by the converter output.
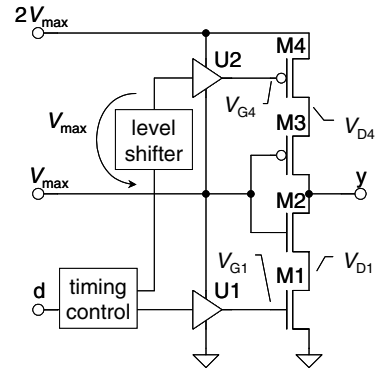

Figure 10. Cascode bridge supporting 2V$_{max}$.

Figure 11 shows the derivation of cascode bridges for higher voltages. The stack of inverters, U1, U2, and a third inverter, U3, together with the level shifter connected to d, form a 2:1 cascode bridge. By adding another rail, $3V_{max}$, the inverters U4, U5, U6, and another level shifter, a 3:1 cascode bridge is formed with the
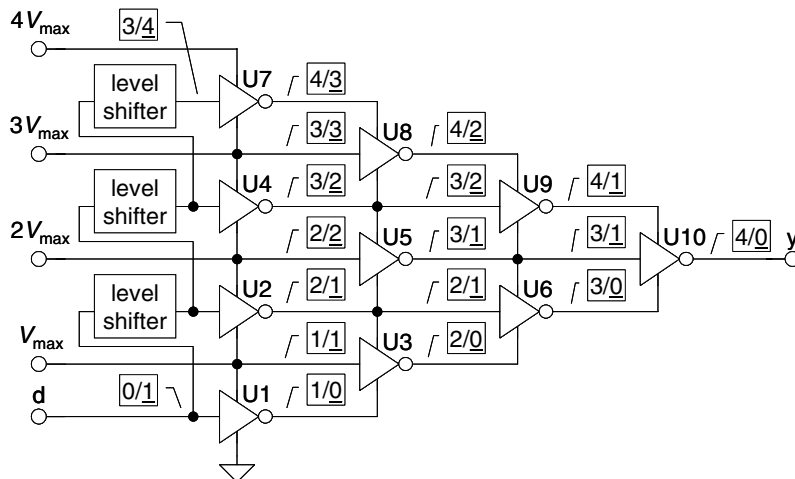

Figure 11. Derivation of an N:1 cascode bridge (*N*=4). The boxed numbers show the voltage levels in multiples of *V*$_{max}$ for the two input states, d=0 (0V) and d=1 (*V*$_{max}$).

output of U6 switching between 0V and $3V_{max}$. Adding rail $4V_{max}$, inverters U7 thru U10, and another level shifter, forms a 4:1 cascode bridge. Note, that all the devices in the inverters may not be required in the final circuit (e.g., the PMOS of U1). Also, separate inputs and level shifters may be used to control the inverter tree.

2.3 Integrated On-Die Inductors

Thin-film regular and coupled inductors need to be fabricated on the microprocessor die or the converter chip. Figure 12 shows cross-section of an inductor with Al or Cu wires, surrounded by insulation and CoZrTa magnetic core material [6]. To reduce eddy current losses, the core should be laminated and/or slotted.
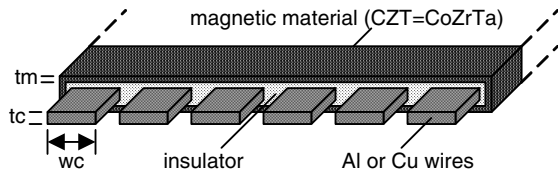

Figure 12 On-die thin-film inductor/transformer.

The inductor can be fabricated using either an existing top-level metal, with limited thickness due to pitch requirements (option A), or with an additional thick metal level (option B). The ratio L/R is the key figure of merit. The two technology options, A and B, with different metal and CoZrTa thicknesses correspond to *L/R* values of 50ns and 200ns respectively (Table I).

Table 1. Thin-film inductor properties.

technology	wire	magn. layer	L/R [ns]	k	I_{max} [mA/μm]
A: 50ns	~1.25μm Cu	2x1.5μm CZT	50	0.98	2.5
B: 200ns	4.0μm Cu	2x2.0μm CZT	200	0.99	2.5

The minimum metal width per unit current is limited to 0.4μm/mA by the saturation field B_{max}=1.4T and relative permeability μ_r =900 of CoZrTa. The wire width w_c is optimized using a 3D EM solver to produce sufficiently high values of *L/R* and *k* (see Figure 13). For a converter-added droop of 5%, w_c=8μm was the optimal value, resulting in *k*=0.98 and L/R=50ns.

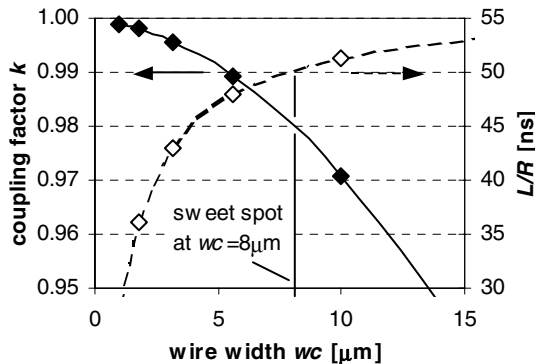**Figure 13. Trade-off between coupling *k* and *L/R* versus wire width w_c (gap *g*=2μm, t_c=1.25μm Cu, t_m=2x1.5μm CZT).**

2.4 Energy Recycling Drivers

Concepts of energy-recycling (adiabatic) bridge drivers and ZVS switching, needed to reduce switching losses, are illustrated in Figure 14. The energy for charging and discharging the bridge input and output capacitances is recycled through the input and output L. Timing parameters ΔT and T_{ZVS} can be tuned to minimize losses.
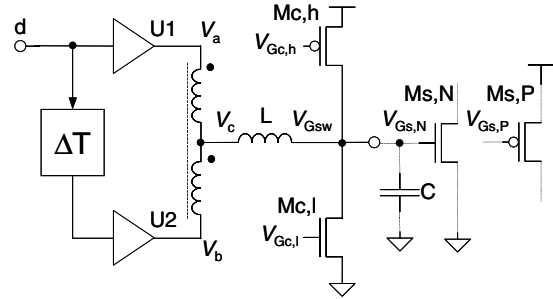

Figure 14. Energy recycling driver circuit.

A delay element, two drivers U1 U2, and an autotransformer drive the inductor L with a stair-case shaped voltage V_c to charge and discharge the load capacitance C (typically the gate capacitance of a bridge transistor Ms,N or Ms,P). The current through the inductor has a half-sine shape. At the end of a transition, the output voltage V_{Gsw} is close to either the ground or the supply and the respective clamping transistor Mcl or Mch is turned on (see Figure 15). In a falling transition, the energy stored in C is transferred through L, the transformer, and the output PMOS of U2 back to the supply rail, i.e., the energy is recycled.
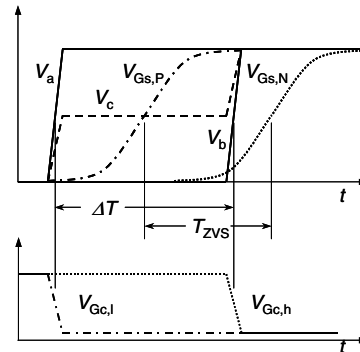

Figure 15. Energy recycling driver timing.

3. MODEL AND ANALYSIS METHOD

The DC-DC converter model assumes following input parameters: (a) converter process technology ranging from 180nm to 90nm, (b) inductor technology choice from Table 1, (c) decoupling cap, (d) conversion ratio, (e) worst-case load current transients, (f) driver configuration, (g) regulation mode, and (h) maximum allowed area overhead. The analysis self-consistently optimizes bridge transistor sizes, frequency *f*, ripple current I_R or *L* value, driver timing parameters ΔT & T_{ZVS} to produce the highest efficiency η, while meeting a worst-case Vcc droop (V_{droop}) requirement of 5%, subject to the specified area overhead constraint.

The efficiency model accounts for the three main power loss components, the resistive loss in the bridge, the capacitive loss in

the bridge and in the drivers and the resistive loss in the inductor, as well as eddy current and hysteresis losses in the magnetic core, and losses in the adiabatic drivers. While most quantities are modeled analytically, the voltage droop is computed by solving a transient differential equation numerically.

4. RESULTS AND DISCUSSION

Optimizations were run for three technology options: (a) 1-chip with L/R=50ns (90nm technology, inductors use existing top-level metal), (b) 1-chip with L/R=200ns (90nm technology, inductors use additional thick metal), and (c) 2-chip with L/R=200ns (130nm technology, thick top-level metal).

A monolithic DC-DC converter using coupled inductors implemented in a 200ns *L/R* inductor technology achieves more than 85% conversion efficiency when $k > 0.9$ for 2:1 conversion ratio with only 10% of the die area used for decaps (Figure 16). No regulation is necessary. Similar efficiency can be achieved for 2:1 conversion ratio by a traditional buck converter ($k = 0$) only if implemented with fast regulation on a second converter chip in the same process technology where all the available area in the second chip is occupied by additional decaps and bridge transistors. In this case, the inductor occupies the entire second chip area, in contrast to only 16% of die area consumed by coupled inductors (Table 2). Furthermore, efficiency will fall below 85% at higher conversion ratios. Converters based on un-coupled inductors require large area and are only viable for current densities much less than 100A/cm^2. Therefore, the following analysis will focus on DC-DC converters utilizing coupled inductors.
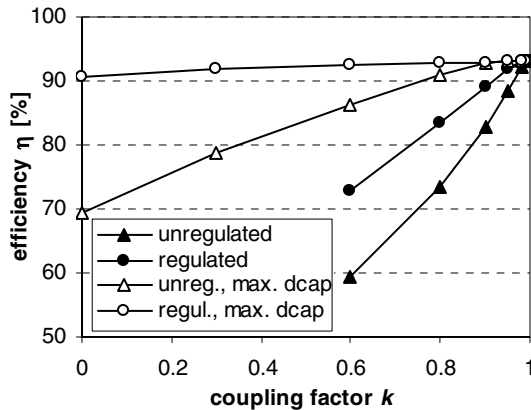

Figure 16. 2:1 converter efficiency vs. coupling factor (2-chip, L/R=200ns).

Table 2. 2:1 converters with and without inductor coupling (2-chip, L/R=200ns).

k	regulated	xtra dcap	η [%]	A_{ind} [%]	L [nH]	f [MHz]	I_R/I_L
0.99	N	0	93.0	16.0	3.39	118	0.75
0	N	max.	**69.3**	8.9	**0.12**	**1241**	2.00
0	**Y**	max.	90.7	**99.8**	1.70	105	1.68

Figure 17 shows efficiency and frequency of 2:1 converters vs. area overhead incurred by bridge transistors and inductor for DC-DC converters based on coupled inductors. The

maximum added V_{droop} is $< 5\%$ for all cases. Larger-area inductors and bridges with smaller switching frequency improve efficiency until the maximum (unconstrained) efficiency is reached. More than 85% efficiency can be achieved by less than 5% overall area impact for monolithic implementations with or without additional metal levels. An extra metal level for inductors improves efficiency by 5% because of lower resistance and relaxed inductor area constraint, at the expense of additional process complexity. The maximum efficiency achievable is similar to the 2-chip implementation.
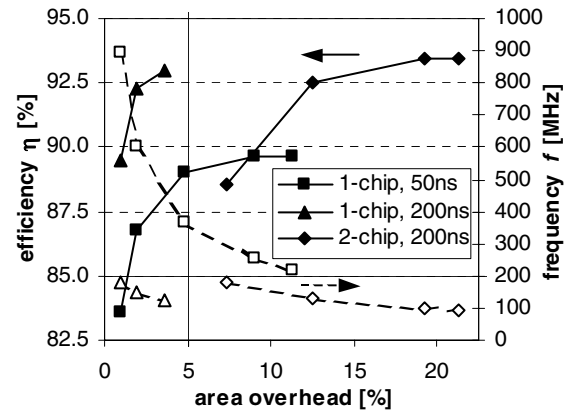

Figure 17. Impact of area constraint in 2:1 converters

Figure 18 shows best achievable efficiency vs. conversion ratio achievable for maximum 5% added voltage droop. Since both capacitance and resistance increase with the cascode bridge stack height and switching losses increase with input voltage V_{in}, efficiency decreases at higher conversion ratios. An efficiency of 85% can be achieved for conversion ratio of 4:1 with 200ns L/R coupled inductor technology implemented either in monolithic scheme or as stacked 2-chip scheme. The second chip can be implemented in an older and less expensive 180nm logic process technology that supports a V_{in} of 3.6V with bridge stack height of two.
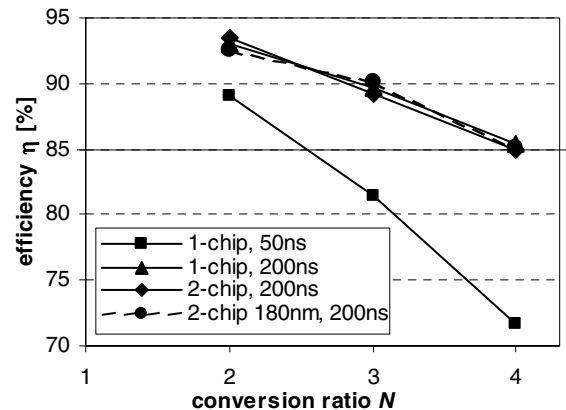

Figure 18. Efficiency vs. conversion ratio.

Table 3 shows the supply current reduction and the external impedance relaxation achievable as a function of conversion ratio. With a 4:1 conversion, current reduces by 0.3X, allowing the external impedance to be relaxed by 6.8X.

Table 3. Primary current reduction I_{ext}/I_L and relaxed primary impedance requirement Z_{ext}/Z_{int} for different conversion ratios

conversion ratio	η [%]	I_{ext}/I_L $1/\eta N$	Z_{ext}/Z_{int} $0.5/\eta N^2$
2:1	93.46	0.53	1.87
3:1	89.27	0.37	4.02
4:1	84.96	0.29	6.80

Table 4 shows the benefit from energy recycling and ZVS bridges applied to converters with coupled inductors. The efficiency improves by 6% for 2:1 conversion ratio and by 20% for 4:1 conversion ratio. Since energy recycling reduces the capacitive losses significantly, the optimal bridge size A_{brdg} increases, leading to lower bridge resistance and smaller switching frequency. Regulation, on the other hand, does not help significantly, since the droop is already small.

Table 4 Efficiency improvement witch adiabatic (energy recycling) drivers

	regulated	adiabatic	η [%]	A_{brdg} [%]	A_{ind} [%]	f [MHz]	I_P/I_L	L [nH]
2:1	Y	N	87.23	4.33	14.23	100	0.56	5.36
	N	**N**	**86.09**	**3.39**	**10.37**	**149**	**0.61**	**3.30**
	Y	Y	93.56	11.79	28.66	83	0.95	3.80
	N	**Y**	**93.46**	**11.64**	**27.08**	**96**	**1.04**	**3.00**
4:1	Y	N	60.32	6.25	13.35	204	0.72	3.06
	N	**N**	**59.85**	**6.87**	**14.24**	**199**	**0.75**	**3.02**
	Y	Y	86.88	21.46	63.29	68	1.15	5.75
	N	**Y**	**84.96**	**16.70**	**42.42**	**111**	**1.26**	**3.22**

5. SUMMARY

We analyzed the feasibility of monolithic and stacked on-die DC-DC converters for microprocessor power delivery, based on a numerical model, and we proposed circuit techniques to support high-voltage switching in a low-voltage CMOS process. Due to the limited available capacitance, inductor coupling is required to meet the droop criterion. The analysis shows that a buck converter with un-coupled inductors requires excessive decoupling capacitance in order to meet the output droop requirement and is feasible only for small load currents. Energy recycling can improve efficiency by as much as 20%. We showed that an efficiency of 85% can be achieved for conversion ratios of up to 4:1 with thin-film inductors using an extra metal layer, either on the processor die (monolithic) or on a separate die (stacked) and that a 2:1 converter requires less than 5% area overhead. With a 4:1 converter, the input current is 3.4x smaller and the off-chip decoupling capacitance is 6.8x smaller when compared with the current power delivery without an integrated DC-DC converter.

6. REFERENCES

[1] Documents 241997, 242323, 290544, 243185, 290607, 243335, 243657, 244452, 245264, and 249657, www.intel.com

[2] F. C. Lee, X. Zhou, *Power Management Issues for Future Generation Microprocessors*, ISPSD 1999, pp. 27-33, 1999.

[3] A. Waizman, C.-Y. Chung, *Resonant Free Power Network Design Using Extended Adaptive Voltage Positioning (EAVP) methodology*, IEEE Trans. Adv. Packaging, pp. 236-244, Aug 2001.

[4] N. X. Sun, et al, *Design of a 4 MHz, 5V To 1V Monolithic Voltage Regulator Chip*, ISPSD 2002, pp. 217-220, 2002.

[5] Y. Ren, et al, *Two-Stage 48V Power Pod Exploration for 64-Bit Microprocessor*, APEC 2003, pp. 426-431, 2003.

[6] A. M. Crawford, D. Gardner, S. X. Wang, *High-Frequency Microinductors With Amorphous Magnetic Ground Planes*, IEEE Trans. Magn., pp. 3168-3170, 2002.

[7] K. H. Kim, et al, *A Megahertz Switching DC/DC Converter Using FeBN Thin Film Inductor*, IEEE Trans. Magn., pp. 3162-3164, 2002.

[8] I. G. Park, S. I. Kim, *Modeling and Analysis of Multi-Interphase Transformers for Connecting Power Converters In Parallel*, 28th Annual IEEE Power Elec. Spec. Conf., pp. 1164-1170, 1997.

[9] Intel 865G/865GV/865PE/865P Chipset Platform Design Guide, ww.intel.com.