

Design Considerations of Monolithically Integrated Voltage Regulators for Multicore Processors

Hesam Fathi Moghadam, Jieun Jang, Michael Dayringer, Thipok Rak-amnonykit, Guanghua Shu, Anatoly Yakovlev, Yue Zhang, David Hopkins
Systems Research Group (SRG), Oracle Labs
Redwood Shores, USA
Email: hesam.fathi.moghadam@oracle.com

Abstract—Presented in this paper are design considerations for a Monolithically Integrated Voltage Regulator (MIVR) targeting a 42mm² multicore processor test chip taped-out in TSMC 28nm process. This is the first work discussing the utilization of on-die magnetic core inductors to support >50A of load current. 64 inductors with switching frequency of 140MHz are strategically grouped into 8 interleaving phases to achieve 85% efficiency and minimize on-die voltage drop.

Keywords—voltage regulation; on-die; magnetic; inductor; IVR

I. INTRODUCTION

The off-chip nature of voltage regulator modules (VRMs) gives rise to many undesirable effects that decrease the total efficiency of power conversion and delivery. Located on the motherboard at a distance from the chips they power, VRMs are often found to be too slow, too coarse, and too inefficient in many recent applications.

A. Background and Motivation

In contrast to VRMs, integrated voltage regulators (IVRs) are located within the same package as their loads. Taking advantage of their proximity to the load, they overcome several limitations of VRMs. Reduction in Ldi/dt , as a result of higher input voltage to the IVR, and higher control loop bandwidth, decreases droop and overshoot of the output voltage. Faster transient response also lowers dynamic voltage scaling (DVS) turn-around time, allowing the voltage regulators to effectively follow the workload and further improve the CPU's efficiency [1][2]. Resistive loss along the power path between regulator and load – including board, package, and guard band circuits – is reduced as the regulator is moved into the package. Lastly, IVRs improve the overall efficiency of power delivery by enabling multiple voltage domain division of the chip's multiple cores, even in the case where board or package resources are limited.

Three types of IVRs are commonly used in practice: buck converter (BC) uses inductors, switched capacitor (SC) uses capacitors, and low drop-out (LDO) regulator which does not use large passives. For high power levels, capacitors in the SC approach consume a large portion of the die area, and offering continuous conversion ratios is challenging. The LDO approach is inefficient at large voltage conversion ratios. The BC approach, with magnetic core on-die inductors, delivers higher power density than the SC approach (for monolithic

process without deep trench isolation), making it more suitable for high power multicore processor applications.

Integrating the inductor into the chip package is crucial to the performance of the BC IVR. Researchers have demonstrated that IVRs with satisfactory efficiency can be implemented with either magnetic core or air core inductors [3][4]. On-package inductors [5] pose a number of drawbacks, such as reduced efficiency due to parasitics between driver output and inductor, as well as between inductor output and load. Moreover, valuable bump resources are consumed getting to the off-chip inductor and back from it to power the on-die load. Lastly, package traces used to create the inductor layers consume package area that could be used to ease I/O routing. Recent 3D integrated on-die magnetic core inductor technologies are capable of reaching more than 30x increase in inductance density compared to 2D trace inductors [6]. Fig. 1a shows the top and side view of a magnetic core inductor and Fig. 1b shows the side view of a chip with silicon, transistors, metal stack-up, inductors implemented with ultra thick metal (UTM) and post passivation interconnect (PPI) layers around magnetic core MU1, and C4 bumps.

Realizing the potential of 3D magnetic core inductors as a key enabling technology for MIVRs [7][8], this technology is adopted and the design of a test chip to demonstrate the feasibility of this technology with on-die load is undertaken and shown in Fig. 1c. The remainder of this paper is structured as follows: Section II discusses design considerations necessary to optimize MIVR performance. Section III presents a design example and Section IV provides some conclusions.

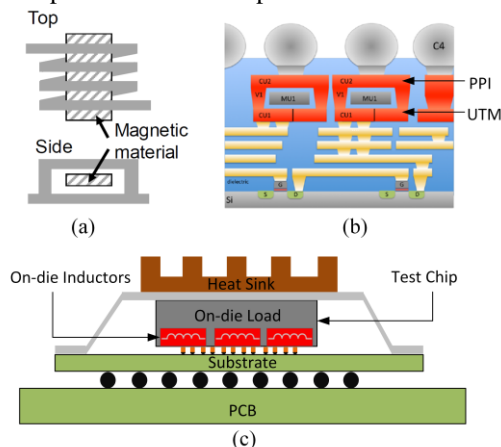


Fig. 1. (a) On-die embedded inductor with magnetic core top and side views. [6]. (b) Cross-sectional view of the die showing embedded inductors, PPI, and UTM [7]. (c) Test chip with IVR and on-die load in the test setup.

II. DESIGN CONSIDERATIONS

The design of an IVR needs to meet certain criteria for it to be integrated into a processor that is already at its thermal-electrical limits. Any decrease in power conversion efficiency can result in significantly more power usage for a large data center. In the case of an IVR for a high power processor, one should be even more cautious about the efficiency, especially at peak load, which is when the processor is at its thermal limit. By moving the board-level VRM on-die, the power density of the CPU increases by 10-20% due to IVR inefficiencies. This makes cooling more challenging.

The current density the IVR is capable of supporting is also important. Modern server-class processors easily consume >200W [9] and average core current density can be on the order of 1-2 A/mm² [10]. Additionally, output ripple is typically constrained to <1% of the supply to minimize noise and voltage margin overhead.

Due to the nature of varying workloads, the processor load requirements change rapidly in a short amount of time. This causes voltage droop and overshoot that have to be filtered by the IVR. Supply voltage droops increase the potential for timing violations. Overshoots increase the risk of transistor break down and reliability issues. To provide faster transient response, interleaving multi-phase regulators [11] are deployed, which will be discussed in the following sections.

BC drivers are typically very wide and consume a lot of valuable die area. Furthermore, this problem is exacerbated by the need of on-die decoupling capacitors around the gates of these devices in order to charge/discharge them effectively. These are all considered to be overhead compared to traditional VRMs, and must be justified when consuming expensive silicon area. Only designs with overhead <5% of total die area were considered.

A. Power Conversion Efficiency

Inductor and driver conduction losses are often dominant loss components due to large load current in high-power processors. Driver conduction losses become more significant if die area constraints limit the driver area.

Inductor DC conduction loss can be addressed by parallelizing multiple inductors at the cost of decreasing effective inductance. This increases total inductor current ripple and AC conduction loss. Inductor AC conduction loss can be reduced by increasing switching frequency, as it decreases inductor current ripple. However, inductor resistance grows rapidly at high frequencies and overwhelms the benefit of smaller current ripple. If the nominal operating duty cycle is around 0.5, then inductor AC conduction loss can be further improved by deploying coupled inductors due to decreased current ripple when opposite phases are driving the coupled inductor pair [7]. This comes at the cost of greater PDN loss to load devices under the inductor as inductor area increases (more discussion in the following section).

Adopting a stacked-switch topology can decrease driver conduction loss. This allows the driver stage to be made of

thin-oxide devices, which can decrease switch resistance by a factor of 5-10 compared to thick-oxide devices occupying the same area. For instance, in TSMC 28nm process, a 0.85V thin oxide device has ~5x less resistance than a 1.5V thick oxide device with the same channel width. Stacked topology does require additional power hungry mid-rail bias voltage generators. The latter can be alleviated by applying a charge-sharing technique, as described in [12], at the cost of local decoupling capacitor area.

Fig. 2a shows estimated power conversion efficiencies with different numbers of inductors. Each design point is optimized for best efficiency with 1.6V input, 1V output, and 50A load current under a switch and gate decoupling area limitation of 5% of total die area. Maximum efficiency saturates beyond the use of 80 inductors. Due to total chip area limitation, 64 inductors are chosen for the test chip presented in Section III without significant efficiency hit. Fig. 2b highlights the major sources of inefficiency, including the inductor AC and DC losses, switch (also known as driver) switching and conduction losses, as well as other losses, which includes controller and dead-time conduction losses. It should be noted that parasitic inductance coupling (PIC) is not considered in this analysis.

B. Inductor Placement and Floorplanning

Lateral power distribution on chip in the high power-density core area can be a difficult task. Many chip designs today focus on having just power and ground bumps in the chip core areas to alleviate routing issues. In this design, current for the core area is supplied from the on-die inductors, so the inductors are spread out evenly to ensure good lateral distribution.

One design tradeoff that must be dealt with early in the design process is choosing the size of the inductor. Small inductors carry large area overheads due to routing keep out zones surrounding the inductor. Large inductors stress the lateral routing and can have large dead-zones for power delivery beneath them due to the thick metal layers being used to construct the inductor. Additionally, inductors cannot have live electrical bumps above them, so choosing a size that works with the desired bump pitch is important for efficient use of metal resources. For this design, the largest inductor that fits within a two-by-two array of missing bumps was selected from a set of standard size inductors provided by the foundry.




Fig. 2. (a) Estimated power conversion efficiencies with different numbers of inductors. (b) Loss sources ($F_{sw}=140\text{MHz}$, $T=27^\circ\text{C}$).

Another design criteria that must be chosen is the orientation and phase assignment of each inductor. Placement of the output port of the inductor can affect output ripple due to the interaction between output currents from different inductors. Pairs of inductors were oriented opposite to each other so that their output ports were as close together as possible. By assigning each inductor in each pair opposite phase, optimal ripple cancellation at the output nodes is ensured. See Section II.C for more discussion on ripple reduction from output current sharing across the PDN.

A modern microprocessor depends on functionality like dynamic voltage and frequency scaling (DVFS) for power savings. To maximize efficiency across a range of output loads, this design supports inductor shedding. Fig. 5b shows the phase distribution for the 64 inductor test chip presented in Section III. All 8 phases, $\Phi 0$ - $\Phi 7$, are distributed across the chip. A traditional phase shedding approach would remove all inductors of each shed phase, e.g., dropping all $\Phi 4$ - $\Phi 7$ to go from full load to half load condition. The latter is not ideal for the presented distributed approach as removal of phases will exacerbate the ripple problems. Conversely, a better approach is keeping all eight phases but removing some inductors from each. For example, half load condition would ensure 4 inductors from each phase remain active such that dead spots are evenly distributed. In Fig. 5b, shaded regions represent active phases.

Floorplanning for the locations of the power switches and control logic is vital as they directly impact overall efficiency and block placement of core logic. Consolidation of all the control logic and drivers is preferable for reducing the control skew to the switches. Additionally, design of the power macro block is simplified by placing the logic into a single, contiguous piece. Consolidation, however, has two major drawbacks: the increased power losses due to routing from the switches to the inductors and the increased sensing errors and delays in the feedback to the control logic. Both of these can significantly reduce the efficiency of the overall power delivery system, which necessitates the distribution of switches and control logic. Consolidating control and carefully managing the distribution of control signals while distributing the drivers to be as close to the inductors as possible balances these issues. Fig. 3 shows the top view of the test chip design. The drivers are placed centrally between each inductor pair.

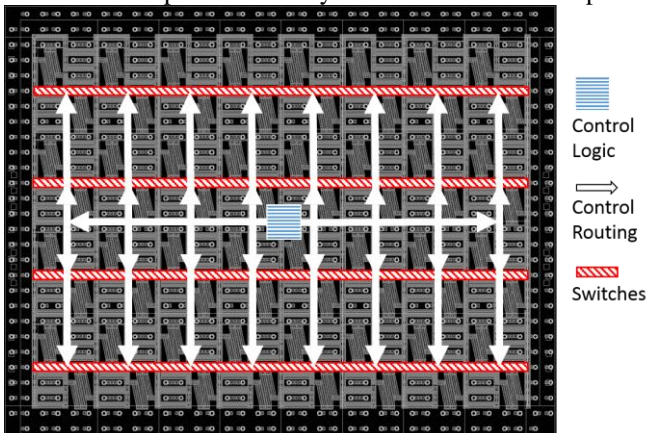


Fig. 3. Test chip floorplan.

Control logic is centralized for the whole chip and control signals are distributed to the center of each column of eight inductors and then further distributed within each column to the drivers.

C. PDN Co-optimization

It is essential to optimize the PDN to minimize the worst case IR drop and voltage ripple. PDN refers to the metal routing for explicit distribution of power supply from output of an inductor to load. Typically, multicore processors rely heavily on thick top metal layers to distribute power and ground uniformly across the chip. The lower metal layers are significantly more resistive, and are primarily used for local current distribution and signal routing. On-die inductors for IVRs utilize the same top-level metal resources as the PDN.

While inductors can vary in size, with some larger inductors having more desirable properties, the total load current density and available metal resources for power distribution under the inductor restrict their areas. Because the effective sheet resistance of metal layers below an inductor can be more than 10x that of thick metal layers, voltage drop can be an order of magnitude higher compared to the drop across thick layer distribution over the same area. With uniform load current distribution, the worst case IR drop would be to the load located under the center point of the inductor. Total load current under the inductor increases linearly with inductor area, proportionally increasing the worst case IR drop [13]. Therefore, doubling the length and width of an inductor approximately quadruples the IR drop to its center.

Additional penalty from higher IR drop arises from overvoltage losses resulting from the voltage margin that is provisioned at the worst-case location on the chip. The higher the regulated voltage necessary to compensate for the worst case IR drop, the higher the power consumption everywhere else on the chip where voltage exceeds the minimum required.

Finally, power distribution is equally important for supply and ground. A reason for considering prioritizing supply distribution over ground distribution is to minimize the connection impedance among different BC phases to ensure high quality current sharing, thereby minimizing voltage ripple, allowing for a smaller voltage margin and lowering the associated overvoltage losses. Therefore, analysis of metal resource allocation between supply and ground and its effects on supply ripple and power loss is often necessary in order to determine appropriate allocation of routing resources.

In order to rapidly verify and optimize phase distribution, inductor geometry choice, and PDN routing, a PDN modeling tool was implemented. Fig. 4 shows normalized transient simulation contour plots of IR drop and ripple across 16 inductors (one quadrant of test chip).




Fig. 4. Contour plot of load voltage across 16 inductors.

The plot is normalized to set a minimum voltage of 1V. Maxima occur at the output of every inductor with approximate overvoltage of 3.5% for this test case.

D. Parasitic Inductive Coupling (PIC) Effects

An on-die inductor, whether magnetic or air core, will have to be simulated within its intended environment. In particular, on-die and on-package metal surrounding the inductor can have an impact on the inductance and AC resistance of the inductor at the converter switching frequency. PIC effects are exacerbated by any physical conductive loops that are co-centered with the inductor core. It can be mitigated by increasing resistance of the surrounding metal or by breaking the loops. PIC mitigation techniques will need to be included in future designs.

E. Control Loop for Distributed Inductors

Besides efficiency considerations described in Section II.B, distributing many inductors over a large area introduces new challenges on designing a control loop. First, a large number of inductors require generating a large number of clock phases as well as matching their duty cycles. [14] It is not practical to have the same number of control phases as the number of inductors. Instead, the inductors can be grouped into a smaller number of control phases. For instance, 64 inductors are grouped into 8 phases (i.e., 8 inductors per individual phase). However, grouping does not ease the duty cycle matching requirement as phase mismatches between each inductor still need to be calibrated. Averaging out mismatches within each inductor group can relax the duty cycle matching requirement at the cost of performance.

Second, a large number of parallel inductors distributed over a large area can degrade phase margin significantly. Parallelizing inductors decreases the effective inductance, and moves the resonance frequency with output capacitance to a higher frequency close to the 0-dB bandwidth. Moreover, sensing points can be as far as die width or height, which gives a long loop delay from a sensing point to a controller and back to a driver, which can further degrade phase margin.

These effects need to be considered early in the design phase. The control loop design should include the loop latency from the sense point to the driver to correctly estimate loop stability. To improve loop stability, loop bandwidth may be trimmed at the cost of transient performance. Alternatively, the number of inductors may be reduced, sacrificing power conversion efficiency.

III. DESIGN EXAMPLE

The high level goal of the test chip, designed and taped-out in TSMC 28nm technology, was to demonstrate an MIVR for area and power levels representative of an Oracle SPARC core cluster. Active die area (not including the I/O ring) was 32mm² and total inductor area was 9.6mm². Fig. 5a shows the simplified block diagram of the test chip. There are 8 parallel inductors per phase and 8 independent phases with the output of all the phases connected directly to the on-die PDN. The four voltage sense points are averaged and fed into the compensator, followed by a multi-phase pulse width

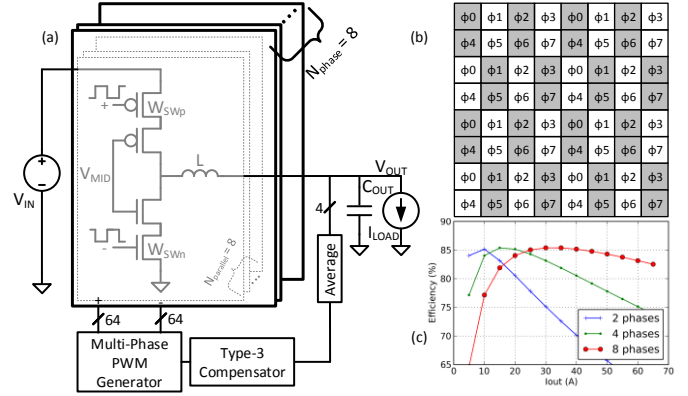


Fig. 5. (a) Test chip block diagram. (b) Phase assignments of the 64 inductors on test chip. (c) Efficiency versus number of phases.

modulation (PWM) generator. The latter produces the top side and bottom side level shifted driver gate voltages. Fig. 5c shows efficiency versus current for different number of phases.

Several important on-die characterization structures were employed to verify the functionality and study the potential influence of integrated magnetics on nearby circuitry. Utility analog circuitry, such as bandgap, bias circuits, and digital logic units were deployed. The test chip contained two phase-locked loops (PLLs), one with an off-chip supply to drive the clock for the MIVR and one powered by the MIVR to compare performance parameters such as jitter and power supply noise rejection.

This test chip was designed to show feasibility of the on-die inductor technology. The next generation of inductor technology shows improved DC/AC resistance and advanced nodes such as 7nm will also improve efficiency as it will be easier to hit area constraints. Furthermore, employing a custom inductor (rather than using one from a table of available inductors like in this test chip) would further improve efficiency and place this technology in line with product-scale efficiency requirements.

IV. CONCLUSION

As power requirements of processors continue to stress the PDN at the board, package, and chip levels, MIVRs show promise as on-die inductor technologies mature. Phase distribution and floorplan have to be optimized together with the traditional IVR parameters in order to get an overall efficient MIVR design. A test chip bearing in mind the design considerations presented in this paper was taped-out showing MIVR feasibility for high power multicore processors.

ACKNOWLEDGMENT

We would like to acknowledge the Oracle Microelectronics silicon design team, CAD team, and mask designers. In particular, Vijay Srinivasan, Junhua Gu, Yuanyuan Gong, Jesse Hsu, Ha Pham, Steven Butler, and Luke Shin. We would also like to acknowledge the support, insight and feedback from Ferric Inc., TSMC Taiwan and TSMC Austin.

REFERENCES

- [1] N. Sturcken *et al.*, "A Switched-Inductor Integrated Voltage Regulator With Nonlinear Feedback and Network-on-Chip Load in 45 nm SOI," in *IEEE Journal of Solid-State Circuits*, vol. 47, no. 8, pp. 1935-1945, Aug. 2012.
- [2] Wonyoung Kim, M. S. Gupta, G. Y. Wei and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," *2008 IEEE 14th International Symposium on High Performance Computer Architecture*, Salt Lake City, UT, 2008, pp. 123-134.
- [3] F. Paillet, G. Schrom and J. Hahn, "A 60MHz 50W Fine-Grain Package-Integrated VR Powering a CPU from 3.3V," in *Advanced Power Electronics Conference*, Palm Springs, CA, 2010.
- [4] J. T. Dibene, et al., "A 400 Amp fully integrated silicon voltage regulator with in-die magnetically coupled embedded inductors," in *Advanced Power Electronics Conference*, Palm Springs, CA, 2010.
- [5] E. A. Burton *et al.*, "FIVR — Fully integrated voltage regulators on 4th generation Intel® Core™ SoCs," *2014 IEEE Applied Power Electronics Conference and Exposition - APEC 2014*, Fort Worth, TX, 2014, pp. 432-439.
- [6] S. Mueller *et al.*, "Design of High Efficiency Integrated Voltage Regulators with Embedded Magnetic Core Inductors," *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, Las Vegas, NV, 2016, pp. 566-573.
- [7] N. Sturcken *et al.*, "Magnetic thin-film inductors for monolithic integration with CMOS," *2015 IEEE International Electron Devices Meeting (IEDM)*, Washington, DC, 2015, pp. 11.4.1-11.4.4.
- [8] N. Sturcken, "DC-DC Power Conversion with CMOS Integrated Thin-Film Inductors," *Proc. 5th International Workshop on Power Supply On Chip (PwrSoC2016)*, Madrid, Spain, October 3-5, 2016.
- [9] A. Sodani *et al.*, "Knights Landing: Second-Generation Intel Xeon Phi Product," in *IEEE Micro*, vol. 36, no. 2, pp. 34-46, Mar.-Apr. 2016.
- [10] B. Bowhill *et al.*, "4.5 The Xeon® processor E5-2600 v3: A 22nm 18-core product family," *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, San Francisco, CA, 2015, pp. 1-3.
- [11] Y. Panov and M. M. Jovanovic, "Design considerations for 12-V/1.5-V, 50-A voltage regulator modules," in *IEEE Transactions on Power Electronics*, vol. 16, no. 6, pp. 776-783, Nov 2001.
- [12] S.-J. Kim, R. K. Nandwana, Q. Khan, R. Pilawa-Podgurski, and P. K. Hanumolu, "A 4-phase 30–70 MHz switching frequency buck converter using a time-based compensator," *IEEE J. Solid-State Circuits*, pp. 2814-2824, Dec. 2015.
- [13] H. Chen, C. K. Cheng, A. B. Kahng, Q. Wang and M. Mori, "Optimal planning for mesh-based power distribution," *ASP-DAC 2004: Asia and South Pacific Design Automation Conference 2004 (IEEE Cat. No.04EX753)*, 2004, pp. 444-449.
- [14] A. Peterchev, J. Xiao, and S. Sanders, "Architecture and IC implementation of a digital VRM controller," *IEEE Trans. Power Electron*, vol. 18, no. 1, pp. 356–364, Jan. 2003.