

A 2.5D Integrated Voltage Regulator Using Coupled-Magnetic-Core Inductors on Silicon Interposer

Noah Sturcken, *Student Member, IEEE*, Eugene J. O'Sullivan, Naigang Wang, *Student Member, IEEE*, Philipp Herget, Bucknell C. Webb, Lubomyr T. Romankiw, Michele Petracca, Ryan Davies, Robert E. Fontana, Jr., Gary M. Decad, Ioannis JohnKymissis, Angel V. Peterchev, *Member, IEEE*, Luca P. Carloni, *Senior Member, IEEE*, William J. Gallagher, *Fellow, IEEE*, and Kenneth L. Shepard, *Fellow, IEEE*

Abstract—An integrated voltage regulator (IVR) is presented that uses custom fabricated thin-film magnetic power inductors. The inductors are fabricated on a silicon interposer and integrated with a multi-phase buck converter IC by 2.5D chip stacking. Several inductor design variations have been fabricated and tested. The best performance has been achieved with a set of eight coupled inductors that each occupies 0.245 mm^2 and provides 12.5 nH with $270 \text{ m}\Omega$ DC. With early inductor prototypes, the IVR efficiency for a $1.8 \text{ V}:1.0 \text{ V}$ conversion ratio peaks at 71% with FEOL current density of $10.8 \text{ A}/\text{mm}^2$ and inductor current density of $1.53 \text{ A}/\text{mm}^2$. At maximum load current, 69% conversion efficiency and $1.8 \text{ V}:1.2 \text{ V}$ conversion ratio the FEOL current density reaches $22.6 \text{ A}/\text{mm}^2$ and inductor current density reaches $3.21 \text{ A}/\text{mm}^2$.

Index Terms—Buck converter, DC-DC power conversion, Integrated Voltage Regulator (IVR), switched inductor, 2.5D integration.

I. INTRODUCTION

THE integration of efficient power converters will enable improved performance-power-watt across the full spectrum of digital computing devices, from high-end servers to smart phones. Increasingly, the processors employed in these computing platforms take the form of multicore architectures where workload is shared or parallelized across multiple independent computing cores on the same chip. Energy efficient operation of these cores requires highly granular, active power

Manuscript received May 01, 2012; revised June 11, 2012; accepted August 31, 2012. Date of publication December 04, 2012; date of current version December 31, 2012. This paper was approved by Guest Editor Uming Ko. This work was supported in part by the U.S. Department of Energy (DE-EE0002892), the National Science Foundation (CCF-1018236 and EECS-0903466), the DARPA LEAP Program, and the SRC Focus Center Research Program.

N. Sturcken, R. Davies, I. Kymissis, L. P. Carloni, and K. L. Shepard are with Columbia University, New York, NY 10027 USA (e-mail: nsturcken@ee.columbia.edu; shepard@ee.columbia.edu).

M. Petracca was with Columbia University, New York, NY 10027 USA, and is now with Cadence Design Systems, Inc., San Jose, CA 95134 USA.

E. J. O'Sullivan, N. Wang, B. C. Webb, L. T. Romankiw, and W. J. Gallagher are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: eosull@us.ibm.com; nwang@us.ibm.com).

P. Herget, R. E. Fontana, Jr., and G. M. Decad are with the IBM Almaden Research Center, San Jose, CA 95120 USA.

A. V. Peterchev is with Duke University, Durham, NC 27710 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSSC.2012.2221237




Fig. 1. Diagram of 2.5D integrated voltage regulator (IVR) chip stack. IC with buck converter and load circuitry flips onto interposer with power inductors, which wirebonds to a ball grid array substrate.

management, where each core (or even parts of a core) require dynamic supply voltage scaling to allow energy-delay trade-offs to be performed in the presence of workload variability [1]–[3]. At times, these active power management schemes will scale supply voltages to less than 0.5 V in order to achieve circuit operation at minimum power consumption. Voltage transitions will occur within nanoseconds, timescales that are relevant to high-performance digital logic, and voltage regulators will suppress large voltage droops that would otherwise cause the processor to skip operating cycles while the power supply recovers.

Current power delivery implementations employ off-chip, board-level voltage regulator (VR) integrated circuits (along with board level passives) to down-convert independent supply voltages, which are then distributed to a processor through board interconnect and IC packaging. Delivering many supply voltages in this manner requires many external VR modules, an implementation that rapidly becomes interconnect limited both in the distribution of a large number of supplies and in losses associated with moving currents over long interconnect distances. Likewise, the slow switching frequencies of board-level VRs extend supply voltage transition periods to microseconds, while the parasitic interconnect impedance between board VR and processor inhibits the regulator's ability to suppress voltage droop during load current transients. In short, current power delivery technologies based on discrete, board-level VRs are unable to support the type of active power management required of future computing devices.

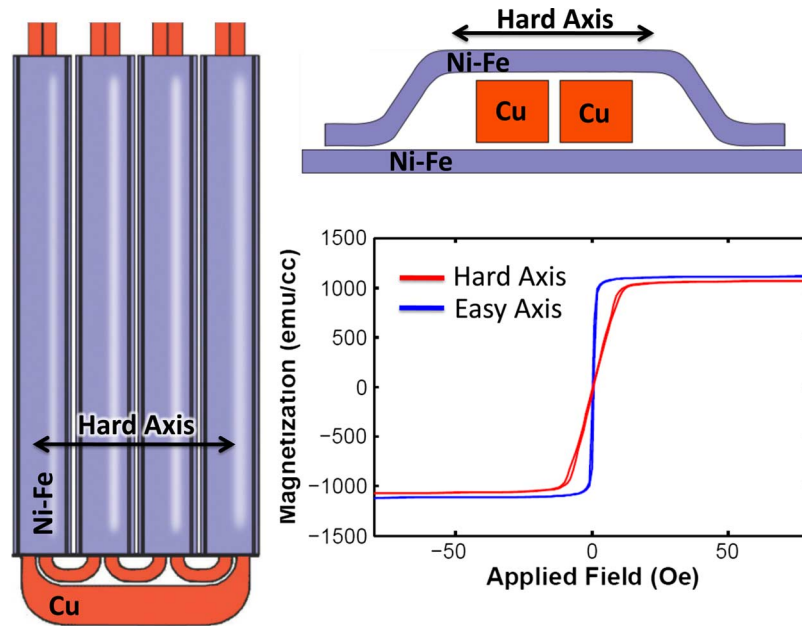


Fig. 2. Top view of four single-turn, coupled power inductors (left), cross-section of magnetic cores and windings (top right) and magnetization curves for the Ni-Fe core material (bottom right).

Switched inductor integrated voltage regulators (IVRs) will operate $\sim 100\times$ faster than board level VRs, enabling dynamic voltage and frequency scaling (DVFS) on nanosecond time scales [4], [5]. Furthermore, IVRs improve efficiency by enabling power delivery at higher voltages in the packaging, reducing I^2R losses in the power delivery network (PDN) and required supply voltage margins. The primary obstacle facing development of IVRs is integration of suitable power inductors that can carry high current levels within a constrained area. This work demonstrates how the challenge of inductor integration can be overcome with an early prototype switched-inductor IVR using thin-film magnetic power inductors that are integrated by 2.5D chip stacking as shown in Fig. 1 [6]. Section II of this paper describes the design and fabrication of the magnetic core inductors that enable the integration of high current density voltage regulators. Section III details the buck converter circuitry that has been designed and fabricated for an IVR test chip. Section IV describes the 2.5D chip stacking method that is used to integrate the magnetic power inductors with the IVR test chip and Section V presents experimental results for the IVR prototype.

II. MAGNETIC CORE INDUCTORS

Efficient power conversion in a switched-inductor IVR requires inductors that deliver both high current density and high effective inductor efficiency [7]. Planar spiral or other inductor topologies that can be constructed using the interconnects of a typical CMOS process are too resistive to provide efficient on-chip power conversion at reasonable current densities [8]. The efficient use of surface mount technology (SMT) air-core inductors, which can provide a current density up to $\sim 1.7 \text{ A/mm}^2$, has been successfully demonstrated [5], [9]–[11]. However, the size and discrete nature of these devices hinders the scalability of any IVR incorporating discrete SMT

inductors. Fortunately, advances have recently been made in the development of integrated magnetic-core power inductors that are highly scalable and capable of delivering current densities as high as 8 A/mm^2 [12]–[14]. These inductors have been included in IVR prototypes by on-chip integration [15] and chip stacking, as described here. The inductors utilized in this work have been described separately in [14], where several types of inductors were fabricated, both for individual electrical test and integration with a buck converter.

A. Inductor Design

The power inductor topology that has been chosen for this work is an elongated spiral inductor where two layers of high permeability magnetic material form a cladding around the copper conductor, as shown in Fig. 2, boosting the inductance of the device. Similar topologies have shown high inductance density and quality factor at relevant frequencies [12], [13]. The magnetic cladding is anisotropic. As a result, the hard axis of magnetization, which typically exhibits a more linear relationship between the applied magnetic field and magnetization, is designed to take the same orientation as the induced magnetic field from the elongated dimension of the inductor.

Fig. 2 shows four inductors, where each inductor is coupled with those on either side of it through the magnetic cladding. The outside inductor wraps around so that all of the inductors are symmetrically coupled with their neighbors. In order to achieve inverse coupling, the inductors are driven by the buck converter such that the DC currents through the windings within a magnetic cladding travel in opposite directions. This inverse coupling helps to avoid magnetic saturation in the cladding, improving maximum achievable current densities [16]. This is possible in the case of a multi-phase buck converter because the DC current through each of the inductors is balanced such that the DC magnetic field from adjacent windings is equal and




Fig. 3. SEM cross section of magnetic power inductors.

opposite, effectively canceling the DC field in the magnetic cladding. This inverse coupling also reduces inductor current ripple and improves transient response [17]. The current density benefit that is obtained from inverse coupling for the case of a two-phase buck converter with our inductor structures is modeled and verified in [18]. Similarly, the improvement in effective inductor efficiency and current density obtained with inverse coupling in multi-phase buck converters is modeled in [7].

B. Inductor Fabrication

The inductors are fabricated on 200 mm silicon wafers in the Microelectronics Research Laboratory (MRL) at the T. J. Watson Research Center. The bottom and top inductor yokes are electroplated galvanostatically in a paddle cell. Physical vapor deposited (PVD) $\text{Ni}_{80}\text{Fe}_{20}$ films, 65 nm thick, are used as the electroplating seedlayers; a bias magnetic field is applied during seedlayer deposition to produce magnetic anisotropy. The magnetic yokes are plated through photoresist-defined molds (thru-mask plating method). This method gave smooth yoke edges, edge smoothness being important to avoid the nucleation of magnetic domains and pinning of domain walls. To ensure good yoke deposit thickness and composition uniformity, the field around the yokes is also plated up at the same time as the yokes, a thin resist frame separating both plated regions. During plating, a dc magnetic field is applied along the longest axis to define the magnetic anisotropy of the yokes. The presence of the field material ensures continuous magnetic flux across the whole 200 mm wafer, which is critical for obtaining good magnetic anisotropy.

After yoke plating, resist mask stripping, and plated field and seedlayer etching, a bilayer of PECVD SiN_x and TEOS dielectric ($\sim 1 \mu\text{m}$ in total thickness) is used to encapsulate the yokes. After bottom yoke fabrication, the magnetic vias, where the top and bottom yokes contact, are opened by reactive ion etching (RIE). Following plating seedlayer deposition, copper coils are electroplated through resist masks to a thickness of about $5 \mu\text{m}$.




Fig. 4. Pictures of magnetic core inductors fabricated for individual electrical test.

After resist mask and seedlayer removal, $6\text{-}\mu\text{m}$ -thick photoresist (AZ Electronics P4620) is used to encapsulate the coils. After patterning, the photoresist is briefly reflowed at 120°C to give sloped sidewalls, ensuring that the top yokes gradually extended to the magnetic via, avoiding the formation of any abrupt angle, which could saturate or pin domain walls. Finally, the photoresist is hard-baked at 200°C for 2 hours to form a rigid encapsulant. The hard-baked photoresist structures exhibit smooth and partially planar surfaces in advance of top yoke plating.

After top yoke fabrication and encapsulation by a bilayer of PECVD SiN_x and TEOS, inductor fabrication is concluded by opening the electrical contacts (Cu pads) using RIE. Fig. 3 shows cross-sections of a single-turn inductor and a magnetic via, while Fig. 4 shows several non-coupled inductors that are fabricated specifically for electrical test.

C. Magnetic Material Properties

The resistivity of the electroplated $\text{Ni}_{45}\text{Fe}_{55}$ is measured to be about $45 \mu\Omega \cdot \text{cm}$ using the four-point probe method, which is double the resistivity of Permalloy ($\text{Ni}_{81}\text{Fe}_{19}$). This higher resistivity helps to reduce eddy current at high frequency. Fig. 2 shows the magnetic hysteresis loops, which are obtained using a vibrating sample magnetometer (VSM), of a plated

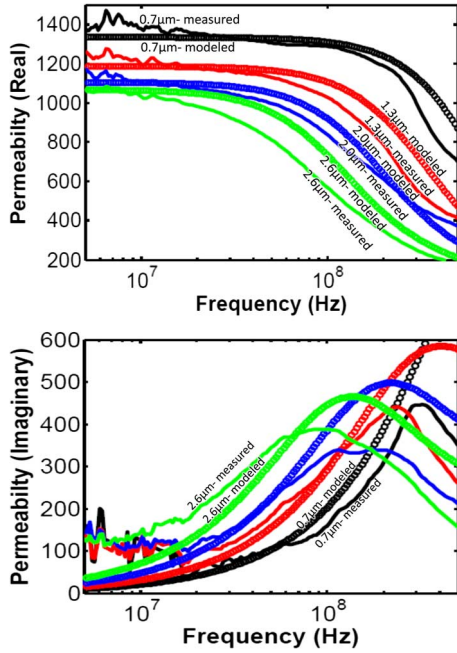


Fig. 5. Complex permeability spectra of $\text{Ni}_{45}\text{Fe}_{55}$ films with different thicknesses. For comparison, theoretical permeability spectra are also shown in the circle lines.

$\text{Ni}_{45}\text{Fe}_{55}$ film with a thickness of $2.0 \mu\text{m}$. The film shows clear anisotropy with a low coercivity of 0.2 Oe along both easy and hard axes. Saturation magnetization and anisotropy fields are 1.5 T and 13 Oe , respectively. Complex permeability spectra are obtained by measuring the impedance of a single-stripe loop fixture loaded with magnetic films. Fig. 5 shows the real and complex permeability spectra of plated $\text{Ni}_{45}\text{Fe}_{55}$ films with different thicknesses. For comparison, theoretical permeability spectra are also shown in the circle lines. The low frequency permeability reaches a value of 1300 for the $0.7 \mu\text{m}$ films. As the thickness of the film increased, the value decreases to about 1000 for the $2.6 \mu\text{m}$ film due to the shape anisotropy induced in the thicker films. In addition, due to eddy current and skin effects, the roll-off frequency decreased from 200 MHz to 50 MHz as the thickness increases.

III. INTEGRATED BUCK CONVERTER

An eight-phase buck converter has been designed to drive the magnetic core power inductors. The controller is designed to accommodate any number of inductor phases up to eight, with variations of inductance values and coupling strengths. The buck converter IC is designed and fabricated in IBM's 45 nm SOI technology. The IC is shown in Fig. 6 and includes the buck converter, a network-on-chip that acts as a realistic on-chip load, an artificial load used for characterizing the buck converter, and input and output decoupling capacitance. Within the buck converter, the control circuitry occupies 0.178 mm^2 , while the bridge FETs occupy 0.1 mm^2 . The bridge FETs are thick-oxide devices, capable of withstanding a maximum V_{ds} of 1.8 V . A total of 48 nF of deep-trench (DT) and thick oxide MOS capacitance decouples V_{OUT} and occupies 0.40 mm^2 , while 21 nF of DT occupying 0.52 mm^2 decouples the 1.8 V input supply to




Fig. 6. Photograph of buck converter IC.

compensate for the large PDN impedance (due primarily to the wirebond connections in this prototype).

A. Feedback Controller

Fig. 10 shows a system level diagram of the IVR. The buck converter control circuitry resides on the IC and is composed of two control loops, a slow voltage-mode outer loop that provides low-frequency regulation and a fast inner loop that responds to high-frequency load transients. The digital pulse-width modulator (DPWM) receives an eight-bit voltage identifier code (VID), from which it derives up to eight pulse-width modulation (PWM) signals with programmable switching frequency, f_s , and phase relationships. The resolution of the DPWM is limited to the 250-ps period of a high-frequency reference clock provided by an on-chip PLL. The DPWM also generates an analog reference voltage, V_{REF} , from a clean 1.8 V for the outer feedback loop. The compensator for the outer feedback loop is a low-pass filter with programmable pole frequency, typically chosen 10 to 16 times lower (depending on inductance value) than the effective switching frequency Nf_s , where N is the number of phases in operation. The outer feedback voltage, $V_{FB,O}$, drives a delay line that modulates the DPWM output to create the reference PWM signal, V_{PWM} , which subsequently drives the fast non-linear inner control block.

The fast inner loop is shown in Fig. 7. Signal V_{PWM} drives an RC filter to generate the inner reference voltage, $V_{REF,I}$, while the bridge output voltage for each phase, V_{BRIDGE} , drives another RC filter to generate the inner feedback voltage, $V_{FB,I}$. The pole in both RC low-pass filters is chosen to be below f_s so that the steady state amplitude of $V_{REF,I}$ and $V_{FB,I}$ is around 150 mV , which gives a small signal feedback gain of $\sim 30 \text{ V/V}$ and ensures stable loop dynamics. In steady state, $V_{FB,I}$ will slew behind $V_{REF,I}$ and the resultant evaluation of the comparator causes V_{BRIDGE} to closely track V_{PWM} . In the event of a large load current transient, the error in the output voltage, V_{OUT} , will couple across C_{FB} onto $V_{FB,I}$ and the comparator will react immediately to reduce overshoot in V_{OUT} . This fast non-linear response can reduce the required decoupling capacitance on the output voltage $V_{FB,I}$.

B. Integrated Network-on-Chip Load

Also residing on the IC is a 64 -tile network-on-chip (NoC) consisting of four parallel, heterogeneous, physical network

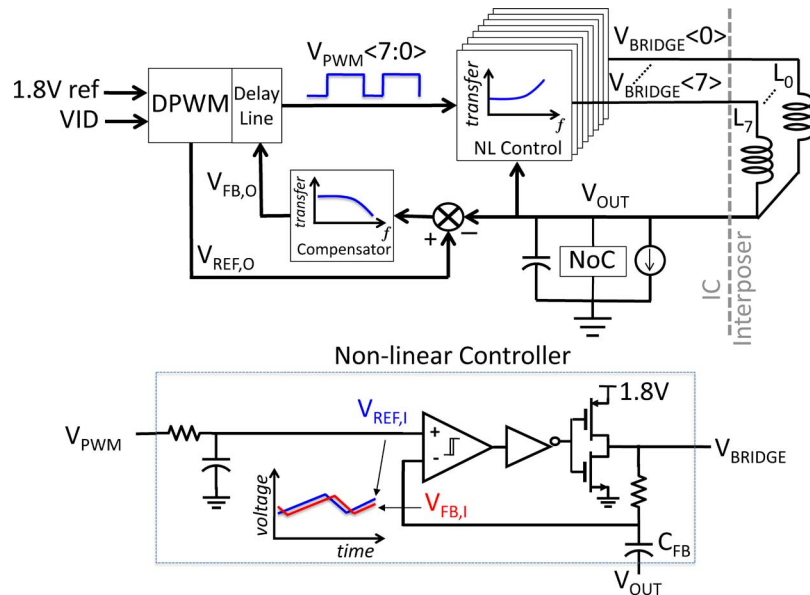


Fig. 7. Complete IVR system overview (top) and fast non-linear control loop (bottom).

planes with independent frequency domains. The NoC provides a highly scalable platform for exploring granular power distributions given the ease with which traffic patterns can be used to modulate load currents and transients. NoCs are becoming the basic interconnect infrastructure for complex SoCs. Since communication plays a key role in SoCs and given the very strict energy and performance requirements imposed on NoCs, recent designs have reserved a separate voltage-clock domain for the NoC alone [2]. The NoC provides realistic load behavior and supports experimentation on supply noise and DVFS. In addition, an artificial load on the IC is capable of generating large current transients with ~ 0.02 A/ps slew for characterization of the feedback controller.

IV. 2.5D CHIP STACKING

Fig. 8 shows a photograph of the complete 2.5D chip stack. The buck converter IC is flip-chip attached to the silicon interposer, which holds the custom fabricated coupled power inductors while breaking out signals and the 1.8 V input power supply to wirebond pads on the perimeter of the interposer. These signal and power nets are wirebonded to a generic BGA laminate, which is subsequently placed in a socket for electrical test. Once the buck converter IC has been attached to the interposer, the bridge FETs on the IC are able to drive current from the 1.8 V input supply through the inductors on the interposer. This current will pass through the inductors and then back into the IC through C4 bumps where it is then distributed to the load across the on-chip power distribution network. Four variations of power inductor have been fabricated on the silicon interposer as shown in Fig. 9: four uncoupled two-turn inductors (type 1), eight single-turn coupled inductors (type 2), eight two-turn coupled inductors (type 3) and two sets of four single-turn coupled inductors (type 4). The C4 footprint of the prototype IC is designed to leave a total of 3.2 mm^2 in the center of the interposer for the inductors, although most inductor variations use less than the available area.

The power inductors are not integrated in the front-end-of-line (FEOL) of the CMOS technology and so the area consumed




Fig. 8. Photograph of packaged 2.5D IVR chip stack.

by these devices comes at substantially reduced cost. For this reason, the primary constraint on the inductor area in a chip stacking integration scheme is set by the scalability of the IVR solution, rather than cost of area consumed by the power inductors. The maximum current density of a candidate inductor topology must match, or exceed the current density of the load. This will allow the inductor to reside within the perimeter of the load, and in the case of a multi-core architecture, would provide perfect scalability, where multiple cores can be stamped across the load IC, with their corresponding set of inductors stamped across the interposer in the same way. In a worst case, modern high performance digital logic can consume current at levels as high as 2 A/mm^2 , current density levels that are exceeded by the power inductors employed here.

A significant downside to the 2.5D chip stacking method is the large impedance of the power delivery network (PDN). The combined impedance in the PDN from the socket, package, wirebonds, and interposer traces is $70 \text{ m}\Omega$ at DC, and increases with frequency due to the inductance of the wirebonds and other traces. The resistive losses from the PDN are a major source of inefficiency for the system, and the high frequency impedance severely impairs the ability of the voltage regulator to suppress




Fig. 9. Photographs of the four types of silicon interposer that were tested with the buck converter IC: Four uncoupled two-turn inductors (type 1—upper left), eight single-turn coupled inductors (type 2—upper right), eight two-turn coupled inductors (type 3—lower left) and two sets of four single-turn coupled inductors (type 4—lower right).




Fig. 10. Inductance (L), coupling coefficient (K) and resistance (R) of coupled single-turn inductors.

voltage droop during load current transients. For this reason, a fully 3D integration approach that incorporates thru-silicon-vias (TSVs) in the interposer would be favorable for high current applications, as it would result in substantially reduced PDN impedance.

V. EXPERIMENTAL RESULTS

A. Magnetic Core Inductors

The inductance, coupling coefficient and resistance of a single turn, $1200\ \mu\text{m}$ inductor with $2\ \mu\text{m}$ thick magnetic layers is shown in Fig. 10. The performance exhibited here is representative of the coupled, single-turn inductors that have been fabricated on

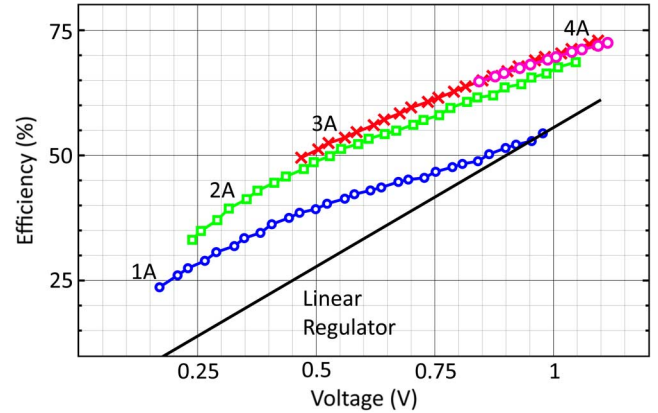


Fig. 11. Efficiency as a function of output voltage at various load currents for the IVR with Type 2 silicon interposer.




Fig. 12. Efficiency as a function of load current at various output voltages for an IVR with Type 2 silicon interposer.

the interposer. The DC inductance of $12.5\ \text{nH}$ is suitable for integrated power conversion, but eddy currents are induced in the magnetic core starting in the $10\text{--}100\ \text{MHz}$ range, which subsequently causes the inductance and eventually the coupling to fall off. At the same time we see that the eddy currents in the core, as well as the winding skin depth and proximity effect result in an increase in the winding resistance over the same switching frequency. This performance is consistent with the measurements of permeability shown in Fig. 5, and is expected to improve significantly with the addition of insulating laminations in the magnetic core, which will suppress eddy currents.

B. Integrated Voltage Regulator

The 2.5D IVR chip stack has been assembled and tested in order to verify functionality. In all DC measurements, the resistive losses from the PDN have been excluded and the input voltage has been compensated, such that the input voltage at the IC is truly $1.8\ \text{V}$. All measurements have been conducted with the silicon interposer carrying eight single-turn coupled inductors, unless otherwise noted, as this inductor configuration provides the best performance.

1) *Efficiency*: The efficiency as a function of output voltage and load current for the IVR is shown in Figs. 11 and 12. The efficiency peaks at 75% with output voltage of $1.2\ \text{V}$ and load current of $3.2\ \text{A}$. The peak efficiency at $1\ \text{V}$ is 71% when the load current is $3\ \text{A}$. The maximum load current that has been




Fig. 13. Efficiency as a function of switching frequency at various load currents for an IVR with Type 2 silicon interposer.

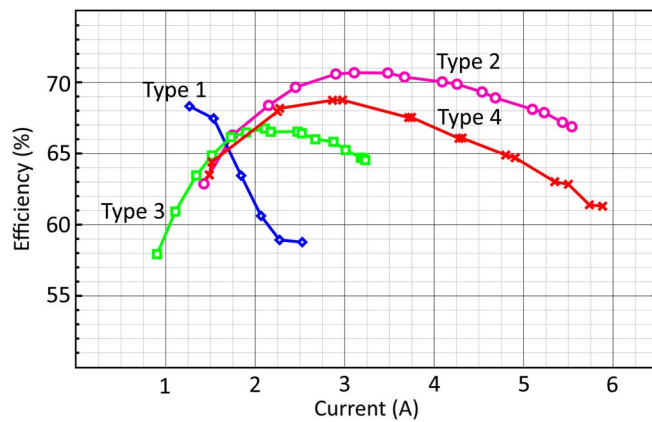


Fig. 14. Efficiency as a function of load current at 1 V output voltage for IVRs with four different kinds of power inductor.




Fig. 15. Loss breakdown and resistive breakdown for an IVR with Type 2 silicon interposer operating with 1 V output and 3 A load current.

measured for the IVR is 6.3 A, limited by the on-chip load. Fig. 13 shows the efficiency of IVR when operated at various switching frequencies and load currents. The optimal switching frequency for the IVR is in the range of 125 to 200 MHz. The efficiency as a function of load current for each of the four inductor variations is shown in Fig. 14, where the eight single-turn coupled inductors configuration (type 2) is clearly the most efficient. A breakdown of the IVR losses is shown in Fig. 15 for the case of a 1 V output voltage and 3 A load current. At this operating point, approximately 40% of all inefficiency is due to

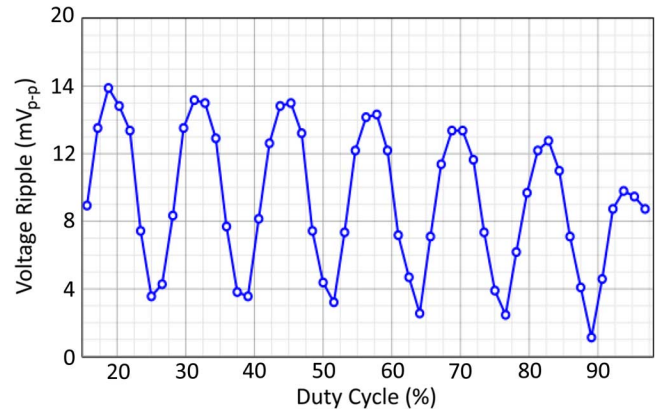


Fig. 16. Output voltage ripple of an IVR with Type 2 silicon interposer operating at 100 MHz switching frequency.

the high-frequency losses in the inductor, which can primarily be attributed to the formation of eddy currents in the magnetic core, as described in Section V-A.

Another major source of loss is the on-chip power distribution network resistance, as described in Section IV. This IVR prototype is designed to act as a flexible platform for testing various power inductor topologies. Therefore, in the case of the power distribution network, optimal design has to be compromised for flexibility. In this prototype, once the load current passes through the inductors it enters the on-chip power distribution network through C4 bumps near the buck converter. The load current then travels across the on-chip power distribution network (approximately 3 mm) to the artificial load, where the output voltage is measured. The on-chip power distribution network resistance of approximately 45 m Ω , is accountable for 25% of the converters losses. The remainder of the conversion loss is attributed to the DC resistance of the inductors and the switching and resistive losses of the bridge FETs.

The IVR that integrates eight single-turn coupled inductors (type 2) down converts with peak efficiency at a load current of 3 A, and achieves a maximum current of at least 6.3 A. The inductors occupy 1.96 mm². Current density for these devices at peak efficiency is 1.53 A/mm², and the peak current density is 3.21 A/mm². The FEOL area consumed by the buck converter, controller, bridge FETs and some input decoupling capacitance is 0.278 mm². At peak efficiency, therefore, the FEOL current density for this IVR is 10.8 A/mm², while the maximum current density is 22.7 A/mm². In this calculation we exclude the area of some input decoupling capacitance, as this capacitance would not be required in a fully 3D integration approach, where the PDN impedance would be lower.

2) *Voltage Ripple*: Fig. 16 shows the output voltage ripple from the IVR as a function of duty cycle when the buck converter is operating at a switching frequency of 100 MHz. The peak voltage ripple is 14 mV peak to peak, this occurs when $1/16 = \text{mod}(D, 1/8)$, where D is the duty cycle. The best case voltage ripple of 3 mV peak to peak occurs when $0 = \text{mod}(D, 1/8)$, when the inductor current ripple from each of the eight phases almost perfectly cancel. The IVR voltage ripple is expected to improve dramatically as insulating laminations are added to the magnetic yoke of the inductors, which will improve high frequency inductance.




Noah Sturcken (S'08) received the B.S. degree *summa cum laude* in electrical and computer engineering from Cornell University, Ithaca, NY, in 2008 and the M.S. degree in electrical engineering from Columbia University, New York, NY, in 2009. He is currently working toward the Ph.D. degree at Columbia University.

He has held several internships, most recently at AMD's Research and Development Laboratory. His research interests include integrated magnetic materials for power conversion, switched mode power supplies, control theory and mixed signal circuit design. Mr. Sturcken is a member of Eta Kappa Nu.



Eugene J. O'Sullivan is a Research Staff Member in the Physical Sciences Department. (He joined the Research Division in 1984. Dr. O'Sullivan received his Ph.D. in electrochemistry from University College, Cork, Ireland, in 1981. From 1980 to 1983 he was a Postdoctoral Fellow at Case Western Reserve University and he was a Senior research Associate at Case from 1983 to 1984.) At IBM, Dr. O'Sullivan's initial research at research included studying the fundamental and applied aspects of electroless metal deposition, e.g. selective deposition of diffusion barrier

alloys. Following a period of leading the fabrication of an integrated, variable reluctance magnetic minimotor, joined the MRAM group in IBM in the late 90s. In the MRAM group, Dr. O'Sullivan works on wet processing, including post-RIE cleans, ion beam etch patterning, and process integration. For the past couple of years, he has also been leading the fabrication effort for an integrated magnetic inductor for on-chip, DC-DC power conversion.



Naigang Wang (S'07) received the B.S. degree from Jilin University, Changchun, China, in 1999, the M.S. degree from the Virginia Polytechnic Institute and State University, Blacksburg, in 2005 and the Ph. D. Degree from University of Florida, Gainesville, in 2010, all in materials science and engineering.

He is currently a Postdoctoral researcher in IBM T. J. Watson Research Center, Yorktown, NY. His research focuses on the fabrication and integration of soft and hard magnetic materials for on-chip inductors, high-frequency power converters, as well as

MEMS sensors and actuators.



Philipp Herget received his M.S. and Ph.D. degrees in electrical and computer engineering from Carnegie Mellon University in 2004. From 2004 to 2005 he was a postdoctoral researcher working in magneto-optical storage systems with the Data Storage Systems Center at Carnegie Mellon University.

In 2006 he joined IBM at the Almaden Research Center in San Jose. His work has focused on modeling, testing, and designing magnetic heads for tape storage systems. For the last two years he has been working in magnetics for on chip power conversion

focused on the measurement and theory of thin film inductors for power conversion. He has authored and co-authored 20 technical papers has 22 issued and pending patents in the areas of data storage and magnetics.



Bucknell C. Webb is a Research Staff Member working in the System on Package (SOP) & 3D Integration Group at IBM's T.J Watson Research Center. He received a B.A. degree in physics from Harvard University and a Ph.D. degree in physics from Cornell University for studies on the far-infrared reflectivity of valence-fluctuating compounds, primarily CePd₃. He joined IBM in 1988, working initially with optical imaging and the high-frequency (100-MHz) magnetic response of magnetic thin films such as in disk drive recording heads. Since 2002, he has been studying stress, fracture and warp in 3D silicon structures. He is an author of more than 28 scientific papers and 53 U.S. patents.



Lubomyr T. Romankiw is an IBM Fellow at the IBM's T.J. Watson Research Center, Yorktown Heights, N. Y. 10598, where he has been affiliated since 1962. He received his B.Sc. in Chem. Eng. from Univ. of Alberta, Canada in 1955, his M.Sc. and Ph. D. in Metallurgy and Materials from MIT in 1962. He holds 68 patents, has over 130 published inventions, has published more than 150 scientific papers, has written six book chapters and edited nine volumes of ECS symposia proceedings in (1) Magnetic Materials, Processes and Devices for Storage and MEMS and in (2) Electrochemistry in Electronics.

His research has dealt with nearly all aspects of electroless plating, electroplating, etching, lithography and micro fabrication used in electronics. He has done pioneering work in combining electro-deposition, electro-etching, lithography, dielectrics and process integration to produce first batch fabricated inductive read/write and merged MR-read inductive write heads. He pioneered the underlying of science, and technology and invented new tools necessary to demonstrate the outstanding capability, and cost performance which electrochemical technology offers. He achieved understanding of the relation between the solution chemistry, nucleation, film growth, metallurgical structure, and magnetic properties of the binary and ternary Ni, Co, and Fe alloys, to produce 1.0 to 2.4 Tesla magnetic saturation films tailored to meet the needs of the magnetic read/write heads and transferred these technologies to manufacturing.

IBM has honored Dr. Romankiw with 14 Outstanding Invention and Contribution Awards and 31 IBM Invention Achievement Awards for his key inventions that have made possible manufacturing of thin film inductive heads, MR heads, merged MR-inductive heads; and his seminal contributions to electrochemical science and technology which led to the wide spread use of electrochemistry in electronics.

Dr. Romankiw received the ECS Vittoria de Nora Medal in 1994, and was named ECS Fellow. He is currently a honorary member of ECS. In 1986 he received the IEEE Morris N. Liebmann Memorial Award and was named an IEEE Fellow. In 1993 he was awarded the Societies of Chemical Industries Perkin Gold Medal. In May of 2012 he was inducted into USA Inventors Hall of Fame for his early patents on magnetic thin film heads.



Michele Petracca received his M.S. and Ph.D. degrees in Electronic Engineering from Politecnico di Torino, Italy. He was a visiting Ph.D. Student and then a Post-Doctorate Research Scientist in the Computer Science Department of Columbia University in the City of New York.

He is currently with Cadence Design Systems. His interests include networking, on-chip communication infrastructures, system-level design, hardware-software interaction and power-management in embedded systems.



Ryan Davies received the B.S. degree in materials science and engineering from the University of Florida (UF) in Gainesville, Florida in 2003, the M.S. degree from UF in 2007, and the Ph.D. degree from UF in 2009. His doctoral research focused on the incorporation of magnetic impurities in compound semiconductor materials for spintronic applications.

In 2010, he joined the Department of Electrical Engineering at Columbia University in New York as a Postdoctoral Research Scientist. His current research interests include magnetic material selection and optimization for power electronic applications, inductor device design, and inductor device fabrication process development. He has co-authored publications and patents involving the design, development, and fabrication of inductors for integrated voltage regulation. Dr. Davies is a member of the Materials Research Society and the American Vacuum Society.



Robert E. Fontana, Jr. is a Senior Research Scientist within the tape head function of the IBM Systems and Technology Group. His technical activities have concentrated on developing thin film processing techniques for fabricating magnetic device structures—TI (1975–1981) with magnetic bubbles, IBM (1981–2002) with thin film magnetic recording heads, Hitachi (2003–2007) with novel flux detecting sensors and nano structure e-beam fabrication, and now, again, (2008–Present) with IBM with tape heads. At IBM, his process group

fabricated the magnetic recording industry's first magneto resistive thin film heads. He subsequently transferred research processing strategies for three generations of magneto resistive thin film heads to manufacturing applications.

Dr. Fontana's current interests are advanced thin film tape head structures, solid state memory devices, nano processing, and technology roadmaps for storage class memories. He has authored 52 papers on magnetic devices and processes and has 102 issued patents in thin film magnetic structures. Dr. Fontana is a member of the National Academy of Engineering (NAE), a Fellow of the IEEE, a past president of the IEEE Magnetics Society, and a past recipient of the IEEE Cleo Brunetti Technical Field Award for excellence in the field of electronic miniaturization.

Gary M. Decad, photograph and biography not available at the time of publication.



Ioannis Kymissis is an electrical engineer teaching at Columbia University. His area of specialization is solid state electronics and device fabrication. His research focuses on thin film devices and systems, especially focusing on optoelectronic and sensing devices based on organic and recrystallized inorganic thin film materials. Current areas of research include investigations into device performance, system integration, fabrication, packaging, and device driving.

John graduated with his SB, M.Eng., and Ph.D. degrees from MIT. His M.Eng. thesis was performed as a co-op at the IBM TJ Watson Research Lab on organic thin film transistors, and his Ph.D. was obtained in the Microsystems Technology Lab at MIT working on field emission displays. After graduation he spent three years as a post-doc in MIT's Laboratory for Organic Optics and Electronics working on a variety of organic electronic devices and as a consulting engineer for QDVision. He joined the faculty at Columbia University in Electrical Engineering in 2006 as an assistant professor. John has won a number of awards for his work, including the NSF CAREER award, the IEEE EDS Paul Rappaport award, the Vodaphone Americas Foundation Wireless Innovation Award, the MIT Clean Energy Prize, and several other foundation, corporate, and best paper awards. He is the editor in chief of the Journal of the Society for Information Display, and is the technical program committee chair for the 2012 Device Research Conference.



Angel V. Peterchev (S'96–M'05) received the A.B. degree in physics and engineering sciences from Harvard University in 1999 and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Berkeley in 2002 and 2005, respectively. He completed postdoctoral training in transcranial magnetic stimulation at Columbia University.

He is presently Assistant Professor in the Departments of Psychiatry & Behavioral Sciences, Biomedical Engineering, and Electrical & Computer Engineering at Duke University. His primary research interests are in the design and modeling of devices and application paradigms for transcranial brain stimulation.

Dr. Peterchev has also published on digital control in power electronics and on architecture and control strategies for microprocessor voltage regulators. Dr. Peterchev has authored over 20 journal papers, 3 book chapters, and multiple conference papers and abstracts.



Luca P. Carloni (S'95–M'04–SM'09) received the Laurea degree (*summa cum laude*) in electrical engineering from the Università di Bologna, Bologna, Italy, in 1995, and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1997 and 2004, respectively.

He is currently an Associate Professor with the Department of Computer Science, Columbia University, New York, NY. He has authored over 90 publications and is the holds one patent. His current research interests include design methodologies and tools for integrated circuits and systems, distributed embedded systems design, and design of high-performance computer systems.

Dr. Carloni received the Faculty Early Career Development (CAREER) Award from the National Science Foundation in 2006, was selected as an Alfred P. Sloan Research Fellow in 2008, and received the ONR Young Investigator Award in 2010. He is the recipient of the 2002 Demetri Angelakos Memorial Achievement Award for "recognition of altruistic attitude toward fellow graduate students." In 2002, one of his papers was selected for the "Best of ICCAD," a collection of the best IEEE International Conference on Computer-Aided Design papers of the past 20 years. He is a Senior Member of the Association for Computing Machinery.



William J. Gallagher (M'83–SM'01–F'03) joined IBM Research in 1978 after receiving his B.S. in Physics *summa cum laude* from Creighton University and his Ph.D. in Physics from MIT. He worked for five years at IBM on scientific and engineering aspects of Josephson computer technology and then managed IBM's Exploratory Cryogenics Research Group for six years. In 1989 Dr. Gallagher participated in the formation of the IBM-AT&T-MIT-founded Consortium for Superconducting Electronics (CSE). He served as a

Director of the CSE until 1995. Since then, he has been leading efforts at IBM and with industrial partners to explore the use of magnetic tunnel junctions for nonvolatile magnetoresistive random access memory, MRAM. Currently he is senior manager of Exploratory Magnetic Memory and Quantum Computing at IBM's Thomas J. Watson Research Center.

Dr. Gallagher is a fellow of the IEEE and of the American Physical Society. He has served as Assistant to the Chairman of the APS Panel on Public Affairs, on the Executive Committee of the APS Forum on Physics and Society, on the Board of Directors of the Applied Superconductivity Corporation, and on numerous university and government lab review panels. He has over 190 technical publications and 20 U.S. patents.



Kenneth L. Shepard (M'91–SM'03–F'08) received the B.S.E. degree from Princeton University, Princeton, NJ, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively. From 1992 to 1997, he was a Research Staff Member and Manager with the VLSI Design Department, IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was responsible for the design methodology for IBM's G4 S/390 microprocessors.

Since 1997, he has been with Columbia University, New York, where he is now Professor of Electrical Engineering and Biomedical Engineering. He also was Chief Technology Officer of CadMOS Design Technology, San Jose, CA, until its acquisition by Cadence Design Systems in 2001. His current research in-

terests include power electronics, carbon-based devices and circuits, and CMOS bioelectronics.

Dr. Shepard was Technical Program Chair and General Chair for the 2002 and 2003 International Conference on Computer Design, respectively. He has served on the Program Committees for IEDM, ISSCC, VLSI Symposium, ICCAD, DAC, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD. He received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992, a National Science Foundation CAREER Award in 1998, and the 1999 Distinguished Faculty Teaching Award from the Columbia Engineering School Alumni Association. He has been an Associate Editor of IEEE TRANSACTIONS ON VERY LARGE-SCALE INTEGRATION (VLSI) SYSTEMS and is current an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS and IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS.