

Package Inductors for Intel Fully Integrated Voltage Regulators

William J. Lambert, *Member, IEEE*, Michael J. Hill, *Senior Member, IEEE*,
Kaladhar Radhakrishnan, Leigh Wojewoda, *Member, IEEE*, and Anne E. Augustine

Abstract—Intel fourth-generation and fifth-generation Core microprocessors are powered by high-frequency integrated switching voltage regulators. The inductors required to implement these regulators are constructed using the routing layers of conventional organic flip chip packaging. This paper provides an overview of the construction of these inductors including representative results from production packages. Measured inductors reported in this paper span from 1 to 6.7 nH under 2.4 mm² and achieve Q of up to 24 at 140 MHz (the switching frequency).

Index Terms—Integrated voltage regulator (VR), package inductor.

I. INTRODUCTION

INTEL'S fourth- and fifth-generation Core microprocessors feature fully integrated voltage regulators (FIVRs) [1], which are 140-MHz multiphase buck regulators integrated into the microprocessor die and the package. In general, microprocessors have been powered by switching regulators on the motherboard, with one voltage regulator (VR) per high-current power domain. As emphasis on fine grain power management has increased, so too has the number, cost, and area of these motherboard VRs. By integrating the VR on the microprocessor die and the package, FIVR substantially reduces the number of motherboard VRs, reducing the motherboard area and the BOM cost while increasing the number of independent high-current voltage domains. These domains improve product performance by reducing guardbands, improving transient response, and minimizing sleep-state entry and exit latency. A detailed discussion of these advantages is available in [1].

FIVR requirements, such as voltage slew rate and load transient response, place an upper limit on the inductance that is well within the range of package air core inductors (ACIs) [2]. These add no cost (an important consideration given is that more than 50 inductors are required for some products), and in many cases achieve comparable or better performance than discrete components. FIVR package inductors achieve quality factors as high as 24 (for a 1.7-mm² footprint) at the VR switching frequency (\sim 140 MHz) with as little as 6-m Ω dc resistance (DCR). When multiple inductors are placed

Manuscript received June 4, 2015; revised October 9, 2015; accepted November 20, 2015. Date of publication December 21, 2015; date of current version January 15, 2016. Recommended for publication by Associate Editor M. Nakhla upon evaluation of reviewers' comments.

The authors are with the Assembly and Test Technology Development Group, Intel Corporation, Chandler, AZ 85226 USA (e-mail: william.j.lambert@intel.com; michael.j.hill@intel.com; kaladhar.radhakrishnan@intel.com; leigh.e.wojewoda@intel.com; anne.e.augustine@intel.com).

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/TCPMT.2015.2505665

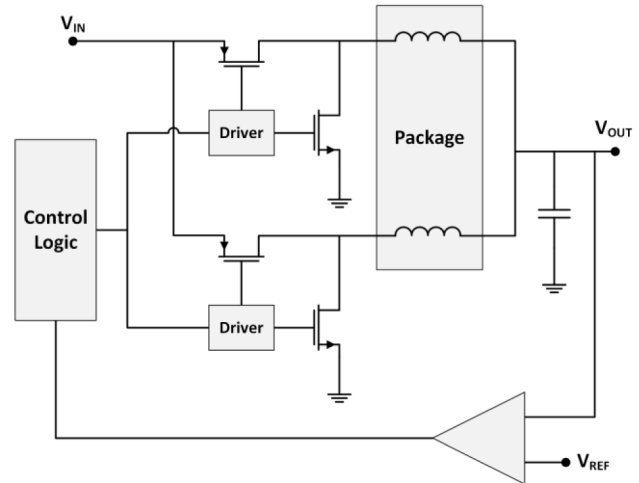


Fig. 1. Circuit representation of a buck regulator.

in parallel, a domain drawing tens of Amperes can see as little as 1-m Ω effective resistance from the ACIs. For FIVRs optimized for low currents, inductors up to 6.7 nH in an area of 2.1 mm² are reported. This paper provides a detailed description of the inductors used for several current microprocessor products along with the measured and simulated data demonstrating their performance.

For reference, a simplified schematic of a two-phase FIVR domain is shown in Fig. 1. Except for the package inductors, all of the remaining circuitry is located on a microprocessor die, which uses Intel 22-nm [3] or 14-nm [4] technology; for additional details on their implementation, see [1].

II. PACKAGE INDUCTOR LAYOUT

Intel microprocessor packaging is customized for individual microprocessor products based on mounting requirements, total design power, thickness and area limitations, and other criteria. At the time of writing, inductors for FIVR are implemented on more than 15 unique package designs already being sold (spanning from low power tablets to high-power servers). Each package design uses five or more different inductor designs. Although inductor designs can be shared between packages with identical design rules and layer count, in many cases, they are custom designed to meet the requirements of each FIVR domain on each individual product. Nearly all FIVRs use two or more buck regulator phases, so the inductors are designed as an array. Fig. 2 shows the bottom of an Intel fourth-generation Core microprocessor, where the arrays of inductors are visible in the center of the package. In Fig. 3,




Fig. 2. Bottom side of a fourth-generation Core LGA package. The package inductors are in the center of the package, which is under the die.

the inductor array for a single CPU core voltage domain is outlined for two separate products. Several different inductor designs are shown in Fig. 3.

Most Intel microprocessors use flip chip ball grid array (BGA) or land grid array (LGA) packaging consisting of a rigid glass cloth core with plated through holes (PTHs) on which multiple pairs of additive routing layers are built up. The cross section of a typical package is shown in Fig. 5(a). Many different inductor topologies are possible in this type of packaging, but practical constraints narrow the options considerably: 1) the inductor must be directly below the power MOSFETs on the die as any lateral offset adds significant DCR; 2) the inductor routing must be on the core or bottom routing layers of the package, so as not to interfere with the output power plane routing and to shield IO traces from the inductor fields; and 3) in many cases, the footprint of the inductors must be equal to or smaller than the footprint of the domain they power in order to avoid interference with inductors on adjacent domains. Although many unique inductors were designed across the different product lines, nearly all of them fell into the four basic topologies covered in Sections II-A–II-D.

A. Uncoupled Solenoid Inductor

The simple uncoupled solenoid inductor structure shown in Fig. 5(a) consistently provided the best results for ac and dc power dissipation. A coordinate axis has been provided to orient the cross section of the structure shown in Fig. 5(b). Connections to the die on either side of the FIVR circuit block are connected directly to PTHs, which provide some series inductance. From there, the bulk of the inductance is realized through parallel turns on the bottom routing layers. The turns terminate in a shorting bar that is connected to the upper power planes that supply the output voltage domain. Fig. 5(a) shows single-turn inductors, but with simple modifications, multiple turns can be created to increase the inductance (for example, in a package with four bottom layers, the first turn would occur on the first and second layers in parallel, and the second turn




Fig. 3. (a) Close-up view of a portion of a 22-nm desktop LGA package with the inductors for an IA Core domain outlined. (b) Same area on a 14 nm BGA package.

on the third and fourth layers in parallel).

Adjusting the number of turns and the area of the void in the center of the turns allows this topology to be tuned from under 1 nH to over 10 nH on a 22-nm product. Using all the bottom routing layers keeps the DCR very low, particularly on high layer count packages. Finally, the topology can easily be tiled to implement an arbitrary even number of phases in a multiphase design.

B. Interleaved Solenoid Inductor

In some cases, locating the output connection of the inductors some distance away from the FIVR circuitry improves the routing of the upper metal layers. A variation in the previous topology that accomplishes this is the interleaved solenoid inductor, shown in Fig. 5(c), with a cross section shown in Fig. 5(d). A coupled inductor is formed by two separate, oppositely oriented windings that are interleaved. The opposite orientation of the windings also introduces negative coupling between the inductors, while the parallel PTHs and common output path introduce a small amount of positive coupling. The net effect is that the two inductors have negative coupling with a coupling factor between -0.5 and -0.1 ; this coupling factor provides a small reduction in ac power dissipation.

Interleaving the windings effectively halves the amount of metal in the dc current path versus the uncoupled inductor, causing increased DCR. As a result, this topology is most frequently used for domains with low average current such as mobile products. Mobile products now almost exclusively use BGA packages with reduced thickness cores, which results in more coupling from the inductor to the motherboard and package planes [as evident from Fig. 5(c)]; this causes decreased inductance and increased eddy current losses.

C. Shielded PTH Loop Inductor

In some cases, the floor plan of the die requires the inductor to be placed over an area of the package with pins or bumps on the bottom layer of the package. Placing the void of the inductor in Fig. 5(a) immediately over a pin or a metal shielding plane dramatically lowers the inductance and quality factor

of the inductor. An alternative topology is the shielded PTH loop inductor shown in Fig. 5(f), along with its cross section in Fig. 5(e), where the inductor loop is formed using only the metal layers on either side of the core. In this topology, the inductor loop (and, hence, the orientation of the B-field in the center of the inductor) is in the xz plane as labeled, as opposed to the previous two topologies that are oriented in the xy plane. To form an array of inductors, the windings are alternated between clockwise and counterclockwise directions to result in negative coupling between phases.

This topology results in a lower quality factor when compared with those of the previous topologies due to coupling to adjacent metal planes on the package. Strategic voiding of these planes can reduce the effect, but in practice, metal density balancing requirements limit the improvement. There is also an R_{dc} penalty, because only two routing layers are used.

D. 3DL

In order to enable extremely thin laptops and tablets, the Core M microprocessor has an aggressive total thickness target of 1.05 mm for the die and the package. Meeting the thickness target requires a thinner core layer than the BGA package cross section shown in Fig. 5(c).

In order to avoid a reduction in the inductor performance due to the reduced package thickness, an alternative inductor implementation is used. A hole is routed in the motherboard of the platform, which is usually at least 800- μm thick. A two-layer, 800- μm -thick Printed Circuit Board (PCB) inductor module (called the 3DL module) is then attached by BGA to the bottom of the microprocessor package. When the package is assembled on the motherboard, the 3DL module is suspended from the package in the hole routed in the motherboard. A cross section of a package with 3DL is shown in Fig. 5(g), and an isometric view of the inductor routing is shown in Fig. 5(h). A Core M package with 3DL is shown in Fig. 4.

Any of the previous topologies could be implemented in the 3DL module using a suitable PCB technology, but an inductor in the xz plane with a similar structure to the enclosed inductor from the previous section was chosen. This topology allows 3DL to be implemented using a two layer, mechanically drilled PCB with coarse design rules that can be manufactured at a very low cost. Since the 3DL module does not require any fine routing, very thick metal is used to minimize the DCR of the inductors.

3DL is the only implemented inductor option that adds cost to the package (due to the PCB module cost and added assembly steps)—the other topologies use only existing routing layers on the package, and hence do not increase cost. It was selected for Core M packaging after simulations showed that 3DL dissipated significantly less power than the uncoupled solenoid and interleaved solenoid topologies at operating points important for that specific product (as shown in Table III).

III. PACKAGE INDUCTOR DESIGN

Design of the inductor roughly follows the procedure for a lower frequency multiphase buck regulator with some addi-




Fig. 4. Cross-sectional view of a package featuring 3DL.

tional constraints due to area limitations and FIVR's high switching frequency (in particular, the voltage ripple does not decrease as a function of the number of phases as it would in a low-frequency system). A detailed block diagram of the system is provided in [1, Fig. 2]. A simplified outline is as follows: the switching frequency is fixed near 140 MHz based on Radio Frequency Interference requirements, the regulator operates in continuous conduction mode with phase shedding, and the amount of output capacitance is roughly fixed based on the area of the voltage domain (additional details are available in [1]). Within these constraints, a lower bound to the acceptable range of inductance is provided by voltage ripple specifications, and an upper bound is required to meet load transient response requirements and to ensure zero-voltage switching at light loads. These inductance requirements are unique for each FIVR based on the circuitry that resides on the voltage domain.

Once bounds are established, then candidate designs are laid out based on compatible topologies from Section II (in general, only one or two of the topologies will be compatible with all the requirements of the package). All the candidate topologies are simulated (along with the output plane routing) using a full-wave simulator (discussed in Section IV), and multiple design-simulation iterations are completed until the desired inductance target is achieved. The optimal inductor is then chosen as the candidate that falls within the specified range while dissipating the least power per

$$P_{\text{Loss}} = I_{\text{out}}^2 (R_{\text{DS,on}} + R_{\text{dc}}) + I_{\text{RMS,ac}}^2 (R_{\text{DS,on}} + R_{\text{ac}}) + P_{\text{sw}} \quad (1)$$

where I_{out} is the dc load current, $R_{\text{DS,on}}$ is the ON-resistance of the MOSFETs, R_{dc} is the DCR of the inductor, $I_{\text{RMS,ac}}$ is the RMS current through the inductor excluding I_{out} , R_{ac} is the ac resistance (ACR) of the inductor at the switching frequency of the VR, and P_{sw} is the power dissipated by the switching circuitry (and is a function of switching frequency, voltage, and other parameters). $I_{\text{RMS,ac}}$ is a function of the conversion ratio; for an uncoupled inductor, a reasonable estimate is given by (2), where V_{in} is the input voltage to the regulator and V_{out} is the output voltage of the regulator

$$I_{\text{RMS,ac}} = \frac{(V_{\text{in}} - V_{\text{out}}) \left(\frac{V_{\text{out}}}{V_{\text{in}}} \right)}{2\sqrt{3}L f_{\text{sw}}} \quad (2)$$

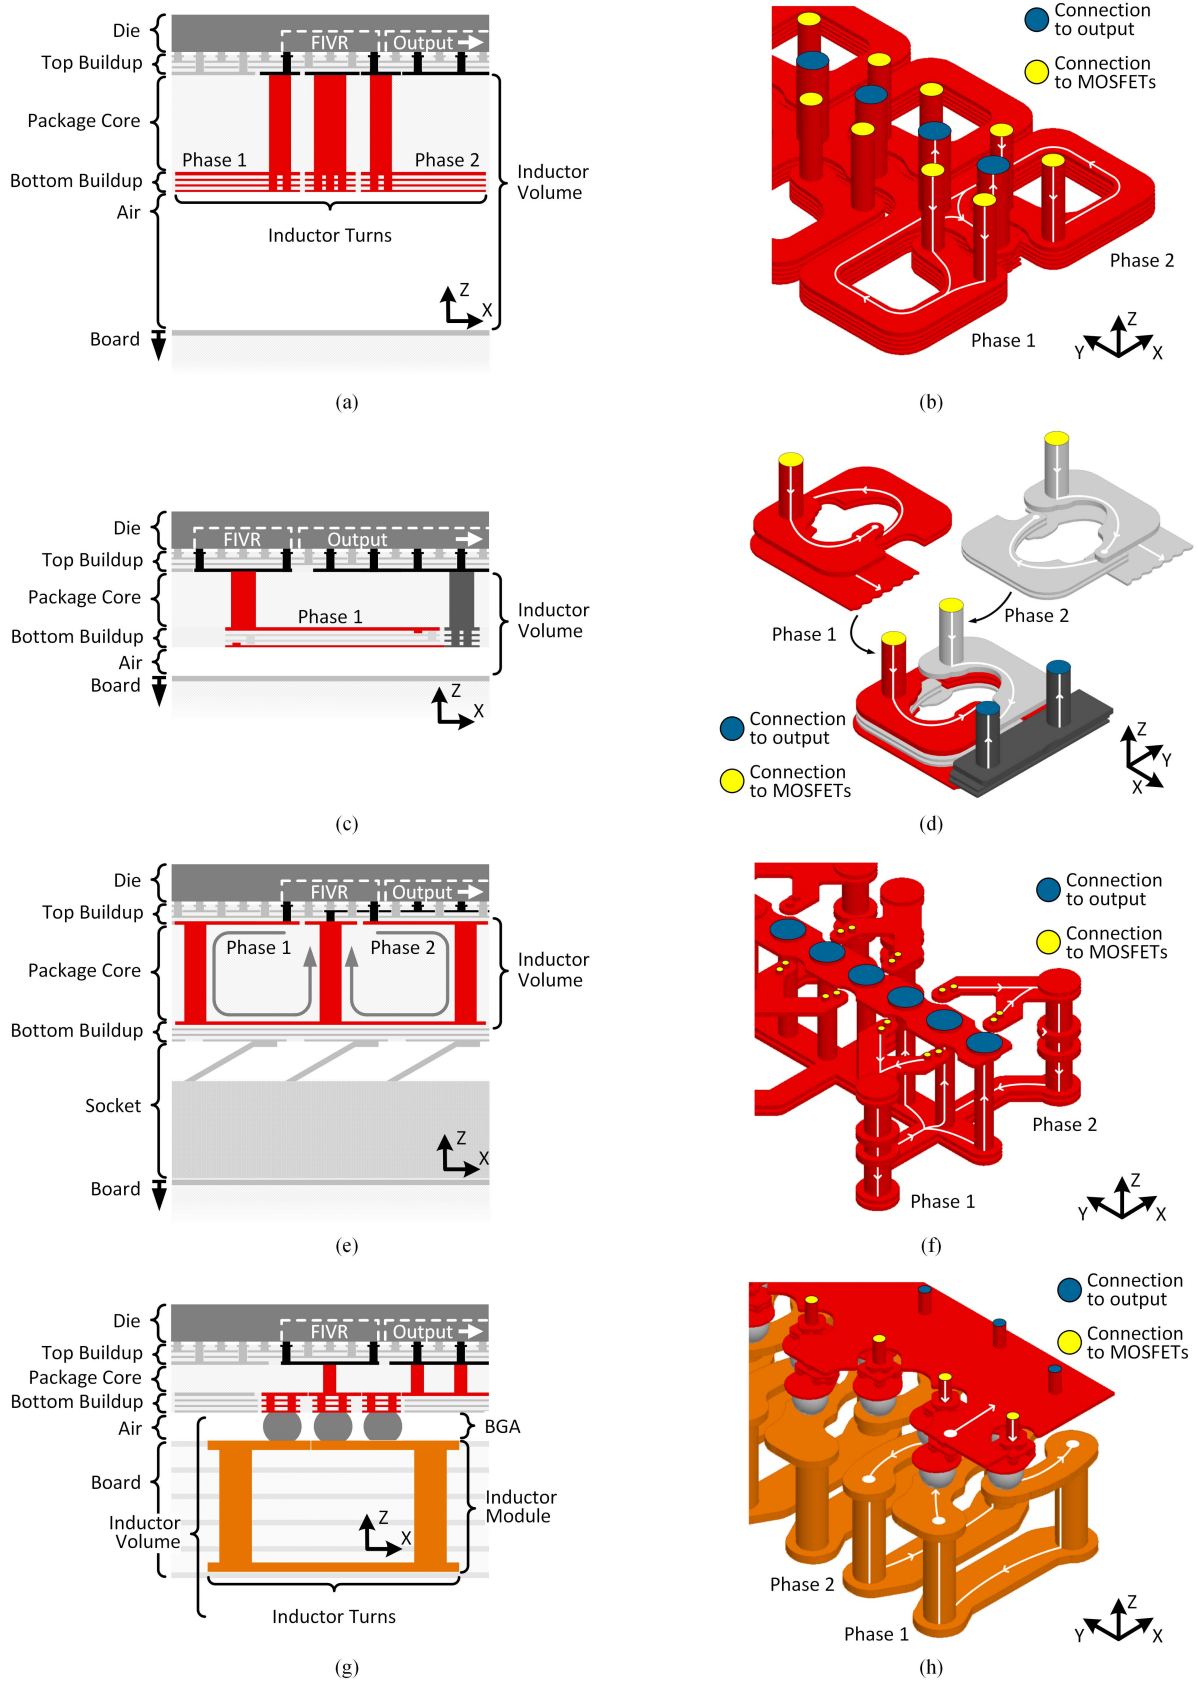


Fig. 5. (a) Cross-sectional view of an uncoupled solenoid inductor array in a 700- μm core LGA package. (b) Isometric view of the same inductor. (c) Cross-sectional view of an interleaved solenoid inductor in a 400- μm core BGA package. (d) Isometric view of the same inductor. (e) Cross-sectional view of a shielded PTH loop inductor in a 700- μm core LGA package in an area over the socket. (f) Isometric view of the same inductor. (g) Cross-sectional view of a 3DL inductor array suspended from a 200- μm core BGA package. (h) Isometric view of the same inductor.

The optimal inductor design is a function of the operating point (V_{in} , V_{out} , and I_{out}), so critical operating points must be provided as part of the design specifications. In general, the best inductor minimizes R_{dc} and $R_{ac}/(L \cdot f_{sw})^2$ [based on (2)]. However, because the inductors are implemented in a package without any magnetic material and with a fixed number of routing layers (adding layers is expensive), it is not possible to increase inductance without also increasing R_{dc} . The design optimization becomes a tradeoff between limiting the R_{ac} losses (by increasing inductance) that dominate (1) when I_{out} is small and the R_{dc} losses that dominate (1) when I_{out} is large.

The R_{dc} of the inductor may impose additional design constraints, because some FIVR inductors must carry several amperes of current. Long-term reliability requirements restrict the maximum current density of the package features, and this reliability limit decreases as operating temperature increases. Therefore, the inductors are designed (in terms of trace width and the number of microvias between layers) in order to meet the current density requirements, and a maximum R_{dc} bound is imposed to keep the self-heating to a specified temperature.

Although the power dissipated due to the ac current in the inductor is proportional to $R_{ac}/(L \cdot f_{sw})^2$, this paper will use the quality factor (3) as the metric of inductor ac performance in order to be consistent with other literature

$$Q(f) = \frac{\text{Im}\{Z_L(f)\}}{\text{Re}\{Z_L(f)\}} = \frac{2\pi f L(f)}{R_{ac}(f)}. \quad (3)$$

IV. INDUCTOR EVALUATION

A. Measurement

Passive inductor measurements are performed using a vector network analyzer (VNA) with RF probes that are landed on the die bumps of bare production microprocessor substrates. Measurements are made in the two-port shunt configuration [5], as the impedance of many structures is under 100 m Ω in the frequency range of interest. By convention, the inductance is measured by probing between the bumps connecting the inductor to the power MOSFET and the closest die bumps on the output plane. Fig. 6(a) shows the RF probes landing on the die bumps on the top of a package and their alignment to the inductor shape on the bottom layers of the package. Fig. 6(b) provides a cross-sectional view of the landing of the probe on the uncoupled solenoid inductor from Fig. 5(a). More detail on the measurement configuration is provided in [2, Fig. 8]. The results inevitably have a small dependence on the probing configuration, because the inductance of the output power plane routing is nonnegligible compared with the small value of the inductor itself. In this paper, the results are reported as is; no effort is taken to artificially deembed the power plane contribution from the inductor itself. In this way, the inductance measured is similar to the inductance seen by the CPU die looking into the inductor structure. For correlation purposes, simulations of the structures use the same nodes for measurement ports.

Active measurements are used for the *in-situ* measurements of the conversion efficiency of the entire FIVR domain (see the

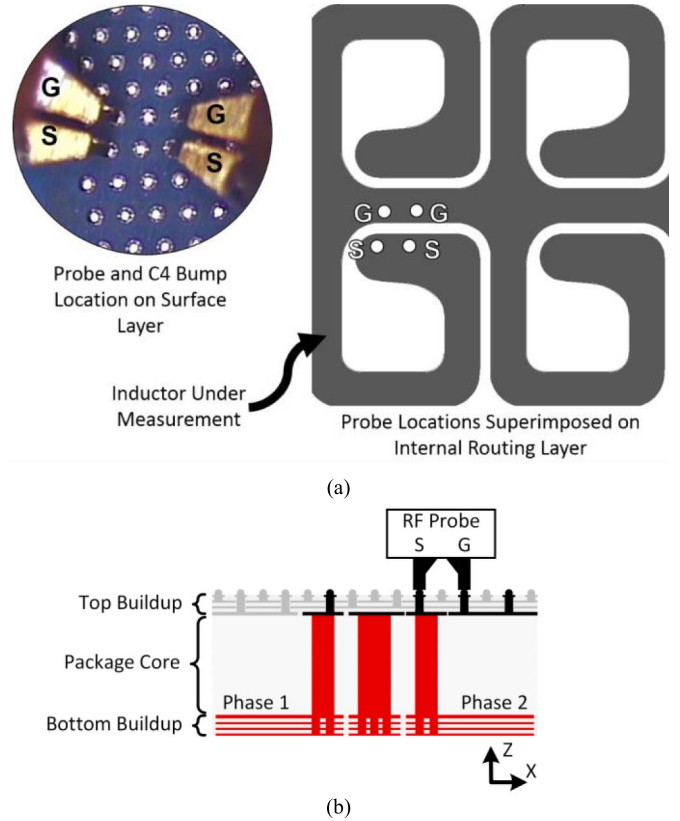


Fig. 6. (a) Landing RF probes on die bumps for inductor measurements. (b) Cross-sectional view of the RF probe measurement.

efficiency results reported in [1] for an example). In general, it is not possible to separate the losses from the inductors from other loss mechanisms. However, certain settings, such as the timing delay between phases with coupled inductors, result in changes to the dissipated power that are a direct consequence of the inductor design, and these results are shared in Section V.

B. Simulation

The frequency range of interest for the simulations spanned from dc to several times the switching frequency (f_{sw}) (which is necessary to accurately capture voltage ripple and power loss at harmonics of the switching frequency). Early analysis showed that the resistance at f_{sw} and above was very sensitive to the details of the inductor routing, including small features in the upper layers, which were not easily accounted for by analytical models. Therefore, full 3-D extractions of the package routing are used for modeling. Full-wave solvers provide the greatest flexibility for configuring excitations, boundaries, and solver settings, and so were used for most simulation tasks.

Multiple versions of two common commercially available full-wave solvers have been tested extensively and provide accurate results when configured as follows. First, because the skin depth at f_{sw} is close to the thickness of the conductors, the fields must be solved inside conductors. Second, to solve inside the conductors accurately with a reasonable mesh size,




Fig. 7. Comparison between the measured and simulated inductances of a package inductor.

first-order or higher basis functions must be used. Third, in order to correctly calculate losses due to induction, conductive objects in close proximity to the inductors (1 mm in practice) must be included in the simulation even if they are not a part of the package.

In order to validate the simulation results, simulations were run in an identical configuration to the VNA measurements. Typical results from a 14-nm BGA package with 3DL are shown in Figs. 7 and 8. The results match well across the entire simulated frequency range and are within 3% at the switching frequency. The largest discrepancy occurs in the resistance at very high frequencies, which is due, at least in part, to the sensitivity limits of the VNA, because the reactance is more than an order of magnitude larger than the resistance.

Changing an FIVR inductor requires manufacturing an entirely new package that adds cost and could delay a product. In order to avoid this, every unique inductor design is simulated in detail.

V. RESULTS

A. Inductors on a Representative Package

This section covers the different types of inductors found on an LGA package from an Intel fourth-generation Core desktop microprocessor. All the inductors on this package use the uncoupled solenoid topology. This choice was made based on the simulation results after comparing with the designs using the interleaved solenoid and shield PTH loop inductor topologies. In Fig. 2, the backside of the package with the inductor routing on the bottom layer clearly visible in the center of the package is shown. A close-up view of the inductors is shown in Fig. 3(a), where the inductor array for a CPU Core voltage domain is outlined in white.

Table I summarizes the performance characteristics of the five distinct inductor designs used in this package. For the two high-current domains (greater than 20-A maximum current), eight or more inductors are connected in parallel to support the maximum current. Transient response requirements mandate a small effective inductance for these domains, resulting in a final value of just around 1 nH, which was implemented using




Fig. 8. Comparison between the measured and simulated resistances of a package inductor.

TABLE I
INDUCTOR DESIGNS ON THE FOURTH-GENERATION CORE LGA PACKAGE

Voltage Rail	Inductor Area	L [nH] @		R [mΩ] @	
		140 MHz	140 MHz	140 MHz	RDC [mΩ]
High Current #1 ^a	1.4 mm x 1.7 mm	0.99	36.9	6.7	
High Current #2 ^a	1.4 mm x 1.2 mm	1.17	42.6	7.1	
Med. Current #1 ^a	1.4 mm x 0.6 mm	2.39	110.6	9.1	
Med. Current #2 ^a	1.4 mm x 1.5 mm	3.17	102.1	12.0	
Low Current ^a	1.4 mm x 1.5 mm	6.68	212.5	36.3	

^a Uncoupled solenoid topology

TABLE II
MEASURED AND MODELED 4 × 4 INDUCTANCE MATRIX

Sim. [L] [nH] - 108 MHz	Meas. [L] [nH] - 108 MHz
$\begin{bmatrix} 1.34 & -0.41 & 0.02 & 0.04 \\ -0.41 & 1.32 & 0.40 & 0.04 \\ 0.02 & 0.40 & 1.33 & -0.39 \\ 0.04 & 0.04 & -0.39 & 1.39 \end{bmatrix}$	$\begin{bmatrix} 1.39 & -0.42 & 0.02 & 0.03 \\ -0.42 & 1.40 & 0.40 & 0.04 \\ 0.02 & 0.40 & 1.38 & -0.40 \\ 0.03 & 0.04 & -0.40 & 1.47 \end{bmatrix}$

the uncoupled topology. A quality factor in the mid-twenties keeps the ac losses low despite the small inductance. When operating at light loads, the FIVR will shed all but two of the phases to further minimize the ac losses. The uncoupled single-turn inductor topology achieves an extremely low DCR. When eight are placed in parallel, the effective resistance is <1 mΩ.

A second turn is added to the inductors on the medium current rails (5–20-A maximum current) to achieve significantly higher inductance. These inductors maintain similar Q values to the high-current inductors, though the DCR grows due to the second turn. However, the ratio of DCR to inductance is actually improved. Finally, for the low-current domain (under 5-A maximum current), a third turn was added to realize a 6.7 nH inductor in order to ensure good light load efficiency and minimal voltage ripple.

TABLE III
COMPARISON OF MEASURED ACI PROPERTIES IN DIFFERENT PACKAGING TECHNOLOGIES

Product	Si Technology	Packaging	Inductor Area	L [nH] @ 140 MHz	R [mΩ] @ 140 MHz	Q @ 140 MHz	RDC [mΩ]
4th Generation Core™ - Desktop ^a	22 nm	700 μm core LGA	1.4 mm x 1.2 mm	1.17	42.6	24.0	7.1
4th Generation Core™ - Mobile ^a	22 nm	400 μm core BGA ^d	1.4 mm x 1.2 mm	1.66	74.5	19.7	8.8
4th Generation Core™ - Mobile ^b	22 nm	400 μm core BGA ^d	1.5 mm x 1.2 mm	1.70	92.7	16.1	16.6
5th Generation Core™ - Mobile ^b	14 nm	400 μm core BGA ^d	1.2 mm x 1.1 mm	1.46	88.6	13.9	15.4
5th Generation Core™ M - Mobile ^c	14 nm	200 μm core BGA	1.8mm x 0.7 mm	1.53	57.1	23.6	10.9

^aUncoupled solenoid topology

^bInterleaved solenoid topology

^c3DL topology

^dTypical inductor to motherboard spacing is less than 300 μm

The results in Table I show that a large range of inductances can be implemented using package inductors with roughly the same footprint. The LGA package described in this section includes a total of 59 custom designed inductors for FIVR without any additional cost to the product.

B. Inductors on Different Package Technologies

Table III lists the measured parameters for the same high-current domain on several different packaging technologies. The inductor shown in Table I corresponds to the first line of the table.

Several trends are evident. First, the inductance target for mobile products is significantly higher than that of the desktop product; this is expected, as mobile products have much lower average and maximum operating power. Second, the Q of the inductor decreases and the DCR of the inductor increases as the area of the inductor and the thickness of the package core are decreased. Fig. 3(b) shows the area for interleaved solenoid Intel Architecture (IA) Core inductors on a 14-nm CPU—the reduction in area is evident when compared with the uncoupled solenoid inductors for a 22-nm IA Core shown in Fig. 3(a). Inductors in the BGA packages pay an additional performance penalty because of their proximity to the motherboard (this is partially offset by a design guide requirement to void the metal of the uppermost motherboard layer under the inductors). Finally, the results from a package with 3DL are provided in the final row of the table. 3DL provides a higher quality factor with a larger inductor than the best 22-nm package inductor, while fitting a footprint that was shrunk by nearly 50% due to scaling the transistors from 22 to 14 nm.

C. Coupling Between Inductors

Table II contains the measured and simulated inductance matrices for two interleaved coupled inductors that are implemented side by side in the same inductor array [as shown in the outlined box in Fig. 3(a)]. Excellent correlation is seen between the measured and modeled results for both the diagonal (self-inductance) and the off-diagonal (mutual inductance) terms. The dominant negative coupling term between inductors one and two and also between inductors three and four produces a beneficial coupling coefficient (ratio of mutual to self inductance) of -0.3 . However, other nonnegligible

TABLE IV
SIX-PHASE INDUCTANCE MATRIX AT 140 MHz

Inductance @ 140 MHz [nH]					
1.35	-0.61	-0.05	0.02	-0.01	0.01
-0.61	1.46	0.35	-0.03	0.09	-0.01
-0.05	0.35	1.44	-0.64	-0.03	0.02
0.02	-0.03	-0.64	1.45	0.36	-0.05
-0.01	0.09	-0.03	0.36	1.45	-0.61
0.01	-0.01	0.02	-0.05	-0.61	1.34

TABLE V
POWER DISSIPATION VERSUS VR PHASE ASSIGNMENT

Phase Assignment	Δ Dissipated Power [mW]
[0° 180° 124° 304° 248° 68°]	Baseline
[0° 180° 180° 0° 0° 180°]	-29
[0° 180° 0° 180° 0° 180°]	75

coupling terms are also present in the matrix due to the proximity of the two pairs of inductors, as seen in the positive coupling term between inductors two and three. The timing of the VR phases, which is programmable, can be optimized to minimize the power dissipation of the whole array due to these coupling terms. The full inductance matrix for a six-phase domain is shown in Table IV. Table V shows the measured delta in power dissipation from an active microprocessor when the phase angles are adjusted.

It is obvious from the proximity of the inductors in Fig. 3 that coupling will also occur between inductors on separate domains. This is undesirable because when one of the two domains is OFF, the other domain will couple energy into the inactive inductor, burning extra power. Careful design of the inductors minimizes this effect.

D. Variation From High Volume Manufacturing

In order to ensure acceptable performance of the FIVRs, the package inductors must meet targets for inductance and resistance variations in high volume manufacturing. To quantify component variability, packages from multiple lots manufactured by two different vendors were measured.

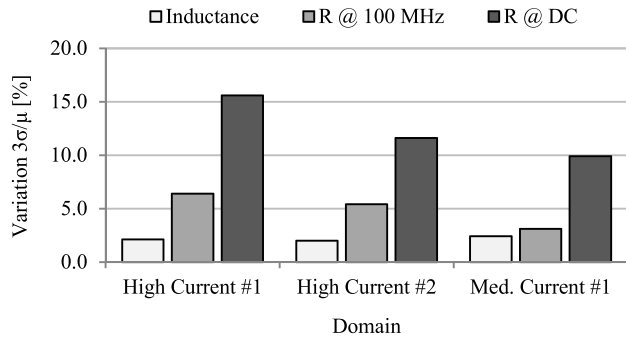


Fig. 9. Variation in measured inductance and resistance for three different inductor designs.

TABLE VI
MANUFACTURING VARIATION OF THE HIGH-CURRENT #1 ACI

Vendor	Variation ($3\sigma/\mu$)		
	L^a	R_{AC}^a	R_{DC}
Vendor A	1.5%	4.1%	6.8%
Vendor B	2.3%	3.3%	10.3%
All Vendors	2.3%	5.4%	15.6%

^aMeasured at 100 MHz

AC performance was verified with 2800 VNA inductance and resistance measurements. The DCR of an even larger number of inductors (over 13000) was measured using the four-wire resistance meter. Units falling outside the expected measurement range were cross-sectioned to find the root cause of the manufacturing defect.

Fig. 9 summarizes the variations in the values of inductance for three inductor domains, and Table VI shows the measured variation of a single inductor on packages supplied by multiple manufacturers. The variation is quantified in terms of three standard deviations divided by the mean value of the measurement ($3\sigma/\mu$), which is an indication of how much the components vary from their mean value. The variation in inductance is very small, less than 2.5% for all domains, and a direct result of the level of process control required to meet the strict IO impedance specifications elsewhere in the package.

Larger variation is seen in the ACR and DCR, as illustrated in Fig. 9. Both ACR and DCR are sensitive to small variations in the thickness of the metal routing layers. The variation in the metal layer thickness results in a large impact on DCR, because the cross-sectional area of the conductor can vary significantly. The metal layer thickness variation does not impact ACR and inductance strongly, because skin effect constrains the currents to the perimeter of the conductor.

For package inductors measured in this paper, manufacturing variations were found to be low, and it was concluded that the measured variation values do not hinder system performance.

VI. CONCLUSION

Package inductors are an important part of enabling integrated voltage regulation on Intel fourth- and fifth-generation Core microprocessors. Integrating the inductors into the package carries some risks because of the difficulty in making changes late in the design, but this paper has demonstrated that well-designed simulation and measurement procedures can accurately predict and measure the inductors, which removes much of this risk. Furthermore, the results achieved show that the inductors that are manufactured have very good performance, even when compared with discrete options.

Aside from reducing cost, the embedded inductors have the additional benefit of allowing the designer to choose the precise value, footprint, and location of the inductor. Affordable discrete passives would have to have been selected from a limited set of values and form factors, and correspond to chip shooter compatible placement rules. A further advantage comes from the extremely tight routing tolerances required on the package for controlled impedance high-speed IO signaling. These tolerances produce embedded inductors with minimal variation between packages even between separate manufacturers, as shown in Seciton V-D.

As the silicon area continues to scale and the current density increases, the effective footprint that is available for the package inductors continues to shrink each generation. The situation is further compounded by the fact that the microprocessors in certain segments, such as Core M, are moving toward thinner and thinner packages. Technologies, such as 3DL, have already been implemented to enable thinner packages without sacrificing inductor performance. In the future, if the current density continues to scale, it may be necessary to introduce magnetic materials to increase inductance density in order to meet product requirements.

REFERENCES

- [1] T. A. Burton *et al.*, "FIVR—Fully integrated voltage regulators on 4th generation Intel Core SoCs," in *Proc. 29th Annu. IEEE Appl. Power Electron. Conf. Expo.*, Fort Worth, TX, USA, Mar. 2014, pp. 432–439.
- [2] W. J. Lambert, M. J. Hill, K. Radhakrishnan, L. Wojewoda, and A. E. Augustine, "Package embedded inductors for integrated voltage regulators," in *Proc. IEEE 64th Electron. Compon. Technol. Conf. (ECTC)*, Orlando, FL, USA, May 2014, pp. 528–534.
- [3] C. Auth *et al.*, "A 22 nm high performance and low-power CMOS technology featuring fully-depleted tri-gate transistors, self-aligned contacts and high density MIM capacitors," in *Proc. Symp. VLSI Technol.*, Honolulu, HI, USA, Jun. 2012, pp. 131–132.
- [4] S. Natarajan *et al.*, "A 14 nm logic technology featuring 2nd-generation FinFET, air-gapped interconnects, self-aligned double patterning and a 0.0588 μm^2 SRAM cell size," in *Proc. IEEE Int. Electron Devices Meeting (IEDM)*, San Francisco, CA, USA, Dec. 2014, pp. 3.7.1–3.7.3.
- [5] L. Wojewoda, D. Athreya, and M. J. Hill, "Use condition characterization of package components," in *Proc. ASME Pacific Rim Tech. Conf. Expo. Packag. Integr. Electron. Photon. Syst.*, Portland, OR, USA, 2011, pp. 633–639.
- [6] N. Kurd *et al.*, "Haswell: A family of IA 22 nm processors," in *IEEE ISSCC Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2014, pp. 112–113.



William J. Lambert (M'05) received the B.S. and M.S. degrees from the Rochester Institute of Technology, Rochester, NY, USA, in 2005, and the Ph.D. degree in electrical engineering from Arizona State University, Tempe, AZ, USA, in 2008.

His work at Arizona State University on advanced power delivery solutions for microprocessors was funded by a fellowship from Intel Corporation, Chandler, AZ, USA, where he was also an Intern from 2005 to 2008. In 2008, he joined the Assembly and Test Technology Development Group, Intel Corporation, where his work is focused on passive and active modeling and measurement of microprocessor integrated voltage regulators.



Michael J. Hill (SM'05) received the B.S., M.S., and Ph.D. degrees from the University of Arizona, Tucson, AZ, USA, all in electrical engineering.

He was involved in the development of high-density active microwave phased arrays for medical applications, and later in the development of wireless networking hardware and characterization systems with an emphasis on medium and low-gain antenna characterization utilizing spherical near-field scanning techniques. He has been with the Electrical Core Competency Group, Intel Corporation, Chandler, AZ, USA, since 2002, where he has been involved in microprocessor power delivery analysis and metrology development. He is currently a Principal Engineer with Intel Corporation. His current research interests include the development of new tools and techniques to allow precise characterization of microprocessor power delivery networks. Most of his recent work has focused on characterization techniques for microprocessor power delivery systems utilizing integrated voltage regulators.

Dr. Hill is a member of Tau Beta Pi and Eta Kappa Nu.



Kaladhar Radhakrishnan received the Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign, Urbana, IL, USA.

He joined Intel Corporation, Chandler, AZ, USA, in 2000, where he is currently a Senior Principal Engineer and manages the Power Delivery team in the Electrical Core Competency Group. His current research interests include power delivery and computational electromagnetics, which was the focus of his dissertation work.



Leigh Wojewoda (M'07) received the B.S. degree in engineering physics and the M.S. degree in mechanical engineering from the University of Arizona, Tucson, AZ, USA.

She joined Intel Corporation, Chandler, AZ, USA, in 1997, where she has been with the Assembly Technology Development Division since 1997, and she is currently a Senior Packaging Engineer and manages the Electrical Measurement Laboratory. She is responsible for development and characterization related to microprocessor packaging. The

laboratory is also active in demonstrating new measurement techniques and correlating measurement results with various Intel suppliers.



Anne E. Augustine received the B.S. degree in electrical engineering from Arizona State University, Tempe, AZ, USA.

She has been with the Assembly Technology Development Division, Intel Corporation, Chandler, AZ, USA, since 2000, and is a Senior Packaging Engineer with the Electrical Core Competency Group. She is involved in simulation methodology for current and next generation power delivery solutions.