FIVR – Fully Integrated Voltage Regulators on 4th Generation Intel® Core™ SoCs

Edward A. Burton, Gerhard Schrom, Fabrice Paillet, Jonathan Douglas
CCDO
Intel Corporation
Hillsboro, OR, USA

William J. Lambert, Kaladhar Radhakrishnan, Michael J. Hill
ATTD
Intel Corporation
Chandler, AZ, USA

***Abstract*—Intel's® 4th generation Core™ microprocessors are powered by Fully Integrated Voltage Regulators (FIVR). These 140 MHz multi-phase buck regulators are integrated into the 22nm processor die, and feature up to 80 MHz unity gain bandwidth, non-magnetic package trace inductors and on-die MIM capacitors. FIVRs are highly configurable, allowing them to power a wide range of products from 3W fanless tablets to 300W servers. FIVR helps enable 50% or more battery life improvements for mobile products and more than doubles the peak power available for burst workloads.**

I. INTRODUCTION

Intel's® 4th generation Core™ microprocessors (code name Haswell) are powered by Fully Integrated Voltage Regulators (FIVR), the industry's first large scale deployment of high current switching regulators integrated into a VLSI die and package. An overview of the scheme is given in Fig. 1(a). A first stage VR, which is on the motherboard, converts from the PSU or battery voltage (12-20V) to approximately 1.8V, which is distributed across the microprocessor die. The second conversion stage is comprised of between 8 and 31 (depending on the product) FIVRs, which are 140MHz synchronous multiphase buck converters with up to 16 phases. A simplified schematic for a two phase FIVR domain is shown in Fig. 1(b). The power FETs, control circuitry and high frequency decoupling are on the die, while the inductors and mid-frequency input decoupling capacitors are placed on the package. Each FIVR is independently programmable to achieve optimal operation given the requirements of the domain it is powering. The settings are optimized by the Power Control Unit (PCU), which specifies the input voltage, output voltage, number of operating phases, and a variety of other settings to minimize the total power consumption of the die.

FIVR is the enabling technology behind key improvements for Intel's® 4th generation Core™ microprocessors including a 50% or more increase in battery life for mobile products, and a 2-3x increase in peak available power (which converts into burst performance). The motherboard voltage regulators eliminated by FIVR free up space that can be used to add platform features or reduce platform dimensions. Details are discussed in Section V.

A. Background

Intel's® 2008 microprocessor microarchitecture introduced the Power Control Unit (PCU) [1], a microcontroller that monitored conditions across the die in real time, and dynamically adjusted a variety of settings to optimally manage power consumption and performance. One of the most important features controlled by the PCU were newly added high current power gates, which provided a significant improvement in CPU energy efficiency by eliminating large leakage losses on idle compute domains. The power gates on high current domains were introduced by adding gate transistors into thin 'cracks' between major functional blocks and represented a very small percentage of the total die area. The die bumps required to support the high current domains required a much larger area than the gate transistors themselves. The large "bump area" posed a hard barrier to productization until a scheme was devised to "borrow" bumps from surrounding circuitry using a thick, low loss routing layer. An extension of this scheme makes FIVR affordable.

A limitation of power gate is that all active domains still operate at the highest voltage required by any individual domain. To create separate voltage domains an entirely new regulator must be added to the motherboard, which adds cost, increases area, and requires extra package pins. An improvement suggested by recent research is the integration of high frequency buck regulators directly on the microprocessor package, or in the die itself [2] [3] [4] [5]. This allows a much larger number of independent power domains, each managed dynamically to match the local computational demand. For example, this would allow one CPU core to run at an elevated voltage and frequency to satisfy a heavy computational load, while other cores execute lower priority code at a much lower voltage and frequency to save power.

Each of the previous works cited had at least one issue making it poorly suited for broad, high volume deployment. The multi-chip approach taken by [2] resulted in an effective current density of 1.3A/mm^2 that would make it expensive to implement because of the silicon area required to support a full power microprocessor product. In [3] another

978-1-4799-2325-0/14/$31.00 ©2014 IEEE

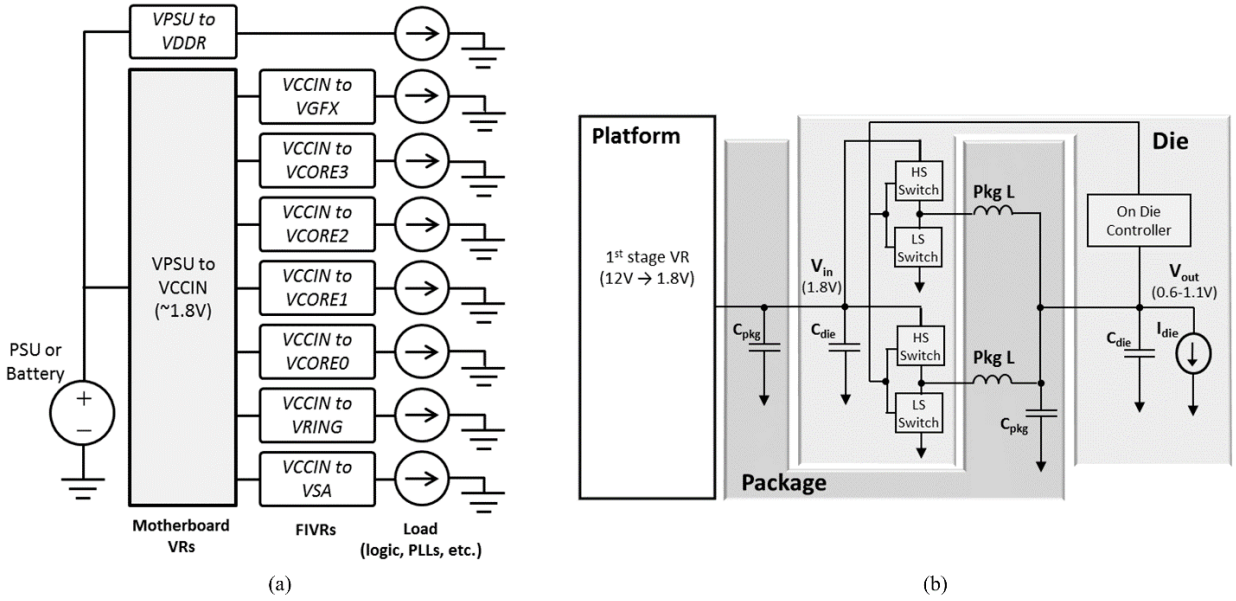

Figure 1. (a) Representative partitioning of the separate high current power domains on a 4th generation Core™ Microprocessor. (b) Simplified schematic of a single FIVR domain, showing the partitioning of the components between the die and the package.

example of the multi-chip approach is demonstrated. This work, which used a 90nm process and inductors integrated onto the die, achieved a higher current density (8A/mm^2) but reported a relatively low efficiency of 76% (compared to 85% for 3.3V to 1.0V conversion in [2]). Instead of the multi-chip approach, the authors of [4] integrated the regulator directly into the die on a 45nm process, but still suffered from relatively low current density (1.7A/mm^2). The authors also report an efficiency of 83% for 1.5V to 1.0V conversion due in part to the quality of the discrete inductors that were used.

FIVR builds on the VR designs presented in [2], while the implementation strategy that makes FIVR affordable, is an extension of the bump "borrowing" scheme developed for the high current power gates [1].

B. Motivation

This paper will show that FIVR addresses the issues in prior work that prevented broad deployment of integrated switching regulators in high volume products. Extending the earlier bump borrowing scheme yields the same current density increase, and corresponding cost decrease, that first made power gates affordable. Improvements in the inductors and transistors yield efficiency in the 90% range for typical high power workloads. The high unity gain frequency (up to 80MHz) allows FIVR to work with just on die MIM for output capacitance.

While these advancements are necessary, they're insufficient to make a reasonable business case for FIVR. To pay the costs of developing and fabricating FIVR, it's important to quantify the actual customer-visible benefits provided in a real implementation. At the start of the design FIVR's expected benefits fell into half a dozen categories. FIVR delivered material benefits in every category, and some benefits were far larger than expected. The benefit categories were: battery life increase, increased available power (for increased burst performance), decreased power required for a given level of performance (or almost equivalently, increased performance for a given power consumed), decreased platform cost and size, improved product flexibility and scalability. See section V for the detailed FIVR impact.

II. Implementation, Design, and Simulation

A. Circuitry

A block diagram representing the circuitry for a single FIVR domain is shown in Fig. 2. The buck regulator bridges are formed by replacing the power gates in previous products with NMOS and PMOS cascode power switches. The cascode configuration allows the power switches to be implemented with standard 22nm logic devices while still handling an input voltage of up to 1.8VDC [2]. This avoids the cost of extra processing steps for high voltage devices, while achieving excellent switching characteristics. The bridge drivers are controlled thru high-voltage level-shifters and support ZVS (zero-voltage switching) and ZCS (zero-current-switching) soft-switching operation. The gates of the cascode devices are connected to the "half-rail", V_{ccdrvn}, regulated to $V_{in}/2$. This is also the negative supply of the PMOS bridge driver as well as the positive supply of the NMOS bridge driver.

The area occupied by the power switches and drivers is small, so they are distributed across the die, immediately above the connection to their associated package inductor which minimizes routing losses. This is illustrated in Fig. 3(a), which shows the location of the package inductors under the die for a four core LGA part. The driver circuitry is interleaved with the power switches in an array which minimizes parasitics to allow for very high switching frequencies. This also allows the size of the bridge to be easily scaled based on the current requirement and optimization points for each power domain.

Each FIVR domain is controlled by a FIVR Control Module (FCM). The FCM contains the circuitry for generating the PWM signals using double-edge modulation, as indicated in

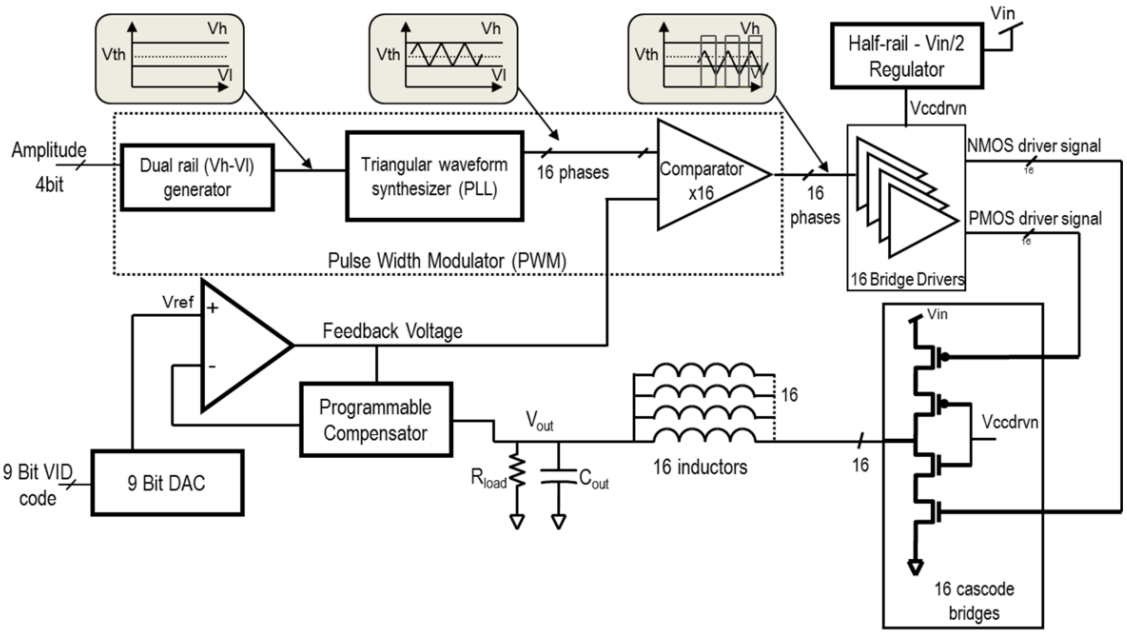

Figure 2. Simplified block diagram of the circuitry for a single representative FIVR domain

Fig. 2 by the dashed box. Separate circuitry not shown in Fig. 2 manages phase current balancing, and the resulting digital PWM signals are distributed from the FCM to individual bridges. The PWM frequency, PWM gain, phase activation, and the angle of each phase are all programmable in fine increments to enable optimal efficiency and minimum voltage ripple across a span of different operating points. Spread-spectrum is used for EMI and RFI (Radio Frequency Interference) control.

The FCM module also contains the feedback control circuitry (compensator). A high-precision 9-bit DAC generates a reference voltage for a programmable, high bandwidth analog fully differential type-3 compensator. Sense lines feed the output voltage back to the compensator. The endpoint of these sense lines is strategically placed to achieve minimum DC error and optimal transient response at an important circuit location in the domain. The compensator is programmed individually for each voltage domain based on its output filter, and can be reprogrammed while the domain is active to maintain optimal transient response as phase shedding occurs.

The key to making FIVR affordable was integrating the power devices directly into the microprocessor die. As with the power gate circuitry discussed in the introduction, the power switching circuitry for FIVR can be placed in small areas between major circuit blocks. The lower current handling of the die bumps makes the die bump area requirements for FIVR much larger than the actual die area required. Since FIVR is integrated into the microprocessor die, routing on the thick metal die layer allows extra bumps to be 'borrowed' from areas over other circuits, which avoids wasting any excess die area due to bump current limits. This makes the effective current density of FIVR 31A/mm^2, a 24x increase over the bump-limited 1.3A/mm^2 reported in [2].

B. Passives

In order to keep the buck output filter small enough to fit on the die and package it is necessary for FIVR to switch at a high frequency – 140 MHz in most cases. This allows the buck output filter inductors to be implemented using only the bottom metal layers of a standard flip-chip package. Power routing is constrained to the top layers of the package as a result, but the proximity of the inductors to the load ensures that minimal power is dissipated on these layers. The inductors are non-magnetic, i.e. Air Core Inductors (ACI). A representative ACI from an 8-phase domain of a product with an LGA package is shown in Fig. 3(b), including the connection points to the power switches, the DC current path through the inductor, and the connection of the inductor to the output plane. Package design rules allow the ACIs to be placed in to close proximity with one another. On a representative LGA package with four CPU cores, this allowed 59 inductors on 10 different voltage rails to be implemented in a 20mm x 8mm area. The package implementation also allows inductor designs to be customized on a per rail basis to meet efficiency, ripple, and transient response requirements.

Decoupling for the input rail is provided by a combination of ceramic package capacitors and on-die MIM capacitors [6]. The on package ceramic capacitors keep the output impedance of the input rail low from approximately 1 MHz to their self-resonance around 20 MHz. The MIM capacitors are on the die along with the power circuitry and provide high frequency decoupling, including at the switching frequency and its harmonics. Decoupling for the output rails is provided primarily by the MIM capacitors, which are sufficient to provide good transient response if wide bandwidth feedback control is used (see the results section). In some cases the MIM capacitors are supplemented with extra package ceramic capacitors. The comparatively low self-resonant frequency of the ceramic capacitors complicates the control loop design and

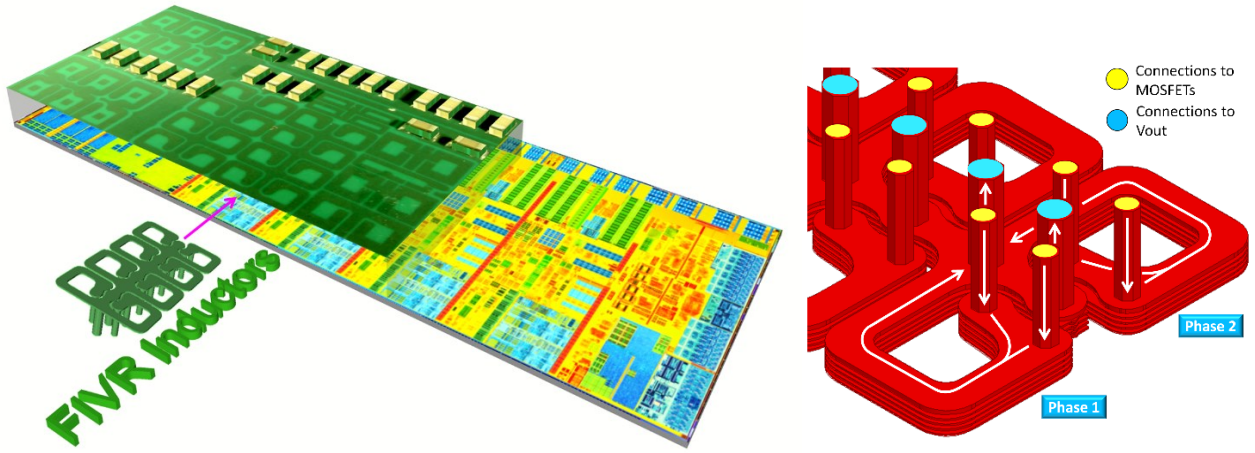

(a) (b)

Figure 3. (a) The bottom of an Intel® 4th generation Core™ microprocessor LGA package is shown along with along with a picture of the corresponding die. A group of eight FIVR inductors is pulled off to the side. (b) An enlarged 3D view of two FIVR inductors is shown with current flow arrows.

does little to attenuate voltage ripple, but the ceramic capacitors can provide a net transient response benefit if they are robustly connected to the output power plane.

C. System Control

In order to minimize losses from FIVR, a modified version of the PCU [1] dynamically configures each FCM based on the current activity level of the domain. The PCU turns each rail on or off based on activity, and specifies an output voltage target to support the desired frequency. It also optimizes the settings discussed in section II.A for the anticipated operating conditions. These settings include the number of active phases (i.e. phase shedding to improve light load performance), the compensator settings (to maintain optimal transient response as the number of phases changes), and the timing of switch drivers (to ensure zero voltage switching at light loads). This allows each FIVR domain to operate at near peak efficiency across a wide range of load conditions from retention to Turbo. An example of the benefit this provides is shown in Section IV.A.

III. Characterization and Performance Testing

Validating and optimizing a voltage regulator requires the measurement of key parameters such as voltage ripple, efficiency, power supply rejection ratio, transient response, and control loop stability margins. Due to FIVR's high level of integration and fast switching frequency many of these standard measurements are difficult or impossible to perform using off the shelf test equipment. For example, an IA Core voltage domain powered by a FIVR capable of supplying over 30A occupies less than 15mm^2 on the package, which is completely covered by the microprocessor die. This renders the attachment of an external load for a full current efficiency measurement impossible. This section describes some key Design For Test (DFT) features that are included in FIVR to allow accurate characterization.

A. Control Loop Transfer Function

To enable characterization of the control, a high frequency programmable signal generator is placed in the feedback network on every FIVR domain. The signal generator is activated in a test mode to inject a known, synchronized signal into the on-die feedback loop. By controlling the test feature and monitoring the output voltage on the package with an oscilloscope, the response of the control loop is directly measured. Repeated measurements are used to tune the compensator to achieve fast response and good stability margins.

B. Load Transient Response and Rejection Ratio

Microprocessors require a nearly constant DC voltage in the presence of large load transients. For characterization purposes, however, it is difficult to create a well behaved step load using only the execution of code in normal operation. The authors in [7] instead use a scheme called Integrated Frequency Domain Impedance Meter (IFDIM), which gates the clock network on and off at a fixed frequency creating a large load transient. The frequency is programmable, the magnitude of the transient can be precisely calibrated using DC measurements, and the load step is known to occur within one clock cycle, so the microprocessor itself is effectively turned into an alternating current load. This feature is included on every FIVR domain, allowing the transient response to a known load to be measured. FIVR domains are characterized across a wide frequency range at multiple operating points for both output impedance and output load coupling across domains.

C. Efficiency

As was previously mentioned, the level of integration makes it impossible to connect a high current load directly to the output of a FIVR rail. An additional complication is that the circuitry on the die cannot be disconnected from the FIVR output, so whenever FIVR is powered some extra current draw due to leakage results. This required the development of a new technique for accurately measuring efficiency. A brief summary of the method is given here. First, a procedure using an external low current load and FIVR operating in a test mode is used to calibrate the leakage. The completely ungated clock tree is then operated at varying frequencies to create a large effective adjustable DC output current. An iterative series of measurements is then used to precisely calibrate the total current drawn by the clocks and the leakage, which allows the efficiency to be measured, when combined with conventional

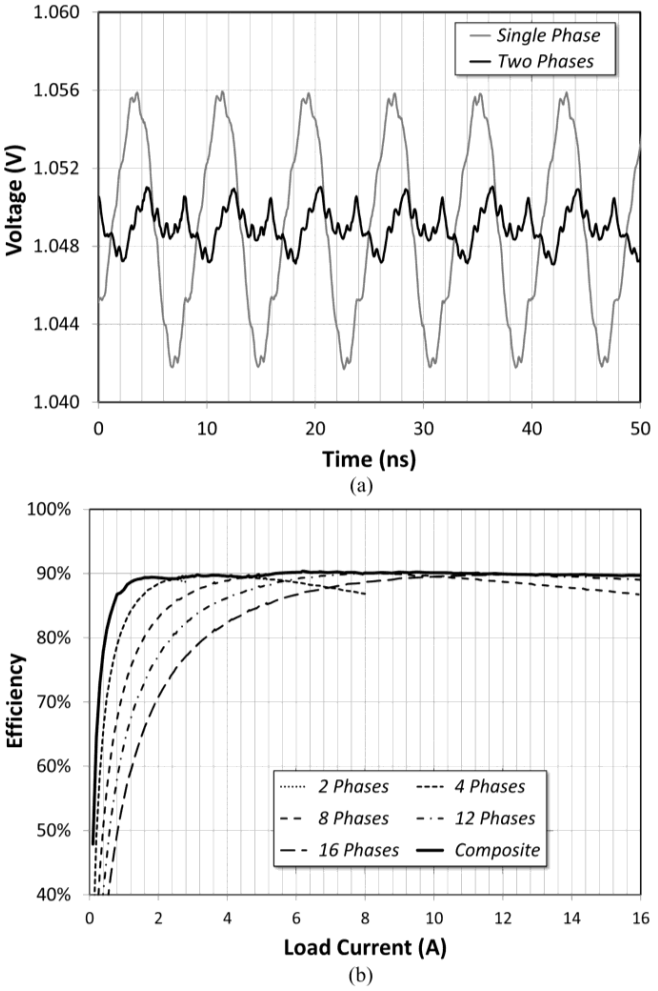

Figure 4. (a) Measured voltage ripple for a low noise domain for single phase and two phase operation (128 averages) (b) Measured efficiency for a voltage domain as a function of the number of active phases

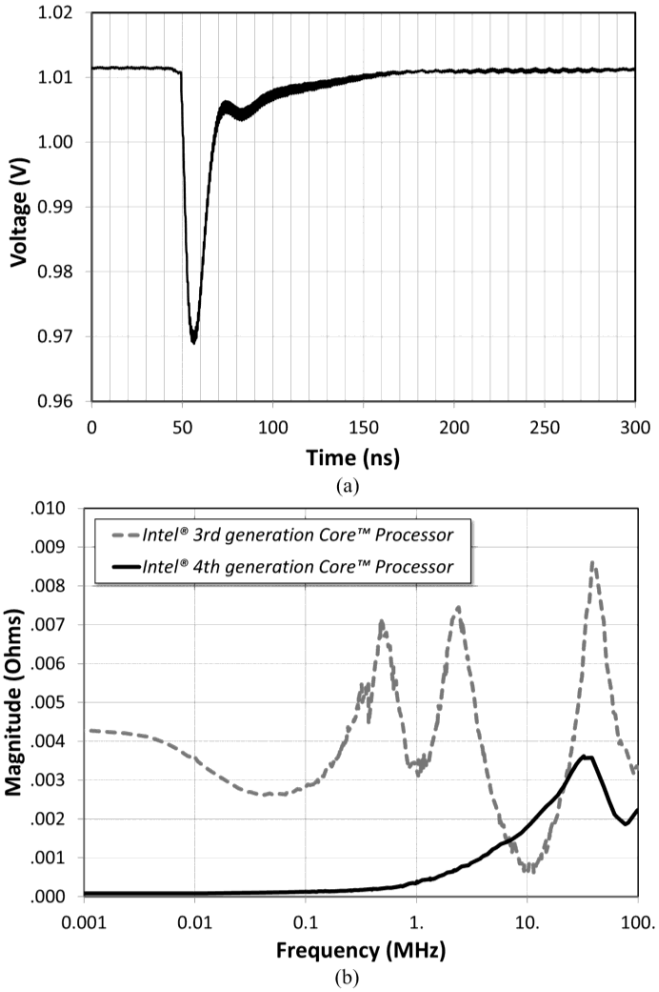

Figure 5. (a) Measured voltage droop on a graphics domain in response to an 8.5A step load (b) Comparison of the effective impedance profile for the graphics voltage domain on a 3rd generation Core™ microprocessor versus a 4th generation Core™ microprocessor

measurements of the voltage and the current at the output of the first stage regulator.

IV. Results

Due to the high switching frequency used, the performance of FIVR is sensitive to the layout of the die and the package which includes the inductors. Each combination of die and package is individually optimized and validated. The following section contains some key validation results from an Intel® 4th generation Core™ microprocessor with four microprocessor cores on an LGA package.

A. Measurements

Fig. 4(a) shows the voltage ripple for a low noise domain measured under the die near the connection of the ACI to the power plane. The measurement was averaged 128 times against the PWM clock (with spread spectrum clocking turned off). To achieve an accurate measurement, a controlled impedance differential sense line was routed on the package from the measurement location to a probe connection point with a matched termination. An active differential oscilloscope probe was then connected to the probe point. This ensures an accurate, wide bandwidth measurement is achieved, as opposed to probing the package power planes directly or using a single ended sense line, which can substantially attenuate the measurement over a distance as short as a few millimeters. In two phase operation less than 4mV (less than 1% of the voltage set point) of ripple is achieved with a rail driven by air core inductors well under 2mm^2 in area.

Fig. 4(b) shows the efficiency as measured using the procedure in section III.C for 1.70V to 1.05V conversion with the bridges configured for hard switching. The efficiency measurement has been repeated for varying numbers of phases, in each case showing a peak efficiency of approximately 90% at 0.75A/phase. By employing a phase shedding scheme it is possible to keep the efficiency of the domain within a few percent of the peak efficiency of the domain from 1A to 15A. This is managed by the PCU which can phase shed when the efficiency can be improved, but also has the intelligence to avoid phase shedding when it could be problematic, for example when a large load transient is possible.

The measured output voltage (averaged 128 times) during an 8.5A load step generated by the IFDIM feature on the graphics voltage rail is shown in Fig 5(a). The measurement

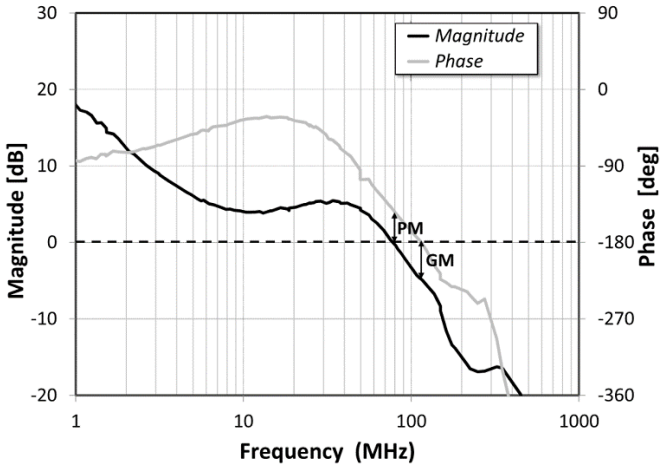

Figure 6. The measured open loop gain and phase of a FIVR showing 78MHz bandwidth with 40° phase magrin

was performed with a similar probing configuration to that used for the voltage ripple measurement. The combination of a high bandwidth feedback loop and on die decoupling keep the voltage droop under 50mV despite a rise time for current step of under 1ns (orders of magnitude faster than normal graphics circuit behavior). The main droop event lasts under 30ns, and the DC voltage is restored within 100ns. A nonlinear control feature saturates the duty cycle when a large transient is detected (not active in Fig. 5(a)). The feature was found to provide up to a 25% reduction in voltage droop for step loads, but the benefit is significantly reduced for certain aperiodic load patterns. The effective output impedance profile for the same rail is shown in Fig. 5(b). The peak impedance demonstrates the fast bandwidth of the compensator. Because the inductors are located immediately below the actual area of the die that consumes current, the DC and low frequency load line is virtually zero. The figure also shows the impedance profile for an Intel® 3rd generation Core™ microprocessor graphics rail, which is powered by a platform VR. For this rail, a DC load line is required due to the distance between the VR and the die. Several resonant peaks occur from the various stages of decoupling capacitors on the motherboard and package, and the parasitic inductance between them and the actual point of current consumption on the die.

Fig. 6 shows the open loop transfer function for a FIVR domain, measured using the signal generator DFT. This rail demonstrates a unity gain bandwidth of 78 MHz while still maintaining 40° phase margin. Robust compensator circuit design and very small propagation delays were necessary to achieve this bandwidth, which, in turn, was required to maintain good transient response on rails with limited output capacitance. The high bandwidth also enables fast voltage transitions. A FIVR rail turning on and turning off are shown in Fig. 7. Both transitions are programmed to about half a microsecond for a full range transition – two orders of magnitude faster than a typical platform-based solution. The fast ramp rate translates into power savings for the system, as the voltage rails can be turned on, used, and turned off again almost instantly.

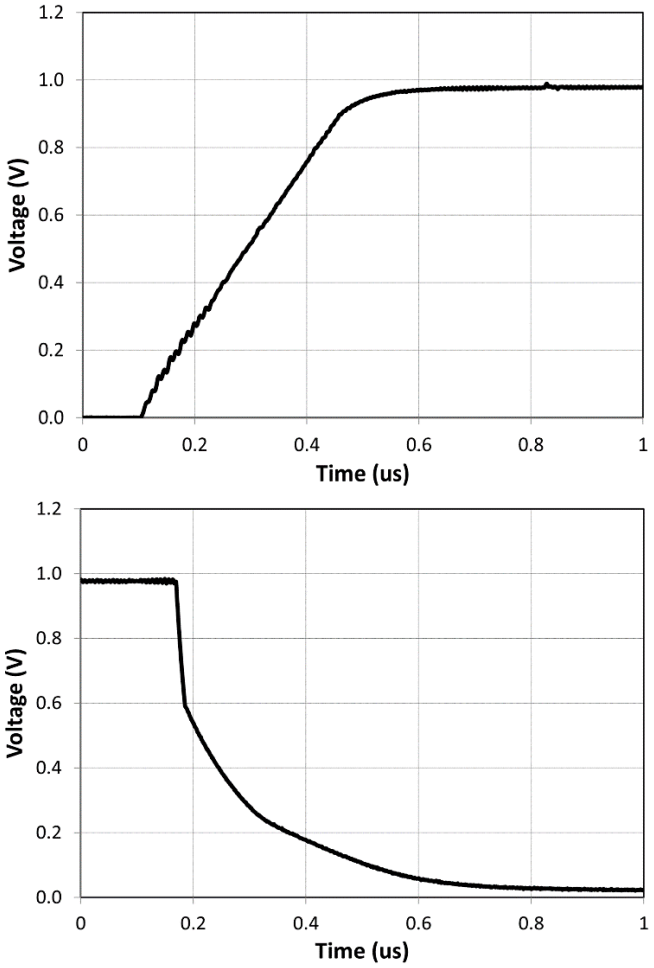

Figure 7. A FIVR rail ramping to its voltage set point from fully off, and then turning off again. Voltage transition times are programmable, but typically set for half a microsecond for a 1V transition.

A large number of additional measurements are taken for validation purposes that are not shown here due to space constraints. These include the output impedance of the V_{in} rail, audio susceptibility measurements, and the coupling noise due to load transients from one rail to another (particularly from very high current domains such as core and graphics to low current system agent domains). EMI/RFI characterization is also performed.

B. Comparison to Previous Work

Table I contains a comparison to previous works discussed in the introduction. FIVR operates at a higher switching frequency than previous works, which is enabled in part by very good gate charge characteristics for the MOSFETs. This allows up to 90% efficiency at a common conversion ratio.

V. FIVR IMPACT TO PRODUCTS

Battery life improvement: Sufficient battery life for a complete work day has long been desired from mobile products. FIVR, combined with power management architecture improvements, has enabled this for Intel® 4th generation Core™ products. Increases of well over 50% have been widely reported (for example, [8] and [9]). FIVR's battery life benefit comes by several means:

TABLE I. COMPARISON OF FIVR TO PREVIOUSLY REPORTED INTEGRATED VOLTAGE REGULATORS

Parameter	G. Schrom et al., 2010 [2]	T. DiBene et al., 2010 [3]	N. Sturcken et al., 2012 [4]	This Work
Process node	130 nm	90 nm	45 nm	22 nm
Switching Frequency	60 MHz	100 MHz	80 MHz	140 MHz
Unity Gain Freq	5 MHz	Not Published	Not Published	80MHz
Efficiency	85-88%, 3.3V:1.0V	76%	83%, 1.5V:1.0V	90%, 1.7V:1.05V
Total Output Imax capability	50 A	Limited by first stage and thermals (Up to 400 A)	1.2 A	Limited by first stage and thermals (Up to 700 A)
Imax/VR die area	1.3 A/mm^2	8 A/mm^2	1.7 A/mm^2	31 A/mm^2
Voltage rail count	4	20	1	8 to 31
Phase count	16	320	4	49 to 360
Integration level	MCM[a]	MCM[a]	Integrated into network die	Integrated into CPU die
Inductor technology	Package trace, & magnetic discrete	Magnetic thin-film on VR die	Discrete wire-wound air core	2D array of package trace
Capacitor type	Ceramic package caps	Ceramic package caps	Die Cap	Die Cap - MIM
Cout per Max Amp	2000 nF/A	not published	15 nF/A	7 nF/A

[a] MCM – Multi Chip Module – the active circuitry is on a separate die assembled on the same package

- Standby current historically consumes a large fraction of the battery's stored energy. FIVR's fast bandwidth allows low frequency supply noise to be rejected, resulting in up to a 90% reduction in decoupling requirements. This allows both the first and second stages of regulation to be power cycled much faster than on previous products, enabling new deep sleep states with up to 20x lower standby power. With the lowered capacitance, power expended, and time wasted entering and exiting the states is similarly reduced. Reduced sleep-state entry/exit time also saves power by increasing sleep-state usage.
- FIVR's fast control loop and integration into the package result in one tenth the peak impedance of prior solutions (see Fig. 5(b)) in the sub-MHz stimulus range most relevant to the graphics architecture. The resulting low frequency supply noise reduction improves power at a given performance by up to 30%.
- FIVR increases the number of voltage rails, allowing each domain to be set at the minimum possible voltage that supports error-free operation, reducing both leakage and dynamic power.
- Replacing multiple high current voltage regulators on the motherboard with a single first stage regulator reduces the PCB footprint of the power delivery solution. This extra space can be used for a larger battery, with some examples demonstrating up to 10% growth.
- Trimming FIVR together with the microprocessor removes manufacturing guard-bands normally required to ensure that every VR will work with every CPU.

Increased available peak power: An illustrative example shows how FIVR can increase the peak power available to the microprocessor. A typical mobile processor platform using the prior generation power delivery scheme had two 30A, 1.1V VRs providing 33W for cores and 33W for graphics. Using the same power FETs and inductors for the FIVR's 1.8V input VR, the part has 108W power rail (30A/phase * 2 phases * 1.8V), which can be dynamically allocated to a combination of FIVRs by the PCU. For core-only workloads, nearly the entire 108W can be allocated for the cores, increasing the available power ceiling by 3x. For graphics workloads, 36W can be partitioned to the cores with the remaining 72W going to graphics – more than double the power available from the 33W platform VR. Because power consumption scales as CV^2F and frequency scales with voltage, the increase in available power could be used to operate the graphics at up to 26% higher frequency than possible with the platform VR. A similar calculation yields a 44% higher core frequency in the core-only scenario. The duration of these scenarios is limited by the thermal capabilities of the platform, but translates into improved speed in many real scenarios.

Decreased power at a given performance: Intel's® Iris™ Pro graphics uses FIVR's higher available power to deliver high end graphics. FIVR's high unity gain bandwidth presents less than a tenth the peak output impedance provided by the prior generation's platform VR in the sub-MHz range important to the graphics load (see fig. 5b). Because FIVR has doubled the graphics power ceiling, few (if any) of our shipping parts would fit within the bounds of the older generation's platform. The higher currents typically imply hundreds of millivolt droops on the older platforms. The combination of high currents with high impedance peaks yields a hypothetical power tax in the 20-30% range (assuming one could, and actually would, ship these high current levels into the old platforms). FIVR avoids that tax.

Improved product flexibility and scalability: FIVR's ability to add voltage rails onto a common shared input rail without package growth or even platform changes brings significant flexibility and modularity into the design space that was not available before. New voltage rails can be added as needed, without any platform change. This ability allowed us to introduce the Iris™ Pro graphics into standard platforms even though new rails were needed for the EDRAM and its high speed OPIO link.

Platform size and cost reduction: Since four platform VR controllers are eliminated, along with the associated decoupling caps, power FETs and inductors, there's a clear platform size and cost advantage. FIVR's total platform BOM cost reduction is expected to be several billion dollars over the product lifetime. The power inductors feeding power to the CPU often show up in the critical thickness cross-section of small form-factor laptops and tablets, and trading FIVR's total component count reduction for a thickness reduction is straightforward. A platform phase count increase results in lower current per phase, and the lower phase current can be satisfied with a lower profile set of inductors.

In the prior generation platforms, some dual-sided PCBs have tall components like ICs and inductors on the primary side and lower profile discrete components on the secondary side. Often the secondary-side components are mostly high frequency decoupling, located immediately underneath front-side ICs, with sparsely populated areas between. In such cases, FIVR eliminates most of the secondary-side components, and frees up space on the primary side to accommodate the rest. The resulting populated PCB thickness is reduced by the height of the tallest removed backside components.

In small systems, the platform size tends to limit the feature set, leading to fewer connectivity options, smaller storage space, etc. FIVR's platform size reductions can provide more space to implement these features.

VI. Conclusion

Consumers expect every generation of mobile computer products to have more compute power, thinner and lighter form factors, and longer battery life than the last. The 4th generation Intel® Core™ power architecture using FIVR provides improvements in all three of these areas. To the author's knowledge, this is the first consumer product to make use of integrated switching regulators on this scale. Furthermore, FIVR's performance is improved versus previously reported prototypes. The authors therefore feel that FIVR is an important advancement in the field of power electronics.

Acknowledgment

We would like to acknowledge the following FIVR team members who were not already listed as authors: the FIVR silicon design team including George Geannopoulos, Keith Hodgson, Narayanan Raghuraman, Alex Lyakhov, Michael W Rogers, Ravi Vunnam, Lan D Vu, Mark S Milshtein, Chiu Keung Tang, Hong Yun Tan, Seh Leon Goh, Samie Samaan, Narayanan Natarajan, Rajan Vijayaraghavan, Ashish Khanna and Munish Chauhan; top-level integration by Pankaj Aswal; modeling and implementation development by Doug Huard and Alex Waizman; layout studies, inductor development and package designs by John Smith, Brad Larson, and Huong Do; mask design by Neal Tanksley and Galyna Burenkova; test and manufacturing support from RJ Hayes and Arun Krishnamoorthy; modeling support by Alex Levin and Anne Augustine.

References

[1] S. Gunther, A. Deval, T. Burton and R. Kumar, "Energy-Efficient Computing: Power Management System on the Nehalem Family of Microprocessors," *Intel Technology Journal,* vol. 14, no. 3, pp. 50-65, 2010.

[2] 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.

[3] 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.

[4] N. Sturcken, et al., "A switched-inductor integrated voltage regulatorwith nonlinear feedback and network-on-chip load in 45nm SOI," *IEEE Journal of Solid-State Circuits,* vol. 47, no. 8, August 2012.

[5] G. Schrom, et al., "A 100 MHz Eight-Phase Buck Converter Delivering 12 A in 25 mm^2 Using Air-Core Inductors," in *Proc. 22nd Annu. IEEE Applied Power Electronics Conf.*, 2007.

[6] C. Auth, et al., "A 22nm High Performance and Low-Power CMOS Technology Featuring Fully-Depleted Tri-Gate Transistors, Self-Aligned Contacts, and High Density MIM Capacitors," in *2012 Symposium on VLSI Technology*, Honolulu, HI, 2012.

[7] A. Waizman, M. Livshitz and M. Sotman, "Integrated Power Supply Frequency Domain Impedance Meter (IFDIM)," in *13th Conference on Electrical Performance of Electronic Packaging*, Portland, OR, 2004.

[8] A. L. Shimpi, "Is Haswell Ready for Tablet Duty? Battery Life of Haswell ULT vs Modern ARM Tablets," 22 July 2013. [Online]. Available: http://www.anandtech.com/show/7117/haswell-ult-investigation. [Accessed 15 November 2013].

[9] R. Baldwin, "How the Haswell Chip Makes the New MacBook Air Last 12 Hours," 10 June 2013. [Online]. Available: http://www.wired.com/gadgetlab/2013/06/haswell-mba/. [Accessed 15 November 2013].