A Comprehensive Evaluation of Power Delivery Schemes for Modern Microprocessors

Jawad Haj-Yahya[1,2] Efraim Rotem[3] Avi Mendelson[4] Anupam Chattopadhyay[1]
[1]Nanyang Technological University (NTU), Singapore
[2]Institute of Microelectronics, Agency for Science Technology and Research (ASTAR), Singapore
[3]Intel Corporation, Israel
[4]Technion - Israel Institute of Technology, Israel

***Abstract*—The continuous quest for energy-efficient computing has led towards the adoption of fine-grained controls in processor sub-systems, of which power delivery network is the most prominent one. Recent industry trends reflect a shift towards on-chip, integrated voltage regulator (IVRs) to that effect. We undertake a thorough and quantitative evaluation of different power delivery networks for modern microprocessors. In contrast to the current trend, we conclude that IVR schemes perform worse compared to the conventional off-chip voltage regulator scheme. Further, we present studies on diverse workloads and Thermal Design Points (TDPs) to highlight the importance of workload-specific power delivery scheme. To the best of our knowledge, this is the first comprehensive study across processors' TDPs and workloads.**

***Index Terms*—**Energy-Efficiency, Integrated Voltage Regulator, Power Delivery, MBVR ,FIVR

I. Introduction

The increasing demand for energy efficient computing, across all market segments, underpins the need for optimizing processors' sub-systems, e.g., high performance processors may consume tens to hundreds of amperes at sub 1V [1]. This demand makes the Power Delivery Network (PDN) an extremely important feature to tackle both the thermal and electrical issues, particularly at advanced technology nodes. In addition, optimizing the PDN and boosting its efficiency helps to reduce the average power, which in turn, increases the battery-life of mobile devices and/or reduces the electricity bills for data-centers.

In recent times, several design houses have added Integrated Voltage Regulators (IVRs) into their product portfolio for mobile and server segments, e.g. Intel Fully Integrated Voltage Regulator [1] (FIVR). Such integration accepts a single input at higher voltage and therefore reduces the current consumption and number of different voltage inputs into the chip, resulting in a smaller motherboard area. While the board area benefits and reduction in current consumption from off-chip voltage regulator may be obvious, it is not clear if the migration towards IVR is indeed fruitful from the overall system perspective, and considering all product segments.

Motivation: Previous studies [1]–[4] claimed that the adoption of IVR increases both battery-life and performance (such as at Intel ® 4th generation Core™ microprocessors code name Haswell [1]), but to the best of our knowledge, there is no comprehensive study that evaluated the tradeoffs of IVR across various microprocessors' TDP and workload scenarios. This forms the motivation of this work.

Architecture: Our study is targeting client products; such as tablets, laptops and desktops; the processor architecture is similar to the one discussed at Intel's paper [1]; the system has two *Cores*, *Graphics* engine (GFX), *Last-Level-Cache* (Ring), *System-Agent* (SA) and *IOs* as depicted in Fig. 1. Variants of these schemes are widely used at modern client processors.

Contributions: In particular, this paper makes the following contributions:

- A thorough evaluation of the power and performance of IVR based power-delivery schemes across various microprocessor segments and workloads.
- Segment-specific analysis of power consumption for two main power delivery schemes, namely direct Mother-Board VR (MBVR) and Integrated Voltage Regulator (IVR).
- Studies on the optimization of power delivery parameters and its impact on the power and performance of power delivery schemes.
- An in-depth analysis of the various components of power delivery losses at modern processors.

II. BACKGROUND

Before deep-diving into building the power models, at this section we give a background on IVR and voltage regulators power losses. In addition, we explain about some of the parameters that affects system level efficiency of PDN, such as, tolerance band and load-line.

A. Integrated Voltage Regulators

Recent trends in power delivery is to bring the power supply closer to the load [10]. A prominent example is Intel Fully Integrated Voltage Regulator (FIVR) [1] - a technology that deploys high current switching regulators integrated into the die and package. There, a two-stage voltage regulation is applied. The first stage Voltage-Regulator, which is on the motherboard, converts from the Power Supply Unit (PSU) or battery voltage (7.2-20V) to typically less than 2V, which is distributed across the microprocessor die. The second conversion stage comprises of several (depending on the product) IVRs, which are synchronous multi-phase buck converters.

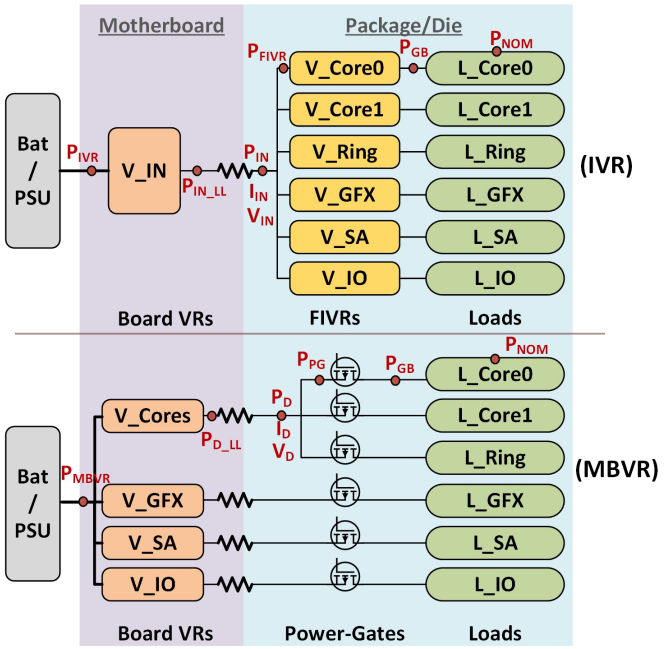

Fig. 1. Two Power-delivery schemes used at modern SoCs: **(IVR)** uses one motherboard VR and six FIVRs. It's used at Intel's® 4th and 5th generation Core™microprocessors ([1] and [5]. **(MBVR)** uses four motherboard VRs and three on-die power-gates. It's used at Intel's® 2nd, 3rd, 6th and 7th generations Core™microprocessors ([6]–[9])

B. Buck Converter Power Losses

The non-idealism of the power devices, such as switches, diodes, inductors, account for the bulk of the power losses in the converter. Both static and dynamic power losses occur in any switching regulator. Static power losses include I^2R (conduction) losses in the wires or board traces, as well as in the switches and inductor, as in any electrical circuit. Dynamic power losses occur as a result of switching, such as the charging and discharging of the switch gate, and are proportional to the switching frequency. The ratio of the total output power of the voltage-regulator to the total input power known as Efficiency (η) . The efficiency of Buck Converter is not constant, rather is a function of the load current (I_{out}) and the input voltage (V_{in}).

C. Tolerance Band

VR tolerance band (TOB) [11] defines the maximum voltage variation for the Voltage Regulator (VR) across temperature, manufacturing variation, and age factors (e.g. $V_{TOB} = 25mV$). The standard VR tolerance band can be sliced into three main categories: controller tolerance, current sense variation, and voltage ripple.

D. Load line

Load line or adaptive voltage positioning (AVP) [12] is a model that describes the voltage and current relationship under a given system impedance (R_{LL}). This relationship is defined in equation 1, where V_{cc} is the voltage at the load and V_{in} is the input voltage to the system. From equation 1, we can see that the voltage at the load input decreases when the current of the load increases. Due to this phenomenon, the input voltage (V_{in}) is increased to a level that keeps the voltage at the load (V_{cc}) above the minimum functional voltage under the highest possible power consumed by the load[1]. At typical application load, the voltage drop is smaller, thereby resulting with higher voltage than actually needed and incurring quadratic power losses. At the FIVR [1] technology, because the inductors are located immediately below the actual area of the die that consumes current, the load-line is very small (negligible). For the main power rail - which is powered by a motherboard VR - a load-line is required due to the distance between the motherboard VR and the die.

$$V_{cc} = V_{in} - V_{TOB} - R_{LL} \cdot I_{cc} \quad (1)$$

III. System Power and Performance Models

We have built a system power and performance models in order to compare two power-delivery schemes (IVR) and (MBVR) as depicted in Fig. 1 and explained in what follows.

IVR scheme: Power-delivery scheme that uses one motherboard VR (V_IN) and six different on die/package FIVRs, this scheme is typical to high-end client processor SoC such as the Intel's® 4th and 5th generation Core™ microprocessors ([1] and [5].

MBVR scheme: Power-delivery scheme that uses four motherboard VRs (V_Cores, V_GFX, V_SA and V_IO) and three on-die power-gates for the Core0/1 and Ring domains, this scheme is intended for a high-end client Processor SoC with non-integrated VRs such as Intel® 2nd, 3rd, 6th and 7th generations Core™microprocessors ([6]–[9]).

To compare the two schemes, we have developed two power models, one for each scheme. The power models reflect the overall power that is consumed from the battery/PSU, denoted by P_{IVR} and P_{MBVR} as shown in Fig. 1. The modeled power begins from the nominal power at each domain (P_{NOM}).

Models Validation: The two power models were validated with experimental data on our lab, showing high modeling accuracy (within 5%) for both models. The methodology used is similar to Intel's Blizzard [13] power/performance modeling, space limitations preclude a detailed discussion on the validation process.

A. IVR Power Model

In order to calculate the total power of the IVR scheme; denoted by P_{IVR}; at the first stage, we calculate the P_{GB}, which is the power after applying the voltage guard-band (V_{GB}) to compensate on *TOB*. The *leakage* and *dynamic* power consumption parts scale differently with voltage increase. The dynamic power part is proportional to the square of the voltage, while the leakage power scales exponentially with voltage and depends on several parameters such as threshold voltage, temperature, and other design and fabrication characteristics [14]. As an approximation, we have used

[1]This power is normally referred to as power-virus (PV).

a scaling model based on polynomial curve fitting, where leakage power scales polynomially with supply voltage. This is further validated with measurements on a commercial processor (Intel® Core™ i7-6660U Processor). Assuming same temperature, the leakage power scales by the power of δ ($\sim$2.8 in our case) proportional to voltage scaling. We assume leakage fraction (F_L) of 22% similar to [15]. Therefore, the P_{GB} power can be calculated as shown in equation 2.

$$P_{GB} = P_N \cdot \left[F_L \cdot (\tfrac{V_N+V_{GB}}{V_{NOM}})^{\delta} + (1-F_L) \cdot (\tfrac{V_N+V_{GB}}{V_{NOM}})^2 \right] \quad (2)$$

The P_{FIVR} is the power after taking into account the FIVR losses. Given the FIVR efficiency η_F, P_{FIVR} can be calculated using equation 3. While the efficiency of the FIVR (η_F) is a function of the load and the input voltage.

$$P_{FIVR} = \frac{P_{GB}}{\eta_F} \quad (3)$$

Due to the voltage drop on the load-line impedance (R_{IN_LL}), we need to compensate on this drop by raising the on-board VR output voltage. Furthermore, the voltage guard-band needs to account for the maximum possible voltage drop, which is attained when the processor consumes the maximum possible power (realistic power-virus [5]). Assuming that P_{IN} is the sum of the power of all domain (obtained from equation 3) connected to V_{IN} voltage-regulator, the corresponding calculation for the voltage and power after taking into the load-line voltage drop is shown at equations 4 and 5, respectively. The total power (P_{IVR}) consumed from the battery/PSU is obtained by dividing the output power on the on-board VR by its efficiency (η_{IN}) as shown in equation 6, while the efficiency of the on-board VR is a function of the load and the input voltage as explained in subsection II-B.

$$V_{IN\_LL} = V_{IN} + \frac{P_{power\_virus}}{V_{IN}} \cdot R_{IN\_LL} \quad (4)$$

$$P_{IN\_LL} = V_{IN\_LL} \cdot I_{IN} = V_{IN\_LL} \cdot \frac{P_{IN}}{V_{IN}} \quad (5)$$

$$P_{IVR} = \frac{P_{IN\_LL}}{\eta_{IN}} \quad (6)$$

B. MBVR Power Model

Using the same approach as for IVR, we calculate the total power (P_{MBVR}) for the MBVR scheme. At the first step we calculate the P_{GB} by applying the additional power due to the increase of the voltage caused by the TOB, essentially following the equation 2. For the domain with power-gates ($L\_Core0/1$ and L_Ring) we need to take the additional voltage drop on the power-gate themselves ($V_{PG} \approx 10mV$). The power consumption (P_{PG}) after taking into account the power-gates voltage guard band can be calculated using equation 2. There, P_D represents the total power of the group of domains that share the same VR (i.e. $\{Core0, Core1, Ring\}$, $\{GFX\}$,$\{SA\}$,$\{IO\}$). P_{D_LL} is the power of a domain after taking into account the voltage drops on the load-line impedance (R_{LL}) for the specific domain, for which, the calculation is shown in equations 7 and 8.

$$V_{D\_LL} = V_D + \frac{P_{domain\_max\_power}}{V_D} \cdot R_{D\_LL} \quad (7)$$

$$P_{D\_LL} = V_{D\_LL} \cdot I_D = V_{D\_LL} \cdot \frac{P_D}{V_D} \quad (8)$$

The total power (P_{MBVR}) consumed from the battery/PSU is obtained by summing the effective power of each domain, which can be calculated by dividing the output power of each on-board VR by its efficiency (η_D) as shown in equation 9.

$$P_{MBVR} = \sum \frac{P_{D\_LL}}{\eta_D} \quad (9)$$

C. Performance Model

For power constrained system, the SoC average power is limited by TDP. Therefore, at iso-power (TDP), the PDN scheme that have higher power losses must run at lower frequency in order to keep the power within TDP limit. Lower operation frequency means lower performance (higher execution-time). The performance scales with frequency proportionally to the performance-scalability[2] factor of the workload, whereas performance scales at higher rate for workloads that are less bounded on memory (have high performance scalability). Our performance model maps the difference in power at both PDN schemes into the equivalent difference in frequency at various TDPs.

The methodology used for performance model is similar to Intel's Blizzard [13] power/performance modeling. Our model predicts the potential increase in frequency of a Core or Graphics for a given power budget, for example a $200mW$ power budget (due to less power losses compared to the other PDN) enables increasing the core frequency by $100MHz$ or graphics frequency by $50MHz$.

The model was built in our lab using a platform with a processor equivalent to Intel® Core™ i7-6660U with DDR4-2133 at 80°C, the platform enables configuring the TDP of the processor (Configurable TDP feature [16]) . For Core performance model we ran highly scalable core workload from SPEC [17] ($456.hmmer$). While for Graphics performance model we used 3DMARK [18] workload.

To build the Core performance model, we ran the Core workload at the minimal frequency and changed the core frequency by steps of $100MHz$ until the frequency reaches the TDP frequency, whereas, at each frequency point we measured the change in power due to the additional $100MHz$ step in frequency. Similarly, for Graphics performance model, the graphics workload was ran and the frequency change was applied to the graphics frequency. This process was repeated for multiple TDPs.

[2]Performance-scalability is performance and frequency correlation; a ratio of 1 means doubling frequency results in doubling performance

TABLE I
DETAILED COMPARISON BETWEEN IVR AND MBVR TOPOLOGIES FOR 15W TDP SEGMENT

	Description	Symbol	Core0	Core1	Ring	GFX	SA	IO
Inputs	Nominal power	P_N (W)	4.00	4.00	2.00	0.00	0.80	1.20
	Max Power	P_{max} (W)	7.86	7.86	3.57	0.00	1.60	3.00
	Nominal Voltage	V_N (V)	0.65	0.65	0.65	0.65	0.65	0.95
	Voltage guard band	V_{GB} (V)	0.02	0.02	0.02	0.02	0.02	0.02
	Fraction of Leakage	F_L	22%	22%	22%	40%	22%	22%
	Additional power after TOB guard-band	P_{GB}(W)	0.27	0.27	0.14	0.00	0.05	0.06
Integrated VR (IVRs)	FIVR Efficiency	η_F	85%	85%	84%	-	86%	87%
	Power after FIVR efficiency looses	P_{FIVR}(W)	5.02	5.02	2.51	0.00	1.00	1.44
	SoC input Voltage	V_{IN} (V)	1.80					
	Total Vin current at package pins	I_{IN} (A)	8.34					
	Power-Virus	P_V (W)	18					
	SoC input load-line	R_{IN-LL}(mΩ)	2					
	Vin voltage after increase by LL	V_{IN_LL} (V)	1.82					
	SoC power after LL additional voltage	P_{IN_LL}(W)	15.18					
	IN VR Efficiency	η_{IN}	90%					
	Power consumed from Battery/PSU	**P_{IVR} (W)**	**16.86**					
Motherboard VR (MBVRs)	Power gate resistance	R_{PG}(mΩ)	1	1	1	1	1	1
	Voltage increase due to PG	V_{PG}(V)	0.01	0.01	0.01	0.00	0.00	0.00
	Voltage of domain after VPG and VGB	V_D(V)	0.68	0.68	0.68	0.67	0.67	0.97
	Power after losses on the power-gates	P_{PG} (W)	0.16	0.16	0.04	0.00	0.01	0.01
	Total power per domain	P_D (W)	11.03			0.00	0.86	1.26
	Total current at domain	I_D (A)	16.17			0.00	1.28	1.30
	SoC input load-line for domain	R_{D-LL} (mΩ)	2			0.01	0.01	0.00
	Domain Voltage increase due to LL	V_{DLL} (V)	0.73			0.67	0.68	0.98
	Domain power after LL voltage increase	P_{DLL} (W)	11.74			0.00	0.88	1.27
	Domain VR Efficiency	η_D	85%			-	85%	87%
	Domain Power after VR efficiency looses	P_{DEFF} (W)	13.81			0.00	1.03	1.46
	Power consumed from Battery/PSU	**P_{MBVR} (W)**	**16.31**					

IV. EXPERIMENTAL RESULTS

In order to evaluate the energy efficiency of the two power delivery schemes, it is imperative to vary the configurations, and also present realistic workload scenarios that can be supported in the modern processors. Note that, recent Core™ Intel microprocessors can scale from near $3W$ of Thermal Design Point (TDP); that exist at passive cooled small systems; up to $95W$ TDP, which is used at high-performance desktop computers [16]. Consequently, the following scenarios are decided upon.

- A processor with TDP near $15W$ for Core workload. In this study we show detailed calculation of both voltage regulators and corresponding power consumption.
- Shmoo plot of the power of the two schemes processors with TDP ranges from $3W$ to $25W$ for Core and Graphics workloads. There, we show how the two schemes behave when running core and graphics performance workloads.
- Compare average power scenarios of the two schemes. In this third study, we examine the two schemes with light workloads that are normally used for system average power evaluation.

The projected power and performance scaling are based on Intel® Core™ i7-6660U Processor with DDR4-2133 at 80°C and 50°C for performance and average power workloads respectively.

A. Evaluation of 15W TDP processor

Here, we evaluate in details the power consumed by a 15W TDP processor in both schemes when running performance workload. The inputs to both models are nominal power (P_{NOM}) that is obtained when running performance benchmark on both cores. We choose $456.hmmer$ from SPEC [17] while other workloads with the same scalability have nearly the same power footprint. The calculation breakdown from our model is shown at Table I .

Our model takes as inputs the nominal power at each one of the six domains, the 15W TDP scenario is core based scenario hence the Graphics is off. In addition, the model takes the maximum possible power at each domain, the nominal voltage and the leakage fraction. The model uses equations (1) to (10) in order to calculate the power at each one of the two schemes. The calculation starts from nominal power until we arrive to

the total power that is consumed from the power sources (as shown in Figure 1 from right to left).

The model first calculates the power after TOB voltage guard-band (which is common for both schemes), and then based on the scheme we follow the calculation using the relevant equation. As previously mentioned, the efficiency of both the IVR and the motherboard VR is a function of the input voltage and the load current. Table I shows that the power consumed by the IVR scheme is higher by $0.55W$ compared to the MBVR scheme, in power constrained system (such as thin and light laptop or tablet) this additional power (due to PDN inefficiency) is directly translated to lower performance as we need to run the cores/ring at lower frequency to keep the power below TDP power as explained in section III-C. Based on our performance model, this additional 0.55W is translated to 7% less performance on $456.hmmer$[3] workload.

B. Evaluation of different segments for Graphics and Core workloads

For this case study we have evaluated different target segments along with Core and Graphics (GFX) workloads. For Core workloads, the power budget is allocated to the cores and ring (cache); after decreasing the rest of the SoC power (SA, IO, etc.); while graphics is idle. On the other hand, at GFX workloads, normally between 10% and 20% (depends on TDP) of the power budget is allocated to the cores, while graphics takes the rest. We have chosen 3DMARK [18] Graphics (GFX) workload and the $454.Calculix$ workload from SPEC [17], hybrid devices (tablet/laptop), laptops and desktops.

Fig. 2 shows the result obtained by our power and performance models described at section III , where the bars represent the total power at both schemes for GFX and Core workloads for $3, 6, 8, 15, 20, 25$ Watts TDPs, while the lines represent the performance difference between both schemes ($Perf_{MBVR} - Perf_{IVR}$) when translating the power difference into performance for these power-limited systems.

In addition, Fig. 2 shows that the MBVR scheme is more power efficient than the IVR scheme for TDP below $15W$ in case of Graphics workload, while at the same TDP the MBVR scheme is more efficient for the Core workload. From the graph at Fig. 2, it can be observed that the for high TDP (e.g. above $20W$) the scheme with Integrated voltage regulator (IVR) is more efficient while for segments with low TDP (e.g. below $15W$) the scheme with direct connection to motherboard VRs is more energy efficient.

Power losses breakdown: in order to understand better the power difference between the two schemes, Fig. 3 shows power-losses breakdown at various TDP segments for core workload at each one of the two schemes. The bars show the percentage contribution of each power loss component (TOB, Power-Gates, IVR in-efficiency, losses due to Shared-Voltage, Load-Line and I^2R and MBVR in-efficiency), while lines show the total system power at each scheme. It can be observed that Power-losses on power-gates (PGs), load-line(LL) and I^2R grow significantly at higher TDP (higher current) for the MBVR scheme making it less efficient at these segments compared to IVR, while the IVR scheme is less affected by the current due to the higher input voltage to the chip.

[3]This workload is highly scalable, while memory bound workloads have less performance degradation.

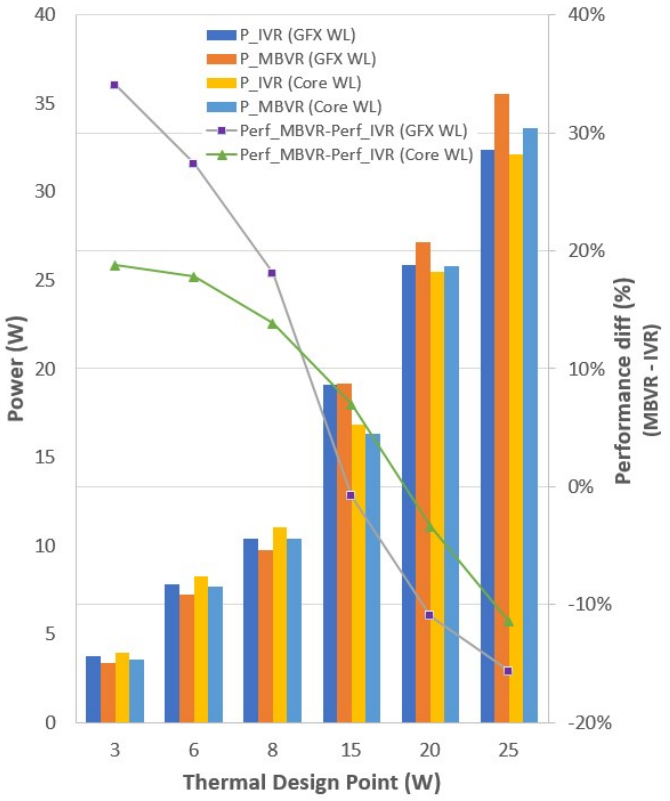

Fig. 2. Power consumption and performance difference at various TDP segments for Graphics (GFX) and Core workloads (WLs) at each one of the two schemes.

C. Evaluation of average power workloads

In this case study we evaluate the two schemes with average power workload. The previous workloads were performance workloads, which have nearly 100% active state residency (i.e. 100% Package C0 state). While average power workloads, such as video-playback and video-conferencing, have idleness periods that enables the SoC to enter to deep package C-states such as package-C7 state [19]. In addition, there are periods at average workloads where the cores and the graphics sub-systems are idle while only the path to the memory is open, this state is called package-C2, for example when one of the components at the SoC want to access the memory with Direct-Memory-Access (DMA) and fill up a local buffer. For this case study we have chosen Video-playback and Video-Conferencing workloads as representative workload for evaluating average power of the system, noting that the other average power workloads have nearly in the same range package c-states residencies. For our modeled system, video-playback have roughly 15%, 5% and 80% of package C0,

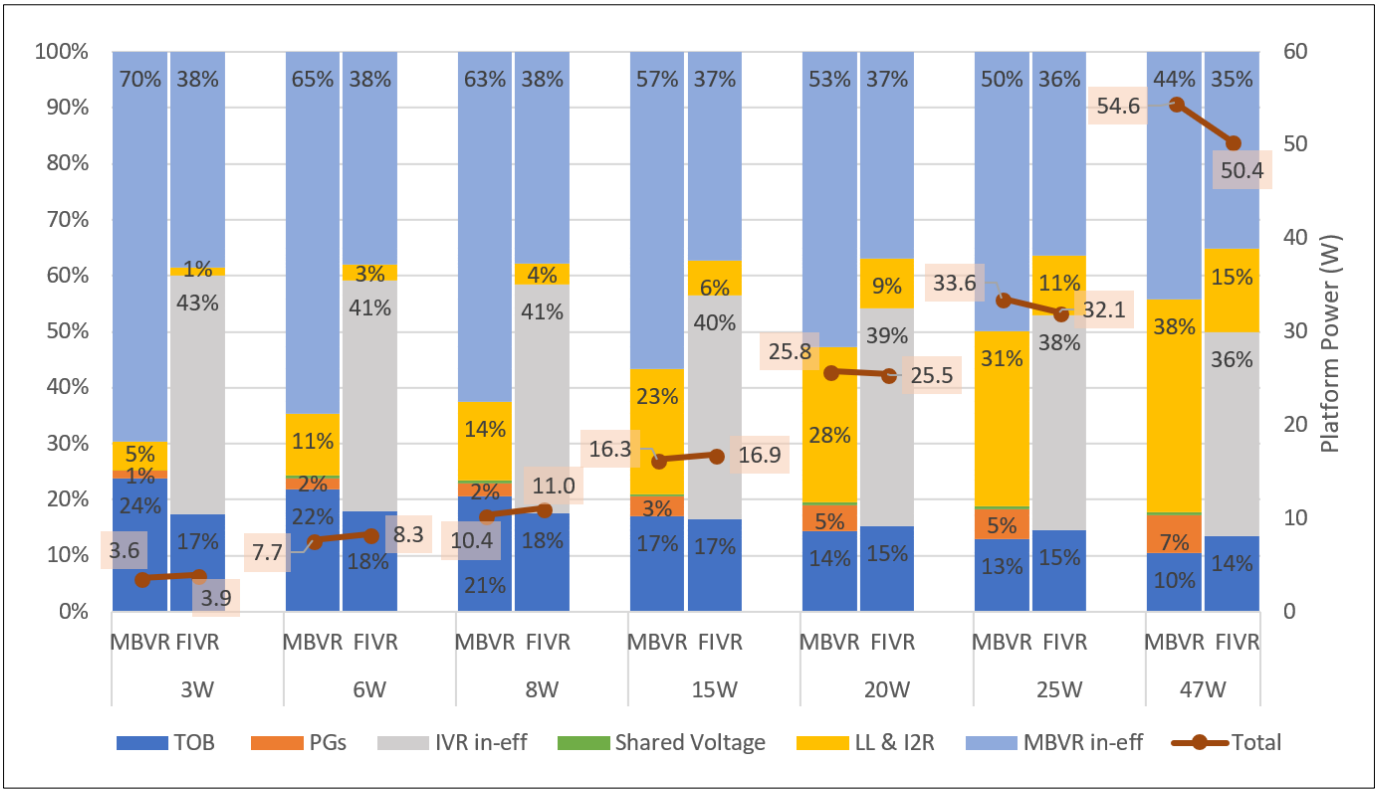

Fig. 3. Power losses breakdown at various TDP segments, for core TDP workload. The bars shows the percentage contribution of each power loss component to the overall power losses, lines show the total system power at each scheme. Power losses on power-gates (PGs), load-line(LL) and I^2R grow significantly at higher TDP (higher input current) for the MBVR scheme making it less efficient at these segments compared to IVR.

C2 and C7 residencies respectively. While video-conferencing have roughly 30%, 10% and 60% package C0, C2 and C7 residencies respectively.

Table II shows the results of the two average power scenarios projected by our models. The table shows the total power consumed from the power source (battery/PSU) for each package C-state (C0, C2, C7) for both schemes. Notice that the power at of this scenario is roughly the same at the various segments with the same configuration and parameter (e.g. caches sizes, material, graphics size, etc.). Multiplying the power at each package state by the residency of the state results in the contribution of the package c-state to the average power of the scenario and the sum of these parts gives the total average power consumed by each scenario.

The results at Table II show that the MBVR scheme have about $140mW$ and $230mW$ less average power than the IVR scheme at the video-playback and video conferencing scenarios respectively. Devices such as the Microsoft Surface Pro 4 [20] that have $38.2Wh$ battery capacity and powered by Skylake [8] processor (i5-6300U), have platform average power (including processor, panel, SSD, memory, etc.) consumption of about $4.1W$ and $4.7W$ [20] for video-playback and video conferencing scenarios respectively. The additional average power consumed due to IVR vs. MBVR is translated to 18.2 minutes (3.3%) less battery life for video-playback scenario and 22.4 minutes (4.6%) less battery life for video-conferencing scenario as shown in Table II. These results show that the MBVR is more energy efficient for light workloads compared to the IVR scheme.

V. Related Works

VRs are fundamental blocks of power delivery systems. In order to realize effective and fine-grained power management, many works have investigated the potentials of integrated VRs. Hazucha et al. [21] proposed a four-phase integrated buck converter using air-core inductors on package which achieves a 80%-87% efficiency. Wang et al. [22], [23] presented an integrated and analytical model of the whole power delivery system called PowerSoC, which involves off-chip and on-chip VRs as well as a power delivery network, providing a platform to evaluate the performance and explore the design space of the entire power delivery system. They discovered that hybrid architectures with both on-chip and off-chip VRs can achieve a better tradeoff between area reduction and efficiency requirements compared to traditional off-chip paradigms. In addition, previous works [1]–[4] claimed that the fully integrated voltage regulator (FIVR) that was first adopted at Intel® 4th [1] processor is raising the performance and increases battery-life. Haoran et al. [24] compared the characteristics of different power delivery for many cores system using on-chip and/or off-chip VRs based on an analytical model. Compared to our work, the study did not cover fully integrated voltage regulator,

TABLE II
COMPARISON BETWEEN IVR AND MBVR SCHEMES WITH VIDEO-PLAYBACK AND VIDEO-CONFERENCING WORKLOADS.

Scheme	MBVR		IVR	
Workload	**Video Playback**	**Video Conf.**	**Video Playback**	**Video Conf.**
C0 Power(W)	4.94		5.48	
C2 Power (W)	1.47		1.82	
C7 Power(W)	0.27		0.32	
C0 %	15%	30%	15%	30%
C2 %	5%	10%	5%	10%
C7 %	80%	60%	80%	60%
Average Power (W)	**1.03**	**1.79**	**1.17**	**2.02**
Workload			**Video Playback**	**Video Conf.**
Average power difference between the two schemes [P_{MBVR}-P_{IVR}] (W)			**-0.14**	**-0.23**
Surface Pro 4 battery capacity (Wh)			**38.2**	
Surface Pro 4 average power (W)			**4.1**	4.7
Battery-life time decrease (min)			18.27	22.47
Battery life decrease percentage (%)			3.3%	4.6%

comparison of different segments and battery-life workloads. Compared to these works, our study included the modeling at various TDP segments showing which scheme is better for each segment. In addition, we evaluated various workloads (core, graphics and average-power) and showed which scheme is more efficient at these workloads.

VI. DISCUSSION

Experimental results showed that MBVR is more energy efficient for Low TDP, while for high TDP the IVR is more energy efficient. In addition, for light workloads, such as average power workloads, the MBVR have better energy-efficiency. A straightforward **Hybrid scheme** that switches between both schemes based on TDP and workload characteristic would ensure better energy-efficiency, on the other hand such solution raises the design complexity and cost. Beside energy efficiency, there are other aspects to consider when designing PDN, space limitations preclude a detailed evaluation of each of them. Table III summarizes additional features to consider for the two schemes and the Hybrid scheme, we explain them in what follows.

Scalability: modern processors have multiple power domains, some of these domains have separate Dynamic-Voltage-Frequency-Scaling, for example, Intel servers have the Per-Core P-States (PCPS [25]) feature that enables each core to have different frequency and voltage work point. Compared to IVR scheme, the MBVR scheme is less scalable with the number of voltage/frequency domains as its limited by the motherboard and package routing resources.

Board area: IVR scheme reduces the board area since the number of high current IVRs are reduced.

Chip area: The IVR scheme adds significant silicon area and package resources for the IVR controllers, MIM (Metal-Insulator-Metal) capacitors and bumps area.

Load transient response: IVR runs at high frequency and closer to the load, compared to the MBVR scheme, hence the feedback loop is much faster this scheme. Nevertheless, both schemes require features to mitigate the third droop [26] due to instantaneous power at modern processor cores that runs at very high speed.

Chip-design complexity: The IVR scheme introduces more complexity to the chip designers, as it integrates new component into the chip, In addition, the IVR require redesign of many components when moving to denser fabrication process, where the IVR analog component doesn't shrink significantly with the newer process making it's relative area high compared to other processor parts that perfectly shrink with the newer process technology (e.g. digital logic).

Dynamic Voltage and Frequency Scaling (DVFS): The IVR enables faster voltage transitions compared to MBVR, whereas it can move the voltage from V_{min} (e.g. 0.5V) to V_{max} (e.g. 1.2V) in less than two microseconds, while it take more than ten microseconds (Assuming slew rate of 50mV/us) for modern VR to do the same voltage transitions. This capability, theoretically, enables finer-grain DVFS feature that potentially results in more power savings.

Off-chip voltage regulator current consumption: The IVR scheme uses higher (about two times) input voltage compared to the MBVR scheme, this results in reduced current consumption from external voltage regulator, that enables smaller and more efficient VR design.

VII. CONCLUSION

In this work we showed that - as many other areas in chip design - there is no one size fit all design choices for PDN. Indeed, there are pros and cons for IVR, in particular for the total power savings. With an IVR, in contrary to the

TABLE III
COMPARISON BETWEEN IVR, MBVR AND HYBRID SCHEMES ON MULTIPLE FEATURES.

Feature	IVR	MBVR	Hybrid
Energy-Efficiency	High TDPs	Low TDPs	All TDPs
Scalability	High	Low	Low
Board Area	Low	High	High
Chip Area	High	Low	High
Load transient response	Fast	Slow	fast/slow IVR/MBVR mode
Chip Design complexity	High	Low	High+
DVFS (V_{min} to V_{max} transition)	<2us	∼15us	2us/15us IVR/MBVR mode
Off-chip VRs current	Low	High	Low/High IVR/MBVR mode

current industry trend, the performance and average power, is not always guaranteed to improve. The total power savings depends on many parameters such as: TDP, efficiency of integrated voltage regulator and efficiency of motherboard voltage regulator. Expectedly, the workload plays a crucial role in the efficiency of PDN, which needs to be taken into consideration as well. On the other hand IVR schemes have higher scalability, less board area, less off-chip current consumption, fast load transient response and enables faster DVFS compared to transitional MBVR schemes. While a straightforward Hybrid scheme guarantees energy-efficiency at all TDPs and workloads, it introduces other challenges and design complexities.

In our future works, we will evaluate more PDN schemes and propose a more cost and area efficient hybrid scheme that maintains most of the advantages of IVR and guarantees energy-efficiency at all TDPs and workloads.

REFERENCES

[1] E. A. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan, and M. J. Hill, "Fivrfully integrated voltage regulators on 4th generation intel® core socs," in *Applied Power Electronics Conference and Exposition (APEC), 2014 Twenty-Ninth Annual IEEE*. IEEE, 2014, pp. 432–439.

[2] D. Kanter, "Haswells fivr extends battery life," *Microprocessor Report, The Linley Group*, 2013.

[3] S. N. D. M. Azmi, A.-L. Kor, C. Pattinson, and N. Bujang, "Energy efficiency of 4th gen intel® core processor versus 3rd gen intel® core processor," in *Green IT Engineering: Components, Networks and Systems Implementation*. Springer, 2017, pp. 135–153.

[4] A. Varma, B. Bowhill, J. Crop, C. Gough, B. Griffith, D. Kingsley, and K. Sistla, "Power management in the intel xeon e5 v3," in *Low Power Electronics and Design (ISLPED), 2015 IEEE/ACM International Symposium on*. IEEE, 2015, pp. 371–376.

[5] A. Nalamalpu, N. Kurd, A. Deval, C. Mozak, J. Douglas, A. Khanna, F. Paillet, G. Schrom, and B. Phelps, "Broadwell: A family of ia 14nm processors," in *VLSI Circuits (VLSI Circuits), 2015 Symposium on*. IEEE, 2015, pp. C314–C315.

[6] E. Rotem, A. Naveh, D. Rajwan, A. Ananthakrishnan, and E. Weissmann, "Power management architecture of the 2nd generation intel® core microarchitecture, formerly codenamed sandy bridge," in *Hot Chips 23 Symposium (HCS), 2011 IEEE*. IEEE, 2011, pp. 1–33.

[7] S. Jahagirdar, V. George, I. Sodhi, and R. Wells, "Power management of the third generation intel core micro architecture formerly codenamed ivy bridge," in *Hot Chips 24 Symposium (HCS), 2012 IEEE*. IEEE, 2012, pp. 1–49.

[8] E. Fayneh, M. Yuffe, E. Knoll, M. Zelikson, M. Abozaed, Y. Talker, Z. Shmuely, and S. A. Rahme, "4.1 14nm 6th-generation core processor soc with low power consumption and improved performance," in *Solid-State Circuits Conference (ISSCC), 2016 IEEE International*. IEEE, 2016, pp. 72–73.

[9] B. Howse and R. Smith, "Tick tock on the rocks: Intel delays 10nm, adds 3rd gen 14nm core product kaby lake," 2015.

[10] F. Waldron, R. Foley, J. Slowey, A. N. Alderman, B. C. Narveson, and S. C. Ó. Mathúna, "Technology roadmapping for power supply in package (psip) and power supply on chip (pwrsoc)," *IEEE Transactions on Power Electronics*, vol. 28, no. 9, pp. 4137–4145, 2013.

[11] V. Regulator-Down, "11.0 processor power delivery design guidelinesfor desktop lga775 socket," 2006.

[12] V. R. Module, "and enterprise voltage regulator-down (evrd) 11.1 design guidelines," *Intel Corp., Santa Clara, CA*, 2009.

[13] K. Anshumali, T. Chappell, W. Gomes, J. Miller, N. Kurd, and R. Kumar, "Circuit and process innovations to enable high-performance, and power and area efficiency on the nehalem and westmere family of intel processors." *Intel Technology Journal*, vol. 14, no. 3, 2010.

[14] R. Jakushokas, M. Popovich, A. V. Mezhiba, S. Köse, and E. G. Friedman, *Power distribution networks with on-chip decoupling capacitors*. Springer Science & Business Media, 2010.

[15] S. Rusu, H. Muljono, D. Ayers, S. Tam, W. Chen, A. Martin, S. Li, S. Vora, R. Varada, and E. Wang, "5.4 ivytown: A 22nm 15-core enterprise xeon® processor family," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*. IEEE, 2014, pp. 102–103.

[16] J. Doweck, W.-F. Kao, A. K.-y. Lu, J. Mandelblat, A. Rahatekar, L. Rappoport, E. Rotem, A. Yasin, and A. Yoaz, "Inside 6th-generation intel core: New microarchitecture code-named skylake," *IEEE Micro*, vol. 37, no. 2, pp. 52–62, 2017.

[17] Standard Performance Evaluation Corporation, "SPEC," online, accessed March 2018, Mar 2018, www.spec.org.

[18] Vantage, "3DMARK," online, accessed March 2018, Mar 2018, http://www.futuremark.com/benchmarks/3dmarkvantage.

[19] A. Nalamalpu, N. Kurd, A. Deval, C. Mozak, J. Douglas, A. Khanna, F. Paillet, G. Schrom, and B. Phelps, "Broadwell: A family of ia 14nm processors," in *VLSI Circuits (VLSI Circuits), 2015 Symposium on*. IEEE, 2015, pp. C314–C315.

[20] Anandtech, "The Microsoft Surface Pro (2017) Review: Evaluation, ," online, accessed March 2018, Mar 2018, https://www.anandtech.com/show/11538/the-microsoft-surface-pro-2017-review-evolution/7.

[21] P. Hazucha, G. Schrom, J. Hahn, B. A. Bloechel, P. Hack, G. E. Dermer, S. Narendra, D. Gardner, T. Karnik, V. De, *et al.*, "A 233-mhz 80%-87% efficient four-phase dc-dc converter utilizing air-core inductors on package," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 4, pp. 838–845, 2005.

[22] X. Wang, J. Xu, Z. Wang, K. J. Chen, X. Wu, and Z. Wang, "Characterizing power delivery systems with on/off-chip voltage regulators for many-core processors," in *Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014*. IEEE, 2014, pp. 1–4.

[23] X. Wang, J. Xu, Z. Wang, K. J. Chen, X. Wu, Z. Wang, P. Yang, and L. H. Duong, "An analytical study of power delivery systems for many-core processors using on-chip and off-chip voltage regulators," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 9, pp. 1401–1414, 2015.

[24] H. Li, X. Wang, J. Xu, Z. Wang, R. K. Maeda, Z. Wang, P. Yang, L. H. Duong, and Z. Wang, "Energy-efficient power delivery system paradigms for many-core processors," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 36, no. 3, pp. 449–462, 2017.

[25] D. Hackenberg, R. Schöne, T. Ilsche, D. Molka, J. Schuchart, and R. Geyer, "An energy efficiency feature survey of the intel haswell processor," in *Parallel and Distributed Processing Symposium Workshop (IPDPSW), 2015 IEEE International*. IEEE, 2015, pp. 896–904.

[26] J. Haj-Yihia, Y. B. Asher, E. Rotem, A. Yasin, and R. Ginosar, "Compiler-directed power management for superscalars," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 11, no. 4, p. 48, 2015.