Variation-Aware Heterogeneous Voltage Regulation for Multi-Core Systems-on-a-Chip with On-Chip Machine Learning

Joseph Riad[†], Jianhao Chen[†], Edgar Sánchez-Sinencio[†] and Peng Li[‡]
[†]Department of Electrical and Computer Engineering, Texas A&M University
[‡] Department of Electrical and Computer Engineering, University of California, Santa Barbara
Email: joseph.riad@tamu.edu

Abstract

Large-scale systems-on-a-chips (SoCs) have stringent power requirements to ensure adequate supply of power to on-die devices and prevent catastrophic timing violations. Heterogeneous voltage regulation (HVR) leveraging a combination of on-chip and off-chip voltage regulators has been advocated for ensuring power integrity with maximum efficiency. However, unavoidable process and temperature variations have not been considered in prior HVR work. In this paper, we present an in-depth evaluation of the impacts of process and temperature variations on HVR. Furthermore, we propose a systemic solution to incorporate variation awareness into the HVR system control policy to add a further improvement of up to 4.28% in system power efficiency with minimal hardware overhead.

Keywords

heterogeneous voltage regulation, variations, multi-core/SoC, machine learning

1. Introduction

With the growing size of systems-on-a-chips (SoCs), power delivery is becoming an increasingly complex problem. The two main metrics for power delivery are power integrity and efficiency; in other words, it is important to ensure that the supply voltage does not fall below a certain specified level to avoid catastrophic timing delays in digital circuits and other failures while minimizing the losses incurred in the power delivery network (PDN).

Power delivery often entails voltage regulation and as such voltage regulators (VRs) can significantly impact power integrity and the overall efficiency of a PDN. It shall be noted that various types of voltage regulators that have been employed in practice possess distinct and complementary characteristics. The bulk of the voltage down-conversion from the system supply (battery voltage for example) to the on-chip supply voltage is usually carried out using a switching VR for its superior efficiency. Typically, single-stage voltage regulation with off-chip switching converters is employed, which may suffer from poor power integrity because of long response time and high board/package parasitics. Introducing on-chip switching converters with faster response time and small footprint as the second stage of regulation could address this problem [1]. However, switching converters require bulky and costly passive elements like inductors. Compared with switching converters, the linear VRs such as low-dropout voltage regulators (LDOs) have high area-efficiency and sub-ns response time. This comes with the benefit to deploy LDOs across the chip and close to power hotspots to supress the voltage noise. As a good example, IBM POWER8 processor employs 1,764 on-chip distributed linear VRs [2].

Recently, a holistic architecture combining the advantages of both switching and linear VRs into a system called heterogeneous voltage regulation (HVR) was proposed [3]. This system consists of a cascade of three types of VRs (see Figure 1) : off-chip and on-chip switching VRs and on-chip

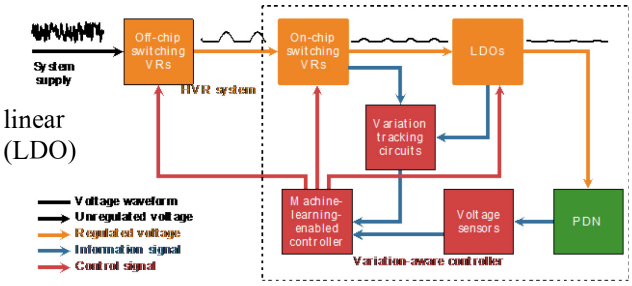

Figure 1: Proposed variation-aware HVR architecture.

VRs. Switching voltage converters convert the supply voltage with high efficiency, while the distributed LDOs suppress the voltage noise in the PDN. A workload-aware management policy was proposed to adjust the system for optimal system power efficiency while providing a guaranteed power integrity.

The system in [3] captures the system workload through distributed voltage and power sensors. To obtain the best system performance, the controller uses an offline trained machine learning (ML) model and efficiency characteristics of switching converters stored in on-chip lookup tables (LUTs). So far, the offline trained ML model and efficiency LUTs are obtained based on the nominal system and no work has studied how the inevitable process variations and temperature drifts affect the HVR system, which is expected to perform sub-optimally when the circuit parameters drift from their nominal values used to train the ML model.

In this work, we study the effect of such variations on the HVR system and propose on-chip measurement circuits tracking parameter variations with minimal hardware overhead, thereby creating a variation-aware HVR system as shown in Figure 1. The main contributions of this work consist of demonstrating the interesting fact that the nominal model has a certain degree of robustness to variations and that implementing our proposed variation-aware architecture can add an improvement of up to 4.28% in overall efficiency while guaranteeing power integrity.

2. Variation-Aware HVR Architecture Description

Figure 2 shows a detailed view of the proposed architecture outlining how different VRs are distributed around the PDN as well as how a controller is used to tune the system. The nominal system control flow and proposed variation-aware control flow are detailed below.

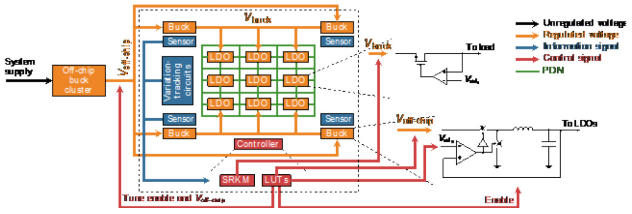

Figure 2: Detailed view of the proposed system architecture.

2.1 Basic system

The proposed system architecture uses three-stage voltage regulation as has been proposed in [3] to capitalize on both the superior efficiency of buck converters and the excellent regulation performance of LDOs. The off-chip buck converter cluster in the first stage is driven by the system power supply and powers the on-chip buck converter clusters in the second stage. The on-chip PDN is divided into multiple power domains, each of which contains one on-chip buck converter cluster and powers one core (Figure 2 shows one power domain). The second-stage-regulated voltage powers a cluster of distributed LDOs that are all connected to the system's PDN that drives the load circuits.

The power loss of an LDO is approximately proportional to its dropout voltage [3], so a low LDO input voltage provides a high regulation efficiency. However, dynamic and non-uniformly distributed workloads may introduce voltage noise in the PDN unless a higher LDO input voltage is used. Therefore, an offline-trained Sparse Relevance Kernel Machine (SRKM) model is used to learn the workload distribution and predict the LDO input voltage (V_{out_on}) online to improve the LDO regulation efficiency without jeopardizing the power integrity.

The intermediate voltage between the off-chip and on-chip buck converters (V_{out_off}) also needs to be tuned as it determines the conversion ratios of the first two stages and consequently affects the system efficiency. In addition to the conversion ratio, the load current can also affect buck converter efficiency. Since multiple converters are used in the first two stages, the number of active on- and off-chip buck converters (N_{on}, N_{off}) can be controlled to adjust the load current for each converter so that they can achieve the best efficiency given the input and output voltage.

To sum up, V_{out_off}, V_{out_on}, N_{off} and N_{on} are defined as system control variables and are tuned by the controller online. The controller searches for the configuration with the highest efficiency according to the system workload and adjusts these control variables. To evaluate the system efficiency, lookup tables (LUTs) are used to store the efficiency characteristics of on-chip and off-chip buck converters, and power and voltage sensors are used to capture workload information. First, the optimal load currents for given input and output voltages are stored and indexed by input and output voltages. N_{on} and N_{off} can be obtained by dividing the estimated total load current by the optimal load current for each converter. The conversion efficiencies are also stored and indexed by input and output voltage and load current so that the overall efficiency may be evaluated. The SRKM model takes readouts of a series of voltage sensors distributed throughout the PDN (not depicted in Figure 2 for clarity) to capture the workload distribution and predict the LDO input voltage. Core-level load currents (I_L) can be estimated using power sensors [4] at the end of each on-chip buck converter cluster.

The response time of on-chip and off-chip buck converters can differ by several orders of magnitudes, thus two different control cycle times, T_{on} and T_{off}, are applied to adjust them separately. Each T_{off} cycle is split into multiple T_{on} cycles. This is shown in Figure 3.

During the T_{off} cycle, the control variables related to off-chip buck converters, V_{out_off} and N_{off}, are determined by sweeping V_{out_off}. In each step, V_{out_on} is estimated using the workload currents in the previous control cycle and the linear dropout voltage model [3]. After the voltages are determined, the optimal load current for each on-chip and off-chip buck converter can be found in the LUTs and N_{on} and N_{off} are determined using the estimated load current (I_L). With all control variables set, the efficiency of the buck converters can be found in the LUTs while the LDO efficiency can be calculated using the dropout voltage and load current [3] . As a result, the system efficiency can be evaluated. When the sweeping is done, the configuration with the highest efficiency is selected.

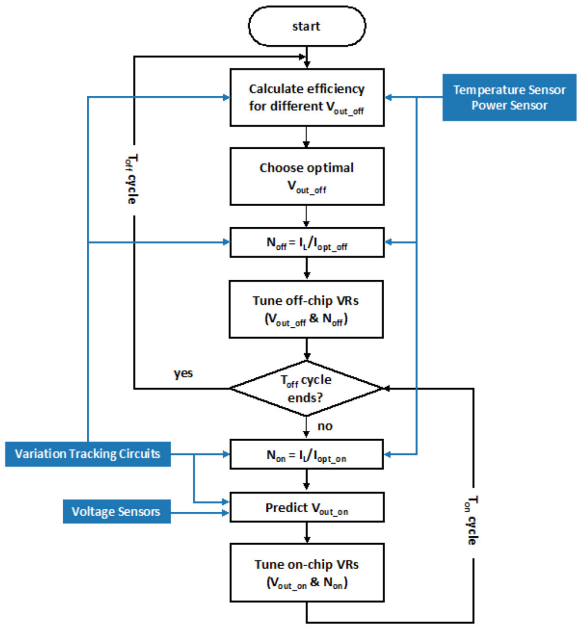

Figure 3: Control flow of HVR and variation tracking mechanisms.

During the T_{on} cycle, the control variables for the on-chip converters are set. First, the SRKM model predicts the LDO input voltage (V_{out_on}). With their input voltage determined in the T_{off} cycle and output voltage (same as the LDO input voltage) determined in the T_{on} cycle, the optimal load current

for the on-chip buck converters can be found using the LUTs, and N_{on} is determined.

2.2 Variation-aware controller

To ensure that the control policy continues to perform correctly in the presence of variations, a variation-aware control flow is proposed and shown in Figure 3. The controller gathers information about the system through a series of voltage sensors distributed throughout the PDN, low-power temperature sensors and variation-tracking circuits that are used to measure the varying circuit parameters *in situ* as detailed in the next section.

The efficiency LUTs of the on-chip buck converters are measured at -40ºC and 125ºC during the setup stage. The efficiency values used for system efficiency evaluation are interpolated from these two LUTs using the measured temperature from the temperature sensor.

To maintain power integrity in the presence of varying load currents and LDO process variations and mismatch, the on-chip SRKM model is modified with the average loop gain and loop bandwidth of LDOs as two extra features. The SRKM module is trained offline using Monte Carlo simulation data at -40ºC, 27ºC and 125ºC with 200 different process variation samples accounting for process and temperature variations. Thus, the SRKM model can predict the correct LDO input voltage during the system's operation.

3. Tracking Variations in Regulator Circuits

This section explains the main performance parameters of interest for both types of VRs and describes the circuits proposed to track their variations.

3.1 Buck converters

In terms of regulation performance, the regulation speed of a buck converter is governed by the loop's dynamics captured by its crossover frequency f_c and its phase margin ϕ_m while the settling error is decided by the loop's low-frequency gain. Note that this work ignores the buck converter's output voltage ripple since the HVR architecture includes LDO regulators to clean this ripple.

A technique for adapting buck converter loop dynamics to process variations has been proposed in [5], using low-cost digital logic circuits to ensure fixed values for f_c and ϕ_m regardless of load current or process variations. Thus, the regulation performance of the buck converters is assumed to be unaffected by process variations. That being said, the modular nature of the proposed solution allows it to be extended to account for these variations if necessary.

The characteristics of the MOS switches and the inductor are the main factors that affect power efficiency. Since these characteristics exhibit process and temperature variations, so does the power efficiency of the buck converters. This work proposes a simple circuit to measure these characteristics during a startup setup stage and store them in the LUTs used by the controller as explained in section 2.

Figure 4 shows the proposed measurement circuit for the efficiency characteristics. The digital logic module sweeps the input and reference voltages. It also sweeps the load current by selecting different values for the load resistance from a resistor bank. As the different parameters are swept, the voltage drops across R_L and R_{sense} are measured using the voltage sensors, digitized using the analog-to-digital converter (ADC) and fed back to the logic module which uses

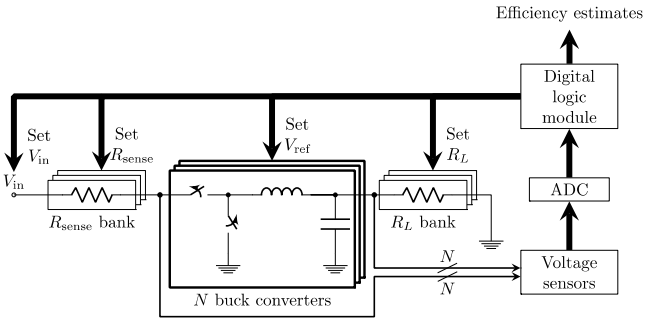

Figure 4: Proposed on-chip circuit to measure buck converter efficiency.

them to estimate the converter efficiency. Note that, as R_L is changed, R_{sense} must be changed as well in order to ensure that the voltage drop across it remains large enough to be captured by the sensors. Of all the components required for these measurements, only the ADC is bulky and expensive but, since ADCs are essential components of large SoCs, the measurement can reuse ADC circuits found on the chip thereby requiring no additional hardware cost.

Knowing the sensor voltages and resistor values, the digital logic can compute an estimate for the efficiency at these operating conditions. Using this setup, an LUT can be constructed for each buck converter thereby accounting for process variations and mismatches.

3.2 LDO regulators

In the case of LDOs, auto-tuning to compensate for variation effects on regulation is impractical. Since the LDOs are the last stage in the power management architecture, the requirements on their regulation performance are quite stringent, requiring them to react quickly to disturbances. This means that LDO control loops tend to have quite high crossover frequencies that are prohibitively expensive to monitor. In contrast, the loop's 3dB bandwidth may be measured more easily (the approach used in this work).

LDO power efficiency is mainly dominated by the difference between the input and output voltage since this difference dictates the power dissipated in the output power device. This means that LDO power efficiency can be controlled by controlling the input voltage so that setting the LDO input voltage sets its efficiency regardless of process and/or temperature variations.

Figure 5 shows the proposed circuit to track the parameters of interest for an LDO. The voltage v_i is used to estimate the loop gain by sweeping v_i and measuring v_x.

This technique is inspired by Middlebrook's technique [6], which allows accurate loop gain measurements in a limited bandwidth. It is equivalent to injecting a voltage between the output of the resistive divider and the input of the error amplifier. The loop gain L can be calculated as

$$L(s) = \frac{R_1}{R_1 + R_2} \cdot \frac{v_i(s)}{v_x(s)} - 1 \qquad (1)$$

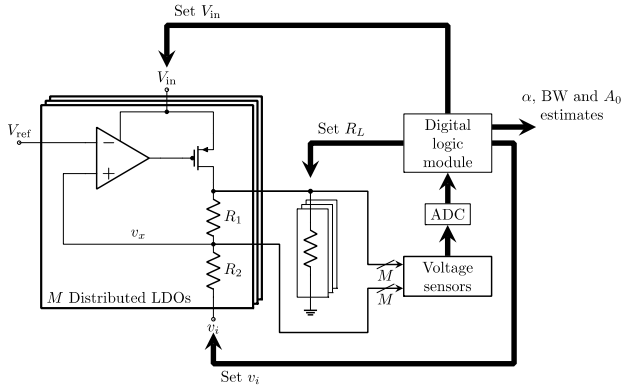

Figure 5: Proposed on-chip circuit to measure LDO parameters.

Using a sinusoidal test signal with varying frequency, $v_i(j\omega)$, the value of $L(j\omega)$ can be measured for different frequencies by measuring $v_x(j\omega)$ and calculating $L(j\omega)$ from (1). Using this procedure, the loop's -3 dB bandwidth (BW) and low-frequency gain (A_0) can be estimated. The parameters BW and A_0 correlate with process and temperature variations of the LDO circuit and can therefore be used to predict the LDO regulation performance. Thus, the values of BW and A_0 are averaged over all LDOs and the resulting two numbers are added as features to the SRKM model in order to account for the effect of variations on the optimal LDO input voltage.

Adjusting the LDO input voltage to optimize efficiency and guarantee power integrity requires an estimate of the LDO dropout voltage V_{do} (as explained in [3]). V_{do} varies with process and temperature variations, so the variation-tracking circuits should monitor it as well. Simulations show that V_{do} is a linear function of load current and may be modeled as $V_{do} \simeq \alpha I_L$. To measure α, v_i is set to 0, then V_{in} and load current (R_L) are swept. Hardware overhead is minimized by sharing the resistors and reusing the existing on-chip ADC.

3.3 Temperature variations

In addition to capturing process information, it is important to measure the temperature as well, so that parameter variations due to temperature drifts do not affect the system efficiency and power integrity.

For the buck converters, the variation of efficiency characteristics with temperature is quite linear and therefore a two-point measurement is sufficient. Thus, temperature information can be incorporated during the initial setup phase when the LUTs are being constructed. It is sufficient to characterize the LUTs at two temperature extremes (-40ºC and 125ºC), During normal operation, an ultra-low-power temperature sensor such as the one reported in [7] is placed near each buck converter as shown in Figure 2 to estimate the temperature and linearly interpolate between the temperature data in the LUT.

As for LDOs, temperature variations are reflected in changing values for the parameters α, BW and A_0. These parameters, therefore, should be measured periodically to account for the change in temperature. This can be accomplished by measuring a replica LDO or performing periodic measurement on the real LDOs and does not constitute a significant overhead as temperature drifts very slowly in comparison to processor speed so the measurements will be repeated with a low frequency.

4. Experimental Setup

This section describes the setup used to model process variations in the different VRs and evaluate their effects on the overall system efficiency.

4.1 Circuit simulations

Since the buck converter output is further regulated by the LDOs, an accurate model for its output voltage ripple is not critical to capture the system behavior. For this reason, the off-chip buck converters were modeled as ideal voltage sources while the on-chip buck converters were implemented with a behavioral model in order to speed up simulation. The behavioral model uses ideal switches and passives and a Verilog-A loop controller.

To model process variations using Monte Carlo simulations, real transistors in a standard 90nm CMOS technology were simulated and exhibited a standard deviation of about 2.5% in their on-state resistance. For the passive devices, a standard deviation of 20% was assumed as their characteristics tend to have a large variance [8]. Using these values and the Verilog-A model, Monte Carlo sampling was used to generate multiple instances of buck converters representing different samples of process variations.

To capture the LDO behavior accurately, an LDO implementation [9] in a standard 90nm CMOS technology was used for simulations. and Monte Carlo simulations sampling random variations were used to generate multiple training cases for the ML model. The ML model was trained with 3000 training points collected from circuit simulation of 200 process variation samples at temperatures of 27ºC, 125ºC and -40ºC. It achieved a normalized mean square error (NMSE) of 0.0343. The model can then be mapped to a hardware accelerator for run-time prediction [3].

4.2 Evaluation setup

We use the full-system architecture simulator GEM5 [10] to run the PARSEC [11] benchmark and generate run-time statistics of the processor, which are then fed to McPAT [12] to generate current traces of the processor. The processor in Table 1 with 4 cores is modeled in the simulator. Each core is divided to 11 function blocks, and the current workload for each block is assumed to be distributed evenly in the corresponding area.

Table 1: Processor Configuration

# Cores	4	**Frequency**	1.8 GHz
Vdd	1 V	I_{max} **(per core)**	25 A
Core Area	40.4 mm^2	**ALU/MUL/FPU**	6/2/6
L1 Cache	32 KB	**L2 Cache**	Shared 2MB

5. Results and Discussion

In this section, the robustness of the nominal control policy in [3] is assessed and the power integrity and the

system efficiency of the nominal and the proposed variation-aware control policies are evaluated and compared.

In modern processors, multiple factors like workload variations and clock gating can cause voltage noise in the PDN. When the voltage in the PDN drops below a particular level (0.9V in this study), a voltage emergency (VE) is said to occur. The power integrity of the PDN is examined by running a time-domain circuit simulation and counting the occurrences of VEs.

Designing the PDN based on the worst case VE to guarantee the power integrity will degrade the system efficiency significantly. Therefore, rare occurrences of VEs are tolerated in exchange for a better system efficiency. It is assumed that the processor is equipped with fail-safe mechanisms such as the rolling-back recovery [13] or adaptive frequency tuning [14] for such rare emergencies. To account for the ML model errors which may cause VEs in the PDN, a small voltage guardband is added to the prediction value. This guardband was adjusted for the nominal control policy and the proposed variation-aware policy such that they have the same level of power integrity measured by the number of VEs for a fair comparison. With the same level of power integrity, the system efficiencies for the nominal and variation-aware control policies are evaluated and compared.

5.1 Nominal control policy robustness

First, the effect of on-chip buck converter variations on the nominal control policy is investigated. The LDO input voltage is set to 1.05V and remains unchanged and the workload current is set to 10A for each core in this part. The system temperature is assumed to be 27ºC. As circuit variations are not considered by the nominal control policy, it will make identical decisions for a given workload even as temperature and process parameters are varied. The red solid curve in Figure 6 shows a sweep of off-chip converter output voltage (V_{out_off}) and the corresponding system efficiency evaluated using the nominal efficiency LUT. The V_{out_off} with the best system efficiency is 1.3V, which is then chosen by

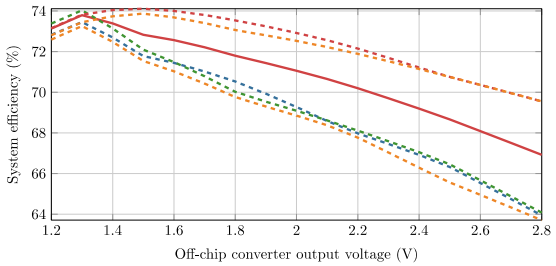

Figure 7: Overall system efficiency as a function of off-chip converter output voltage. The solid curve corresponds to the nominal model and the dashed curves represent 5 variation samples.

the nominal control policy regardless of circuit variations. However, the true efficiency curves in the presence of variations would deviate from the red solid curve. The dashed curves in Figure 6 are efficiency curves evaluated using variation-aware LUTs for 5 process variation samples.

For 2 out of the 5 samples, the control policy using the nominal LUTs will fail to select the best output voltage for the off-chip converters. However, the variations only shift the optimal solution by 200mV, resulting in an efficiency degradation of < 1%. For the other 3 samples, the optimal solution is at the same output voltage as the nominal control policy so that, even though the system will estimate efficiency incorrectly, it will still manage to find the optimal control point. Thus, the system exhibits some robustness to variations but, since the system can perform sub-optimally in some cases, it is to be expected that adding variation awareness to the controller should improve efficiency.

Next, the effect of LDO variations on the nominal control policy is considered. The nominal and variation-aware ML models were simulated with the same workload segment from the Blackscholes benchmark. The simulation waveforms are shown in Figure 7. Both models were simulated with the same LDO network and simulation temperature (125ºC). The guardbands for both models were adjusted to provide the same level of power integrity.

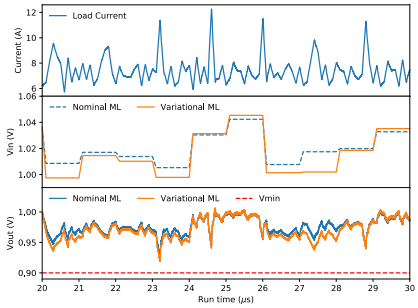

Figure 6: Predictions of the nominal and variation-aware ML models.

At 23μs, a current surge occurs, leading to the output voltage drop in both waveforms. In this case, the ML models should raise the LDO input voltage in the cycle between 24μs and 25μs. As shown in Figure 7, both ML models react this way. Another current surge occurs after 24μs, which drives both models to increase the LDO input voltage further in the cycle between 25μs and 26μs.

These waveforms indicate that the nominal ML model can correctly function in the presence of variations; *i.e.,* it can react to low voltage sensor readouts and raise the LDO supply voltage to prevent further VEs. However, the increment applied to the LDO input voltage might not be optimal due to the variations. As shown in Figure 7, the variation-aware ML model could predict lower voltages with the same power integrity level, thus increasing the system efficiency.

5.2 Efficiency improvement with the variation-aware controller

After ensuring both control policies are adjusted for the same power integrity (as explained above), the power efficiencies of four different optimization schemes are evaluated in the presence of variations. The 1st scheme is the nominal control policy proposed in [3], denoted by "Nominal". The LUTs and the ML model for the nominal case are used in the control policy. In the 2nd and 3rd schemes, the control policy is augmented with either the variation-aware LUTs ("LUT_OPT") or the variation-aware ML model

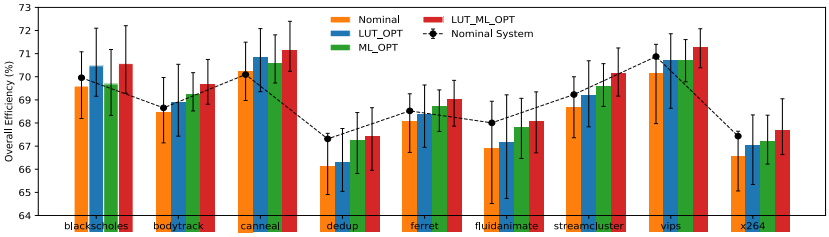

Figure 8: Average system efficiency for different optimization schemes.

("ML_OPT"). In the 4th scheme, both optimizations are applied ("LUT_ML_OPT").

To compare these schemes, the system efficiency using each scheme is evaluated for 4 different process variation samples at temperatures of 27ºC, 125ºC and -40ºC. Figure 8 shows the average efficiency for each scheme over all 12 evaluation cases. The error bars represent the maximum and minimum efficiency for each scheme. As shown in the figure, the proposed control policy outperforms the nominal one in every benchmark. The average efficiency improvement has a maximum of 4.28% and an average of 1.13%.

Furthermore, the efficiency of the nominal control policy applied to the nominal system (no process variations) is evaluated (shown as "Nominal System" in Figure 8). Thus, Figure 8 shows that, on average, the actual efficiency in a system with variations is lower than the efficiency in the nominal system. Note also that the effect of variations on system efficiency is limited, with the maximum efficiency deviation from the nominal system being 3.28%, This is further proof of the nominal control policy robustness described in section 5.1.

6. Conclusion

A detailed study of the effects of process and temperature variation on an HVR power distribution system was presented, wherein it was found that the system has some degree of robustness to variations. Two circuits were proposed to enable the tracking of process and temperature variation effects on the different VRs found in the HVR system with minimal hardware overhead. Using the proposed circuits, a scheme was developed and implemented to add variation awareness to the HVR system controller, which resulted in an efficiency increase of up to 4.28% and 1.13% on average with the same level of power integrity.

7. Acknowledgment

This work was funded in part by the National Science Foundation under Grant No. ECCS-1810125.

8. References

[1] E. A. Burton and others, "FIVR- Fully Integrated Voltage Regulators on 4th Generation Intel® Core SoCs," in *APEC*, 2014.

[2] V. Zyuban and others, "IBM POWER8 Circuit Design and Energy Optimization," *IBM Journal of Research and Development,* vol. 59, pp. 9-1, 2015.

[3] X. Zhan, J. Chen, E. Sánchez-Sinencio and P. Li, "Power Management for Multicore Processors via Heterogeneous Voltage Regulation and Machine Learning Enabled Adaptation," *TVLSI,* pp. 1-14, 2019.

[4] M. Ware, K. Rajamani, M. Floyd, B. Brock, J. C. Rubio, F. Rawson and J. B. Carter, "Architecting for Power Management: the IBM® Power7 Approach," in *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*, 2010.

[5] J. Morroni, R. Zane and others, "Adaptive Tuning of Digitally Controlled Switched Mode Power Supplies based on Desired Phase Margin," in *IEEE Power Electronics Specialists Conference*, 2008.

[6] R. D. Middlebrook, "Measurement of Loop Gain in Feedback Systems," *International Journal of Electronics Theoretical and Experimental,* vol. 38, pp. 485-512, 1975.

[7] W. Arshad, R. Ramzan and others, "Comparison and Design of VCOs for Ultra-Low Power CMOS Temperature Sensors," in *ICCSPA*, 2015.

[8] P. Napolitano, M. Poncik and others, "Effect of Parasitics and Process Variations in Buck Converters Operated in Voltage-Mode," in *2018 29th Irish Signals and Systems Conference (ISSC)*, 2018.

[9] S. Lai and P. Li, "A Fully On-Chip Area-Efficient CMOS Low-Dropout Regulator with Fast Load Regulation," *Analog Integrated Circuits and Signal Processing,* vol. 72, pp. 433-450, 01 8 2012.

[10] N. Binkert, B. Beckmann and others, "The GEM5 Simulator," *ACM SIGARCH Computer Architecture News,* vol. 39, pp. 1-7, 2011.

[11] C. Bienia and K. Li, Benchmarking Modern Multiprocessors, Princeton University Princeton, 2011.

[12] S. Li, J. H. Ahn and others, "McPAT: an Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures," in *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, 2009.

[13] H. Akkary, R. Rajwar and others, "Checkpoint Processing and Recovery: Towards Scalable Large Instruction Window Processors," in *Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture*, 2003.

[14] K. A. Bowman and others, "A 16 nm All-Digital Auto-Calibrating Adaptive Clock Distribution for Supply Voltage Droop Tolerance across a Wide Operating Range," *JSSC,* vol. 51, pp. 8-17, 2016.