

# Voltage Adaptation under Temperature Variation

Hussam Amrouch\*, Behnam Khaleghi† and Jörg Henkel\*

\*Karlsruhe Institute of Technology, Chair for Embedded Systems (CES), Karlsruhe, Germany

†University of California, San Diego, USA

amrouch@kit.edu

(Special Session)

**Abstract**—The maximum delay of a processor is subject to both temperature ( $T$ ) and voltage ( $V$ ) jointly. Therefore, there is the prospect of reducing  $V$  at runtime whenever  $T$  is below the worst-case temperature without violating the predetermined timing constraints, instead of employing throughout a conservative  $V$  that corresponds to  $T_{worst}$ . Hence, the efficiency can be maximized, while still accounting for temperature variation at runtime, without sacrificing reliability. In this work, we focus on the key challenge of how the *small, yet sufficient voltage* can be accurately, at design time, obtained in which timing constraints at runtime are assuredly fulfilled. To achieve that, we model the delay of processor under the joint effects of  $T$  and  $V$  through creating cell libraries that contain the detailed delay information of cells at a wide variety of  $T$  and  $V$  cases. Libraries are compatible with commercial EDA tools. Hence, standard tool flows, even though they were not designed for that purpose, can be employed to seamlessly obtain the correlation between  $T$  and  $V$  in any circuit regardless its complexity. After modeling of the correlation between  $T$  and  $V$  from the physical level to circuit level, we implement at the system level a temperature-guided voltage adaptation technique that tunes the voltage at runtime following temperature variations leading to a considerable reduction in power.

## I. INTRODUCTION

Unsustainable temperatures, caused by high on-chip power densities [1], [2], have become one of the major obstacles that challenges designers in the current technology nodes. To overcome the effects of temperature increase on prolonging the circuit’s delay, a timing guardband that corresponds to the worst-case temperature ( $T_{worst}$ ), is typically employed on top of the critical path ( $CP$ ) delay of circuit [3], [4]. This guarantees that the circuit will be always clocked at a sustainable frequency regardless temperature changes at runtime. Even though such an approach is safe, it is, in fact, inefficient as the temperature may, most of the operation time, stay below  $T_{worst}$  due to dynamics in running workloads. On the other hand, voltage has an opposite impact to temperature on delay (i.e. when  $V$  decreases, the circuit’s delay increases). Hence, there is the prospect of reducing  $V$  at runtime whenever  $T < T_{worst}$  without violating timing constraints. In this work, we introduce a novel technique that continuously tunes the voltage based on the operation temperature of processor. *It continuously eliminates any excess in voltage margin (i.e. avoiding wasted energy) following temperature dynamics at runtime towards maximizing the efficiency.* We call our technique *voltage adaptation* to differentiate it from conventional voltage scaling techniques (e.g., [5]) in which a controller predicts the performance requirement of the next task and accordingly scales  $V$  to save energy. Such techniques often necessitate having the maximum temperature that will be

generated by the next task a priori, which might not be possible in some (e.g., real time) embedded systems.

**Temperature and Voltage Correlation:** When the temperature rises, lattice vibrations in silicon increase leading to increasing the probability of an electron being scattered [6]. Thus, the carrier mobility ( $\mu$ ), which is the drift velocity of a particle in an electric field, decreases. In turn, a lower  $\mu$  leads to a lower  $I_{ON}$  in transistors [6] and hence larger delay (i.e. transistors become slower) [7], demonstrates. Therefore, a temperature increase results in increasing the  $CP$ ’s delay of circuit and vice versa. To ensure clocking the circuit with a sustainable frequency at runtime, a timing guardband needs to be typically included. Opposite to temperature, voltage has a contrary impact on transistor delay. When  $V$  decreases, the  $CP$ ’s delay becomes larger as the delay of its individual transistors increases. *Therefore, a delay improvement due to a temperature reduction can be equivalently translated into a lower voltage without violating timing constraints.*

**The focus of this work** is to accurately model the correlation between  $T$  and  $V$  with respect to the timing behavior of complex circuits like processors. This enables us to *adaptively* employ an equivalent voltage ( $V_{adapted} < V_{nominal}$ ) whenever ( $T_{operation} < T_{worst}$ ) in which the efficiency is maximized while timing constraints are still fulfilled. Reducing  $V$  results in less dynamic and leakage power. Therefore, energy and temperature can be optimized as well as reliability because reductions in  $T$  and  $V$  decelerate the underlying mechanisms of aging [8]. Note that the focus of our work is minimizing the efficiency loss akin to temperature guardband that inevitably required to overcome temperature variations at runtime. Other kinds for runtime variation like aging [9] are out of the scope and they necessitate their own guardbands to be employed.

## II. KEY CHALLENGES IN MODELING TEMPERATURE AND VOLTAGE EFFECTS ON PROCESSOR’S DELAY

Temperature and voltage effects originate from the physical level and then propagate all the way up to the system level where they finally cause timing violations. In the following we briefly summaries these effects:

**(a) Physical and Transistor Levels:**  $T$  and  $V$  dependencies in MOSFETs are very complex as they *differently* impact its key parameters like energy band gap, velocity saturation, carrier mobility, threshold voltage, leakage current, etc. [6]. A higher  $T$  reduces  $\mu$ , but it concurrently reduces  $V_{th}$ , which has a contradictory impact compared to  $\mu$  (i.e. it increases  $I_{ON}$ ). Therefore, the overall impact of  $T$  and  $V$  changes on transistor delay inevitably necessitates employing physics-based modeling. To account for that, the industrial compact model of MOSFET BSIM [7] is typically employed as it accurately

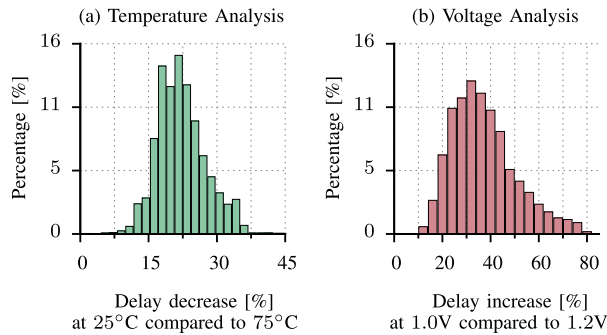


Fig. 1. The delay of gates within a standard cell library can be *non-uniformly* shifted under the same temperature (a) or voltage (b) change.

captures the effects that  $T$  and  $V$  have on all of the MOSFETs parameters along with the existing interdependencies.

**(b) Gate Level:** Even if the impact of  $T$  and  $V$  changes on pMOS/nMOS transistors is accurately estimated, the corresponding impact on the delay of gates will still be ambiguous because gates can be unevenly affected. Furthermore, the delay of the same gate itself may also be differently influenced based on the existing operating conditions ( $OPCs$ ) in that gate (i.e. its input signal slews and output load capacitance). To demonstrate that, we analyze the impact of  $T$  reduction from  $75^\circ\text{C}$  to  $25^\circ\text{C}$  on the gates of a 45nm standard cell library [10]. As shown in Fig. 1(a), the same  $T$  reduction (i.e.  $50^\circ\text{C}$ ) *non-uniformly* impacts the delay of gates. The delay improvement (i.e. decrease) can vary from merely 0.6% to 49%. Analogously, we also analyze the impact of  $V$  decrease from 1.2V to 1.0V. As shown in Fig. 1(b), the same  $V$  reduction of 200mV also *non-uniformly* impacts the delay of gates with a large variance between 2.5% and 81%.

**(c) Circuit Level:** Since the same change in  $T$  or  $V$  *non-uniformly* shifts the delay of gates, the prospect that the path which was formerly (i.e. before the  $T$ ,  $V$  change) non-critical to become critical needs to indeed be considered.

**In summary:** Analyzing *accurately* the processor’s delay under  $T$  and  $V$  effects, towards implementing a reliable temperature-guided voltage adaptation, is challenging as it necessitates involving all of the physical, transistor and circuit levels jointly. Considering the fact that the paths of a processor may switch their role w.r.t criticality necessitates, additionally, studying the entire processor’s netlist. Hence, it is indispensable to perform the required timing analysis by means of existing commercial tool flows. Including their optimized and verified algorithms therein can effectively and correctly handle the large complexity in such sophisticated designs.

### III. OUR TEMPERATURE-GUIDED VOLTAGE ADAPTATION

Our implementation of the proposed temperature-guided voltage adaptation consists of the following key steps:

- Creating cell libraries for the operational ranges of  $T$ ,  $V$ .
- Measuring the processor’s delay under  $T$  and  $V$  effects.
- Modeling the timing of processor correlating temperature reductions to the corresponding equivalent voltages.
- Implementing a runtime technique that adapts voltage based on the processor’s temperature without violating timings.

*Linking the physical and system levels:* step 1 covers the physical, transistor and gate levels. Then, steps (2 and 3) cover the circuit level. Finally, step 4 covers the system level where our technique is implemented. Note that because our cell libraries created in step 1 are compatible with standard tool flows, we are able to employ in step 2 a commercial Static Timing Analysis tool (STA) in which its mature algorithms automatically handle the challenge of  $CP$  switching when  $T$  or  $V$  changes (details in Section II).

**Cell Libraries Creation under  $T$  and  $V$  Effects:** Similar to the presented steps in [4], [11], [12], which showed how degradation-aware cell libraries can be created to analyze degradation effects (e.g., aging and temperature), we create in the following cell libraries that consider  $T$  and  $V$  effects. We employ SPICE netlists of the combinational and sequential gates within the Nangate 45nm open-cell library [10]. Every netlist includes the parasitics information extracted from the gate’s layout. The Predictive Technology Modeling (PTM) [13] for the high-performance 45nm technology together with the BSIM4 model [7] is employed. BSIM4 enables us to model the complex dependencies of MOSFET’s parameters. Thus, the effects of  $T$  and  $V$ , at the physical and transistor levels, can be properly considered. Then, we employ the HSPICE circuit simulator to accurately measure the delay and power (leakage and dynamic) of every gate. To take the impact of  $OPCs$  into account, we consider  $7 \times 7$  input signal slews and output load capacitances. The measured 49 values of delay and power of each gate are then stored using the standard “liberty” format. This ensures that our libraries are compatible with existing commercial tool flows, e.g., Synopsys.

Based on the aforementioned process, we create a cell library for every  $1^\circ\text{C}$  for the range of  $T \in [25, 85]^\circ\text{C}$  under the voltage of 1.2V. Likewise, we also create a cell library for every 10mV for the range of  $V \in [0.9, 1.2]\text{V}$  under the temperature of  $25^\circ\text{C}$ . As a result, we create 91 cell libraries. Both  $T$  and  $V$  ranges are typical for processors in the 45nm and the selected fine granularity provides us with a high certainty when we later analyze the processor’s delay. Note that selecting a smaller  $T$  step would not be helpful as the accuracy of on-chip temperature sensors (which we later rely on to perform the online adaptation of  $V$ ) is typically  $>1^\circ\text{C}$ . Similarly, performing  $V$  adaptation at less than 10mV would also not be beneficial due to voltage fluctuation induced uncertainty. We select  $T_{worst}$  as  $85^\circ\text{C}$ , which is normal for commercial processors. Other  $T_{worst}$  can be easily considered.

### IV. MODELING THE TEMPERATURE-VOLTAGE CORRELATION IN PROCESSORS

To model the  $T$ - $V$  correlation, we study four kinds of processors to take various architectures and complexities into account: The state-of-the-art industry-competitive out-of-order processor from Berkeley [14] and three processors (RISC with 5 and 6 pipelines and VLIW CPU) generated from the Synopsys Processor Designer tool. In the following, we take the BOOM CPU as an example and later summarize the analysis of others. First, we synthesize the RTL using the Synopsys Design Compiler based on the baseline library (1.2V,  $25^\circ\text{C}$ ). Then, we estimate the maximum delay in the entire processor’s netlist for the whole range of  $T_{operation} \in [25, 85]^\circ\text{C}$  with

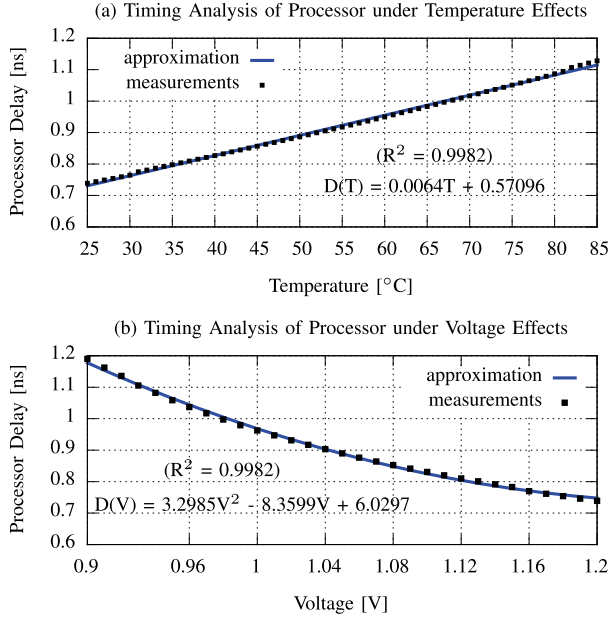


Fig. 2. Modeling the delay of BOOM processor under  $T$  and  $V$  effects.

a  $1^\circ\text{C}$  step using the Synopsys STA tool and our created libraries. Any switch in the  $CP$  due to  $T$ ,  $V$  effects will be automatically captured by the STA tool. Based on the obtained data, we model the impact of  $T$  on the processor's delay  $D(T)$ . As Fig. 2 (a) shows, the linear approximation exhibits a good agreement with data. The average error is merely 0.42% and the root-mean-square error (RMSE) is 0.0048. Analogously, we perform timing analysis of processor's netlist for whole range of  $V \in [0.9, 1.2]\text{V}$  with a 10mV step to model the impact of  $V$  on the processor's delay  $D(V)$ . As shown in Fig. 2(b), a second-order polynomial approximation exhibits a good agreement with obtained data and the average error is 0.54% along with an RMSE of 0.0056. Finally,  $T$  variation ( $\Delta T = T_{\text{worst}} - T_{\text{operation}}$ ) is correlated to the corresponding equivalent voltage as Eqs.(1, 2, 3) demonstrate. Fig. 3 shows  $V_{\text{adapted}}$  for every  $T_{\text{operation}}$ . All curve's points result in a delay that does not violate the timing constraint (i.e., the processor's delay at  $T_{\text{worst}} = 85^\circ\text{C}$  and  $V_{\text{nominal}} = 1.2\text{V}$ ). Note that only the slope of  $T$  and  $V$  curves (see Fig. 2(a, b)) contribute in modeling the  $T$ - $V$  correlation (see Eq. 2). Thus, the initial values (i.e.  $\alpha_1$  and  $\beta_2$ ) are not important. The fitting parameters vary for different processors, as Table I summarizes (all errors are within 0.55%). Thus, the processor's architecture plays a role in how the  $T$  and  $V$  affect the maximum delay.

$$D(T) = \alpha_0 T + \alpha_1, D(V) = \beta_0 V^2 + \beta_1 V + \beta_2 \quad (1)$$

$$\Delta D(T) = D(T_{\text{worst}}) - D(T_{\text{operation}}) = \alpha_0 \Delta T > 0$$

$$\Delta D(V) = D(V_{\text{nominal}}) - D(V_{\text{adapted}}) < 0$$

$$\Delta D(V) = \beta_0 (-\Delta V^2 + 2V\Delta V) + \beta_1 \Delta V = \Delta D(T) \Rightarrow$$

$$\Delta V = \frac{(2\beta_0 V + \beta_1) + \sqrt{(2\beta_0 V + \beta_1)^2 - 4\beta_0 \Delta D}}{2\beta_0}; \beta_0 \neq 0$$

$$\Delta V = \frac{(2\beta_0 V + \beta_1) + \sqrt{(2\beta_0 V + \beta_1)^2 + 4\beta_0 \alpha_0 \Delta T}}{2\beta_0} \quad (2)$$




Fig. 3. The corresponding equivalent voltage for every operation temperature in which the timing constraint is not violated.

$$V_{\text{adapted}} = V_{\text{nominal}} - \Delta V \quad (3)$$

## V. VOLTAGE ADAPTATION AT RUNTIME USING THE $T$ - $V$ CORRELATION

Once the  $T$ - $V$  correlation for the targeted processor is modeled at design time, a look-up table (LUT) that links each operation temperature with its corresponding equivalent voltage (see Fig. 3) can be built. Then, the LUT is employed at runtime to adapt the  $V$  based on current  $T$ . To achieve that, the processor's temperature is periodically read at every  $t_{\text{read}}$  from an on-chip sensor, which is often available in most of processor chips to provide the maximum  $T$  of processor at runtime.  $V$  adaptation can be rapidly performed as recent on-chip voltage regulators switches  $V$  within  $< 1\mu\text{s}$  like Intel on-the-fly regulators [15]. To sustain reliability, the  $V$  adaptation must take into account the maximum potential increase in  $T$  until the next sensor reading arrives based on its sampling rate. Otherwise, timing violations might occur due to applying a smaller  $V$  than what can be tolerated leading to unsustainable frequency. To this end, we add a safety margin ( $SM$ ) on top of the read temperatures. The  $SM$  can be estimated at design time from analyzing the maximum possible  $T$  increase when the maximum power rise occurs. Such a  $SM$  means that the selected  $V_{\text{adapted}}$  is always higher than what is actually needed. This additionally ensures that any small uncertainty in modeling (reported in Table I) will be compensated.

## VI. EVALUATION AND COMPARISONS

State of the art (e.g., [16]) typically employs a ring oscillator ( $RO$ ) as a sensor to predict the effects of  $T$  and  $V$  on delay. Therefore, we first evaluate whether such an approach could be alternatively used instead of creating cell libraries under  $T$  and  $V$  effects, as we propose. As demonstrated in Section II and Fig. 1, the same  $T$ ,  $V$  change can *differently* impact the gates' delay. This brings up the critical question of "how good is the  $RO$  in predicting the impact that  $T$  or  $V$  change actually

TABLE I  
REQUIRED PARAMETERS IN EQ. 2 FOR DIFFERENT PROCESSORS

Processor	$\alpha_0$	$\beta_0$	$\beta_1$	$D(T)$ Err.	$D(V)$ Err.
BOOM	0.0064	3.2985	-8.3599	0.42%	0.54%
RISC-5P	0.0057	2.5167	-6.5215	0.36%	0.34%
RISC-6P	0.0065	3.1048	-7.9779	0.39%	0.39%
VLIW	0.0063	3.0396	-7.7905	0.39%	0.40%




Fig. 4. Evaluating how accurate a  $RO$ -based approach in predicting the impact of temperature and voltage changes.

has on the processor’s delay?”. To investigate that, we study a typical  $RO$  of 15 stages [17]. We compare its predictions against the actual delay increase of the BOOM processor. The latter is accurately obtained by the Synopsys STA. As shown in Fig. 4, in both  $T$  and  $V$  scenarios there is a noticeable error in  $RO$ -based predictions reaching 23% and 28.5%, respectively. Such large errors demonstrate why one cannot rely on  $RO$ -based predictions to *reliably* perform temperature-guided voltage adaptation. The  $RO$  fails in representing how  $T$  and  $V$  affect the processor’s delay because  $T$  and  $V$  *non-uniformly* impact the delays of gates (see Fig. 1).

To evaluate the effectiveness of our technique compared to employing throughout a non-adaptive voltage, we build the following simulation toolchain: The cycle-accurate simulator gem5 [18] of an out-of-order processor extracts the activities of varied applications towards estimating the resulting temperature trace of each application using the Hotspot thermal simulator [19] together with the McPAT power simulator [20]. Afterwards, we integrate our technique proposed in Section V to perform the voltage adaptation based on the  $T$ - $V$  correlation modeling. For a proper compatibility with the simulated architecture in gem5, we employ the BOOM processor in our evaluation since it is an out-of-order processor. In our implementation we consider  $t_{read} = 1\text{ms}$  and we select  $SM$  as  $2^\circ\text{C}$  to be added on top of the read temperature. The  $SM$  is obtained after analyzing the transient temperature traces of all studied applications to determine the maximum possible  $T$  increase within the selected  $t_{read}$  (i.e. max  $\Delta T$  within 1ms). We employ diverse benchmarks from the PARSEC [21] and SPEC2006 [22] benchmark suites. Each benchmark is simulated on top of Linux kernel for realistic analysis.

The obtained results showed that our technique reduces the consumed power by 21%, on average, and up to 24%. This is because that it enables the processor to operate during the majority of operation time at lower voltages than the nominal (1.2V). On average, the applied  $V_{adapted}$  is around 0.97V. Reductions in power directly lead to a lower power densities and thus less generated heat. On average, the temperature is reduced by 12% and up to 14%. Fig. 5 demonstrates the dynamics in  $T$  and  $V_{adapted}$  for the “dedup” benchmark as an example. As shown, our technique reacts to the changes in temperature and accordingly adapts the voltage. This leads to employing lower  $V$  most of the time and thereby generating lower temperatures, compared to employing throughout a constant voltage of 1.2V. Note that reductions in  $T$  due to lower  $V$  may allow further reduction in  $V$  until  $T$  converges.




Fig. 5. Demonstration of the thermal behaviour of “dedup” benchmark when employing our temperature-guided voltage adaptation.

## VII. CONCLUSION AND SUMMARY

In this work, we investigated how the effects of temperature and voltage on the delay of processor can be modeled towards implementing a temperature-guided voltage adaptation technique. The proposed technique adapts the voltage during the operation of processor in which the minimum, yet sufficient voltage is selected under the effects that temperature variation has, leading to a more efficient computing.

## REFERENCES

- [1] H. Amrouch and J. Henkel, “Lucid infrared thermography of thermally-constrained processors,” in *ISLPED*, 2015, pp. 347–352.
- [2] H. Khdr, H. Amrouch *et al.*, “Aging-aware boosting,” *IEEE Transactions on Computers*, 2018.
- [3] J. Keane and C. H. Kim, “Transistor aging,” *IEEE Spectrum*, 2011.
- [4] H. Amrouch, K. Behnam *et al.*, “Optimizing temperature guardbands,” in *DATE*, 2017.
- [5] L. Yuan and G. Qu, “Analysis of energy reduction on dynamic voltage scaling-enabled systems,” *TCAD*, vol. 24, no. 12, 2005.
- [6] D. Wolpert and P. Ampadu, “Temperature effects in semiconductors,” in *Managing Temperature Effects in Nanoscale Adaptive Systems*, 2012.
- [7] Y. Chauhan, S. Venugopalan *et al.*, “BSIM,” in *ESSCIRC*, 2012.
- [8] K. Joshi, S. Mukhopadhyay *et al.*, “A consistent physical framework for N and P BTI in HKMG MOSFETs,” in *IRPS*, 2012.
- [9] V. M. van Santen, J. Martin-Martinez *et al.*, “Reliability in super- and near-threshold computing: A unified model of rtn, bti, and pv,” *TCAS-I*, vol. 65, no. 1, pp. 293–306, 2018.
- [10] “Nangate, Open Cell Library,” <http://www.nangate.com/>.
- [11] H. Amrouch, S. Mishra *et al.*, “Impact of bti on dynamic and static power: From the physical to circuit level,” in *IRPS*. IEEE, 2017, pp. CR–3.
- [12] H. Amrouch, B. Khaleghi *et al.*, “Reliability-aware design to suppress aging,” in *DAC*, 2016, pp. 1–6.
- [13] W. Zhao and Y. Cao, “New generation of predictive technology model for sub-45 nm early design exploration,” *Electron Devices, IEEE Transactions on*, vol. 53, no. 11, Nov 2006.
- [14] C. Celio, D. A. Patterson *et al.*, “The Berkeley Out-of-Order Machine (BOOM): An Industry-Competitive, Synthesizable, Parameterized RISC-V Processor,” Berkeley, Tech. Rep., Jun 2015.
- [15] E. Burton, G. Schrom *et al.*, “FIVRFully integrated voltage regulators on 4th generation Intel® Core SoCs,” in *APEC*, 2014.
- [16] C. R. Lefurgy, A. J. Drake *et al.*, “Active management of timing guardband to save energy in power7,” in *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture*, 2011.
- [17] G.-Y. Wei and M. Horowitz, “A fully digital, energy-efficient, adaptive power-supply regulator,” *JSSC*, vol. 34, no. 4, 1999.
- [18] N. Binkert, B. Beckmann *et al.*, “The Gem5 simulator,” *SIGARCH Comput. Archit. News*, 2011.
- [19] M. R. Stan, K. Skadron *et al.*, “Hotspot: a dynamic compact thermal model at the processorarchitecture level,” *Microelectronics Journal*, 2003.
- [20] S. Li, J. H. Ahn *et al.*, “The McPAT framework for multicore and many-core architectures: Simultaneously modeling power, area, and timing,” *ACM Trans. Archit. Code Optim.*, 2013.
- [21] C. Bienia, S. Kumar *et al.*, “The PARSEC benchmark suite: Characterization and architectural implications,” in *PACT*, 2008.
- [22] J. L. Henning, “SPEC cpu2006 benchmark descriptions,” *SIGARCH Comput. Archit. News*, 2006.