International Journal of Control and Automation
Vol.7, No.6 (2014), pp.159-168
http://dx.doi.org/10.14257/ijca.2014.7.6.16

Power Estimation for Alpha 21264 Using Performance Events and Impact of Ambient Temperature

Biying Zhang[1,2,*] and Gang Cui[1]

[1]*Department of Computer Science and Technology, Harbin Institute of Technology, Harbin, 150001, China*
[2]*College of Computer and Information Engineering, Harbin University of Commerce, Harbin, 150028, China*

[*] *Corresponding author: biyingzhang@163.com*

Abstract

Higher power has become one of primary obstacles for improving the performance of processors. To manage the power of processor dynamically, it is necessary to obtain the power timely and accurately. Power estimation using performance events is a better way in terms of timeliness and accuracy. We analyzed the correlation between performance events and powers of microprocessors. Only one performance event was selected for the power estimation of one functional unit based on the correlation analysis. The power of each functional unit was estimated based on simple regression analysis. In addition, the impact of ambient temperature on the accuracy for power estimation was presented. It can be seen from the experimental results that (1) the power of each functional unit has close linear correlation with a performance event and can be estimated accurately; (2) the ambient temperature has the negative impacts on the accuracy for power estimation.

Keywords: *Power estimation; performance event; dynamic power management; Alpha 21264*

1. Introduction

Higher power has become one of primary obstacles for improving the performance of processors. Dynamic power management (DPM) is a widely-used technique for reducing the power of processors [1]. In order to manage the power of processor dynamically, it is necessary to obtain the power timely and accurately. There are two ways for acquiring the power at run time: direct measurement and indirect estimation. Usually, the direct measurement is used in the laboratory environment, and rarely utilized in the final product, because current, voltage and charge measurement devices are required to be incorporated into processors [2]. Indirect estimation of power is widely used due to its low area overheads. Power can be estimated using (1) performance events captured from performance counters, (2) temperatures captured from thermal sensors. Due to the thermal inertia in microprocessor packaging, detection of temperature changes may occur significantly later than the power events causing them [3]. Hence, performance events are more effective than temperature for the estimation of power in terms of timeliness.

Generally, powers are various between functional units in the microprocessor, which leads to the variation of temperature between functional units. In order to manage power and temperature more efficiently, the power should be estimated at the micro-architectural level, *i.e.*, the power of each functional unit should be estimated. As the power estimation is carried

ISSN: 2005-4297 IJCA

out at run time, the higher speed of estimation is required. Although a large number of performance events are available in processor, most of them are tedious for the power estimation of a certain functional unit. The tedious performance events not only decrease the accuracy of power estimation but also reduce the estimation speeds. In addition, the ambient temperature of microprocessor has significant impact on the estimation accuracy.

In this paper, we analyzed the correlation between performance events and powers of microprocessors. We used Alpha 21264 as the target processor, used PTScalar as the simulation tool, and some tasks from SPEC CPU2000 and MiBench as the benchmarks. As for each functional unit, only one performance event was selected based on correlation analysis. The performance event having the maximum correlation coefficient with the functional unit was used for the estimation of this functional unit. The power of each functional unit was estimated based on simple regression analysis. We also analyzed the impact of ambient temperature on the correlation. The ambient temperatures were set at the different values, and the total powers and dynamic powers of each functional unit at each ambient temperature were traced. And then the simulated and estimated powers were compared.

The remainder of this paper is organized as follows. In Section 2, related work on the power estimation of microprocessor using performance events is presented. In Section 3, the experimental methodology is described. In Section 4, the selection of performance events using correlation analysis is explained. In Section 5, the power estimation using regression analysis is presented. In Section 6, the impact of ambient temperature on the estimation accuracy of power is shown. In section 7, this research is concluded.

2. Related Work

The power estimation of microprocessor using performance events is not a new concept. Zamani *et al.* [4] examined the relationship between power and performance events from a stochastic perspective. The autoregressive moving average (ARMA) models were used for modeling various trends between performance and power. Zamani *et al.* [5] also studied the relationship of performance events with power consumption in the context of single performance event and multiple performance events. The selection results of single event for each test application were presented. In addition, the variability of measurement for correlation coefficients was investigated. Bircher *et al.* [3] proposed the use of microprocessor performance counters for online measurement of complete system power consumption. The approach takes advantage of the trickle-down effect of performance events in microprocessors. The power consumptions in memory and disk and other subsystems outside of the microprocessor were estimated using performance events. Contreras *et al.* [6] demonstrated a first-order, linear power estimation model that uses performance counters to estimate run-time CPU and memory power consumption of the Intel PXA 255 processor. The model used a set of power weights that map hardware performance counter values to processor and memory power consumption. Kim and Chao *et al.* [2, 7] proposed an on-chip bus (OCB) performance monitoring unit (PMU) that directly captures on-chip and off-chip component activities by snooping the OCB. The actual power values were converts by online software using counter values. An optimization algorithm also was introduced to reduce the number of counters in the OCB PMU. Oh et al. [8] proposed a dynamic power estimation method using a single thermal sensor for each core in multi-core processors. And then a die temperature was reconstructed using the estimated power.

3. Experimental Methodology

We used the Alpha 21264 processor as the target processor in our experiments, which is a four-issue superscalar microprocessor with out-of-order and speculative execution [9]. The Alpha 21264 processor is composed of the following components: decoder, branch predictor, register alias table (RAT), register update unit (RUU), load store queue (LSQ), four integer ALUs (IntALU1, IntALU2, IntALU3, IntALU4), floating point adder (FPAdd), floating point multiplier (FPMul), integer registers (IntReg), floating point registers (FPReg), instruction translation lookaside buffer (ITLB), data translation lookaside buffer (DTLB), L1 instruction cache (IL1), L1 data cache (DL1) and L2 cache (L2).

PTScalar [10] was used as the simulation tool, and the power trace of each component in the processor can be obtained directly. However, the performance trace can not be tracked directly. PTScalar was modified to obtain not only the power trace but also the performance trace. Firstly, the following performance events were tracked using the modified PTScalar: (1) executed cycles; (2) the numbers of total instructions, controlling instructions, integer instructions, floating point instructions, memory access instructions and trap instructions; (3) the accessed and hit numbers of DL1, IL1, ITLB, DTLB and L2. And then, the numbers of events per cycle were obtained by dividing executed cycles.

The running frequency of PTScalar simulator was set at 3 GHz. The supply voltage was set 1.5 V. In order to analyze the impact of ambient temperature on powers, the ambient temperature was set 35℃, 45℃, 55℃ and 65℃ respectively. The performance and power trace were sampled every 5ms. Except for these parameters, others were set default values of the simulator.

We selected some tasks from SPEC CPU2000 [11] and MiBench [12] as the benchmarks. The behavioral characteristics of task can be represented by the number of performance events per cycle. Table 1 shows the average number of performance events per cycle for each benchmark. Table 2 shows the average power of the whole processor and various functional units for each workload. As the space of page is limited, some less important performance events are not given in Table 1, and some functional units with very low power are not presented in Table 2.

Table 1. Average Number of Performance Events per Cycle for each Benchmark

Bench	Tot.	Mem.	Bran.	Int.	FP	Dl1	Il1	Dl2	Dtlb	Itlb
mesa	1.07	0.5	0.11	0.35	0.11	0.48	1.2	0.08	0.5	1.2
ammp	0.72	0.35	0.15	0.2	0.02	0.33	0.76	0.05	0.35	0.76
quake	1.58	0.54	0.31	0.7	0.03	0.53	1.75	0.04	0.56	1.75
bzip	1.42	0.77	0.11	0.55	0	0.64	1.46	0.05	0.77	1.46
mcf	0.97	0.45	0.21	0.31	0	0.41	1.15	0.07	0.46	1.15
math	1.06	0.25	0.19	0.61	0.01	0.25	1.26	0.08	0.26	1.26
qsort	1.44	0.35	0.3	0.78	0.01	0.35	1.55	0.01	0.36	1.55

Table 2. Average Power of Various Elements for each Benchmark

Bench	Tot.	Bran.	RUU	LSQ	IALU	FAdd	FMul	IReg	IL1	DL1	L2
Mesa	36.88	5.41	3.59	0.95	3.53	1.32	1.27	1.92	6.41	3.67	7.92
Ammp	31.87	6.19	2.49	0.73	3.14	1.17	1.15	1.45	4.54	2.99	7.4
Quake	46.2	9.6	5.01	1.03	4.33	1.18	1.17	2.98	8.75	3.88	7.29
Bzip	40.6	5.15	4.84	1.39	4.13	1.13	1.13	2.68	7.56	4.3	7.42
Mcf	37.88	7.56	3.39	0.94	3.54	1.13	1.13	2.07	6.17	3.35	7.88
Math	36.77	6.95	3.21	0.63	3.77	1.15	1.16	1.97	6.72	2.62	7.86
Qsort	42.48	9.41	4.21	0.77	4.42	1.15	1.14	2.61	7.92	3.06	6.93

4. Selection of Performance Events

Although many performance events are available in processor, most of them are tedious for the power estimation of a certain functional unit. The tedious performance events not only decrease the accuracy of power estimation but also reduce the estimation speeds. Therefore, in our experiments, only one performance event was selected for one functional unit. For a functional unit, the performance event which has the maximum correlation with it is selected.

Set m and n be the number of functional units in the processor and the number of performance events given by processor respectively. Set P_i be the power trace of the ith functional unit, E_j be the performance event trace of the jth performance event. The performance event which has the maximum correlation with the ith functional unit can be denoted by

$$f(i) = \arg\max_{1 \le j \le n} (coeff(P_i, E_j)) \quad (1)$$

where $coeff(P_i, E_j)$ is a function which compute the correlation coefficient between the power of the ith functional unit and the jth performance event.

Table 3. Performance Event Selection for each Functional Unit

Units	Events	Coeff	Units	Events	Coeff
Decode	il1_hits	0.9981	FPAdd	Float	0.9971
Branch	Branch	0.9993	FPMul	Float	0.9597
RAT	Il1_hits	0.9980	IntReg	Total	0.9837
RUU	Total	0.9770	FPReg	Float	0.9995
LSQ	Memory	0.9774	ITLB	Itlb_hits	0.9997
IntALU1	Total	0.9741	IL1	Il1_hits	0.9997
IntALU2	Total	0.9808	DTB	Dtlb	0.9990
IntALU3	il1_hits	0.9510	DL1	dl1	0.9985
IntALU4	Integer	0.4851	L2	dl2	0.9951

The power traces which is used for the selection of performance events in this section and the regression analysis in next section were obtained on condition that the ambient temperature was set 35℃. Table 3 shows the selected performance event for each functional

unit and the correlation coefficient between them. It can be seen that each functional unit has a corresponding performance event such that the correlation between the power of the functional unit and the corresponding performance event is significantly high. The correlation coefficients for each couple except for IntALU4 are higher than 0.95, and the highest correlation coefficient is 0.999. It is obvious for the correlation between most of functional units and events. For example, the power of branch predictor is determined by controlling instructions executed per cycle, the power of floating point adder is determined by floating point instructions executed per cycle, and so forth. However, it can be seen that the performance events having highest correlation with the power of IntALU1, IntALU2 and IntALU3 are not integer instructions executed per cycle but total instructions and hitting number of L1 instruction cache per cycle. This results from the fact that not only integer instructions but also controlling instructions, memory access instructions, etc will make IntALUs active. There is no performance event having high correlation with the power of IntALU4 because most of integer operations are performed by IntALU1, IntALU2 and IntALU3, IntALU4 is inactive in most of time.

5. Regression Analysis

As shown in Table 3, a functional unit's power has close correlation with a performance event. Hence, the power of a functional unit can be estimated with the performance event having the closest correlation with it. The basic equation for estimating the power of the *i*th functional unit is expressed by

$$P_i = aE_{f(i)} + b \tag{2}$$

where P_i is the power of the *i*th functional unit, $f(i)$ is the performance event having the closest correlation with the *i*th functional unit which is presented by (1), $E_{f(i)}$ is the number of the performance event per cycle, *a* and *b* are regression coefficients.

Table 4. Regression Coefficients for each Functional Unit

Units	a	b	Units	a	b
Decode	0.2530	0.0729	FPAdd	1.8663	1.1271
Branch	21.4780	2.9556	FPMul	1.2982	1.1316
RAT	0.0416	0.0120	IntReg	1.7189	0.2167
RUU	3.0470	0.2606	FPReg	1.0991	0.1436
LSQ	1.5552	0.2053	ITLB	0.0799	0.0215
IntALU1	0.5971	0.6170	IL1	4.3018	1.2592
IntALU2	0.4670	0.5012	DTB	0.0713	0.0232
IntALU3	0.2468	0.4945	DL1	4.4474	1.5078
IntALU4	0.2971	0.4913	L2	2.2058	0.9598

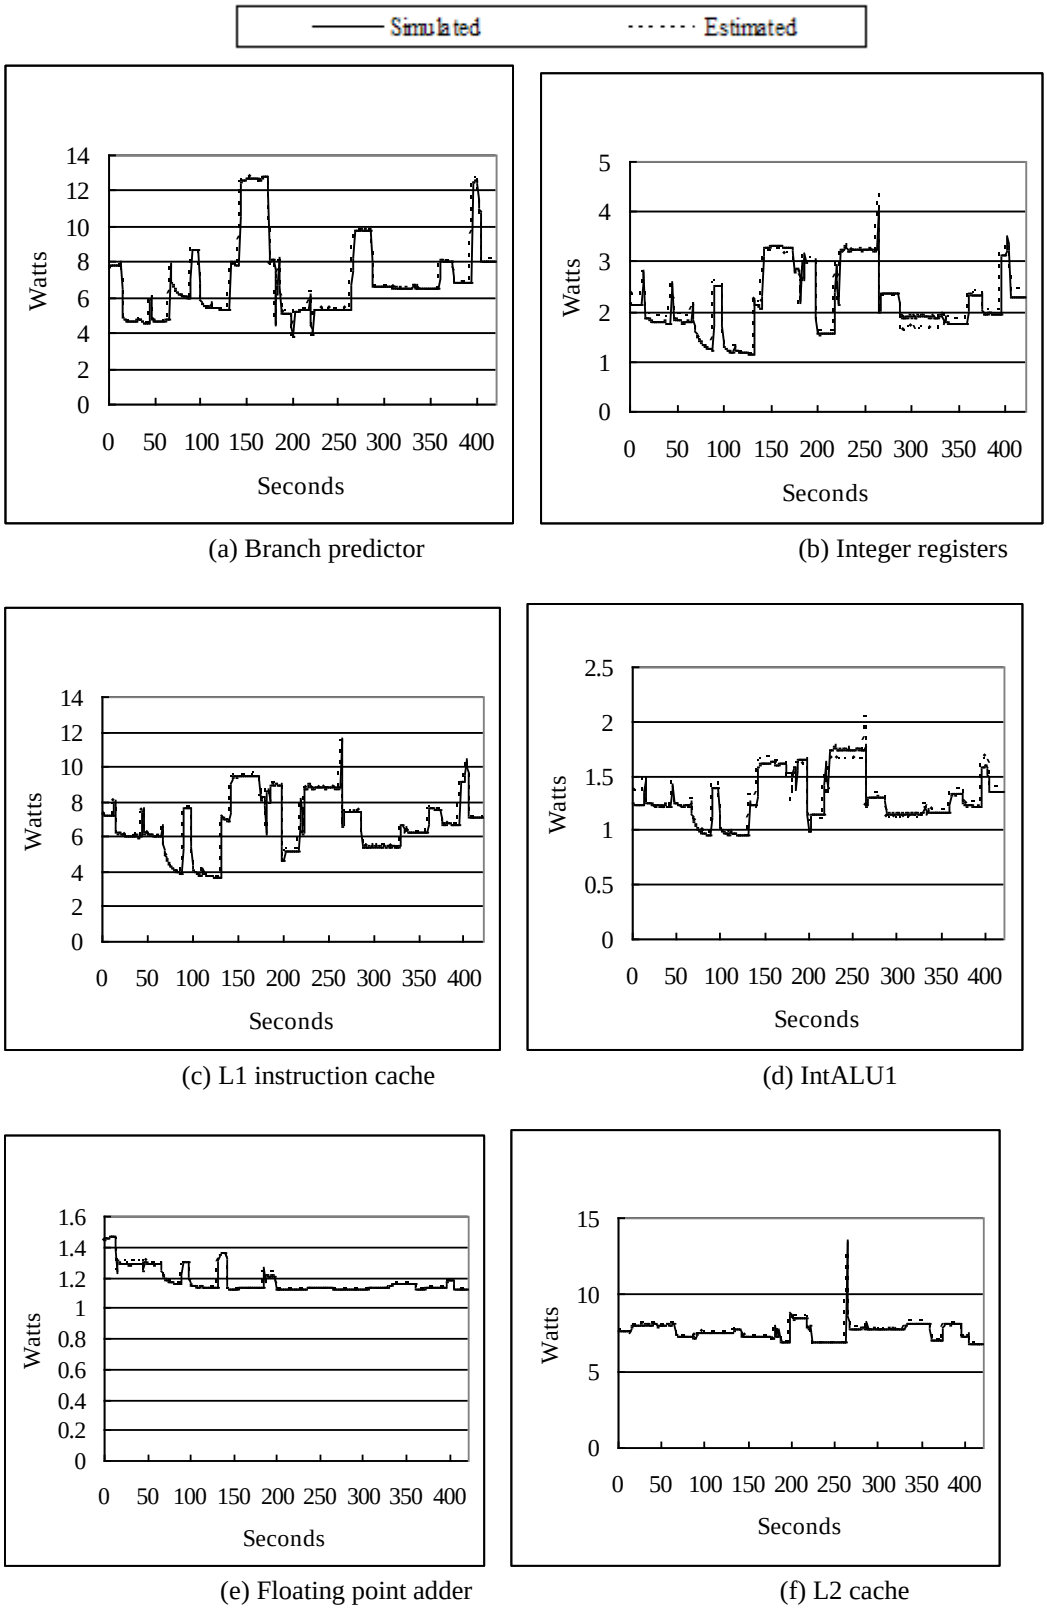

Figure 1. Comparison between Simulated and Estimated Total Power

We used the simple regression analysis to obtain the regression coefficients *a* and *b*. The regression coefficients for each functional unit are shown by Table 4.

All seven benchmarks were executed on the processor, and the simulated power and estimated power of each functional unit were traced together. Since the space of page is limited, we only presented the comparative results between the simulated and estimated power traces of six representative functional units, as shown by Figure 1. It can be seen that the estimated power trace is very close to the simulated power trace for each functional unit. Different functional unit dissipates different power. Branch predictor, L1 instruction cache and L2 cache dissipate much more power than other functional units. In addition, it also can be seen that the power fluctuation with time between different functional units is different. The power fluctuation of branch predictor, integer register, L1 instruction cache is more than that of floating point adder and L2 cache. Especially, the power of floating point adder almost does not change with time.

6. Impact of Ambient Temperature on Estimation Accuracy

In order to explore the impact of ambient temperature on the correlation between the power and the performance events, the ambient temperature of the simulator PTScalar was set 35℃, 45℃,55℃ and 65℃ respectively. The powers of functional units at various ambient temperatures were traced and were compared with the estimated power. The error rate between the simulated and the estimated power was used as the metric to represent the impact of the ambient temperature on the correlation between the power and the performance events. Figure 2 shows the error rates of estimation at the different ambient temperature. As the ambient temperature increases, the error rates of estimation for each functional unit also do. It is proved that the ambient temperature has the reversed impact on the estimation. It also can be seen that the sensitivities of error rate of each functional unit to the temperature are different. For example, the sensitivity of L2 cache to temperature is highest. The error rate of L2 cache at 35℃ is only 0.46%, while that at 65℃ reaches 27.16%. The sensitivity of RUU to temperature is smallest. The error rates of RUU at all temperatures are almost same. In addition, it can be seen that the estimation accuracy for different functional units is different. The estimation accuracy for ITLB is highest, while that for L2 cache is smallest.

The dynamic powers of each unit at each ambient temperature were also traced, and estimated with the performance events. The error rates of estimation of dynamic power at different ambient temperature are depicted by Figure 3. It can be seen that the error rates of estimation at the different temperature are same for each functional unit. In other words, the temperature has no impact on the dynamic power. The error rates are different with respect to functional units. Except for RUU and IntALU4, the error rates for each functional unit are less than 5%. The minimum error rate is only 0.36%, which happens on the floating point adder.

Figure 4 depicts the proportion of the dynamic power to the total power for each functional unit. It can be seen that the proportion of the dynamic power to the total power decreases as the ambient temperature increases. As presented by Figure 3, the dynamic power is independent from the ambient temperature, hence, it can be deduced that the static power depend on temperature. Higher temperature will result in more static power, and then more total power.

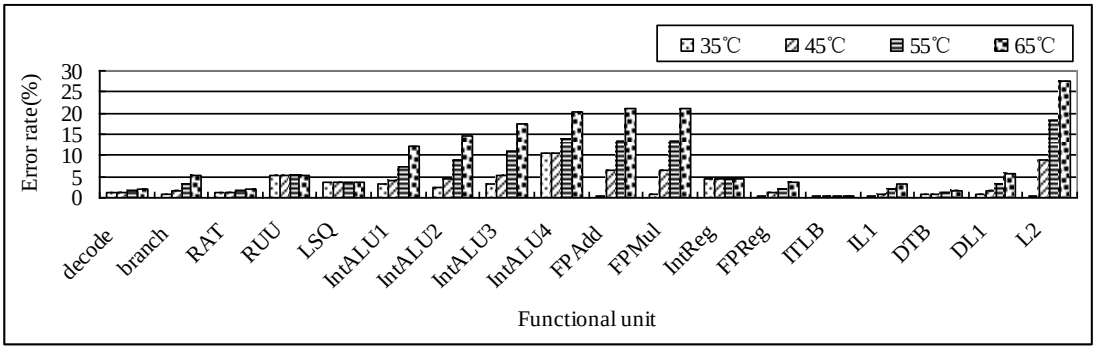

Figure 2. Error Rates of Power Estimation at the Different Ambient Temperature

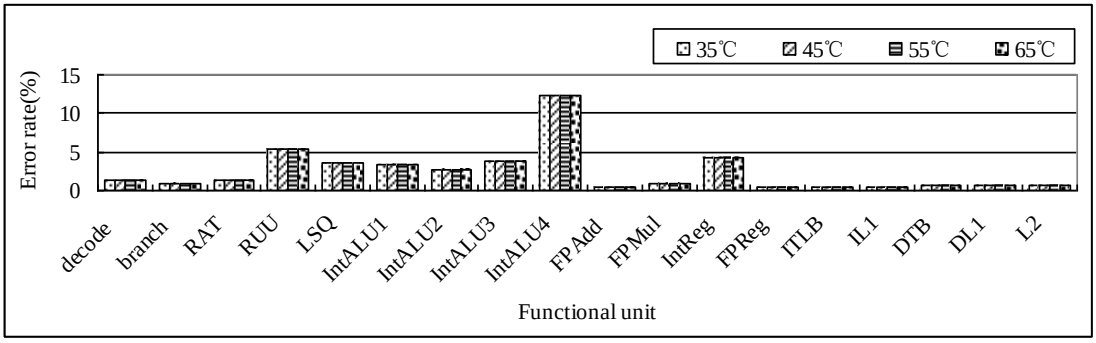

Figure 3. Error Rates of Dynamic Power Estimation at the Different Ambient Temperature

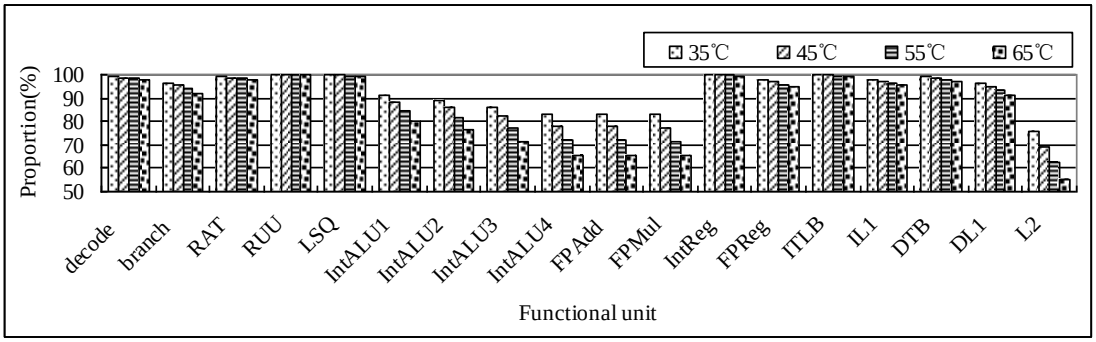

Figure 4. Proportion of the Dynamic Power to the Total Power

7. Conclusion

The correlations between the power and the performance events are analyzed through experiments. From the experimental results, it can be seen that the power of a functional unit in the processor has higher linear correlation with one performance event. Therefore, the power of each unit can be estimated accurately using the performance events based on simple regression analysis. In addition, the ambient temperature has great impacts on the estimation accuracy of total powers and has no impacts on that of dynamic powers. Higher ambient temperature will reduce the estimation accuracy.

References

[1] H. Shen, Y. Tan, J. Lu, Q. Wu and Q. Qiu, "Achieving Autonomous Power Management Using Reinforcement Learning", ACM Transactions on Design Automation of Electronic Systems, vol. 18, no. 2, **(2013)**.
[2] Y. Kim, S. Park, Y. Cho and N. Chang, "System-Level Online Power Estimation Using an On-Chip Bus Performance Monitoring Unit", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 30, no. 11, **(2011)**, pp. 1585-1598.
[3] W. L. Bircher and L. K. John, "Complete System Power Estimation Using Processor Performance Events", IEEE Transactions on Computers, vol. 61, no. 4, **(2012)**, pp. 563-577.
[4] R. Zamani and A. Afsahi, "Adaptive estimation and prediction of power and performance in high performance computing", Computer Science - Research and Development, vol. 25, no. 3, **(2010)**, pp. 177-186.
[5] R. Zamani and A. Afsahi, "A Study of Hardware Performance Monitoring Counter Selection in Power Modeling of Computing Systems", Proceedings of International Green Computing Conference (IGCC), **(2012)**, pp. 1-10.
[6] G. Contreras and M. Martonosi, "Power Prediction for Intel XScaler Processors Using Performance Monitoring Unit Events", Proceedings of International Symposium on Low Power Electronics and Design (ISLPED), **(2005)**, pp. 221-226.
[7] Y. Cho, Y. Kim, S. Park and N. Chang, "System-level power estimation using an on-chip bus performance monitoring unit", Proceedings of IEEE/ACM International Conference on Computer-Aided Design (ICCAD), **(2008)**, pp. 149-154.
[8] D. Oh, N. S. Kim, C. C. P. Chen, A. Davoodi and Y. H. Hu, "Runtime Temperature-Based Power Estimation for Optimizing Throughput of Thermal-Constrained Multi-Core Processors", Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC), **(2010)**, pp. 593-599.
[9] R. E. Kessler, "The Alpha 21264 microprocessor", IEEE Micro, vol. 19, no. 2, **(1999)**, pp. 24-36.
[10] W. Liao, L. He and K. M. Lepak, "Temperature and Supply Voltage Aware Performance and Power Modeling at Microarchitecture Level", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 24, no. 7, **(2005)**, pp. 1042-1053.
[11] SPEC CPU2000 V1.3 [Online], Available: http://www.spec.org.
[12] MiBench Version 1.0 [Online], Available: http://www.eecs.umich.edu.

Authors

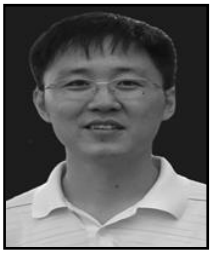

Biying Zhang

Biying Zhang received M.S. degree in computer software and theory from Harbin Engineering University, Harbin, China, in 2005, and is pursuing the Ph.D. degree in computer architecture from Harbin Institute of Technology, Harbin, China. He is currently a lecturer at College of Computer and Information Engineering, Harbin University of Commerce, China. His research interests include thermal/power optimization, thermal-aware performance evaluation and thermal-aware scheduling. **(Corresponding author of this paper)**

Gang Cui

Gang Cui is currently a professor at Department of Computer Science and Technology, Harbin Institute of Technology, Harbin, China. He serves as a senior member of Fault-Tolerance Computing Committee under China Computer Federation (CCF). His research interests include computer architecture, high-dependability computing, fault-tolerance computing, embedded computing, and internet of things.