

Accurate Characterization of the Variability in Power Consumption in Modern Mobile Processors

Bharathan Balaji, John McCullough, Rajesh K. Gupta, Yuvraj Agarwal
University of California, San Diego
{bbalaji, jmccullo, gupta, yuvraj} @ cs.ucsd.edu

Abstract

The variability in performance and power consumption is slated to grow further with continued scaling of process technologies. While this variability has been studied and modeled before, there is lack of empirical data on its extent, as well as the factors affecting it, especially for modern general purpose microprocessors. Using detailed power measurements we show that the part to part variability for modern processors utilizing the Nehalem microarchitecture is indeed significant. We chose six Core i5-540M laptop processors marketed in the same frequency bins – thus presumed to be identical – and characterized their power consumption for a variety of representative single-threaded and multi-threaded application workloads. Our data shows processor power variation ranging from 7% - 17% across different applications and configuration options such as Hyper-Threading and Turbo Boost. We present our hypotheses on the underlying causes of this observed power variation and discuss its potential implications.

1 Introduction

Variability in microelectronic manufacturing leads to variation in the threshold voltages of transistors in a chip, which in turn affects the power consumption and the maximum operating frequency of the resulting chips. Variability in manufacturing stems not only from the physical processes involved but also secondary factors such as processing temperature, wafer properties, wafer polishing and wafer placement [4]. Variability in microelectronics continues to grow [8].

Process variation has already affected several aspects of processor design, manufacture, testing and also their use. Transistor models, for example, have to now consider the variations caused by different effects in the underlying device models [5, 16]. Researchers and circuit designers have developed test frameworks to measure

and characterize the effects of process variation, as well as proposed techniques to reduce the effects of variability [3, 7, 19]. While these techniques are geared towards reducing the effects of variation, there is little empirical information – especially that is publicly available – about their effectiveness in reducing power variability in modern microprocessors.

We make the following contributions in this paper. First, we present fine-grained power measurements for six instances of a modern Intel Nehalem based Core i5-540M microprocessor, across a wide variety of modern application workloads. We show that the variability across parts is indeed significant – ranging from 7% to 17% – and is dependent on application characteristics as well as processor configuration options such as Hyper-Threading and Turbo Boost. Second, we present our hypotheses for the causes of this underlying variation.

2 Related work

Nassif [14] provides a general overview of the trends in process variation and also identifies reasons why characterization is key for advancement of processor design. Several studies have performed measurements on custom test chips [1, 15, 21], while system level characterization of variability has also been done in a few cases [20]. Wanner et al. [20] characterize the power variability across multiple ARM Cortex M3 based microcontrollers, reporting 5x variation in sleep power and <10% variation in active power, which can be leveraged to adaptively duty-cycle sensor nodes. In our prior work [13], we have identified limitations of power models due to the variability in the processor and memory components. We proposed finer grained instrumentation in platforms for online characterization of power consumption.

Since the focus of this paper is on commercially available Intel processors, we differentiate from the prior work which has also studied similar processors. Charles et al. [6] investigate the effectiveness of Turbo Boost on

numerous SPEC2006 benchmarks running on a Core i7 Nehalem processor by reporting the speedup obtained due to Turbo Boost for both CPU intensive as well as memory intensive benchmarks. There was, however, no measurement study of the power consumption of the processor. Le Seur et al. [11] provide a comparison study between executing applications at a lower frequency state but for a longer time period (slowdown) as opposed to a higher frequency state and transitioning to sleep sooner (shutdown). They also study the performance gains obtained in different P-states (frequency of operation) as well as in Turbo Boost, and compare the power consumption of the processor with different levels of C-states (sleep modes) enabled. The authors report total system power, and use it to estimate the processor power.

In contrast to prior work, this paper examines the part to part power variability in modern laptop class microprocessors and the factors affecting it. To the best of our knowledge, this is the first attempt at such a detailed variability characterization.

3 Methodology

In this section we describe our overall measurement setup and methodology. Isolating processor power from power consumed by the entire platform is a challenging problem. In the presence of variations, this is even more challenging and must rise above any measurement errors.

Our measurements are based on six identical dual-core Intel Core i5-540M parts that feature both Hyper-Threading and Turbo Boost technologies. The processor utilizes the Nehalem architecture (32nm), supports six sleep states (C-states) and ten frequencies (P-states) ranging from 1.33Ghz to 2.53Ghz and has maximum thermal design power (TDP) of 35W. Our test setup comprises of two identical laptop-class development platforms, called *Calpella*, from Intel Labs which are highly instrumented with over fifty current sense resistors to isolate subsystem power. To isolate the CPU power we combine the measurements from three independent supply lines feeding different parts of the processor. As shown in Figure 1, we collect power measurements and initiate experiments from separate ‘Measurement’ and ‘Harness’ machines respectively so as to reduce any experimental errors. We acquire the voltage and current measurements using a high precision USB 6218 NI-DAQ at 250kSA/s multiplexed across different channels.

Our test harness uses Linux `userspace` CPU governor to control the frequency of the target processor and uses Linux `cset` to set the core affinity to run our single core experiments on specific cores. For our Core i5-540M process we use four out of a total ten available frequencies – 1.2Ghz (lowest), 1.73Ghz, 2.13Ghz and 2.63Ghz (highest) – to reduce experimentation time, with




Figure 1: Experimental setup.

corresponding voltages at each frequency determined by the hardware. We use both SPEC CPU 2006 [17] and Parsec [2] benchmarks for our characterization.

3.1 Eliminating Measurement Errors

We have taken special care to ensure that the variation in values observed does not manifest due to measurement errors. For example, when we switch processors it is difficult to ensure uniform application of thermal paste on the processor package. We account for any effects due to differences in thermal dissipation by re-running the experiments for every processor after removing and plugging it back in, a procedure we call *processor swap*. Furthermore, despite using multiple CPU sets it is possible that OS scheduling decisions or interaction with other system processes may lead to variation, and hence we performed a system *reboot* and repeated the experimental runs for each processor. These additional experimental runs for processor swap and reboots are included in the results and thus are part of the standard deviation. Finally, to account for any effects due to the platform itself, we performed the experiments on identical Capella platform. We observed that the results obtained from the two platforms are consistent, with each processor giving very similar power values for each of the benchmarks.

4 Results

For the purposes of our experiments, we consider the percentage variation for a benchmark as the largest difference between the mean power consumption of each processor execution of that benchmark. We analyze variation across both single threaded benchmarks as well as more complex multi-threaded workloads.

4.1 Serial Benchmarks - SPEC CPU2006

Turbo Off, C States On: We start by characterizing the processor power variation for the simplest configuration by disabling Hyper-Threading and Turbo Boost modes in the BIOS, while still enabling all sleep states (C-states).




Figure 2: Power consumption of six Intel Core i5-540M processors for SPEC CPU 2006 benchmarks with Turbo Boost and Hyper-Threading disabled, C states enabled at 2.53GHz. Power variation ranges from 12% to 17%.

	bzip2	milc	povray	soplex	sphinx3
P1	0.455	0.014	0.275	0.081	0.014
P2	0.025	0	0	1.4E-8	0
P3	0.263	0	4.9E-8	0	0
P4	0.135	0	3.4E-5	0.001	0
P5	0.709	0	0.016	1.2E-5	9.6E-7
P6	0.455	1	0.827	0.662	1

Table 1: P-values for benchmarks with non-zero standard deviation in Figure 2.

We run benchmarks on one core at a time using `cset` and utilize the userspace governor to fix the operating frequency to 2.53Ghz (highest). We ran nineteen SPEC CPU 2006 benchmarks on our six test processors.

Figure 2 plots the mean power consumption of each processor for the different benchmarks. We observe that for the benchmarks with high variations (*bzip2*, *povray*, *soplex*) the standard deviation across runs is also high, and thus, the actual process variation may be lower than the measurements indicate. Table 1 shows the p-values of the benchmarks which have non-zero standard deviation in Figure 2. The p-values are calculated by comparing the power measurements for each processor against the average power measurement of all six processors for each benchmark. We can see that p-values for all benchmarks except *bzip2* are low, indicating that the variation we observe is independent of any measurement errors. Furthermore, we ensured that our observations are indeed caused by process variation since processors perform consistently across benchmarks – e.g. processors P2, P5 and P6 have relatively high power consumption and P3 has a lower than average power consumption.

Figure 3 shows the change in power variation for all the benchmarks in different processor P-states. Each

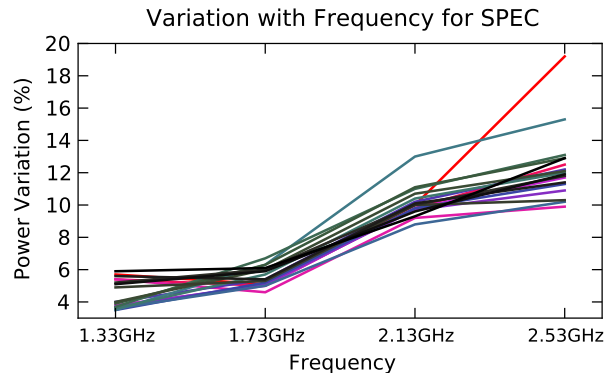


Figure 3: Power variation across six Intel Core i5-540Ms with change in P-States for SPEC CPU 2006 benchmarks. Turbo Boost and Hyper-Threading disabled.

line represents a single benchmark with the maximum variation across all the processors plotted for each frequency (P-state). We observe that all of the benchmarks show a similar trend of increased power variation with increase in frequency. At 2.53GHz, the variation is maximum – ranging from 12% to 20% depending on the particular benchmark. In higher P-states, the supply voltage also increases, which in turn causes the leakage power to grow exponentially while the dynamic power increases quadratically. It has been shown that leakage power is a more dominant factor than active power consumption as a cause for higher power variability [20]. We speculate that the increase in power variation at higher frequencies, and hence, higher voltages is because of this disproportional increase in leakage power.

Turbo On, C States On: With Turbo Boost on, the results are shown for operating frequency of 2.53GHz. The results are shown as the “Turbo on, C-states on” case in




Figure 4: Power variation across six Intel Core i5-540M processors for SPEC CPU 2006 benchmarks with different CPU configurations. Hyper-Threading is disabled.

Figure 4. Interestingly, we see noticeably lower power variation than that at the highest frequency with Turbo Boost disabled (Figure 3). Based on the trend in variation across the different frequencies of operation (P-states) with Turbo Boost disabled, we would have expected a higher variation at the increased frequency (and voltage) of operation with Turbo Boost enabled.

We hypothesize that the non-intuitive results we observe is because in Turbo Boost mode the processor aggressively switches off other cores to increase the voltage and frequency of a single core in order to maintain the thermal budget of the chip. As more circuits have been disabled, the leakage power dissipated by these circuits has been eliminated. Therefore, as the leakage component of power consumption has reduced, the power variation has decreased as well. To test our hypothesis, we disable sleep states (C-states) in the next set of measurements which should increase the parts of the chip that remain powered on, leading to a significant increase in the leakage power consumed. The higher leakage power will in turn lead to an increase in power variation.

C states Off: In this set of experiments, we keep the settings of the Turbo Boost mode On/Off experiments, except disabling all processor C-states. As expected, the power consumption increases, with the increment varying from 1% to 17% depending on the benchmark and the processor. Figure 4 plots the power variation across benchmarks with four CPU configurations - Turbo Boost on/off, C-states on/off. We observe that there is a noticeable increase in power variation when C-states are disabled when Turbo Boost is On, supporting the hypothesis that the variation is mainly caused due to the leakage power component of the total power consumption. However, when Turbo Boost is Off, there are a few benchmarks which show exception to the rule.

4.2 Parallel Benchmarks

While the SPEC CPU results presented earlier are sequential workloads meant to stress the CPU, we now




Figure 5: Power variation across six Intel Core i5-540M processors for PARSEC benchmarks with different CPU configurations. C-states are enabled.

present power variability results for a suite of twelve parallel benchmarks from the PARSEC suite which are representative of modern multi-threaded workloads. Figure 5 shows the power variation for the PARSEC suite for three CPU configurations as mentioned below.

Turbo Boost On, Hyper-Threading On: We start with both Turbo Boost (TB) and Hyper-Threading (HT) enabled, which provides a total of four cores and is the highest performance configuration. Even with thread scheduling uncertainties, the trends seen in power values are consistent with those of the SPEC CPU benchmarks. We observe an average variation of about 10% between the six tested processors across different benchmarks.

Turbo Boost On, Hyper-Threading Off: We next disable HT and measure power for PARSEC benchmarks. Surprisingly, the power variation increases with HT disabled as seen in Figure 5. We believe this is because of the drop in frequency of operation due to HT; however this hypothesis remains to be validated by measuring the processor frequencies during benchmark runs.

Turbo Boost Off, Hyper-Threading Off: With Turbo boost mode disabled, we again observe the expected decrease in power. Power variation increases to 14% - 17%, consistent with the results we see with the serial benchmark experiments. Figure 5 summarizes the variation seen from the three sets of experiments.

5 Discussion

We have measured processor power variation of 7-17% depending on configuration and application between identical processors at the same frequency of operation. Moreover, the variation we have observed is likely to get worse with future technology processors [8].

Based on our measurements on laptop processors, it is likely that such power variability will exist for server processors using the same process technologies. If that is the case, variation aware job scheduling, especially in multi-socket servers could lead to improvements in energy efficiency. For battery powered devices such as lap-

tops, application adaptations that account for CPU power variation could lead to better battery lifetimes as has been previously proposed for small sensor nodes [20].

Prior work has explored mechanisms to exploit variation using system level simulators [9, 18]. These efforts utilize models to characterize the underlying process variations [16]. Given the inherent lack of accuracy in such models due to variation, experimental measurements are essential to understanding variation and variation induced architectural innovations.

In particular, awareness of variation effects can be used at multiple levels to improve energy efficiency. Optimizations based on application knowledge have been shown to be effective in previous work [10, 12]. Using our data, researchers can get insights into how applications perform in various processor configuration modes, and the information can be leveraged to build more optimized power management techniques. The dataset for this paper is available at: http://mes1.ucsd.edu/site/pubs/HotPower12_dataset.tgz.

6 Conclusion

We have presented a characterization of the variation in power consumption of Intel Nehalem class processors, in different configurations, for various representative benchmarks. Our data analysis reveals several surprising results. First, we show that multiple instances of parts binned into the same frequency bin exhibit power variations of 12% to 17%, for the base configuration depending on the characteristics of the benchmark. Second, we observe that changing P-states as well as enabling or disabling architectural features such as Turbo Boost, Hyper-Threading and C-states affect power variation differently. Specifically, we observe higher variation at higher P-states, ranging from 5% at the lowest to 17% at the highest frequency. Disabling Turbo Boost and C-states also cause the power variation to increase significantly. Since process variation is expected to increase, detecting its extent and harnessing it using adaptive operating systems and applications will become essential to improve system performance and energy efficiency.

7 Acknowledgements

We wish to thank Jaideep Chandrashekar for providing access to the instrumented Intel platforms. This work is supported in part by NSF grants SHF-1018632 and CCF-1029783, a Calit2 grant and a grant from Intel Research.

References

- [1] AGARWAL, K., AND NASSIF, S. Characterizing Process Variation in Nanometer CMOS. In *Proc. of ACM/IEEE DAC* (2007).
- [2] BIENIA, C. *Benchmarking Modern Multiprocessors*. PhD thesis, Princeton University, January 2011.
- [3] BORKAR, S. Designing Reliable Systems from Unreliable Components: The Challenges of Transistor Variability and Degradation. In *Proc. of IEEE MICRO'05* (2005).
- [4] BOWMAN, K., DUVALL, S., AND MEINDL, J. Impact of Die-to-Die and Within-Die Parameter Fluctuations on the Maximum Clock Frequency Distribution for Gigascale Integration. *IEEE Journal of Solid-State Circuits* (2002).
- [5] CAO, Y., AND CLARK, L. Mapping Statistical Process Variations Toward Circuit Performance Variability: An Analytical Modeling Approach. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2007).
- [6] CHARLES, J., JASSI, P., ANANTH, N., SADAT, A., AND FEDOROVA, A. Evaluation of the Intel[®] Core[™] i7 Turbo Boost Feature. In *Proc. of IEEE IISWC'09* (2009).
- [7] CHEN, T., AND NAFFZIGER, S. Comparison of Adaptive Body bias and Adaptive Supply Voltage for Improving Delay and Leakage under the Presence of Process Variation. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* (2003).
- [8] DARCY, D., AND KEMERER, C. International Technology Roadmap for Semiconductors, 2009 Edition.
- [9] HERBERT, S., AND MARCULESCU, D. Variation-Aware Dynamic Voltage/Frequency Scaling. In *Proc. of IEEE HPCA'09* (2009).
- [10] JAVAID, H., SHAFIQUE, M., HENKEL, J., AND PARAMESWARAN, S. System-level Application-aware Dynamic Power Management in Adaptive Pipelined MPSoCs for Multimedia. In *Proc. of IEEE/ACM ICCAD'11* (2011).
- [11] LE SUEUR, E., AND HEISER, G. Slow Down or Sleep, That is the Question. In *Proc. of USENIX ATC'11* (2011).
- [12] LIU, X., SHENOY, P., AND CORNER, M. Chameleon: Application-level Power Management. *IEEE Transactions on Mobile Computing* (2008).
- [13] MCCULLOUGH, J., AGARWAL, Y., CHANDRASHEKAR, J., KUPPUSWAMY, S., SNOEREN, A., AND GUPTA, R. Evaluating the Effectiveness of Model-based Power Characterization. In *Proc. of USENIX ATC'11* (2011).
- [14] NASSIF, S. Process Variability at the 65nm Node and Beyond. In *Proc. of IEEE CICC'08* (2008).
- [15] PANG, L., QIAN, K., SPANOS, C., AND NIKOLIC, B. Measurement and Analysis of Variability in 45nm Strained-Si CMOS Technology. *IEEE Journal of Solid-State Circuits* (2009).
- [16] SARANGI, S., GRESKAMP, B., TEODORESCU, R., NAKANO, J., TIWARI, A., AND TORRELLAS, J. Varius: A Model of Process Variation and Resulting Timing Errors for Microarchitects. *IEEE Transactions on Semiconductor Manufacturing* (2008).
- [17] SpecCPU2006. <http://www.spec.org/cpu2006>.
- [18] TEODORESCU, R., AND TORRELLAS, J. Variation-aware Application Scheduling and Power Management for Chip Multiprocessors. *ACM SIGARCH Computer Architecture News* (2008).
- [19] UNSAL, O., TSCHANZ, J., BOWMAN, K., DE, V., VERA, X., GONZALEZ, A., AND ERGIN, O. Impact of Parameter Variations on Circuits and Microarchitecture. In *Proc. of IEEE MICRO'06* (2006).
- [20] WANNER, L., BALANI, R., ZAHEDI, S., APTE, C., GUPTA, P., AND SRIVASTAVA, M. Variability-Aware Duty Cycle Scheduling in Long Running Embedded Sensing Systems. In *Proc. of IEEE DATE'11* (2011).
- [21] ZHAO, W., LIU, F., AGARWAL, K., ACHARYYA, D., NASSIF, S., NOWKA, K., AND CAO, Y. Rigorous Extraction of Process Variations for 65nm CMOS Design. *IEEE Transactions on Semiconductor Manufacturing* (2009).