T-DVS: Temperature-aware DVS based on Temperature Inversion Phenomenon

Jinsoo Park, Hojung Cha
Dept. of Computer Science,
Yonsei University, Korea
{jspark, hjcha}@yonsei.ac.kr

ABSTRACT
Dynamic Voltage and Frequency Scaling (DVFS) is a widely used methodology to reduce the power consumption of mobile devices. This scheme performs frequency scaling in accordance with a specific governor and sets an operating voltage to be paired with the frequency. Temperature is one of the critical parameters affecting device operation. Practically, a guard-band exists in the operating voltage to ensure safe processor operation even at the worst temperature. DVFS can be optimized in terms of operating voltage under nominal conditions. In this paper, we propose a Temperature-aware DVS (T-DVS) that aggressively reduces the voltage guard-band. We explore the opportunity of providing the minimum operating voltages for frequencies at different temperatures and realize a dynamic voltage control scheme to optimize power consumption. The effectiveness of T-DVS is validated under various thermal conditions by using multi-core application processor. We experimentally observe that T-DVS leads to voltage gain without performance degradation regardless of both thermal conditions and chip characteristics. We show by using off-the-shelf smartphones that the voltage gain achieved by the scheme results in battery lifetime increment.

Keywords
DVFS; Power; Temperature.

1. INTRODUCTION
Recent mobile devices employ high performance multi-core application processors (AP). The performance increment is accompanied by additional power consumption which constrains mobile devices energy-wise. Power optimization is therefore a key issue and diverse methodologies have been developed to reduce power consumption. Dynamic Voltage and Frequency Scaling (DVFS) is a widely employed methodology that performs frequency scaling by using a specific governor in an AP. The scheme pairs pre-determined operating voltages with selected frequencies. The operating voltage encompasses a guard-band to ensure stable execution under various operating conditions, such as temperature, voltage fluctuation, and circuit aging. The guard-band is set conservatively to support the worst-case operating conditions.

ISLPED '16, August 08-10, 2016, San Francisco Airport, CA, USA
© 2016 ACM. ISBN 978-1-4503-4185-1/16/08…$15.00.
DOI: http://dx.doi.org/10.1145/2934583.2934631

Voltage reduction, too, is an effective means for improving energy efficiency. Prior work have attempted to lower operating voltage by minimizing the guard-band via diverse techniques such as dynamic voltage reduction using error correction [1], [2], and guard-band reduction with predictions of voltage noise caused by program behavior [9].

In mobile devices, temperature is one of the critical parameters affecting device operation. Most mobile devices adopt Dynamic Thermal Management (DTM) to protect the system from various temperature-related problems. The mobile AP in a smartphone encounters various temperatures resulting from device-internal and –external conditions during device operation. The guaranteed temperature usually ranges from -25 to 125°C [12]. In DVFS, a thoroughly-tested unique voltage is adopted for each frequency in this temperature range. This unique voltage guarantees safe system operation even at the worst-case temperature. However, in general use cases of mobile devices, the AP remains at normal temperatures for most of the time, while the power consumed by the guard-band that guarantees the safe voltage is potentially wasteful.

In this paper, we attempt to minimize this unnecessary voltage guard-band by considering the relationship between voltage and temperature based on the temperature inversion phenomenon [8]. According to the temperature inversion phenomenon, the operating voltage is changed based on the on-chip temperature, which provides an opportunity to reduce power consumption. To achieve our goal, we seek solutions to two challenging problems. First, we obtain the minimum operating voltage for various temperature conditions such that system stability is unaffected. Second, for ensuring robustness of the solution, we develop means to compensate for various chip characteristics. We propose a Temperature-aware DVS (T-DVS) as a solution to the aforementioned problems. Given frequency and temperature, T-DVS estimates appropriately the optimum operating voltage despite chip variation. We show that T-DVS operates adequately on the commercial AP and demonstrate that the rescaled voltage, indeed, leads to a decrease in power consumption.

2. TEMPERATURE INVERSION
All clock-driven semiconductors are designed to target a frequency with a driving voltage called nominal voltage. Mobile APs typically operate over a wide frequency range that includes the target frequency. The available frequency range is determined by a characterization test after chip manufacturing. *Over-drive* is the case where frequency and voltage higher than the design-target are used. *Under-drive* is the opposite case. According to the temperature inversion phenomenon, the minimum voltage level that drives silicon circuits changes according to circuit

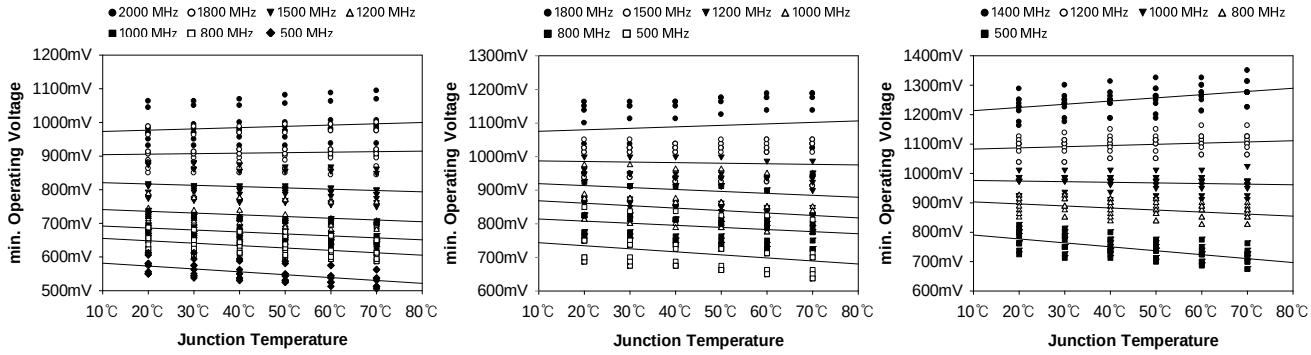

(a) Exynos 7 (14nm) **(b) Exynos 5 (20nm)** **(c) Exynos 4 (28nm)**

Figure 1 Operating voltage changes vs. temperature changes on Exynos APs

temperature [8]. In the over-drive case, operating voltage must be increased when temperature increases. In the under-drive case, operating voltage can be decreased when temperature increases.

To reproduce temperature inversion in a real device, we conducted experiments using commercial APs. The experiments involved measuring the minimum operating voltage V_{min} of a CPU at various on-chip temperatures and observing the changes in V_{min}. We used three types of mobile APs, Samsung Exynos 4412, Exynos 5433, and Exynos 7420. They are manufactured using 28 nm, 20 nm and 14 nm processes and have nominal voltages of 1.0V, 0.95V, and 0.8V, respectively. We used the AP tester equipment and AP evaluation boards for the experiment (see Section 4.1 for details). The AP tester controls the test procedure using the AP evaluation board and facilitates ambient temperature adjustment. Figure 1 shows the changes in V_{min} for specific frequencies according to on-chip temperatures. Temperature inversion, indeed, appears in the graph; V_{min} increases with an increase in temperature at high frequencies (over-drive), and it increases with a decrease in temperature at low frequencies (under-drive). The phenomenon was observed commonly in the three aforementioned APs, which were manufactured using different processes.

The experiments indicate that temperature inversion occurs regardless of the chip process variation. In terms of AP operating conditions, the highest temperature is the worst condition in over-drive and the lowest temperature is the worst in under-drive. In conventional DVFS, the V_{min} that is determined at the worst temperature is used in the voltage table as the unique voltage for a frequency. This fixed voltage is sufficient for system operation in the worst case, but it is *excessive* at normal temperatures. According to the experiment, the average operating temperature is about 40°C, and the expected voltage gains based on the average temperature are 12.5-37.5mV at high frequencies and 25-50mV at low frequencies.

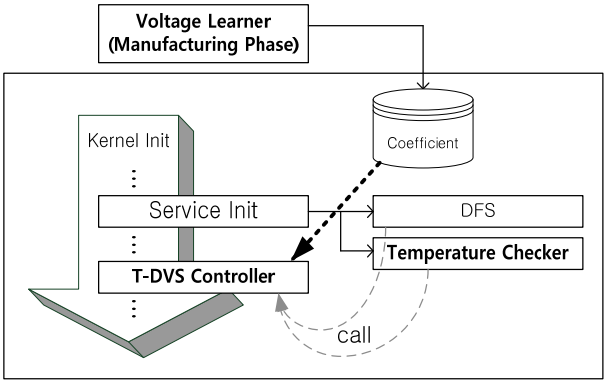

Figure 2 Overview of T-DVS software interconnection

3. T-DVS SYSTEM

T-DVS performs aggressive voltage control for power saving in temperature ranges. In the preceding experiment, we observed that V_{min} varies depending on chip variations, hence, we should consider the characteristics of individual chipsets to estimate voltage gain. T-DVS is designed to control voltage adaptively considering chipset characteristics.

T-DVS consists of three functional modules; voltage gain learner, runtime temperature checker, and DVS controller. Figure 2 shows an overview of the scheme. The voltage gain learner obtains the coefficient representing various characteristics of chipsets. The runtime temperature checker measures CPU temperature in runtime via thermal sensors in an AP. The DVS controller makes the decision to change voltage with reference to given parameters such as frequency, temperature, and coefficient of chip characteristics.

3.1 Voltage Gain Learner

Voltage gain learning is conducted in the chip manufacturing phase to acquire a voltage gain by considering semiconductor process variations. The semiconductor manufacturing line provides an adequate environment for voltage gain learning because the test equipment in the manufacturing phase allows for temperature control. Figure 3 shows an overview of voltage gain learning. First, the four V_{min} VH1, VH2, VL1, and VL2 are

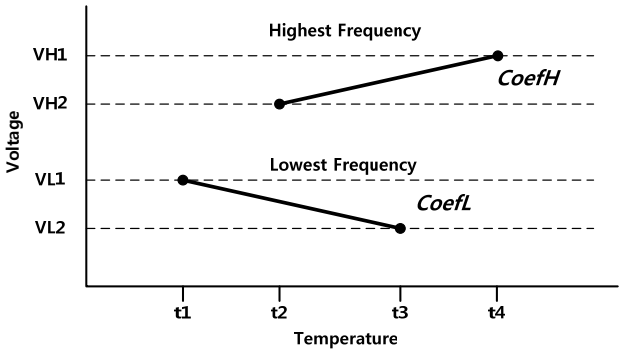

Figure 3 Voltage gain learning

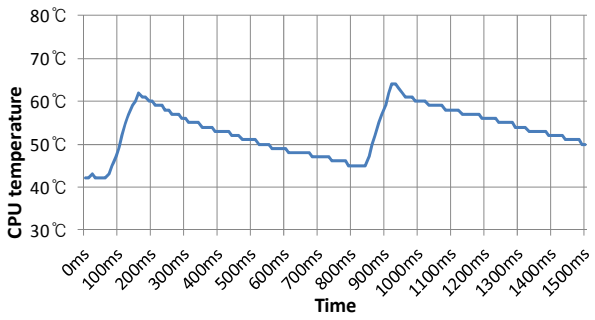

Figure 4 CPU Temperature change with maximum load

measured at four different temperatures t1, t2, t3, and t4; the lowest frequency F_{min} with low temperature t1, the highest frequency F_{max} with low temperature t2, F_{min} with high temperature t3, and F_{max} with high temperature t4, respectively. Then, the coefficients $Coef_H$ and $Coef_L$ are calculated and stored in the non-volatile memory of the AP.

$$Coef_H = \frac{VH1-VH2}{t4-t2} \ldots \textit{for high frequency}$$

$$Coef_L = \frac{VL1-VL2}{t3-t1} \ldots \textit{for low frequency}$$

$Coef_H$ and $Coef_L$ are essential parameters for calculating the common coefficient, $Coef_X$, for any frequency.

3.2 Temperature Checker

The runtime temperature checker monitors the changes of on-chip temperature in real-time and notifies the DVS controller to avoid unstable situations due to rapid temperature changes. Abrupt changes in on-chip temperature occur frequently in high performance APs, and the time required for these change is shorter than the default DVFS interval in the normal Android system. Therefore, independent temperature checking is safe and preferred over synchronization with the DVFS interval. Figure 4 shows the CPU temperature change for the Zip-Unzip operation, which causes the maximum load. The temperature increases by almost 20°C in 100 ms. Based on the test results, we decided that a 10 ms interval for temperature check is sufficient to cover the fastest temperature change. The maximum temperature change is 3°C in 10 ms, and the voltage change during this time is insignificant. The runtime temperature checker invokes the DVS controller if the current temperature is different from the previous temperature.

3.3 DVS Controller

The DVS controller is separated from the frequency scalier of conventional DVFS. The DVS controller is invoked when a new frequency is configured by the frequency scalier or a temperature change is observed by the temperature checker. The DVS controller estimates the appropriate operating voltage using the given temperature, frequency, and calculated coefficient. The voltage estimation procedure is shown in Figure 5. $Coef_H$ is a positive value while $Coef_L$ is negative. $Coef_X$ for a frequency F_X is between $Coef_H$ and $Coef_L$, and it is calculated as follows using linear interpolation.

$$Coef_X = \frac{Coef_L \times F_{max} - Coef_H \times F_{min} - F_x \times (Coef_H - Coef_L)}{F_{max} - F_{min}}$$

With temperature change, the voltage corresponding to the frequency F_X altered based on $Coef_X$. The voltage offset ΔV is then defined as follows.

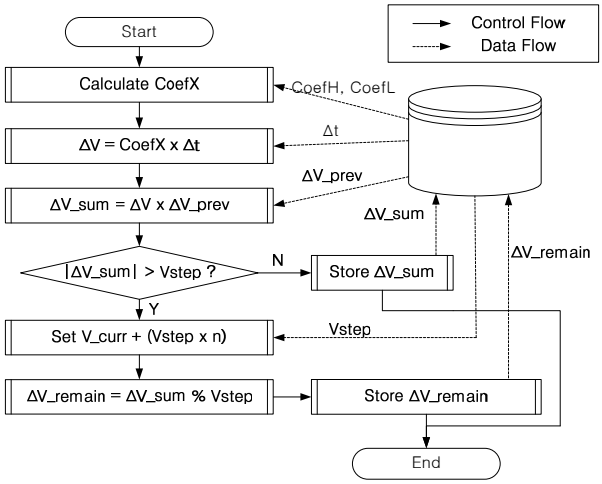

Figure 5 T-DVS voltage estimation procedure

$$Coef_X = \frac{\Delta V}{\Delta t} \rightarrow \Delta V = Coef_X \bullet \Delta t$$

The calculated ΔV is added to the previous remainder of ΔV. In the next step, the accumulated ΔV is compared with the minimum voltage control step of the Power Management IC (PMIC) to decide whether it is adequate to change the PMIC voltage level. If the accumulated ΔV is greater than the PMIC control step, The DVS controller changes the PMIC voltage, and stores the remainder. Else, the accumulated ΔV is stored until the next calculation, and there is no change in the PMIC voltage. If $Coef_X$ is positive, the operating voltage is increased at high temperatures. If $Coef_X$ is negative, the operating voltage is increased at low temperatures.

4. EVALUATION

We validated the effectiveness of T-DVS using commercial APs.

4.1 Experimental Setup

The experimental setup for evaluation was similar to that of the temperature inversion test. We used the AP tester equipment and the AP evaluation board shown in Figure 6. The AP evaluation board is designed to evaluate AP functionalities and provides multiple power measurement points. The voltage changes and power consumption are readily obtained using the measurement points on the board. The on-chip temperature is also collected by reading the Android sysfs which provides the measured temperature from the embedded thermal sensors in AP. The AP tester enables multi-chip tests with programmable test scripts that control chip handling, power supply, and ambient temperature. Fine-grained temperature control should be available for T-DVS

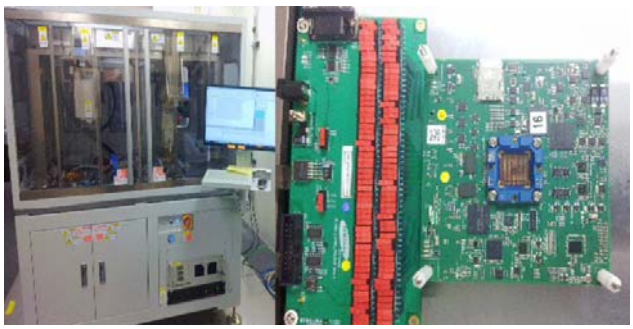

Figure 6 AP tester (left) and AP evaluation board (right) for Exynos 5433

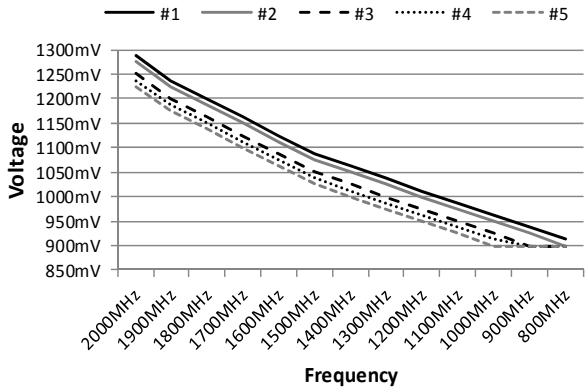

Figure 7 Default DVFS operating voltage of five Exynos 5433 chips

evaluation. For the evaluation, we used the Exynos 5433 AP which have four Cortex-A57 big cores and four Cortex-A53 little cores. The highest and lowest operating frequencies of the big cores are 2.0 GHz and 800 MHz respectively. We tested five AP chips with different speeds and leakage characteristics to verify whether T-DVS works well with different chips. Figure 7 shows that each chip has different operating voltages depending on its characteristics. We intended to check whether T- DVS can obtain appropriate voltage gains for different chips under the same test condition.

We conducted tests at diverse ambient temperatures that reflect real environments where mobile devices are operated. The temperature was set to cold (0°C), room (25°C), and hot (60°C). We disabled DTM to avoid the thermal throttling being activated in high temperatures. The test repeated a scenario that runs three Android applications: web browsing with scrolling, scrolling on the menu screen, and the Angry Bird game. Data were collected to analyze the behavior of T-DVS in real-time. These included on-chip CPU temperatures, CPU operating frequency, and power consumption. The test was performed without T-DVS as well for comparing the results with those of the proposed scheme.

4.2 Accuracy of Estimated Voltage

The accuracy of $Coef_X$ is crucial for voltage gain estimation. In the event of a wrong decision, the system malfunctions immediately. We conducted an experiment to validate the accuracy of V_{min} estimation. Figure 8 (a) shows the comparison of the V_{min} estimated using $Coef_X$ and the V_{min} measured on the Exynos 5433 evaluation board. The dotted lines in the graph denote the estimated V_{min} s for specific frequencies and temperatures. The solid lines denote the compensated voltages considering the controllable voltage step of the PMIC. The compensated voltage is set in the PMIC. The red dots in the graph denote the real V_{min}s measured repeatedly with the same sample. The maximum values of the red dots do not exceed the solid lines for all temperatures and frequencies. This means that the error in V_{min} estimation is less than the control step of the PMIC, and the estimated voltage is valid for application with the rounded up compensation.

We performed voltage tracing on the Exynos 5433 evaluation board with Android. The test involved touch and scroll actions in the Android menu for one minute. The trace result is shown in Figure 8 (b). We assumed that the maximum values of the

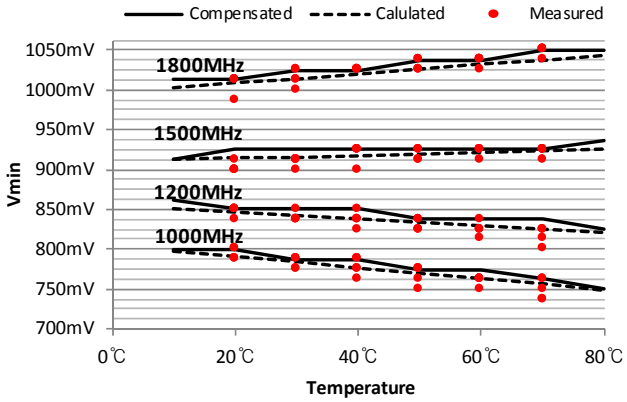

(a) Calculated voltage by $Coef_X$ and measured voltage.

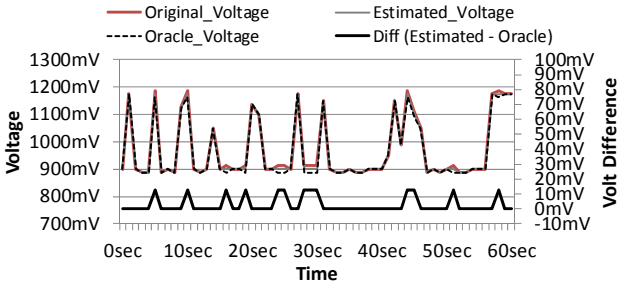

(b) Comparison of estimated voltage and oracle voltage (max. of measured V_{min}) during simple scenario run.

Figure 8 Comparison of estimated V_{min} and measured V_{min}

measured V_{min} s were the oracle voltages represented by the dotted line in the graph. Because the oracle voltage is the baseline for safe operation, voltages lower than the oracle voltage are not allowed. The red line in the graph indicates the original voltage defined in the default voltage table. We found that all estimated voltages were between the original voltages and the oracle voltage, and no estimated voltage is less than the oracle voltage. This experimental result shows the validity of the proposed voltage estimation procedure.

4.3 Voltage Gain

Figure 9 shows the utilization of CPU frequencies and the average voltage during repeated execution cycles of three CPU-centric applications. The frequency use ratio is shown in Figure 9 (a). The lowest frequency of the big CPU, 800 MHz, was used for approximately 48% of the time, and the high frequencies between 2.0 GHz and 1.6 GHz were used approximately 46% of the time. The use ratios represent general operation with On-demand governor and touch-booster, which use the lowest frequency most of the time and use high frequencies for specific events such as a touch action. As observed in the case of temperature inversion, the low and the high frequencies provide higher voltage gains than the mid-range frequencies. Thus effective power saving is available with the frequency ratio shown in Figure 9 (a). Figures 9 (b), (c), and (d) show the normalized voltages for all frequencies. The voltages for each chip are the averages during the test, and they include the down-controlled voltage by T-DVS. In other words, these voltages indicate the effectiveness of voltage gains in the scenario.

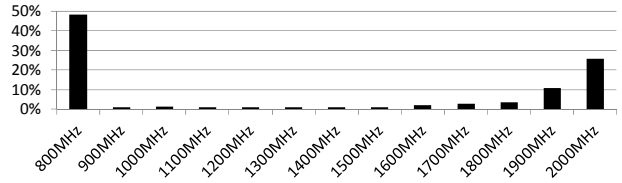

(a) Average frequency use ratio during the test

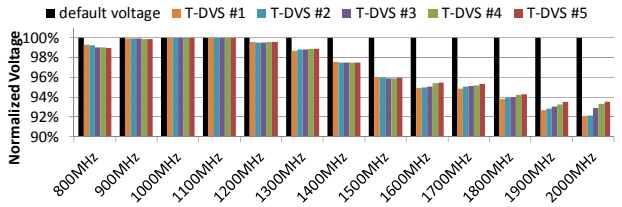

(b) Normalized average CPU voltage @ COLD temp. (0°C)

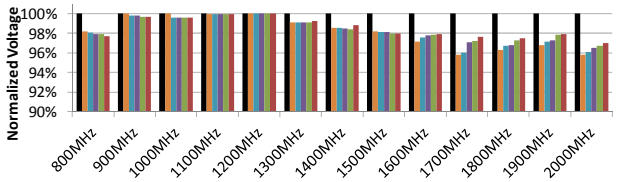

(c) Normalized average CPU voltage @ ROOM temp. (25°C)

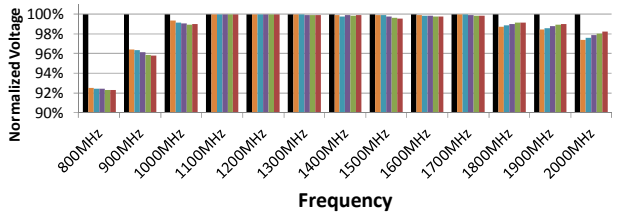

(d) Normalized average CPU voltage @ HOT temp. (60°C)

Figure 9 Big CPU voltage gains for each frequency

We found a difference between the voltage gains in terms of ambient temperature. At the high temperature, the voltage gain at low frequency was higher than at high frequency. This situation was reversed at low temperatures. The average on-chip temperatures under the cold, room, and hot conditions were 12, 41, and 78°C, respectively. The different on-chip temperatures led to different voltage gains. We also noticed difference in the graphs that seems to have resulted from chip variation. Chip #1 exhibits higher voltage gains than the others at high frequency, whereas Chip #5 exhibits higher voltage gains than the others at low frequencies. This phenomenon could be explained using the leakage current characteristics of the chips. In descending order of leakage current, Chip #5 is the first, Chip #4 is the second, and so on. If a chip has a large leakage current, it is more sensitive to temperature changes. Accordingly, the temperature of Chip #5 increased quicker than those of the other chips, and reverse is true for Chip #1. The voltage gain of each chip, therefore, varies as shown in Figure 9.

4.4 Power Consumption

The voltage gain achieved using T-DVS leads to direct power saving. Figure 10 shows the measured power at three temperatures With T-DVS, the power consumptions were lower than those without T-DVS by 6.6%, 10.9%, and 9.6% at the cold, room, and hot temperatures, respectively. We found that the effectiveness of

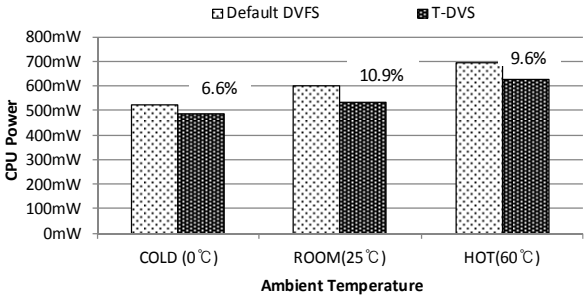

Figure 10 CPU power consumption gain with T-DVS

power reduction at room temperature is better than that at the other temperatures. At the cold temperature, the total power reduction is lower than that at other temperature owing to the low power saving at low frequencies. At the hot temperature, owing to the small voltage gain at high frequency which increases power consumption, the total power saving is lower than that at room temperature. This experimental result indicates that the effect of T-DVS is maximized when the temperature is maintained in a specific range, and the best output is achieved at mid-range temperature. Since mobile devices are mostly used at mid-range temperature, we expect approximately 10% of CPU power saving with T-DVS.

We calculated the runtime overhead with T-DVS. The runtime overhead is caused by the temperature checker that runs every 10 ms, and the DVS controller that is invoked if a temperature or frequency change is detected. The overall execution time of the temperature checker and DVS controller is 45 µs with the lowest CPU clock. This is 0.45% of the CPU running time and is indeed a trivial overhead.

4.5 Battery Life-time Improvement

In the real world, battery lifetime increment is an indicator for the effectiveness of a power saving mechanism. We conducted additional experiments to understand how much improvement in battery lifetime is achieved by T-DVS. We used the Exynos 5422 AP-based Galaxy S5 and wrote a Monkey script [3] consisting of nine Android applications: Web Browser, Camera, Gallery, Google Play Store, Settings, SMS, YouTube, Candy Crush, and

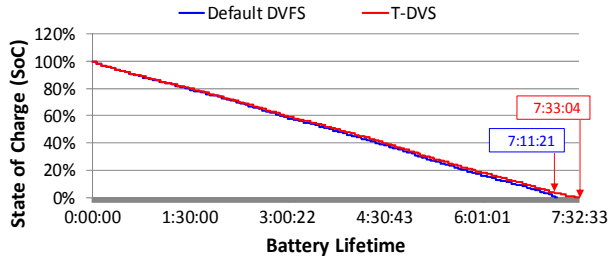

(a) Battery SoC changes of with and without T-DVS

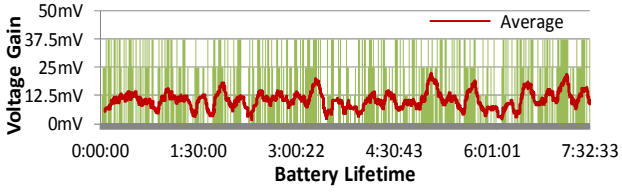

(b) Voltage gains by T-DVS during battery lifetime test

Figure 11 Battery lifetime improvement with T-DVS

Angry Birds. We collected voltage gain and State of Charge (SoC) information for the experiment. Battery lifetime improvement was then calculated by comparing the cases with and without T-DVS. The test was started with the battery in the fully charged and run until complete battery discharge. Figure 11 (a) shows that the smartphone was running for 7 hours 33 minutes with T-DVS, and 7 hours and 11 minutes without T-DVS. Overall, a 4.9% or 22 minutes increase in battery lifetime was achieved with T-DVS. Figure 11 (b) shows that the sustained voltage gains by T-DVS were acquired during the lifetime.

Furthermore, we conducted an aging test to observe whether T-DVS runs stably on a smartphone for an extended period. We used Galaxy S5 with a DC power supply as a substitute for the battery and ran the same aforementioned Monkey script. The test lasted for two weeks without any functional problem.

5. RELATED WORK

DVFS is an effective scheme for power management of mobile devices [10], [11]. Many studies have attempted to optimize frequency to reduce power consumption while preserving performance. A few work even studied voltage optimization. Bacha et al. [1], [2] achieved dynamic voltage reduction using on-chip Error Correction Code (ECC)-based voltage speculation. A benchmark test with particular firmware demonstrated that approximately 20% power saving is possible with aggressive voltage reduction on Intel's Itanium processor. Leng et al. [9] tried to reduce the voltage guard-band that exists in conventional GPU operation. The work pointed that the guard-band for voltage noises comes from various behaviors of programs and optimized the guard-band with off-the-shelf GPUs used in desktops. The study, however, underrated the influence of temperature on the guard-band because a desktop has strong means of system cooling. Mobile devices have no explicit cooling control and temperature typically affects device power consumption significantly. Regarding the thermal management of mobile devices, Kim et al. [6] and Lee et al. [7] studied temperature-aware DVFS aimed to enhance DTM by managing the temperature effectively with frequency control. In general, frequency optimization is accompanied by performance degradation, an undesired outcome. On the contrary, voltage optimization itself does not affect system performance. This is, in fact, a major merit of T-DVS.

Meanwhile, Near-Threshold-Voltage (NTV) technology emphasizes the importance of reduced operating voltage. Kaul et al. [5] introduced the concept of NTV, and demonstrated that the maximum energy efficiency can be achieved using NTV. In fact, Intel demonstrated an NTV Pentium processor [4] that runs on very low power. Voltage down by software techniques is helpful for the NTV scheme. We expect that T-DVS would lead to the development of techniques focusing on lower operating voltages and has synergistic effect on technologies such as NTV. In the case of combined use of T-DVS and other methods, its benefits will be retained without any loss.

6. CONCLUSION

We demonstrated that T-DVS can achieve power saving with aggressive voltage control by using the relationship between operating voltage and temperature. We believe that T-DVS would set a milestone in active voltage control for mobile device power management. The effectiveness of T-DVS was achieved solely by using software approaches and employing existing hardware features. If AP manufacturers develop additional hardware features to assist T-DVS, the T-DVS scheme would become more powerful and convenient. In the present work, we validated the scheme with only big CPU cores, but we plan to extend the scheme to other hardware components of APs.

7. ACKNOWLEDGMENTS

This work was supported by Samsung Research Funding Center of Samsung Electronics under project number SRFC-TB1503-02.

8. REFERENCES

[1] Bacha, A. and Teodorescu, R. 2013. Dynamic reduction of voltage margin by leveraging on-chip ECC in Itanium II processor. In *Proceedings of the 47th Annual International Symposium on Computer Architecture – ISCA `13, Association for Computing Machinery (ACM).*

[2] Bacha, A. and Teodorescu, R. 2014. Using ECC feedback to guide voltage speculation in low-voltage processors. In *Proceedings of the 40th Annual International Symposium on Microarchitecture—MICRO `14, Association for Computing Machinery (ACM).*Google, Android, "UI/Application Exerciser Monkey" http://developer.android.com/intl/ko/tools/help/monkey.html

[3] Google, Android. UI/Application Exerciser Monkey. http://developer.android.com/intl/ko/tools/help/monkey.html

[4] Intel. A Solar-Powered NTV Processor. http://blogs.intel.com/intellabs/2011/09/15/ntvp/

[5] Kaul, H., Anders, M., Hsu, S., Agarwal, A., Krishnamurthy, R., and Borkar, S. 2010. Near-Threshold Voltage (NTV) Design. In *Proceedings of the 49th Design Automation Conference – DAC-49, Association for Computing Machinery (ACM).*

[6] Kim, J. M., Kim, Y. G., and Chung, S. W. 2015. Stabilizing CPU frequency and voltage for temperature-aware DVFS in mobile devices, *IEEE Trans. Comput., Institute of Electrical & Electronics Engineers (IEEE)*

[7] Lee, J. S., Skadron, K., and Chung, S.W. 2010. Predictive temperature-aware DVFS. *IEEE Trans. Comput., Institute of Electrical & Electronics Engineers (IEEE).*

[8] Lee, W., Wang, Y. Cui, T., Nazarian, S., and Pedram M. 2014. Dynamic thermal management for FinFET-based circuits exploiting the temperature effect inversion phenomenon. In *Proceedings of the 2014 International Symposium on Low Power Electronics and Design - ISLPED `14, Association for Computing Machinery (ACM).*

[9] Leng, J., Buyuktosunoglu, A., Bertran, R., Bose, P., and Reddi, V. J. 2015. Safe limits on voltage reduction efficiency in GPUs. In *Proceedings of the 48th International Symposium on Microarchitecture—MICRO-48, Association for Computing Machinery (ACM).*

[10] Pillai, P. and Shin, K. G. 2001. Real-time dynamic voltage scaling for low-power embedded operating systems. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles—SOSP `01, Association for Computing Machinery (ACM).*

[11] Pouwelse, J., Langendoen, K., and Sips, H. 2001. Dynamic voltage scaling on a low-power microprocessor. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking—MobiCom`01, Association for Computing Machinery (ACM).*

[12] Samsung Electronics Co. Ltd. 2014. Exynos 5433 Application Process User's Manual