



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

User-specific Skin Temperature-aware DVFS for Smartphones

Citation for published version:

Egilmez, B, Memik, G, Ogrenci-Memik, S & Ergin, O 2015, User-specific Skin Temperature-aware DVFS for Smartphones. in *Proceedings of the 2015 Design, Automation Test in Europe Conference Exhibition*. EDA Consortium, San Jose, CA, USA, pp. 1217-1220. <<http://dl.acm.org/citation.cfm?id=2757094>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of the 2015 Design, Automation Test in Europe Conference Exhibition

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



User-Specific Skin Temperature-Aware DVFS for Smartphones

Begum Egilmez*, Gokhan Memik*, Seda Ogren-ci-Memik*, Oguz Ergin⁺

*Dept. of Electrical Engineering and Computer Science
Northwestern University
Evanston, IL, USA

⁺Department of Computer Engineering
TOBB University of Economics and Technology
Ankara, TURKEY

Abstract— Skin temperature of mobile devices intimately affects the user experience. Power management schemes built into smartphones can lead to quickly crossing a user's threshold of tolerable skin temperature. Furthermore, there is a significant variation among users in terms of their sensitivity. Hence, controlling the skin temperature as part of the device's power management scheme is paramount. To achieve this, we first present a method for estimating skin and screen temperature at run-time using a combination of available on-device thermal sensors and performance indicators. In an Android-based smartphone, we achieve 99.05% and 99.14% accuracy in estimations of back cover and screen temperatures, respectively. Leveraging this run-time predictor, we develop User-specific Skin Temperature-Aware (USTA) DVFS mechanism to control the skin temperature. Performance of USTA is tested both with benchmarks and user tests comparing USTA to the standard Android governor. The results show that more users prefer to use USTA as opposed to the default DVFS mechanism.

Keywords— Thermal modeling, mobile device, skin temperature

I. INTRODUCTION

Smartphones and tablets have become a vital part of daily life. Users are often in touch with the back cover and screen for prolonged durations. Thus, heat dissipation intimately affects user experience. Heat pain is experienced by most humans when they touch an object hotter than 45°C [1, 2], which is not uncommon in mobile devices. There have been complaints regarding the outer temperature of iPad 3 while playing graphics intensive games and it is reported that the hottest spot reaches 47°C [3]. In fact, complaints of excessive temperatures are common to many smartphones [4, 5].

Smartphones are equipped with a number of temperature sensors monitoring the CPU and the battery, but lack any tracking of the *skin temperature*¹. In this paper, we are presenting a new system for optimizing user experience by minimizing discomfort due to extreme skin temperatures. Our user studies presented in Section 3 confirm the large variation in discomfort and perception of heat across users. Motivated by these results, we develop User-specific Skin Temperature-Aware (USTA) Dynamic Voltage and Frequency Scaling (DVFS). We developed a light-weight and accurate run-time predictor for skin temperature. We combine system level events and on-device physical thermal measurements in a

systematic framework to develop highly accurate prediction models. The average prediction error for the skin and screen are 0.95% and 0.86% respectively. USTA uses this predictor to control the skin temperature of the device in a user-aware manner. The effectiveness of USTA is tested both with benchmarks and real users. During a heavy workload, such as a video call, the skin temperature of the smartphone is reduced by up to 4.1°C, while the average CPU frequency is reduced by 34% compared to the default power management system. Our results show that the users are not affected by this performance change and have rated USTA higher than the default DVFS in overall experience. Hence, USTA is capable of maintaining lower skin temperature using proactive DVFS while incurring minimal performance degradation.

The rest of the paper is organized as follows. In the next section, related works are summarized and compared to our approach. In Section 3, we present our prediction model and DVFS mechanism. Section 4 delivers an evaluation of the accuracy of our prediction model along with final user study and benchmark tests to highlight the effectiveness of USTA. We conclude with a summary in Section 5.

II. RELATED WORK

The need for more sophisticated thermal management for smartphones, including special consideration to the user experience, has been recognized in industry [6]. Lee et al. [7] studied the effects of different materials used for the casing of the phone. Xie et al. [8] propose a thermal analyzer, which generates steady-state temperature maps of entire smartphone including the cover. Another study demonstrates the thermal coupling between the application processor and the battery [9]. Our work differs from these works since it has a system-level approach and provides a solution that does not require any hardware design change. Besides, to the best of our knowledge we present the first publicly available results of user studies that demonstrate the variability among users by means of temperature perception.

III. SKIN TEMPERATURE-AWARE THERMAL MANAGEMENT

High skin temperatures can cause dissatisfaction in users. First, we aim to determine the variance in the skin temperature tolerance among users. Then, we will utilize this information to control skin temperature. To this end, we developed a prediction model to monitor skin temperature during runtime. Then, we built our DVFS scheme utilizing the user study data gathered along with the prediction model.

Our experiments were conducted with a Google Nexus 4

¹ In this paper, we use the term “skin temperature” to indicate the temperature at the middle of the back cover (a location that is commonly touched by the user). We use the term “screen temperature” to indicate the operating temperature in the middle of the screen.




Figure 1. Our preliminary user study indicating individual levels of comfort with respect to skin and screen temperatures.

device running Android operating system v.4.3, Jelly Bean. Nevertheless, our proposed skin temperature prediction and skin temperature-aware DVFS can be applied to all current and emerging smartphone platforms. The smartphone has two built-in thermal sensors to report the CPU and battery temperatures. We have extended this set with two sensors on back cover (upper section and midsection) and one on screen. We have used thermistors for temperature measurements. We have utilized an application to periodically log system level information, such as CPU temperature, battery temperature, CPU utilization, and CPU frequency.

We performed a user study with 10 participants (5 males, 5 females), who are graduate or undergraduate students. Subjects were asked to hold the phone in their hands and to report the instant when the discomfort due to skin temperature was unacceptable to them. The most tolerant subject ended test in seven minutes. This is consistent with a realistic environment, where an average user could easily be expected to touch the smartphone for several minutes.

A computationally heavy application Antutu Tester is used during this user study. Even though the CPU temperature does not exceed the threshold for the built-in power management system to intervene, we observed that this very same thermal state causes the skin temperature to exceed comfortable limits for all the participants. Another important result of our study is the sensitivity variation across users. Figure 1 shows that the minimum skin temperature exceeding the comfort limit for a user is 34.0°C, while the maximum is 42.8°C.

A. Framework to Derive a Prediction Model

The prediction system is based on empirical data gathered from thermal measurements as well as the Android OS. The measurements collected from attached external thermistors were used as the expected value for the skin temperature and screen temperature. Specifically, we have used the CPU temperature, utilization, and frequency, as well as battery temperature collected in our log file. We have used a total of thirteen benchmarks to collect a large set of data. Eleven of these benchmarks can be downloaded from the PlayStore [10]. These are AnTuTu Benchmark Set, AnTuTu Tester, GFXBench 3.0 3D, Vellamo, Skype, Youtube, and The Legend of Holy Archer. Since AnTuTu Benchmark Set is customizable, we derived different benchmarks from it. The other two benchmarks are built-in functionalities.

We have leveraged machine learning schemes of WEKA [11] to build the prediction model. We identified four algorithms that match our data type. These algorithms are linear regression, multilayer perceptron, MSP, and REPTree. Each algorithm is tested with 10-fold cross-validation. The performance metric to find the best performing algorithm has been defined as the average error rate of all predictions in that benchmark. The error rate is calculated with the following equation:

$$Error\ rate = \frac{|expected\ value - predicted\ value|}{expected\ value} \times 100 \quad (1)$$

Experiments presented in the next section reveals that our prediction mechanism is highly accurate. In order to demonstrate the generality of our model, we also analyzed the impact of human touch on the prediction accuracy. We first measured skin temperature when the device is turned-off and not touched. Next, the user holds the device in her/his palm while it was still turned-off. Third, we measured the skin temperature when the Antutu tester application was running and the phone was not touched. Finally, users held the phone while the application is running. When holding the phone, we made sure that the palm of the hand always touched the back of the phone. It is observed that human touch does not alter exterior temperature values of the device significantly especially when the phone is actively used.

B. User Specific Skin Temperature-Aware DVFS

The aforementioned prediction model is used by our DVFS system to keep skin temperature under user's desired threshold, thereby minimize discomfort. We have experimented with both using a default threshold representative of an average user as well as configuring it for each individual's threshold.

In our configuration, USTA performs skin temperature prediction every 3 seconds and intervenes to enforce a DVFS decision on the system only if skin temperature needs to be controlled. Otherwise, the baseline DVFS performs its function for power optimization only. The baseline DVFS is the default Android on-demand governor and it scales the frequency of the processor according to CPU utilization. When utilization is at the maximum, the frequency is also set at the maximum level. The reduction in frequency can be steep if the utilization is very low or it could be in steps if the utilization is below a threshold (around 80%), but above a minimum (around 20%). For Nexus 4, there are twelve frequency levels between 384MHz and 1.512GHz.

USTA has a threshold for activation which is set to 2°C below the skin temperature limit of the user. If the difference between the predicted skin temperature and the temperature limit is between 1°C and 2°C, the maximum allowed CPU frequency is decreased by one level (i.e., from the highest frequency to the one below). If the difference between the prediction and the temperature limit is between 0.5°C and 1°C, then, the maximum allowed CPU frequency is decreased by two levels. Finally, if the prediction is closer than 0.5°C to the limit or it is exceeding the limit, then, the maximum CPU frequency is set to the minimum frequency level.




Figure 2. Percentage of time that will be spent over the temperature threshold for each user in the experiments described in 3.1.

Figure 2 presents the percentage of time where the user’s comfort threshold has been exceeded during half hour video call in Skype. We present data for eleven different settings for the skin temperature limit. Ten of those cases are configured according to a specific individual (as derived from our user studies) and the eleventh case represents a “default” user. The limit of the default user has been set as the average of the ten users’ reported discomfort limits. For the default user, 15.6% of the total execution time is spent outside the comfort limit.

IV. EXPERIMENTAL RESULTS

In this section, we present our results to evaluate the effectiveness of our skin temperature prediction model as well as USTA, which is guided by this model.

A. Evaluation of the Prediction Model

Figure 3 presents the average error rate for the four prediction models we have experimented with. We must highlight that for all the target applications, we have developed a single global model. In other words, all the data from all our target applications are gathered and then we generated a single model for this global set (WEKA performs the 10-fold cross validation and then lists the expected values and predicted values from which we calculate average error rates for the cross validation). Linear regression and multilayer perceptron are relatively poor in accuracy when compared to other methods. M5P and REPTree make nearly 100% accurate estimations. REPTree gives the best results with average of 0.95% error rate for skin temperature estimations and 0.86% error rate at screen estimations. M5P algorithm has 0.96% error rate for skin estimations and 0.89% error rate for screen estimations. However, when we ignore temperature differences less than 1°C (as humans are less sensitive in that range), the M5P algorithm gives better results. Average skin




Figure 4. Recorded temperature values during 0.5 hour Skype video call.

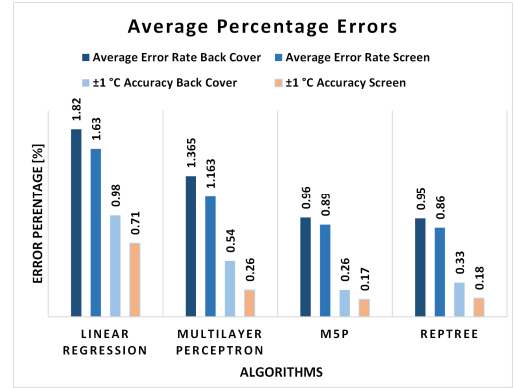


Figure 3. Average error rates for the four prediction models used.

temperature prediction error rate decreases to 0.26% and error rate of screen temperature prediction decreases to 0.17%.

From computation cost perspective, REPTree builds faster than M5P and does not cause halting. Thus, we have chosen REPTree to implement and utilized WEKA libraries for that. The run-time overhead of predicting the skin temperature and screen temperature are 5.603ms and 6.708ms, respectively. This incurs a total cost of 12.383ms in a window 3 seconds. Although this cost is already very small (~0.4%), it could further be reduced by; a) increasing the duration between predictions and b) selectively predicting the screen and/or the back cover temperature (for example, screen temperature prediction can be used only during a phone call).

B. Evaluation of USTA

The performance of USTA is tested both with benchmarks and by real users. The temperature limit for USTA was set to 37°C, which is calculated by finding the average discomfort limit reported by the users in our experiments. We observe that despite the equal initial temperatures, baseline DVFS leads to higher peak skin temperatures for all benchmarks.

The results for the Skype video call benchmark are presented in Figure 4. Even though on occasion USTA cannot remain below the comfort limit, USTA succeeds in maintaining a more steady temperature, near that limit. The peak skin temperature of the baseline DVFS is 4.1°C higher than the skin temperature experienced with USTA. Skin and screen temperatures, as well as average processor frequencies for all benchmarks using USTA versus baseline DVFS are listed in Table 1. USTA reduces peak temperatures without incurring significant performance penalties.

Finally, we have conducted one more user study. We

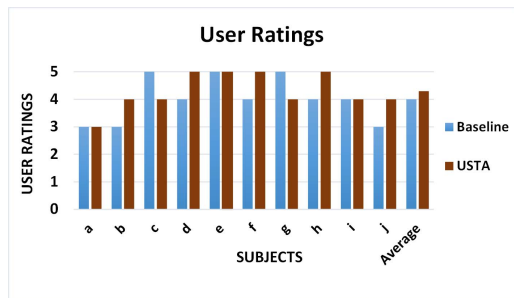


Figure 5. User ratings for baseline DVFS and USTA on a scale of 1 to 5.

		CPU	CPU-GPU-RAM	User Exp.	AnTuTu	CPU (1.5 hours)	AnTuTu Tester	GFX Bench	Vellamo	Skype	Youtube	Record	Charging	Game
Baseline DVFS	Max Screen Temp [°C]	33.4	32.5	28.5	30.5	35.1	34.3	26.3	28.6	40.5	28.0	32.8	29.0	33.3
	Max Skin Temp [°C]	37.9	36.3	31.9	34.0	39.3	42.8	29.3	31.0	42.8	30.4	37.1	31.7	36.6
	Average Freq. [GHz]	1.04	1.01	1.22	1.11	1.09	1.16	0.85	0.97	1.09	0.80	0.86	0.45	1.14
USTA	Max Screen Temp [°C]	31.7	31.4	29.2	31.5	34.9	34.9	28.5	29.7	35.4	30.0	32.5	29.9	31.7
	Max Skin Temp [°C]	35.1	35.1	32.7	34.0	38.8	41.1	34.8	32.1	38.7	32.9	36.6	32.3	35.1
	Average Freq. [GHz]	1.22	0.91	1.05	0.99	0.69	0.89	1.16	0.96	0.72	0.64	0.81	0.39	0.63

Table 1. Maximum screen and skin temperatures, and average frequency with all benchmarks for the baseline DVFS and USTA. USTA thermal limit is set according to the “default” user, i.e., 37°C. In all applications where the temperature is within 2°C or exceeds this threshold for the default DVFS, USTA is able to reduce the peak temperature.

invited the same group of users who participated in our initial study and asked them to hold the device during a Skype video call for an hour: 30 minutes was controlled with the baseline DVFS while the remaining 30 minutes was using USTA. Participants were unaware of the scheme controlling the frequency. For each user, we have configured USTA to the comfort limit reported by that same user in the previous study. Then, we asked the users to grade their satisfaction using this smartphone on a scale of 1 to 5. The results are shown in Figure 5. Most of the users either favor USTA or gave the same rating to both systems. *On average, the satisfaction with the default DVFS is 4, while it is 4.3 for USTA, showing that USTA indeed improves the overall user experience.*

We also asked the participants which one of these systems they would prefer taking also the device’s performance into consideration, because USTA is changing the frequency, and may cause degradation in performance. However, none of the users reported dissatisfaction with USTA’s performance. 4 out of 10 subjects (users a, d, e, and i) reported that they did not observe any noticeable change in two systems. For these users, we observe that a good fraction of them indicated a high temperature threshold initially. Therefore, it is natural that these users did not see a difference between these two systems, since USTA did not take any action for them. Out of the 6 remaining users, 2 preferred the baseline (users c and g) while 4 preferred USTA (users b, f, h, and j). The users c and g have not indicated their reasons for selecting the baseline. In fact, the user g had a very high threshold, and did not require USTA to take action at all, yet the user preferred the default. We must note that at the time of the experiment the users did not know which one of the systems they are using, so the fact that a larger fraction of the users have chosen USTA over the baseline scheme shows that USTA can reduce the temperature without a significant impact on the performance and hence improves the overall user experience.

V. CONCLUSION

In this paper, we have developed a method to control the skin temperature in smartphones. USTA is capable of estimating

skin temperature with 99.05% accuracy. Using this prediction USTA scales frequency of application processor by taking user comfort limit into consideration. Our final user survey shows that a larger fraction of our participants prefer using USTA instead of the baseline DVFS: on average, USTA has a rating of 4.3 and the baseline DVFS has a rating of 4.

VI. ACKNOWLEDGEMENTS

This work was partially supported by DOE (DE-SC0012531), NSF (CCF-1422489, CCF-0747201), and Intel URO Energy Smart Soc Rapid Prototyping ISRA. We also would like to thank the participants in our experiments.

REFERENCES

- [1] E. Arens, H. Zhang, “The Skin’s Role in Human Thermoregulation and Comfort”, *Indoor Environ. Qual.*, Oct. 2006.
- [2] G. L. Wasner and J. A. Brock, “Determinants of Thermal Pain Thresholds in Normal Subjects,” *Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol.*, vol. 119, no. 10, pp. 2389–2395, Oct. 2008.
- [3] J. A. Kaplan (2012) “New Apple iPad hits 116 degrees, Consumer Reports says”. <http://www.foxnews.com/tech/2012/03/20/ipads-not-overheating-apple-says/>.
- [4] (2014) “Sony Xperia Z2 Smartphone Gets Too ‘Hot’ to Use”, <http://www.maxconsole.com/news/ANDROID>
- [5] J. Herrman (2011) “Why Is My Phone So Hot?”, <http://www.popularmechanics.com/technology/how-to/tips/why-does-my-phone-get-so-hot>.
- [6] K. Sekar (2013) “Power and Thermal Challenges at MobilePhones” http://www.sigmobile.org/mobicom/2013/MobiCom2013_IndustrialTalks_KrishnaSekar.pdf
- [7] J. Lee, et al., “Parametric Thermal Modeling of Heat Transfer in Handheld Electronic Devices”, *ITHERM*, 2008.
- [8] Q. Xie, et al., “Therminator: A Thermal Simulator for Smartphones Producing Accurate Chip and Skin Temperature Maps”, *ISLPED*, 2014.
- [9] Q. Xie, et al., “Dynamic Thermal Management in Mobile Devices Considering the Thermal Coupling between Battery and Application Processor”, *International Conference on Computer-Aided Design*, 2013.
- [10] play.google.com/store/
- [11] M. Hall, et al., “The WEKA Data Mining Software: An Update”, *SIGKDD Explorations*, Vol. 11, Issue 1, 2009.