Fire in Your Hands: Understanding Thermal Behavior of Smartphones

Soowon Kang*, Hyeonwoo Choi†, Sooyoung Park*, Chunjong Park‡, Jemin Lee§, Uichin Lee*, Sung-Ju Lee*

*KAIST †Samsung Electronics ‡University of Washington §ETRI

ABSTRACT

Overheating smartphones could hamper user experiences. While there have been numerous reports on smartphone overheating, a systematic measurement and user experience study on the thermal aspect of smartphones is missing. Using thermal imaging cameras, we measure and analyze the temperatures of various smartphones running diverse application workloads such as voice calling, video recording, video chatting, and 3D online gaming. Our experiments show that running popular applications such as video chat, could raise the smartphone's surface temperature to over 50°C in only 10 minutes, which could easily cause thermal pain to users. Recent ubiquitous scenarios such as augmented reality and mobile deep learning also have considerable thermal issues. We then perform a user study to examine when the users perceive heat discomfort from the smartphones and how they react to overheating. Most of our user study participants reported considerable thermal discomfort while playing a mobile game, and that overheating disrupted interaction flows. With this in mind, we devise a smartphone surface temperature prediction model, by using only system statistics and internal sensor values. Our evaluation showed high prediction accuracy with root-mean-square errors of less than 2°C. We discuss several insights from our findings and recommendations for user experience, OS design, and developer support for better user-thermal interactions.

MobiCom '19, October 21–25, 2019, Los Cabos, Mexico

ACM ISBN 978-1-4503-6169-9/19/10…$15.00
https://doi.org/10.1145/3300061.3300128

CCS CONCEPTS

• **Hardware** → **Thermal issues**; Analysis and design of emerging devices and systems; • **Human-centered computing** → **Smartphones**; User studies; • **General and reference** → *Measurement*.

KEYWORDS

Smartphones; Thermal characteristics; Thermal imaging camera; Thermal modeling; User guidance

ACM Reference Format:
Soowon Kang ,Hyeonwoo Choi, Sooyoung Park, Chunjong Park, Jemin Lee, Uichin Lee, Sung-Ju Lee. 2019. Fire in Your Hands: Understanding Thermal Behavior of Smartphones. In *The 25th Annual International Conference on Mobile Computing and Networking (MobiCom '19), October 21–25, 2019, Los Cabos, Mexico.* ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3300061.3300128

1 INTRODUCTION

Electronic devices, including smartphones contain electric elements that generate heat. In a competitive smartphones market, the manufacturers constantly enhance the computing performance and diversify functionalities, which could result in generating a great amount of heat. Users have recognized the smartphone thermal problem, as many reviews and articles [4, 38] empirically highlight the smartphone overheating issues.

Heat, if not handled properly, can not only degrade the processors' performance and damage the battery, but also degrade user experiences [57], aggravate a thermal-regulatory disorder [65], and pose health threats (e.g., thermal pain [37], skin burns [33], skin aging [54]).

Overheating problem is well-known in the field of electronic devices and various cooling technologies have been proposed (e.g., using proper thermal conductivity materials, dynamic voltage and frequency scaling (DVFS), etc.). Despite such techniques, overheating in smartphones is challenging due to small form factors, high power density, and close physical contact with the human body [53], and thus needs careful investigation. Prior work highlighted the importance and challenges of thermal management in mobile and wearable devices [53]: thermal models of smartphones [27], thermal management schemes [24], and thermal characteristics of smart glasses [39]. However, what is missing is systematically exploring when and how much heat smartphones (and each of their components) generate in practical scenarios and what impact thermal issues have on user experiences. We aim to deepen our understanding of the thermal issues of smartphones with systematic thermal measurements under various usage scenarios.

Using thermographic cameras [3], we first investigate the thermal characteristics of recent smartphones. We measure the surface temperature of a wide range of smartphones such as Android reference phones (Nexus 5, Nexus 5X, Nexus 6, Pixel, and Pixel 2), iOS phones (iPhone 7, iPhone 7+, and iPhone 8), and Android phones (Galaxy S7 and Huawei P20) by considering representative workloads such as video chatting, video recording, voice calling, and gaming. We believe our extensive measurement is the first of its kind on smartphones. We discover that our test smartphones easily reach over 45°C, which can cause thermal pain. The temperature rises in short duration, especially for the camera usage scenarios; e.g., only 6 minutes for video chat. Applications that require video encoding especially generate excessive heat quickly.

Our component-specific analyses reveal that intensive AP (Application Processor), power management circuit, and camera usage are the major heat sources. An AP generates the most heat as it houses multiple sub-components such as CPU, GPU, and multimedia codec, whereas GPS does not generate much heat despite its relatively high power consumption. Furthermore, we report that recent ubiquitous scenarios such as augmented reality, virtual reality and mobile deep learning have considerable thermal issues.

In addition to systematic measurements, we analyze how smartphone overheating affects user experiences through a user study. A majority of our participants (75%) reported considerable thermal discomfort while playing a mobile game and complained that overheating disrupted interaction flows and caused interaction errors. We find the major reasons for thermal discomfort are concerns on health risks and sweating related issues, including sensor malfunction and user input errors due to slip caused by sweat.

With the understanding of thermal issues in user experiences, we propose a model that estimates the surface temperature of a smartphone using only system statistics (e.g., CPU usage) and internal sensors (e.g., battery temperature), without kernel-level hardware usage information (e.g., current video recording resolution, clock frequency information). After selecting the dominant features, we train a time-series regression model each for Nexus 5X and Galaxy S7. For each smartphone, our model shows high accuracy (less than 1.2°C root-mean-square error (RMSE) for Nexus 5X and 0.523°C RMSE for Galaxy S7) in both real-time and one minute look-ahead predictions. Our model requires only one-time training, and thus is easily applicable to other smartphone models.

We stress that no prior studies systematically examine the overheating patterns under various workloads and their impact on user experiences. Our experiments show that various real-world workloads generate excessive heat, and our user study confirm that overheating degrades user experiences. Furthermore, emerging ubiquitous applications such as AR/VR and mobile deep learning also have significant thermal issues. Our results indicate that DVFS alone is not sufficient for managing smartphone surface temperature. There should be a holistic thermal management framework that simultaneously considers multiple components' thermal characteristics (e.g., PMIC, cameras, sensors, and wireless chipsets). We call for further studies on this important and difficult thermal problem.

2 BACKGROUND

2.1 Thermal Concerns on Mobile Devices

Thermal concerns are broadly classified as follows: performance degradation, negative user experiences, and health risks. System performance degrades as the temperature increases as mobile thermal management algorithms throttle operating CPU clock frequencies and switch off cores when devices overheat. In addition, overheating could hamper user experience and might even result in skin damage. Discomfort due to overheating is a concern that smartphone users often complain [4]. A recent user experience research [57] showed that user burdens, including physical discomfort and pain, negatively influence overall user experiences. Prior human perception studies showed that a person's thermal threshold for warm sensation is in the range of 33.0°C to 35.0°C [28]. At a temperature higher than this threshold, a user is likely to feel discomfort and starts to feel *pain* at around 42–45°C [37]. Beside such discomfort, there are several health concerns. Heat can accelerate premature skin aging, which gradually happens over time due to lack of awareness [54]. In addition, long-term exposure to such thermal condition of mild heating could lead to *erythema ab igne*, known as the toasted skin syndrome [33]. These temperature thresholds are also used in the European Standard (EN563) that provides the ergonomic temperature limits for surfaces [17].

2.2 Heat Transfer Basics

In electric devices, there are two major heat transfer routes: *conduction* and *convection* depending on whether heat transfer happens through a solid medium or a fluid (e.g., liquid/gas). According to thermodynamics principles [19], the heat transfer rates of conduction and convection ($\dot{q}$) can be modeled as: $\frac{kA}{d} \times \Delta T$, and $hA \times \Delta T$, respectively. Here, k is the medium's conductivity, and h is the convection heat transfer coefficient, which are dependent on material/fluid properties (e.g., copper's conductivity is 10,000 times larger than that of plastic). A is the area, ΔT is the temperature difference (e.g., between the phone body and its surroundings), and d is the thickness of the medium. According to these equations, the heat transfer rate between two points is proportional to the temperature difference, the area through which heat transfer occurs, and the heat transfer coefficient of the material/fluid.

Smartphones are comprised of multiple layers of electric components (integrated circuits and wires) with different thermal conductivities that generate heat due to power dissipation and Joule heating. Because of the layers' wide surface area and thinness, heat transfer mostly occurs in vertical directions across multiple layers, while relatively minor transfer occurs in horizontal directions. This means that more heat transfer occurs via conduction than convection, and the air gap between layers acts as an insulator. Overall, the heat generated from a number of heat sources is transferred through the adjacent materials and is finally released to the atmosphere. Typically, the rate of temperature change steadily decreases and reaches the steady-state (called the steady-state temperature).

2.3 Mobile Thermal Managements

Various cooling methods have been proposed that are generally categorized into passive and active cooling. Passive cooling [59] includes the techniques that use the heat pipe [26], heat sink [22], and thermal-interface materials. Active cooling [45] includes forced air and forced liquid cooling techniques. However, the small form factor and light weight requirements of smartphones prevent adopting active cooling and smartphones thus mostly rely on passive cooling.

Besides these cooling techniques, thermal management also happens at an electronic component level. Modern application processors (APs) support DVFS and CPU hot-plug to manage power consumption and protect devices from overheating [16, 58, 61] . A thermal protection mechanism suppresses temperature increase by reducing power consumption via lowering operating voltage/frequency of a CPU, or turning off CPU cores.

The thermal management in Android smartphones is carried out by the Linux kernel's *thermal engine*. This engine monitors CPU temperatures from on-die sensors and controls the temperature by mainly using two algorithms: dynamic control and threshold control. The dynamic control algorithm checks whether the core temperatures exceeds the threshold values and determines whether to throttle clock frequency or turn off cores. The threshold values are manufacturer-dependent and set by the threshold control algorithm. Android also implements an additional thermal protection mechanism by which a device is shut down when the battery temperature exceeds the threshold (e.g., 68°C). In practice, device manufactures may configure different thresholds owing to the heterogeneity of temperature sensors and device characteristics (e.g., Qualcomm DragonBoard 410c's threshold of 70°C for dynamic controlling [50]).

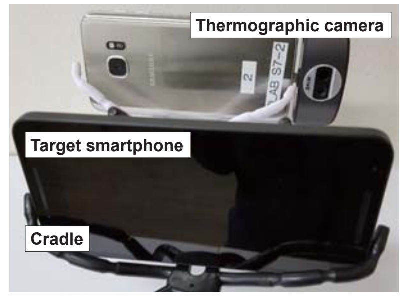

Figure 1: Measurement setup.

3 PRELIMINARY STUDY

We performed a preliminary measurement study to evaluate whether smartphones show thermal issues that could influence user experiences and even cause health problems.

3.1 Measurement Setup

We used FLIR ONE [3], a thermographic camera (also known as infrared (IR) camera or thermal imaging camera) as the temperature measurement device. An IR camera is popularly used in research communities [14, 36, 44, 52] as it shows comparable accuracy and precision to those of thermocouple, which is known to provide the most accurate temperature measurements. We used FLIR ONE's Android SDK to implement an app for automated data collection.

We conducted experiments to investigate the thermal characteristics of the smartphones, shown through the surface of both the front and back. We fastened a cradle on the target smartphone and the distance between the thermal camera and the target smartphone was fixed, as shown in Figure 1. The main goal of this preliminary study is to identify whether well-known use cases of smartphones could generate significant amount of heat, which could provide thermal discomfort, or even skin damage to the users; recall that users start to feel heat pain at 42–45°C [37].

To check thermal issues across different phones, we chose Android reference phones (Nexus 5, Nexus 5X, Nexus 6, Pixel, and Pixel 2), iOS phones (iPhone 7, iPhone 7+, and iPhone 8), and Android phones (Galaxy S7 and Huawei P20 lite). We considered experimental applications often used in our daily lives, ranging from instant messaging to cameras and mobile games. Our focus was on the major use scenarios that utilize various hardware components, such as the application processor, camera, and wireless network chipset: real-time video conversations, video recording, mobile gaming, and voice calling. For video conversations, we chose Google Hangouts [5] and Microsoft Skype [11], as they are two of the most popular video chat apps. For mobile gaming, we selected Abyssrium [12] and PUBG MOBILE (PUBG) [10], two of the popular games in recent years. For video recording and voice calling, we used the default app on the smartphones.

We considered the mobile operating systems of Android 6.0.1 and iOS 10.2. Note that Galaxy S7 uses Android 7.0, Pixel and P20 lite use Android 8.0.0, Pixel 2 uses Android 8.1.0, and iPhone 8 uses iOS 11.4.1; these are the default operating systems shipped. We controlled the running environments as follows. There were no background processes. All experiments started with the battery fully charged. Auto-update, battery saving, and adaptive screen brightness control modes were disabled. For video recording, we set the same video resolution (1080p) across multiple devices. For wireless connectivity, we used Wi-Fi unless otherwise noted. We restricted touch interactions, which happen only when starting or ending an app. The mobile gaming scenarios required more touch interactions than others, but we made our best effort to minimize the interactions (e.g., by leveraging autoplay modes). The ambient temperature of the laboratory was maintained between 24°C and 28°C. As the heat transfer equation shows, the amount of heat transferred due to convection (air cooling) is much smaller than that due to direct conduction across multiple components from the heat sources. Thus, a 4°C variation of room temperature had a negligible effect on our measurement result.

We ran each app for 30 minutes, because all our tested smartphones reached steady-state temperature within 30 minutes. The steady-state temperature is considered as the mean value of the samples measured for 2 minutes starting from the peak temperature in each 30-minute experiment (a total of 3 measurements performed). Since our measurement tool indicates a 0.3 Hz sampling rate on average, every steady-state temperature is calculated from at least 100 samples.

3.2 Results

Figure 2 presents the overall thermal characteristics of the smartphones under various use cases. The apps for each category show similar temperature change tendencies. After the initial temperature ramp-up, a smartphone quickly reaches its steady-state. After the running app is closed, it then gradually dissipates its heat. The smartphones under consideration converged to their steady-state temperatures within 20 minutes. Table 1 summarizes the steady-state temperature for each smartphone while running various scenarios. Our tested smartphones reached the steady-state temperatures that were even higher than 45°C in some cases (i.e., video chatting, video recording, and mobile gaming). The results clearly show that thermal issues are prevalent in modern smartphones, and users could be exposed to high temperatures that lead to considerable thermal discomfort.

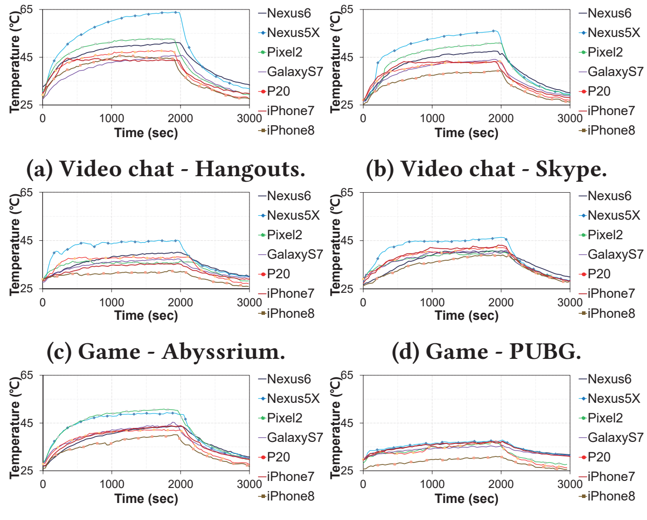

(a) Video chat - Hangouts. (b) Video chat - Skype.

(c) Game - Abyssrium. (d) Game - PUBG.

(e) Video recording. (f) Voice calling.

Figure 2: Temperature changes of the hottest region on each phone while running various target apps.

Table 1: Steady-state temperature (°C) for each smartphone while running various target scenarios (standard Deviation in parenthesis).

Temp. (°C)	Skype	Hangouts	Abyssrium	PUBG	Video rec.	Voice call
Nexus 5	56.30 (2.76)	53.30 (0.67)	43.90 (0.53)	46.20 (0.62)	49.44 (1.77)	36.86 (1.83)
Nexus 6	51.88 (2.66)	49.34 (2.95)	44.13 (1.96)	40.66 (0.08)	47.64 (1.31)	36.88 (0.39)
Nexus 5X	58.13 (3.10)	62.28 (3.50)	45.94 (1.18)	45.19 (0.76)	50.79 (0.52)	33.02 (1.61)
Pixel	49.90 (1.45)	50.02 (1.19)	45.58 (0.25)	43.74 (1.52)	46.61 (2.51)	38.78 (2.54)
Pixel 2	51.12 (0.62)	52.65 (0.12)	36.16 (0.44)	38.27 (0.76)	51.64 (0.75)	36.63 (0.72)
Galaxy S7	46.93 (1.44)	45.69 (0.59)	42.68 (0.90)	40.69 (0.51)	45.12 (1.63)	36.52 (0.96)
P20 lite	43.36 (0.56)	47.33 (0.46)	37.63 (0.63)	41.48 (1.33)	43.12 (0.94)	36.94 (0.51)
iPhone 7	45.85 (2.10)	44.29 (0.89)	39.52 (0.46)	42.32 (0.52)	45.72 (1.21)	35.61 (2.32)
iPhone 7+	46.53 (2.99)	44.34 (0.92)	36.16 (0.71)	41.92 (0.81)	46.57 (1.32)	36.39 (1.32)
iPhone 8	39.69 (0.27)	45.30 (0.36)	33.59 (0.42)	38.27 (0.76)	40.05 (0.90)	31.86 (0.23)

4 IN-DEPTH THERMAL MEASUREMENT

We learned from Section 3 that thermal concerns are not confined to specific smartphone models, as all tested smartphones exhibit thermal issues. To deepen our understanding on thermal characteristics of smartphones, we performed in-depth thermal measurements by focusing on an Android reference phone (i.e., Nexus 5X). This model has hardware information provided by the kernel, and it is one of the recent reference phones with modern thermal management techniques. We also performed measurements at the component level (e.g., AP, power management IC, camera modules) to identify the root causes of excessive heat generation. In

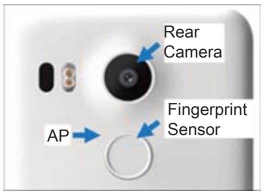

(a) Visible light image.

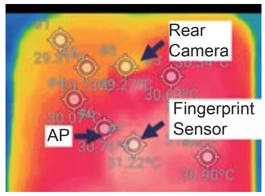

(b) Infrared image.

Figure 3: The relative positions of the rear camera, the fingerprint sensor, and the AP of Nexus 5X.

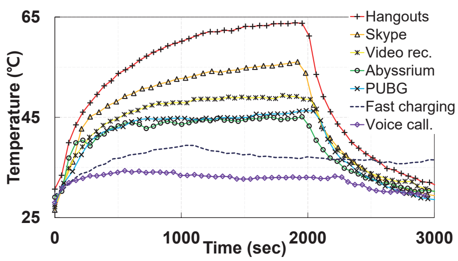

Figure 4: The temperature changes at the fingerprint sensor while running each application.

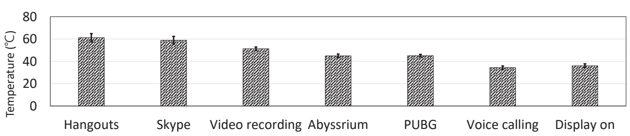

Figure 5: Mean steady-state temperatures of each scenario with three Nexus 5X phones. The error bar shows standard deviation.

addition, we measure the surface temperature of various ubiquitous usage scenarios.

4.1 Surface Temperature Analysis

Figure 3 shows the rear side of Nexus 5X including an example thermographic image captured by the IR camera. We measured the temperature of the surface, both front and back, when running a video chat (Hangouts [5], Skype [11]), video recording, voice calling, or 3D mobile games (Abyssrium [12], PUBG [10]), and when charging the battery.

Figure 4 shows the temperature change of the smartphone surface while running each application. For each experiment, once the temperature reaches the maximum, it stabilizes. When the application is closed, it slowly cools down. In general, the areas near the fingerprint sensor show higher temperatures than the other regions, as the fingerprint sensor has lower thermal resistance than its surroundings. We measured the surface temperature changes of *three different Nexus 5X phones* for all the cases to check whether they showed consistent trends. Our results in Figure 5 confirmed that steady-state temperature differences were within the error range, and only minor temperature variations were observed among different phones of the same model.

Video chat: It took only 6 minutes to exceed the thermal pain threshold, with Hangouts reaching 62°C (SD: 3.50) and Skype 58°C (SD: 3.10) in 30 minutes. When we used the front camera during a video chat, the surface temperature slightly decreased to 55°C (SD: 1.59). Since the rear-side camera is located near the fingerprint sensor, the heat from the rear camera affects the temperature measured at the fingerprint sensor more than that from the front camera. For this reason, the video chat using the rear camera showed a higher temperature than that of using the front camera. Regardless, the temperatures in both cases far exceeded the thermal pain threshold. Note that when we turned off the cameras, the temperature was only 38°C during a 30-minute experiment. We additionally ran Skype over LTE, and it showed similar temperature changes (58°C, SD: 1.04) as Skype over Wi-Fi.

Video recording: Video recording showed the second highest temperature in our experiment. The temperature took 8 minutes to exceed the thermal pain threshold, and the value reached 51°C (SD: 0.52) in 20 minutes. We examined video recording with different resolutions. Since higher resolution requires more computation, it shows slightly higher temperatures than a lower resolution. According to a prior study [41], video encoding typically requires more computation than decoding. When we simply turned on the built-in camera without recording, the temperature reached approximately 46°C in 15 minutes, which is almost similar to that of the video recording (48°C at the 15th minute). In contrast, playing a 1080p video (39°C) or turning on the screen without running any apps (34°C) showed a much lower temperature than in the camera-enabled case.

Game: We tested Abyssrium and PUBG. As the 3D games show diverse graphic effects, the processors must handle various graphic rendering operations. The results revealed that the temperatures settle around 45°C (SD: 0.76), with several minor peaks and valleys, depending on the occurrence of in-game events. As 45°C is the pain threshold temperature, this result demonstrates that game playing on smartphones for a long duration can also be a source of thermal discomfort.

Voice call: Voice calling generated the least amount of heat in our experiment, approaching only 33°C (SD: 1.61), which is lower than when we only turned on the screen without running any applications (34°C). However, unlike other usage scenarios, the skin contact region for voice calling is the user's cheek, which is known to have a high level of thermal sensation [46].

Charging: Simply charging the battery also showed moderate temperature (about 37°C), even when we applied fast charging until the smartphone is fully charged (100 mins

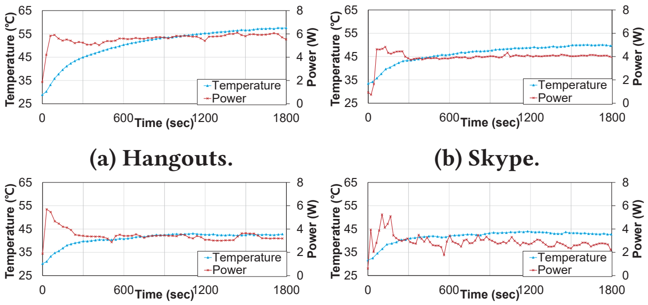

(a) Hangouts.

(b) Skype.

(c) Abyssrium.

(d) PUBG.

Figure 6: Temperature and power consumption.

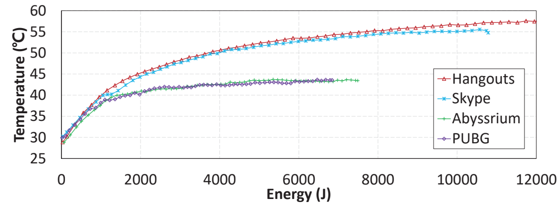

Figure 7: Energy vs. surface temperature.

from 2% to 100%). Therefore, battery charging by itself would likely not cause thermal problems.

However, in our experiment, we found an interesting phenomenon in the usage scenarios while charging, i.e., running video chat or video recording at a low battery level. In the beginning phase of the experiment (only 7% battery level), the temperature at the fingerprint sensor showed similar changes to that of the non-charging experiment. At this moment, the rate of discharging due to app usage was faster than the rate of charging, so the battery level was falling. Interestingly, when the battery level reached 0%, it suddenly accelerated the rate of charging and in consequence, the rate of heat generation significantly increased. In the end, the temperature of the fingerprint sensor rose to 72°C, which was 14°C higher than that of the non-charging experiment.

We observed a similar pattern when we ran video recording while charging at 5% battery level. According to its technical manual, Nexus 5X uses a special charging algorithm [8] that selects the charging rate based on the remaining battery level and the computational intensity of the running application. This phenomenon shows that charging algorithms could be another heating factor that influences thermal characteristics of smartphones.

Cooling down: Although each app converges at different temperatures at different heating phases, they tend to show similar cooling down trends. After the heating phase ends, the smartphone enters the sleep mode, consuming minimal energy. The heat energy then spreads from the high temperature regions to nearby low temperature materials. The rate of heat dispersion into the nearby solid object is faster than that into the air, and thus the hottest object shares the heat energy with the neighboring solid materials. As the object is cooling off, it finally arrives at the same temperature as the surrounding materials. Afterward, it loses the remaining heat energy at the same rate until the temperature becomes the same as the ambient temperature.

Energy vs. surface temperature: In order to examine how energy consumption patterns are related to corresponding thermal behaviors, we recorded the power consumption and temperature values simultaneously while running four most heat generating apps (i.e., Skype, Hangouts, Abyssrium and PUBG). We developed a measurement app that records voltage and current values at an approximately 0.3 Hz sampling rate by accessing the Android system's proc files [42].

The results in Figure 6 show two interesting observations. There was a surge of power consumption in the beginning, possibly due to app initialization (e.g., data fetching and content loading). Furthermore, there is a tendency of increasing power consumption even with proactive thermal management (e.g., DVFS). This phenomenon might be attributed to the fact that the power leakage increases as the temperature rises [23, 53].

Figure 7 shows the relationship between energy consumption and surface temperature. It shows that four apps consumed different amount of energy to reach the same temperature, but similar apps showed resembling energy consumption patterns. While the amount of heat generated is generally proportional to energy consumption, surface temperature increment before reaching the steady state is closely related to the heat transfer rate, which is proportional to the temperature difference (or ΔT) as shown in Section 2.2. This means that high power consumption associated with a specific hardware component (e.g., AP) can significantly increase a component's temperature and thus expedite surface overheating due to large temperature differences.

Impact of app restarting: The power graphs in Figure 6 hint that app restarting may result in different temperature increment trends due to high power consumption during the initialization phase. To validate this intuition, we measured the temperature changes of two apps (i.e., video recording and PUBG) by restarting the apps when the surface temperature dropped by 50% of the temperature difference between the steady-state temperature and the baseline temperature. As discussed in Section 2.2, we can model the temperature at time t as follows [39]: $T(t) = T_\infty - (T_\infty - T_0)e^{-\tau t}$ where T_∞ is the peak temperature, T_0 is the initial temperature, and τ is a coefficient that represents how fast the temperature changes. The curve fitting results showed that restarting has a larger coefficient τ: video (45 vs. 62) and PUBG (61 vs. 73) (in units of $10^{-4}/s$).

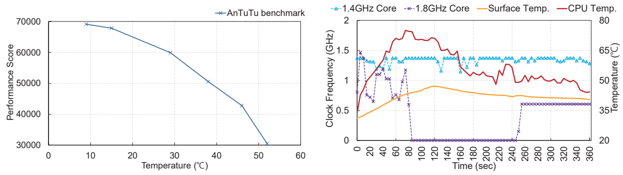

(a) AnTuTu benchmark performance vs. ambient temperature.

(b) Clock frequency scaling behaviors (with an ambient temperature of 28°C).

Figure 8: AnTuTu benchmark test results of Nexus 5X.

Impact of a protection case: We measured the temperature changes of Skype with and without a plastic protection case, which is made of Thermoplastic Poly-Urethane (TPU) and Polycarbonate (PC) (thickness=1.5mm). Heat transfer principles hint that the thicker the medium, the slower the heat propagation. Our measurement results confirmed that surface temperature rose slowly; e.g., to reach 45°C, it took 418 seconds (SD: 22) without the case and 643 seconds (SD: 46) with the case. Furthermore, the steady-state temperature with the case (48°C, SD: 2.13) was slightly lower than that without the case (54°C, SD: 0.21).

Performance implications: Figure 8a shows the AnTuTu benchmark [2] performance at different ambient temperatures. The AnTuTu benchmark score represents the overall performance of the smartphone as it measures computation capability, video frame variance, etc. The result shows that the performance is significantly influenced by the temperature. The key reason for performance degradation is frequency scaling as shown in Figure 8b. The clock frequencies of two cores vary in accordance with the increase of CPU temperature. The 1.8 GHz cores are switched off when the temperature exceeds the threshold as determined by DVFS.

4.2 Component-specific Analysis

We next focus on analyzing the component-specific temperature characteristics of smartphones to identify which components contribute to heat generation. There are several heat sinks, and each covers one or more components. To identify the main heat source in each scenario, we removed the heat sinks and measured the temperature changes on both the front and back of the printed circuit board (PCB), by using two IR cameras. We ran video chat (Skype), video recording, and a 3D mobile game (Abyssrium), and these applications combined with fast charging. We also measured the component usages, such as CPU and GPU utilization, as well as the network data rate and clock frequency changes.

Figure 9 illustrates the steady-state temperature of the components while running each application and Table 2 summarizes the corresponding CPU/GPU utilization. The components include application processor (AP), power management integrated circuits (PMICs), rear camera, and battery. Here, PMIC refers to a class of integrated circuits that include various functions related to power requirements such as DC to DC conversion, battery charging, power-source selection, and voltage scaling. Nexus 5X has two PMICs; one is a regulator for power supply to other components, and the other is a battery charger for controlling the charging rate. Figure 9 shows that AP and PMIC are generally the main contributors of heat generation.

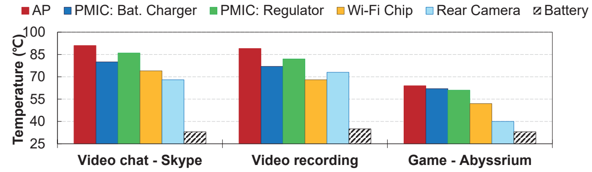

Figure 9: The steady-state temperature of each component while running each application.

Table 2: Utilization statistics of each component when running each application.

Utilization (%)	Video chat - Skype	Video recording	Game - Abyssrium
CPU	69 (SD: 5)	32 (SD: 2)	20 (SD: 11)
GPU	4 (SD: 2)	0 (SD: 1)	58 (SD: 17)

In Nexus 5X, the fingerprint sensor is located on top of the AP and PMIC. The fingerprint sensor directly connects into the PCB, and there is no air gap in between. Other regions have additional plastic shields to isolate the heat sources and the smartphone surface. The fingerprint sensor therefore shows a higher temperature than other regions.

Video chatting and recording: Both video chatting and recording display high temperatures as they involve heavy computation and camera usage. With Skype, CPU utilization was 69% on average, and the AP generated the most heat (91°C). Video recording also generated considerable heat similar to Skype on the AP, even though CPU utilization was about 30%. This is because the multimedia codec built into the AP, which encodes/decodes videos, generates significant heat. Regardless of Wi-Fi or LTE, the networking components did not show notable temperature changes. We measured the data rates of Skype using the Android monitor, and they were Tx: 698 kb/s and Rx: 396 kb/s on average.

Game: Similar to the camera-based applications, AP is the main source of heat generation in mobile gaming. During the game play, the CPU usage rate was only 20% on average, while the GPU usage rate increased to 58%. As both GPU and CPU are built into the AP, the AP remains a main heat source. Due to adaptive thermal management, the heat was much lower than other application usage scenarios.

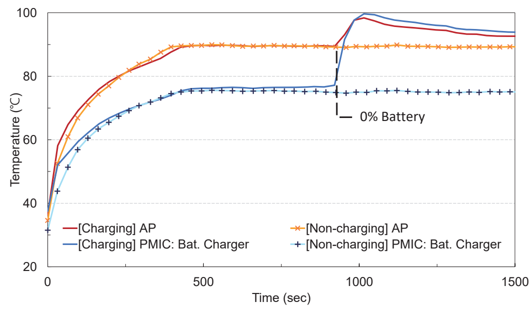

Figure 10: The AP and PMIC temperature changes while video recording, with and without charging.

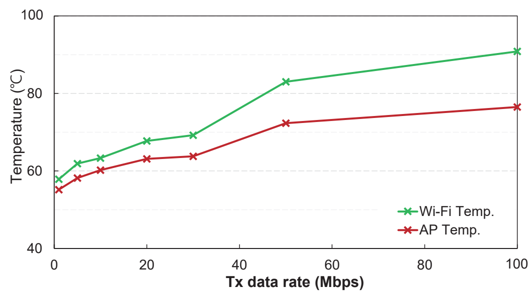

Figure 11: The steady-state temperatures of AP and Wi-Fi chip during 10 minute network transmissions.

Running applications while charging: Figure 10 shows that the temperature changes of both AP and PMIC while charging look similar to when non-charging. We started the experiment at 10% battery level and the discharging rate was faster than the charging rate. When the battery level dropped to 0%, the charging rate increased, causing the temperatures of both AP and PMIC to suddenly leap to nearly 100°C. Unlike other scenarios, PMIC exhibits a higher temperature than the AP. The charging IC and its paired PMIC modulate the voltage and the current. It appears that these components can be overloaded when we simultaneously use and charge the smartphone. This may be attributed to a charging algorithm controlled by the PMIC when the battery level reaches 0% [8].

Wi-Fi communication: We measured the temperature changes of the application processor (AP) and the Wi-Fi chip while transmitting UDP data packets with *iPerf*. Figure 11 reveals the higher the Tx data rate, the higher the temperature. The Wi-Fi chip showed a higher temperature than the AP, suggesting that the Wi-Fi chip could be a main heat source when the smartphone uses high-speed wireless communication.

CPU intensive vs. video processing workload: APs typically contain multiple modules such as CPUs, GPUs, multimedia codec, and memory. This means that simultaneous

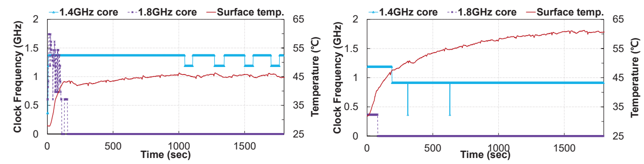

(a) CPU-intensive workload. **(b) Video rec. workload.**

Figure 12: Nexus 5X's surface temperature vs. frequency scaling behaviors.

use of multiple modules would likely cause surface overheating. To check this, we conducted two experiments: CPU-intensive workload (90% CPU utilization), and high-quality video recording. The results are presented in Figure 12.

Interestingly, CPU intensive workload alone did not cause significant overheating. The surface temperature did not exceed 46°C due to frequency scaling and core disabling. During video recording, however, the surface temperature continued to rise and reached 60°C in 20 minutes, despite frequency scaling and core disabling. This overheating originates largely from the video codec in the multimedia processing component built into recent mobile APs such as Qualcomm Snapdragon 808 in Nexus 5X. A hardware accelerator, such as a video codec in a mobile AP typically does not have its own thermal management capability as it simply runs the required workload. Thus, the higher the workload, the more heat it generates.

4.3 Ubiquitous Usage Analysis

Our smartphone temperature measurement studies helped us to characterize the surface temperature of well-known app usage scenarios and to identify the major heat sources. Beyond these scenarios, we measured the surface temperature in a variety of mobile and ubiquitous scenarios: virtual reality (VR) gaming, mobile, augmented reality (AR) gaming, GPS driving navigation, high-speed wireless data transfer, and mobile deep learning. Note that for VR gaming and high-speed wireless data transfer scenarios, we use Samsung Galaxy S7, because Gear VR only supports Samsung phones, and Nexus 5X's Wi-Fi Direct does not fully support high-speed data transfer (i.e., 300 Mb/s).

Virtual reality: We measured the temperature changes of Galaxy S7 connected to a Gear VR headset. We considered a scenario of playing a 3D runner game, called Temple Run 3. We measured the temperature of the air confined in the VR headset by using a K-type thermocouple attached inside the headset. After 30 minutes of game playing, the temperature of the smartphone surface rose to 42°C, and the air temperature reached 38°C. Although this level of temperature is less concerning than that of other scenarios, the confined space between the face and the device got wet with sweat due to

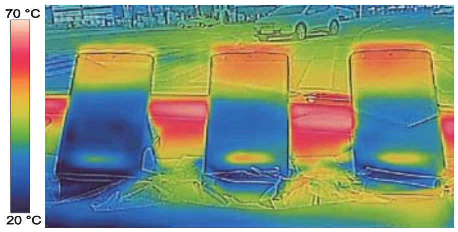

Figure 13: Dashboard navigation measurement.

the heat and the lack of ventilation, causing considerable thermal discomfort.

Augmented reality gaming: Pokémon Go [9], a location-based augmented reality game was considered for measurement due to its popularity. This game consistently uses diverse hardware components, such as AP, Wi-Fi, LTE, GPS, and camera, which could generate considerable heat. In particular, the AR mode displays the image captured by the camera and renders 3D objects over the background. If this mode was enabled, our measurement showed that the temperature quickly reached 53°C at the hot-spot of the surface. In addition, over half of the surface reached higher than 45°C. When the temperature was over 51°C in the AR mode, we observed unnatural and discontinuous rendering of 3D graphics due to thermal throttling. When the AR mode was disabled, the surface temperature was around 46°C.

GPS navigation: Navigation is one of the most popular vehicular smartphone applications. Since GPS is one of the main sources of energy consumption in smartphones [18], we hypothesized that continuous use of GPS and map rendering could generate considerable heat. Furthermore, due to direct sunlight exposure, there could be significant environmental influence on the surface temperature. To inspect the temperature/power relationship and an environmental factor, we considered the following GPS navigation applications: KakaoNavi [6], NaverMap [7], and Waze [13]. Note that Google and Apple Maps do not provide turn-by-turn navigation services in the locations where the experiments were performed.

We conducted an experiment by measuring the temperature of three smartphones (Fig. 13), each running a different navigation app mentioned above, using thermal cameras mounted on the dashboard in a car with air conditioning (22–24°C). We drove about 12 km for 30 minutes at four different times in the evening. Waze, KakaoNavi, and NaverMap reached on average 44°C, 42°C, and 41°C, respectively. Similar to the Pokémon Go results without the AR mode, these results show that running the navigation applications does not appear to cause serious thermal concerns, and GPS, while consuming a lot of energy, does not generate excessive heat.

We then examine the effect of sunlight, by measuring the surface temperature of the smartphones mounted on the dashboard under direct sunlight without air conditioning during the daytime. In this case, the temperature quickly rose up to 64°C in just 15 minutes. It appears that in-vehicle usage of GPS navigation can cause considerable thermal issues mainly due to the sunlight rather than the workload.

High-speed wireless communication: High-speed wireless networking scenarios such as local file sharing are becoming common as wireless networking technologies evolve such as Wi-Fi Device-to-Device (D2D). We examine how high speed data transmission affects the smartphone surface temperature. For accurate temperature measurement, we exchange a large video file using Galaxy S7 with the maximum throughput of 300 Mb/s. When transmitting the file, the temperature of Galaxy S7 using Wi-Fi Direct quickly rose to 51°C. The location where the Wi-Fi module resides showed the highest temperature, which is different from when video chatting using Wi-Fi, where the application processor generated the most heat. As discussed in Section 4.2, the network throughput of Skype was less than 1 Mb/s. We thus infer that the temperature of Wi-Fi module increases considerably as network throughput increases. Furthermore, the transmitter exhibits higher temperature (51°C) than the receiver (41°C). This can be partly explained based on the fact that the transmitter consumes more energy than the receiver [15].

GPU-based mobile deep learning: There is a growing demand of building deep learning platforms in mobile devices (e.g., Caffe2Go [25] and DeepMon [30]) to support continuous vision applications such as real-time style transfer and image recognition. As deep learning requires intensive computation, we measure how much heat deep learning applications generate.

For this measurement, we considered DeepMon, a mobile deep learning inference system that runs various deep learning algorithms on a mobile device and supports mobile GPU acceleration [30]. Using the DeepMon framework, we built and ran a simple deep neural network (DNN) based image processing app. We used Galaxy S7 for this measurement as the DeepMon framework is optimized for Galaxy S7's GPUs. Our app continuously captures and classifies images for object detection from the camera using the You Only Look Once (YOLO) model [51]. Our test app runs the DNN image processing method every three seconds, and the processing time of single instance takes about a second.

Our results showed that running this deep learning algorithm generated a temperature (54°C) that is 9°C higher than simply capturing images from the rear camera (45°C). While it took 21 minutes to reach this temperature, it easily exceeds the highest temperature that we measured from Galaxy S7 (47°C when running video chat). Given that the processing time of the DNN algorithm is less than about 33% of the total processing time, we expect that more heat could be generated depending on the GPU workload. Our results

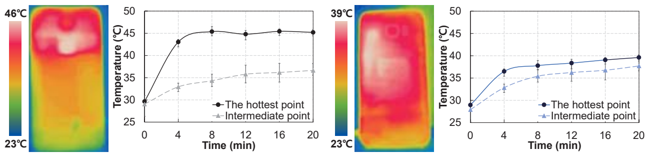

(a) Nexus 5X. **(b) Galaxy S7.**

Figure 14: Surface temperature during a gameplay.

clearly show that the chipset vendors and smartphone manufacturers should pay more attention to overheating issues particularly when GPUs are highly utilized as in the deep learning scenarios.

5 USER EXPERIENCE STUDY ON HEAT SENSATION AND DISCOMFORT

We conducted a user experience study to understand how surface heat affects user experiences: (i) how much heat and discomfort users feel when they use a smartphone for a prolonged time (e.g., 20 minutes of a game play), and (ii) what are user reaction patterns due to the surface heat.

5.1 Experiment Design

We recruited 20 participants (10 male and 10 female), with the average age of 21.45 (min=19, max=26). We asked the participants to play Yokai Saga [1], a popular 3D mobile action game. This game was chosen because (i) it received a good user rating (i.e., fun to play), (ii) its controls are intuitive and simple to handle, (iii) it requires frequent screen touches, and (iv) it caused smartphone surface heating with consistent trends. We used Nexus 5X and Galaxy S7 smartphones to account for variations on heat sensation and discomfort across different devices. As shown earlier, Nexus 5X had a slightly higher surface temperature (44.33°C) than Galaxy S7 (39.14°C), and took much shorter time (352 sec) than Galaxy S7 (606 sec) to reach the steady-state temperature (see Figure 14). When participants played the game, we recommended them to hold the phone in the landscape-mode with their hands; this recommendation was made to avoid the case of laying the phone on the desk.

We conducted an IRB-approved within-subjects experiment. The participants performed in two conditions, changing the smartphone models (i.e., Nexus 5X and Galaxy S7). To reduce the order effect, we counter-balanced the order such that a half of them used Nexus 5X first, and the other half used Galaxy S7 first. Each participant was rewarded with a $10 gift voucher.

We asked the participants to rate the level of heat sensation and discomfort every 4 minutes on a 5-point Likert scale: 1: Not at all, 2: Little, 3: Somewhat, 4: Very much, and 5: Extremely. These questions have been widely used in the field

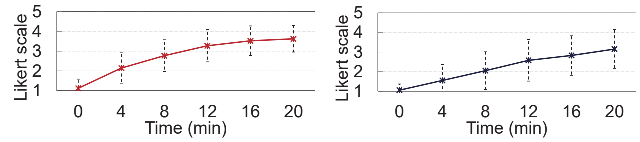

(a) Heat sensation. **(b) Heat discomfort.**

Figure 15: Heat sensation and discomfort during gameplay (error bars denote standard deviation).

of ergonomics [62, 63]. While heat sensation and discomfort levels were known to be correlated [63, 64], it is interesting to study the overall user experiences with smartphone usage scenarios. During the experiment, we video recorded the grabbing patterns to investigate their reactions to surface heat. After game plays, we conducted a follow-up interview.

5.2 User Study Results

Figure 15 shows the mean levels of heat sensation and discomfort during gameplay. The level of heat sensation increased at a similar rate to the temperature increase shown in Figure 14. A few participants indicated that the smartphone was very hot for the first 8 minutes during which the temperature rapidly increased. As the temperature increased so did the heat sensation and discomfort levels. Unlike heat sensation, we observed that the discomfort level increased almost linearly. As a result, 75% of the participants responded that they felt discomfort due to surface heat (levels 3 to 5 in Fig. 15b). P1 commented *"Because the smartphone was hot, my fingers became numb just like when the electricity flowed"* and P20 mentioned *"It was uncomfortable to keep holding the smartphone as it got slippery."*

To verify the time-series differences, we applied paired samples t-tests between two consecutive measurements (e.g., the 0-th minute and the 4-th minute measurements). The levels of heat and discomfort from the 0-th minute to the 16-th minute were significantly different. However, at the 16-th and the 20-th minutes, the levels of heat sensation were not different ($p=0.159$, Cohen's $d=0.143$), whereas the levels of thermal discomfort were different ($p=0.003$, Cohen's $d=0.323$). This shows that the participants did not perceive a significant difference in heat sensation at the 16th minute mark, but felt *more discomfort as time passed*. Despite slightly lower temperature, prolonged usage still can cause considerable discomfort, thereby negatively influencing user experiences.

In the post-interview, the participants reported thermal discomfort for various reasons: (i) heat sensation, (ii) concerns on skin burn, (iii) sweating, (iv) sensor malfunction caused by sweat, and (v) slip caused by sweat. Most participants said that Galaxy S7 was especially slippery because of its glass-based surface material. While discomfort levels were not significantly different, our participants consistently reported that using Galaxy S7 was more uncomfortable than

Table 3: Measured data types for prediction modeling.

System Statistics (21)	
CPU Statistics (17)	(1) Time spent on un-niced user processes, (2) Time spent on niced user processes (UserTime), (3) Time spent on kernel processes (SysTime), (4) Time spent in kernel idle handler, (5) Time waiting for I/O completion, (6/7) Time spent servicing HW/SW interrupts, (8) Total elapsed time for above seven items, (9) Percentage of time running user space (User%), (10) Percentage of time running kernel space (System%), (11) Percentage of time running I/O completion, (12) Percentage of time running HW interrupts, (13) Percentage of total CPU time (CPU%), (14) CPU's clock freq. (CPUfreq), (15) Total accessible address space of a process, (16) Total memory held in RAM for a process, (17) Number of threads
On-device Sensors (2)	(1) CPU temp. (cpuTemp), (2) Battery temp. (batTemp)
Network (2)	(1) Tx data rate (netTx), (2) Rx data rate (netRx)

Nexus 5X even though Nexus 5X showed a higher temperature during the experiments. Participants mentioned that the whole surface of Galaxy S7 was heated, whereas only one side of the surface (near the fingerprint sensor) of Nexus 5X was heated.

Our video recording analysis showed that as this game requires landscape-mode user interaction, 75% of the participants grabbed the smartphone with both hands, but some users played with one hand (left: 20%, right: 5%).

All participants were hence exposed to the hottest areas of each device during gameplay. We observed that most (95%) participants repeatedly changed their grips and stretched their hands. At the end of the experiment, we asked them why they changed the grip and 35% said that it was because the smartphone was too hot or their hands were sweaty. The other participants (65%) responded that it was a less conscious behavior and they wanted to resolve the discomfort in keep holding the smartphone.

6 INFERRING SURFACE TEMPERATURE

Our experiments have shown that smartphones could quickly overheat during popular usage scenarios and cause user discomfort or even injury. Estimating and predicting high surface temperatures could be a first step in lowering user discomfort due to heat. However, current smartphones only monitor the temperature of specific components and do not provide surface temperature measurement. We explore the feasibility of estimating surface temperature by using only Android system statistics and internal sensors (e.g., CPU and battery temperature).

6.1 Model Design

We collected two datasets from Nexus 5X and Galaxy S7: the surface temperature of each device and system statistics while running several usage scenarios. We selected features to identify the dominant factors that relate to surface temperature. We ran pre-processing to handle the time correlation factor to better predict upcoming temperature changes.

We present a prediction model for each device by means of smartphone system statistics without any additional hardware. We collected 21 system statistics by using the *top* command of the Android debug bridge and the Android system API (see Table 3) and measured the corresponding surface temperature every 5 seconds. The maximum CPU utilization of the logging app was about 3% for both Nexus 5X and Galaxy S7, suggesting that the logging process has a negligible overhead to collect proper amount of data.

To collect a dataset, we categorized a variety of popular smartphone usage scenarios and selected 11 representative scenarios. We selected six apps to cover the usage scenarios for training: games (HIT, Yokai Saga), cameras (Skype, video recording), networking (iPerf), and video streaming (YouTube). Each dataset was collected as follows. For each scenario, we performed temperature measurements four times, 30 minutes each. After removing erroneous measurements in some scenarios, the number of total measurements was 40 instances, and thus the total duration was 1,200 minutes. For evaluation, we used a leave-one-scenario-out method, by randomly choosing one scenario for testing and the rest for training.

Different smartphones might have different thermal management methods and their heat dissipation could be different due to the differences in internal components. For each phone, we thus must obtain representative features that could best predict temperature changes. For feature selection, we used the correlation-based feature selection (CFS) method [29], which selects the features that are highly correlated with the predictive variable but are less correlated with one another. We used *CorrelationAttributeEval* in Weka 3.8.2 for feature selection and chose the top 8 features each for Nexus 5X and Galaxy S7.

For prediction, while multiple regression could be a viable choice, it is known that simple multiple regression is less accurate when there is time correlation of neighboring data rows [20]. To address this concern, we performed preprocessing called time lagging (i.e., attaching neighboring data pairs) to make our data attain time dependency information as follows: $y_t = \vec{\alpha}^T X_t + \vec{\beta}^T X_{t-L} + \epsilon_t$ where y_t denotes the temperature at time t, X_t denotes the system log vector at time t, L is a lag number, $\vec{\alpha}$ and $\vec{\beta}$ are regression coefficient vectors, and ϵ_t denotes the corresponding error term. Note that $L = 0$ represents a regular multiple regression.

6.2 Evaluation

The dominant features are listed in Table 4. There are seven common features (e.g., the percentage of time running user/kernel

Table 4: The dominant features for each smartphone in descending order of correlation coefficient.

	Selected Features
Nexus 5X	batTemp, cpuTemp, CPUfreq, System%, User%, UserTime, SysTime, CPU%
Galaxy S7	batTemp, cpuTemp, CPU%, SysTime, UserTime, User%, System%, netTx

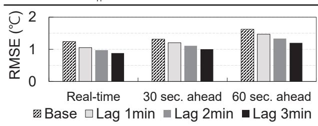

(a) Nexus 5X.

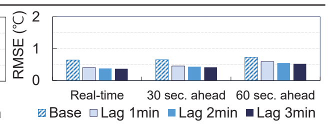

(b) Galaxy S7.

Figure 16: Prediction accuracy with time lagging.

space, total CPU utilization, CPU time spent in user/kernel space, CPU temperature, and battery temperature) for both smartphone models, and one distinct feature for each smartphone. The high utilization and run-time means that the CPU is intensively working and hence, the temperature will rise. The clock frequency has negatively correlated with the surface temperature due to DVFS. While Nexus 5X actively manipulated the clock frequency as shown in Figure 12b, Galaxy S7 did not show such frequency scaling behavior. Thus, the clock frequency feature was not included in Galaxy S7. Instead, the data transmission rate showed a high correlation with the surface temperature.

With these dominant factors as shown in Table 4, we trained our model for each smartphone to estimate the surface temperature. To do this, we assembled all training data from multiple apps to build a single model for each smartphone. The real-time prediction performance of our multiple regression model is shown in Figure 16. As an accuracy measure we used *RMSE* that denotes the root mean square error. Our model accurately predicts the surface temperature, with errors less than 2°C in most cases. This performance is superior to the model that uses only internal thermometers; a model using only *cpuTemp* and *batTemp* showed R^2: 0.87 and *RMSE*: 3.81°C for Nexus 5X. This pattern is also true for Galaxy S7: R^2: 0.82 and *RMSE*: 1.51°C. Since our measurement tool, the IR camera, has approximately 2°C error bounds at most, we conclude that our prediction method estimates the temperature changes accurately.

Figure 17 summarizes the model performance of predicting temperature changes. The overall prediction error was 1.24°C *RMSE* for Nexus 5X and 0.64°C *RMSE* for Galaxy S7. The models could also predict the temperature changes while running several new test apps that were not included in the training apps. Each model of Nexus 5X and Galaxy S7 showed 1.96°C and 0.983°C *RMSE* for *Google Hangouts*, and 2.46°C and 0.676°C *RMSE* for *Abyssrium*, respectively. We can improve the performance with the time lagging technique as shown in Figure 16. The *Lag t min* means each temperature datum is attached with two system log of *t* minutes away.

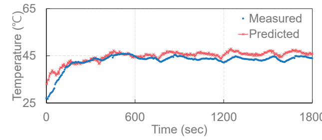

(a) 5X: Game HIT.

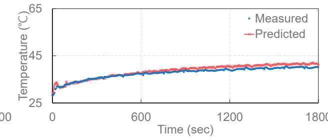

(b) S7: Game HIT.

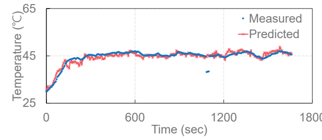

(c) 5X: Game Yokai Saga.

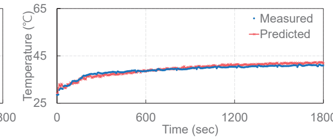

(d) S7: Game Yokai Saga.

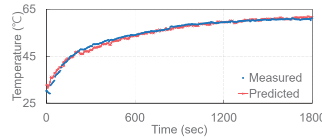

(e) 5X: Video chat Skype.

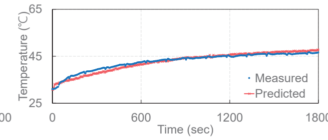

(f) S7: Video chat Skype.

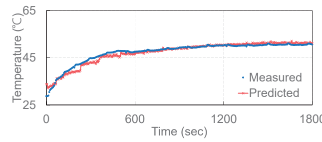

(g) 5X: 1080p video rec.

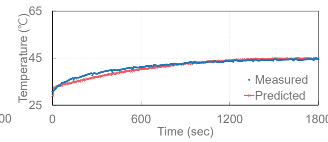

(h) S7: 1080p video rec.

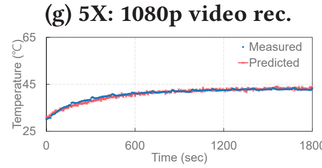

(i) 5X: Tx 1Mbps networking.

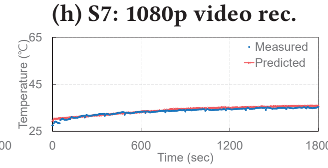

(j) S7: Tx 1Mbps networking.

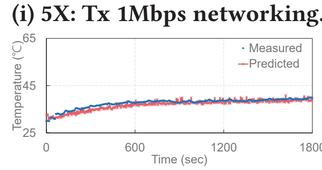

(k) 5X: 4K stream YouTube.

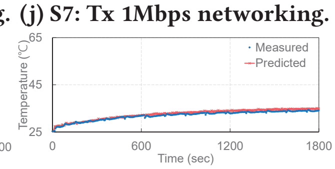

(l) S7: 4K stream YouTube.

Figure 17: Model performance over various apps.

It is natural that the farther the model predicts, the lower performance it shows. However, we observe that our model yields better performance when we attach the farther lagging data. Thus, when we choose 3 minutes lagging data, we can estimate the temperature within 1.20°C RMSE for Nexus 5X and 0.523°C RMSE for Galaxy S7.

7 DISCUSSION

7.1 Implications and Recommendations

User experiences: While user interface research extensively studied how heat can be used as an output modality (e.g., subtle information delivery) [32], no prior studies systematically examined the overheating patterns under various workloads

and their impact on smartphone user experiences. Our experiment showed that various real-world workloads generate excessive heat, and our user study confirmed that overheating degrades user experiences. Emerging ubiquitous applications such as AR/VR and mobile deep learning also showed significant thermal issues. Indeed, overheating might incur serious discomfort or even pain since the smartphones can directly touch heat-sensitive facial regions [46]. Our VR results indicate that wearable devices (e.g., glasses, virtual reality headsets) are likely to have more serious thermal issues, as they are worn and attached to the body. LiKamWa et al.'s work on thermal concerns of smart glasses [39] emphasized the importance of placement and power control of heat sources.
OS design: Mobile operating systems can provide options for users: system-forced regulation that the device cannot exceed the temperature threshold, and user-imposed regulation whether to take a (mildly) high temperature per user's setting. Currently only few phones show warning messages before they turn off due to high temperature; others switch off even without any warning. It would be helpful for users to be notified and make them respond in advance before their smartphones shut down. Our prediction model would be very useful to enable these options. Furthermore, we suggest a new evaluation criterion of smartphones and apps for thermal awareness. As in PowerForecaster [43] that allows users to check power efficiency before app downloading, we can display thermal ratings for informed app downloading.

There was an attempt to consider a smartphone's surface temperature for DVFS based thermal management [24]. However, our results demonstrated that DVFS alone is not sufficient for managing surface temperature. There should be a holistic thermal management framework that simultaneously considers multiple components' thermal characteristics (e.g., PMIC, cameras, sensors, and wireless chipsets). Our prediction model can be integrated into this framework.
Developer support: Our participants commented that overheating disrupted interaction flows and degraded interaction performance (e.g., input errors due to sweat). We recommend that app developers systematically consider user interaction as well as user reaction patterns (e.g., phone grab patterns, body part contact patterns). Overall, our thermal model can be useful for understanding the thermal impacts of their apps, including user interaction patterns. Our prediction model only requires *one-time training* per phone model and works in a real-time system without using external thermometers. As in several power profilers (e.g., Eprof [49] and PowerForecaster [43]), our model can track which components generate heat and also can be integrated into the Android IDE. This allows app developers to simulate the thermal characteristics of their apps. Besides such profilers, OS-level API supports for real-time surface temperature measurement will enable an array of user level applications including temperature warning.

7.2 Limitations and Future Work

For generalizability of our findings, we need further measurement studies on a wide range of recent phones and application workloads. We believe that our measurement of ten recent phones provided ample evidence of thermal issues in smartphones. Our deep learning results demand further studies on emerging mobile AI applications (e.g., real-time image processing and intelligent assistants) [35, 40] as they also require high computation power and large data transfer. Note that 3D gaming requires careful measurement considerations because various in-game events may cause heterogeneous CPU/GPU workloads. In this case, surface heating happens over a longer period of time (say at least several minutes) as opposed to temporal variations of in-game events. Thus, we hypothesize that the impact of in-game events and their variations would not be significant. Nonetheless, further measurements are required to accurately understand the impact of diverse in-game events.

Our prediction model is device-dependent and requires new model building with new phones. This model can be applicable to Android smartphone models, and it only requires *one-time training* per model. The model can be extended to predict multiple points by adding more measured points during the training stage. In our work, we have not considered running multiple apps at the same time, as the Android system prioritizes resource allocation to the foreground app. It is likely that background apps are mostly CPU-intensive workloads. According to our results, their thermal issues may not be significant. However, examining diverse coexistent scenarios including app restarting would still be a very interesting direction for future work. Furthermore, it would be interesting to explore various methods of improving prediction accuracy, by considering diverse hardware stats (e.g., GPU utilization), and leveraging advanced prediction algorithms (e.g., recurrent neural networks).

8 RELATED WORK

A number of studies proposed thermal models and management methods for mobile and wearable devices. Sekar [53] discussed major challenges in power and thermal management of mobile devices and emphasized the critical role of heat in the power context, calling for further studies on thermal measurement, modeling, management, and user experience research.

Singla et al. [55] presented a dynamic thermal and power management algorithm for heterogeneous multiprocessor systems-on-chip (MPSoCs) powering mobile platforms. Im et al. [31] designed a circuit-level temperature prediction method for electronic devices. Kwon et al. [34] proposed a thermal prediction method for N-App N-Screen (NANS) services on smartphones by considering the multi-core CPU and the display interface chipsets. All of these studies focused on measuring and predicting the temperature changes of APs, and none of the studies measured the surface temperatures, which could directly impact user experiences. Egilmez et al. [24] showed the feasibility of surface temperature estimation using CPU usage and CPU/battery temperature and proposed a novel DVFS method that considers surface temperature. This work however did not consider various system statistics information, and only basic linear regression was considered for modeling. Our model includes the basic regression model as well (i.e., no lagging case). Our results showed that the temporal series modeling with lagging considerably improves the prediction accuracy, when compared with the basic regression case. Park et al. [47, 48] considered power consumption characteristics of hardware components to predict surface temperature, but this model requires kernel-level hardware usage information (e.g., current video recording resolution, clock frequency information) for model construction. Our prediction model differs from these prior studies in that we use a variety of system stats and on-device sensors that are accessible at user-level processes to train time-series regression models, and we validate the accuracy of our real-time prediction model by considering diverse application workloads with different smartphones.

Xie et al. [60] developed a compact-thermal-modeling-based thermal simulator that simulates how heat is dissipated from its source to the surface. This work focused on modeling the thermal resistance network of smartphones, while our work mainly investigates how different applications affect the thermal conditions of the smartphones and the impact on user experiences. Chiriac et al. [21] developed a heat spreading metric, namely coefficient of thermal spreading that quantifies how well the generated heat is spread on smartphone surface. This metric helps hardware designers to better analyze the thermal design effectiveness. The authors used surface temperature data to estimate the spreading metric, but they did not report comprehensive measurement results with diverse workloads.

Several studies investigated the relationship between various thermal conditions of mobile/wearable devices and user experiences. Suh et al. [57] proposed the user burden scale that quantifies how user burdens including physical discomfort affect user experiences. In practice, physical burdens due to overheating can result in a negative effect on initial adoption, retention, and overall user experience. Zhang et al. [63] studied how laptop temperature is related to user discomfort by conducting thermal measurements of laptops at different relative positions. LiKamWa et al. [39] measured the energy and thermal characteristics of Google Glass. They found that video chatting reached around 50°C in 10 minutes, which far exceeded the thermal pain threshold. Thus, they suggested that head-mounted devices carefully regulate surface temperature. Straume et al. [56] analyzed the health risks of mobile phone use due to RF radiation and heat exposure by measuring a user's skin temperature in the head.

While these studies emphasized the effect of thermal issues in mobile/wearable devices on user experiences and health risks, there is still a lack of systematic analyses of surface and component-specific temperature measurements of various recent smartphones under various application workloads, and the overall impact of surface temperature on user experiences.

9 CONCLUSIONS

With emerging energy demanding applications, smartphones with powerful processors are being introduced, thereby challenging thermal management practices. Our research provides a first step towards investigating thermal issues associated with smartphones under various workloads. We performed extensive measurements and analyzed whether practical smartphone use gives rise to thermal concerns. We discovered that many applications, especially those using the camera, generated excessive heat, often surpassing the thermal pain threshold. The major heat sources include application processors, camera, and power management modules. Furthermore, we conducted a user study to examine how users perceive heat and discomfort while using smartphones. 75% of our study participants reported discomfort due to heat while playing a mobile game. Overheating disrupted user interaction flows and sometimes caused interaction errors. We also built a system statistics-based smartphone surface temperature prediction model and showed that our model yielded a mean squared error less than 2°C.

Many of us use smartphones often and for long duration. We make direct physical contact with smartphones, and thus thermal problems affect user experiences. We believe that the thermal issue associated with smartphones is an important and difficult problem. We hope our study raises awareness of this topic, and the mobile computing communities, both research and industry, actively work together to improve smartphone thermal computing.

10 ACKNOWLEDGMENTS

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (No.2016R1A2B4014068) and the Next-Generation Information Computing Development Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science and ICT (NRF-2017M3C4A7065960).

REFERENCES

[1] Yokai saga, 2017. https://goo.gl/pjGMd9.

[2] Antutu benchmark - know your android better, 2018. https://goo.gl/vEagph.

[3] Flir one, 2018. https://goo.gl/gdvO2Q.

[4] Google pixel 2, pixel 2 xl owners complain of overheating and poor battery life, 2018. https://goo.gl/viiYhq.

[5] Hangouts - android apps on google play, 2018. https://goo.gl/rpLL8W.

[6] Kakaonavi - apps on google play, 2018. https://goo.gl/YajfPM.

[7] Naver map - android apps on google play, 2018. https://goo.gl/h0iymG.

[8] Nexus 5x, 2018. https://goo.gl/UbuosR.

[9] Pokémon go - android apps on google play, 2018. https://goo.gl/xo7oBb.

[10] Pubg mobile - android apps on google play, 2018. https://goo.gl/6bHv1o.

[11] Skype - free im & video calls - android apps on google play, 2018. https://goo.gl/gBKjl6.

[12] Tap tap fish - abyssrium - android apps on google play, 2018. https://goo.gl/CxP5wp.

[13] Waze - gps, maps & traffic - android apps on google play, 2018. https://goo.gl/dtWUuL.

[14] Y. Abdelrahman, M. Khamis, S. Schneegass, and F. Alt. Stay cool! understanding thermal attacks on mobile-based user authentication. In *Proceedings of the 35th Annual ACM Conference on Human Factors in Computing Systems, ser. CHI*, volume 17.

[15] M. Anand, E. B. Nightingale, and J. Flinn. Self-tuning wireless network power management. *Wireless Networks*, 11(4):451–469, 2005.

[16] D. Brooks and M. Martonosi. Dynamic thermal management for high-performance microprocessors. In *High-Performance Computer Architecture, 2001. HPCA. The Seventh International Symposium on*, pages 171–182. IEEE, 2001.

[17] Safety of machinery – Temperatures of touchable surfaces – Ergonomics data to establish temperature limit values for hot surfaces. Standard, British-Standard-Institution, 1994.

[18] A. Carroll, G. Heiser, et al. An analysis of power consumption in a smartphone. In *USENIX annual technical conference*, volume 14, pages 21–21. Boston, MA, 2010.

[19] Y. A. Cengel. *Heat transfer : a practical approach*. McGraw-Hill New York, NY, 2004.

[20] C. Chatfield. *The analysis of time series: an introduction*. CRC press, 2016.

[21] V. Chiriac, S. Molloy, J. Anderson, and K. Goodson. A figure of merit for smart phone thermal management. *Electronics COOLING*, 2015.

[22] J. R. Culham, W. A. Khan, M. M. Yovanovich, and Y. S. Muzychka. The influence of material properties and spreading resistance in the thermal design of plate fin heat sinks. *Journal of electronic packaging*, 129(1):76–81, 2007.

[23] K. DeVogeleer, G. Memmi, P. Jouvelot, and F. Coelho. Modeling the temperature bias of power consumption for nanometer-scale cpus in application processors. In *Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIV), 2014 International Conference on*, pages 172–180. IEEE, 2014.

[24] B. Egilmez, G. Memik, S. Ogrenci-Memik, and O. Ergin. User-specific skin temperature-aware dvfs for smartphones. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, pages 1217–1220. EDA Consortium, 2015.

[25] I. Facebook. Delivering real-time ai in the palm of your hand., 2016. https://goo.gl/M37BSs.

[26] A. Faghri. *Heat pipe science and technology*. Global Digital Press, 1995.

[27] S. P. Gurrum, D. R. Edwards, T. Marchand-Golder, J. Akiyama, S. Yokoya, J.-F. Drouard, and F. Dahan. Generic thermal analysis for phone and tablet systems. In *Electronic Components and Technology Conference (ECTC), 2012 IEEE 62nd*, pages 1488–1492. IEEE, 2012.

[28] L. G. Hagander, H. A. Midani, M. A. Kuskowski, and G. J. Parry. Quantitative sensory testing: effect of site and skin temperature on thermal thresholds. *Clinical neurophysiology*, 111(1):17–22, 2000.

[29] M. A. Hall. Correlation-based feature selection for machine learning. 1999.

[30] L. N. Huynh, R. K. Balan, and Y. Lee. Deepmon: Building mobile gpu deep learning models for continuous vision applications. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, pages 186–186. ACM, 2017.

[31] Y.-H. Im, P. Kyol, and T.-J. Cho. System and method for predicting the temperature of a device, Nov. 1 2016. US Patent 9,482,584.

[32] L. A. Jones and M. Berris. The psychophysics of temperature perception and thermal-interface design. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2002. HAPTICS 2002. Proceedings. 10th Symposium on*, pages 137–142. IEEE, 2002.

[33] A.-G. Kibbi and Z. Tannous. Skin diseases caused by heat and cold. *Clinics in dermatology*, 16(1):91–98, 1998.

[34] O. Kwon, W. Jang, G. Kim, and C.-G. Lee. Accurate thermal prediction for nans (n-app n-screen) services on a smart phone. In *2018 IEEE 13th International Symposium on Industrial Embedded Systems (SIES)*, pages 1–10. IEEE, 2018.

[35] N. D. Lane, P. Georgiev, and L. Qendro. Deepear: robust smartphone audio sensing in unconstrained acoustic environments using deep learning. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 283–294. ACM, 2015.

[36] E. Larson, G. Cohn, S. Gupta, X. Ren, B. Harrison, D. Fox, and S. Patel. Heatwave: thermal imaging for surface user interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2565–2574. ACM, 2011.

[37] J. Lawrence and J. Bull. Thermal conditions which cause skin burns. *Engineering in Medicine*, 5(3):61–63, 1976.

[38] R. M. Leith Huffadine. Pregnant woman suffers burns after falling asleep on her iphone, 2016. https://goo.gl/FQAcnq.

[39] R. LiKamWa, Z. Wang, A. Carroll, F. X. Lin, and L. Zhong. Draining our glass: An energy and heat characterization of google glass. In *Proceedings of 5th Asia-Pacific Workshop on Systems*, page 10. ACM, 2014.

[40] C. Liu, L. Zhang, Z. Liu, K. Liu, X. Li, and Y. Liu. Lasagna: towards deep hierarchical understanding and searching over mobile sensing data. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 334–347. ACM, 2016.

[41] C. X. Mavromoustakis, E. Pallis, G. Mastorakis, et al. *Resource management in mobile computing environments*, volume 3. Springer, 2014.

[42] Y. Michalevsky, A. Schulman, G. A. Veerapandian, D. Boneh, and G. Nakibly. Powerspy: Location tracking using mobile device power analysis. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 785–800, Washington, D.C., 2015. USENIX Association.

[43] C. Min, Y. Lee, C. Yoo, S. Kang, S. Choi, P. Park, I. Hwang, Y. Ju, S. Choi, and J. Song. Powerforecaster: Predicting smartphone power impact of continuous sensing applications at pre-installation time. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 31–44. ACM, 2015.

[44] K. Mowery, S. Meiklejohn, and S. Savage. Heat of the moment: characterizing the efficacy of thermal camera-based attacks. In *Proceedings of the 5th USENIX conference on Offensive technologies*, pages 6–6. USENIX Association, 2011.

[45] B. Nagendran, A. Raghupathy, and W. Maltz. Thermal management challenges in forced convection tablets. In *Thermal Measurement, Modeling & Management Symposium (SEMI-THERM), 2015 31st*, pages 37–40. IEEE, 2015.

[46] M. Nakamura, T. Yoda, L. I. Crawshaw, S. Yasuhara, Y. Saito, M. Kasuga, K. Nagashima, and K. Kanosue. Regional differences in temperature sensation and thermal comfort in humans. *Journal of Applied Physiology*, 105(6):1897–1906, 2008.

[47] J. Park, S. Lee, and H. Cha. Accurate prediction of smartphones' skin temperature by considering exothermic components. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2018*, pages 1500–1503. IEEE, 2018.

[48] J. Park, S. Lee, and H. Cha. App-oriented thermal management of mobile devices. In *Proceedings of the International Symposium on Low Power Electronics and Design*, page 36. ACM, 2018.

[49] A. Pathak, Y. C. Hu, and M. Zhang. Where is the energy spent inside my app?: fine grained energy accounting on smartphones with eprof. In *Proceedings of the 7th ACM european conference on Computer Systems*, pages 29–42. ACM, 2012.

[50] Qualcomm Technologies, Inc. Dragonboard™ 410c based on qualcomm® snapdragon™ 410e processor thermal debugging guide. 2016. https://goo.gl/2Xfzg1.

[51] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.

[52] A. Sahami Shirazi, Y. Abdelrahman, N. Henze, S. Schneegass, M. Khalilbeigi, and A. Schmidt. Exploiting thermal reflection for interactive systems. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 3483–3492. ACM, 2014.

[53] K. Sekar. Power and thermal challenges in mobile devices. In *Proceedings of the 19th annual international conference on Mobile computing & networking*, pages 363–368. ACM, 2013.

[54] J. Y. Seo and J. H. Chung. Thermal aging: a new concept of skin aging. *Journal of Dermatological Science Supplement*, 2(1):S13–S22, 2006.

[55] G. Singla, G. Kaur, A. K. Unver, and U. Y. Ogras. Predictive dynamic thermal and power management for heterogeneous mobile platforms. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition*, pages 960–965. EDA Consortium, 2015.

[56] A. Straume, G. Oftedal, and A. Johnsson. Skin temperature increase caused by a mobile phone: a methodological infrared camera study. *Bioelectromagnetics*, 26(6):510–519, 2005.

[57] H. Suh, N. Shahriaree, E. B. Hekler, and J. A. Kientz. Developing and validating the user burden scale: A tool for assessing user burden in computing systems. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 3988–3999. ACM, 2016.

[58] D. Suleiman, M. Ibrahim, and I. Hamarash. Dynamic voltage frequency scaling (dvfs) for microprocessors power and energy reduction. 2005.

[59] G. Wagner and W. Maltz. Thermal management challenges in the passive cooling of handheld devices. In *Thermal Investigations of ICs and Systems (THERMINIC), 2013 19th International Workshop on*, pages 344–347. IEEE, 2013.

[60] Q. Xie, M. J. Dousti, and M. Pedram. Therminator: a thermal simulator for smartphones producing accurate chip and skin temperature maps. In *Low Power Electronics and Design (ISLPED), 2014 IEEE/ACM International Symposium on*, pages 117–122. IEEE, 2014.

[61] Q. Xie, J. Kim, Y. Wang, D. Shin, N. Chang, and M. Pedram. Dynamic thermal management in mobile devices considering the thermal coupling between battery and application processor. In *Computer-Aided Design (ICCAD), 2013 IEEE/ACM International Conference on*, pages 242–247. IEEE, 2013.

[62] H. Zhang, E. Arens, C. Huizenga, and T. Han. Thermal sensation and comfort models for non-uniform and transient environments, part ii: local comfort of individual body parts. *Building and Environment*, 45(2):389–398, 2010.

[63] H. Zhang and A. Hedge. Laptop heat and models of user thermal discomfort. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 58, pages 1456–1460. Sage Publications Sage CA: Los Angeles, CA, 2014.

[64] H. Zhang, A. Hedge, and D. Cosley. Thermal sensation, rate of temperature change, and the heat dissipation design for tablet computers. *Applied Ergonomics*, 62:197–203, 2017.

[65] L. Zhong, A. Bellemer, H. Yan, K. Honjo, J. Robertson, R. Y. Hwang, G. S. Pitt, and W. D. Tracey. Thermosensory and nonthermosensory isoforms of drosophila melanogaster trpa1 reveal heat-sensor domains of a thermotrp channel. *Cell reports*, 1(1):43–55, 2012.