Power-Temperature Stability and Safety Analysis for Multiprocessor Systems

GANAPATI BHAT, Arizona State University
SUAT GUMUSSOY, IEEE Member
UMIT Y. OGRAS, Arizona State University

Modern multiprocessor system-on-chips (SoCs) integrate multiple heterogeneous cores to achieve high energy efficiency. The power consumption of each core contributes to an increase in the temperature across the chip floorplan. In turn, higher temperature increases the leakage power exponentially, and leads to a positive feedback with nonlinear dynamics. This paper presents a power-temperature stability and safety analysis technique for multiprocessor systems. This analysis reveals the conditions under which the power-temperature trajectory converges to a stable fixed point. We also present a simple formula to compute the stable fixed point and maximum thermally-safe power consumption at *runtime*. Hardware measurements on a state-of-the-art mobile processor show that our analytical formulation can predict the stable fixed point with an average error of 2.6%. Hence, our approach can be used at runtime to ensure thermally safe operation and guard against thermal threats.

CCS Concepts: • **Computer systems organization** → **System on a chip**; • **Hardware** → **Temperature simulation and estimation**; **Platform power issues**;

Additional Key Words and Phrases: Power-temperature stability analysis, dynamic thermal and power management, multi-core architectures, mobile platforms

ACM Reference format:

Ganapati Bhat, Suat Gumussoy, and Umit Y. Ogras. 2017. Power-Temperature Stability and Safety Analysis for Multiprocessor Systems. *ACM Trans. Embed. Comput. Syst.* 16, 5s, Article 145 (September 2017), 19 pages. <https://doi.org/10.1145/3126567>

1 INTRODUCTION

Power consumption and the resulting heat dissipation are among the major problems faced by mobile platforms. Besides draining the battery, high temperature deteriorates the reliability and user experience [8, 17]. Recent evidence also shows that uncontrollable temperature increase in one part of the chip poses serious safety risks [29]. To mitigate these risks, commercial chips typically have a hard-coded maximum safe temperature. Thermal sensors monitor the temperature at critical hotspots. If the observed temperature exceeds the maximum limit, the system throttles

This article was presented in the International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS) 2017 and appears as part of the ESWEK-TECS special issue.

This work was supported partially by Semiconductor Research Corporation (SRC) task 2721.001 and National Science Foundation grant CNS-1526562.

Authors' addresses: G. Bhat and U. Y. Ogras, School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, 85287; emails: {gmbhat, umit}@asu.edu; S. Gumussoy, Boston, MA, USA; email: suat@gumussoy.net.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM 1539-9087/2017/09-ART145 \$15.00

<https://doi.org/10.1145/3126567>

Fig. 1. Illustration of the power consumption temperature trajectory for two power consumption levels using experiments performed on Odroid XU3 board [12].

the computational resources or shuts down the platform depending on the severity of the violation. However, these techniques are triggered only after the damage becomes observable.

There is a well-known positive feedback between the power consumption and the temperature [23, 32]. Power consumption drives the chip temperature up through thermal resistance and capacitance networks. Higher temperature, in turn, leads to an exponential increase in leakage power. This nonlinear dynamics leads to a positive feedback which increases both the temperature and power consumption. *When a stable fixed point exists*, it attracts all the temperature trajectories within the region of convergence to itself. Therefore, the steady increase in power consumption and temperature continues until the stable fixed point is reached. Otherwise, a thermal runaway occurs, as we prove in this paper.

Although the consequences of thermal runaway are detrimental, to the best of our knowledge, there are no models that can analyze the existence and stability of fixed points at runtime. More importantly, even if a stable fixed point exists, it may be well above the maximum safe temperature limit. A static bound on the maximum power consumption can neither avoid thermal violations, nor provide insight on the expected time before a thermal violation occurs. Therefore, it is critical to monitor the stability and safety of power-temperature dynamics at *runtime*, and detect potential violations *before any damage is done*.

To illustrate the problem addressed by this paper, we measured the power-temperature profile of a commercial SoC for a complete heat-up/cool-down cycle, as shown in Figure 1. In particular, the inner trajectory, denoted by black $\square$ markers, starts with a power consumption slightly larger than 0.2 W. When we increase the dynamic activity, the power consumption quickly rises to 0.9 W. As a result, the temperature starts ramping up during this period marked as “1”. Then, we keep the dynamic activity constant, but the temperature continues to increase towards a fixed point (segment 2). The corresponding rise in the power consumption reveals the impact of the leakage power, since the dynamic activity is kept constant. Eventually, the power consumption and temperature converge to (1.2 W, 69°C)¹. The existence of this fixed point (i.e., the upper right corner) is necessary to avoid a thermal runaway, but it is *not sufficient* to ensure a thermally safe operation. For example, when we repeat the same experiment with a higher dynamic activity, we observe

¹After the dynamic activity is reduced, first, the power consumption drops. Then, the temperature starts decreasing (segments 3 and 4).

the other trajectory denoted by Δ markers. The second trajectory also converges to a stable fixed point given by (2.6 W, 93°C). However, this point is larger than the temperature which triggers throttling. If a demanding application or a power virus drive the system to this fixed point, throttling can deteriorate the performance. In contrast, thermally safe operation without performance penalties can be achieved, if we can compute the fixed point and the expected time to reach it as a function of the dynamic activity.

The major contributions of this paper are as follows:

- (1) We first show that the power-temperature dynamics have either no fixed point or two fixed points, as a function of the system parameters and the dynamic power consumption (Section 4.1).
- (2) We prove that the no-fixed-point case is unstable and causes thermal runaway (Section 4.2).
- (3) When there are two fixed points, we prove that one of these fixed points is stable and we give the region of convergence, i.e., the temperature interval where any temperature inside it converges to the stable fixed point. We also prove that the second fixed point is unstable. We derive the region of convergence and the intervals for which the temperature diverges (Section 4.2).
- (4) We derive an analytical formula to compute the maximum dynamic power consumption allowed to guarantee that the temperature does not exceed a thermally safe value (Section 4.3).
- (5) To validate the proposed approach, we present thorough experimental evaluations on an 8-core big.LITTLE platform [12] using single-threaded, multi-threaded and concurrent applications. We demonstrate that the average and maximum prediction errors are 2.6% and 6.2%, respectively. We also show that the total computational overhead of all proposed computations is 75.2 μ s of 100 ms control interval, i.e., $\approx 0.075\%$ (Section 5).

Potential Impact: This paper lays the theoretical foundation for power-temperature stability analysis and presents experimental validation of the contributions summarized in the enumerated list above. As we demonstrate in Section 5, our power-temperature stability analysis has a very efficient and practical implementation despite the complexity of the derivations. Therefore, it can be employed by other researchers in dynamic thermal and power management (DTPM) algorithms to determine if the power-temperature dynamics is stable or not. If any instability is detected, immediate corrective actions can be taken by the system. Otherwise, the proposed approach can be used to predict the stable fixed point and the expected time to reach it. DTPM algorithms can use this prediction to determine the urgency and degree of the response. Finally, the proposed approach can accurately compute the maximum power consumption that can be tolerated before violating the thermal constraints. This insight can be used by DTPM algorithms to make informed decisions. For example, if a power-hungry application is driving the system beyond a safe temperature, the DTPM algorithm can selectively isolate the application or terminate it. The proposed approach can also be applied as an effective built-in test to determine reliability and thermal violation risks.

The rest of this paper is organized as follows. We present the related work in Section 2. We give an overview of the proposed methodology and detail the theoretical derivations in Section 3 and Section 4, respectively. Finally, we present the experimental results in Section 5, and summarize our conclusions in Section 6.

2 RELATED WORK AND NOVELTY

Thermal modeling and analysis have recently attracted significant attention due to large power densities and the impact of temperature on reliability [22, 32]. These studies can be broadly

classified as design time and runtime approaches. Design time approaches primarily focus on a full-chip thermal analysis such that parameters like thermal design power can be determined [16, 17, 34, 35]. For instance, Hotspot [17] models the thermal behavior of the entire chip as a function of the floorplan, technological parameters and packaging. Then, power consumption traces obtained using common benchmarks are used to simulate the thermal behavior. Similarly, authors in [34] propose a tool which does the full-chip thermal analysis during the synthesis of a chip. These models are very useful for early design stages, however, the high fidelity of these approaches comes at the expense of computational complexity. It is possible to extract high-level models from these tools, and evaluate them iteratively to analyze the thermal behavior. However, iterative approaches are time-consuming and not accurate as we demonstrate in Section 5.7.

Virtually all commercial products have a mechanism to throttle performance, or shut down the whole system in case of thermal violations. However, reactive approaches penalize performance, and respond only after the fact [18, 28]. This led to predictive approaches for dynamic thermal and power management. Predictive approaches first develop computationally efficient thermal models which can be used at runtime [2, 11]. These models are used to predict the temperature as a function of the power consumption to guide the DTPM algorithms [6, 21, 27, 31]. For instance, the authors in [21] propose a hybrid thermal management algorithm which uses hardware and software techniques for temperature control. In particular, they employ clock gating and thermal-aware scheduling to improve the performance of the algorithm. Similarly, the work presented in [33] characterizes the thermal system parameters offline, considering the coupling between various components. Then, this model is used at runtime to predict the temperature, and control the frequency to minimize thermal violations. The authors in [4, 25] propose methods that consider the transient thermal effects and use them for thermal management [25]. While these models work for short prediction intervals, the error increases considerably when larger prediction windows are used [31]. Furthermore, they do not analyze the existence and stability of thermal fixed points. In contrast, our approach can accurately estimate the fixed point and maximum allowed power consumption at runtime with a low computational overhead. Therefore, it can be utilized by DTPM algorithms.

Recent studies have also proposed techniques to calculate a thermally sustainable power budget [5, 26], and maintain it at runtime [9]. In particular, the authors in [26] propose a method to calculate a thermally safe power such that thermal constraints of the system are not violated. However, it does not consider the positive feedback between leakage power and temperature. The work in [5] proposes a framework called TSocket which evaluates the sustainable power budget for different threading strategies in a multiprocessor system. These studies do not address the problem of calculating the thermal fixed point and the conditions on existence of a fixed point. They also employ mainly simulation tools such as HotSpot. In contrast, our approach of finding the maximum safe power is implemented and validated on a real hardware platform.

A number of studies analyze the positive feedback effect between power consumption and temperature [15, 23, 32]. In particular, the authors of [23] show that a thermal runaway is implied when the second order derivative of temperature with respect to time is positive. As pointed out by the authors, this criterion can be successfully applied during design time analysis and simulation. However, it cannot be used as a preventive measure at runtime, since it is satisfied only after the thermal runaway kicks off. Similarly, the technique presented in [32] uses a simple junction-to-ambient heat removal model to predict a thermal runaway during burn-in reliability screening before shipping the chip. The authors in [15] use the temperature dependence of leakage to increase thermally sustainable power dissipation through activity migration. However, neither one of these techniques addresses the problem of calculating the fixed point when there is no thermal runaway. Our work addresses this problem by first deriving the conditions for the existence of

Table 1. Summary of Major Parameters

Symbol	Description
N, M	The number of thermal hotspots and processing elements (resources) in the SoC, respectively.
$T[k]$	$N \times 1$ array where $T_i[k]$, $1 \leq j \leq N$ denotes the temperature of the i^{th} hotspot.
T	The maximum (scalar) steady state temperature over all thermal hotspots.
T^*, P_C^*	The maximum thermally safe temperature and power, respectively.
$P[k]$	$M \times 1$ array where $P_i[k]$, $1 \leq i \leq M$ denotes the power consumption of the i^{th} resource.
$\kappa_{1,i}, \kappa_{2,i}$	Technology dependent parameters of the leakage power for the i^{th} resource.
a, b	Parameters of the single input single output model which describes the thermal dynamics.
$\tilde{T}, \alpha, \beta$	Auxiliary parameters introduced in Equation (10).

a fixed point. When the fixed point exists, we provide the region of convergence for the power-temperature dynamics. Then, we predict the stable fixed point of the system. Hence, it can be used to guard against power attacks that aim to induce damage by elevating the temperature [7, 13, 20].

3 PRELIMINARIES AND OVERVIEW

This section first presents the power consumption and temperature models required for the proposed analysis. Readers familiar with these models can jump to Section 3.2, where we summarize the challenges and give an overview of the proposed approach.

3.1 Power and Temperature Models

Suppose that there are M processors in the target system, as summarized in Table 1. We can express the power consumption of processor i as the sum of dynamic and leakage power consumption:

$$P_i = C_{sw,i} V_i^2 f_i + V_i I_{leak,i} \quad (1)$$

where $C_{sw,i}$ is the switching capacitance, V_i is the supply voltage and f_i is the operating frequency. The leakage current $I_{leak,i}$, which depends on the temperature T , can be approximated as the sum of the gate leakage and subthreshold current as:

$$I_{leak,i} = I_{g,i} + A_s \frac{W_i}{L} \frac{kT^2}{q} e^{\frac{q(V_{GS,i} - V_{th,i})}{nkT}} \quad (2)$$

where $I_{g,i}$ is the gate leakage, A_s is a technology constant, W_i/L is the ratio of the effective channel width to channel length, k is the Boltzmann constant, q is the electron charge, $V_{GS,i}$ is the gate to source voltage, $V_{th,i}$ is the threshold voltage, and n is the sub-threshold swing coefficient [19]. For notational convenience, we consolidate the technology and design parameters by introducing the following constants:

$$\kappa_{1,i} = A_s \frac{W_i}{L} \frac{k}{q}, \quad \kappa_{2,i} = \frac{q(V_{GS,i} - V_{th,i})}{nk} \quad (3)$$

Note that $\kappa_{1,i} > 0$, while $\kappa_{2,i} < 0$ since $V_{GS,i} < V_{th,i}$ for sub-threshold voltages. In summary, the power consumptions of the processors in the target system can be denoted by the $M \times 1$ vector

Fig. 2. Overview of the proposed approach.

$\mathbf{P} = [P_1, P_2, \dots, P_M]^T$, where P_i is obtained using Equations (1)–(3) as:

$$P_i = C_{sw,i} V_i^2 f_i + V_i I_{g,i} + V_i \kappa_{1,i} T_i^2 e^{\frac{\kappa_{2,i}}{T_i}}, \quad 1 \leq i \leq M \quad (4)$$

Suppose that there are N thermal hotspots of interest. The dynamics of the temperature can be expressed using the power consumption vector $\mathbf{P}$ and thermal capacitance and conductance matrices [17, 30]. We employ the following discrete-time state-space system to model the temperature, since the measurements and control decisions are made in periodic intervals:

$$\mathbf{T}[k+1] = \mathbf{A}\mathbf{T}[k] + \mathbf{B}\mathbf{P}[k] \quad (5)$$

This equation expresses the temperature of N cores as a function of the power consumption of M sources, e.g., the big/little CPU clusters, GPU and memory. The $N \times N$ matrix $\mathbf{A}$ describes the effect of temperature in time step k on the temperature in the next time step. The $N \times M$ matrix $\mathbf{B}$, on the other hand, describes the effect of each power source on the temperature in the next time step. When we plug the $M \times 1$ power consumption vector $\mathbf{P}$ from Equation (4) into Equation (5), we obtain the following system of nonlinear equations:

$$\mathbf{T}[k+1] = \mathbf{A}\mathbf{T}[k] + \mathbf{B}[P_1[k], P_2[k], \dots, P_M[k]]^T, \text{ where} \quad (6)$$

$$P_i[k] = C_{sw,i} V_i^2 f_i + I_{g,i} V_i + V_i \kappa_{1,i} T_i[k]^2 e^{\frac{\kappa_{2,i}}{T_i[k]}}, \quad 1 \leq i \leq M$$

3.2 Challenges and Needs Assessment

The nonlinear system given in Equation (6) shows the positive feedback between the power consumption of M processors and N temperature hotspots. Solving this problem at runtime presents *three major challenges*. First, a stable solution may not even exist due to the nonlinear positive feedback. Second, the power consumption and temperature at time k depend on their values at time $k-1$ due to the temperature dependence of the leakage power. This dependency requires an iterative solution, which is intractable for runtime analysis, as shown in Section 5.7. Finally, we need to find the maximum power consumption P_i^* that *guarantees* a thermally safe operation. A simple iteration can find temperature using the power inputs, but the opposite direction requires a rigorous approach.

The proposed approach addresses each of these challenges one by one as outlined in Figure 2. We first determine whether a stable fixed point exists using the current power and temperature

measurements. Then, we describe how the stable fixed points can be computed efficiently. Finally, we derive compact analytical formulae to compute the time to reach the fixed point and the constraints on the maximum power consumption P_i^* to avoid thermal violations. The proposed approach is called with the default frequency governors, which typically have a period of 100 ms.

4 THERMAL FIXED POINT ANALYSIS

Let the scalar T denote the maximum steady state temperature over all thermal hotspots, i.e.:

$$T = \max_{1 \leq i \leq N} \lim_{k \rightarrow \infty} T_i[k]$$

Since the thermal safety is determined by the maximum temperature, we focus on the hotspot with the highest temperature. At steady state ($k \rightarrow \infty$), we can model the temperature of *each hotspot* using the following single input single output (SISO) system:

$$T = aT + bP \quad (7)$$

where $0 < a < 1$ and $b > 0$ are the parameters of the reduced order system. We emphasize that the reduced order model is obtained through system identification, as described in Section 5.2. Using a SISO model does not mean that we consider only one core. Unlike a crude approximation that directly uses the corresponding entries in A and B matrices, the coefficient a in our model reflects the thermal coupling between different hotspots, and coefficient b reflects the impact of multiple power sources. We employ a SISO model, since it enables an in-depth theoretical analysis with powerful insights for practical scenarios. We discuss the solution to the multi-input multi-output (MIMO) case at the end of this section.

To isolate the impact of the temperature, we rewrite the total power consumption given in Equation (4) as:

$$P = P_C + V\kappa_1 T^2 e^{\frac{\kappa_2}{T}} \quad (8)$$

where $P_C = C_{sw}V^2f + VI_g$ represents the *temperature-independent* component, and subscript i is dropped to simplify the notation. Substituting Equation (8) into Equation (7) gives:

$$(1 - a)T - bP_C = bV\kappa_1 T^2 e^{\frac{\kappa_2}{T}} \quad (9)$$

If this equation has feasible solution(s), we can say that fixed points exist. Since Equation (9) is the *focal point* of the subsequent analysis, we introduce the following change of variables to leave the exponential term alone and facilitate the subsequent analysis:

$$\tilde{T} = -\frac{\kappa_2}{T}, \quad \alpha = \frac{b}{a-1} \frac{P_C}{\kappa_2} > 0, \quad \beta = \frac{a-1}{b} \frac{1}{V\kappa_1\kappa_2} > 0 \quad (10)$$

With this change of variables, we rewrite Equation (9) as:

$$\beta\tilde{T}(1 - \alpha\tilde{T}) = e^{-\tilde{T}} \quad (11)$$

where $\alpha > 0, \beta > 0$. We will first derive the conditions on α and β such that Equation (11) has a solution. Then, we will go back from the transformed domain to the original parameters. Finally, we will show how to compute the constraint on the maximum power consumption P_C^* required to avoid thermal violations, given a maximum temperature constraint T^* .

4.1 Necessary and Sufficient Conditions for the Existence of Fixed Point(s)

The domain of the auxiliary temperature is given by $\tilde{T} \in (0, \infty)$, since $\tilde{T} = -\kappa_2/T$ where $\kappa_2 < 0$. Hence, the right-hand side of Equation (11) lies in the interval $(0, 1)$. That is,

Fig. 3. Illustration of $\mathcal{F}(\tilde{T})$ when there is no fixed point and when there are two fixed points, respectively.

$0 < \beta\tilde{T}(1 - \alpha\tilde{T}) = e^{-\tilde{T}} < 1$. Since this condition requires that $\tilde{T} < 1/\alpha$, we can take the logarithm of both sides while the equality holds, i.e.,

$$\ln \beta + \ln \tilde{T} + \ln(1 - \alpha\tilde{T}) = -\tilde{T}$$

Equation (11) has the same fixed points as the following equation:

$$\mathcal{F}(\tilde{T}) \triangleq \ln \beta + \ln \tilde{T} + \ln(1 - \alpha\tilde{T}) + \tilde{T} = 0 \quad (12)$$

The important properties of $\mathcal{F}(\tilde{T})$ employed in our analysis are summarized in the following lemma.

LEMMA 4.1. $\mathcal{F}(\tilde{T})$ given in Equation (12) satisfies the following properties:

- (1) $\mathcal{F}(\tilde{T})$ is a concave function in the interval $\tilde{T} \in (0, 1/\alpha)$.
- (2) $\mathcal{F}(\tilde{T})$ has a unique maxima at $\tilde{T}_m$, which is given by:

$$\tilde{T}_m = \frac{1}{2\alpha} - 1 + \sqrt{\frac{1}{4\alpha^2} + 1} \quad (13)$$

- (3) $\mathcal{F}(\tilde{T})$ is an increasing function in the interval $(0, \tilde{T}_m)$ and a decreasing function in $(\tilde{T}_m, 1/\alpha)$.

PROOF. The proof is provided in Appendix A.1, while an informal explanation is provided below to maintain a smooth flow. $\square$

Figure 3 sketches $\mathcal{F}(\tilde{T})$ for $\tilde{T} \in (0, 1/\alpha)$. As the first two properties of Lemma 4.1 state, $\mathcal{F}(\tilde{T})$ is concave and has a unique maxima at $\tilde{T} = \tilde{T}_m$. Furthermore, it is increasing in the interval $(0, \tilde{T}_m)$ and decreasing in $(\tilde{T}_m, 1/\alpha)$, as shown in Figure 3. Hence, Equation (12) ($\mathcal{F}(\tilde{T}) = 0$) has two solutions if and only if the maxima is non-negative, i.e., $\mathcal{F}(\tilde{T}_m) \geq 0$. The fixed points coincide when $\mathcal{F}(\tilde{T}_m) = 0$, and there are no fixed points for $\mathcal{F}(\tilde{T}_m) < 0$. To be practical, these conditions should be expressed in terms of the design parameters α and β . This is achieved by Theorem 4.1, which summarizes our first major result.

THEOREM 4.1. The maxima of $\mathcal{F}(\tilde{T})$ is given by:

$$\mathcal{F}(\tilde{T}_m) = \ln \beta - \ln \left(\frac{2}{\tilde{T}_m} + 1 \right) e^{-\tilde{T}_m} \quad (14)$$

Hence, Equation (11) has two fixed points if and only if $\beta \geq (\frac{2}{\tilde{T}_m} + 1)e^{-\tilde{T}_m}$ where $\tilde{T}_m$ depends only the parameter α and it is defined in Equation (13). Otherwise, it has no solution.

PROOF. The proof is provided in Appendix A.2. $\square$

At runtime, we first compute $\tilde{T}_m$ using Equation (13). Then, we check the condition given in Theorem 4.1. If it is not satisfied, we conclude that there will be a thermal runaway. This knowledge can be used to throttle the cores aggressively or enter an emergency state. Otherwise, we proceed to compute the maximum allowed power consumption that will avoid thermal violations.

4.2 Stability of the Fixed Points

The stability of the fixed points is determined by the behavior of $\tilde{T}$ as function of $\mathcal{F}(\tilde{T})$. To provide a smooth flow, we summarize this behavior using the following lemma.

LEMMA 4.2. *The value of $\tilde{T}$ in the temperature iteration increases when $\mathcal{F}(\tilde{T}) < 0$, and decreases when $\mathcal{F}(\tilde{T}) > 0$.*

PROOF. The proof is provided in Appendix A.3. □

This lemma allows us to determine the stability characteristics of fixed points by inspecting the sign of the function $\mathcal{F}(\tilde{T})$. The following theorem summarizes the stability results using this lemma, which is also illustrated with the arrows in Figure 3(a).

THEOREM 4.2. *The stability of the fixed points is as follows.*

- (1) *When Equation (12) has no solution, the temperature iteration diverges, i.e., $\tilde{T} \rightarrow 0$ ($T \rightarrow \infty$), as illustrated in Figure 3. Hence, there is a thermal runaway.*
- (2) *When there are two fixed points, $\tilde{T}_u \in (0, \tilde{T}_m)$ is unstable and $\tilde{T}_s \in (\tilde{T}_m, \frac{1}{\alpha})$ is stable.*
- (3) *In the latter case, any temperature iteration starting in $(0, \tilde{T}_u)$ diverges, i.e., $\tilde{T} \rightarrow 0$ and $T \rightarrow \infty$ leading to a thermal runaway. However, any temperature iteration starting in $(\tilde{T}_u, \frac{1}{\alpha})$ converges to the stable fixed point $\tilde{T}_s$.*

PROOF. The proof is provided in Appendix A.4, while an informal explanation is provided below. □

Theorem 4.2 states that the temperature proceeds along the arrows shown in Figure 3 during fixed point iterations. When $\mathcal{F}(\tilde{T}_m) < 0$, i.e., no fixed point exists, $\tilde{T}$ decreases at each temperature iteration no matter where it starts. Therefore, there is a thermal runaway as illustrated in Figure 3(a). When $\mathcal{F}(\tilde{T}_m) > 0$, there are two fixed points denoted by $\tilde{T}_u$ and $\tilde{T}_s$. Any iteration starting in the interval $(0, \tilde{T}_u)$ diverges to $\tilde{T} \rightarrow 0$, since $\mathcal{F}(\tilde{T}) < 0$ in that interval. Conversely, any iteration starting in the interval $(\tilde{T}_u, \tilde{T}_s)$ will converge to $\tilde{T}_s$, since $\mathcal{F}(\tilde{T}) > 0$. That is, any iteration starting in the interval $(\tilde{T}_u, \frac{1}{\alpha})$ will also converge to $\tilde{T}_s$, as illustrated in Figure 3(b). Therefore, we conclude that $\tilde{T}_u \in (0, \tilde{T}_m)$ is unstable, while $\tilde{T}_s \in (\tilde{T}_m, \frac{1}{\alpha})$ is stable.

In summary, we derived the conditions for the existence of fixed points and their stability regions in terms of the auxiliary temperature $\tilde{T}$, α and β . Next, we will derive the constraints on the dynamic power consumption required for thermally safe operation.

4.3 From Temperature to Power Constraint

Suppose that the thermally safe temperature is given by T^* . We need to work backwards starting with this constraint to find the maximum allowable P_C^* (i.e., the sum of dynamic and gate leakage power consumption). To this end, we first show the existence of P_C^* , and then, we provide a *compact formula* to compute P_C^* in terms of T^* and the system parameters.

Existence of P_C^* : We know that the system converges to a stable fixed point when there is no dynamic activity ($P_C \rightarrow 0$). This means that the necessary and sufficient condition given in Theorem 4.1 is satisfied, i.e., $\beta \geq (\frac{2}{\tilde{T}_m} + 1)e^{-\tilde{T}_m}$. Now, consider the other extreme where P_C grows indefinitely due to heavy dynamic activity. Equation (10) shows that α will also grow, since it increases

linearly with P_C . Growing α , in turn, implies that $\tilde{T}_m \rightarrow 0$ according to Equation (13). Hence, as the dynamic activity (P_C) increases, the term $(\frac{2}{\tilde{T}_m} + 1)e^{-\tilde{T}_m}$ increases monotonically. This suggests that there exists a $\tilde{T}_m$ where $\beta = (\frac{2}{\tilde{T}_m} + 1)e^{-\tilde{T}_m}$ holds. At the same time, this corresponds to the maximum power $P(\tilde{T}_m)$ with two fixed points. Furthermore, Equation (14) shows that $\mathcal{F}(\tilde{T}_m)$ decreases monotonically with decreasing $\tilde{T}_m$, as illustrated by Figure 3(a). Therefore, $\mathcal{F}(\tilde{T})$ decreases monotonically as P_C increases, and crosses the x -axis zero at P_C^* . Due to the monotonic behavior of stable fixed points as a function of P_C , we conclude that there exists a maximum allowable P_C^* whose thermal fixed point does not exceed T^* .

P_C^* for a given T^* is computed as follows at runtime:

- (1) Compute the auxiliary temperature that corresponds to T^* using Equation (10) as $\tilde{T}^* = -\kappa_2/T^*$
- (2) Given $\tilde{T}^*$, find the corresponding α^* using Equation (11): $\alpha^* = \frac{1}{\tilde{T}^*}(1 - \frac{e^{-\tilde{T}^*}}{\beta\tilde{T}^*})$
- (3) Finally, substitute α^* into Equation (10) to find P_C^* : $P_C^* = \frac{a-1}{b}\alpha^*\kappa_2$

In summary, we conclude that any power value such that $P_C < P_C^*$ has a thermal stable fixed point less than T^* .

4.4 Time to Reach the Stable Fixed Point

The existence and specific value of the fixed point can enable DTPM algorithms to determine whether the power/temperature trajectory moves towards a dangerous operating point. In addition to this, the expected time to reach to that point reveals how soon the dangerous zone will be reached. Therefore, DTPM algorithms can utilize the timing information to determine if there is any imminent possibility of a thermal violation. Moreover, this estimate can also be used to decide how long the current power consumption can be sustained without violating the thermal limit. To obtain a computationally efficient estimation, we employ the following exponential model:

$$T[kT_s] = T_{init} + (T_{fix} - T_{init})(1 - e^{-\frac{kT_s}{\tau}}) \quad (15)$$

where T_{init} is the initial temperature, T_{fix} is the stable fixed point and τ is the time constant. We use the discrete time stamp kT_s , since the temperature is sampled with a period of T_s . Two temperature readings separated by a delay DT_s , i.e., $T[kT_s]$ and $T[(k-D)T_s]$, can be used to estimate the time constant τ as:

$$\tau = \frac{t[kT_s] - t[(k-D)T_s]}{\ln\left(\frac{T[(k-D)T_s] - T_{fix}}{T[kT_s] - T_{fix}}\right)} \quad (16)$$

In our experiments, $D = 10$ and the time to fixed point computation, like the other computations in the proposed technique, are repeated with $T_s = 100$ ms.

4.5 Solution for the MIMO Case

The SISO model given in Equation (7) can be extended to explicitly show the impact of each power source on the thermal hotspot of interest T_i , as follows:

$$\begin{aligned} T_i &= [a_{i1}, a_{i2}, \dots, a_{iN}]\mathbf{T}[\mathbf{k}] + [b_{i1}, b_{i2}, \dots, b_{iM}]\mathbf{P}[\mathbf{k}] \\ &= [a_{i1}, \dots, a_{iN}][T_1, \dots, T_N]^T + [b_{i1}, \dots, b_{iM}][P_1, \dots, P_M]^T \end{aligned}$$

where P_i terms are defined in Equation (4). In order to find the fixed point temperatures, we need to solve the set of equations given by:

$$\mathbf{f}(\mathbf{T}) = \mathbf{f}(T_1, \dots, T_N) = \begin{bmatrix} f_1(T_1, \dots, T_N) \\ \vdots \\ f_N(T_1, \dots, T_N) \end{bmatrix} = 0 \quad (17)$$

where $f_i(T_1, \dots, T_N) = [a_{i1}, \dots, a_{iN}][T_1, \dots, T_N]^T + [b_{i1}, \dots, b_{iM}][P_1, \dots, P_M]^T - T_i, \forall 1 \leq i \leq N$.

This is a nonlinear equation due to the exponential terms in P_i . Generalizing our solution to obtain similar closed-form formulae for the MIMO model requires solving the nonlinear equations in Equation (17). The standard approach to this problem is to find a good initial point, and use a root search algorithm for nonlinear equations. By following this approach, we use our SISO solution for each hotspot as the “*initial point*”. Then, we employ an iterative search to find the roots of nonlinear equations. More specifically, we solve Equation (17) via Newton’s method [1] where our SISO solution is used as the initial point. Our experimental measurements show that the closed loop solutions presented in Sections 4.1–4.4 are very close to the roots of the MIMO system. Hence, iterations converged to the roots of the MIMO system in less than five iterations in the worst case. Moreover, our simulations show that the region of convergence is quite large. This means that even if our initial points are not accurate, we will still converge to the solution.

5 EXPERIMENTAL EVALUATION

5.1 Experimental Setup

The proposed fixed point prediction scheme is evaluated on the Odroid XU3 board which employs Samsung Exynos 5422 SoC [12]. Exynos 5422 is a single ISA big.LITTLE SoC consisting of 4 Cortex-A15 (big) cores, 4 Cortex-A7 (little) cores and a Mali GPU. The Odroid XU3 board also comes with thermal sensors to measure the temperature of each big core and the GPU. It also includes current sensors to measure the power consumption of the big cluster, little cluster, the memory and the GPU. The proposed approach is invoked with the default frequency governors every 100 ms. That is, we periodically sample four power consumption and five temperature values. The proposed technique takes $75.2 \mu\text{s}$ of this 100 ms period, as detailed in Section 5.7. We evaluated the proposed approach on commonly used benchmarks, including four from the MiBench suite [10], three from the PARSEC suite [3], one SPEC [14], and a matrix multiplication kernel. These benchmarks are listed in the first column of Table 2.

5.2 Parameter Characterization

The stability and safety analysis presented in Section 4 does not make any assumptions about specific parameter values. However, we need to characterize the system parameters, such as the leakage current and the state-space model parameters, in order to evaluate the accuracy of our analysis. As the first step, we characterized the leakage current parameters I_g, κ_1, κ_2 for the Odroid XU3 board using a methodology similar to those presented in [24, 31]. More precisely, we placed the target device in a furnace and collected data at five different temperatures ranging from 40°C to 80°C , while running a light workload to make sure that the temperature is not increased due to the dynamic activity. Since the dynamic power consumption remained constant, the measured power consumption difference between two experiments gives the difference in leakage power at two known temperatures. We used measurements at five different temperatures to find the unknowns $I_{g,i}, \kappa_{1,i}, \kappa_{2,i}$ for the big core cluster and the GPU as these are the dominant sources of leakage in our board.

Table 2. Summary of Results for Fixed Point Prediction. The Foreground Application is Listed in the First Column. Since these Experiments were Performed Under Android OS, There were More than 100 Background Applications Running at all Times

Benchmark	Avg. Total Power (W)	Avg. Dyn. Power (W)	Empirical Fixed Point Temp. (°C)	Comput. Fixed Point Temp. (°C)	Abs. Pred. Error (°C)	Perc. Pred. Error (%)	Runtime (s)
Idle @ 1.3 GHz	0.38	0.31	51.8	52.2	0.4	0.8	3979
Idle @ 1.5 GHz	0.47	0.38	54.4	55.5	1.1	2.0	4045
Idle @ 1.8 GHz	0.70	0.59	60.2	60.1	0.1	0.2	4171
Idle @ 2.0 GHz	1.01	0.87	66.0	66.6	0.6	0.9	3413
Vortex	1.73	1.55	80.0	81.4	1.4	1.7	1989
Matrix Mult.	1.84	1.65	83.0	83.8	0.8	1.0	521
CRC32	2.04	1.83	85.0	88.5	3.5	4.1	907
Patricia	2.20	1.97	89.0	91.8	2.8	3.1	900
Blackscholes	2.42	2.17	94.0	96.6	2.6	2.7	785
Streamcluster	2.48	2.22	94.0	97.9	3.9	4.1	614
Basicmath	2.49	2.24	93.0	98.3	5.3	5.7	492
Fluidanimate	2.50	2.26	93.0	98.8	5.8	6.2	550
FFT	2.71	2.41	101.0	102.5	1.5	1.5	729
Streamcluster + CRC32	3.14	2.79	109.0	111.5	2.5	2.3	1132

Fig. 4. Comparison of predicted temperature with the measured temperature.

The second step is estimating the A and B matrices used in Equation (5). In order to characterize the A matrix, we ran the big cluster at a fixed frequency level until the temperature rises to a steady state value. Once the steady state was reached, we turned off the big core and let the device run until the board cools down. We repeated this for multiple frequency levels as shown in Figure 4. We also repeated the complete experiment multiple times and then averaged the results to account for any variations in the system behavior over time. The data from this experiment helps us determine how the temperature at the current time step affects the temperature in the next time step. In a typical run, different workloads lead to varying utilization of system resources, which lead to different operating frequencies. To determine the effect of each power source on the temperature, we excited all the frequency levels of the A7 cores, A15 cores and the GPU in a pseudo-random bit sequence (PRBS), and recorded the temperature. This helps us to identify the effect of each power source on the temperature under varying conditions, i.e., characterizing the B matrix. This data is then used in conjunction with the data from the previous experiment to jointly identify A and B . The A and B matrices are used to predict the temperature given an initial temperature and power consumption at each time step. As we can see in Figure 4, the computed temperature closely follows the measured temperature. In this work, we performed characterization on one Samsung Exynos 5422 chip. The parameters obtained for leakage current and the thermal model are used to

evaluate the accuracy of the fixed point analysis on three different instances. In general, the chip manufacturers divide their products into multiple bins, such as low power and high performance parts, and sell them using different stock keeping units (SKUs). Therefore, our analysis can be performed for each SKU of a given commercial chip.

5.3 Validation of Fixed Point Approximation

Evaluation at fixed frequency: To evaluate the accuracy of the proposed fixed point analysis, we first performed experiments at various power levels ranging from 0.38 W to 1.01 W, while running a light load on the CPU. After running the system until a steady-state fixed point is attained, we recorded the average power consumption and final temperature. We also estimated the average dynamic power consumption for each experiment by subtracting the estimated leakage power from the measured total power. The first three columns in Table 2 provide these values. Then, we used the estimated dynamic power consumption and the ambient temperature to analyze whether a stable fixed point exists or not. After confirming that the analysis result is correct, we computed the fixed point temperature using the proposed technique. The fourth column of Table 2, labeled as “Comput. Fixed Point Temp.,” lists the results for each power consumption level. As summarized in Table 2, the fixed point prediction is within 1°C of the empirical result for four power levels. The largest observed difference was 1.1°C, which implies 2.0% error with respect to the measured fixed point. This error is quite acceptable for our system, as the temperature sensors operate at an integer precision, which can introduce an error in the measurement.

Evaluation on benchmarks: We also evaluated our analysis technique on commonly used benchmarks which represent real-world applications. We ran these applications for minutes to capture the temperature dynamics, as summarized in the last column of Table 2. During these experiments, we overwrote the safe temperature limit on the board such that the temperature could rise beyond 110°C. The results of the fixed point evaluation for the benchmarks are summarized in the lower part of Table 2. We observe that the fixed point prediction error increases slightly compared to the fixed frequency experiments. This is expected, since there are larger variations in the power consumption. However, the average prediction error is still only 3.0°C, and in the worst case, the error is 5.8°C. Finally, the average prediction error across all experiments is 2.3°C.

5.4 Power-Temperature Trajectory

We also performed experiments to compare the measured power-temperature trajectory against the fixed point predicted by our analysis. Trajectories for different initial power consumption and temperature are shown in Figure 5(a). The initial points are denoted by black $\circ$ markers and simulated trajectories are plotted using dashed black lines. For each of the initial points, the dynamic power starts with a small value and then rises to a steady value depending on the workload. For example, consider the trajectory that starts from the initial point (0.5 W, 50°C). The dynamic power increases until about 0.87 W, and then it remains steady at 0.87 W. The arrows show how each trajectory converges to the stable fixed point at (1.02 W, 66.6°C) shown by a larger black $\bigcirc$. These trajectories are obtained by solving the set of nonlinear equations given in Equation (6) iteratively. Given the same assumptions, the proposed approach finds the stable fixed point as (1.02 W, 66.6°C), as denoted by the red $\blacktriangle$ marker on the same figure. Hence, the prediction exactly matches the simulated trajectory. Furthermore, we performed one more set of experiments on the board by imposing a dynamic power consumption close to the value used in the simulation. The trajectory measured during this experiment is plotted using solid blue line in Figure 5(a). This trajectory almost overlaps with the simulation and converges to (1.01 W, 66.0°C). As a result, the difference between the empirical and theoretical fixed points is only 0.6°C, which agrees with the results in Table 2.

(a) Analytical, experimental and simulation results when the fixed point is *lower* than the thermally safe temperature.

(b) Analytical, experimental and simulation results when the fixed point is *higher* than the thermally safe temperature.

Fig. 5. Power-temperature trajectories with multiple initial power and temperature values. Black $\circ$ markers are the initial points and the dotted black line shows the trajectory followed by power and temperature. Black $\bigcirc$ and red $\blacktriangle$ markers show the simulated and *computed* fixed point respectively. We observe that the simulation converges to the computed fixed point for each initial point. The trajectory of the real experiment (shown using blue lines) follows the simulated trajectory closely.

In order to illustrate the case where the temperature converges to a value beyond the thermal limit of our board, we performed one more set of experiment and simulation, as shown in Figure 5(b). We disabled the thermal throttling to let the temperature rise beyond 110°C . Similar to Figure 5(a), the dynamic power starts with a small value until it increases to 4.23 W. The dashed black lines show the simulated trajectory followed by the power and temperature for each initial point. We note that the simulation for each initial point converges to the fixed point at (5.14 W, 149.5°C) marked with a larger black $\bigcirc$. With these conditions, the proposed approach finds the fixed point as (5.14 W, 149.6°C), as denoted by the red $\blacktriangle$ marker. The difference between the analytical solution and simulated trajectory is 0.1°C , hence, they almost overlap in Figure 5(b). Moreover, we performed one instance of the experiment on the board by choosing the initial point as (0.4 W, 40°C). We also forced the dynamic power consumption close to 4.20 W, i.e., the value used in the simulation. The solid blue line in Figure 5(b) shows that the actual trajectory closely follows the simulated trajectory until the temperature reaches 93°C . At that point, we stopped the experiment to avoid damage to our board.

Power-Temperature with Thermal Throttling: To further evaluate the validity of our fixed point prediction, we performed two sets of experiments: one with and the other without throttling. We used the FFT benchmark from the MiBench suite, as it exhibits a representative behavior and leads to higher temperature fixed point than the other applications. The proposed fixed point calculation is performed every 100 ms in the Linux kernel. This allows us to analyze how the fixed point prediction evolves over time. Red $\blacktriangle$ markers in Figure 6 show that the fixed point predictions vary from about 95.0°C to 105.0°C . This variation is expected, since the dynamic power changes during the execution of the application. However, our analysis results still match closely with the measured trajectory which approaches the predicted fixed point.

We repeated the previous experiment, this time by incorporating a thermal throttling policy. The power consumption starts from a small value of 0.48 W and then increases to about 2.40 W after the benchmark starts execution. When the power consumption is 2.40 W, the fixed point is

Fig. 6. Evaluation of the power-temperature trajectory and fixed point evaluated every 100 ms for the FFT benchmark.

Fig. 7. Using the proposed analysis for thermal throttling, and its impact on power-temperature trajectory.

predicted as 99.0°C. As the power consumption steadily increases due to the increased leakage, the fixed point prediction increases to 102.0°C, as shown by the region A in Figure 7. The solid blue line shows that the measured trajectory advances towards the predicted fixed point, as in the previous experiment. This time, the DTPM policy is triggered to throttle the frequency as soon as the temperature reaches 85.0°C. Throttling starts reducing the power consumption, which in turn slows down the temperature increase. Figure 7 shows that the reduction in the power consumption is reflected in our fixed point prediction. More precisely, the proposed technique updates the fixed point prediction as (2.07 W, 87.7°C). At the same time, the measured power-temperature trajectory changes its course. We observe that it starts converging to (2.06 W, 86.0°C), which matches very well with our prediction. This experiment illustrates that our analysis can adapt to changes in the dynamic power consumption and predict the fixed point accurately.

5.5 Power Constraint from Temperature

We also evaluated the change in the maximum power constraint P_C^* when the temperature constraint T^* is varied. In particular, we swept the value of T^* from 50°C to 105°C and calculated the value of P_C^* . The black line in Figure 8 shows the maximum the power constraint P_C^* found using the analytical approach outlined in Section 4.3. To validate the analysis results, we set the temperature constraint T^* as the theoretical fixed point. Then, we empirically found the power constraint P_C^* on the target board. The red Δ markers in Figure 8 show that the measured results are indeed on the trend found by the proposed technique. This result can be easily used as a guideline to decide the maximum power level that a chip can be operated based on a given temperature constraint.

5.6 Evaluation of the Time to Reach Fixed Point

We used Equation (16) to estimate the time at which the temperature will reach the fixed point for the FFT benchmark. Then, we compared this estimate to the actual time to reach the fixed point. Figure 9 shows that our estimate provides a lower bound for the time to reach the fixed point. A lower bound is useful, since it can be used safely to avoid thermal violations. We also see that the estimation improves in accuracy as the benchmark continues to run. In summary, this estimation can be used by DTPM policies to decide how long the current power consumption level can be sustained without violating the thermal limit.

Fig. 8. Variation of the maximum power constraint P_C^* for different values of T^* .

Fig. 9. Comparison of the actual time taken to reach fixed point and the predicted time to reach fixed point.

5.7 Implementation Overhead

Our theoretical analysis and proofs enable us to derive analytical solutions that can be implemented with negligible overhead. To quantify this overhead, we implemented the proposed solutions on Android 4.4.4 / Linux 3.10.9 kernel user space, and measured the overhead on the Odroid XU3 mobile platform. Our implementations are invoked periodically with the default frequency governors, i.e., every 100 ms. We observed that reading the sensors takes $13.8 \mu\text{s}$, while computing the fixed point estimate for the SISO case takes $6.8 \mu\text{s}$. We can achieve this low overhead since both $1/\alpha$ and $\tilde{T}_m$ have closed form solutions. Once we have the fixed point estimate, the Newton's method to solve the MIMO case takes about $50 \mu\text{s}$. Similarly, it takes $1.1 \mu\text{s}$ to compute the maximum allowable power consumption P_C^* given a temperature constraint. This small overhead is enabled by the three closed form equations framed at the end of Section 5.7. Finally, computing the time to reach the stable fixed point takes $3.5 \mu\text{s}$. The combined overhead of all three computations is about $75.2 \mu\text{s}$ out of 100 ms, i.e., 0.075%. When the implementation is moved to the kernel, the execution time of the functions reduces by about 30%. In contrast, an iterative approach cannot determine the existence and stability of fixed points. Furthermore, it can predict the temperature given the power consumption, but it cannot compute the maximum allowable power consumption P_C^* . Finally, temperature prediction over an interval of 1000 s alone takes about $550 \mu\text{s}$. The iterative approach also has an average error of more than 10°C , which is higher than that of our approach.

6 CONCLUSION

This paper presents a theoretical analysis of the stability of the power consumption and temperature dynamics. First, we show that the power-temperature dynamics have either no fixed point or two fixed points, as a function of the system parameters and the dynamic power consumption. When there are two fixed points, we prove that one of the fixed points is stable, while the second one is unstable. We also determine the region of convergence, which is important for safe thermal operation. Third, we derive an analytical formula to compute the maximum dynamic power consumption that guarantees a thermally safe operation. Experiments and simulation results show that our analysis can be used to predict the fixed point within 0.1°C to 5.8°C accuracy with only 0.075 ms computational overhead. Hence, the proposed approach can be used to take *proactive* DTPM decisions, and detect security threats which force the system to operate beyond the thermal limit.

APPENDIX

A.1 Proof for Lemma 4.1

Note that $\mathcal{F}(\tilde{T})$ approaches $-\infty$ at both end points. Now, the first and second derivatives of $\mathcal{F}(\tilde{T})$ with respect to $\tilde{T}$ are evaluated as:

$$\mathcal{F}'(\tilde{T}) = \frac{1}{\tilde{T}} - \frac{\alpha}{1 - \alpha\tilde{T}} + 1, \quad \mathcal{F}''(\tilde{T}) = -\frac{1}{\tilde{T}^2} - \frac{\alpha^2}{(1 - \alpha\tilde{T})^2}$$

Since $\mathcal{F}''(\tilde{T}) < 0$ for $\tilde{T} > 0$, this function is concave. By setting $\mathcal{F}'(\tilde{T}) = 0$, we can show that the maxima of $\mathcal{F}(\tilde{T})$ occurs when $\tilde{T}_m = \frac{1}{2\alpha} - 1 \pm \sqrt{\frac{1}{4\alpha^2} + 1}$. Since the temperature is positive, the following relations hold at the maximum point:

$$\tilde{T}_m = \frac{1}{2\alpha} - 1 + \sqrt{\frac{1}{4\alpha^2} + 1}, \text{ i.e., } \alpha = \frac{\tilde{T}_m + 1}{\tilde{T}_m^2 + 2\tilde{T}_m} \quad (18)$$

Moreover, due to the concavity of $\mathcal{F}(\tilde{T})$ is an increasing function on $(0, \tilde{T}_m)$ and decreasing function on $(\tilde{T}_m, \frac{1}{\alpha})$, as depicted in Figure 3.

A.2 Proof for Theorem 4.1

PROOF. Function $\mathcal{F}$ reaches its peak at $\tilde{T} = \tilde{T}_m$, and its solution contains two points if and only if $\mathcal{F}(\tilde{T}_m) \geq 0$ (Figure 3(b)). Otherwise, if $\mathcal{F}(\tilde{T}_m) < 0$, it does not intersect the x-axis and there is no solution (Figure 3(a)). The condition $\mathcal{F}(\tilde{T}_m) \geq 0$ is equivalent to:

$$\begin{aligned} \mathcal{F}(\tilde{T}_m) &= \ln \beta + \tilde{T}_m + \ln(\tilde{T}_m(1 - \alpha\tilde{T}_m)) \\ &= \ln \beta + \tilde{T}_m + \ln\left(\tilde{T}_m\left(1 - \frac{\tilde{T}_m + 1}{\tilde{T}_m^2 + 2\tilde{T}_m}\tilde{T}_m\right)\right) \\ &= \ln \beta + \tilde{T}_m - \ln\left(\frac{2}{\tilde{T}_m} + 1\right) \geq 0. \end{aligned} \quad (19)$$

Hence, $\beta \geq (\frac{2}{\tilde{T}_m} + 1)e^{-\tilde{T}_m}$ follows. $\square$

A.3 Proof for Lemma 4.2

PROOF. The temperature iteration equation is:

$$T[k+1] = aT[k] + b(P_C + V_{\kappa_1}T[k]^2)e^{\frac{\kappa_2}{T[k]}}.$$

We can rewrite this equation by a change of variable, i.e., $T[k] = \frac{-\kappa_2}{\tilde{T}[k]}$, as:

$$\frac{-\kappa_2}{\tilde{T}[k+1]} = -\frac{a\kappa_2}{\tilde{T}[k]} + bP_C + bV_{\kappa_1}\frac{\kappa_2^2}{\tilde{T}[k]^2}e^{-\tilde{T}[k]}.$$

After substituting the definitions of α and β in (10) and rearrangement of terms, we obtain:

$$\begin{aligned} \frac{1}{\tilde{T}[k+1]} &= \frac{a}{\tilde{T}[k]} - \frac{bP_C}{\kappa_2} - \frac{bV_{\kappa_1}\kappa_2}{\tilde{T}[k]^2}e^{-\tilde{T}[k]}, \\ &= \frac{a}{\tilde{T}[k]} - (a-1)\alpha - \frac{(a-1)}{\beta\tilde{T}[k]^2}e^{-\tilde{T}[k]}, \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\tilde{T}[k]} + \frac{a-1}{\tilde{T}[k]} - (a-1)\alpha - \frac{(a-1)}{\beta\tilde{T}[k]^2} e^{-\tilde{T}[k]}, \\
&= \frac{1}{\tilde{T}[k]} - \frac{(1-a)}{\beta\tilde{T}[k]^2} \left(\beta(1-\alpha\tilde{T}[k])\tilde{T}[k] - e^{-\tilde{T}[k]} \right).
\end{aligned}$$

Note that when $\mathcal{F}(\tilde{T}) < 0$, we can show that $\beta(1-\alpha\tilde{T})\tilde{T} - e^{-\tilde{T}} < 0$ using Equation (12). By inspecting the last equation, this means that each temperature iteration decreases $\tilde{T}[k]$ when $\mathcal{F}(\tilde{T})$ is negative. In contrast, $\mathcal{F}(\tilde{T}) > 0$ implies that $\beta(1-\alpha\tilde{T})\tilde{T} - e^{-\tilde{T}} > 0$. That is, each fixed point iteration increases $\tilde{T}[k]$ when $\mathcal{F}(\tilde{T})$ is positive. $\square$

A.4 Proof for Theorem 4.2

PROOF. When Equation (11) has no solution, Equation (12) has no solution and the sign of the function $\mathcal{F}(\tilde{T})$ is negative since its terms are concave and affine functions. By Lemma 4.2, at each fixed point iteration, the value of $\tilde{T}[k]$ decreases towards to 0. Due to the fact that $T[k] = -\kappa_2/\tilde{T}[k] \rightarrow \infty$, thermal runaway occurs as illustrated with the arrows in Figure 3.

When Equation (11) (equivalently Equation (12)) has a solution, there are two fixed points of the function $\mathcal{F}(\tilde{T})$, i.e., $\tilde{T}_u$ and $\tilde{T}_s$ such that $0 < \tilde{T}_u < \tilde{T}_s < \frac{1}{\alpha}$. Since $\mathcal{F}(\tilde{T})$ is a concave down function, the sign of this function at $\tilde{T} \in (0, \tilde{T}_u)$ is negative and by Lemma 4.2, the value of $\tilde{T}$ decreases to 0 on $\tilde{T} \in (0, \tilde{T}_u)$ just like the no-solution case, and results in thermal runaway. On the other hand, the sign of the function $\mathcal{F}(\tilde{T})$ is positive on $\tilde{T} \in (\tilde{T}_u, \tilde{T}_s)$ and negative on $\tilde{T} \in (\tilde{T}_s, \frac{1}{\alpha})$. By Lemma 4.2, the value of $\tilde{T}$ increases in $\tilde{T} \in (\tilde{T}_u, \tilde{T}_s)$ and decreases in $\tilde{T} \in (\tilde{T}_s, \frac{1}{\alpha})$ towards to $\tilde{T}_s$ on both intervals. Hence, we conclude that $\tilde{T}_s$ is a stable and $\tilde{T}_u$ is unstable fixed point. $\square$

REFERENCES

- [1] Kendall E Atkinson. 2008. *An Introduction to Numerical Analysis*. John Wiley & Sons.
- [2] Francesco Beneventi, Andrea Bartolini, Andrea Tilli, and Luca Benini. 2014. An effective gray-box identification procedure for multicore thermal modeling. *IEEE Trans. Comput.* 63, 5 (2014), 1097–1110.
- [3] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC benchmark suite: Characterization and architectural implications. In *Proc. 17th Int. Conf. on Parallel Arch. and Compilation Techniques*. 72–81.
- [4] David Brooks, Robert P. Dick, Russ Joseph, and Li Shang. 2007. Power, thermal, and reliability modeling in nanometer-scale microprocessors. *IEEE Micro* 27, 3 (2007), 49–62.
- [5] Guoqing Chen, Yi Xu, Xing Hu, Xiangyang Guo, Jun Ma, Yu Hu, and Yuan Xie. 2016. TSocket: Thermal sustainable power budgeting. *ACM Trans. Des. Autom. of Electron. Syst.* 21, 2 (2016), 29:1–29:22.
- [6] Ryan Cochran and Sherief Reda. 2013. Thermal prediction and adaptive control through workload phase detection. *ACM Trans. Des. Autom. of Electron. Syst.* 18, 1 (2013), 7:1–7:19.
- [7] Puyan Dadvar and Kevin Skadron. 2005. Potential thermal security risks. In *21st IEEE Semicond. Thermal Meas. & Management Symp.* 229–234.
- [8] Begum Egilmez, Gokhan Memik, Seda Ogrenci-Memik, and Oguz Ergin. 2015. User-specific skin temperature-aware DVFS for smartphones. In *Proc. 2015 Design, Autom. & Test in Europe Conf. & Exhibition*. 1217–1220.
- [9] Ujjwal Gupta, Raid Ayoub, Michael Kishinevsky, David Kadjo, Niranjana Soundararajan, Ugurkhan Tursun, and Umit Ogras. 2017. Dynamic power budgeting for mobile systems running graphics workloads. *IEEE Trans. Multi-Scale Comput. Syst.* (2017).
- [10] M. R. Guthaus, J. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge, and R. B. Brown. 2001. Mibench: A free, commercially representative embedded benchmark suite. In *Proc. Int. Workshop on Workload Char.* 3–14.
- [11] Vinay Hanumaiah, Sarma Vruthula, and Karam S. Chatha. 2011. Performance optimal online DVFS and task migration techniques for thermally constrained multi-core processors. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.* 30, 11 (2011), 1677–1690.
- [12] Hardkernel. 2014. ODROID-XU3. http://www.hardkernel.com/main/products/prdt_info.php?g_code=g140448267127 Accessed 07/14/2017. (2014).
- [13] Jahangir Hasan, Ankit Jalote, T. N. Vijaykumar, and Carla E. Brodley. 2005. Heat stroke: Power-density-based denial of service in SMT. In *11th Int. Symp. on High-Performance Computer Arch.* 166–177.

- [14] John L. Henning. 2006. SPEC CPU2006 benchmark descriptions. *ACM SIGARCH Comput. Arch. News* 34, 4 (2006), 1–17.
- [15] Seongmoo Heo, Kenneth Barr, and Krste Asanović. 2003. Reducing power density through activity migration. In *Proc. 2003 Int. Symp. on Low Power Electron. and Design*. 217–222.
- [16] Pei-Yu Huang and Yu-Min Lee. 2009. Full-chip thermal analysis for the early design stage via generalized integral transforms. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 17, 5 (2009), 613–626.
- [17] Wei Huang, Shougata Ghosh, Siva Velusamy, Karthik Sankaranarayanan, Kevin Skadron, and Mircea R. Stan. 2006. HotSpot: A compact thermal modeling methodology for early-stage VLSI design. *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.* 14, 5 (2006), 501–513.
- [18] Canturk Isci, Alper Buyuktosunoglu, Chen-Yong Cher, Pradip Bose, and Margaret Martonosi. 2006. An analysis of efficient multi-core global power management policies: Maximizing performance for a given power budget. In *Proc. Int. Symp. on Microarchitecture*. 347–358.
- [19] Nam Sung Kim et al. 2003. Leakage current: Moore’s law meets static power. *Computer* 36, 12 (2003), 68–75.
- [20] Joonho Kong, Johnsy K. John, Eui-Young Chung, Sung Woo Chung, and Jie Hu. 2010. On the thermal attack in instruction caches. *IEEE Trans. Depend. and Sec. Comput.* 7, 2 (2010), 217–223.
- [21] Amit Kumar, Li Shang, Li-Shiuan Peh, and Niraj K. Jha. 2008. System-level dynamic thermal management for high-performance microprocessors. *IEEE Trans. Comput.-Aided Design of Integr. Circuits Syst.* 27, 1 (2008), 96–108.
- [22] Peng Li, Lawrence T. Pileggi, Mehdi Asheghi, and Rajit Chandra. 2004. Efficient full-chip thermal modeling and analysis. In *Proc. Int. Conf. on Comput.-Aided Design*. 319–326.
- [23] Weiping Liao and Lei He. 2003. Coupled power and thermal simulation with active cooling. In *Int. Workshop on Power-Aware Comput. Syst.* 148–163.
- [24] Yongpan Liu, Robert P. Dick, Li Shang, and Huazhong Yang. 2007. Accurate temperature-dependent integrated circuit leakage power estimation is easy. In *Proc. Conf. on Design, Autom. and Test in Europe*. 1526–1531.
- [25] Zao Liu, Tailong Xu, Sheldon X.-D. Tan, and Hai Wang. 2013. Dynamic thermal management for multi-core microprocessors considering transient thermal effects. In *18th Asia and South Pacific Design Autom. Conf.* 473–478.
- [26] Santiago Pagani, Heba Khdr, Jian-Jia Chen, Muhammad Shafique, Minming Li, and Jörg Henkel. 2017. Thermal safe power (TSP): Efficient power budgeting for heterogeneous manycore systems in dark silicon. *IEEE Trans. Comput.* 66, 1 (2017), 147–162.
- [27] Alok Prakash, Hussam Amrouch, Muhammad Shafique, Tulika Mitra, and Jörg Henkel. 2016. Improving mobile gaming performance through cooperative CPU-GPU thermal management. In *Proc. of the Design Autom. Conf.* 47:1–47:6.
- [28] Onur Sahin and A. Coskun. 2016. QScale: Thermally-efficient QoS management on heterogeneous mobile platforms. In *Proc. Int. Conf. on Comput.-Aided Design*. 125:1–125:8.
- [29] Samsung Electronics Co. 2016. Samsung Expands Recall to All Galaxy Note7 Devices. <http://www.samsung.com/us/note7recall/>. Accessed 07/14/2017. (2016).
- [30] Shervin Sharifi, Dilip Krishnaswamy, and Tajana Simunic Rosing. 2013. PROMETHEUS: A proactive method for thermal management of heterogeneous MPSoCs. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.* (2013), 1110–1123.
- [31] Gaurav Singla, Gurinderjit Kaur, Ali K. Unver, and Umit Y. Ogras. 2015. Predictive dynamic thermal and power management for heterogeneous mobile platforms. In *Proc. 2015 Design, Autom. & Test in Europe Conf. & Exhibition*. 960–965.
- [32] Arman Vassighi and Manoj Sachdev. 2006. Thermal runaway in integrated circuits. *IEEE Trans. Device Mater. Rel.* 6, 2 (2006), 300–305.
- [33] Qing Xie, Jaemin Kim, Yanzhi Wang, Donghwa Shin, Naehyuck Chang, and Massoud Pedram. 2013. Dynamic thermal management in mobile devices considering the thermal coupling between battery and application processor. In *Proc. Int. Conf. on Comput.-Aided Design*. 242–247.
- [34] Yonghong Yang, Zhenyu Gu, Changyun Zhu, Robert P. Dick, and Li Shang. 2007. ISAC: Integrated space-and-time-adaptive chip-package thermal analysis. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.* 26, 1 (2007), 86–99.
- [35] Yong Zhan and Sachin S. Sapatnekar. 2005. A high efficiency full-chip thermal simulation algorithm. In *Proc. Int. Conf. on Comput.-Aided Design*. 635–638.

Received April 2017; revised June 2017; accepted June 2017