

Thermal-Aware Resource Management for Embedded Real-Time Systems

Youngmoon Lee, *Student Member, IEEE*, Hoon Sung Chwa, *Member, IEEE*, Kang G. Shin, *Life Fellow, IEEE*
Shige Wang, *Senior Member, IEEE*

Abstract—With an increasing demand for complex and powerful System-on-Chips (SoCs), modern real-time automotive systems face significant challenges in managing on-chip-temperature. We demonstrate, via real experiments, the importance of accounting for *dynamic ambient temperature* and *task-level power dissipation* in resource management so as to meet both thermal and timing constraints. To address this problem, we propose RT-TRM, a real-time thermal-aware resource management framework. We first introduce a task-level dynamic power model that can capture different power dissipations with a simple task-level parameter called the *activity factor*. We then develop two new mechanisms, *adaptive parameter assignment* and *online idle-time scheduling*. The former adjusts voltage/frequency levels and task periods according to the varying ambient temperature while preserving feasibility. The latter generates a schedule by allocating idle times efficiently without missing any task/job deadline. By tightly integrating the solutions of these two mechanisms, we can guarantee both thermal and timing constraints in the presence of dynamic ambient temperature variations. We have implemented RT-TRM on an automotive microcontroller to demonstrate its effectiveness, achieving better resource utilization by 18.2% over other runtime approaches while meeting both thermal and timing constraints.

Index Terms—Thermal-aware resource management, embedded real-time systems, dynamic ambient temperature, task-level power dissipation.

I. INTRODUCTION

Thermal-aware resource management has become critically important to modern embedded real-time systems like automotive controls and smartphones as they are increasingly realized on powerful computing platforms with exponentially increasing power density. High on-chip temperatures shorten the platform lifetime and severely degrade its performance and reliability, risking safety (e.g., vehicle breakdown or smartphone explosion). We must, therefore, keep the processor temperature below the peak temperature constraint while satisfying all application timing constraints.

Manuscript received April 3, 2018; revised June 8, 2018; accepted July 2, 2018. This work was supported in part by the U.S. Office of Naval Research under Grant N00014-18-1-2141, and in part by the Global Research Laboratory Program through the National Research Foundation of Korea funded by the Ministry of Science and ICT under Grant AQ1 NRF-2013K1A1A2A02078326. This paper was recommended by Conference Chairperson B. Brandenburg. (Corresponding authors: Hoon Sung Chwa; Kang G. Shin.)

Y. Lee and K. G. Shin is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA.

H.S. Chwa was with the University of Michigan and is now with the Department of Information and Communication Engineering, DGIST, Daegu 42988 Republic of Korea.

S. Wang is with General Motors Global Research and Development Center, Warren, MI 48093 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>. Digital Object Identifier 10.1109/TCAD.2018.2857279

There are two key thermal issues to consider for embedded real-time systems: (1) dynamically *varying ambient temperature* and (2) task-level *power dissipation*. Our experimental evaluation has shown the ambient temperature of an automotive electronic module to vary highly and dynamically even during a single driving event, and its seasonal temperature also varies widely up to a difference of 28°C. Moreover, the average power consumed by each application differs up to 140%, according to our evaluation of various automotive benchmark applications (to be detailed in §III).

These dynamic thermal features pose significant challenges in meeting the application timing constraints. In particular, the maximum available computation power varies with the ambient temperature, as the temperature of a processor depends on its ambient. According to our experimental evaluation (§III), an increase of 14.9°C in ambient temperature results in a maximum reduction of 28.8% in the processor's computation power. In such a case, the processor's thermal constraint may be violated if it executes a task whose average power dissipation exceeds a certain limit. One may reduce the processor temperature by idling or slowing it down, but such an action may also lead to task/job deadline miss(es), thus violating the app timing constraint. This calls for adaptive resource management that considers dynamically varying ambient temperature and task-level power dissipation so as to meet both the processor's thermal and the app's timing constraints.

A significant amount of work has been done on real-time thermal-aware scheduling (see [1] for a survey). Existing approaches usually employ DVFS scheduling [2], [3], idle-time scheduling [4], or task scheduling [5] to minimize the peak temperature while guaranteeing the timing constraint. Worst-case temperature analyses [6], [7] have also been proposed for offline guarantees to meet thermal constraints. The concept of thermal utilization [8] was introduced to capture the thermal impact of periodic real-time tasks on processors. Most of these existing solutions, however, assume either *fixed* ambient temperature or *constant, task-independent* power dissipation. Although there exist real-time feedback thermal controllers that minimize the error between the current processor temperature and the desired temperature by regulating task utilization [9] or frequency [10], they are limited to guaranteeing thermal constraints due to temperature overshooting.

In this paper, we propose a new real-time thermal-aware resource management framework, called RT-TRM, that captures not only varying computation power bounds due to the variations of ambient temperature but also different power demands by different tasks. RT-TRM adaptively makes a parameter assignment (voltage/frequency levels and a task period assignment) and builds a schedule (the processor idling

or task-execution, and priority ordering of tasks) so as to meet both thermal and timing constraints.

We first propose a *task-level* dynamic power model that uses a simple task-level parameter, called the *activity factor*, to capture different power dissipations by different tasks. Building on this dynamic power model, we study the effect of task-execution or processor-idling on the processor temperature. Our model is also validated experimentally with several automotive benchmarks running in various realistic environments. Second, we propose the notions of *dynamic power demand* of a task set and *dynamic power bound* at a given ambient temperature and derive the feasibility conditions for a parameter assignment with respect to both thermal and timing constraints. We then develop a runtime adaptive strategy that can preserve feasibility under ambient temperature variations by adjusting the parameter assignment. Third, building on a feasible parameter assignment, we develop an online scheduling policy that determines not only the processor state (active or idle) but also the order of executing tasks in active state. Our scheduling algorithm can reclaim slack (spare capacity) at runtime and allocate it to tasks *in proportion* to their power demands by considering the fact that a task with higher power dissipation should be assigned more idle time. This way, we can meet both thermal and timing constraints with much fewer preemptions. Finally, we have implemented RT-TRM on an automotive microcontroller to demonstrate its effectiveness in dealing with ambient temperature variations. RT-TRM is shown to improve resource utilization by 18.2% over the existing runtime feedback thermal controllers while guaranteeing both thermal and timing constraints.

This paper makes the following main contributions:

- Demonstration of the importance of accounting for dynamic ambient temperature and task-level power dissipation for thermal-aware resource management (§III);
- Development of a dynamic power model that captures different power dissipations with a simple task-level parameter called the *activity factor* and its experimental validation (§IV);
- Development of an adaptive parameter-assignment framework under varying ambient temperature while preserving feasibility (§V);
- Development of an online idle-time scheduling algorithm that enables dynamic idle-time allocation with much fewer preemptions while guaranteeing both thermal and timing constraints (§VI);
- Implementation and evaluation of the effectiveness of RT-TRM on an automotive microcontroller (§VII).

II. RELATED WORK

Significant work has been done on thermal management at both hardware (e.g., architecture design, floorplan) and OS level (e.g., thermal-aware DVFS, scheduling) [1]. The focus of this paper is on OS-level thermal-aware resource management for hard real-time systems, such as cars.

Thermal-aware real-time scheduling has been an active subject of research trying to meet timing and thermal constraints in a *constant* environment. DVFS scheduling determines the voltage and frequency of a processor to minimize the

power consumption [11], [12] and peak temperature subject to timing constraints on single-core [2], [13], [14] or multi-core platforms [3]. Multi-core task scheduling [5] determines task-to-core assignment and scheduling to minimize the peak temperature. Thermal shaping puts idle periods during task execution to reduce the temperature without missing deadlines [4], [15].

Researchers focused on different task-level power dissipations to reduce the peak temperature [16], [17] or maximize throughput [18], [19] by interleaving the execution of hot and cold tasks. By analyzing such task-level power variations, the peak temperature was derived to meet the thermal constraint [6]. The concept of thermal utilization was introduced in [8], [20] to capture the different thermal impacts of real-time tasks. A few researchers considered adaptive thermal-aware resource management for real-time tasks to cope with dynamic environments. Feedback controllers regulate the processor temperature by adjusting processor utilization [9], [21] or operating frequency [10] subject to the timing constraint.

While researchers developed task-level scheduling and processor-level thermal control techniques to deal with both the thermal and timing requirements, they have not yet addressed both large environmental variations and peak temperature caused by task workloads together. To meet this need, we first verify the significance of these factors in automotive systems, then develop and validate a task-level thermal model that can capture individual tasks' different power dissipations. Building upon the task-level thermal model, we propose a new thermal-aware resource management scheme that (i) jointly adapts task periods and processor frequency in response to the varying ambient temperature and (ii) schedules tasks to meet both thermal and timing constraints.

III. TARGET SYSTEM, CHALLENGES, AND SOLUTION OVERVIEW

This section describes our target system (§III-A) and introduces the challenges faced therein (§III-B) followed by an overview of our approach (§III-C).

A. Target System

We consider an embedded real-time system running a set of real-time tasks on a computing platform.

Processor and task model. Our target system is a uniprocessor platform that provides dynamic voltage and frequency scaling (DVFS) with a separate set of discrete frequency/voltage levels. If an operating frequency f is adjustable within the specified range of $[f_{min}, f_{max}]$, its corresponding voltage V is determined according to a typical implementation principle [22]. We also consider a task set τ composed of implicit-deadline periodic tasks. Each task $\tau_i \in \tau$ is characterized by period p_i and its worst-case execution time (WCET) $e_i(f)$ as a function of operating frequency f . We assume that p_i is adjustable within $[p_i^{min}, p_i^{max}]$ based on typical application elasticity [23], [24]. Note that, for a task whose period is fixed, we set its period range as $p_i^{min} = p_i = p_i^{max}$. Such τ_i is assumed to generate a sequence of jobs, once every

p_i time-units, with each job needing to complete $e_i(f)$ within a relative deadline of p_i time-units.

Power and thermal model. We consider dynamic power management where the processor is in either *idle* or *active* state. The processor is said to be active if it is currently executing a job, or idle otherwise. Its power dissipation (P_{proc}) is then expressed as $P_{proc} = P_{leak} + P_{dyn}$, where P_{leak} is the leakage power for the processor to stay ready (in active or idle state) for execution of jobs, and P_{dyn} is the additional dynamic power to execute a job (in active state). The term P_{leak} is modeled as [25]: $P_{leak} = V \cdot (\beta_1 \cdot T + \beta_0)$, where β_1 and β_0 are processor-dependent constants, and T is the processor's temperature. Note that P_{dyn} depends on the task currently running on the processor, and its detailed model will be described in §IV-A.

To translate the processor's power dissipation to its temperature, we use a well-known thermal circuit model [26]. If average processor power and ambient temperature are $P_{proc}(t)$ and $T_{amb}(t)$, respectively, over a time period t , then the processor temperature $T(t)$ at the end of this period is

$$T(t) = T(0) \cdot e^{-\frac{t}{R \cdot C}} + (T_{amb}(t) + P_{proc}(t) \cdot R) \cdot (1 - e^{-\frac{t}{R \cdot C}}), \quad (1)$$

where R and C are the thermal resistance and capacitance, respectively, and $T(0)$ is the initial temperature of the processor. We can observe from Eq. (1) that the temperature will increase/decrease towards and eventually reach $T_{amb}(t) + P_{proc}(t) \cdot R$. We define the steady temperature $T(\infty)$ of the processor as

$$T(\infty) = T_{amb}(t) + P_{proc}(t) \cdot R. \quad (2)$$

B. Problem Statement and Motivation

Problem Definition. We want to address the following real-time thermal-aware resource management problem.

Definition 1: Given a task set τ running on a uniprocessor, determine (i) the voltage/frequency (V/f) level, (ii) the period $\{p_i\}$ of task $\tau_i \in \tau$ parameters, and (iii) the schedule of jobs such that (a) temperature $T(t)$ never exceeds the peak temperature T_{max} (thermal constraint), and (b) all jobs of $\tau_i \in \tau$ meet their deadlines for all possible legitimate job arrival sequences (timing constraint).

To generate a job schedule, we need to determine not only the processor state (active or idle) but also the order of executing jobs in active state. From real embedded systems (e.g., cars and smartphones), we found two key thermal characteristics: (1) dynamic changes in the ambient temperature and (2) different power dissipations by different tasks. These are the primary motivation behind RT-TRM.

Dynamic changes on ambient temperature. Unlike desktops or data-centers, embedded real-time systems experience a wide range of environmental variations (especially the ambient temperature) during their operation/life. To confirm this fact, we measured the ambient temperature of a vehicle infotainment module embedded in the dashboard over days and months, and the results are plotted in Fig. 1. When the car was driven for several days, the change in the ambient temperature was highly dynamic and fluctuating between 0°C




Figure 1: Ambient temperature variations over time and the corresponding available computation power.




Figure 2: Average power consumptions for various automotive applications.

and 23°C. During a single driving event on Feb. 21, 2018 the ambient temperature was increased by up to 180%. The seasonal variation of the ambient temperature is also very wide. A similar phenomenon was also reported in [27] for car engine and transmission control units.

Under such a varying ambient temperature, real-time thermal-aware resource management becomes much more challenging as the processor's temperature is affected by the ambient temperature. We can use Eq. (2) to calculate the maximum processor's power dissipation without exceeding T_{max} for a given ambient temperature T_{amb} as $\frac{T_{max} - T_{amb}}{R}$. The change in the maximum processor's computation power under a varying ambient temperature is then plotted in Fig. 1 (the gray line).¹ For example, on Feb. 21, 2018 as the ambient temperature increased by 14.9°C from 8.3°C, the available processor computation power was decreased by 28.8%.

Task-level power dissipations. We also measured the processor's average power consumption to run various automotive benchmarks [29], and plotted the results in Fig. 2. Each app is shown to consume a different amount of power. For example, a table lookup task consumes 1726mW, while a bit manipulation task does 2348mW at the maximum processor frequency.²

In summary, the available processor's computation power varies with the ambient temperature. In addition, the execution of each task imposes a different power demand on the processor. So, we need to consider different task-level power dissipations and make adaptive parameter assignments and job schedules according to the varying ambient temperature so as to meet both thermal and timing requirements.

¹We set $T_{max} = 60^\circ\text{C}$ and $R = 22^\circ\text{C/W}$ based on the specification of an automotive microcontroller [28].

²A table lookup operation is used by an engine control module to find an output value corresponding to an input value (e.g., the ignition angle). A bit-manipulation operation is used by a display module where the pixels are moved into a display buffer until the entire buffer is displayed.

C. Overview of the Proposed Approach

To solve the real-time thermal-aware resource management problem while considering the varying ambient temperature and diverse task-level power dissipations, we address the following questions:

- Q1.** How to model power dissipations of different tasks and analyze the impact of their execution on the thermal behavior?
- Q2.** How to make adaptive parameter assignments under dynamically changing ambient temperature while meeting both thermal and timing constraints?
- Q3.** How to derive a job schedule meeting both thermal and timing constraints based on parameter assignment?

To answer Q1, we develop a *task-level* dynamic power model by using a simple task-level parameter called the *activity factor*. Based on this task-level dynamic power model, we analyze the effect of task execution on the processor's temperature, i.e., whether it increases or decreases the processor's temperature. Moreover, we empirically determine the activity factors for several automotive apps and verify our model in various environments, i.e., under different processor-frequency/task-utilization settings, varying ambient temperature, and executing multiple tasks (§IV).

To address Q2, we define a *dynamic power demand* of a task set τ that represents the total dynamic power demand by τ at the processor's steady temperature. We also define a *dynamic power bound* function of T_{amb} that represents the maximum processor's dynamic power P_{dyn} at T_{amb} without violating the thermal constraint. Based on these concepts, we derive the feasibility conditions of a task set and formulate an optimization problem that finds a feasible parameter assignment for a given T_{amb} . We also develop a runtime adaptive strategy that can preserve feasibility by adapting the parameter assignment to ambient temperature changes. We determine a tolerable ambient temperature range for parameter adaptation by considering the trade-off between adaptation overhead and resource efficiency. With our adaptive parameter assignment, the steady temperature is guaranteed not to exceed the peak temperature limit without missing any deadlines in the presence of ambient temperature variations (§V).

To answer Q3, building on a feasible parameter assignment derived by answering Q2, we develop an online scheduling policy. A task schedule may affect the *transient temperature*, potentially violating the thermal constraint before reaching the steady temperature. To avoid this, we calculate the *minimum idle-time* required for the execution of each job with respect to the thermal constraint. We then develop an idle-time scheduling algorithm that can reclaim unused resources at runtime and utilize them to allocate idle-time efficiently while meeting all deadlines with the minimum idle-time for each task. As a result, our algorithm can guarantee both thermal and timing constraints with much fewer preemptions (§VI).

IV. TASK-LEVEL POWER MODEL

We present a task-level power-consumption model that captures different dynamic power dissipations by individual tasks. In particular, we use a simple task-level activity factor

to capture each task's dynamic power dissipation (§IV-A) and empirically validate the model using an automotive platform and workloads (§IV-B).

A. Task-Level Dynamic Power Model

For automotive workloads, power dissipation is found to vary significantly with the executing task (Fig. 2). Since individual tasks programmed with distinct sets of instructions generate different switching activities and dynamic power dissipations, we used a task-level activity factor α_i to capture such different dynamic power P_i consumed by each task τ_i as $P_i = V^2 \cdot f \cdot \alpha_i$. Using this task-level dynamic power model, we can analyze how the processor's temperature changes with the execution of each job/task. Let $T(t)$ ($T_i(t + e_i(f))$) be the temperature at the beginning (end) of the execution of a job of τ_i . Using Eqs. (1) and (2), $T_i(t + e_i(f))$ can be written as:

$$T_i(t + e_i(f)) = T(t) \cdot e^{-\frac{e_i(f)}{R \cdot C}} + T_i^\infty(T_{amb}) \cdot (1 - e^{-\frac{e_i(f)}{R \cdot C}}), \quad (3)$$

where $T_i^\infty(T_{amb})$ is the steady temperature associated with the execution of τ_i that would be reached if the processor executes τ_i continuously. $T_i^\infty(T_{amb})$ can be expressed as:

$$T_i^\infty(T_{amb}) = T_{amb} + (P_i + P_{leak}) \cdot R. \quad (4)$$

We observe from Eqs. (3) and (4) that (i) if $T(t) < T_i^\infty(T_{amb})$ then the temperature increases towards $T_i^\infty(T_{amb})$, and (ii) if $T(t) \geq T_i^\infty(T_{amb})$ then the temperature decreases towards $T_i^\infty(T_{amb})$. A task τ_i is said to be *hot* if $T_i^\infty(T_{amb}) > T_{max}$, or *cold* otherwise. Depending on T_{amb} , τ_i can become hot or cold.

To consider the effect of idling the processor on its temperature, we let $T_0(t+l)$ denote the temperature at the end of an idle period of length l . Similarly to Eq. (3), $T_0(t+l)$ can be written as:

$$T_0(t+l) = T(t) \cdot e^{-\frac{l}{R \cdot C}} + T_0^\infty(T_{amb}) \cdot (1 - e^{-\frac{l}{R \cdot C}}), \quad (5)$$

where $T_0^\infty(T_{amb}) = T_{amb} + P_{leak} \cdot R$ is the processor's steady temperature in idle state.

So far, we have discussed the thermal effect of continuous execution (idling) of a single job (a processor). Now, let's consider the impact of a schedule of periodic tasks and idle-times. Let $T(t, W(t))$ denote the temperature at the end of a schedule $W(t) = \{w_i(t)\}$, where $w_i(t)$ is the total workload scheduled in $(0, t]$. Then, $T(t, W(t))$ can be written as:

$$T(t, W(t)) = T(0) \cdot e^{-\frac{t}{R \cdot C}} + (T_{amb} + (\sum_{\tau_i} P_i \cdot \frac{w_i(t)}{t} + P_{leak}) \cdot R) \cdot (1 - e^{-\frac{t}{R \cdot C}}), \quad (6)$$

where $\sum_{\tau_i} P_i \cdot \frac{w_i(t)}{t}$ is the average dynamic power consumed by $W(t)$. Note that every task τ_i generates a sequence of jobs executing $e_i(f)$ every p_i time-units, consuming an average dynamic power of $P_i \cdot \frac{e_i(f)}{p_i}$. We can then define the steady temperature $T(\infty, \tau)$ of a task set τ as:

$$T(\infty, \tau) = T_{amb} + (\sum_{\tau_i} P_i \cdot \frac{e_i(f)}{p_i} + P_{leak}) \cdot R. \quad (7)$$

Table I: Identifying activity factors and maximum errors

Task	Angle	Bit	Table	Edge	FFT	PID
T_i^∞ ($^\circ\text{C}$)	66.1	73.6	59.9	61.6	69.5	67.8
α_i	0.355	0.446	0.284	0.304	0.435	0.377

Note that the steady temperature of a task set is independent of its schedule, which serves as a basis for the feasibility condition presented in §V.

Identifying task-level activity factors. To identify the activity factor of each task, we ran automotive benchmarks, one at a time, with 100% resource utilization at the maximum frequency under the room temperature. We then measured the steady temperature and determined each task’s activity factor by using Eq. (4), and presented them in Table I.³ The activity factor varies greatly with tasks by up to 65%. For example, a table-lookup task with a large number of conditional switches and I/O accesses shows a low activity factor, while a bit-manipulation task with a high instruction-per-cycle rate shows a high activity factor.⁴

B. Model Validation

To confirm that the task-level dynamic power model and its thermal effect represent real hardware behaviors, we measured the steady temperature of the processor and compared it with our model’s estimation under various settings. In particular, we validated our model under (i) different processor-frequency/task-utilization settings, (ii) running multiple tasks together, and (iii) different ambient temperatures.

First, as shown in Fig. 3, we varied the task period to achieve the processor utilization ranging from 10% to 90% with a 10% increment as well as the frequency level from 0.4GHz to 1GHz for each task and measured the steady temperature.⁵ Fig. 3 shows the measured steady temperature as dotted points while the estimations with our model are plotted as lines with an error up to 0.65 $^\circ\text{C}$.

Second, as shown in Fig. 4, we ran two tasks — bit manipulation and angle-time conversion — together on a single core at 1GHz by varying their utilizations. The result shows that the steady temperature of the processor is linearly increased with each task’s utilization (plotted as a linear surface) as formulated in Eq. (7). We also confirmed that the same tendency is observed for different numbers of tasks (i.e., 4 and 8 tasks) with an error up to 1.2 $^\circ\text{C}$.

Finally, we validated our model for different ambient temperatures from 20 $^\circ\text{C}$ to 35 $^\circ\text{C}$. For each configuration, we repeated this 10 times with sufficient intervals, showing an error of up to 0.98 $^\circ\text{C}$.

V. ADAPTIVE PARAMETER ASSIGNMENT

We now present how to adjust a voltage/frequency level and task period assignment according to the varying ambient temperature, called *Adaptive Parameter Assignment Framework* (APAF). Specifically, we derive feasibility conditions for

³The detailed experimental setup will be given in §VII.

⁴The activity factor α is normalized by the maximum power, i.e., $\alpha = 1$ means the maximum power dissipation.

⁵See §VII for more details of the experimental setup.

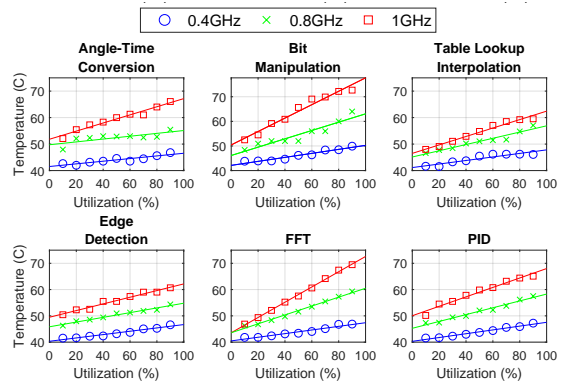



Figure 3: Model validation with varying utilizations.

Figure 4: Model validation with two periodic tasks (*Bit manipulation, Angle-time Conversion*).

a parameter assignment, formulate a parameter optimization problem, and introduce a runtime strategy for adapting to the varying ambient temperature.

A. Parameter Assignment

We first consider the *feasible parameter assignment* problem for a given ambient temperature.

Definition 2 (Feasible parameter assignment): Given a task set τ and the ambient temperature T_{amb} , determine V , f and p_i for every $\tau_i \in \tau$ such that if τ is feasible (i.e., meeting the thermal and timing constraints), it remains feasible even with the new parameter assignment.

To solve this problem, we introduce two conditions for a parameter assignment to be feasible with respect to thermal and timing constraints for a given ambient temperature and formulate an optimization problem to find a feasible parameter assignment.

Feasibility condition. Recall that for a given task set τ , the processor temperature will eventually reach the steady temperature $T(\infty, \tau)$ of τ (defined in Eq. (7)) regardless of its schedule. Therefore, to meet the thermal constraint, the steady temperature $T(\infty, \tau)$ should be lower than or equal to the peak temperature limit T_{max} ,

$$C1: T(\infty, \tau) \leq T_{max}. \quad (8)$$

We define a *dynamic power demand* $PD(\tau)$ of τ as the total dynamic power demand by τ at the steady temperature, which is:

$$PD(\tau) = \sum_{\tau_i} P_i \cdot \frac{e_i(f)}{p_i} = V^2 \cdot f \cdot \sum_{\tau_i} \alpha_i \cdot \frac{e_i(f)}{p_i}. \quad (9)$$

We also define a *dynamic power bound* $PB(T_{amb})$ of T_{amb} as the maximum processor's dynamic power at T_{amb} without exceeding T_{max} . We can derive $PB(T_{amb})$ by solving $T(\infty, \tau) = T_{max}$:

$$PB(T_{amb}) = \frac{T_{max} - T_{amb}}{R} - V \cdot (\beta_1 \cdot T_{max} + \beta_0). \quad (10)$$

Using these, the feasibility condition $C1$ with respect to the thermal constraint can be re-written as

$$\overline{C1}: PD(\tau) \leq PB(T_{amb}). \quad (11)$$

To meet the timing constraint, we use a well-known exact feasibility analysis by Liu and Layland [30]:

$$C2: \sum_{\tau_i} \frac{e_i(f)}{p_i} \leq 1. \quad (12)$$

If a parameter assignment satisfies both $\overline{C1}$ and $C2$, the steady temperature of τ is guaranteed not to exceed T_{max} without missing any task deadline when a task set is scheduled by an optimal scheduling algorithm. However, as can be seen from Eq. (6), a job schedule may affect a transient temperature $T(t, W(t))$, potentially violating the thermal constraint before reaching the steady temperature. To avoid this situation, we define the minimum idle-time $I_i^{min}(T_{amb})$ required for the execution of each job without violating the thermal constraint and include the term in $C2$ (to be detailed in §VI). Then, the feasibility condition $C2$ can be extended to:

$$\overline{C2}: \sum_{\tau_i} \frac{e_i(f) + I_i^{min}(T_{amb})}{p_i} \leq 1. \quad (13)$$

So, if there exists a parameter assignment satisfying both $\overline{C1}$ and $\overline{C2}$, we can guarantee that a task set τ is feasible with respect to both thermal and timing constraints.

Parameter optimization. We formulate the parameter assignment problem as an optimization problem subject to the feasibility conditions ($\overline{C1}$ and $\overline{C2}$):

$$\text{maximize}_{f, p_i} \sum_{\tau_i} w_i \cdot \frac{1}{p_i} \quad (14)$$

$$\text{s.t.} \quad \overline{C1}: PD(\tau) = V^2 \cdot f \cdot \sum_{\tau_i} \alpha_i \cdot \frac{e_i(f)}{p_i} \leq PB(T_{amb}) \quad (15)$$


$$\overline{C2}: \sum_{\tau_i} \frac{e_i(f) + I_i^{min}(T_{amb})}{p_i} \leq 1 \quad (16)$$

$$f \in [f_{min}, \dots, f_{max}]. \quad (17)$$

$$\forall \tau_i \quad p_i^{min} \leq p_i \leq p_i^{max} \quad (18)$$

As an optimization goal, a QoS function associated with resource usage can be used as in [23], [24]. Our objective in Eq. (14) is to maximize the weighted sum of each task-rate $\frac{1}{p_i}$.⁶ Eq. (17) specifies the discrete frequency scaling levels available on the processor. Eq. (18) specifies the minimum

⁶The value of w_i can be determined by the importance of each task.



(a) Different Adaptation Interval ΔT (b) Impact of ΔT on Overhead

Figure 5: Runtime adaptation with (a) different different adaptation intervals and (b) the trade-off between adaptation overhead and resource efficiency

and maximum bounds of an allowable task period within $[p_i^{min}, p_i^{max}]$. We use linear programming to determine a task period assignment for a given voltage/frequency level starting from the maximum level. If there is no solution, we lower the voltage/frequency level until a feasible solution is found. The computational complexity is $O(m \cdot n^{3.5})$ for n tasks and m frequency scaling levels [31].

B. Runtime Parameter Adaptation

We propose a runtime parameter adaptation strategy that samples ambient temperature variations and dynamically adjusts the voltage/frequency level and period assignment. For this, we need to determine when and how to adjust the parameter assignment. We set fixed points of the ambient temperature threshold $\{TS_{amb}(k)\}$, which are determined by

$$TS_{amb}(k+1) = TS_{amb}(k) + \Delta T, \quad (19)$$

where ΔT is a tolerable ambient temperature range.

Our runtime adaptation strategy periodically estimates the ambient temperature and adjusts the parameter assignment whenever the sampled ambient temperature is out of the range $(TS_{amb}(k), TS_{amb}(k+1)]$ for any k . The parameter assignment in each range is determined by solving the optimization problem in Eq. (14) with the ambient temperature of $TS_{amb}(k+1)$. The challenge is how to choose ΔT and estimate the ambient temperature.

Determining a tolerable ambient temperature range.

There exists a trade-off between resource utilization and adaptation overhead in choosing ΔT . A smaller ΔT can achieve efficient resource utilization with prompt response upon small ambient temperature changes at the expense of high adaptation overhead. If the adaptation interval ΔT is too large, coarse-grained parameter adaptation incurs resource utilization loss.

To determine the optimal value of ΔT , we analyze the ambient temperature trace in Fig. 1a and compare the runtime overhead and resource efficiency depending on ΔT . Fig. 5a illustrates how our parameter adaptation responds to the varying ambient temperature for different values of ΔT . From the trace, we obtain the adaptation overhead and resource utilization loss for each value of ΔT .

Fig. 5b shows the trade-off between resource efficiency and adaptation overhead, where the adaptation overhead (dotted line) decreases, but the resource utilization loss (grey line) increases as ΔT increases. The adaptation overhead in our

experimental setup (§VII) was about 27ms. When adapting every sampling period ($\Delta T = 0$), processor utilization overhead incurred is 2.7%. (Fig. 5b). We set the optimal value of ΔT to the point where the sum of the adaptation overhead and resource utilization loss (solid line) is minimized, which is $\Delta T = 1^\circ\text{C}$.

VI. ONLINE IDLE-TIME SCHEDULING

So far, we have discussed how to adaptively adjust the processor's voltage/frequency and the task periods under the varying ambient temperature. We now consider how to schedule task/job executions and idle-times in order to meet both thermal and timing requirements. Specifically, we want to address the following problem, which we call the *schedule-generation* problem.

Definition 3 (Schedule generation): Given the assignment of V , f , and $\{p_i\}$ (with APAF), determine a schedule of job executions and idle-times such that the processor temperature $T(t)$ does not exceed T_{max} at any time t while all jobs of all tasks $\tau_i \in \tau$ meet their deadlines.

To solve this problem, we must consider two key issues: 1) transient temperature $T(t)$ varies with the task running at any given time, and 2) ambient temperature T_{amb} also affects $T(t)$. Suppose that the processor has reached T_{max} (i.e., $T(t) = T_{max}$) and two tasks — a cold task τ_1 and a hot task τ_2 — are ready to run at time t . If a cold task τ_1 is scheduled, the temperature will decrease since $T_1^\infty(T_{amb}) \leq T_{max}$. On the other hand, if a hot task τ_2 starts to run immediately, the temperature will increase (since $T_2^\infty(T_{amb}) > T_{max}$), and the thermal constraint will be violated. To avoid the processor temperature exceeding T_{max} , we must idle the processor to drop its temperature to a *safe* temperature before executing τ_2 . With this safe temperature, continued execution of τ_2 will not violate the thermal constraint. The main challenge is then how to derive a safe temperature and schedule idle-times to reach the temperature before executing each hot task. Note that each task has a different power dissipation, so the safe temperature may vary with task. Moreover, the amount of idle time required to reach a safe temperature varies with the ambient temperature. Without a proper idle-time scheduling decision, it may end up with some undesirable situations, such as those where (a) the temperature exceeds T_{max} and/or (b) a task/job deadline miss occurs.

To resolve such issues, we develop a thermal-aware on-line idle-time scheduling policy that determines idle-times between the execution of tasks to meet both thermal and timing constraints. We assume that tasks is priority-ordered according to the earliest deadline first (EDF) policy. We calculate the *minimum idle-time* required for the execution of each task to avoid the situation (a) and take the minimum idle-time into account in our adaptive parameter assignment to avoid the situation (b). Our proposed online scheduling algorithm then makes the trade-off between the total amount of required idle-time and preemption overhead. In particular, it updates available slack at runtime and effectively utilizes it to allocate more idle-time with much fewer preemptions while guaranteeing both thermal and timing constraints.

Calculation on the minimum idle-time. Now, we describe the relation between the amount of necessary idle-time and the number of preemptions. We first consider the case of executing a hot task τ_i for $e_i(f)$ units without any preemption. We define the safe temperature of τ_i to execute for $e_i(f)$ units at T_{amb} (denoted by $T_i^{safe}(e_i(f), T_{amb})$) as the initial temperature at which the temperature reaches T_{max} after the execution of $e_i(f)$ units. The safe temperature can then be derived by solving the term $T(t)$ in Eq. (3) when $T_i(t + e_i(f)) = T_{max}$:

$$T_i^{safe}(e_i(f), T_{amb}) = T_i^\infty(T_{amb}) - \frac{T_i^\infty(T_{amb}) - T_{max}}{e^{-\frac{e_i(f)}{R \cdot C}}}. \quad (20)$$

Similarly, we can calculate the idle-time necessary to reach $T_i^{safe}(e_i(f), T_{amb})$ (denoted by $t_{idle}(e_i(f), T_{amb})$) by solving the term l in Eq. (5) when $T_0(t + l) = T_i^{safe}(e_i(f), T_{amb})$ and $T(t) = T_{max}$:

$$t_{idle}(e_i(f), T_{amb}) = R \cdot C \cdot \ln\left(\frac{T_{max} - T_0^\infty(T_{amb})}{T_i^{safe}(e_i(f), T_{amb}) - T_0^\infty(T_{amb})}\right). \quad (21)$$

Now, let's consider the case where preemption is allowed, i.e., splitting each task τ_i into multiple — m_i ($m_i > 1$) — sub-tasks and inserting idle-time in between. Likewise, by using Eqs. (20) and (21), we can calculate the safe temperature and idle-time required for executing each sub-task for $\frac{e_i(f)}{m_i}$ units. Then, the cumulative idle-time to execute m_i sub-tasks at T_{amb} can be calculated as $m_i \cdot t_{idle}(\frac{e_i(f)}{m_i}, T_{amb})$.

Fig. 6 shows the cumulative idle-time as m_i increases. It is important to observe that the more sub-tasks, the less cumulative idle-time required, as was also observed in [19]. In Fig. 6(a), we can see that the amount of required idle-time also depends the ambient temperature. As shown in Fig. 6(b), each task requires a different amount of idle-time. Note that cold tasks — e.g., a table-lookup task — does not require idle-time. The results shown in Fig. 6 imply that the cumulative idle-time can be reduced by splitting each task into more sub-tasks with frequent idling of the processor. However, the benefit of frequent idling becomes saturated as m_i increases, and the preemption overhead can no longer be ignored. Considering this, we derive the minimum idle-time for a task τ_i (denoted by $I_i^{min}(T_{amb})$) as follows. We calculate a decreasing amount of cumulative idle-time by taking derivative $\frac{\partial}{\partial m_i}(m_i \cdot t_{idle}(\frac{e_i(f)}{m_i}, T_{amb}))$, and find the value of m_i (denoted by m_i^{max}) where the value of the derivative becomes closest to the preemption cost for switching between active and idle states. Then, the minimum idle-time of τ_i can be calculated as

$$I_i^{min}(T_{amb}) = m_i^{max} \cdot t_{idle}(\frac{e_i(f)}{m_i^{max}}, T_{amb}). \quad (22)$$

Note that it is sufficient to update the minimum idle-time of each task only when there exists any parameter change caused by our adaptive parameter assignment.

Guarantee of thermal and timing constraints. For every invocation of a task τ_i , if the minimum idle-time is correctly scheduled before finishing the execution of τ_i , we can guarantee that the thermal constraint is never violated. The question then becomes how to guarantee the timing constraint when

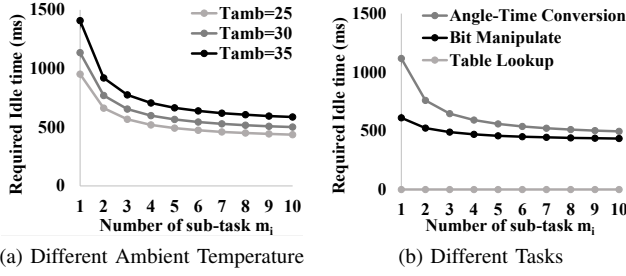


Figure 6: Cumulative idle-time for (a) different ambient temperature and (b) different tasks decreases with the number of subtasks m_i

all tasks are scheduled together with their minimum idle-time. To address this, we derive a new feasibility condition by incorporating the minimum idle-time for each task. In order for a task set τ to be feasible under both thermal and timing constraints, every job of each task τ_i should have its minimum idle-time (for at least $I_i^{min}(T_{amb})$) and finish its execution (for at most its WCET $e_i(f)$) before its deadline. Then, a new feasibility condition can be derived by extending the utilization-based exact feasibility analysis by Liu and Layland [30]:

$$\sum_{\tau_i} \frac{e_i(f) + I_i^{min}(T_{amb})}{p_i} \leq 1. \quad (23)$$

We include the feasibility condition (Eq. (23)) in the optimization formulation for the parameter assignment presented in §V-A. This way, RT-TRM can guarantee both thermal and timing constraints.

Online idle-time scheduling. Building upon the parameter assignment obtained by APAF, if we divide each task τ_i into $m_i(I_i^{min}(T_{amb}))$ sub-tasks and evenly distribute the idle-time $I_i^{min}(T_{amb})$ between the execution of each sub-task, we can schedule all tasks without violating thermal and timing constraints. However, such *static* idle time allocation under pessimistic assumptions cannot efficiently utilize all available slack resources at runtime, which may, in turn, incur unnecessary preemption overheads. Therefore, we develop an online idle-time scheduling algorithm that reclaims unused resources and utilize them to allocate *dynamic* idle-time for each task in an efficient way. As a result, our algorithm can meet both thermal and timing requirements with much fewer preemptions.

Described below is our online idle-time scheduling algorithm. The scheduler is invoked upon (i) release of a new job (JOB_RELEASE), (ii) completion of a job (JOB_COMPLETION), or (iii) update of frequency by APAF (FREQ_UPDATE). The scheduler keeps track of the worst-case remaining execution time, $e_{left_i}(f)$ for the active job of τ_i . This is set to $e_i(f)$ on JOB_RELEASE, decremented as the job executes, updated according to frequency change on FREQ_UPDATE, and set to 0 on JOB_COMPLETION. Upon each invocation (either JOB_RELEASE, JOB_COMPLETION, or FREQ_UPDATE), the scheduler updates available slack $S(t_{cur}, d_1(t_{cur}))$ for the interval of $[t_{cur}, d_1(t_{cur}))$, where t_{cur} is the current time instant and $d_1(t_{cur})$ is the earliest

Algorithm 1 Slack calculation

```

1:  $U = \sum_{\tau_i} \frac{e_i(f) + I_i^{min}(T_{amb})}{p_i}$ 
2:  $p = 0$ 
3: for  $i = n$  to 1,  $\tau_i \in \{\tau_1, \dots, \tau_n | d_1(t_{cur}) \leq \dots \leq d_n(t_{cur})\}$ 
   do
4:    $\triangleright$  In reverse EDF order of tasks
5:    $U = U - \frac{e_i(f) + I_i^{min}(T_{amb})}{p_i}$ 
6:    $q_i = \max(0, e_{left_i}(f) + I_i^{min}(T_{amb}) - (1 - U) \cdot (d_i(t_{cur}) - d_1(t_{cur})))$ 
7:    $U = \min(1.0, U + \frac{e_{left_i}(f) + I_i^{min}(T_{amb}) - q_i}{d_i(t_{cur}) - d_1(t_{cur})})$ 
8:    $p = p + q_i$ 
9: end for
10:  $S(t_{cur}, d_1(t_{cur})) = d_1(t_{cur}) - t_{cur} - p$ 

```

absolute deadline among all released jobs whose deadline is after t_{cur} . Then, the scheduler assigns slack $S(t_{cur}, d_1(t_{cur}))$ to tasks in *proportion* to their average power dissipation (i.e., $P_i \cdot \frac{e_i(f)}{p_i}$). The rationale for such a proportional slack distribution is that a task with higher power dissipation requires more idle-time. In this way, each task is assigned an amount of idle-time equal to

$$I_i(t_{cur}) = I_i^{min}(T_{amb}) + S(t_{cur}, d_1(t_{cur})) \cdot \frac{P_i \cdot \frac{e_i(f)}{p_i}}{\sum_{\tau_i} P_i \cdot \frac{e_i(f)}{p_i}}. \quad (24)$$

Based on the assigned idle time $I_i(t_{cur})$ and the remaining execution time $e_{left_i}(f)$, the scheduler splits τ_i into $m_i(I_i(t_{cur}))$ sub-tasks and alternates the processor to be idle for $\frac{I_i(t_{cur})}{m_i(I_i(t_{cur}))}$ units and task execution for $\frac{e_{left_i}(f)}{m_i(I_i(t_{cur}))}$ units.

Let's consider how to calculate slack $S(t_{cur}, d_1(t_{cur}))$. Our goal is to find the maximum amount of slack time, which may be available during the interval $[t_{cur}, d_1(t_{cur}))$, while guaranteeing 1) at least the minimum idle-time for all future jobs and 2) all future deadlines ($\geq t_{cur}$) to be met. Algorithm 1 presents our slack calculation method. At time t_{cur} , we look at the interval until the earliest absolute deadline $d_1(t_{cur})$ among all tasks and examine all tasks in reverse EDF order, i.e., latest deadline first (Line 4). Note that tasks are indexed in EDF order (i.e., for τ_i and τ_k where $i < k$, $d_i(t_{cur}) \leq d_k(t_{cur})$). We assume that future task invocations require the worst-case execution and minimum idle-times, and thus their utilization is $\sum_{\tau_i} \frac{e_i(f) + I_i^{min}(T_{amb})}{p_i}$ (Line 1). We try to defer as much execution/idling as possible beyond $d_1(t_{cur})$ and compute the minimum amount of execution/idling p that must execute before $d_1(t_{cur})$ in order to meet all future deadlines (Lines 5–8). This step is repeated for all tasks. To calculate p , we use the similar approach as in [32], [33]. Then, the slack is set to the remaining time slots except for p over the interval $[t_{cur}, d_1(t_{cur}))$ (Line 10). The underlying principle behind our slack calculation is that EDF will determine a feasible schedule if the utilization in Eq. (23) is ≤ 1.0 at any time [34].

Runtime complexity. At each invocation (either JOB_RELEASE, JOB_COMPLETION, or FREQ_UPDATE), our scheduling algorithm updates the slack by Algorithm 1 with the complexity of $O(n)$, where n is the number of tasks. Then, our algorithm allocates the slack to a job with the

Table II: Thermal parameters of i.MX6 processor [22], [28]

R ($^{\circ}\text{C}/\text{W}$)	C ($\text{J}/^{\circ}\text{C}$)	β_1 ($\text{mA}/^{\circ}\text{C}$)	β_0 (mA)	P_{max} (mW)
22	0.0454	0.435	611	3860

Table III: WCET and min/maximum periods

(s)	Angle	Bit	Table	Edge	FFT	PID
e_i	2.51	1.03	0.919	0.872	0.456	0.151
p_i^{min}, p_i^{max}	15, 30	6, 12	6, 12	5, 10	2.5, 5	1, 2

earliest deadline according to Eq. (24) with the complexity of $O(1)$. Thus, the total complexity is $O(n)$.

VII. EVALUATION

We have implemented and evaluated RT-TRM on a commercial embedded processor for automotive and infotainment applications. Our evaluation focuses on how it guarantees thermal and real-time constraints under various conditions.

Experimental setup. Our evaluation platform is i.MX6 [28] with ARM A9 supporting 3 discrete frequency levels (1GHz, 0.8GHz, 0.4GHz) and the corresponding voltage levels (1.25V, 1.15V, 0.95V). The chip is equipped with an on-chip thermal sensor with precision of 0.4°C . Table II specifies the power and thermal parameters of our target platform. We set the peak temperature constraint T_{max} to 60°C .⁷ For the purpose of demonstration, we use realistic automotive workloads obtained from MiBench [29], including *Angle-time Conversion*, *Bit Manipulation*, *Table Lookup*, *Edge Detection*, *FFT*, *PID*. The configuration of each workload is provided in Table III.⁸ We use real-time kernel [36] to periodically execute above benchmark applications. For idle-time scheduling, we generate a kernel idle thread to preempt task execution.

Handling ambient temperature variation. To illustrate how RT-TRM adapts to various environmental conditions to meet the thermal constraint, we conducted a set of experiments at different ambient temperatures (25, 30, and 35°C). Fig. 7a plots the real-time traces of processor temperature, frequency and task-rate. The task-rate is defined in Eq. (14) and normalized by the maximum rate. The results show that RT-TRM effectively regulates the processor temperature below T_{max} . At ambient temperature 25°C (dotted line), RT-TRM is shown to be able to maintain the 1GHz processor frequency and 91.5% of the maximum task-rate. At 30°C (grey line), the processor frequency is switched between 1 and 0.8GHz, and the task-rate is dynamically adjusted, achieving 82.6% of the maximum task-rate. At 35°C (solid line), the processor frequency had to be reduced at time around 150s to meet the thermal constraint, resulting in 65.6% of the maximum task-rate. We also look closer the results in Fig. 7a in a shorter time interval (0, 40] and present the execution behavior of the hottest task (bit manipulation) and its corresponding temperature variation under RT-TRM as shown in Fig. 8. In the figure, we observe that the processor temperature is increased

⁷According to the mean-time-to-failure (MTTF) model [35], the thermal constraint of 60°C can cover a typical vehicle warranty period of 10 years.

⁸Note that MiBench does not specify task period and execution time. We thus measure the worst-case execution time of each task in our experimental setup and synthetically assign a period range of each task proportional to the WCET as done similarly in [18].

Table IV: The number of preemptions and idle-time per job

	Preemption	Used idle-time (s)
Static minimum idle-time	8.77	0.221
Online idle-time scheduling	1.18	0.275

whenever the bit manipulation task executes. When the ambient temperature is 35°C , a job of the bit manipulation task is invoked every 12 seconds, while it is invoked every 8 seconds (6 seconds) when the ambient temperature is 30°C (25°C). RT-TRM can adaptively adjust the processor frequency and task periods under different ambient temperatures while meeting both thermal and timing requirements.

Handling different thermal constraints. Fig. 7b shows the results of RT-TRM under different thermal constraints (55, 60, and 65°C). Under the thermal constraint of 65°C , RT-TRM achieves a higher task-rate of 91% without reducing core frequency. Under the thermal constraint of 55°C , it achieves a lower task-rate of 64.6% by reducing core frequency to 0.8GHz in order to meet the thermal constraint. RT-TRM is shown to effectively control the processor temperature close to the thermal constraint and maximize the resource utilization.

Handling different power dissipations. Fig. 7c shows the results of RT-TRM for different power dissipation workloads under the thermal constraint of 60°C . For the low power dissipation workload, RT-TRM achieves a higher task-rate of 88.9% with the maximum core frequency. For the high power dissipation workload, it dynamically adjusts task-rate in order to meet the thermal constraint, achieving a task rate of 79.3%.

Effect of online slack usage. We also analyze the effect of slack usage on online idle-time scheduling. During the above-mentioned experiment, we measure the total idle-time and the number of preemptions per job, as shown in Table IV. Our online idle-time scheduling algorithm can assign more idle-time by 0.054s by efficiently utilizing runtime slack, and thus reduce the number of preemptions by 7.4x, compared to the static minimum idle-time allocation method. We observe that a small amount of additional idle-time can dramatically reduce the number of preemptions. By reclaiming the available slack at runtime, RT-TRM uses 24.4% more idle-time to reduce 86.5% of preemptions without violating both thermal and timing constraints.

Performance evaluation. We have demonstrated how RT-TRM handles the dynamically changing ambient temperature and uses runtime slack to reduce the number of preemptions while satisfying thermal and timing constraints. We now focus on resource-efficiency and compare RT-TRM with two baseline approaches:

- EDF: static processor frequency and task period assignment under EDF⁹
- RT-MTC : dynamic processor frequency scaling using feedback control under EDF [10]
- RT-TRM: adaptive parameter assignment (§V) and online idle-time scheduling (§VI)

Under EDF, we consider two static parameter assignments: one assumes the average ambient temperature of 25°C (EDF-A); and the other assumes the worst-case ambient temperature of

⁹Parameters are assigned by solving the optimization problem Eq. (14)




Figure 7: Experimental results of RT-TRM showing processor temperature, frequency, task-rate traces under (a) different ambient temperatures, (b) thermal constraints and (c) power dissipations.

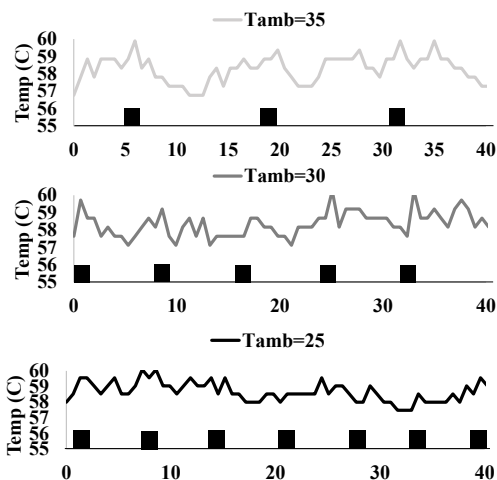


Figure 8: Job schedule of a task (*Bit manipulation*) and corresponding temperature variation by RT-TRM

35°C (EDF-W). Under RT-MTC, if the processor utilization exceeds the schedulable utilization by lowering the processor frequency, task periods are scaled to meet deadlines. We use two metrics: (1) the percentage of time during the thermal constraint is violated and (2) the task-rate. Higher task-rate indicates higher resource-efficiency.

Fig. 9 compares the processor temperature, frequency and task-rate for three different thermal management schemes. EDF-A assigns the processor frequency of 1GHz and the task-rate of 100% whereas EDF-W assigns the frequency of 0.8GHz and the task-rate of 63.5% (Fig. 9a). Under EDF-A, the maximum temperature is 71.5°C, and thus the thermal constraint is violated for 76.7% of the time. Under EDF-W, on the other hand, the thermal constraint is satisfied for all the time with the maximum temperature of 59.1°C, but resources are severely under-utilized.

Under RT-MTC in Fig. 9b, when the processor temperature hits the threshold at time 250s, the processor frequency is lowered to 0.8GHz. The temperature still exceeds the limit, so the frequency is lowered again to 0.4GHz at time 750s. Due to the reduced frequency to the lowest level, the task-rate for RT-MTC is reduced to 67.2%. While the feedback

controller regulates the temperature close to the set point, it violates the thermal constraint for 3% of the time with the maximum temperature of 60.5 °C.

Fig. 9c shows that RT-TRM maintains the maximum processor frequency for most of the time by adaptively adjusting the task periods, achieving the task-rate of 79.4% — an 18.2% improvement over RT-MTC. Note that parameters are adjusted every 3 seconds on average. By efficiently scheduling idle-time, RT-TRM can always meet the thermal constraint with the maximum temperature of 59.6 °C. The runtime overheads of parameter adjustment and idle-time scheduling are 27ms and 1ms, respectively.

VIII. DISCUSSION

Thus far, we have presented a task-level power/thermal model and developed RT-TRM that guarantees both thermal and timing constraints in the presence of dynamic ambient temperature variations. To demonstrate the importance of accounting for dynamic ambient temperature and different power dissipations, we have considered simple model and platform as an example — a task-level linear power model and a uniprocessor platform.

We now discuss the applicability of RT-TRM to general models and multi-core platforms.

Task-level power variations. We have assumed a task-level dynamic power model where power dissipation is constant across the jobs of a task and during the execution of a job. To guarantee the feasibility of RT-TRM without this assumption, we have used the maximum power dissipation among all jobs as task-level dynamic power.¹⁰

Linear power/thermal model. According to [25], [26], we have assumed that the leakage power P_{leak} (the thermal resistance R) has a linear (no) relation with the processor temperature. We have validated that such relations hold in a small temperature range (i.e., 20°C–35°C), but this may not hold in a wider temperature range. For example, the leakage power is known to increase exponentially as the temperature increases from 20°C to 120°C [37]. To apply RT-TRM in

¹⁰Characterization of precise job-level power dissipation is part of our future work.




Figure 9: Experimental results of different schemes showing the processor temperature, frequency, and task-rate traces

a wider temperature range, the leakage power and thermal resistance can be approximated by a piecewise linear model, i.e., the operating temperature range can be divided into multiple sub-ranges, each of which can be approximated by a linear model as shown in [37].

Multi-core platform. To apply RT-TRM for multi-core platforms, we can consider partitioned or global scheduling. In the case of partitioned scheduling, RT-TRM can be directly applied once a task-to-core assignment is made. Since the task-to-core assignment is known to be NP-hard [38], we can use well-known heuristics for the assignment. In the case of global scheduling, we need to extend our slack calculation in Alg. 1 to consider the concurrent execution of multiple tasks on a multi-core platform (which is part of our future work). Note that the calculation of the minimum idle-time for each task has nothing to do with a task schedule, and hence it can be directly applied for multi-core platforms. Besides, on multi-core platforms, tasks scheduled on a core could affect the temperature of its neighboring cores. For this situation, we need a new thermal model that can capture the thermal effect between neighboring cores. We also need to calculate the new idle-time.

IX. CONCLUSION

Emerging embedded real-time systems, such as connected cars and smartphones, pose new challenges in meeting the timing constraints under the processors' thermal constraints. Such a system should consider a new *dynamic computation power bound* in addition to the conventional schedulable utilization bound. To address this problem, we have developed a new thermal model that captures individual tasks' heat generations as their *activity factors*. We then developed two new mechanisms, *adaptive parameter assignment* and *online idle-time scheduling*. By tightly coupling the solutions of these two mechanisms, we can guarantee both thermal and timing constraints in the presence of dynamic ambient temperature variations. Our evaluation of RT-TRM on a realistic micro-controller using automotive benchmarks has demonstrated the validity of the proposed thermal model and the effectiveness of RT-TRM in meeting both real-time and thermal constraints.

REFERENCES


- [1] J. Kong, S. W. Chung, and K. Skadron, "Recent thermal management techniques for microprocessors," *ACM Computing Surveys (CSUR)*, vol. 44, no. 3, p. 13, 2012.
- [2] J.-J. Chen, S. Wang, and L. Thiele, "Proactive speed scheduling for real-time tasks under thermal constraints," in *RTAS*, 2009.
- [3] N. Fisher, J.-J. Chen, S. Wang, and L. Thiele, "Thermal-Aware Global Real-Time Scheduling on Multicore Systems," in *RTAS*, 2009.
- [4] P. Kumar and L. Thiele, "Cool shapers: Shaping real-time tasks for improved thermal guarantees," in *DAC*, 2011.
- [5] T. Chantem, X. S. Hu, and R. P. Dick, "Temperature-aware scheduling and assignment for hard real-time applications on mpsoes," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2011.
- [6] P. Kumar and L. Thiele, "System-level power and timing variability characterization to compute thermal guarantees," in *CODES+ISSS*, 2011.
- [7] L. Schor, I. Bacivarov, H. Yang, and L. Thiele, "Worst-case temperature guarantees for real-time applications on multi-core systems," in *RTAS*, 2012.
- [8] R. Ahmed, P. Ramanathan, and K. K. Saluja, "Necessary and sufficient conditions for thermal schedulability of periodic real-time tasks," in *ECRTS*, 2014.
- [9] Y. Fu, N. Kottenstette, Y. Chen, C. Lu, X. D. Koutsoukos, and H. Wang, "Feedback Thermal Control for Real-time Systems," in *RTAS*, 2010.
- [10] Y. Fu, N. Kottenstette, C. Lu, and X. D. Koutsoukos, "Feedback thermal control of real-time systems on multicore processors," in *EMSOFT*, 2012.
- [11] H. Aydin, R. Melhem, D. Mossé, and P. Mejía-Alvarez, "Power-aware scheduling for periodic real-time tasks," *IEEE Transactions on Computers*, vol. 53, no. 5, pp. 584–600, 2004.
- [12] Y. Zhu and F. Mueller, "Feedback edf scheduling exploiting dynamic voltage scaling," in *RTAS*, 2004, pp. 84–93.
- [13] S. Wang and R. Bettati, "Delay analysis in temperature-constrained hard real-time systems with general task arrivals," in *RTSS*, 2006.
- [14] J.-J. Chen, C.-M. Hung, and T.-W. Kuo, "On the minimization of the instantaneous temperature for periodic real-time tasks," in *RTAS*, 2007.
- [15] P. Bailis, V. J. Reddi, S. Gandhi, D. Brooks, and M. Seltzer, "Dimetrodon: processor-level preventive thermal management via idle cycle injection," in *DAC*, 2011.
- [16] R. Jayaseelan and T. Mitra, "Temperature aware task sequencing and voltage scaling," in *ICCAD*, 2008.
- [17] R. Ahmed, P. Ramanathan, and K. K. Saluja, "Temperature minimization using power redistribution in embedded systems," in *VLSI Design*, 2014.
- [18] S. Zhang and K. Chatha, "Thermal aware task sequencing on embedded processors," in *DAC*, 2010.
- [19] H. Huang, G. Quan, J. Fan, and M. Qiu, "Throughput Maximization for Periodic Real-Time Systems Under the Maximal Temperature Constraint," in *DAC*, 2011.
- [20] R. Ahmed, P. Ramanathan, and K. K. Saluja, "On thermal utilization of periodic task sets in uni-processor systems," in *RTCSA*, 2013.
- [21] Y. Ma, T. Chantem, X. S. Hu, R. P. Dick, and N. Dame, "Improving Lifetime of Multicore Soft Real-Time Systems through Global Utilization Control," in *GLVLSI*, 2015.
- [22] Freescale, "i.mx 6dual/6quad power consumption measurement: Table 1. vddarm, vddsoc, vddpu voltage levels," 2012. [Online]. Available: https://cache.freescale.com/files/32bit/doc/app_note/AN4509.pdf

- [23] D. Seto, J. P. Lehoczky, L. Sha, and K. G. Shin, "Trade-off analysis of real-time control performance and schedulability," *Real-Time Systems*, vol. 21, no. 3, pp. 199–217, 2001.
- [24] T. Cucinotta, L. Palopoli, L. Abeni, D. Faggioli, and G. Lipari, "On the integration of application level and resource level qos control for real-time applications," *IEEE Transactions on Industrial Informatics*, 2010.
- [25] Y. Liu, R. P. Dick, L. Shang, and H. Yang, "Accurate temperature dependent integrated circuit leakage power estimation is easy," in *DATE*, 2007.
- [26] T. L. Bergman, *Introduction to heat transfer*. John Wiley & Sons, 2011.
- [27] R. W. Johnson, J. L. Evans, P. Jacobsen, J. R. Thompson, and M. Christopher, "The changing automotive environment: High-temperature electronics," *IEEE Transactions on Electronics Packaging Manufacturing*, vol. 27, 2004.
- [28] Freescale, "Technical data applications processors: Table 5. fcpbga package thermal resistance," 2012.
- [29] M. R. Guthaus, J. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge, and R. B. Brown, "MiBench: A free, commercially representative embedded benchmark suite," in *Workshop on Workload Characterization*, 2001.
- [30] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of the Association for Computing Machinery*, vol. 20, no. 1, pp. 46–61, 1973.
- [31] N. Megiddo, "Linear Programming in Linear Time When the Dimension Is Fixed," *Journal of the ACM*, 1984.
- [32] P. Pillai and K. Shin, "Real-time dynamic voltage scaling for low-power embedded operating systems," in *SOSP*, 2001.
- [33] H. S. Chwa, K. G. Shin, H. Baek, and J. Lee, "Physical-state-aware dynamic slack management for mixed-criticality systems," in *RTAS*, 2018.
- [34] S. A. Brandt, S. Banachowski, C. Lin, and T. Bisson, "Dynamic integrated scheduling of hard real-time, soft real-time and non-real-time processes," in *RTSS*, 2003.
- [35] T. Chantem, Y. Xiang, X. Hu, and R. Dick, "Enhancing multicore reliability through wear compensation in online assignment and scheduling," in *DATE*, 2013.
- [36] S. Oikawa and R. Rajkumar, "Linux rk: A portable resource kernel in linux," in *RTSS*, 1998.
- [37] Y. Liu and H. Yang, "Temperature-aware leakage estimation using piecewise linear power models," *IEICE transactions on electronics*, vol. 93, no. 12, pp. 1679–1691, 2010.
- [38] K. W. Tindell, A. Burns, and A. J. Wellings, "Allocating hard real-time tasks: an np-hard problem made easy," *Real-Time Systems*, vol. 4, no. 2, pp. 145–165, 1992.



Kang G. Shin (LF'92) is the Kevin & Nancy O'Connor Professor of Computer Science in the Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor. His current research focuses on QoS-sensitive computing and networking as well as on embedded real-time and cyber-physical systems. He has supervised the completion of 82 PhDs, and authored/coauthored more than 900 technical articles, one a textbook and more than 30 patents or invention disclosures, and received numerous best paper awards, including the

Best Paper Awards from the 2011 ACM International Conference on Mobile Computing and Networking (MobiCom'11), the 2011 IEEE International Conference on Autonomic Computing, the 2010 and 2000 USENIX Annual Technical Conferences, as well as the 2003 IEEE Communications Society William R. Bennett Prize Paper Award and the 1987 Outstanding IEEE Transactions of Automatic Control Paper Award. He has also received several institutional awards, including the Research Excellence Award in 1989, Outstanding Achievement Award in 1999, Distinguished Faculty Achievement Award in 2001, and Stephen Attwood Award in 2004 from The University of Michigan (the highest honor bestowed to Michigan Engineering faculty); a Distinguished Alumni Award of the College of Engineering, Seoul National University in 2002; 2003 IEEE RTC Technical Achievement Award; and 2006 Ho-Am Prize in Engineering (the highest honor bestowed to Korean-origin engineers). He has chaired several major conferences, including 2009 ACM MobiCom, 2008 IEEE SECON, 2005 ACM/USENIX MobiSys, 2000 IEEE RTAS, and 1987 IEEE RTSS. He is the fellow of both IEEE and ACM, and served on editorial boards, including IEEE TPDS and ACM Transactions on Embedded Systems. He has also served or is serving on numerous government committees, such as the US NSF Cyber-Physical Systems Executive Committee and the Korean Government R&D Strategy Advisory Committee. He has also co-founded a couple of startups. He received his BS in Electronics Engineering from Seoul National University (1970); MS in Electrical Engineering from Cornell University (1976); PhD in Electrical Engineering from Cornell University (1978)



Youngmoon Lee (S'17) received his BS in Electrical and Computer Engineering from Seoul National University (2014); MS in Computer Science and Engineering from University of Michigan (2016). He is a PhD candidate at the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, USA. His research focuses on embedded and cyber-physical systems and networked systems.



Shige Wang (SM'11) received his BE and MS degrees in computer engineering from the Northeastern University, Shenyang, China, and the PhD degree in computer science and engineering from the University of Michigan, Ann Arbor, Michigan in 2004. He is currently with General Motors Global Research and Development Center in Warren, Michigan. He has been leading research projects on model-based software engineering, multicore and parallel computing for embedded vehicle control system. His current research focuses on embedded software architecture and analysis of embedded image processing and automated driving. He is a senior member of the IEEE.



Hoon Sung Chwa (M'16) received the BS, MS, and PhD degrees from Korea Advanced Institute of Science and Technology, all in computer science, in 2009, 2011, and 2016, respectively. He is an assistant professor in the Department of Information and Communication Engineering, DGIST, Korea. He has been a research fellow in the Department of Electrical Engineering and Computer Science, the University of Michigan, Ann Arbor, USA, until 2018. His research interests include system design and analysis with timing guarantees and resource management in real-time embedded systems and cyber-physical systems. He won two best paper awards from the 33rd IEEE Real-Time Systems Symposium (RTSS) in 2012 and from the IEEE International Conference on Cyber-Physical Systems, Networks, and Applications (CPSNA) in 2014.