Leakage Conscious DVS Scheduling for Peak Temperature Minimization

Vivek Chaturvedi Gang Quan
Electrical & Computer Engineering Department
Florida International University
Miami, Fl, 33174
e-mail: vchaturv@fiu.edu, gaquan@fiu.edu

Abstract— In this paper, we incorporate the dependencies among the leakage, the temperature and the supply voltage into the theoretical analysis and explore the fundamental characteristics on how to employ dynamic voltage scaling (DVS) to reduce the peak operating temperature. We find that, for a specific interval, a real-time schedule using the lowest constant speed is not necessarily the optimal choice any more in minimizing the peak temperature. We identify the scenarios when a schedule using two different speeds can outperform the one using the constant speed. In addition, we find that the constant speed schedule is still the optimal one to minimize the peak temperature at the temperature *stable status* when scheduling a periodic task set. We formulate our conclusions into several theorems with formal proofs.

I. Introduction

As semiconductor technology continues to scale down, the chip temperature increases dramatically due to the exponentially growing power consumption. The escalating heat has directly led to high packaging/cooling costs and also potentially degrades the performance, life span, and reliability of computing systems [1, 2]. Left unchecked, the thermal issue will severely handicap the computing systems in the near future. The severity of the problem is further highlight by the Intel's acknowledgement of hitting a "thermal wall".

For the past several decades, a closely related problem, i.e. the power/energy reduction problem, has been researched extensively. While lowering power/energy consumption does help to reduce the heat generation, the power aware computing problem and the thermal aware computing problem are distinctly different, as exhibited in many previous work (e.g. [3, 4]). As a result, existing power/energy reduction techniques cannot be readily applied to address the thermal issues most effectively.

When studying the thermal aware dynamic voltage scaling problem, it is imperative to take the dependency among the leakage, the temperature, and the supply voltage into considerations. It is a well known fact [5] that the leakage power increases with the temperature and the supply voltage. Due to down-scaling of technology, the leakage power consumption is becoming comparable or even surpasses the dynamic power consumption [6]. Therefore, a DVS technique cannot be effective without considering the relations of the leakage, the temperature, and the supply voltage.

The goal of this paper is to explore the characteristics and guidelines that can be exploited in development of effective thermal aware design techniques. Since the thermal management problem is closely related to the power reduction problem, we started our research by investigating how effective the basic principles for energy reduction can be, when applied for thermal management. We begin with the two well-known principles [7, 8], for dynamic energy reduction,

- *Principle 1*: Using the lowest constant speed is the schedule that consumes the minimum dynamic energy;
- *Principle 2*: If a single lowest constant speed is not available, then using the two closest neighboring speeds is the optimal solution in dynamic energy reduction.

The question becomes: when considering the complex relationship among the leakage power, the temperature, and the supply voltage, is it still true that a real-time schedule employing the lowest constant speed will lead to the lowest peak temperature within the interval? We find that, for a specific workload and interval, the schedule that uses the lowest constant speed is not necessarily optimal anymore in reducing the peak temperature. We study the scenarios when a schedule that uses two different speeds may in fact lead to lower peak temperature. We also find that principles similar to the two listed above do exist to minimize the peak temperature during the temperature stable status when scheduling a periodic task. We formulate our observations into several theorems with formal proofs. The significance of this paper is that it has provided a number of fundamental principles for the development of effective DVS techniques in maximum temperature reduction.

The rest of the paper is organized as follows. System models used in this paper are described in Section II. Section III presents our empirical research results to motivate our research. Fundamental principles are formulated and proved in Section IV and V. Section VI introduces the related work, and Section VII concludes the paper.

II. System models

We first introduce the system models used in this paper. The *real-time system* we consider consists of a number of real-time tasks with the same period. We can thus simplify this model by assuming that the real-time system has only one periodic task. The period of the task is denoted as p and its worst-case

978-1-4244-7516-2/11/$26.00 ©2011 IEEE

execution cycle as c. We assume that the deadline of the task equals its period.

The *thermal model* used in this paper is the same as that in the similar research (e.g. [9, 10, 11]). Specifically,

$$\frac{dT(t)}{dt} = aP(t) - bT(t), \tag{1}$$

where $T(t)$ is the temperature at time t, $P(t)$ denotes the power consumption (in $Watt$) at time t, $a = 1/C_{th}$ and $b = 1/R_{th}C_{th}$, and R_{th}, C_{th} denote the thermal resistance (in $J/^oC$) and thermal capacitance (in $Watt/^oC$).

The processor can run in n different modes, with each mode as (v_i, f_i), $i = 0, 1, ..., n-1$. where v_i is the supply voltage and f_i is the working frequency in mode i. We assume that $v_i < v_j$, if $i < j$. The power consumption of the processor consists of the dynamic power P_{dyn} and the leakage power P_{leak}. P_{leak} changes with both temperature and supply voltage. Specifically, as shown in [12], the leakage current for a single transistor I_{leak} can be formulated as follows:

$$I_{leak} = I_s \cdot (\mathcal{A} \cdot T^2 \cdot e^{((\alpha \cdot V_{dd} + \beta)/T)} + \mathcal{B} \cdot e^{(\gamma \cdot V_{dd} + \delta)}) \tag{2}$$

where I_s is the leakage current at certain reference temperature and supply voltage, T is the temperature, $\mathcal{A}, \mathcal{B}, \alpha, \beta, \gamma, \delta$ are empirically determined constants. Liu et al. [13] found that using linear approximation method to model the leakage/temperature dependence is reasonably accurate and several approaches (e.g. [10, 11, 14]) adopt this method to simplify the leakage temperature dependency. In our work, we adopt a leakage power model as follows:

$$P_{leak}(k) = (C_0(k))v_k + C_1 T, \tag{3}$$

where $C_0(k)$ and C_1 are constants. And the total power consumption at processor mode k is formulated as

$$P(k) = C_0(k)v_k + C_1 \cdot T + C_2 v_k^3. \tag{4}$$

From equation (4) and (1), when a processor running in mode k for interval $[t_0, t_e]$, let the starting temperature be T_0, then solving equation (1), the ending temperature can be formulated as below:

$$\begin{aligned} T_e &= \frac{A(k)}{B} + (T_0 - \frac{A(k)}{B})e^{-B(t_e - t_0)} \\ &= G(k) + (T_0 - G(k))e^{-B(t_e - t_0)}. \end{aligned} \tag{5}$$

where

$$A(k) = a(C_0(k)v_k + C_2 v_k^3), \tag{6}$$

$$B = b - aC_1, \tag{7}$$

and

$$G(k) = \frac{A(k)}{B}. \tag{8}$$

Equation (5)-(8) play a critical role in our analytical analysis.

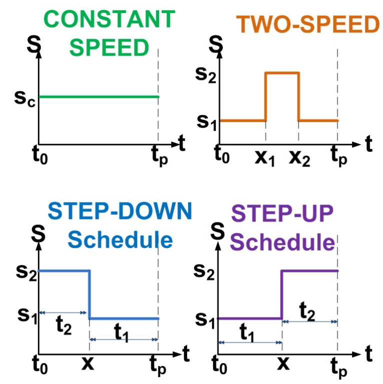

Fig. 1. The constant-speed and the two-speed schedule.

III. The Empirical Studies

Considering that the leakage power changes with both temperature and supply voltage, is the constant speed schedule still the optimal choice in minimizing the peak temperature within a specific interval? Before we drew any conclusions, we first launched a number of empirical studies to obtain some intuitions. For the ease of our presentation, we first formally define several representative real-time schedules as follows.

Definition 1 *The* constant-speed schedule $\hat{S}(S_c)$ *within an interval* $[t_0, t_p]$ *is the schedule that employs the lowest constant processor speed* S_c *to complete the workload within the interval.*

Definition 2 *A* two-speed schedule $\hat{S}(S_1, S_2)$ *within an interval* $[t_0, t_p]$ *is the schedule that uses speed* S_2 *during interval* $[x_1, x_2](t_0 \le x_1 < x_2 \le t_p)$*, and* S_1 *for the rest of the interval* $[t_0, t_p]$*, or vice versa to complete the workload with the interval.*

Definition 3 *A* step-up schedule $\hat{S}(S_1, S_2)$ *within interval* $[t_0, t_p]$ *is a two-speed schedule that uses* S_1 *during interval* $[t_0, x]$*, and uses* S_2 *in interval* $[x, t_p]$ *with* $S_1 \le S_2$*.*

Definition 4 *A* step-down schedule $\hat{S}(S_1, S_2)$ *within interval* $[t_0, t_p]$ *is a two-speed schedule that uses* S_2 *during interval* $[t_0, x]$*, and uses* S_1 *in interval* $[x, t_p]$ *with* $S_1 \le S_2$*.*

Figure 1 shows an example of different speed schedules defined. In what follows, we present two sets of results to obtain some intuitions on the applicability of the above two power reduction principles in context of peak temperature minimization. We also present some empirical results on the characteristics of the parameter, i.e. $G(k)$, which helps to justify the assumption made in our theorems.

Empirical study 1 First, we wanted to verify if two principles listed before are still valid in terms of the peak temperature minimization for a given interval.

We constructed our processor and power model similar to one shown in [11], based on the 65nm technology and with the conventional air cooling option of $Rth = 0.8K/W$, $Cth = 340J/K$ [1]. The ambient temperature was set to 25^oC. We selected three available processor speeds with corresponding supply voltages as 0.75V, 0.80V, and 0.85V respectively. Three different types of schedules, i.e. the constant-speed schedule, the step-down and the step-up schedules that can complete same workload were constructed. We then varied the interval length, collected the highest temperature by each schedule within the interval, and plotted in Figure 2(a).

As can be seen from Figure 2(a), while the maximum temperature using the step-up schedule is always higher than that by the constant-speed schedule, the peak temperature by the step-down schedule can in fact be lower sometimes. For instance, when period is equal to 700 seconds, the peak temperature of the step-down, the constant and the step-up speed schedules are $33.45^{o}C$, $34.04^{o}C$ and $35.16^{o}C$ respectively. This result clearly contradicts the conclusion made in [15] that the constant-speed schedule is the optimal schedule in terms of minimizing the peak temperature with a given interval. In the meantime, we can also observe that the peak temperature by the step-up schedule is indeed consistently higher than that of the constant-speed schedule.

We further studied if using the two closest neighboring speeds is the best choice in terms of peak temperature reduction within a given interval. Two step-down speed schedules *Sa* and *Sb* were constructed. *Sa* uses the speed corresponding to supply voltages 0.75V & 0.85V for low and high speed respectively, and *Sb* uses speeds corresponding to 0.75V & 0.90V. When comparing their peak temperatures, as shown in Figure 2(b), step-down schedule *Sa* is not necessarily always better than *Sb*. When the period is set to 700 seconds, the peak temperature for *Sa* is $33.45^{o}C$, while for *Sb* it is $33.10^{o}C$. This result seems to also imply that the second principle is not valid in terms of the peak temperature reduction either.

Empirical study 2 We next studied the maximum temperature minimization problem when the processor temperature reaches its stable status.

We constructed the three different schedules, i.e., the constant-speed schedule, the step-down and the step-up schedules the same as above, and ran each schedule not one time but periodically until the temperature became stable. We then varied the periods, collected the maximum temperature for each case and plotted in Figure 3(a). From this figure, we can clearly see that the constant-speed schedule always lead to the lowest peak temperature in our experiment, and the peak temperatures by the step-up and step-down schedule eventually become the same.

In addition, if we constructed two step-down schedules *Sa* and *Sb* as above and ran them periodically. From Figure 3(b), we can see that the step-down schedule *Sa* using the two closest neighboring speeds is always better than *Sb*. Our empirical results here seem to suggest that the two principles listed before may be valid in terms of minimizing peak temperature when the processor reaches the stable status.

Empirical study 3 To formulate a theorem formally and conduct analytical analysis based on temperature dynamics in equation (5), it is highly desirable that the characteristics of $G(k)$ is known. However, since $G(k)$ is determined essentially by the curve-fitting constants C_0 and C_1, it is difficult to analytically study its properties. Therefore we study its attributes empirically.

Figure 4 plot the characteristics of function G_k under different supply voltages based on the conventional air cooling option. Similar results were obtained with the conventional water cooling option (we omitted the results due to the page limit).

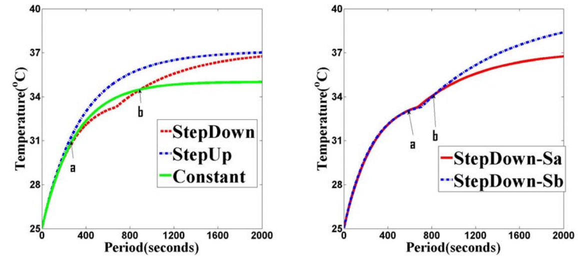

(a) (b)

Fig. 2. Peak temperature within a given interval

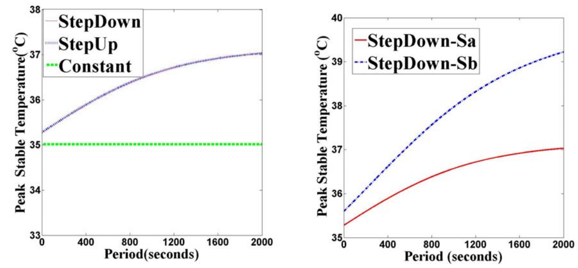

(a) (b)

Fig. 3. Peak temperature at stable status

As illustrated in this figure, we can clearly see that, function G_k is a positive, monotonic increasing, and convex function of the supply voltage. Also, from equation (7), we can see that it is necessary that $B > 0$, or the temperature will run away otherwise. Therefore, in what follows, unless otherwise specified, we assume that (*i*) G_k (or $G(v_k)$) is a *positive*, *monotonic increasing*, and *convex function* of k (or v_k), respectively and (*ii*) $B > 0$.

IV. Peak temperature minimization within a specified interval

It is clear from above empirical results that a constant speed schedule is no longer the optimal method for peak temperature reduction. In this section, we take a closer look at the conclusions drawn from our motivational examples and formulate several theorems. Specifically Theorem 1 characterizes the peak temperature obtained using the step-up schedule, within a given interval.

Theorem 1 *Given two available processor speeds S_1 and S_2, let the processor temperature equal to its ambient temperature, then the step-up schedule ($\hat{S}(S1,S2)$) has the highest peak temperature, among all two-speed schedules within a given interval.*

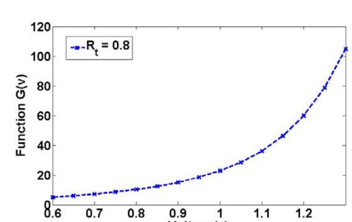

Fig. 4. Function $G(v)$ with different supply voltage.

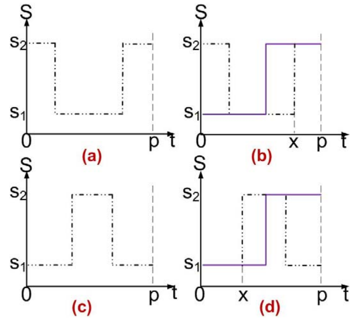

Fig. 5. Comparison of peak temperature of two-speed and step-up schedule.

Proof sketch: Consider the step-up and step-down schedules in Figure 1. Without loss of generality, we assume the initial temperature to be zero. Let T_u be the peak temperature of step-up schedule, which always occur at the end of the interval, then

$$T_u = G_2(1-e^{-Bt_2}) + G_1(1-e^{-Bt_1})e^{-Bt_2} \tag{9}$$

Let T_d be peak temperature for step-down which can occur either at switching point x or at end of the interval i.e. t_p in Figure 1.

- Case 1: T_d at x, or $T_d = G_2(1-e^{-Bt_2})$. Then $T_d \leq T_u$ since $G(k)$ and B are positive functions and $G_2(1-e^{-Bt_2}) \leq G_2(1-e^{-Bt_2}) + G_1(1-e^{-Bt_1})e^{-Bt_2}$.
- Case 2: T_d at t_p, or $T_d = G_1(1-e^{-Bt_1}) + G_2(1-e^{-Bt_2})e^{-Bt_1}$. Then $T_d \leq T_u$ since $G_1 \leq G_2$ and $G_1(1-e^{-Bt_1})(1-e^{-Bt_2}) \leq G_2(1-e^{-Bt_2})(1-e^{-Bt_1})$

We now compare a step-up schedule with an arbitrary two-speed schedule that uses the same two speeds. We first compare it with the schedule shown in Figure 5(a). Let T_f be the maximum temperature within interval $[0,t_p]$. In Figure 5(b), let $T_u(x)$ and $T_f(x)$ denote the peak temperature within the interval [0,x] by the step-up schedule and the two-speed schedule, respectively. Then $T_f(x) \leq T_u(x)$, since we have proved that the step-up schedule always incurs a higher peak temperature than that of a step-down schedule. Since both schedules use the same speed within interval $[x,p]$, we conclude that $T_f \leq T_u$. Similar conclusions can be proved for the two speed schedule shown in Figure 5(c). □

It is worthy of mentioning that the conclusion is true only when the processor starts from the ambient temperature. The effects of higher initial temperature on the schedules will be considered in our future work.

Furthermore, our empirical study shows that the conclusion [15] that the constant-speed schedule is the optimal choice in terms of peak temperature minimization within a given interval is flawed. With a closer look at the proof presented in [15], we found that the proof shows only the constant-speed schedule is always better than a step-up schedule, which conforms to our empirical research results, but not an arbitrary two-speed schedule and is thus flawed. Even though our empirical results shows that in most cases the constant-speed schedule is a better choice, it can be inferior to a step-down schedule sometimes. Specifically, Theorem 2 describes some conditions such that a constant-speed becomes worse than a step-down schedule in terms of peak temperature minimization within an interval. (The proof is omitted due to the page limit.)

Theorem 2 *Given a constant-speed schedule $\hat{S}(S_1)$ and a step-down schedule $\hat{S}(S_0,S_2)$ for a hard real-time job J and let $T_m(\hat{S}(S_1))$ and $T_m(\hat{S}(S_0,S_2))$ represent the peak temperature by $\hat{S}(S_1)$ and $\hat{S}(S_0,S_2)$ within the interval [0,p] and let the processor starts from the ambient temperature, then*

$$T_m(\hat{S}(S_1)) \geq T_m(\hat{S}(S_0,S_2)) \tag{10}$$

if and only if,

- $\frac{1}{B}ln(\frac{G_2}{G_2-G_0}) \leq x \leq \frac{1}{B}ln(\frac{G_2}{G_2-G_1(1-e^{-B_1p})})$;

or,

- $x < min(\frac{1}{B}ln(\frac{G_2}{G_2-G_0}), p - \frac{1}{B}ln(\frac{G_2-G_0}{(G_1-G_0)+(G_2-G_1)e^{-Bp}}))$.

where $S_1p = S_2x + S_0(p-x)$, and B and G_i are defined in equation (6) and (8), respectively.

While the equations in Theorem 2 seem complex, they identify the scenarios when a constant speed or a step-down schedule should be applied to better reduce the peak temperature. From Theorem 2, once the workload and processor characteristics are given, we can immediately decide if the corresponding constant-speed schedule or a step-down schedule should be used to best minimize the peak temperature within the given interval.

V. Peak temperature minimization at the stable state

We next investigate the validity of applying *Principle* 1 and 2 in the context of minimizing peak temperature when scheduling a periodic task set. When running a real-time task set periodically, unless the processor temperature "runs away" [12], the processor temperature is eventually stabled. The stable status is defined as below.

Definition 5 *[11] When running a periodic task with period p, the temperature at the processor is called to be* stable *if for a given threshold, i.e.* $0 < \varepsilon << 1$,

$$|T((n+1)p) - T(np)| < \varepsilon, \tag{11}$$

where $n \geq 0, n \in Z$, and $T(t)$ is the temperature at t.

We first present two theorems that act as the basis in formulating the key principles of peak temperature minimization when the processor temperature becomes stable. (The proofs are omitted due to page limitation).

Theorem 3 *Given a hard real-time periodic task τ, the maximum temperature when the processor temperature reaches its stable status does not depend upon the initial temperature.*

Theorem 4 *Given a real-time periodic task τ and two processor speeds, then the maximum temperature at the stable status with any two-speed schedule using the same speeds are the same.*

Based on the conclusions from Theorem 3 and 4, we can now formulate an important theorem for the problem of scheduling hard real-time periodic task, with the goal of peak temperature minimization.

Theorem 5 *Given a real-time periodic task τ, the maximum temperature at the stable state is minimized when running τ using the lowest constant-speed.*

Proof sketch: Let T_c^∞ and T_u^∞ denote the maximum stable temperature for the constant-speed (S_1) and the step-up schedule ($S_0 < S_2$) respectively. Without loss of generality, we can assume $p = 1$. Also from the conclusion of Theorem 3, we know the stable temperature does not depend on the initial temperature, therefore we assume the initial temperature to be zero. Then we have

$$T_c^\infty = \frac{G_1(1-e^{-B})}{1-e^{-B}} = G_1 \tag{12}$$

$$T_u^\infty = \frac{G_2(1-e^{-B(1-x)}) + G_0(1-e^{-Bx})e^{-B(1-x)}}{1-e^{-(B(1-x)+Bx)}} \tag{13}$$

To show that $T_c^\infty \leq T_u^\infty$, we only need to show that

$$G_1 \leq kG_0 + (1-k)G_2, \tag{14}$$

where

$$k = \frac{e^{-B(1-x)} - e^{-B}}{1-e^{-B}}, 1-k = \frac{1-e^{-B(1-x)}}{1-e^{-B}}. \tag{15}$$

Since

$$S_1 = S_0x + S_2(1-x), \tag{16}$$

and $B > 0$ and G_i is a convex function, we have

$$G_1 \leq xG_0 + (1-x)G_2. \tag{17}$$

Therefore, to show that equation (14) holds, we only need to show that

$$xG_0 + (1-x)G_2 \leq kG_0 + (1-k)G_2, \tag{18}$$

or

$$(G_0 - G_2)(x-k) \leq 0. \tag{19}$$

As $G_0 \leq G_2$, we only need to prove that

$$x \geq k = 1 - \frac{1-e^{-B(1-x)}}{1-e^{-B}}. \tag{20}$$

Or, equivalently,

$$\frac{1-e^{-B(1-x)}}{1-e^{-B}} \geq 1-x. \tag{21}$$

Now consider function

$$F(z) = \frac{1-e^{-Bz}}{1-e^{-B}} - z. \tag{22}$$

with $0 \leq z \leq 1$. We can readily show that function $F(z)$ is a concave function since $F''(z) < 0$. Note that the curve $F(z)$ passes two points, i.e. $(0,0)$ and $(1,0)$, as $F(0) = 0$ and $F(1) = 0$. Let $H(z)$ be the line that crosses these two points. Since $F(z)$ is concave, we have $F(z) \geq H(z) \geq 0$ for $0 \leq z \leq 1$.

We therefore prove that the constant speed schedule always outperforms a step-up periodic schedule in minimizing the peak temperature when the temperature reaches the stable status. In Theorem 4, we have already proved that at stable status the peak temperature of step-up and any other two-speed are the same. Hence we can immediately conclude that at the stable state the constant-speed outperforms any two-speed schedule in peak temperature reduction. □

Moreover, since there are only a small number of processor speeds available. A constant-speed schedule is not always available and two or more speeds have to be used. In that case, we show that a principle that is similar to *Principle 2* stated previously can also be established.

Theorem 6 *If a* two-speed schedule *is used for a hard real-time periodic task, then the one that uses the two closest neighboring speeds minimizes the maximum temperature at the stable state .*

Note that, even though Theorem 5 and Theorem 6 look very similar to the two basic principles that have been widely used for dynamic energy reduction, it does not necessarily imply that the existing energy reduction techniques can be readily applied for the purpose of peak temperature minimization. On the other hand, Theorem 1 to Theorem 6 present some fundamental guidelines in the development of new schedule techniques that can minimize the peak temperature.

VI. Related work

While there have been extensive researches published on power/energy reduction using DVS techniques, existing work (e.g [3, 4, 16, 17]) has clearly shown that the power/energy aware problem and the thermal aware problem have drastically different characteristics. An optimal technique in power/energy reduction is not an optimal technique with regard to the peak temperature minimization.

An increasing number of researches have been published on thermal aware real-time scheduling, for both single and multiple processor platforms. However, most of the papers (e.g. [4, 16, 17, 18, 19]) have not taken the dependencies among leakage, temperature, and supply voltage into considerations. A few researches incorporate the complex circuit level leakage model into system analysis ([20, 21]). However, due to the non-linear and high-order magnitude terms in equation, such a model or tool can be too complex and cumbersome to be used for more rigorous real-time analysis and scheduling technique development.

A number of researches (such as [10, 22, 9, 23] assume that the leakage current changes linearly of quadratically *only* with temperature. However, leakage varies not only with temperature but also supply voltage as well. Quan et al. [11] introduced a leakage/temperature model that is more practical. Based on this model, they developed three methods to check the feasibility for a real-time scheduling under the maximum temperature constraints. When scheduling a task within a given interval, Quan et al. claimed that [15] a constant speed is always the optimal in minimizing the peak temperature. Unfortunately, as shown in our empirical results and theoretical analysis, this conclusion is flawed. Chaturvedi et al. developed [24] a so-called "m-oscillating" scheduling method to minimize the peak temperature for a periodic task set. However, an m-oscillating schedule is no longer a two-speed schedule as defined in our paper.

VII. Summary

In deep sub-micron domain, the thermal management is becoming a critical issue in design of modern computing systems. Also the leakage is also becoming a prominent problem that needs to be addressed effectively.

In this paper, we incorporate the leakage/temperature/supply voltage dependency into the real-time scheduling analysis that aims at minimizing the peak temperature. We show that a constant-speed schedule is not always the optimal schedule in terms of peak temperature minimization for a given interval. We further show that for a given periodic task, the lowest constant-speed is the optimal schedule among all two-speed schedules to minimize the peak temperature at the stable state. If this constant-speed is not available, then the schedule that uses the two closest neighboring speeds is the best choice. These new findings and theorems form the basis for the future study of developing more effective power and thermal aware scheduling techniques for more complicated architectures and real-time systems.

Acknowledgments

We thank anonymous reviewers for their comments and suggestions which contributed significantly to improving quality of this paper. This work is supported in part by NSF under projects CNS-0545913, CNS-0917021, and CNS-1018108.

References

[1] K. Skadron, M. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," *ICSA*, pp. 2–13, 2003.

[2] L.-T. Yeh and R. C. Chu, *Thermal Management of Microelectronic Equipment: Heat Transfer Theory, Analysis Methods, and Design Practices*. New York, NY: ASME Press, 2002.

[3] K. Skadron, M. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware computer systems: opportunities and challenges," *IEEE Micro*, vol. 23, no. 6, pp. 52–61, 2003.

[4] N. Bansal, T. Kimbrel, and K. Pruhs, "Speed scaling to manage energy and temperature," *Journal of the ACM*, vol. 54, no. 1, pp. 1–39, 2007.

[5] J. Rabaey, A. Chandrakasan, and B. Nikolic, *Digital Integrated Circuits: A Design Perspective*. Prentice Hall, 2003.

[6] ITRS, *International Technology Roadmap for Semiconductors (2008 Edition)*. Austin, TX.: International SEMATECH, http://public.itrs.net/.

[7] F. Yao, A. Demers, and S. Shenker, "A scheduling model for reduced cpu energy," in *FOCS*, 1995, pp. 374–382.

[8] T. Ishihara and H. Yasuura, "Voltage scheduling problem for dynamically variable voltage processors," in *ISLPED*, August 1998, pp. 197–202.

[9] T. Chantem, X. S. Hu, and R. Dick, "Online work maximization under a peak temperature constraint," in *ISLPED*, 2009, pp. 105–110.

[10] J.-J. Chen, S. Wang, and L. Thiele, "Proactive speed scheduling for real-time tasks under thermal constraints," *RTAS*, vol. 0, pp. 141–150, 2009.

[11] G. Quan and Y. Zhang, "Leakage aware feasibility analysis for temperature-constrained hard real-time periodic tasks," *ECRTS*, pp. 207–216, 2009.

[12] W. Liao, L. He, and K. Lepak, "Temperature and supply voltage aware performance and power modeling at microarchitecture level," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, no. 7, pp. 1042 – 1053, 2005.

[13] Y. Liu, R. P. Dick, L. Shang, and H. Yang, "Accurate temperature-dependent integrated circuit leakage power estimation is easy," in *DATE*, 2007, pp. 1526–1531.

[14] G. Quan and V. Chaturvedi, "Feasibility analysis for temperature- constraint hard real-time periodic tasks," *IEEE Transaction on Industrial Informatics*, vol. 6, no. 3, pp. 329–339, 2010.

[15] G. Quan and S. Ren, "Leakage-aware real-time scheduling for maximal temperature minimization," *SIGBED Review*, vol. 7, p. 1, 2010.

[16] G. Quan, Y. Zhang, W. Wiles, and P. Pei, "Guaranteed scheduling for repetitive hard real-time tasks under the maximal temperature constraint," *ISSS+CODES*, 2008.

[17] S. Zhang and K. S. Chatha, "Approximation algorithm for the temperature-aware scheduling problem," in *ICCAD*, 2007, pp. 281–288.

[18] S. Wang and R. Bettati, "Reactive speed control in temperature-constrained real-time systems," *ECRTS*, pp. 161–170, 2006.

[19] R. Rao, S. Vrudhula, C. Chakrabarti, and N. Chang, "An optimal analytical solution for processor speed control with thermal constraints," *ISLPED*, 2006.

[20] L. Yuan, S. Leventhal, and G. Qu, "Temperature-aware leakage minimization technique for real-time systems," in *ICCAD*, 2006, pp. 761–764.

[21] M. Bao, A. Andrei, P. Eles, and Z. Peng, "Temperature-aware idle time distribution for energy optimization with dynamic voltage scaling," *Date*, 2010.

[22] N. Fisher, J.-J. Chen, S. Wang, and L. Thiele, "Thermal-aware global real-time scheduling on multicore systems," *RTAS*, vol. 0, pp. 131–140, 2009.

[23] C. Yang, J. Chen, L. Thiele, and T. Kuo, "Energy-efficient real-time task scheduling with temperature-dependent leakage," *Date*, 2010.

[24] V. Chaturvedi, H. Huang, and G. Quan, "Leakage aware scheduling on maximal temperature minimization for periodic hard real-time systems," *ICESS*, 2010.