

# Dynamic Thermal Management for Single and Multicore Processors Under Soft Thermal Constraints

Bing Shi, Yufu Zhang, Ankur Srivastava

Dept. of Electrical and Computer Engineering, University of Maryland, College Park, MD, U.S.A.

{bingshi, yufuzh, ankurs}@umd.edu

## ABSTRACT

In this paper, we investigate Dynamic Thermal Management (DTM) policies under soft thermal constraint that allows the thermal constraint to be violated for a user specified period. For single core processor, we develop analytical expression for the optimal frequency policy under the soft constraint such that maximal performance can be extracted. We extend this problem to multi-core processor and provide optimal frequency policy when all cores run at same frequency. We also present LP based approximated formulation that generates frequency policies where each core has separate frequency control and considers leakage power. We use frequency legalization to approximate the frequency into discrete values. Experimental results indicate that 10°C increase in core temperature for 100sec results in 13% performance gain for single core processor, and 30% performance gain for two-core processor. Without  $T_{max}$  constraint, the performance improves almost 100% for two-core processor.

## Categories and Subject Descriptors

B.7.2 [Integrated Circuits]: Design Aids; J.6 [Computer Applications]: Computer-Aided Engineering

## General Terms

Algorithm, Management, Performance

## Keywords

Dynamic thermal management, multi-core processor

## 1. INTRODUCTION

Loss of performance and reliability due to un-predictable thermal hotspots has become a major issue. Dynamic thermal management, where the chip operation is controlled during runtime for curtailing thermal emergencies has become an active topic of research. Thermal management can be achieved by controlling processor knobs such as core frequency and supply voltage, scheduling of tasks etc, which in effect, control the power dissipation in different parts of the chip. Many works have made significant contributions on DTM [1][2][3][4][7][8].

The thermal constraint used to test the thermal feasibility of any control policy is usually a ballpark figure. Most existing thermal management schemes assume this to be

an absolute hard constraint which cannot be violated under any circumstance. The impact of temperature on reliability and lifetime is an average phenomenon. If we persistently keep core temperature above the constraint, then the performance and reliability will certainly suffer. But sporadic violation of thermal constraint will have little or no impact. Moreover, in some mission critical systems, allocating as much processing power to a task might be more desirable than maintaining acceptable chip thermal state while missing task deadlines. In this paper, we investigate DTM policies that allow the thermal constraint to be violated for a user specified time period under the constraint that the processor be brought back into acceptable thermal state at the end of this period. Our techniques can also enforce an absolute maximum thermal constraint which can never be violated. Using our methods, the user can choose to violate the thermal constraint whenever he/she chooses to do so. Our objective here is not to refute the detrimental impact of temperature on reliability, but to provide the option of violating the thermal constraint, should the designer choose to do so. The choice of the degree and duration of this thermal violation is up to the designer and depends strongly on the criticality of the task deadline, the nature of the task and the cost to replace the system should there be a failure.

We investigate four problem instances:

1. Single Core: Given a time period  $t_f$ , we allow the temperature to rise above the regular thermal constraint in this period. We also give an absolute thermal constraint  $T_{max}$  which can never be violated. We develop analytical expressions for the optimal frequency policy such that maximal performance can be extracted. Also, the final temperature at the end of  $t_f$  is within acceptable thermal constraint. Our optimal policy uses only two distinct frequency and one shutdown state, thereby ensuring practical implementability.
2. Multi-core with same frequency: We extend the single core problem into multi-core case, in which all the cores run at the same frequency. We first do not consider the absolute thermal constraint  $T_{max}$ , and develop analytical expression for the optimal frequency policy which achieves the best performance while ensuring the final temperature requirement. In this formulation, the optimal frequency policy has only one frequency and one shutdown state.
3. Multi-core same frequency with  $T_{max}$ : We place absolute maximum temperature constraint  $T_{max}$  on the multi-core same frequency problem, and develop the frequency policy for this problem. In this policy, the frequency should change continuously, so we develop frequency legalization mechanism to approximate the continuous frequency policy into a set of discrete values. This ensures the practicality of our approach.
4. Multi-core multi-frequency with  $T_{max}$ : We generalize previous formulation to include independent frequency control and separate  $T_{max}$  for each core. **We also incorporate leakage into the power model.** We wish to assign a frequency policy for each core

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED'10, August 18–20, 2010, Austin, Texas, USA.

Copyright 2010 ACM 978-1-4503-0146-6/10/08 ...\$10.00.

such that optimal total performance is obtained and the core temperatures, while not violating their respective  $T_{max}$  constraints, are brought back into acceptable thermal levels at time  $t_f$ . Since a closed form analytical solution does not exist for this formulation, we approximate the optimization problem using linear programming. Hence, near optimal solution can be obtained efficiently. We also explore frequency legalization methods to ensure practicality.

The key novelty of our work is that we investigate optimal processor core frequency policies when temperature is a soft constraint. The analytical expressions and the LP based formulations ensure the practicality of our schemes. [7], [8] derive similar analytical expressions for single and multicore systems but their objective is to maximize performance under an absolute thermal constraint. Besides, the general concept of soft thermal constraints can be applied to sensor based DTM as well. The experimental results show that 10°C increase in the core temperature for 100 seconds results in 13% performance gain for single core processor. A similar thermal increase in a two-core processor results in 30% performance gain and 100% performance increase without  $T_{max}$  constraint.

The organization of this paper is as follow. In section 2, we explore the problem of placing soft thermal constraint on single core processor. In section 3, we extend this problem to multi-core processor assuming all cores have the same frequency. We generalize this formulation to include separate frequency control for each core in section 4. In section 5, we discuss the influence of leakage power. Section 6 illustrates frequency legalization strategy and the experimental result is given in section 7.

## 2. SINGLE CORE DTM

### 2.1 Single core thermal model

The thermal behavior of a single core chip can be modeled by an RC circuit (figure 1). In this thermal RC model, voltage represents the temperature, and current represents the power dissipation in the processor. The resistance represents a potential heat transfer path through the packaging, while the capacitors indicate the ability of the processor to store heat [2].




Figure 1: RC thermal model

Based on the RC thermal model in figure 1, we obtain the following relationship between core temperature  $T$  and power consumption  $P$ :

$$\frac{dT}{dt} = -\frac{T}{RC} + \frac{P}{C} \quad (1)$$

where  $R$  (K/W) is the thermal resistance, and  $C$  (J/K) is the thermal capacitance. The dynamic power consumption is  $P = kf$ , where  $f$  is the frequency and  $k$  is a constant. Here we ignore leakage power first and assume that the power dissipation  $P$  is linearly dependent on the frequency. Later on, we'll discuss the influence of leakage in section 5

Assuming  $T_m$  is the initial temperature, we obtain the core temperature  $T$ :

$$T = (T_m - Rkf)e^{-t/RC} + Rkf \quad (2)$$

For this RC model, the temperature can be maintained at  $T_m$ , if it is running at the natural frequency  $f_{nat} = T_m/Rk$ . If the processor is running at a frequency  $f < f_{nat}$ , the temperature will decrease. Otherwise, the temperature increases.

### 2.2 Soft thermal constraint for single-core processor ( $SCT_{max}$ )

Here we call this problem " $SCT_{max}$ " problem. Suppose  $T_m$  is the manufacturer specified thermal constraint. In the rest of the discussion, we assume that we start from the initial condition  $T(0) = T_m$ . Our formulation can be easily generalized to other steady state initial conditions as well. As we have discussed, instead of placing strict constraint on the maximum temperature  $T_m$ , we investigate the possibility of increasing the temperature and allow the temperature to violate this constraint for a while, as long as this period is within some bound. That is, starting from the steady state  $T_m$ , for some period  $t_f$ , say several minutes, we allow the temperature to increase above  $T_m$ , as long as at the end of  $t_f$ , the temperature goes back to the steady state  $T_m$ . Meanwhile, although we have the flexibility to violate the constraint  $T_m$ , it is undesirable for the temperature to increase significantly above  $T_m$ . So we impose an absolute constraint on maximum temperature  $T_{max}$  within  $[0, t_f]$ . This value can be selected based upon how much the user allows the temperature to increase beyond  $T_m$ . Besides, the core frequency should be always smaller than the maximum frequency  $f_{max}$ . Under these constraints, we would like to maximize the total frequency within this period. So we arrive at the following optimization problem:

$$\begin{aligned} \max \quad & \int_0^{t_f} f(t) dt \\ \text{s.t.} \quad & T(0) = T(t_f) = T_m, \quad T(t) \leq T_{max} \\ & \frac{dT}{dt} = -\frac{T}{RC} + \frac{kf(t)}{C} \\ & 0 \leq f(t) \leq f_{max} \end{aligned} \quad (3)$$

**Optimal Solution:** The optimal frequency policy here is (figure 2(a)): first, let the processor run at  $f_{max}$ , until the temperature reaches  $T_{max}$  (see figure 2(b)); then the processor runs at  $T_{max}/Rk$ , so that the temperature maintains at  $T_{max}$ ; finally, shut down the processor so that temperature goes back to  $T_m$ . The first switching point  $t_{s1}$  is the moment when temperature reaches  $T_{max}$ , while the second switching point  $t_{s2}$  is the moment when we should shut down the processor so that its temperature could reduce to  $T_m$  at time  $t_f$ :

$$\begin{aligned} t_{s1} &= RC \ln \frac{Rkf_{max} - T_m}{Rkf_{max} - T_{max}} & t_{s2} &= t_f - RC \ln \frac{T_{max}}{T_m} \end{aligned} \quad (4)$$


Figure 2: (a) Optimal frequency strategy for  $SCT_{max}$ , (b) Corresponding temperature

It is noteworthy that this optimal strategy requires only two distinct frequency  $f_{max}$  and  $T_{max}/Rk$ , plus one shutdown state, which can be easily realized in practice. Also, the analytical nature of our solution makes the determination of the switching point very easy.

## 3. MULTICORE SAME FREQUENCY

### 3.1 Multi-core processor thermal model

We use the thermal RC circuit model in [7] as our multi-core processor thermal model. Figure 3(a) shows the RC model of a four-core processor, each core can be represented by an RC pair, while  $R_p$  and  $C_p$  pair is analogous to the package, whose value is much larger than  $R$  and  $C$ .  $R_l$  is the lateral resistance, which represent the inter-core lateral heat flow. We first explore the problem in which all processors run at the same frequency under the same maximum temperature constraint  $T_{max}$ .



**Figure 3: (a) RC model of a 4-core multi-core processor, (b) Simplified model**

Compared with the vertical resistances  $R$  that connect the core to package, lateral resistance  $R_l$  between each core is much larger [7]. So the lateral heat flow is much smaller than the vertical heat flow. **Here we first use a simplified model that ignores lateral heat flow (figure 3(b)), and then add lateral resistance in the generalized model in section 4.** Based on this simplified multi-core RC thermal model, the relationship between the core temperature  $T_i$ , frequency  $f_i$  of each core, and package temperature  $T_p$  is in equation 5. (Here we still assume linear dependency between the power dissipation  $P$  and frequency  $f$ , and we assume there are  $n$  cores.)

$$\begin{aligned} \frac{dT_i}{dt} &= -\frac{(T_i - T_p)}{RC} + \frac{kf_i}{C}, \quad \forall \text{core}(i) \\ \frac{dT_p}{dt} &= -\frac{T_p}{R_p C_p} + \frac{\sum_{i=1}^n (T_i - T_p)}{RC_p} \end{aligned} \quad (5)$$

Starting from the steady state, which occurs when both  $T_i$  and  $T_p$  do not change, all the cores have the same initial temperature that  $T_i(0) = T_m$  and the initial package temperature is  $T_p(0) = \frac{nR_p}{R+nR_p}T_m$ , which represents the steady state.

### 3.2 Soft thermal constraint for multi-core with same frequency (MCSF)


In this version of problem, we assume all processors run at the same frequency. For the sake of simplicity, we also assume thermal steady state at  $t=0$  with  $T_i(0) = T_m$  and  $T_p(0) = \frac{nR_p}{R+nR_p}T_m$  (the core temperature is exactly the thermal constraint). Our methods can easily be generalized to any valid initial steady state as well. Besides, since all processors have the same frequency and initial temperature, they will have the same temperature profile  $T_i(t)$  based on the simplified model in 3(b). So the optimization problem for this multi-core case is:

$$\begin{aligned} \max \quad & n \int_0^{t_f} f(t) dt \\ \text{s.t.} \quad & T_i(0) = T_i(t_f) = T_m, T_p(0) = \frac{nR_p}{R+nR_p}T_m, \forall i \\ & \frac{dT_i}{dt} = -\frac{(T_i - T_p)}{RC} + \frac{kf(t)}{C} \\ & \frac{dT_p}{dt} = -\frac{T_p}{R_p C_p} + \frac{n(T_i - T_p)}{RC_p} \\ & 0 \leq f(t) \leq f_{max} \end{aligned} \quad (6)$$

Starting from the steady state in which  $T_i(0)$  is  $T_m$ , and ends up at  $T_i(t_f) = T_m$ , the purpose is to maximize the total frequency within  $[0, t_f]$ . The first constraint is the constraint on initial and final processor temperature and initial package temperature. The next two constraints specifies how the processor and package temperature change with respect to the frequency. *Note, there is no constraint on the maximum temperature in this problem formulation. In the next subsection, we will incorporate a maximum core temperature constraint  $T_{max}$  as well.*

**The Optimal Solution:** The optimal frequency strategy

is first all processors run at the maximum frequency  $f_{max}$  and then, at the switching point  $t_s$ , shut down simultaneously (as shown in figure 4).



**Figure 4: (a) Optimal frequency strategy for MCSF, (b) Corresponding temperature**

The switching point  $t_s$  is the unique real number satisfying:

$$\begin{aligned} & \frac{T_m C}{f_{max} k} (\alpha_1 (1 - e^{\alpha_2 t_f}) - \alpha_2 (1 - e^{\alpha_1 t_f}) + \frac{e^{\alpha_1 t_f} - e^{\alpha_2 t_f}}{(R + nR_p)C}) \\ & + \gamma_1 (e^{\alpha_2 t_f} - e^{\alpha_2 (t_f - t_s)}) - \gamma_2 (e^{\alpha_1 t_f} - e^{\alpha_1 (t_f - t_s)}) = 0 \\ & \text{where } \alpha_1 = -\beta_1 + \beta_2, \quad \alpha_2 = -\beta_1 - \beta_2 \\ & \gamma_1 = (R + nR_p)C\alpha_1 + 1, \quad \gamma_2 = (R + nR_p)C\alpha_2 + 1 \\ & \beta_1 = \frac{1}{2} \left( \frac{1}{RC} + \frac{1}{R_p C_p} + \frac{n}{RC_p} \right), \quad \beta_2 = \frac{1}{2} \sqrt{4\beta_1^2 - \frac{4}{R_p C_p RC}} \end{aligned} \quad (7)$$

This optimal solution has only one distinct frequency ( $f_{max}$ ) and one shutdown state, which can be very easily realized in practice.

### 3.3 Imposing constraints on maximum temperature (MCSFT<sub>max</sub>)

Similar with the single core case, even though we allow the temperature to violate the constraint, we don't want the core temperature to be too high. So we impose maximum temperature constraint  $T_{max}$  here and the problem formulation becomes:

$$\begin{aligned} \max \quad & n \int_0^{t_f} f(t) dt \\ \text{s.t.} \quad & T_i(0) = T_i(t_f) = T_m, T_p(0) = \frac{nR_p}{R+nR_p}T_m \end{aligned} \quad (8)$$

$$\begin{aligned} \frac{dT_i}{dt} &= -\frac{(T_i - T_p)}{RC} + \frac{kf(t)}{C} \\ \frac{dT_p}{dt} &= -\frac{T_p}{R_p C_p} + \frac{n(T_i - T_p)}{RC_p} \\ T_i(t) &\leq T_{max}, 0 \leq f_i(t) \leq f_{max} \end{aligned}$$


**Solution:** The frequency strategy for this problem is (as shown in figure 5): first, let all processors run at  $f_{max}$ , until their temperature  $T_i$  reaches  $T_{max}$ ; then the frequency changes so that the temperature maintains at  $T_{max}$ ; finally, shut down all the processors so that temperature go back to  $T_m$  as fast as possible. Basically this solution follows the same intuition as the single core problem formulation  $SCT_{max}$ . The first switching point  $t_{s1}$  is the moment when core temperature  $T_i$  reaches  $T_{max}$ . The second switching point  $t_{s2}$  is the moment that we should shut down the processors so that  $T_i$  could reduce to  $T_m$  at time  $t_f$ .

Between  $[t_{s1}, t_{s2}]$ , the processors should run at the frequency which could let  $T_i$  maintain at  $T_{max}$ . Since we don't want the core temperature to change,  $\frac{dT_i}{dt} = 0$  and  $T_i = T_{max}$ . Hence the frequencies should be such that

$$\begin{aligned} \frac{dT_i}{dt} &= -\frac{(T_i - T_p)}{RC} + \frac{kf(t)}{C} = 0, \quad t_{s1} \leq t \leq t_{s2} \\ \Rightarrow f(t) &= \frac{T_i - T_p(t)}{Rk} = \frac{T_{max} - T_p(t)}{Rk} \end{aligned} \quad (9)$$

Also,  $T_p$  follows the following dynamics:

$$\frac{dT_p}{dt} = -\frac{T_p}{R_p C_p} + \frac{n(T_{max} - T_p)}{RC_p} \quad (10)$$



**Figure 5: (a) Optimal frequency strategy for MCSFT<sub>max</sub>, (b) Corresponding temperature**

So the frequency  $f(t)$  for  $t_{s1} \leq t \leq t_{s2}$  is:

$$f(t) = -\frac{1}{Rk}(T_p^{s1} - \frac{nR_p}{R+nR_p}T_{max})e^{-\sigma(t-t_{s1})} + \frac{T_{max}}{(R+nR_p)k} \quad (11)$$

Between  $t_{s1}$  and  $t_{s2}$ , the frequency changes continuously, which is undesirable since continuous frequency change is hard to implement. In Section 6, we will present techniques of legalizing the continuous frequency into pre-defined number of discrete values.

#### 4. MULTICORE DISTINCT FREQUENCY

So far, we have assumed that for the multi-core processor, all the cores have the same frequency and is constrained by same maximum temperature  $T_{max}$ . Also, we used a simplified model that ignores lateral resistance. In this section, we investigate the problem where each core has a separate frequency control and has different maximum temperature constraint. We also use the more accurate model as in figure 3(a) that incorporate lateral resistance. This adds complexity to the problem and analytical solutions are hard to formulate. Therefore, we propose a Linear Programming (LP) based approximate problem formulation.

##### 4.1 Approximate thermal model for multi-core processor

In our model,  $t_f$  is divided into  $l$  slots and each slot, whose length is  $\Delta t$ , contains two parts.


*First part:* for  $0 \leq t < 10RC$ , the package temperature  $T_p$  is assumed to be constant. This is because  $R_p C_p$  is much larger than  $10RC$ , so within this period, the change in  $T_p$  is so small compared to the change in core temperature  $T_i$  that  $T_p$  could be regarded as constant. The lateral heat flow is also ignored here, since it is much smaller than vertical heat flow. So during this short period, the change in power consumption of core  $i$  mostly influences  $T_i$ , but has little impact on other cores. The paper [7] has a similar approximated model. The model is described with figure 6(a) and equation 12 (for simplicity but without lose of generality, we assume two cores here).

*Second part:* for  $10RC \leq t < \Delta t$ , the capacitors  $C$  get charged. Although  $T_i$  changes, it changes due to the change of  $T_p$ . So the capacitors  $C$  can be removed from the model. Also, during this period, power consumption of core  $i$  will influence the temperature of other cores through lateral heat flow, so we incorporate lateral resistance here. The resulting model is shown in figure 6(b), and equation 13 describes its thermal dynamics. Note that the paper [7] has similar model.

$$\frac{dT_i}{dt} = -\frac{(T_i - T_p)}{RC} + \frac{kf_i}{C}, \quad \forall i \quad (12)$$

$$\begin{aligned} \frac{dT_p}{dt} &= -\frac{T_p}{R_p C_p} + \frac{k(f_1 + f_2)}{C_p} \\ T_1(t) &= T_p(t) + R(kf_1 + I_l) \\ T_2(t) &= T_p(t) + R(kf_2 - I_l) \\ T_2 - T_1 &= R_l \times I_l \end{aligned} \quad (13)$$

*The overall model:* As mentioned earlier, the duration  $t_f$  is divided into slots of length  $\Delta t$ . Let  $f_1^j(f_2^j)$  denotes the frequency of the first(second) core in  $j$ 'th slot and in



**Figure 6: (a) LP based thermal model for  $t < 10RC$ , (b) LP based thermal model for  $t \geq 10RC$**

each slot, the frequency is assumed constant. Therefore the core temperature  $T_1^j$  and  $T_2^j$ , and package temperature  $T_p^j$  of each slot is monotonically increasing or decreasing. Let the core temperature  $T_1^{j-1}$  and  $T_2^{j-1}$ , and the package temperature  $T_p^{j-1}$  be the thermal state at the end of slot  $j-1$ , these values form the initial condition to the differential equations describing the thermal dynamics in slot  $j$  (with  $f_1^j$  and  $f_2^j$  be the core frequencies). By solving equation 12 we can get the core and package temperature at  $10RC$  in slot  $j$ :

$$\begin{aligned} T_i^{10RC} &= (T_i^{j-1} - T_p^{j-1} - Rkf_i^j)e^{-10} + T_p^{j-1} + Rkf_i^j, \quad \forall i \\ T_p^{10RC} &= 0.5 \times (T_1^{10RC} + T_2^{10RC} - Rk(f_1^j + f_2^j)) \end{aligned} \quad (14)$$

Note that for the purpose of simplifying the model, we had assumed  $T_p$  to be a constant before  $10RC$ . In reality  $T_p$  does change a bit and we estimate the value of  $T_p$  at  $10RC$  by taking the average of  $(T_1^{10RC} - R(kf_1^j + I_l^j))$  and  $(T_2^{10RC} - R(kf_2^j - I_l^j))$ .

Now using the temperature at  $10RC$  as the initial temperature, we get the temperature at the end of slot  $j$  by solving equation 13:

$$\begin{aligned} T_p^j &= (T_p^{10RC} - R_p k(f_1^j + f_2^j))e^{-\frac{\Delta t - 10RC}{R_p C_p}} + R_p k(f_1^j + f_2^j) \\ T_1^j &= T_p^j + Rk(\frac{R + R_l}{2R + R_l} f_1^j + \frac{R}{2R + R_l} f_2^j) \\ T_2^j &= T_p^j + Rk(\frac{R}{2R + R_l} f_1^j + \frac{R + R_l}{2R + R_l} f_2^j) \end{aligned} \quad (15)$$

From equation 14 15, we can see that the core and package temperature at the end of any slot is a linear function of the frequency. We simulated this approximate model in SPICE and found our results to be very accurate.

##### 4.2 LP based multi-core multi-frequency optimization problem

In the problem formulation for multi-core multi-frequency, we generalize the previous problem so that each core can have separate control of frequency (different  $f$ ) and can heat up to difference maximum temperature (different  $T_{max}$ ). We also use the more accurate model incorporating lateral resistance in this problems. Here we use two cores to develop this problem, but without lose of generality, we can extend it to  $n$  cores.

Starting from the steady state, the initial temperature  $T_1^0$  and  $T_2^0$  is  $T_m$ , and the initial package temperature  $T_p^0$  is  $\frac{2R_p}{R+2R_p}T_m$ . As indicated in equations 14 and 15, with the initial temperature, we can get the temperature in the first slot. Then, using the temperature at the end of the first slot  $T_1^1$ ,  $T_2^1$  and  $T_p^1$  as the initial temperature for the second slot, we can get the temperature at the second slot, and so forth. So with the initial temperature, we can get the temperature at each slot. Since the frequency in each slot is constant and the temperature in each slot is monotonic,

we just need to place constraint on the temperature at the end of each slot. So the problem formulation here is:

$$\begin{aligned}
\max \quad & \sum_{\forall \text{ slots: } j} (f_1^j + f_2^j) \\
\text{s.t.} \quad & T_1^0 = T_2^0 = T_m, \quad T_p^0 = \frac{2R_p}{R + 2R_p} T_m \\
& T_1^j \leq T_{max1}, \quad T_2^j \leq T_{max2}, \quad j = 1 \dots l-1 \\
& T_1^l \leq T_m, \quad T_2^l \leq T_m \\
& 0 \leq f_1^j \leq f_{max1}, \quad 0 \leq f_2^j \leq f_{max2}
\end{aligned} \tag{16}$$

where  $j$  is the index of slot,  $T_{max1}$  and  $T_{max2}$  are the maximum temperature constraints for the two cores. The first constraint requires the initial temperature to start from the steady state. The second constraint requires that the temperature at the end of each slot (except the last slot) should not exceed  $T_{max1}$  or  $T_{max2}$ . The next constraint requires the temperature reduce to  $T_m$  at the end of the last slot.

As indicated in equation 14 and 15, each temperature is a linear function of the frequency in each slot. Therefore, this problem can be solved efficiently with any standard linear programming methods. If we divide  $t_f$  into smaller slots, the LP formulation has more variables, but we can potentially get better solution. Contrarily, if we use less number of slot, the problem will be solved faster, but the result might be worse. Hence by controlling the sampling degree of  $t_f$ , we can obtain better or faster results.

## 5. IMPACT OF LEAKAGE

The total power consumption consists of dynamic power  $P_d = kf$  and leakage power  $P_l$ . The leakage power, which depends on temperature, is  $P_l = A \times T^2 \times e^{-\frac{a}{T}} + B$  (A, B, a are constants)[5][11]. The dynamic power  $P_d$  is a linear function of frequency. As it has been highlighted in LP model,  $t_f$  is divided into slots. In each slot, frequency is constant. The leakage power follows the change of temperature  $T(t)$ . Since the length of slot is very small, the change in temperature in each slot is very small. So is the change in leakage power consumption. Therefore, we approximate the leakage in each slot that it just depends on the initial temperature of the slot. So the power consumption of core  $i$  in  $j$ 'th slot  $P_i^j$  is:

$$P_i^j = P_d^j + P_l^j = kf_i^j + A \times (T_i^{j-1})^2 \times e^{-\frac{a}{T_i^{j-1}}} + B \tag{17}$$

In equation 18, we re-derive the temperature  $T_i^{10RC}, T_p^j, T_i^j$  by replacing the dynamic power  $kf_i^j$  in equations 14, 15 with total power consumption  $P_i^j$  in equations 17:

$$\begin{aligned}
T_i^{10RC} &= T_i^{j-1} e^{-10} + T_p^{j-1} (1 - e^{-10}) + R(1 - e^{-10}) P_i^j \\
T_p^j &= (0.5e^{-10} (T_1^{j-1} + T_2^{j-1}) + (1 - e^{-10}) T_p^{j-1}) e^{-\frac{\Delta t - 10RC}{R_p C_p}} \\
&\quad + (R_p - (R_p + 0.5Re^{-10}) e^{-\frac{\Delta t - 10RC}{R_p C_p}}) (P_1^j + P_2^j) \\
T_1^j &= T_p^j + R \left( \frac{R + R_l}{2R + R_l} P_1^j + \frac{R}{2R + R_l} P_2^j \right) \\
T_2^j &= T_p^j + R \left( \frac{R}{2R + R_l} P_1^j + \frac{R + R_l}{2R + R_l} P_2^j \right)
\end{aligned} \tag{18}$$

From equation 17, we can prove that the total power consumption of core  $i$  in slot  $j$   $P_i^j$  is a linear function of the frequency and convex function of temperature  $T_i^{j-1}$ . Here  $T_i^{j-1}$  is the temperature of core  $i$  at the end of slot  $j-1$ , which is also the initial temperature of core  $i$  at slot  $j$ . So the new  $T_i^{10RC}$ ,  $T_p^j$  and  $T_i^j$  in equation 18, in which the coefficients before  $P_i^j$  are all positive, are also linear function of frequency and convex function of the temperature

at the end of slot  $j-1$  ( $T_i^{j-1}$ ). The optimization problem formulation is still described by equation 16, in which the objective function is linear and the constraints are convex. As a result, this problem can be solved with convex optimization methods.

## 6. FREQUENCY LEGALIZATION

In the optimal frequency strategy for multi-core single frequency case, to let the temperature maintain at  $T_{max}$ , the frequency needs to change continuously, which is impractical in real implementation. The LP formulation also generates solutions with arbitrary frequency profiles. In general, most processors cannot efficiently implement a continuous range of frequencies and are constrained to operate in a pre-decided set of discrete frequencies. So we wish to legalize the continuous frequency  $f(t)$  into  $r$  discrete levels, where  $r$  depends on how many distinct frequencies we wish to have on a processor. The approximation is basically approximating the continuous frequency function by  $r$  lower bound piecewise constant values. Basically, we will divide the frequency function into  $r$  slots. In each slot, we run the processor at the highest possible lower bounding frequency within this slot. The problem is to determine the length of each slot. Once the slot length is fixed, we can execute the processor at the smallest frequency in the function  $f(t)$  in this slot. This ensures that the thermal constraints are never violated.

In multi-core single frequency model, during period  $[t_{s1}, t_{s2}]$ , the frequency profile (see equation 11) is of the form  $e^{-\sigma t}$ . Hence in slot  $(t_{j-1}, t_j)$ , the highest lower bounding frequency is  $f(t_j)$ . The objective of legalization is to find  $r$  slots while minimizing the total error between the legalized frequency and the original frequency:

$$\min \sum_{j=1}^r \left( \int_{t_{j-1}}^{t_j} f(t) dt - (t_j - t_{j-1}) f(t_j) \right) \tag{19}$$

where  $t_0$  is  $t_{s1}$ ,  $t_r$  is  $t_{s2}$  and the  $j$ 'th slot is from  $t_{j-1}$  to  $t_j$ . Also  $f(t)$  is the actual frequency profile and  $(t_j - t_{j-1})$  is the length of the  $j$ -th slot and  $f(t_j)$  is the lower bound approximation in that slot. It can be proved that the position of the splits that minimizes the total error is given by the following equation:

$$e^{-\sigma(t_{j+1} - t_j)} = \sigma(t_{j-1} - t_j) + 1, \quad 1 \leq j \leq r \tag{20}$$

For the frequency legalization of LP based model, we could also use a similar approach.

## 7. EXPERIMENTAL RESULTS


We obtained the parameters for the equivalent RC circuit  $R, C, R_p, C_p, R_l$  from [9] [10]. Assuming linear dependency between dynamic power and frequency,  $P = kf$ , the coefficient  $k$  is obtained by  $k = C_{eff} V_{dd}^2$ . Simulating the  $0.18\mu m$  technology, we set  $C_{eff}$  as  $1.11 \times 10^{-9}$ , and  $V_{dd} = 1.6V$  [6]. The normal thermal constraint  $T_m$  is  $70^\circ C$ . We experiment within the period  $t_f = 100sec$ .

### 7.1 Single core

In single core model, the value of  $R$  and  $C$  is a combination of both the core and package as a whole. We set thermal resistance  $R = 0.34K/W$  and the thermal capacitance  $C = 340J/K$ . If the initial temperature is  $T_m$ , and we don't wish to violate the thermal constraint then the highest frequency we can execute for the entire duration  $t_f$  is  $f_{nat} = \frac{T_m}{Rk}$  [8]. As shown in figure 7(a), if the absolute temperature constraint  $T_{max} = 80^\circ C$ , as we increase the maximum frequency  $f_{max}$  from  $f_{nat}$  to  $20f_{nat}$ , the ratio between the average frequency of the optimal solution  $(f_{max} \times t_s + 0 \times (t_f - t_s)) / t_f$  and the natural frequency increase from 1 to about 1.13. Interestingly, we can see that, after  $f_{max}$  increase to about  $5f_{nat}$ , the frequency gain almost remains the same. That means, increasing  $f_{max}$  beyond a certain level is not beneficial since the system is



then constrained by the maximum temperature  $T_{max}$ . We also tested the frequency gain for different  $T_{max}$ . As we expected, when  $T_{max}$  increases, the frequency gain will also increase.



**Figure 7: (a) Frequency gain for  $SCT_{max}$  with different  $T_{max}$ , (b) Frequency gain for  $MCSFT_{max}$ , (c) Frequency gain for  $MCSFT_{max}$  with different  $T_{max}$**

## 7.2 Multi-core with same frequency

In the multi-core model, we set thermal resistance for each core  $R = 0.42$  K/W and their thermal capacitance  $C = 0.024$  J/K. For the package,  $R_p = 1$  K/W and  $C_p = 140.449$  J/K. We assume there are  $n = 2$  cores, and the natural frequency is  $f_{nat} = \frac{T_m}{(R+nR_p)k}$ . Conceptually, for a multicore processor with the initial core temperature  $T_i(0) = T_m$  and package temperature  $T_p(0) = \frac{nR_p}{R+nR_p}T_m$ , the maximum core frequency such that the core temperature stays within  $T_m$  is the natural frequency  $f_{nat}$ . As we increase the maximum frequency  $f_{max}$  from  $f_{nat}$  to  $20f_{nat}$ , the ratio between the average frequency of the optimal solution and the natural frequency (frequency gain) increase from 1 to about 2, if there is no constraint on the maximum temperature  $T_{max}$  (figure 7(b)).

If we also place the maximum temperature constraint, e.g.,  $T_{max} = 80^\circ\text{C}$  ( $MCSFT_{max}$  problem), the resultant frequency gain ranges from 1 to 1.3 (as shown in figure 7(c)). Interestingly also, after  $f_{max}$  increases to about  $5f_{nat}$ , the frequency gain almost stays unchanged. This is also because of the constraint on  $T_{max}$ . So with the maximum temperature constraint, increasing  $f_{max}$  to very high does not help in improving the performance.

As shown in figure 7(c), in which  $f_{max}$  still changes from  $f_{nat}$  to  $20f_{nat}$ , and  $T_{max}$  changes from  $75^\circ\text{C}$  to  $100^\circ\text{C}$ , when  $T_{max}$  increases, the frequency gain also increases.

## 7.3 LP based multi-core multi-frequency

For the LP based multi-core multi-frequency model, we use more accurate model where lateral resistance is added and simulate two core case. We set lateral resistance  $R_l = 3.75$  K/W. Each core has different  $T_{max}$ , and can control the frequency separately. Here, we set  $T_{max1} = 80^\circ\text{C}$ ,  $T_{max2} = 90^\circ\text{C}$  and there are  $n = 2$  cores. The initial temperature  $T_m = 70^\circ\text{C}$ .


First, we obtain the optimal frequency policy through linear programming. The optimal frequency policy and the corresponding temperature is shown in figure 8(a), 8(b). In the optimal frequency policy, firstly, the two processors run at their maximum frequency  $f_{max1}$  and  $f_{max2}$ , respectively, until each of them reaches its maximum temperature one after another; and when each processor reaches its maximum temperature, the frequency of this processor changes so that the temperature of this processor maintains at its maximum temperature. Then, they shut down simultaneously and their temperature reduce to  $T_m$  at  $t_f$ .

When  $f_{max}$  increases, its frequency gain is shown in figure 8(c) (solid line). Assume both core have the same  $f_{max}$ , which increase from  $f_{nat}$  to  $20f_{nat}$ , the frequency gain is from 1 to about 1.46. And after about  $2f_{nat}$ , the frequency gain almost stays unchanged.

## 7.4 Frequency legalization

Using the frequency legalization strategy in section 6 for our LP based model, the resultant frequency gain is shown in figure 8(c) (dotted line). Here we use 4 levels to approximate the continuously changed frequency. Compared

with the average frequency without legalization, there is about 8% frequency loss. But even after frequency legalization, there is still about 1.4 performance gain. More importantly, the legalized frequency strategy is practical and can be used in real implementation.



**Figure 8: (a) Frequency policy for LP model, (b) Corresponding temperature, (c) Frequency gain and frequency legalization for LP model**

## 8. CONCLUSION

In this paper, we explored dynamic thermal management problem under soft constraints for both single and multi-core processor. We obtained analytical expression of the optimal frequency policy. We also developed linear programming based approximation for multi-core processor. Finally, we used frequency legalization method so that our frequency policy can be used in realistic scenario.

## 9. ACKNOWLEDGMENTS

We would like to thank NSF grant CCF 0937865 for supporting part of this research.

## 10. REFERENCES

- [1] D. Brooks and M. Martonosi. Dynamic thermal management for high-performance microprocessors. In *Proc. of the 7th Intl. Symp. on High-Performance Computer Architecture (HPCA'01)*.
- [2] A. Cohen, L. Finkelstein, A. Mendelson, R. Ronen, and D. Rudoy. On estimating optimal performance of cpu dynamic thermal management. *IEEE Computer Architecture Letters*, 2:6, 2003.
- [3] A. K. Coskun, T. S. Rosing, and K. C. Gross. Proactive temperature management in mpsoes. In *Proc. of Intl. Symp. on Low Power Electronics and Design (ISLPED'08)*.
- [4] A. K. Coskun, T. S. Rosing, and K. C. Gross. Temperature management in microprocessor socs using online learning. In *Design Automation Conference (DAC'08)*.
- [5] L. He, W. Liao, and M. R. Stan. System level leakage reduction considering the interdependence of temperature and leakage. In *Design Automation Conference (DAC'04)*.
- [6] S. M. Martin, K. Flautner, T. Mudge, and D. Blaauw. Combined dynamic voltage scaling and adaptive body biasing for lower power microprocessors under dynamic workloads. In *IEEE/ACM Intl. Conf. on Computer Aided Design (ICCAD'02)*.
- [7] R. Rao and S. Vrudhula. Performance optimal processor throttling under thermal constraints. In *Proc. of Intl. Conf. on Compilers Architectures and Synthesis for Embedded Systems (CASES'07)*.
- [8] R. Rao, S. Vrudhula, and N. Chang. An optimal analytical solution for processor speed control with thermal constraints. In *Proc. of Intl. Symp. on Low Power Electronics and Design (ISLPED'06)*.
- [9] K. Skadron, T. Abdelzahery, and M. R. Stan. Control-theoretic techniques and thermal-rc modeling for accurate and localized dynamic thermal management. In *Proc. of Intl. Symp. on High-Performance Computer Architecture (HPCA'02)*.
- [10] K. Skadron, M. R. Stan, K. Sankaranarayanan, W. Huang, S. Velusamy, and D. Tarjan. Temperature-aware microarchitecture: Modeling and implementation. *ACM Trans. on Architecture and Code Optimization*, 1:94–125, 3.
- [11] L. Yuan, S. Leventhal, and G. Qu. Temperature-aware leakage minimization technique for real-time systems. In *IEEE/ACM Intl. Conf. on Computer Aided Design (ICCAD'06)*.