

An Optimal Analytical Solution for Processor Speed Control with Thermal Constraints*

Ravishankar Rao, Sarma Vrudhula,
Chaitali Chakrabarti
Consortium for Embedded Systems
Arizona State University
Tempe, AZ 85281, USA
{ravirao, vrudhula, chaitali}@asu.edu

Naehyuck Chang
Computer Systems Lab
School of Computer Science and Engineering
Seoul National University
Seoul, 151-010, South Korea
naehyuck@cslab.snu.ac.kr

ABSTRACT

As semiconductor manufacturing technology scales to smaller device sizes, the power consumption of clocked digital ICs begins to increase. Dynamic voltage and frequency scaling (DVFS) is a well-known technique for conserving energy. Recently, it has also been used to control the CPU temperature as part of Dynamic Thermal Management (DTM) techniques. Most works in these areas assume that the optimum speed profile (for either minimizing energy or maximizing performance) is a constant profile. However, in the presence of thermal constraints, we show that the optimal profile is in general, a time-varying function. We formulate the problem of maximizing the average throughput of a processor over a given time period, subject to thermal and speed constraints, as a problem in the calculus of variations. The variational approach provides a powerful framework for precisely specifying and solving the speed control problem, and allows us to obtain an exact analytical solution. The solution methodology is very general, and works for any convex power model, and simple lumped RC thermal models. The resulting speed profiles were found to consist of up to three segments, of which one of them is a decreasing function of time, and the others are constant. We analyze the effect of different parameters like the initial temperature, thermal capacitance and the maximum rated speed on the nature and the cost of the optimum solution. We also propose a two-speed solution that approximates the optimal speed curve. This solution was found to achieve a performance close to that of the optimum, and is also easier to implement in real processors.

Categories and Subject Descriptors

C.4 [Performance of systems]: modeling techniques, performance attributes; G.1.6 [Optimization]: constrained optimization

*We gratefully acknowledge the support for this work by the Consortium for Embedded Systems at the Arizona State University and by a grant from the National Science Foundation, grant number CNS-0509540.

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 '06, October 4–6, 2006, Tegernsee, Germany.
Copyright 2006 ACM 1-59593-462-6/06/0010 ...\$5.00.

General Terms

algorithms, performance, theory

Keywords

optimal control, thermal management, temperature, DTM, DVFS

1. INTRODUCTION

Over the past decade, microprocessor power consumption has increased exponentially [5]. Higher power consumption means greater energy costs, both for powering the chip, and for cooling it. Further, as chip form factors shrink, the power density increases too. This makes it difficult to design thermal solutions that ensure that the chip temperature is always within safe limits. Even if achieved, designing for the worst case could result in prohibitively expensive or bulky systems. Hence, thermal designers target a power consumption value smaller than the worst case power consumption – they ensure that the case temperature θ remains below a specified maximum value θ_{\max} whenever the processor is operating at or below the Thermal Design Power (TDP) P_m . Such a design allows the processor to operate temporarily at higher speeds that dissipate more power than the TDP. If the TDP is exceeded for too long, however, the chip could heat up dangerously beyond θ_{\max} , and the thermal solution, by design, cannot prevent it.

To handle such scenarios, processor manufacturers usually provide a hardware mechanism to quickly reduce the power consumption whenever an on-chip temperature sensor detects a thermal emergency [1, 5, 7]. These mechanisms include varying the duty cycle of the clock, throttling the instruction fetch unit of the processor (fetch throttling), and reducing the clock frequency and/or voltage scaling (DVFS). They constitute the simplest of a class of techniques known as Dynamic Thermal Management (DTM). Brooks and Martonosi [2] provide a comprehensive overview and comparison of different DTM techniques. All these techniques involve some kind of processor slowdown, and as pointed out by Skadron *et al* [14], the challenge of DTM is to adaptively control the processor speed to minimize the slowdown but still meet thermal constraints.

The DTM techniques implemented in current processors [1, 7] are fixed-response reactive heuristics i.e. (i) the response does not vary in proportion to the thermal emergency, and (ii) they are invoked only when the temperature crosses a certain threshold. Skadron *et al* proposed a feedback control system that adaptively varies the fetch toggling rate of the processor to maintain the chip temperature just under the maximum allowed temperature T_{\max} . This allowed the processor to achieve safe thermal operation while suffering significantly lower performance penalty. Srinivasan and Adve [15] proposed predictive DTM algorithms for multimedia applications.

These algorithms are able to extract higher performance than reactive DTM techniques by using profiled temperature information to achieve maximum performance under thermal constraints.

A problem central to dynamic thermal management that has not been addressed so far, is the derivation of the shape of the performance optimal speed profile (a function of time), while satisfying specified thermal constraints. Consider the following simple problem in processor speed control: (X-MAX): maximize a processor's performance (measured in number of clock cycles executed) over a given time duration.¹ The optimal solution for this problem is not a function of time – one would simply operate at the largest supported speed u_{\max} . However, this solution does not account for the effect of thermal constraints on the optimal speed profile.

An ideal speed control technique should to be able to achieve a desired combination of performance, energy consumption, temperature, acoustic noise (from cooling fans), and reliability. This paper takes an important first step towards that goal by revisiting the simplest problem in DTM mentioned above (X-MAX), but this time, with the thermal constraint $T \leq T_{\max}$. Using the theory of the calculus of variations, we obtained an exact analytical solutions to this problem. The thermal constraints were observed to cause the optimal speed profile to be, in general, a piecewise non-linear function of time. The solution to this problem provides a theoretical upper bound to the maximum performance improvement that can be achieved by using speeds that dissipate power larger than the TDP. We also devised a two-speed (piecewise constant) profile that achieves nearly the same performance as the continuously time-varying performance optimum curve. This result makes it easier to implement the performance-optimum speed profile and to extend it for more complex scenarios.

2. NOTATION, MODELS, AND PROBLEM FORMULATION

Throughout this paper, the notation $x' = dx/dt$ will be used to denote the (total) derivative, always w.r.t. time t . Different glyphs will be used to differentiate between a variable (e.g. P) and a function assigned to a variable (e.g. $P = P(u)$).

2.1 Processor model

In this work, we restrict ourselves to a single DTM mechanism, namely DVFS, so that there are potentially two control variables, the supply voltage and clock speed u . Our processor model is similar to that of many works on DVFS like [11]. The processor can attain any speed u between 0 and u_{\max} cycles/s. At each speed, the minimum feasible voltage [8, 12, 13] is assumed to be used. Hence, the processor's power consumption P is a function of only the speed. This function $P(u)$, is assumed to be increasing, strictly convex, and once differentiable. The power slope is defined as $S(u) = dP(u)/du$. Although it is usually not emphasized in the notation, u , P , and S are understood to be functions of time.

In recent times, there has been considerable research into developing accurate processor power models, particularly for modeling leakage currents [8, 12] and leakage dependence on temperature [10]. Our work is independent of the specific form of the power model, and is applicable to most leakage models including [8, 12]. Currently, however, our work does not include the effect of temperature on the leakage currents. The performance of a processor x over a given time interval $[t_i, t_f]$ is defined to be the number of clock cycles (amount of work done) executed by the processor over that interval and is given by $x = \int_{t_i}^{t_f} u dt$. Note that $x' = u$.

¹The number of clock cycles executed over a given duration is proportional to the average throughput, which, for a given Instructions Per Clock (IPC) is proportional to the clock frequency [10].

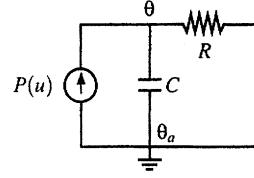


Figure 1: Equivalent thermal RC circuit.

2.2 Thermal model

Using a well-known duality between heat transfer and electrical phenomena [3], the relationship between the chip temperature θ and the processor power P can be modeled by an equivalent thermal R-C circuit. A second-order RC model is usually adequate to accurately model the transient thermal behavior of the chip package and heat sink [4]. In the interest of clarity however, we have used a lumped first order lumped RC model shown in Figure 1. This thermal model provides the following relationship between the power consumption P and the chip junction temperature θ (measured relative to the fixed ambient temperature):

$$RC\theta' + \theta - RP(u) = 0 \quad (1)$$

where R (unit: $^{\circ}\text{C}/\text{W}$) is the thermal resistance from junction to ambient, and C (unit: $\text{J}/^{\circ}\text{C}$) is the thermal capacitance.

The thermal resistance is a measure of the temperature θ_{ss} to which the chip junction will heat up to in steady state, for a given constant power dissipation P , i.e. $R = \theta_{ss}/P$. In most electronic systems, the main modes of heat transfer are through conduction (from chip junction through the packaging and heat sink to the heat sink's outer surface) and convection (from heat sink to ambient). Each mode has its own equivalent thermal resistance, or $R = R_{\text{cond}} + R_{\text{conv}}$. Typically, R_{cond} is determined by the thermal conductivity and geometry of the packaging and heat sink (fixed) while R_{conv} by the cooling mechanism [3, 6]. If an adaptive cooling mechanism like a fan is used, R_{conv} is not fixed [3]. As we are interested in controlling only the processor speed in this work, we assume that the cooling solution operates at a fixed “speed”, so that R_{conv} and hence R does not vary with time.

The thermal capacitance C is a measure of the heat storage capacity of the system. It is chiefly determined by the mass and specific heat capacity of the packaging and heat sink (fixed). Larger values of C allow the temperature to rise more slowly for the same power dissipation. This is an important factor in optimal speed control under thermal constraints as it determines the amount of time a chip can operate at powers larger than the Thermal Design Power (TDP). For a given initial temperature θ_0 , and a constant speed u , the time taken t_m to reach the maximum temperature can be obtained using (1) as

$$t_m(u) = -RC \log \left[\frac{RP(u) - \theta_{\max}}{RP(u) - \theta_0} \right] \quad (2)$$

The temperature profile generated by a constant speed u can be found using (1) and is given by

$$\theta = RP(u) + [\theta_0 - RP(u)] e^{-t/(RC)}. \quad (3)$$

If $\theta_0 < P(u)$, the temperature profile is increasing, else, non-increasing. In either case, the steady state temperature $\theta_{ss} = P(u)R$. To ensure that the chip reliability is not compromised, it is essential to always maintain the chip temperature below a maximum value θ_{\max} that is specified by the processor manufacturer [1, 7, 14]. As $P(u)$ is an increasing function, there is a unique solution to the equation $P(u)R = \theta_{\max}$, which is denoted as u_m . The corresponding power consumption $P(u_m)$ is denoted as P_m . The speed u_m

is the maximum speed that can, *in steady state*, maintain the chip temperature below the maximum temperature θ_{\max} , i.e. if $u > u_m$, $\theta_{ss}(u) > \theta_{\max}$.

Note that if at any instant, the temperature reaches θ_{\max} , reducing the speed immediately to u_m will ensure that the temperature remains at θ_{\max} . Hence, we assume that an ideal DTM technique will scale down the processor speed to u_m at the onset of a thermal emergency $\theta = \theta_{\max}$. The thermal RC system equation (1), together with the constraint $\theta(t) \leq \theta_{\max}$ then constitute the thermal constraints that need to be added to the original speed control problem (X-MAX).

2.3 Problem formulation

The performance optimization problem under thermal constraints (henceforth called X-MAX-T) involves controlling the speed profile u to maximize the number of clock cycles executed over a time duration $[0, T]$ subject to (1), the temperature upper bound, and speed constraints. It can be expressed as

$$\max_u \quad X = \int_0^T u \, dt \quad (4)$$

$$\text{subject to} \quad RC\theta' + \theta - RP(u) = 0, \quad \theta(0) = \theta_0 \quad (5)$$

$$\theta \leq \theta_{\max}, \quad (6)$$

$$\text{and} \quad 0 \leq u \leq u_{\max}. \quad (7)$$

The above formulation involves computing an unknown function $u(t)$ that optimizes a function of this unknown function ($\int_0^T u \, dt$), subject to constraints on other functions of u . This kind of optimization problem does not fit the traditional realm of optimization problems because the unknowns are functions, and not simple variables. Instead, they are best approached using the theory of the calculus of variations [16] that was developed particularly for solving such problems. Using this theory, we now show how the optimal speed profile can be obtained for X-MAX-T.

3. THE OPTIMAL SPEED PROFILE

Clearly, if the thermal constraints (5) and (6) were omitted, the optimum solution to the problem is $u = u_{\max}, 0 \leq t \leq T$. We call this the *unconstrained solution*. As the chip is not designed for worst case power dissipation by assumption, $u_{\max} > u_m$, or $RP(u_{\max}) > \theta_{\max} \geq \theta_0$. This means that the temperature will eventually reach θ_{\max} in a time $\tilde{t}_m \equiv t_m(u_{\max})$ given by (2). If $\tilde{t}_m \geq T$, the unconstrained solution is feasible. Otherwise, the unconstrained solution violates the thermal constraint $\theta \leq \theta_{\max}$ over the interval $[\tilde{t}_m, T]$. Then, we use a result from [16] that states that in the true (constrained) solution, the thermal constraint is binding over a sub-interval $[t_m, T]$ of $[\tilde{t}_m, T]$, or $u = u_m$ for $t_m \leq t \leq T$ and the optimum speed profile has two segments. We now need to find the optimum speed profile u over the segment $[0, t_m]$ and also find the optimum “switch point” t_m between the segments.

3.1 Solution without binding speed constraints

We first assume that the speed constraint $u \leq u_{\max}$ is not binding over any portion of the speed profile. The original problem can now be recast as a variational calculus problem over the interval $[0, t_m]$ with an unknown endpoint t_m and a terminal cost $u_m(T - t_m)$, as follows:

$$\min_u \quad J = - \left[\int_0^{t_m} u \, dt + u_m(T - t_m) \right] \quad (8)$$

$$\text{subject to} \quad \psi = RC\theta' + \theta - RP(u) = 0, \quad (9)$$

$$\text{and} \quad \theta(0) = \theta_0, \quad \theta(t_m) = \theta_{\max}. \quad (10)$$

The solution methodology for this problem is detailed in the Appendix. We summarize the final optimal solution here:

$$u^*(t) = \begin{cases} S^{-1} \left((1/(\alpha R)) e^{-t/(RC)} \right), & 0 \leq t < t_m \\ u_m, & t_m \leq t \leq T, \end{cases} \quad (11)$$

where the parameters α and t_m are determined using the continuity conditions

$$u^*(t_m^-) = u_m, \quad \theta^*(t_m^-) = \theta_{\max}. \quad (12)$$

The optimum solution consists of a time-varying, and a constant segment. The time-varying portion is expressed in terms of the inverse power slope function S .² At t_m , the optimum solution switches to the constant segment $u = u_m$. This has the effect of maintaining the temperature at θ_{\max} during $[t_m, T]$. If the solution process produces a value of t_m greater than T , we must re-solve the problem by setting $t_m = T$. The optimum solution will then have a single time-varying segment. The optimum value of the parameter α is now determined by the temperature continuity condition $\theta(t_m) = \theta(T) = \theta_{\max}$.

3.2 Solution with binding speed constraints

If the solution obtained above violates the speed constraint $u \leq u_{\max}$, it must do so over an interval $[0, \tilde{t}_u]$ where $\tilde{t}_u < t_m$. This is because the optimum speed profile is a non-increasing function of time. Using the same reasoning as before, we must have the true optimum solution such that $u = u_{\max}$ over some interval $[0, t_u]$ such that $t_u \leq \tilde{t}_u$. To find the optimum switch point t_u and compute the true optimum solution, we must now re-solve the problem described by (8)–(10) with the new objective

$$\min_u \quad - \left[u_{\max} t_u + \int_{t_u}^{t_m} u \, dt + u_m(T - t_m) \right]. \quad (13)$$

It can be shown that³ the optimum solution to this problem is

$$u^*(t) = \begin{cases} u_{\max}, & 0 \leq t \leq t_u, \\ S^{-1} \left((1/(\alpha R)) e^{-(t-t_u)/(RC)} \right), & t_u < t \leq t_m, \\ u_m, & t_m < t \leq T, \end{cases} \quad (14)$$

where the unknown parameters α , t_u and t_m are given by the continuity conditions

$$u^*(t_u^+) = u_{\max}, \quad u^*(t_m^-) = u_m, \quad \theta^*(t_m^-) = \theta_{\max}. \quad (15)$$

Now, if the above solution process results in a value of t_u that exceeds $t_m(u_{\max})$, the optimum solution is a two-speed profile $u^*(t) = u_{\max} \forall 0 \leq t \leq t_m(u_{\max})$, and u_m otherwise.

3.3 A two-speed approximate solution

As real processors cannot change speed smoothly, we propose a piecewise constant speed policy, CONST that approximates the optimal solution. This speed policy operates at a constant speed u_1 until $\theta = \theta_{\max}$ and then changes to u_m . The optimum value of u_1 must maximize $x(u_1) \equiv u_1 t_m(u_1) + u_m(T - t_m(u_1))$. It is found by solving the following optimization problem:

$$\max_{u_1} \quad -(u_1 - u_m) RC \log \left[\frac{RP(u_1) - \theta_{\max}}{RP(u_1) - \theta_0} \right], \quad (16)$$

$$\text{subject to} \quad 0 \leq u_1 \leq u_{\max}. \quad (17)$$

²As the power is a strictly convex function of speed, the powerslope is an increasing function of speed. Hence, the optimum speed profile in this segment is a decreasing function of time.

³The solution proceeds on similar lines as the derivation in the Appendix. We have omitted the details for brevity.

This is a non-linear optimization problem in one variable with two simple linear constraints. The logarithm and the non-linear nature of the $P(u)$ function make it difficult to obtain an analytical solution. However, the problem is easily solved numerically.

4. EXPERIMENTAL SETUP

We obtained an analytical power model for the full chip power consumption of a 70 nm CMOS processor based on data from [8, Table 1]. We made three changes to the parameters used in [8] to model a high performance high power processor. (i) we chose L_g , the number of devices as 200 million based on the transistor count of typical 90 nm and 65 nm processors, (ii) we set P_{on} , the power required to keep the processor on as 10 W, and (iii) we computed the coefficient C_{eff} , the switched capacitance per device as $1 \times 10^{-16} \text{ J V}^{-2} \text{ cycle}^{-1}$. This was obtained by dividing the value for C_{eff} of a 10 inverter chain (20 devices) from [12, Table 1] by 20.

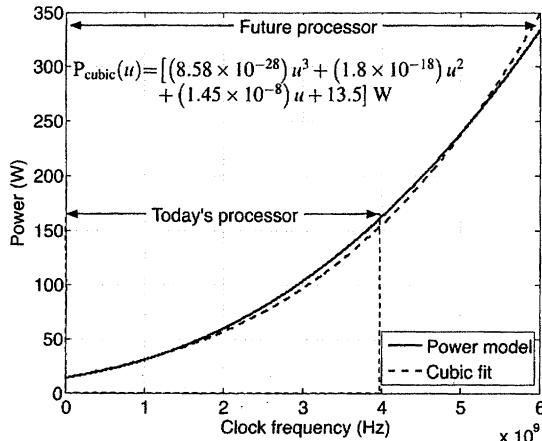


Figure 2: Processor power model, and a cubic approximation.

Our aim was to obtain a representative model of power consumption in existing high-end processors to illustrate our solution methodology, which by itself, is independent of the actual power model. We also wish to extrapolate this power model to model a future high-performance processor. To achieve this, and to be able to make a fair comparison to the power model of the current processor, we use the same power speed relationship for both, but limit the maximum speeds to 4 GHz for the current processor and 6 GHz for a future processor.

Figure 2 shows a plot of this power-speed relationship (solid curve). The power model $P(u)$ presented in [8] has a power-slope function $S(u) = dP/du$ that is not invertible. To illustrate our analytical solution, it is convenient to approximate the power-speed relationship to a simpler form (we chose a cubic) that has an invertible power slope. Figure 2 shows the cubic fit to the earlier power model. The RMS error over the range of frequencies 0 to 6 GHz was found to be 5.39 W.

5. NUMERICAL RESULTS

In this section, we study the effect of the initial temperature θ_0 , and the maximum speed u_{\max} on the nature and cost of the optimum solution. We also compare the performance of the optimum profile with the two-speed approximation and a “safe” speed policy that always operates within the TDP.

5.1 Effect of initial temperature

The initial temperature has a strong effect on the performance of the optimal solution. Figure 3 plots the optimum speed profiles for

different values of θ_0 , for a fixed maximum speed $u_{\max} = 4 \text{ GHz}$. For the lowest possible value $\theta_0 = \theta_{\text{ambient}}$, the thermal limit is not reached over the duration $T = 140 \text{ s}$, and the unconstrained optimum solution $u(t) = u_{\max}$ is the optimum solution. This constitutes the upper bound on the achievable performance. As the initial temperature rises, the optimum speed profile is forced to scale down towards u_m , which happens faster as θ_0 increases.

The resulting performance degradation is almost linear in θ_0 . This translates roughly to a reduction in average throughput of 33.3 million cycles/s per °C increase in θ_0 . This dependence of performance on θ_0 presents opportunities for balancing the temperature over successive workloads executed over a processor. As one workload’s final temperature is the initial temperature of the next, optimizing the performance or energy consumption over a set of workloads must consider the effect of each tasks final temperature on the cost of future tasks.

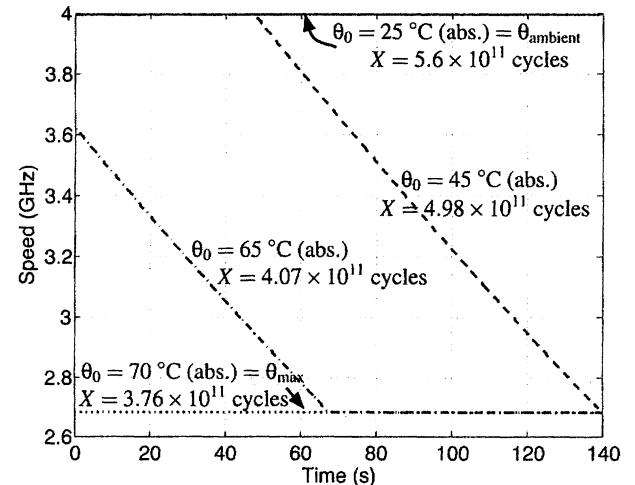


Figure 3: Effect of initial temperature on optimum solution.

5.2 Effect of the maximum speed

If the maximum speed of a processor were arbitrarily large, the resulting performance optimum solution would constitute an upper bound on the performance that can be achieved by using speeds larger than u_m . The maximum speed used by such a speed profile can be valuable for a designer as speeds larger than that speed will not help achieve higher performance. Figure 4 shows the optimal speed profiles for different values of u_{\max} , for a fixed initial temperature of 45°C (abs). It can be seen that for $u_{\max} > 4.63 \text{ GHz}$, the maximum speed no longer constrains the optimum performance profile. The corresponding performance value is the maximum performance that can be achieved using the given power-speed relationship if all parameters other than u_{\max} are kept constant.

Currently, the maximum power consumption allowed in a processor is about 10% to 40% above the TDP. This results in an even smaller difference between u_{\max} and u_m . For the processor model we study here, we can see from Figure 4 that as u_{\max} gets closer to u_m , the optimum performance sharply falls by up to 15%. While the power consumption of processors is rising quickly with each technology generation, the thermal solutions do not scale correspondingly. Hence, a practical solution for future processors may be to design speeds $> u_m$ that can only be used temporarily. This then creates a need for developing intelligent ways to use those larger speeds to extract maximum performance.

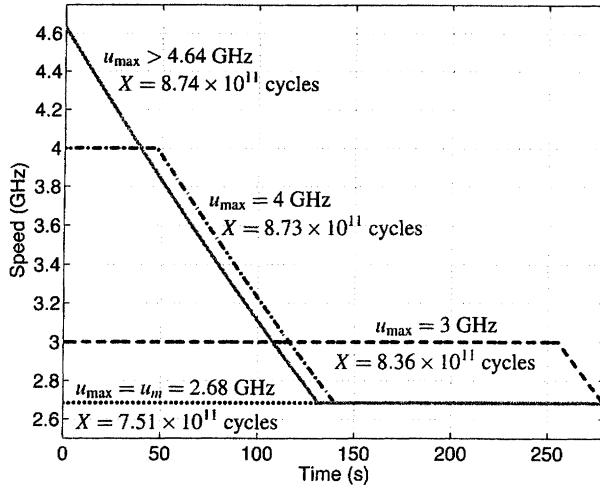


Figure 4: Effect of maximum speed on the optimum solution.

5.3 Effect of thermal capacitance

The thermal capacitance of the package-heat sink combination is an important factor to determine the time taken to reach the maximum temperature after which a forced slowdown to u_m occurs. As seen from (2), this time is directly proportional to the thermal capacitance. To study the effect of thermal capacitance on the optimum solution, we computed the optimum profiles for different values of the thermal capacitance, by keeping $u_{\max} = 5 \text{ GHz}$, and $\theta_0 = 45^\circ\text{C}$ (abs.). The resulting solutions are plotted in Figure 5. The value of 340 J/K that is typical of most processor-heat sink packages is able to stay above u_m for only two minutes.

There exist novel packaging solutions using removable aluminum plates [9] and phase-change material based heat sinks [17] that can provide much larger thermal capacitances. The use of such heat sinks for smoothing out spikes in electronic workloads has been studied by the electronic packaging community [9, 17]. From Figure 5, it can be seen that using such solutions for increasing the thermal capacitance can increase both, the time taken to reach θ_{\max} and the performance. The increase in both cases is linear. The performance increase translates to an increase in average throughput of about 0.5 million cycles/s per J/K increase in C .

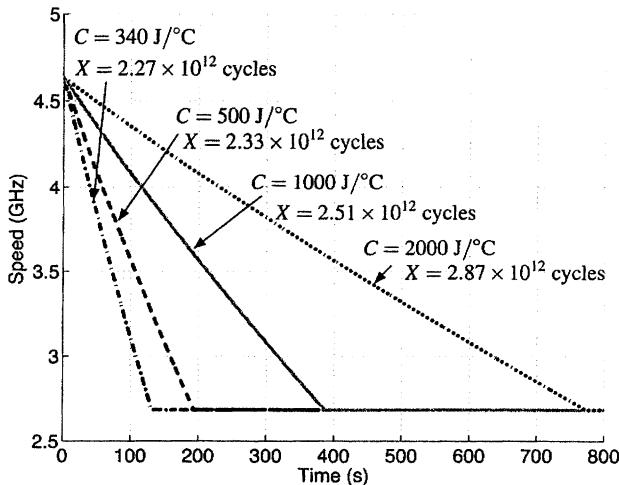


Figure 5: Effect of thermal capacitance on optimum solution.

5.4 Comparison of OPT and CONST policies

We wish to compare the performance of the optimum speed profile (OPT) with the two-speed approximation (CONST) described in Section 3.3. We also examine the policy SAFE, which avoids thermal emergencies altogether by always operating at the speed u_m . Clearly, SAFE will have a lower performance than MAX. But we wish to see how much performance improvement can be gained over it by operating at “unsafe speeds”, for a modern processor with $u_{\max} = 4 \text{ GHz}$.

Figure 6 shows the optimum speed and temperature profiles for OPT, CONST and SAFE for a time duration $T = 120 \text{ s}$. The analytical solution for the optimum speed profile OPT is given by,

$$u^*(t) = \begin{cases} 4.00 \text{ GHz}, & 0 \leq t < 16.44 \text{ s}, \\ (-0.7059 + 2.2768\sqrt{-1 + 5.7107e^{-0.0049*t}}) \text{ GHz}, & 16.44 \text{ s} < t \leq 108.80 \text{ s}, \\ 2.68 \text{ GHz}, & 108.8 \text{ s} < t \leq 120 \text{ s} \end{cases} \quad (18)$$

while the optimum solution for the two-speed profile CONST is given by

$$u_{\text{CONST}}^* = \begin{cases} u_1^* = 3.75 \text{ GHz}, & 0 \leq t \leq 73.98 \text{ s} \\ u_m = 2.68 \text{ GHz}, & 73.98 \text{ s} < t \leq 120 \text{ s} \end{cases} \quad (19)$$

The resulting performance for the three speed policies are $X_{\text{SAFE}} = 3.221 \times 10^{11} \text{ cycles}$, $X_{\text{CONST}} = 4.003 \times 10^{11} \text{ cycles}$, and $X_{\text{OPT}} = 4.010 \times 10^{11} \text{ cycles}$. An intelligent speed control policy like OPT could extract up to 25% more performance than SAFE. However, the continuous speed variation required by OPT is not practical. The two-speed solution CONST was able to achieve up to 0.62% of the performance of the optimum policy. Additionally, this solution has the advantage that it is easier to implement in a real processor. The optimum speed value u_1^* can be pre-computed for different values of the time duration T by numerically solving the single-variable optimization problem described in 3.3. The resulting solutions can then be stored in a look-up table for easy access at run-time.

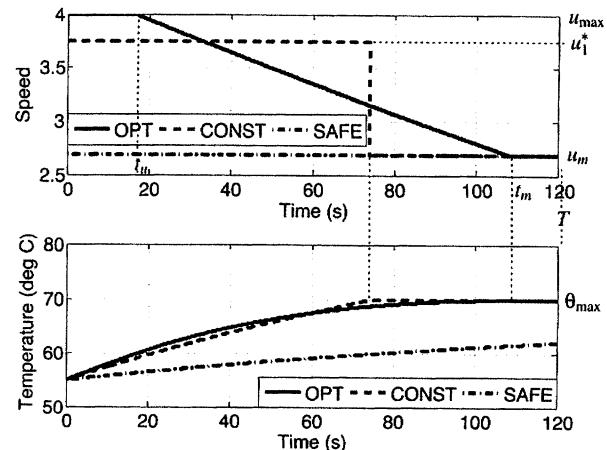


Figure 6: Optimum speed profile of different policies.

6 CONCLUSION

The disparity between thermal solutions and chip heat dissipation is growing just like the one between battery capacity and chip power consumption. The technological solutions (battery capacity increase, improved thermal design) alone cannot support the

increased capabilities of future processors. In the case of batteries, researchers responded by devising high-level battery models to account for non-linear effects and developed system-level load shaping techniques that varied the current demand from the battery to extract maximum charge. We need an analogous approach to the problem of system-level thermal management. While accurate thermal and power models already exist, the optimum manner to control the processor speed to address thermal concerns has only recently begun to be studied [2, 14, 15].

In this paper, we address a fundamental problem in system-level thermal management: What is the best way to control the speed of a processor to get maximum performance, while satisfying thermal and speed constraints? We used a novel formulation and solution methodology based on the calculus of variations, and were able to obtain an analytical solution. This solution serves as an upper bound for the maximum performance that can be extracted from a processor by speed control. We obtained a cubic power-speed relationship based on the literature, and studied the effect of different parameters like the initial temperature, thermal capacitance and maximum speed on the optimum solution. The performance was found to drop almost linearly with rising initial temperatures, and increase linearly with increasing thermal capacitance. Allowing larger maximum clock speeds helps improve the performance, but only up to a certain limit. The cubic power model also enabled us to solve for the optimum speed profile in closed form. For future work, we are working on extending the proposed solution for higher order thermal models, and including terminal temperature constraints ($\theta(T) \leq \theta_f$) in our formulation.

7. REFERENCES

- [1] Advanced Micro Devices. *CPU Thermal Management: Application Note*.
- [2] D. Brooks and M. Martonosi. Dynamic thermal management for high-performance microprocessors. In *Intl' Symp. High Speed Computer Arch. (HPCA)*, pages 171–182, 2001.
- [3] Y. A. Cengel. *Heat Transfer*. McGraw Hill, 2 edition, 2002.
- [4] B. M. Guenin. Calculation corner: Simplified transient model for IC packages. *Electronics Cooling*, 8(3), August 2002.
- [5] S. H. Gunther, F. Binns, D. M. Carmean, and J. C. Hall. Managing the impact of increasing microprocessor power consumption. *Intel Technology Journal*, Quarter 1 2001.
- [6] W. Huang, M. R. Stan, K. Skadron, K. Sankaranarayanan, S. Ghosh, and S. Velusamy. Compact thermal modeling for temperature-aware design. In *Proc. Design Automation Conf. (DAC)*, pages 878–883, 2004.
- [7] Intel Corp. *Intel Pentium 4 Processor 6x1 Sequence: Datasheet*, January 2006.
- [8] R. Jejurikar, C. Pereira, and R. Gupta. Leakage aware dynamic voltage scaling for real-time embedded systems. In *Proc. Design Automation Conf. (DAC)*, pages 275–280, 2004.
- [9] S. Krishnan and S. V. Garimella. Thermal management of transient power spikes in electronics—phase change energy storage or copper heat sinks? *Journal of Electronic Packaging*, 126(3):308–316, 2004.
- [10] W. Liao, L. He, and K. M. Lepak. Temperature and supply voltage aware performance and power modeling at microarchitecture level. *IEEE Trans. Computer-Aided Design*, 24(7):1042–1053, July 2005.
- [11] J. R. Lorch and A. J. Smith. PACE: A new approach to dynamic voltage scaling. *IEEE Trans. Computers*, 53(7):856–869, July 2004.
- [12] 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 *Proc. Intl' Conf. Computer Aided Design (ICCAD)*, pages 721–725, 2002.
- [13] J. M. Rabaey, A. Chandrakasan, and B. Nikolic. *Digital Integrated Circuits*. Prentice Hall, 2 edition, 2002.
- [14] K. Skadron, T. Abdelzaher, and M. R. Stan. Control-theoretic techniques and thermal-RC modeling for accurate and localized

dynamic thermal management. In *Proc. Intl' Symp. High Perf. Comp. Arch. (HPCA)*, pages 17–28, 2002.

- [15] J. Srinivasan and S. V. Adve. Predictive dynamic thermal management for multimedia applications. In *Proc. Intl' Conf. Supercomputing (ICS)*, pages 109–120, 2003.
- [16] F. Y. M. Wan. *Introduction to the Calculus of Variations and its Applications*. Chapman and Hall, 1995.
- [17] D. won Yoo and Y. K. Joshi. Energy efficient thermal management of electronic components using solid-liquid phase change materials. *Device and Materials Reliability, IEEE Transactions on*, 4(4):641–649, 2004.

APPENDIX

Derivation of the performance optimal speed profile: A necessary condition to extremize the objective (8) is that its first variation δJ must vanish.⁴ As the system equation (9) is satisfied at every instant over $[0, t_m]$, its variation must also vanish, or $\delta\psi = 0$. Now, we can use a Lagrangian multiplier function $\lambda(t)$ to form the augmented performance index

$$I = J - \int_0^{t_m} \lambda \psi \, dt = \int_0^{t_m} (-u - \lambda \psi) \, dt - u_m (T - t_m) \quad (20)$$

As $\delta J = \delta\psi = 0$, we must have $\delta I = 0$, or

$$\delta I = \int_0^{t_m} [\{ -\delta u \} - \lambda \{ -RS \delta u + RC \delta \theta' + \delta \theta \}] \, dt + u_m \, dt_m = 0$$

It can be shown that on integrating by parts, and noting that $\delta\theta' = (\delta\theta)'$, the above equation reduces to

$$\begin{aligned} \delta I = & \int_0^{t_m} [\{ -1 + \lambda RS \} \delta u + \{ -\lambda - (-RC\lambda') \} \delta \theta] \, dt - u_m \, dt_m \\ & + [-u(t_m) - \lambda(t_m) \{ \theta_{\max} - RP(u(t_m)) \}] \, dt_m = 0 \end{aligned}$$

where the last term with dt_m arises because t_m is a variable endpoint [16]. As δI must vanish for all variations δu , $\delta \theta$ and dt_m , we require their respective coefficients to vanish. This results in the following the two Euler-Lagrange equations

$$S(u) = 1/(R\lambda), \quad \lambda = \alpha e^{t/(RC)},$$

and a transversality condition

$$u(t_m) - u_m = \alpha [RP(u(t_m)) - \theta_{\max}] e^{t_m/(RC)} \quad (21)$$

The two Euler-Lagrange equations reduce to

$$S(u(t)) = [1/(\alpha R)] e^{-t/(RC)}, \quad (22)$$

which gives us the form of the optimum solution. To determine the actual optimum speed profile, we must determine two unknowns α and t_m , and hence, we need two equations. One equation is provided by the transversality condition (21), and another is obtained by the condition [16] that the variable involved in the inequality constraint θ must be continuous at the switch point, or $\theta(t_m) = \theta_{\max}$.⁵ Note that $u(t_m) = u_m$ always satisfies the transversality condition, as $P(u_m) = P_m = \theta_{\max}/R$.

⁴The variation can be considered the variational analogue of the differential operator in differential calculus. The theoretical underpinnings of the variational operator have been explained in [16].

⁵This makes intuitive sense as the temperature (unlike the speed) cannot vary suddenly due to the non-zero thermal capacitance.