# On Estimating Optimal Performance of CPU Dynamic Thermal Management

Aviad Cohen, Lev Finkelstein, Avi Mendelson, Ronny Ronen, Dmitry Rudoy
Microprocessor Research Intel Labs, Haifa, Israel
{aviad.cohen,lev.finkelstein,avi.mendelson,ronny.ronen,dmitry.rudoy}@intel.com

*Abstract*— **In this paper we focus on dynamic thermal management (DTM) strategies that use dynamic voltage scaling (DVS) for power control. We perform a theoretical analysis targeted at estimating the optimal strategy, and show two facts: (1) when there is a gap between the initial and the limit temperatures, it is best to start with a high (though not necessarily maximal) frequency and decrease it exponentially until the limit temperature is reached; (2) when being close to the limit temperature, the best strategy is to stay there. We use the patterns exhibited by the optimal strategy in order to analyze some existing DTM techniques.**

*Index Terms*— **DTM, DVS, optimal control**

## I. INTRODUCTION

Modern microprocessor designs are becoming *thermally limited*. This means that the power densities generated by various workloads may limit maximal operating frequency and necessitate prohibitively expensive thermal solutions. The predicted increase in power densities and the steeply rising cost of the required thermal solutions are expected to exacerbate this trend in future designs [1][2].

The gap between peak power and typical power [2] has led to the following approach: avoid the high cost of thermal solutions by designing them for typical power, and rely on *dynamic thermal management (DTM)* mechanisms to handle peak power. In recent years DTM is becoming an integral part of modern microprocessors (e.g. [2][3]).

A DTM mechanism monitors the maximal junction temperature by means of on-die thermal sensors and regulates power accordingly. Since all power regulating mechanisms slow down some activities in the processor when it gets hot, they inevitably entail a performance loss. Hence optimization of DTM performance loss has become a research focus [4][5][6][7][8][9].

We may discern two ingredients in a DTM technique: a *response mechanism*, which is the means for changing the input power (e.g. DVS), and a *control strategy*, which determines when the response mechanism is invoked and what its magnitude is. In seeking to optimize DTM one may focus on either ingredient. In this paper we are concerned with the control strategy. With regard to the response mechanism, we will just remark that various microarchitectural and software-based response mechanisms have been considered as alternatives to DVS [3][4][7][8][9][10].

The typical control strategies used in the industry are based on an ON/OFF scheme: a power decreasing response mechanism is turned on as soon as the temperature reaches a fixed threshold, and turned off when some indication that the chip has cooled off is registered. In seeking to improve this scheme, several different directions have been taken. [7] proposed a *predictive* DTM strategy which determines the response mechanism based on a prediction of the future power profile. However, this strategy is tailored for multimedia applications and remains to be extended to general workloads. Within the mainstream group of *reactive* strategies, i.e., strategies which determine the response mechanism based on the current temperature, two schemes have been suggested: *multi-response control* [5] and *PID control* [6]. Both vary the response according to the current temperature. The multi-response control does so by applying a dynamically changing set of power decreasing response mechanisms simultaneously, whereas PID controllers change continuously the magnitude of a single response with the goal of maintaining the temperature close to the limit.

The growing focus on DTM performance leads naturally to the following questions: Are there optimal strategies? What is their form? Note, that while an optimal strategy may not necessarily be implementable, it can still be used to analyze existing strategies, understand their limitations, suggest improvements, and provide guidelines for the design of new ones. We can (1) evaluate the performance of existing (or new) strategies by comparing their performance to the (computed) optimum, and (2) analyze strategies by checking whether their reactions to various events (e.g., a sudden drop in temperature as a consequence of a throttling activity or a sudden drop in workload power) follow the reaction patterns exhibited by the optimal strategy.

The goal of this paper is to present a theoretical analysis of the optimal strategy in the case of DVS response mechanism, and show how the conclusions of the analysis can be used to evaluate existing DTM strategies and indicate potential improvements. The key idea is to try to formulate the DTM scenario as an optimization problem and then use the theory of *state-constrained optimal control* to seek an optimal strategy that would keep the microprocessor below the limit temperature while maximizing performance.

We show the existence of an optimal strategy under some commonly used assumptions – the system's thermal behaviour is modeled by a thermal resistor-capacitor pair; the number of cycles is taken as a measure of performance; and the

Manuscript submitted: 9 Sept. 2003. Manuscript accepted: 8 Oct. 2003. Final manuscript received: 15 Oct. 2003.

power is proportional to some power of the frequency. We characterize the optimal strategy and establish in particular that (1) when there is a gap between the initial and the limit temperatures, it is best to start with a high (though not necessarily maximal) frequency and decrease it exponentially until the limit temperature is reached, and (2) when being close to the limit temperature, the best strategy is to stay there rather than apply short bursts of high frequencies interspersed with intervals of low frequencies. Our model assumes an ideal environment that allows continuous voltage and frequency scaling.

We evaluate the PID control strategy and the multi-response strategy by comparing them to the optimal strategy. We conclude that PID control follows the patterns exhibited by the optimal strategy. Based on the analysis, we suggest a potential improvement in the implementation of [5].

For want of space we cite our results without proofs. These will appear in a more detailed work.

## II. DTM AS AN OPTIMIZATION PROBLEM

In this section we set up a mathematical model for the analysis. The microprocessor is considered as a thermal system whose input is the generated power, and whose output is the maximal junction temperature. We consider a fixed time interval $[0, t_f]$ and assume that in this interval the workload produces a constant power profile[1]. The initial junction temperature is $T_0$, and the limit junction temperature is $T_m$. The response mechanisms is taken to be DVS. The voltage and the frequency are varied simultaneously in order to control the power, but since their values are related, we choose the frequency $f$ as the control and require that it be bounded by a maximal value $f_m$. The control strategy is specified by presenting the frequency as a function of time $f(t)$.

Our goal is to seek among the set of all strategies $f(t)$ that keep the system below the limit temperature $T_m$, the one which yields the best performance. To achieve this we must first specify in precise terms:

1) The relation between the control $f(t)$ and the power $P(t)$.
2) The relation between the control $f(t)$ and the performance.
3) A model for the thermal behaviour of the system from which one can obtain for each power profile $P(t)$ the junction temperature profile $T(t)$.

To specify 1, we assume that the power of the system is proportional to the square of the voltage and to the frequency, i.e., $P \sim V^2 f$ (we neglect leakage power in this paper). We assume also that the voltage is proportional to some power of the frequency, i.e., $V \sim f^\beta$, where $\beta \geq 0$ is an exponent which depends on the technology and the control strategy. Thus, we obtain $P = kf^\alpha$, where $\alpha = 1 + 2\beta \geq 1$ and $k$ is a proportionality constant[2].

To specify 2, we take the performance as a number of cycles in the interval $[0, t_f]$, which can be written as $\int_0^{t_f} f(t)dt$.

Finally, to specify 3, we model the CPU thermal behaviour by an RC thermal pair connected to a fixed temperature source (the ambient temperature) as depicted in Figure 1. Without loss of generality, we take the ambient temperature to be zero. Applying Kirchoff and Ohm laws, we obtain

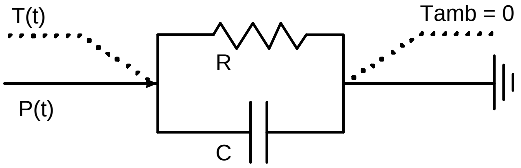

Fig. 1. The RC-pair thermal model

$$\frac{dT}{dt} = \frac{P}{C} - \frac{T}{\tau} , \qquad \text{where } \tau = RC.$$

Note that in this model, once the system is at temperature $T_m$, it can be maintained there by applying the frequency $\sqrt[\alpha]{T_m/kR}$. We denote this frequency by $f_{natural}$. Note that if $f_m \leq f_{natural}$, then the optimal strategy is to run always at $f_m$. Therefore we assume $f_m > f_{natural}$.

Thus, we arrive at the following optimization problem:

$$\int_0^{t_f} f(t)dt \rightarrow \max, \tag{1a}$$

$$T(0) = T_0, \tag{1b}$$

$$\frac{dT}{dt} = \frac{kf^\alpha}{C} - \frac{T}{\tau}, \tag{1c}$$

$$f(t) \leq f_m, \quad T(t) \leq T_m. \tag{1d}$$

This formulation is an instance of the optimization problems addressed by the theory of *state-constrained optimal control*. This theory deals with dynamical systems with a state $X(t) \in \Re^n$ whose rate of change $\frac{dX}{dt}$ is a function of its own current value $X(t)$, of an independent input $U(t) \in \Re^m$ called the *control*, and possibly of the time $t$. There may exist constraints imposed on either the control $U$ or the state itself (for example, they may be required to be bounded). The goal of optimal control is to find among the set of controls satisfying the constraints a member that maximizes the time integral of some *criterion function* $\Omega$ which depends on $X$, $U$ and possibly on $t$. In mathematical notation, we can write it as follows:

$$\int_0^{t_f} \Omega(X(t), U(t), t)dt \rightarrow \max, \tag{2a}$$

$$X(0) = X_0, \tag{2b}$$

$$\frac{dX}{dt} = \Psi(X(t), U(t), t), \tag{2c}$$

$$C(U(t), t) \geq 0, \quad S(X(t), t) \geq 0, \tag{2d}$$

where (2a) stands for the optimization target, (2b) for the initial conditions, (2c) for the state evolution, and (2d) for the control and the state constraints respectively. It can be seen that the optimization problem (1) is an instance of (2).

Optimal control theory supplies theorems for the existence and characterization of optimal solutions (for a detailed survey see [11]), and these can be used to solve our problem.

---

[1] This is not a restricting assumption since the power profile of benchmarks can usually be broken into time intervals along which power is approximately constant.

[2] Strictly speaking, the case of $\alpha = 1$ is referred to as "clock throttling" [2] rather than DVS, but it is subsumed under our analysis as a special case.

## III. FORMAL ANALYSIS

In this section we describe the solution of the optimization problem (1). The solution was obtained by applying the theory of optimal control. In all the figures appearing here, solid lines show the frequency as a function of time whereas dashed lines represent the corresponding evolution of the temperature. In all figures we take $T_0 = 0$ and $T_m = 50$, and unless otherwise stated, $\alpha = 3$.

We start with the clock throttling case ($\alpha = 1$) which turns out to be simpler. The behaviour of the optimal strategy is described by the following theorem:

*Theorem 1 (constant voltage optimal strategy):* The optimal strategy runs with $f = f_m$ until $t = t_f$ or $T = T_m$, whichever occurs first. In the second case, it continues with $f = f_{natural}$, thus maintaining $T = T_m$, until the end of the interval.

Theorem 1 is illustrated in Figure 2.

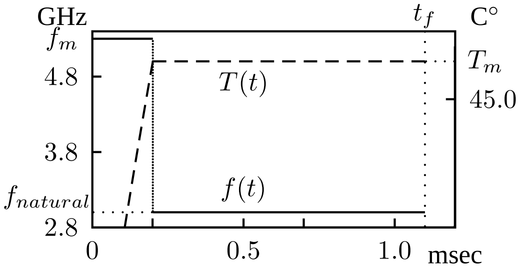

Fig. 2. Optimal strategy for clock throttling

The DVS case ($\alpha > 1$) is more complex. The optimal strategy depends on $T_0$ and $t_f$, and is therefore denoted by $\psi(T_0, t_f, t)$. [3] In unraveling its structure it seems best to study first the case where there is no upper bound on the frequency, and then move on to the general case. The optimal strategy for the unconstrained case is denoted by $\psi_u(T_0, t_f, t)$.

The form of $\psi_u(T_0, t_f, t)$ depends on whether the time interval is long or short. We consider the time interval to be long if it is greater than a certain *threshold* value $t_{th}$ which depends on $T_0$ (we indicate how to compute $t_{th}$ in Theorem 2). The case $t_f > t_{th}$ is shown in Figure 3. The frequency starts at a value $f_0$, and decreases exponentially with the factor $e^{-\frac{t}{(\alpha-1)\tau}}$ until $t = t_{th}$. The initial frequency $f_0$ is so chosen that the temperature reaches $T_m$ for the first time at $t = t_{th}$ (we indicate how to compute $f_0$ in Theorem 2) . Then in the time interval $[t_{th}, t_f]$ the temperature is maintained constant at $T = T_m$ by applying the frequency $f_{natural}$. Note that the optimal strategy is *continuous*.

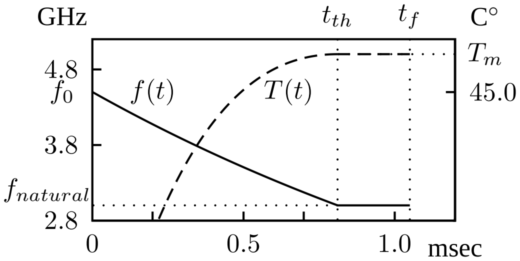

Fig. 3. Unconstrained optimal strategy for $t_f > t_{th}$

Observe also, that the exponential decrease of the strategy appears almost linear. This is due to the fact that $t_{th}$ is bounded by $(\alpha-1)\ln(\alpha/(\alpha-1))\tau$ as can be seen from the formula for $t_{th}$ given in Theorem 2. This implies that for the case $\alpha = 3$ which is shown in the figure, the exponent $\frac{t}{(\alpha-1)\tau}$ ranges from 0 to at most $\ln(\alpha/(\alpha-1)) \sim 0.4$. Using the Taylor expansion of $e^{-x}$ it can be verified that in this range $e^{-\frac{t}{(\alpha-1)\tau}}$ can be approximated by $1 - \frac{t}{(\alpha-1)\tau}$ with an average relative error bounded by $0.04$. [4]

The case $t_f \le t_{th}$ is shown in Figure 4. The frequency starts at a value $f_0$ and decreases exponentially with the factor $e^{-\frac{t}{(\alpha-1)\tau}}$ until the end of the interval $t = t_f$. The initial frequency $f_0$ is so chosen that the temperature reaches $T_m$ for the first time at $t = t_f$.

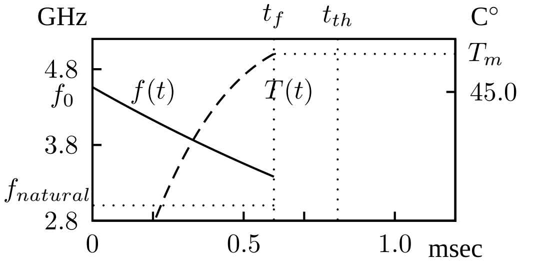

Fig. 4. Unconstrained optimal strategy for $t_f < t_{th}$

All this is summed up in the following theorem:

*Theorem 2 (unconstrained optimal strategy for $\alpha > 1$):* Let $t_{th}$ be given by

$$t_{th} = -\frac{\alpha-1}{\alpha}\tau\ln(x)$$

where $x$ is the unique real number satisfying

$$x - \left(\frac{1}{\alpha}\left(\frac{T_0}{T_m}\right)x + \frac{\alpha-1}{\alpha}\right)^{\alpha} = 0, \qquad x \le 1.$$

and set $t^* = \min(t_f, t_{th})$. The unconstrained optimal strategy is

$$\psi_u(T_0, t_f, t) = \begin{cases} f_0(T_0, t^*) \cdot e^{-\frac{t}{(\alpha-1)\tau}}, & t \in [0, t^*), \\ f_{natural}, & t \in [t^*, t_f], \end{cases}$$

where $f_0(T_0, t^*)$ is given by

$$f_0(T_0, t^*) = \left(\frac{1}{(\alpha-1)Rk} \cdot \frac{T_m - T_0 e^{-\frac{t^*}{\tau}}}{e^{-\frac{t^*}{\tau}} - e^{-\frac{\alpha}{\alpha-1}\frac{t^*}{\tau}}}\right)^{1/\alpha}$$

*Corollary 1:* The number of cycles for the optimal unconstrained strategy $c_u(T_0, t_f)$ is

$$\int_0^{t_f} \psi_u(T_0, t_f, t)dt = (\alpha-1)\tau \cdot f_0(T_0, t^*) \cdot \left(1 - e^{-\frac{t^*}{(\alpha-1)\tau}}\right) + f_{natural} \cdot (t_f - t^*)$$

Note that if $T_0 = T_m$ then $t^* = 0$, and we obtain the following corollary:

---

[3] $\psi$ depends also on $T_m$ and $R, C, k$, but these are assumed fixed for a given system.

[4] Since the Taylor series of $e^{-x}$ is alternating, we have $|e^{-x} - (1-x)| < \frac{x^2}{2}$ so the average error is bounded by $\frac{1}{0.4}\int_0^{0.4} \frac{x^2}{2}/(1-x)dx$.

*Corollary 2:* When being at the limit temperature $T_m$, the optimal strategy is to stay there.

Let us now consider the general case. If the initial frequency prescribed by $\psi_u$ does not exceed $f_m$, then $\psi_u$ is the optimal strategy. Otherwise, the optimal strategy starts by running at $f_m$. If the time interval is short enough, $f_m$ is maintained until $t = t_f$ without violating the temperature constraint. Otherwise the strategy has the form depicted in Figure 5. The optimal strategy proceeds with $f = f_m$ until it reaches a time $t_{\text{sw}}$, at which the initial frequency prescribed by the unconstrained strategy is exactly equal to $f_m$. At this point we switch to the unconstrained optimal strategy. This is summed up by the following theorem:

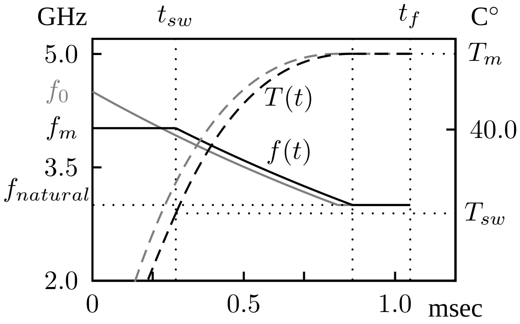

Fig. 5. Optimal strategy - black curves stand for general case, grey curves for the unconstrained case

*Theorem 3 (constrained optimal strategy):* There exists a switch time $t_{\text{sw}} \in [0, t_f]$, such that the optimal strategy for the general case $\psi(T_0, t_f, t)$ is given by

$$\psi(T_0, t_f, t) = \begin{cases} f_m, & t \in [0, t_{\text{sw}}), \\ \psi_u(T_{\text{sw}}, t_f - t_{\text{sw}}, t - t_{\text{sw}}), & t \in [t_{\text{sw}}, t_f], \end{cases}$$

where $T_{\text{sw}}$ is the temperature at $t = t_{\text{sw}}$.

*Corollary 3:* The optimal number of cycles at the general case $c(T_0, t_f)$ is

$$\int_0^{t_f} \psi(T_0, t_f, t)dt = t_{\text{sw}} \cdot f_m + c_u(T_{\text{sw}}, t_f - t_{\text{sw}})$$

## IV. DISCUSSION AND CONCLUSIONS

In this paper we characterize the optimal strategy for DVS-based DTM, assuming a simple thermal model and using the number of cycles as a measure of performance. In general, the optimal strategy may consist of up to three stages. In the first stage $f = f_m$. In the second stage $f$ decreases exponentially, but in many cases the decrease can be approximated by a linear function. In the third stage $f = f_{natural}$. Depending on $T_0$, $t_f$ and $\alpha$, some of the above stages may not be present. We also provide formulae for the maximal number of cycles, and this can be used to estimate the distance from optimality.

Our analysis lends a theoretical support to the PID approach since PID control seems to follow the patterns of the optimal strategy. First, PID control reacts to a temperature drop below the allowed limit by increasing the power to a high value and then decreasing it gradually as the temperature rises. Second, once the system gets close to the limit temperature, the PID control would strive to maintain the temperature there.

The multi-response control as implemented in [5] reacts in a different way. When the temperature drops below some threshold, the controller may actually increase the power gradually, removing each time one power decreasing response mechanism. The following modification would make it follow the pattern of the optimal strategy: upon a temperature drop, remove all power decreasing response mechanisms, and then add them one by one as the temperature rises. This may be considered as a potential improvement.

This paper focuses on a methodology for a theoretical analysis of optimal DTM and demonstrates its application with some simplifying assumptions. Even in this form the analysis may yield preliminary design guidelines and initial estimates for optimal DTM (by using effective system parameters such as $k$, $R$, and $C$). The theoretically optimal behaviour may then be used as a reference in the analysis of various DTM techniques. For a more accurate evaluation one needs to refine the thermal model and incorporate leakage into the analysis. This may lead to a more complex setup and necessitate using numerical optimization algorithms in conjunction with the theoretical approach. The effect of some of the assumptions, such as assuming a constant power profile, is unclear, and needs to be analyzed in detail.

Future directions of our research consist of extending the results to more accurate thermal models and evaluating the proximity of existing strategies to the optimum.

## ACKNOWLEDGMENT

The authors would like to thank Barukh Ziv for many insightful discussions and for his help in the preparing this document.

## REFERENCES

[1] S. Borkar, "Design challenges of technology scaling," *Micro IEEE*, vol. 19, no. 4, pp. 23–29, Jul-Aug 1999.
[2] S. H. Gunther, F. Binns, D. M. Carmean, and J. C. Hall, "Managing the impact of increasing microprocessor power consumption," *Intel Technology Journal*, vol. Q1, 2001.
[3] H. Sanchez, B. Kuttanna, T. Olson, M. Alexander, G. Gerosa, R. Philip, and J. Alvarez, "Thermal management system for high performance PowerPC microprocessors," in *Proceedings of IEEE Compcon'97 Digest of Papers*, 1997, pp. 325–330.
[4] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in *HPCA'01*, January 2001, pp. 171–182.
[5] M. Huang, J. Renau, S. M. Yoo, and J. Torrellas, "The design of DEETM: a framework for dynamic energy efficiency and temperature management," *Journal of Instruction Level Parallelism*, vol. 3, 2002.
[6] K. Skadron, T. Abdelzaher, and M. R. Stan, "Control-theoretic techniques and thermal RC modeling for accurate and localized dynamic thermal management," in *Proceedings of HPCA'02*, February 2002, pp. 17–28.
[7] J. Srinivasan and S. V. Adve, "predictive dynamic thermal management for multimedia applications," in *Proceedings of ICS'03*, June 2003, pp. 23–26.
[8] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *Proceedings of ISCA'03*, June 2003.
[9] E. Rohou and M. D. Smith, "Dynamically managing processor temperature and power," 2nd workshop on feedback-directed optimization, November 1999.
[10] C. H. Lim, W. R. Daasch, and G. Cai, "A thermal-aware superscalar microprocessor," in *Proceedings of ISQED'02*, March 2002, pp. 517–522.
[11] R. F. Hartl, S. P. Sethi, and R. G. Vickson, "A survey of the maximum principles for optimal control problems with state constraints," *SIAM Review*, vol. 37, no. 2, pp. 181–218, 1995.