

A New Thermal-Conscious System-Level Methodology for Energy-Efficient Processor Voltage Selection

Yongpan Liu, Yu Wang, Feng Zhang, Rong Luo, Hui Wang
Department of Electronic Engineering, Tsinghua University, Beijing, 100084, P.R.China
ypliu99@mails.tsinghua.edu.cn

Abstract— In this paper, we propose a thermal-conscious system-level methodology to make energy-efficient voltage selection (VS) for nanometer processors under real-time constraints. New modeling parameters, such as power composition ratio (PCR), thermal resistance, are integrated and considered in our system models, and their impacts on energy consumption are explored. The interdependence between temperature and power is modeled in an iterative way and two nonlinear programming formulations are presented to determine the optimal energy-efficient supply voltages. Our experiment results show that the energy estimation discrepancy given by the thermal-conscious and traditional system models can reach up to 50% in 65nm *PTM* CMOS technology, which again underscores the necessity of a thermal-aware power model. Furthermore, our thermal-conscious voltage selection (TCVS) approach can achieve up to 12% further energy savings and much lower peak temperature than the traditional approach. Finally, our specific temperature voltage selection (STVS) approach can reduce the computation complexity greatly with acceptable energy overheads compared with TCVS.

I. INTRODUCTION

Energy reduction techniques, such as dynamic voltage scaling (DVS), and dynamic power management (DPM) [1-5] are widely adopted in battery-powered digital systems in order to acquire a longer duty life. With the semiconductor manufacturing technology coming into nanometer era, the ratio of leakage power in total power consumption increases steadily. According to Intel's data [6], leakage power contributes about 40% in its latest high performance microprocessor. Since the processor's operation temperature may vary greatly due to different power profiles of tasks running on it, leakage power also varies greatly because of its exponential relationship with temperature. Therefore, the optimal energy-efficient operation voltage would change according to different leakage power under corresponding processor temperatures, which makes temperature dependent power models a must for energy-efficient voltage selection of each task.

Plenty of research works have been done to deploy DVS techniques for single or multiple DVS-enable processors aiming at energy minimization. [1] proposed a joint approach to combine both DVS and ABB (Adaptive body biasing) techniques so as to reduce both dynamic and leakage power consumption of a low power microprocessor effectively. [2] combined DVS with procrastination to minimize the total energy consumption. [4] used mathematics programming approaches to choose optimal energy-efficient supply and body biasing voltage. Given the power-composition profile of each task, [3] presented a co-synthesis methodology to reduce the system implementation cost of DVS

and ABB. However, all above research works neglect the thermal effects on leakage power consumption during task execution, which may lead to great energy and temperature estimation errors in nanometer processors.

Recently, [5] presented a temperature dependent power model at micro-architecture level, and discussed the optimal supply voltage for maximum performance under different cooling conditions. However, neither energy-related issues, such as DVS, nor VS problems under real-time constraints are covered. In this paper, to our best knowledge, we are the first to propose a system-level methodology for optimal VS based on thermal-conscious system model. Our main contributions are listed as below:

I). A novel thermal-conscious system modeling methodology is proposed to evaluate thermal effects on energy consumption and VS, which includes iterative power, thermal and delay models. It is shown that up to 50% energy consumption is underestimated if thermal effect is omitted.

II). New modeling parameters, such as *PCR* and thermal resistance are integrated and considered in our system models, and their impacts on energy consumption and VS are also explored.

III). Based on our thermal-conscious system models, two voltage selection procedures: TCVS and STVS are proposed, which are both formulated as nonlinear programming problem. According to our experiment results, TCVS could lead to up to 12% further energy savings and lower peak temperature compared with traditional VS method. Meanwhile, STVS reduces computation complexity greatly with acceptable energy overheads compared with TCVS.

This paper is organized as below: the motivation of this work is given out in section 2; the formulation of voltage selection problem and system-level models are then illustrated in section 3. Section 4 gives out two energy-efficient voltage selection procedures. Finally, experiment results and future work are reported in section 5.

II. MOTIVATION

Previous VS research works [1-4] didn't consider the operational temperature for tasks, thus they omitted the thermal effects on task's energy estimation and optimization. It may lead to quite large energy estimation deviations in future nanometer processors, where leakage power contributes a substantial part of total power budget. As we can see from [10], the power profile of different tasks for Intel's latest Montecito processor has a very large power variation, from 68W to 126W. Assuming a typical air cooling condition [9], Fig.1 gives out the temperature curve of the processor executing a task set, which contains different tasks with random power profile ranges from 10 to 130W. It is shown that the processor's temperature varies in a range from 60 to 110°C. Such big temperature variations actually lead to quite different leakage

Supported by National NSF of China (No. 90207001, No. 60506010) and National 863 project of China (No. 2005AA1Z1230)

power consumption and accordingly affect the energy-efficient VS. Therefore, thermal-conscious model must be used in order to acquire accurate energy estimation and VS.




Fig.1 Thermal profiles of running a task set

In traditional VS problem, each task is assumed to be operated at the same temperature under different supply voltages, thus the interdependence between power and temperature is omitted. Based on iterative models, which accurately describe one task's power and thermal profiles, Fig.2 has illustrated the thermal-conscious power profile of one task at different supply voltages. As we can see, the cross and square curves give out the upper and lower bounds of traditional leakage power trends assuming constant temperatures, respectively. The dot curve gives out the accurate iterative leakage power profile. There are a big difference between the traditional trend and the iterative one.

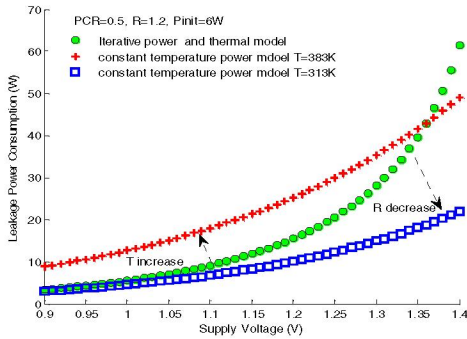


Fig.2 Thermal-conscious leakage power vs supply voltage

In this paper, a thermal-conscious system model is proposed to revisit the traditional VS problem, and thermal impacts on it will be thoroughly examined. As [8] has shown, typical tasks' power spectrum mainly focuses on a fixed DC component, which can be considered to keep constant during its execution. Therefore, it is firstly assumed that each task's operational temperature only varies within a small margin during its execution, thus we can use the average temperature value to calculate each task's energy consumption without sacrificing its accuracy. Secondly, each task's execution time is assumed at least two magnitudes of thermal constant, which locates in the microsecond magnitude [9], so the processor reach its steady temperature most of the execution time. We may satisfy this by choosing appropriate task granularity.

III. PRELIMINARIES

This section presents a thermal-conscious system-level methodology for energy and temperature analysis, and contains the following subsections: VS formulation, power and delay model, thermal model and iterative modeling method.

A. Voltage Selection Formulation

The energy-efficient VS problem is to find out an optimal operational voltage for each task assigned to certain processor so as to minimize the total energy of multiprocessor system, which is shown in Fig.3. Here, a key difference from traditional approaches is that thermal-conscious parameters, such as T , PCR , R (refer to the following section), are integrated in system models.




Fig.3 Voltage selection Formulation

The task graph is denoted as $G = (V, E)$, which is a directed acyclic graph composed of a set of N vertices $V = \{v_1, v_2, \dots, v_N\}$, with each vertex representing a specific task of the graph G , and a set E of edges where each edge is denoted by an ordered pair of tasks (v_i, v_j) , which represents the precedence constraint between these two tasks. For each task v_i , four energy related parameters are defined: the number of execution cycles EC_i , the average switching probability d_i , the operational temperature T_i , and the power composition ratio $PCR_i^{T=T_0, V=V_{min}}$, which is defined as the ratio between leakage and total power at minimal voltage and ambient temperature for a task. Furthermore, there is a deadline $deadline_i$ associated with each leaf node v_i in the task graph G , which specifies the real-time constraints. For each processor P_p , there are also five technology parameters, namely total switch capacitance, peak running frequency, thermal resistance, maximal and minimal operating voltage, denoted as C_{eff}^p , f_{max}^p , R_p , V_{max}^p and V_{min}^p , respectively.

B. Power and Delay Model

Next, we will derive the temperature dependent power model for DVS-enabled processors. The dynamic power consumption of a CMOS microprocessor is temperature insensitive [1], thus it can be acquired by the following formula:

$$P_{dyn} = C_{eff} \cdot V_{dd}^2 \cdot f^i \quad (1)$$

where C_{eff} is the effective switch capacitance, f^i is the processor running frequency. Due to our simulation reports [11], the leakage power model is based on the following formula:

$$P_{leakage} = V_{dd} \cdot I_s(T_0, V_{dd}^{min}) \cdot (A \cdot T^2 e^{(\alpha V_{dd} + \beta V_{th} + \gamma)/T} + B \cdot e^{-\mu V_{dd}}) \quad (2)$$

, where I_s is the average leakage current under ambient temperature and minimum voltage, T is the operational temperature of the processor, which will be discussed in the next section and other parameters are empirical constants for different circuit types, technologies and designs. All of the technology constants are acquired by using the curve-fitting method. In our experiments [11], 65nm *PTM* models are adopted to derive those constants. The delay of VLSI circuits is decided by the following formula [1]:

$$f^{-1} = t = \frac{L_d K_6 V_{dd}}{(V_{dd} - V_{th})^\sigma} = \frac{L_d K_6 V_{dd}}{((1 + K_1) \cdot V_{dd} + K_2 \cdot V_{bs} - V_{th})^\sigma} \quad (3)$$

C. Thermal Model

According to the heat transfer theory, thermal phenomena can be modeled as an equivalent RC thermal circuit. The R and C values for a specific cooling configuration are derived using Hotspot [9]. The ambient temperature of the processor is denoted as T_A , and T_{init} is the initial temperature of the processor when executing the i th task: $task_i$, we derive the following formula to calculate the operational temperature T_i :

$$T_i = T_A + (T_{init} - T_A) \cdot e^{-\frac{t}{RC}} + P_i \cdot R \cdot (1 - e^{-\frac{t}{RC}}) \quad (4)$$

where P_i is the average power consumption of $task_i$. As section 2 has illustrated, assume that $t \gg RC$, the following formula can be derived:

$$T_i = T_A + P_i \cdot R \quad (5)$$

Since the temperature distribution in a processor is variable, the temperature in the formula should be an average one. It will introduce (<5%) leakage power estimation error due to detail simulation results [11].

D. Iterative Modeling

Given initial dynamic and static power consumption, each processor converges to a steady temperature under its specific thermal condition. Since operational temperature for each task varies greatly, traditional approach leads to quite large estimation error, which may generate suboptimal VS solution. In this paper, we propose an accurate iterative approach to calculate the steady power consumption and temperature. Since it is encapsulated in TCVS procedure, it will be explained in the next section.

IV. THERMAL-CONSCIOUS VOLTAGE SELECTION

In this section, we will give out two VS formulations based on the thermal-conscious system models.

A. TCVS

The energy-efficient voltage selection problem under real-time constraints is formulated as the following expression.

$$\left\{ \begin{array}{l} \text{Min} \sum_i E_{tot}^i(V_{dd}^i) = P_{tot}^{i(n)} \cdot exec^i \\ \text{s.t.} \\ \text{c1: } exec^i = EC^i(f^i)^{-1} = \frac{EC^i \cdot K_3 \cdot V_{dd}^i}{((1 + K_1) \cdot V_{dd}^i + K_2 \cdot V_{bs}^i - V_{th})^\sigma} \\ \text{c2: } \begin{cases} \text{for}(k=1, T^{(0)} = T_{ambient}; k < n+1; k++)\{ \\ P_{tot}^{i(k)} = \alpha^i \cdot C_{eff}^j \cdot (V_{dd}^i)^2 \cdot f^i + V_{dd}^i \cdot I_{sg}^{i(k)}; \\ I_{sg}^{i(k)} = I_{nor} \cdot [A \cdot (T^{i(k-1)})^2 \cdot e^{(\alpha V_{dd}^i + \beta V_{th}^i + \gamma) / T^{i(k-1)}} + B \cdot e^{\mu V_{dd}^i}]; \\ T_{me}^{i(k)} = T_A + P_{tot}^{i(k-1)} \cdot R_j; \} \end{cases} \\ \text{c3: } start^i + exec^i \leq start^{i+1} \\ \text{c4: } start^i + exec^i \leq deadline^i \\ \text{c5: } start^i \geq 0 \\ \text{c6: } V_{dd}^{\min} \leq V_{dd}^i \leq V_{dd}^{\max} \end{array} \right. \quad (6)$$

The optimization variables for this problem are the task execution times $exec^i$, the task start times $start^i$, and the operational voltages V_{dd}^i . The goal of this problem is to minimize the total execution energy of all tasks. For each task, constraint c1 is based on the delay model, which gives out the relationship between its

execution time and supply voltage. Constraint c2 describes the iterative relationship between power and temperature, where I_{nor} is a real number for each task to make the ratio between dynamic power and leakage power equal to PCR at minimal voltage V_{min} and ambient temperature $T_{ambients}$. R_j is the thermal resistance of the assigned processor, and n is a user-specific integer number. The processor's steady power and temperature when executing $task_i$ will be given out after n th iterations. In our experiments, n is usually less than 5 and only 0.01% errors are observed. Constraint c3 gives out the precedent relationships among tasks in the task set. Since the processor can only execute one task at the same time, resource constraints between two tasks assigned to the same processor are also considered in c3. Constraint c4 satisfies the leaf nodes' deadline. Constraint c5 and c6 are the bounds of variables.

B. STVS

Next, we propose another formulation to the VS problem. The total VS problem is divided into two phases. In the first phase, a constant temperature is used to estimate the power consumption for each task, and then voltage selection for each task is done. In the second phase, the assigned supply voltage for each task is used as the input for iterative models to calculate its accurate energy consumption. Compared with TCVS, this approach uses the energy under a user-specific constant temperature, thus it may acquire a suboptimal voltage schedule. However, according to our experiment results, energy overheads are usually acceptable. Since the iterative energy and temperature estimation are removed from the nonlinear programming optimization procedure, it greatly reduces the NLP problem's computation complexity.

V. EXPERIMENT RESULTS AND ANALYSIS

To demonstrate thermal effects on the energy estimation and voltage selections, several experiments are used, and they are solved on a 1.4Ghz Centrino™ laptop running Linux. For the largest case, the run time does not exceed 300s.

The first experiment is to show thermal impacts on VS under different parameters. Supply voltage must be decided for two tasks running on two processors. Parameters for $task_1$ are set as: power=40W, $PCR=0.5$, and it is assigned to processor P_1 with $R=1.2$; $task_2$ has two variable parameters: namely power, PCR , and it is assigned to processor P_2 with variable thermal resistance; $task_2$'s execution can't start until the completion of $task_1$. Fig.4 gives out the total energy consumption trends under different parameters using different VS methods. Since there are three variables, namely power, PCR and R , only one is changeable in each diagram. The energy estimation using non-thermal-conscious (ENT) approach gives out the energy consumption by traditional approach. Since it omits thermal impacts on the energy consumption, it always gives out lower estimation. The curve here is mainly to show how large energy discrepancy between thermal-conscious and traditional model. As we can see, up to 50% energy estimation errors are observed, so it further underscores the necessity of thermal-conscious power modeling method. The STVS curve is drawn by doing voltage selection using a constant user-specific temperature for all tasks, thus its voltage schedule may be suboptimal. However, since its energy consumption would be recalculated by iterative thermal-conscious models, its energy estimation is accurate.

The most energy efficient VS results are given by TCVS curve because of its accurate thermal-conscious energy model. As we can see from Fig.4, TCVS is always better than STVS under all

different cases. In Fig.4b and Fig.4c, when the PCR (or R) difference between $task_2$ and $task_1$ is maximal, the largest energy savings can be acquired. Since only PCR (or R) is different between $task_1$ and $task_2$, in STVS both tasks are considered equally, while in TCVS, the task with large PCR (or R) has a higher priority to be scaled down, because of sharper energy increase with higher voltage. Therefore, the maximum difference in PCR (or R) leads to up to 12% energy savings. However, for the power parameter P , both STVS and TCVS set high priority to the task with higher initial power consumption, thus the energy difference is trivial.

Table.1 Energy Consumption under Different VS Approaches

	Nodes /Edges/PE	STVS /Tmax	TCVS /Tmax	ENT Errors	Further Savings
TS1	13/12/2	245099/362	243655/365	20.5%	0.6%
TS2	16/15/2	280488/414	279389/398	13.3%	0.4%
TS3	37/40/2	757749/456	747530/397	16%	1.3%
TS4	66/81/2	1534056/380	1518576/380	15%	1.0%

In order to illustrate the thermal effects in more general cases, we further testify our VS procedure on random generated benchmark set so as to show its effectiveness, the experiment results are listed in Tab.1. The task graph is generated by TGFF [12]. The PCR and R values of all tasks' are automatically generated random numbers with a uniform distribution, and they locate in the range of [0.05, 0.5] and [0.9, 1.0], respectively.

As we can see, the energy estimation error between ENT and TCVS is up to 20.5%. Furthermore, TCVS always gets further energy savings compared with STVS, but STVS provides a feasible VS solution with very small energy overheads. Since the PCR and R difference in the random generated benchmarks is not big enough, it makes the energy savings from TCVS not so large as those in the two-tasks case. In addition, TCVS approach can generate selection with lower peak temperature. We can see these in the cases of TS2 and TS3. It is due to the fact that TCVS adopts a thermal-conscious model for VS, which presents a sharp energy increase with higher temperature above certain inflexion. Therefore, TCVS assigns a lower voltage to that task to reduce the energy and accordingly acquires a lower operation temperature. While in STVS, the feedback of energy and temperature is not considered in VS, and the scheduler may assign equal supply voltage to tasks with same power but different PCR and R . Thus, processors become quite hot in some cases. Another interesting phenomenon is that though the maximum temperature of the

processor varies greatly according to different VS solution, their energy consumption is quite close. This reveals that a thermal-conscious optimization voltage selection with small energy consumption overhead can be achieved. Since temperature has an exponential impact on the system reliability [7], our TCVS is quite useful because it not only reduces the energy consumption, but also improves the reliability of nanometer processors.

VI. CONCLUSIONS

In this paper, we presented a thermal-conscious system-level methodology for VS problem. Iterative methods are used to accurately estimate task energy consumption and processor temperature. New modeling parameters, such as power composition ratio and thermal resistance are integrated and considered in our system models, and their impacts on energy consumption are explored. Based on the thermal-conscious model, energy-efficient voltage selection problem are solved by two nonlinear programming approaches: TCVS and STVS. Compared with traditional VS approaches, our methods can achieve more accurate energy estimation and further energy savings.

REFERENCES

- [1] Martin, M., etc: Combined dynamic voltage scaling and adaptive body biasing for lower power microprocessors under dynamic workloads. ICCAD2002, p 721-725.
- [2] Jejurikar, R; Pereira, C; Gupta, R.: Leakage aware dynamic voltage scaling for real-time embedded systems. DAC, June, 2004, p 275-280.
- [3] Wu, D. etc: Power-composition profile driven co-synthesis with power management selection for dynamic and leakage energy reduction. DSD, 2005, p 34-40.
- [4] Andrei, A. Schmitz, M.; Eles, P.; Peng, Z.; Al-Hashimi, B.M.: Overhead-conscious voltage selection for dynamic and leakage energy reduction of time-constrained systems. IEE Proceedings., v 152, n 1, January, 2005, p 28-38.
- [5] Liao, W.P.; He, L.; Lepak, K.M.: Temperature and supply voltage aware performance and power modeling at microarchitecture level, IEEE Transactions on CAD, v 24, n 7, July 2005, p 1042-53.
- [6] <http://www.intel.com>.
- [7] <http://public.itrs.net/Common/2005ITRS/Home2005.htm>.
- [8] Li, H.; Liu, P.; Qi, Z.Y.; Jin, L.L.; Wu, W., etc: Efficient thermal simulation for run-time temperature tracking and management, ICCD2005, p 130-133.
- [9] Skadron, K.; Stan, M. R.; Huang, W., etc.: Temperature-aware microarchitecture. Conference Proceedings, ISCA, 2003, p 2-13.
- [10] Naffziger, S.etc.: The implementation of a 2-core, multi-threaded titanium family processor. JSSC, vol. 41, no. 1, Jan. 2006, p197-209.
- [11] Zhang, F., Power modeling for microprocessors, BS thesis. Tsinghua Univ. 2006.
- [12] Dick, R., etc., TGFF: task graphs for free, Proc. CODES, pages 97-101, 1998.




Fig.4.a Energy vs power of $task_2$




Fig.4.b Energy vs PCR of $task_2$




Fig.4.c Energy vs R of $task_2$