

Energy optimization of parallel programs in a heterogeneous system by combining processor core-shutdown and dynamic voltage scaling

Zhuowei Wang,¹ Hao Wang,^{2,*} Wuqing Zhao,³ Lianglun Cheng,¹

¹ School of Computers, Guangdong University of Technology, Guangzhou 510006, China

² Department of ICT and Natural Sciences, Norwegian University of Science and Technology, Norway

³ Dingxin Information Technology Co., Ltd, Guangzhou, 510006, China

*Corresponding author: hawa@ntnu.no

Abstract: Reducing power consumption and improving efficiency are important aspects of the development of supercomputers into large-scale systems. As a result, heterogeneous systems have become an important development trend in high-performance computing. From the perspective of heterogeneous systems, this study establishes a model for energy optimization of parallel programs (EOPP) and puts forward a method of using it. By considering the energy overheads caused by re-synchronization, voltage switching, and operations in critical sections, the model effectively combines processor core-shutdown and dynamic voltage scaling technologies, which can be applied in a heterogeneous system to guide the optimization process. The results show that the proposed model can effectively reduce the energy consumption of parallel programs. Moreover, increasing the proportion of operations in the critical section enhances the optimal frequency of a processor while decreasing the probability of conflicts in the critical section. It can thus provide optimization space for reducing the frequency of a processor which ultimately reduces the energy overhead of the system.

Keywords: Energy optimization; Parallel program; Heterogeneous system; Processor core-shutdown; Dynamic voltage scaling

1. Introduction

Increasing growth in the demand for information is promoting a rapid development in high-performance computers. In order to enhance the performance of high-performance computer systems and solve large-scale computing problems, the development of heterogeneous systems which integrate general-purpose processors and accelerated processing units (APUs) has been one of the important trends in high-performance computing. Although heterogeneous systems show high peak computation speeds and efficiency, the problem of high power consumption still exists. High power consumption greatly challenges various aspects of large-scale heterogeneous and parallel computing systems (including reliability and heat dissipation) and causes significant consumption of energy. In this context, the power consumption problem is causing unprecedented concern [1].

Power optimization research encompasses a broad range of approaches ranging from hardware (bottom-level) to software (upper-level) [2–3]. Hardware approaches mainly include the use of low-power optimization technologies, e.g. using low-power logic devices [4-5], interconnected power optimization [6-7], control of leakage currents [8-9] and layout packaging [10-11]. Hardware optimization methods have reached a certain level of maturity and it is hard to further tweak them to satisfy the increasing demand for power optimization. Therefore, concern has shifted to power optimization at the software level [12–15].

Dynamic voltage scaling and processor core shutdown are two commonly used technologies to reduce the energy consumption. By reducing the voltage to reduce the energy consumption, Dynamic voltage scaling technology mainly uses the relationship between the dynamic power and the square of the voltage. The processor core shutdown technology is to reduce the energy consumption by closing the processor in the idle state. In recent

1 years, a large number of international research work focused on 57
2 the low energy optimizations using the two technologies 58
3 respectively. However, Combining the two technologies can get 59
4 better energy optimization effect. At the same time, during the 60
5 parallel program execution process, energy overheads caused by 61
6 the re-synchronization, voltage switching, and operations in the 62
7 critical section will inevitably bring extra energy consumption. 63

8 Overall, the contributions made by this study can be 64
9 summarized as follows: 65

10 (1) We establish an energy optimization of parallel program 66
11 (EOPP) model based on a heterogeneous system. During the 67
12 modeling stage, all-round power modeling is conducted on the 68
13 heterogeneous system from a program level perspective (fully 69
14 considering the effect of energy overheads caused by 70
15 re-synchronization, voltage switching, and operations in the 71
16 critical section) to increase the accuracy of the power model. 72

17 (2) Using the EOPP model, we study a strategy that combines 73
18 processor core-shutdown and DVS. Processor core-shutdown is 74
19 used in the serial segment of the program — the other processor 75
20 cores are shut down (i.e. all but the host processor core 76
21 performing the serial program) in order to reduce energy 77
22 consumption. On the other hand, the parallel segments of the 78
23 program are processed using DVS technology. According to the 79
24 loads on various processor cores, the voltage/frequency of each 80
25 processor core is scaled to reach to an ideal voltage/frequency 81
26 balance. By doing so, energy wasted by performing at maximum 82
27 voltage/frequency can be avoided. 83

28 Section 2 presents a formalized description of the EOPP mode 84
29 in a heterogeneous system. Section 3 gives an analysis of the 85
30 EOPP model in a heterogeneous system. On this basis 86
31 considering the influence of energy overheads caused by 87
32 re-synchronization, voltage switching, and operations in the 88
33 critical section, we propose a strategy for energy optimization 89
34 that combines both processor core-shutdown and DV 90
35 technology. Section 4 presents the experimental results, and the 91
36 last section summarizes the conclusions and future work. 92

37 2. Related work 93

38 A great deal of research has been carried out on technologies 95
39 aimed at optimizing the power consumption of software running 96
40 in a multi-core processor. Therein, processor core-shutdown [16- 97
41 17] and dynamic voltage scaling (DVS) [18–19] are the two main 98
42 technologies aimed at optimizing the power used by software. 99

43 From the perspective of compiler optimization technology with 100
44 multi-threading parallel and low power consumption, Zhao et al. 101
45 [20] proposed a low-power optimization model integrating 102
46 fine-grained multi-threading division and dynamic frequency 103
47 adjustment based on two multi-thread system structures (Ch 104
48 Multi-processing and Simultaneous Threading). Their model 105
49 performs well in reducing the power consumption of a processor 106
50 during operation as much as possible on the premise that 107
51 instruction-level and thread-level parallelism is influenced little 108
52 Grochowski et al. [21] discussed the problem of trade-off between 109
53 the throughput capacity of a microprocessor under speed and 110
54 power constraints. These workers applied DVS, asymmetrical 111
55 and sizeable processor cores, and inferential control technology 112
56 dynamically change the energy consumed during instruction

execution under parallelization instruction of software. Their
results indicate that the optimal choice is to comprehensively
utilize DVS with asymmetrical processor cores.

Kadayis et al. [22] proposed a method capable of shutting down
idle processor cores to reduce the energy consumption of nested
loops. However, their energy optimization method does not utilize
DVS technology. In subsequent research [23], they proposed
employing DVS to decrease the voltages of processor cores with
lightly-loaded threads. In this way, the degree of load equilibrium
of the programs can be adjusted to further save power and reduce
energy consumption. However, their research only made a
comparison between DVS and processor core-shutdown and
failed to favorably combine the two low-power technologies.

Li et al. [24–25] have advised that a two-dimensional space
defined for parallel programs in a multi-core structure should be
optimized (one dimension corresponding to change in number of
active processors; the other to conducting DVS on each processor
core). However, a parallel program tends to contain both parallel
and serial segments and there are certain time and energy
overheads associated with conducting DVS and processor
core-shutdown. When a program progresses from a serial
segment to a parallel segment, it is assumed that closing idle
processor cores does not waste extra time and energy. However,
when the operations in a serial segment complete and enter into
the next parallel segment, it is necessary to re-activate the
processor cores in the dormant state and then make the processor
core recover to the normal work state. In this context, there is a
re-synchronization overhead with respect to time and energy.
Moreover, the voltage switching overhead in time and energy
always exists during each DVS process. Thus, the associated
power consumption cannot be ignored when frequently
conducting DVS operations. Additionally, data must inevitably be
shared during parallelization of programs. Therefore, in order to
maintain data consistency, shared data need to be operated on
within a critical section. Thus, the synchronization operation on
the critical section has an effect on power optimization.

This study establishes a model for energy optimization of
parallel programs. By considering the energy overheads caused
by re-synchronization, voltage switching, and operations in
critical sections, the model effectively combines dynamic voltage
scaling and processor core-shutdown technologies, which can
effectively reduce the energy consumption of parallel programs.

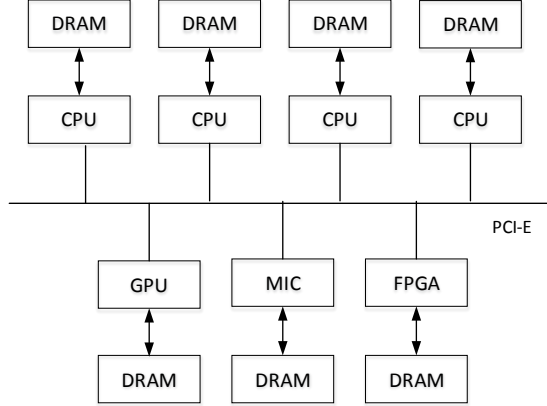
3. Formal description of the optimization problem

3.1 Architecture of a heterogeneous system

Before giving a description of the problem, the architecture of a
heterogeneous system is first abstracted to establish a basis for
the energy model for parallel programs. Fig. 1 shows the typical
architecture of a heterogeneous parallel system. Such a system
contains multiple computing resources: central processing units
(CPUs), graphics processing units (GPUs), medium interface
connectors (MICs), and field programmable gate arrays (FPGAs).

In general, those parts used to accelerate processing (the
APUs) are only charged with performing specific computing tasks
and are not equipped with accomplished task management and

1 scheduling mechanisms. Therefore, most of the APUs need to be
 2 executed under the control of a general-purpose microprocessor
 3 (host processor). The host processor is generally connected to the
 4 APU through an external bus and they have their own (off-chip)
 5 dynamic random-access memory (DRAM). In this way, the two
 6 processors can realize data communication in the form of direct
 7 memory access (DMA).



8
 9 Fig.1. The typical architecture of a heterogeneous parallel system.

with competitive critical sections within unit time is in accordance with the binomial distribution.

It can be considered that the assumption 1 and 2 are true and reasonable. Processor core shut-down, as well as the data transmission between the master processor and the accelerated processor will impose additional time and energy consumption. However, it is important to point out that this article focuses on the energy consumption of re-synchronization, voltage switching, and operation in critical sections, the time and energy consumption overhead caused by processor core shutdown are not the main component of the execution time and total system energy consumption, so it can be considered that the assumptions 3, 4, 5, 6 are reasonable. In addition, the energy consumption analysis of operations in the critical sections is based on the literature [26], therefore the assumption 7 is also reasonable.

3.3 Parameters used in the EOPP model for the heterogeneous system

The following parameters are defined for the problem in the heterogeneous system:

- N_C : the number of serial segments;
- N_G : the number of parallel segments;
- P_C : the number of host processor cores;
- P_G : the number of APU cores;
- tc_i : the execution time at the i th serial segment;
- IC_i^C : the instruction cycles of the i th serial segment;
- t_s : time of each synchronization overhead;
- $f_{C_{max}}$: the maximum clock frequency of a host processor core in an active state;
- $f_{G_{max}}$: the maximum clock frequency of an APU core in an active state;
- $V_{C_{max}}$: the maximum voltage of a host processor core;
- $V_{G_{max}}$: the maximum voltage of an APU core;
- $P_{dyn}(f)$: the dynamic power consumption when the frequency of the processor is f ;
- P_{stat} : the static power consumption of a processor core in the shutdown state;
- E_{syn} : the energy overhead of each re-synchronization step;
- $tG_{j,m}$: the execution time of the m th processor at the j th serial segment;
- tG_j : the longest execution time of P processors at the j th parallel segment;
- $IC_{j,m}^G$: (profiling data) the instruction cycles of the m th processor at the j th parallel segment;
- E_C : the energy consumption of a serial segment;
- E_G : the energy consumption of parallel segment;
- E_{total} : the total energy consumption before optimization;
- E_{opt} : the total energy consumption after optimization;
- $t_{swi}(V_{j-1,m}, V_{j,m})$: the time overhead due to switching voltage from $V_{j-1,m}$ to $V_{j,m}$;
- $f_{j,m}$: the optimal clock frequency of the m th processor at the j th parallel segment;
- $V_{j,m}$: the optimal voltage of the m th processor at the j th parallel segment;
- $E_{swi}(V_{j-1,m}, V_{j,m})$: the energy overhead when switching the voltage from $V_{j-1,m}$ to $V_{j,m}$.

3.2 Problem description and conditional hypotheses

The problem can be described as follows. Consider a heterogeneous system consisting of a host processor and an APU and a parallel program containing N_C serial and N_G parallel segments. It is hypothesized that the architecture of the heterogeneous system contains P_C host processor cores and P_G APU cores, and that the voltage and frequency allocated to the m th processor at the j th stage (serial or parallel) are expressed in the form $(V_{j,m}, f_{j,m})$. We are required to find the optimal voltage/frequency of each processor core that reduces the total energy consumed by the whole parallel program to the largest extent without incurring performance loss.

In order to accurately define the energy optimization problem for parallel programs executing on a heterogeneous system, some further hypotheses are made about the program, system architecture, and circuit realization:

- (1) The logistical behavior of the program does not change when the frequency changes.
- (2) The frequency of a processor is continuous and adjustable.
- (3) The serial and parallel parts of the parallel program are separately fulfilled by the host processor and APU. The time and energy overheads resulting from data transmission between the host processor and APU can be ignored.
- (4) The program flows from parallel to serial segments and there are no extra time and energy overheads caused by closing idle processors.
- (5) The program operates from serial to parallel segments taking into account the re-synchronization overheads in time and energy.
- (6) The time and energy overheads caused by converting the frequencies of the processors are considered.
- (7) The probability of a critical section occurring within a thread conforms to a uniform distribution while that in different thread levels is completely independent. The number of threads

4. Analysis of the EOPP model for a heterogeneous system

4.1 The total energy consumed by parallel programs that are not optimized

Complementary metal-oxide semiconductor (CMOS) transistors are the basic devices that make up a computer. The power consumption of a CMOS is mainly comprised of dynamic and static components and the dynamic component P_{dyn} arises due to the changes in state of the CMOS as it performs work. Thus, we can write:

$$P_{dyn} = \alpha CV^2 f, \quad (1)$$

where α , and refer to is the switching activity factor (in the range of 0–1), C the switching capacitance, V the supply voltage, and f the clock frequency. The static component arises due to leakage currents in the idle state and can be estimated according to a certain proportion.

According to basic physics, the energy consumed (E) can be found by integrating power consumption level with respect to time. A simple calculation can therefore be made by multiplying the average power consumption (P) by time (t):

$$E = Pt. \quad (2)$$

The parallel program is assumed to consist of N_C serial segments and N_G parallel segments. Therefore, the total energy consumed by running the parallel program can be expressed as the sum of the energies consumed by the serial and parallel segments:

$$E_{total} = E_C + E_G. \quad (3)$$

If DVS is not conducted, all the processor cores (including the host processor and APU cores) operate at maximum voltage/frequency. Therefore, according to the hypotheses in Section 2.2 (the serial and parallel parts of the parallel program are separately completed by the host processor and APU, and the time and energy overheads caused by data transmission between the host processor and APU are ignored), the total energy consumed by all the processor cores due to the N_C serial and N_G parallel segments is given by:

$$E_{total} = \sum_{i=1}^{N_C} (P_C \cdot tc_i \cdot \alpha CV_{C_{max}}^2 f_{C_{max}}) + \sum_{j=1}^{N_G} (P_G \cdot tG_j \cdot \alpha CV_{G_{max}}^2 f_{G_{max}}). \quad (4)$$

4.2 Combining DVS and processor core shutdown

The parallel parts of the program are to be subjected to energy optimization. On the one hand, $P_C - 1$ host processors (which are not utilized in the execution of the serial segment) are shut down. On the other hand, it is hypothesized that the optimal voltage/frequency of the m th APU core during the execution of the j th parallel segment is $(V_{j,m}, f_{j,m})$. Therefore, failing to consider the time and energy overheads caused by state switching, the ideal total energy consumed by the parallel program can be expressed as:

$$E_{opt} = \sum_{i=1}^{N_C} (tc_i \cdot \alpha CV_{C_{max}}^2 f_{C_{max}} + (P_C - 1) \cdot tc_i \cdot P_{stat}) + \sum_{j=1}^{N_G} \sum_{m=1}^{P_G} (tG_j \cdot \alpha CV_{j,m}^2 f_{j,m}). \quad (5)$$

It is further hypothesized that the ratio of the power consumed by a processor core in the inactive state to that in the active state is ε . When $P_C - 1$ processor cores are shut down during the execution of the serial segment, the power consumed by the processor cores in the shutdown state is given by:

$$P_{stat} = \varepsilon \cdot \alpha CV_{C_{max}}^2 f_{C_{max}}. \quad (6)$$

For the sake of convenience, the frequency and voltage of the processor cores can be standardized. For a parallel segment of the program, it is hypothesized that the frequency of the APU core due to DVS is $f_{j,m}$ and the maximum allowable frequency is f_{max} . The standardized frequency is now given by

$$f_{j,m}' = \frac{f_{j,m}}{f_{max}} \quad (m = 1, \dots, P_G, j = 1, \dots, N_G). \quad (7)$$

In a similar way, it can be hypothesized that the voltage of a processor core due to DVS is $V_{j,m}$ and the maximum allowable voltage is V_{max} so that a standardized voltage can be defined as

$$V_{j,m}' = \frac{V_{j,m}}{V_{max}} \quad (m = 1, \dots, P_G, j = 1, \dots, N_G). \quad (8)$$

According to the relationship between supply voltage and clock frequency, it can be further seen that:

$$V'' = k_1 + k_2 \cdot f'. \quad (9)$$

Eq. (7) conforms to industrial standards, according to the literature [27]. By analyzing the technical indices available, it can be speculated that k_1 and k_2 should be offset to make them 0.3 and 0.7, respectively.

For the parallel segments of the program, ignoring the synchronous waiting and intercommunication between parallel parts, the voltage/frequency of each processor core can be scaled to different levels according to the different execution times. The longest execution time tG_j of the processor core can be standardized to 1 so that those of the other $P_G - 1$ processor cores all satisfy $tG_{j,m} \leq tG_j$. According to the various different execution times, $tG_{j,m}$, the frequency of the m th coprocessor core is therefore scaled by $f_{j,m}'$. By doing so, the following formula can be obtained, where:

$$f_{j,m}' = \frac{tG_{j,m}}{tG_j}. \quad (10)$$

The two parameters $tG_{j,m}$ and tG_j in Eq. (10) can be measured using a performance test tool (data profiling) and, based on Eq.(9), the following formula can be acquired:

$$V_{j,m}' = k_1 + k_2 \cdot f_{j,m}'. \quad (11)$$


4.3 Re-synchronization and voltage switching overheads

Each DVS operation can cause additional time and energy overheads due to voltage switching. Therefore, the total energy consumed by the parallel program needs to have the overheads caused by re-synchronization and voltage switching to be separately taken into account.

(1) Energy overhead caused by re-synchronization

During the execution of the parallel program, it is necessary to re-activate the processor cores in a dormant state to put them back into the working state. This occurs when the execution process in a serial segment has completed and we return to the next parallel segment. In this context, the whole system exhibits re-synchronization overheads in terms of time and energy. Time overheads induced by re-synchronization can be avoided, however, by using a pre-activation strategy [28]. Pre-activation

1 refers to when a processor core in a shutdown state is activated in 4 any time overhead. The ideal pre-activation strategy is shown in
 2 advance before reutilization is required so that the processor core 5 Fig.2.
 3 is in an active state when it is needed. This can obviously avoid




6
7
8 Fig.2. The ideal pre-activation strategy showing: (a) the initial state, (b) re-synchronization overhead caused by processor core-shutdown, and (c)
9 pre-activation to eliminate the synchronization overhead

10 The pre-activation strategy can remove the time delay but fail
11 to avoid an energy overhead caused by re-synchronization. On
12 purpose of this study is to take re-synchronization overheads into
13 account when calculating the total energy overhead of the
14 parallel program. To that end, the total energy overhead is
15 written in the form:
16 $E_{opt}^1 = \sum_{i=1}^{N_c} (tc_i \cdot \alpha CV_{c_{max}}^2 f_{c_{max}} + (P_c - 1)((tc_i - t_s) \cdot P_{stat} +$
17 $E_{syn}) + \sum_{j=1}^{N_g} \sum_{m=1}^{P_g} (tG_j \cdot \alpha CV_{j,m}^2 f_{j,m}')$ (12)
18 (2) Energy overhead caused by voltage switching
19 During the execution of the parallel program, each DVS step
20 can incur a voltage switching overhead in time and energy. In
21 this case, the time overhead can once again be avoided by, this
22 time, using a method of pre-switching. In the CPU-GPU
23 heterogeneous system, the parallel parts of the parallel program
24 are completed using APUs and, therefore, DVS is generally

25 aimed at processes going on in the APU cores. According to the
26 various loads on each APU core, the voltages/frequencies are
27 rescaled to ensure the different parallel parts are completed in
28 the same amount of time, thus eliminating any null cycles.

29 Fig.3 shows the strategy used to avoid time overheads using
30 DVS. Fig.3(a) shows the load conditions when each processor core
31 operates at its maximum voltage/frequency and DVS running
32 consumes certain time. If the optimal voltage/frequency is
33 calculated based only on the different loads, a time delay is
34 accumulated once the paralleled intervals are increased which
35 causes performance loss, as shown in Fig.3(b). In order to avoid a
36 time delay, the time overhead in this part is subtracted when
37 calculating the voltage/frequency of each processor core, as shown
38 in Fig.3(c).



39
40
41 Fig.3. Diagram to illustrate DVS overheads: (a) without conducting DVS, (b) carrying out DVS after considering the voltage switching overhead, and (c)
42 carrying out DVS to avoid performance loss.

43 As a result of rescaling, the frequency of the APU core can be
44 expressed in the form:

$$45 \quad f_{j,m}'' = \frac{tG_{j,m}}{tG_j - t_{swi}(V_{j-1,m} - V_{j,m})} \cdot f_{G_{max}} \quad (13) \quad 56$$

46 It has been shown that the time and energy overheads caused
47 by voltage switching can be expressed via the relationships [29]:

$$48 \quad t_{swi}(V_{j-1,m}, V_{j,m}) = \frac{2C}{I_{max}} |V_{j-1,m} - V_{j,m}|, \quad j = 2, \dots, N_g \quad (14)$$

$$49 \quad E_{swi}(V_{j-1,m}, V_{j,m}) = (1 - u) \cdot C \cdot |V_{j-1,m}^2 - V_{j,m}^2|, \quad (15) \quad 60$$

50 where C , and refer to is the capacitance, u the energy
51 effectiveness factor of the voltage regulator, and I_{max} the
52 maximum allowable current. Substituting Eq. (14) into Eq. (13),
53 allows us to acquire the optimal frequencies to which each
54 processor core needs to be scaled to using Eqs.(7)-(9): $f_{j,m}' =$ 65

$$\frac{tG_{j,m} \cdot f_{G_{max}}}{tG_j - \frac{2C \cdot k_2 \cdot V_{G_{max}}}{I_{max} \cdot f_{G_{max}}} |f_{j-1,m} - f_{j,m}'|}$$

57 Considering the re-synchronization overhead part of the total
58 energy overhead of the parallel program, the total energy
overhead can be calculated using:

$$61 \quad E_{opt}^2 = \sum_{i=1}^{N_c} (tc_i \cdot \alpha CV_{c_{max}}^2 f_{c_{max}} + (P_c - 1)((tc_i - t_s) \cdot P_{stat} + E_{syn})$$

$$62 \quad + \sum_{j=1}^{N_g} \sum_{m=1}^{P_g} (tG_j \cdot \alpha CV_{j,m}^2 f_{j,m}'') + \sum_{j=1}^{N_g-1} \sum_{m=1}^{P_g} (1 - u) \cdot C \cdot |V_{j,m}^2 - V_{j+1,m}^2|. \quad (16)$$

(3) Energy overheads due to operations in critical sections

63 In the process of parallelizing programs, the need to share data
64 between parallel segments becomes inevitable. In order to
65 guarantee the consistency of the data, operations involving
shared data need to be performed in a critical section. The effect

of using critical sections on the power consumption of the program must therefore be accounted for.

It is hypothesized that the parallel parts need to undergo parallel processing using j th class processor in which the task load, operating proportion, and probability of conflict in the critical section are S_j , σ , and c , respectively. The number of j th class processor r_j is recorded as N_j and $\sum_{r_j \in R} N_j$.

The work in this section is based on certain hypotheses appearing in the literature [30]. In particular, it is hypothesized that the probability that a critical section occurs within a thread conforms to a uniform distribution and that the probabilities for different threads are completely independent. The number of threads with competitive critical sections per unit time follows binomial distribution. According to conflict models established in the literature [24], the time a processor takes to execute a critical section is composed of the time taken for execution of the local critical section and the waiting time in the critical section. The waiting time is proportional to the execution time, execution probability, and conflict probabilities of the critical sections of the other processors. Therefore, the execution time in the critical section of the j th class processor can be expressed as:

$$t_{j,m} = \frac{\sigma S_j}{f_{j,m} v_j N_j} + \left((N_j - 1) \frac{\sigma S_j}{f_{j,m} v_j N_j} + \sum_{r_k \in R - \{r_j\}} N_k \frac{\sigma S_k}{f_{k,m} v_k N_k} \right) \sigma c. \quad (17)$$

The second item in Eq.(17) represents the average waiting time of the j th class processor. Considering that the contribution made by $\sigma^2 S_j c / f_{j,m} v_j N_j$ to $t_{j,m}$ is small, the execution time of the critical sections of the j th class processor can be approximated as:

$$t_{j,m} = \frac{\sigma S_j}{f_{j,m} v_j N_j} + \sum_{r_k \in R} \frac{\sigma^2 S_k c}{f_{k,m} v_k}. \quad (18)$$

When we include the overheads in the critical section, the execution time of the j th class processor is therefore given by the expression:

$$t_{G_{j,m}} = \frac{(1-\sigma)S_j}{f_{j,m} v_j N_j} + \frac{\sigma S_j}{f_{j,m} v_j N_j} + \sum_{r_k \in R} \frac{\sigma^2 S_k c}{f_{k,m} v_k} = \frac{S_j}{f_{j,m} v_j N_j} + \sum_{r_k \in R} \frac{S_k}{f_{k,m} v_k} \sigma^2 c. \quad (19)$$

Now, $\forall r_j, r_k \in R$, $t_{G_{j,m}} = t_{G_{k,m}}$. Therefore, $S_j/S_k =$; Owing to $f_{j,m} v_j N_j / f_{k,m} v_k N_k$. As $\sum_{r_k \in R} S_j = S$, it can be seen that $S_j =$

Thus $f_{j,m} v_j N_j / \sum_{r_k \in R} f_{k,m} v_k N_k$, so that,

$$t_{G_{j,m}} = \frac{S_j}{f_{j,m} v_j N_j} + \sum_{r_k \in R} \frac{S_k}{f_{k,m} v_k} \sigma^2 c = \frac{1 + \sigma^2 cn}{\sum_{r_k \in R} f_{k,m} v_k N_k} S. \quad (20)$$

Therefore, according to Eq. (10), it can be seen that:

$$f_{j,m}''' = f_{j,m} (1 + \sigma^2 cn). \quad (21)$$

Based on the foregoing analysis, the following conclusion can be drawn. For a given program, the operations in the critical section cause a fractional increase in the operating frequency of a

processor of $\sigma^2 cn$ (compared to the situation in which critical sections are not used). This further increases the total energy overhead of the program.

5. Experiments and analysis

5.1 Experimental platform

A heterogeneous system consisting of an Intel Core i7 920 quad-core CPU and AMD 4870 GPU was used to form an experimental platform for testing purposes. In the system, the CPU and GPU have their own separate independent memory spaces and are connected via the PCI-E bus to realize data communication.

The theoretical analysis of the energy consumption model proposed in this study suffers from certain limitations. More precisely, it fails to accurately describe some of the uncertain behavior of the processors during the execution process and to simulate the energy consumed in each part. In practical CPU-GPU heterogeneous systems, the GPU does not provide perfect support for dynamic scaling of the voltage/frequency (i.e. a few frequency tuning ranges are used). This is disadvantageous when it comes to conducting theoretical research and verifying the behavior of the EOPP model under the combined effect of DVS and processor core-shutdown. Therefore, we employ a GPU power simulator, GPGPUSim, to facilitate experimental verification.

GPGPUSim, however, fails to operate in a GPU environment and therefore we made modifications based on GPGPUSim for this study. A simple simulator was established in the application layer (CPU end of CUDA program) to simulate the simultaneous execution of multi-GPUs by practice driving. In addition, GPGPUSim configures the CPU by reading configuration files during operation. Therefore, the multi-GPU environment of the heterogeneous system can be simulated by dynamically modifying the configuration files of the program.

The GPU power simulator used in this study was developed using the Wattch power model (based on GPGPUSim). To favorably realize scaleable dynamic voltages, the following modifications were made to the simulator:

(1) The frequency of each processor core is independently scaled by changing the latency of all the functional parts in the processor cores to scale the clock frequency of each core.

(2) The voltage is updated according to the scaled clock frequency.

(3) The power consumption of each functional part in the Wattch model is updated based on the new clock frequency and voltage.

(4) When the voltage state switches, the voltage switching overhead in time and energy are calculated according to Eqs. (14) and (15).

Table 1 Parameters of the processors used in the test platform

Intel Core i7 920 CPU	AMD 4870 GPU-H/GPU-L
-----------------------	----------------------

Processor frequency (GHz)	2.67, 2.4, 2.0, 1.6	0.75, 0.65, 0.55
Memory frequency (GHz)	1.33 (DDR3)	0.9/0.7/0.5 (GDDR5)
Cache	L1: I 32 KB, D 32 KB; L2: 256 KB; L3: 8 MB	—
Memory (GB)	8	1

1

2 5.2 Test cases

3 A great deal of attention has been paid in the scientific
4 computing field to multiplying a vector by a sparse matrix using
5 parallel computing. For this study, we chose four sparse matrices
6 from the Harwell11 sparse matrix set [31-32] and used them as
7 test cases to verify the results from the EOPP model. Table 2
8 displays details of the sparse matrices chosen.

Table 2. Details of the test cases used.

Application name	Scale	Description
Matrix 1	1030	21*21*5 irregular grid
Matrix 2	886	21*21*5 irregular grid
Matrix 3	1080	35*11*13 grid
Matrix 4	5005	16*23*3 grid

9 5.3 Compiler implementation

10 To unify the programming models of the CPU and GPU and
11 simplify the migration of existed applications to CPU-GPU
12 heterogeneous systems, we have extended four GPU oriented
13 instruction commands based on the OpenMP language [33]. The
14 commands are used to instruct the compiler to convert the
15 OpenMP code executed on the original CPU to Brook+ code
16 executed on the GPU. At the same time, a source-to-source
17 compiler called MPtoStream is designed and implemented based
18 on the GCC compiler. It can complete the aforementioned code
19 conversion process. Fig.4 shows the MPtoStream compiler
20 framework. Firstly, the expanded OpenMP program is analysed
21 by the Parser module, and the intermediate syntax tree is
22 generated. Then, the OpenMP program is converted through a
23 series of syntax trees to generate the intermediate syntax tree
24 structure for the Brook+ language. Finally, the reverse output
25 process is called to generate the OpenMP codes which are
26 executed on CPU processor and Brook+ codes which are executed
27 (separately) on a GPU processor. Based on this MPtoStream
28 framework, we expand the parallel loop scheduling module of
29 sparse matrix vector multiplication (SpMV) for CPU-GPU
30 heterogeneous system, as shown in the grey boxes in Fig. 4.

31

32 5.4 Code example

33 Fig. 5 shows that an example of parallel loop scheduling code.
34 This program implements a sparse matrix vector multiplication
35 algorithm (SPMV). The outermost layer loop index variable i
36 represents the line number of the matrix, and the inner layer loop
37 is calculated only for non-zero elements. $N(i)$ indicates the number
38 of non-zero elements in line i . It can be seen that the outermost
39 layer loop is composed of TotalLine iterations. In this paper, the
40 whole procedure is divided into one serial segment and two
41 consecutive parallel program segments. The serial segment was
42 mainly used to execute initialisation of the program (processed by
43 the CPU). The parallel program segment was divided the sparse
44 matrix into even two parts according to the row number, which
45 were separately executed using two segments (processed by the
GPU).

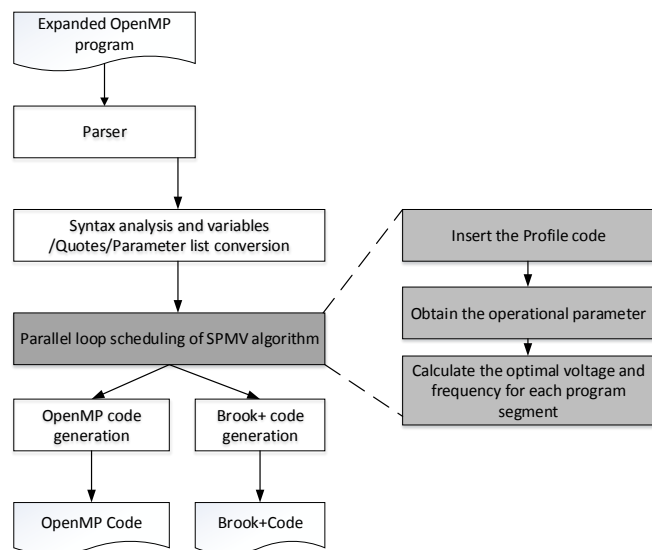


Fig. 4 MPtoStream compilation framework

The first loop iteration algebra on GPU is N_{gpu} . Therefore, the iterative subset $(1, N_{gpu})$, $(N_{gpu}+1, TotalLine)$ is allocated to the GPU as the two continuous parallel program segments, and the rest of the programs are allocated to the GPU. As shown in Fig. 5, in the left upper dashed box, the expanded compile command identity maps the parallel segment to GPU for execution. The right-hand dashed box is the CPU code that has been converted. Firstly, the GPU data flow space is declared, and the data loading process is completed. Then the iteration space which mapped on the GPU is formulated (through the `spmv_kernel` program implementation). Finally, the computing process is executed on

59

1	75620	531381/ 521680	531743/ 531916	540119/ 540320	531736/ 531993	531295/ 531531	539931 /539572
2	90160	393208/ 393354	393557/ 393492	400844/ 400983	393620/ 393688	393480/ 393688	400383/ 401101
3	72590	580267/ 583660	588839/ 592460	580181/ 583581	588839/ 592078	580181/ 583140	588839/ 591924
4	72600	12502099/ 12546222	12543121/ 12547473	12544017/ 12508724	12544815/ 12548586	12547937/ 12547808	12511428/ 12549846

Table 4. The optimal scaled voltages and frequencies in each parallel segment.

Test case	Parallel segment	Optimal voltage/frequency (V/MHz).					
		P1	P2	P3	P4	P5	P6
1	1	2.0268/ 753.82	2.0277/ 754.49	2.0500/ 770	2.0277/ 754.48	2.0495/ 769.65	2.0270/ 753.66
	2	2.02271/ 754.02	2.0277/ 754.45	2.0500/ 770	2.0279/ 754.59	2.0267/ 753.74	2.0481/ 768.63
2	1	2.0227/ 750.95	2.0239/ 751.82	2.0500/ 770	2.0240/ 751.98	2.0236/ 751.63	2.0481/ 768.85
	2	2.0223/ 750.70	2.0229/ 751.08	2.0496/ 769.73	2.0235/ 751.35	2.0234/ 751.53	2.0501/ 770
3	1	2.0290/ 755.33	2.0498/ 769.89	2.0287/ 755.19	2.0498/ 769.89	2.0287/ 755.19	2.0498/ 769.88
	2	2.0288/ 755.24	2.0500/ 770	2.0286/ 755.11	2.0492/ 769.40	2.0276/ 754.40	2.0489/ 769.10
4	1	2.0466/ 769.26	2.0493/ 769.45	2.0492/ 769.78	2.0495/ 769.34	2.0499/ 769.90	2.0457/ 767.23
	2	2.0496/ 769.53	2.0567/ 769.65	2.0451/ 766.55	2.0495/ 769.82	2.0496/ 769.72	2.0497/ 769.86

1

2 The voltage/frequency scaling results were then obtained
3 according to the method given in the analysis for setting the
4 optimal voltage (Table 4). These values were then used together
5 with the energy calculation formulae, Eqs. (4) and (16), to
6 calculate the energy used before and after the optimization
7 process. Thus, the resulting energy saving can be obtained taking
8 into account the energy overheads caused by re-synchronization
9 voltage switching, and use of critical sections. 30

10 Fig. 6 shows the energies consumed in the serial and parallel
11 segments, as well as the total energy consumed, in the four test
12 cases (before and after optimization). The energy savings are also
13 shown (as percentages). It can be seen from the results that the
14 total energy saving was as large as 10.5%. Also, the energy saved
15 in the serial segments was significantly larger than that saved in
16 the parallel segments. The main reason for this is that the
17 processor cores that were not being utilized were shut down
18 when the programs entered the serial segment which greatly
19 reduces energy consumption. Moreover, the results obtained can
20 be related to the degree of parallelization and parallel
21 programming loads in the test cases. 42

22 (2) Analysis of the EOPP real results 43

The theoretical analysis of the energy consumption model
proposed in this study suffers from certain limitations. More
precisely, it fails to describe some of the uncertain behaviours of
the processors during the execution process and to simulate the
energy consumed in each part. In this section, the energy
optimization methods of parallel programs are implemented
under the real heterogeneous system platform, so, we can obtain
more realistic energy optimization results.

The test platform in this study is a system composed of an
Intel Core I7 920 Quad-Core CPU and two AMD 4870 GPUs. To
examine the efficiency of the algorithms proposed on this system,
the frequency of the storage of one GPU kernel is adjusted from
900 MHz to 700 MHz so as to obtain two different GPU kernels
with different performances. Thereinto, the kernel with the
higher performance is recorded as GPU-H, while that with the
lower performance is designated GPU-L. The specific parameters
for this test platform are listed in Table 1.

The voltage/frequency information supported by the
processor is obtained. The current mainstream CPU processor
supports DVS technology, such as Intel's SpeedStep [34], AMD's
PowerPlay [35] and other power management technologies. In

1 addition, the power management modules for the specific
2 processor are integrated in the operating system (OS). The OS in
3 this study is OpenSUSE v10.3. The operating frequency
4 information supported by CPU can be obtained through the
5 ACPI [36] interface provided by the system, and the operating
6 frequency can be dynamically adjusted. The controlled power
7 method of the GPU is also gradually improved. AMD provides an
8 ADL library (AMD Display Library)[37], which can dynamically
9 access and modify the operating voltage and frequency of GPU
10 through the ADL_Overdrive5_ODPerformanceLevels_Get and
11 ADL_Overdrive5_ODPerfromanceLevels_Set interfaces. 70

12 The power consumption of the system is measured by the
13 external HIOKI 3344 power tester, and the power consumption is
14 read through the RS-232 serial port mechanism. The
15 measurement error of the instrument is within the ± 1 digit
16 range of the measured value. 74

17 It is difficult to measure the energy consumption of the serial
18 segment and parallel segment separately in real heterogeneous
19 systems. Therefore, this article only compares the total energy
20 consumption of the system before and after the optimization; and
21 presents the statistics pertaining to the total energy consumption
22 savings, as shown in Fig.7. During the execution of parallel
23 programs, the voltage and frequency of the CPU and GPU are set
24 to the maximum value (2.0 V/2.67 GHz; 2.0 V/750 MHz) in the
25 initial case. After entering the parallel section, the voltage and
26 frequency of the GPU are adjusted respectively (750 MHz, 650
27 MHz, 550 MHz). The experiment results show that using
28 processor core shut-down and DVS technologies can reduce the
29 energy consumption of parallel programs effectively, and the
30 maximum energy saving can reach about 7.2%. The error rate in
31 the energy saving and simulation energy saving is within $\pm 5\%$. 88

32 (3) The effect of critical sections on energy consumption 89

33 In order to suitably test the influence of critical sections on the
34 frequency and consumption of energy using parallel programs
35 we carried out an investigation and analysis using typical
36 parallel programming cases (Fig. 8). 93

37 As can be seen from the figure, the parallel segment of the
38 parallel program consists of a completely-parallelized parallel
39 segment and a critical section. The operating proportion of the
40 critical section, σ of critical section is, was scaled without
41 changing the computing amount in the parallel segment and the
42 results used to estimate the effect of the critical section operation
43 on the energy optimization of the parallel segments. Once again
44 the Intel Core i7 920 quad-core CPU was used as the
45 experimental platform and the program compiled using gfortran
46 (v4.2.1). 103

47 (a) Power overhead when the critical section is in a waiting
48 state 105

49 In order to avoid overheads caused by other operations, the
50 parallel segments are formed under the operation of a critical
51 section ($\sigma = 1$) and therefore only one processor executes the
52 effective calculation at any time. By modifying the environment
53 variable OMP_NUM_THREADS, the number of threads used in
54 the concurrent execution can be adjusted. 111

55 Fig. 9 shows the change in the dynamic power overhead of the
56 processor as the thread number changes. As can be seen from the
57 figure, the total dynamic power consumed by the processor varies
58 insignificantly. This implies that the processor does not cause
59 significant power overhead in the waiting state. 116

60 (b) Relationship between σ and optimal frequency 117

By adjusting the operating proportion of the critical section, σ ,
used in the test cases, we can detect the subsequent changes that
occur in the frequency of the processor. The execution time when
the processor operates in a non-critical state ($\sigma = 0$) using the
lowest operating frequency is taken as the time constraint. In this
work, we make a comparison of three different frequency values:
(i) the ‘theoretical’ value (i.e. the frequency of the processor
obtained using the analysis model), (ii) the ‘physical’ value (i.e.
the actual frequency used by the processor in practice —
corresponding to the lowest discrete operating frequency of the
processor above the theoretical frequency), and (iii) the ‘optimal’
value (i.e. the ideal frequency according to the relationship
between frequency and execution time under a discrete frequency
value).

As shown in Fig. 10, the optimal frequency of the processor
gradually rises with increasing σ value. When σ reaches 50%, the
theoretical frequency exceeds the highest frequency allowed by
the processor and therefore the physical value cannot be
displayed.

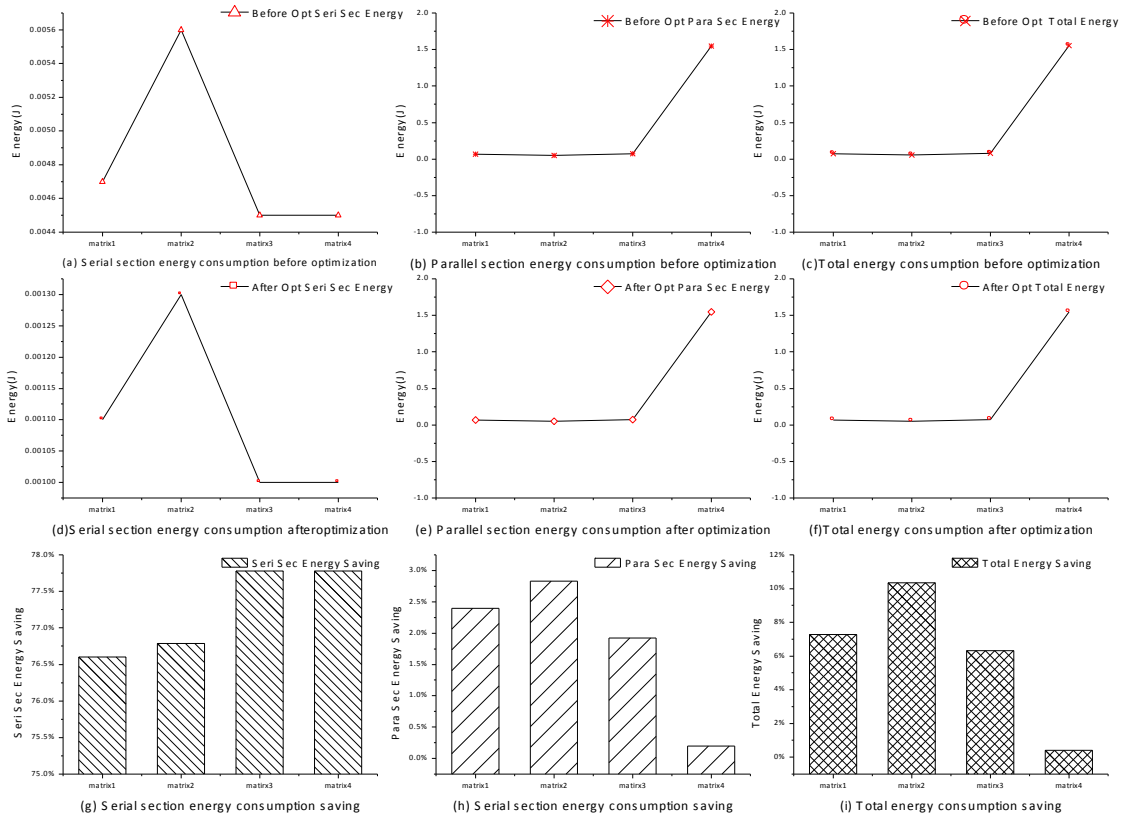
(c) Reducing energy consumption by reducing conflict
probability

The above analysis has shown that, compared to not using
critical sections, the operating frequency of the processor is
 $(1 + \sigma^2 cn)$ times that of the original frequency (under the
constraint that the execution time is the same). In this context,
the effective computing time (eliminating the waiting time of the
critical section) is $(1 + \sigma^2 cn)^{-1}$ that of the original time.
Therefore, it can be seen that decreasing the conflict probability
of the critical section (c) can effectively reduce the energy
overhead, assuming σ is constant. For this study, we partitioned
the critical section of the test program into two critical sections
that can be executed concurrently using different processors. In
this case, the conflict probability of the critical section can fall to
about 50% which further saves energy by allowing the operating
frequency of the processor to be decreased. Fig. 11 displays the
energy optimization results obtained (all results are normalized
values compared to the initial amount of energy consumed). The
figure shows that reducing the conflict probability provides space
for optimization to be made. This allows the frequency of the
processor to be lowered which subsequently reduces the energy
overhead of the system.

(3) Scalability of the solution

A new energy optimization of parallel programs (EOPP) model
is proposed in this paper and its use can be extended to
large-scale heterogeneous systems with multiple processing
units, eg, data center [38-40]. At present, we have applied our
theoretical results to the Tianhe-2A supercomputer. Based on the
system load, energy consumption distribution, and hardware
features, a power optimization management system is designed.
The system architecture is composed of three layers (Fig.12),
which are: the energy consumption decision layer, the perception
control layer, and the hardware platform layer, respectively. The
energy consumption decision layer is composed of six parts: a
monitoring module, a job management module, a resource
management module, a low-power compiling module, a
peripheral device control module, and a power management
interface module.

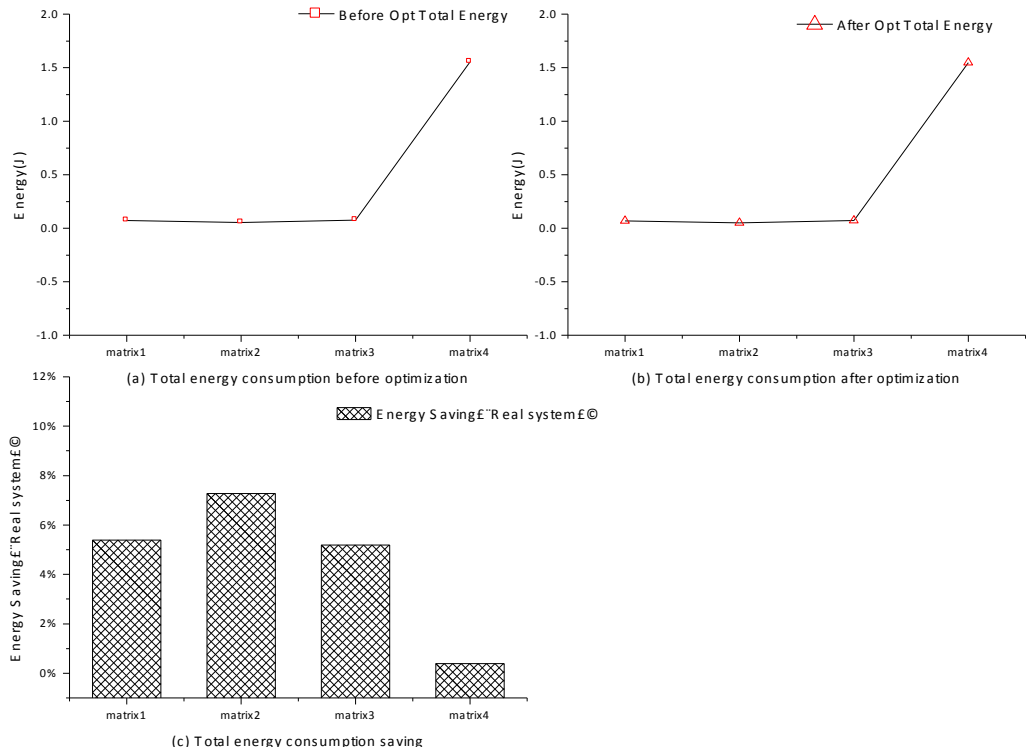
1



2

3

Fig. 6. The energy savings gained using the optimization process.



4

5

Fig. 7. Actual energy consumption optimization


```
Initial input array A
```

```
!$omp parallel
```

```
do i = 1, n
```

```
    //full parallel code
```

```
    update A
```

```
    //critical section
```

```
    !$omp critical
```

```
        update A
```

```
    !$omp end critical
```

```
enddo
```

```
!$omp end parallel
```

Fig. 8. Code illustrating the synchronization tests carried out on the critical sections.




Fig. 9. The change in the power overhead with thread number.

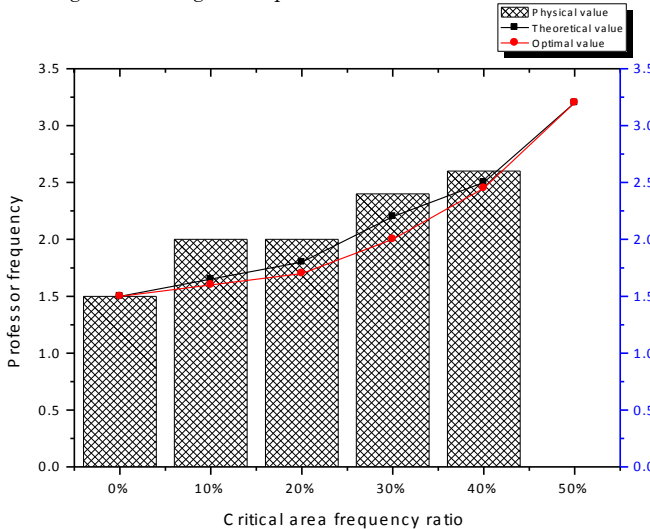


Fig. 10. Relationship between the operating frequency of the processor and the operating proportion of the critical section.

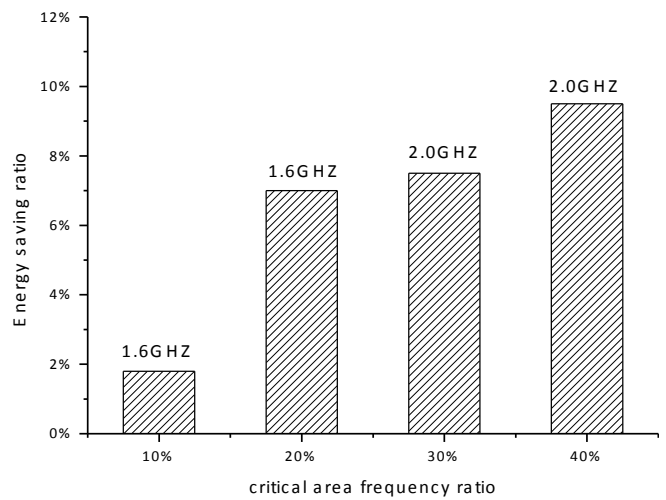


Fig. 11. Decreasing the energy consumption of the system by reducing the conflict probability of the critical section.

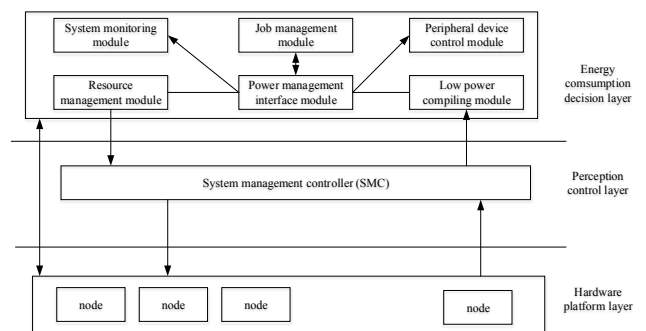


Fig. 12 Power management system architecture

The main function is to collect all data related to power consumption. The power optimization management system is integrated with monitoring, strategy making, power adjustment, and equipment control. It tries to reduce the whole system power while guaranteeing the performance, which is then combined with the methods of resource management, operation

1 management, low-power consumption compilation, and dynamic
2 voltage adjustment.

3 6. Conclusions and future work

4 This study presents an energy optimization model (EOPP) for
5 a heterogeneous system. By jointly utilizing processor
6 core-shutdown and DVS, the model can instruct parallel
7 programs to conduct energy optimization taking into account the
8 energy overheads caused by re-synchronization, voltage
9 switching, and the operation of critical sections.

10 Our experimental results show that the EOPP model is able to
11 reduce the energy consumed by parallel programs, realizing a
12 total energy saving of up to ~10.5%. Based on an experimental
13 analysis of the effect of the operations in critical sections on
14 energy consumption, it is speculated that increasing the
15 proportion of the critical section causes a corresponding
16 enhancement to be made to the optimal frequency of a processor.
17 However, lowering the conflict probability of the critical section
18 can provide optimization space which lowers the frequency of the
19 processor and finally results in reduced energy overheads.

20 Due to the parallel program exists the load imbalance and data
21 dependency during the actual execution process, the modelling
22 of energy optimization for parallel program becomes very
23 complicated. In the future work, the characteristic of the program
24 parallelization will be further investigated, and the
25 corresponding energy analysis and optimization will be carried
26 out.

28 Acknowledgments

29 This work was sponsored by National Natural Science
30 Foundation of China (grant number 61672168, 61300029,
31 61502110, 61672172, 61772143), Guangzhou Major Science and
32 Technology Projects (201604010096)

34 References and links

35 1. Geist A. Paving the Roadmap to Exascale
36 SciDAC Review. 2010,16:52-59. 97
37 2. Holmbacka S, Keller J, Eitschberger P, Lilius J
38 Accurate energy modelling for many-core static
39 schedules. Euromicro International Conference on
40 Parallel Distributed & Network-Based Processing
41 Turku, 2015:525-532 102
42 3. Zhao Y, Li X, Ju L, Zong Z. Dependency-based
43 energy-efficient scheduling for homogeneous multi-core
44 clusters. IEEE International Conference on Trust
45 security and privacy in computing and
46 communications. Melbourne, 2013:1299-1306. 107
47 4. Amilifard B., Fallah F., Pedram M. Low-power fanout
48 optimization use in multiple threshold voltage
49 inverters. Proceedings of the international symposium
50 on low power electronics and design. New York, NY,
51 USA, 2015:95-98. 112
52 5. Calhoun B H, Chandrakasan A. Characterizing and
53 modeling minimum energy operation for subthreshold
54 circuits. International Symposium on Low Power
55 Electronics and Design. IEEE, 2004:90-95. 116

6. Donno M, Ivaldi A, Benini L, et al. Clock-tree power
optimization based on RTL clock-gating. Design
Automation Conference. ACM, 2003:622-627.
7. Wason V., Banerjee K. A probabilistic framework for
power-optimal repeater insertion in global
interconnects under parameter variations. Proceedings
of the international symposium on Low power
electronics and design. New York, NY, USA,
2005:131-136.
8. Kim N S, Blaauw D, Mudge T. Leakage Power
Optimization Techniques for Ultra Deep Sub-Micron
Multi-Level Caches. Ieee/acm International Conference
on Computer-Aided Design. IEEE Computer Society,
2003:627.
9. Ananthan H, Kim C H, Roy K. Larger-than-Vdd
Forward Body Bias in Sub-0.5V Nanoscale CMOS.
International Symposium on Low Power Electronics
and Design. IEEE, 2004:8-13.
10. Powell M D, Schuchman E, Vijaykumar T N. Balancing
Resource Utilization to Mitigate Power Density in
Processor Pipelines. Ieee/acm International
Symposium on Microarchitecture, 2005. Micro-38.
Proceedings. IEEE, 2005:294-304.
11. Jayaseelan R. Application-specific thermal
management of computer systems. Singapore: National
University of Singapore, 2009.
12. Wang Z, Xu X, Xiong N, et al. Energy cost evaluation of
parallel algorithms for multiprocessor systems. Cluster
Computing, 2013, 16(1):77-90.
13. Wang Z, Xiong N, Wang H, et al. Whole procedure
heterogeneous multiprocessors low-power optimization
at algorithm-level. Cluster Computing, 2018(1):1-17.
14. Wang Z, Zhao W, Wang H, et al. Three-level
performance optimization for heterogeneous systems
based on software prefetching under power constraints.
Future Generation Computer Systems, 2018, 86:51-58.
15. Wang Z, Cheng L, Zhao W, et al. An architecture - level
graphics processing unit energy model. Concurrency &
Computation Practice & Experience, 2016,
28(10):2795-2810.
16. Kadayif I, Kandemir M, Karakoy M. An energy saving
strategy based on adaptive loop parallelization. Design
Automation Conference. ACM, 2002:195-200.
17. Kim H S, Vijaykrishnan N, Kandemir M, et al.
Adapting instruction level parallelism for optimizing
leakage in VLIW architectures. Acm Sigplan Notices,
2003, 38(7):275-283.
18. Hsu C H, Kremer U, Hsiao M. Compiler-directed
dynamic voltage/frequency scheduling for energy
reduction in microprocessors. Low Power Electronics
and Design, International Symposium on. IEEE,
2001:275-278.
19. Kim N S, Flautner K, Blaauw D, et al. Circuit and
microarchitectural techniques for reducing cache
leakage power. IEEE Transactions on Very Large Scale
Integration Systems, 2004, 12(2):167-184.
20. Zhao R.C, Tang Z.M, Zhang Z.Q, et al. A multithreaded
compiler optimization technology with low power.
Journal of software, 2002, 13(1):1123-1129.
21. Grochowski E, Ronen R, Shen J, et al. Best of both
latency and throughput. Computer Design: VLSI in

- 1 Computers and Processors, ICCD 2004. Proceedings62
- 2 IEEE International Conference on. IEEE63
- 3 2004:236-243. 64
- 4 22. Kadayif I, Kandemir M, Sezer U. An integer linea65
- 5 programming based approach for parallelizing66
- 6 applications in On-chip multiprocessors. Desigr67
- 7 Automation Conference, 2002. Proceedings. IEEE68
- 8 2002:703-708. 69
- 9 23. Kadayif I, Kandemir M, Kolcu I. Exploiting Processor70
- 10 Workload Heterogeneity for Reducing Energy71
- 11 Consumption in Chip Multiprocessors. Conference on72
- 12 Design, Automation and Test in Europe. IEEE73
- 13 Computer Society, 2004:21158. 74
- 14 24. Li N J, Martinez J F. Power-Performance Implications75
- 15 of Thread-level Parallelism on Chip Multiprocessors.
- 16 IEEE International Symposium on PERFORMANCE
- 17 Analysis of Systems and Software. IEEE Computer
- 18 Society, 2005:124-134.
- 19 25. Li J, Martinez J F. Dynamic power-performance
- 20 adaptation of parallel computation on chip
- 21 multiprocessors. In Proceedings of the International
- 22 Symposium on High Performance Computer
- 23 Architecture. 2006:77-87.
- 24 26. Eyerman S, Eeckhout L. Modeling critical sections in
- 25 Amdahl's law and its implications for multicore design.
- 26 Acm Sigarch Computer Architecture News, 2010,
- 27 38(3):362-370.
- 28 27. Nowka K, Carpenter G, Donald E M, et al. A 0.9 V to
- 29 1.95 V dynamic voltage-scalable and
- 30 frequency-scalable 32 b PowerPC processor.
- 31 Solid-State Circuits Conference, 2002. Digest of
- 32 Technical Papers. ISSCC. IEEE International.
- 33 2002(1):340-341.
- 34 28. Delaluz V, Kandemir M, Vijaykrishnan N, et al.
- 35 Energy-oriented compiler optimizations for partitioned
- 36 memory architectures. International Conference on
- 37 Compilers, Architecture, and Synthesis for Embedded
- 38 Systems. ACM, 2000:138-147.
- 39 29. Burd T D. Design issues for dynamic voltage scaling. In
- 40 Proceedings of 2000 International Symposium on Low
- 41 Power Electronics and Design, 2000:9-14.
- 42 30. Stijn Eyerman, Lieven Eeckhout. Modeling critical
- 43 sections in Amdahl's law and its implications for
- 44 multicore design. International Symposium on
- 45 Computer Architecture. ACM, 2010:362-370.
- 46 31. Du I S, Grimes R G, Lewis J G, et al. User's Guide for
- 47 the Harwell-Boeing Sparse Matrix Collection. Tech.
- 48 Report TR-PA-92-96. Toulouse Cedex, France, Oct
- 49 1992.
- 50 32. <http://www.cise.ufl.edu/research/sparse/HBformat/HB/>.
- 51 33. Yang X J, Tang T, Wang G B, et al. MPtostream: an
- 52 OpenMP compiler for CPU-GPU heterogeneous
- 53 parallel systems. Science China(Information Sciences),
- 54 2012, 55(9):1961-1971.
- 55 34. Enhanced Intel SpeedStep Technology for the Intel
- 56 Pentium M Processor, White Paper. March, 2004.
- 57 35. Powerplay Technology.
- 58 [http://www.amd.com/us/products/technologies/ati-powe](http://www.amd.com/us/products/technologies/ati-power-play/Pages/ati-power-play.aspx)
- 59 [r-play/Pages/ati-power-play.aspx](http://www.amd.com/us/products/technologies/ati-power-play/Pages/ati-power-play.aspx).
- 60 36. Advanced Configuration and Power Interface.
- 61 <http://www.acpi.info/>
37. AMD Display Library.
- [http://developer.amd.com/GPU/ADLSDK/Pages/default](http://developer.amd.com/GPU/ADLSDK/Pages/default.aspx)
- [.aspx](http://developer.amd.com/GPU/ADLSDK/Pages/default.aspx).
38. Shuja J, Bilal K, Madani S A, et al. Survey of
- Techniques and Architectures for Designing
- Energy-Efficient Data Centers. IEEE Systems Journal,
- 2016, 10(2):507-519.
39. Shuja J, Gani A, Shamshirband S, et al. Sustainable
- Cloud Data Centers: A survey of enabling techniques
- and technologies. Renewable & Sustainable Energy
- Reviews, 2016, 62:195-214.
40. Shuja J, Bilal K, Madani S A, et al. Data center energy
- efficient resource scheduling. Cluster Computing,
- 2014, 17(4):1265-1277.