Low Power Compiler Optimization
for Pipelining Scaling

Jen-Chieh Chang, Cheng-Yu Lee, Chia-Jung Chen, and Rong-Guey Chang

Department of Computer Science, National Chung Cheng University, Chia-Yi, Taiwan
{cjc99p,lcyu95m,ccj98p,rgchang}@cs.ccu.edu.tw

Abstract. Low power has played an increasingly important role for embedded systems. To save power, lowering voltage and frequency is very straightforward and effective; therefore dynamic voltage scaling (DVS) has become a prevalent low-power technique. However, DVS makes no effect on power saving when the voltage reaches a lower bound. Fortunately, a technique called dynamic pipeline scaling (DPS) can overcome this limitation by switching pipeline modes at low-voltage level. Approaches proposed in previous work on DPS were based on hardware support. From viewpoint of compiler, little has been addressed on this issue. This paper presents a DPS optimization technique at compiler time to reduce power dissipation. The useful information of an application is exploited to devise an analytical model to assess the cost of enabling DPS mechanism. As a consequence we can determine the switching timing between pipeline modes at compiler time without causing significant run-time overhead. The experimental result shows that our approach is effective in reducing energy consumption.

1 Introduction

Since most embedded systems are portable, reducing energy consumption to extend the lifetime of batteries has become a crucial issue. In recent years, many techniques have been proposed to address this issue. DVS is the famous one, which has been demonstrated by much work to be very effective [8,2,12]. It adjusts dynamically voltage and frequency to save power, as indicated in Equation 1.

$$E \propto f \times C \times V^2, \tag{1}$$

where A $\propto$ B means A is in direct ratio to B. However, DVS has no effect on energy saving when the voltage reaches its low bound because it becomes a constant [10]. Fortunately, with reference to Equation 2, energy is in direct ratio not only to the clock frequency and the square of the voltage, but also to instruction-per-cycle (IPC).

$$E \propto f \times V^2 \times t \propto f \times V^2 \times \frac{I_t}{f \times IPC} \propto \frac{V^2}{IPC} \tag{2}$$

Thus, we can reduce energy dissipation at low-voltage level based on IPC. Equation 1 and Equation 2 reveal that IPC is the key to power dissipation at low-voltage level. This fact shows that power will increase in the opposite direction of IPC and motivates our low-power idea to devise a DPS technique to evaluate the IPC and determine the

Fetch   Decode   Execution   Writeback

deep mode

shallow mode

disabled   disabled   disabled disabled

Fig. 1. Pipeline modes of DPS

switching timing between pipeline modes at compiler time. In DPS, the pipeline consists of deep mode and shallow mode, as shown in Figure 1. The deep mode is the default pipeline mode and the shallow mode is designed by dynamically merging adjacent pipeline stages of deep mode, where the latches between pipeline stages are made transparent and the corresponding feedback paths are disabled. In theory IPC is in inverse ratio to the pipeline depth, the IPC of deep mode may be smaller than that of shallow mode [6]. Therefore, executing applications in shallow mode will lead to the reduction of power dissipation. But this statement is not always true. In reality, many factors in deep pipeline mode will influence IPC [15,13]. Here are three examples.

(a) In deep mode, the branch penalty is about twice as large as that of shallow mode. It leads to the reduction of IPC, and then deep mode may consume more power than shallow mode.

(b) The deeper pipeline increases the execution latency. As a result, IPC will become smaller, but power consumption will be larger.

(c) In deep mode, the issue queue must be long and the number of loads of the reorder buffer must be large so that the reorder buffer can hold many in-flight instructions. In this case, this window pressure makes IPC smaller and raises the power dissipation of deep mode.

Hence, if we want to apply DPS to save power at low-voltage level, we must decide when the pipeline enter deep mode or shallow mode depending upon the IPC. Consider the voltage characteristic shown in Figure 2. In the first stage, DVS is applied to save power at high-voltage level. Although reducing voltage is very effective to low power, DVS fails in the second stage when the voltage reaches its lower bound. At the final stage, we can enable DPS to switch the pipeline modes based on IPC. Since IPC is affected by some factors, we can consider their impact on IPC to determine the switching timing between pipeline modes to save energy.

Previous work on DPS was proposed by architects with architectural support [10,7,4,14,3]. However, the research about how to solve this issue with compilation techniques remains open. In this paper, we present an optimization technique to enable DPS with respect to IPC at compile time. We first partition an application into many regions and then calculate the IPC of each region to determine the switching timing between pipeline modes based on our evaluating model. Since our work is performed at compiler time, the run-time overhead will be small and the hardware cost and complexity will be as minimal as possible. The experimental results prove that the energy reduction really benefit from our work.

Fig. 2. Voltage characteristic

Fig. 3. The proposed DPS compilation system

The rest of this paper is organized as follows. Section 2 gives the overview of our work and then presents our approach in detail. The experimental results are shown in Section 3. Finally, we conclude our paper in Section 4.

2   The Proposed Approach

In this section, we focus on how our DPS approach is applied to applications to save energy at compiler time. We first introduce our basic idea in Section 2.1. In section 2.2, we depict the method to partition a code into regions and then present the evaluating function to decide the switching between pipeline modes. The mechanism to enable DPS is given in Section 2.3.

2.1   Basic Idea

Figure 3 shows our compilation system composed of SUIF framework [16], our proposed engine, and Wattch simulator. First, an application is compiled by SUIF as a

control flow graph (CFG) and a data flow graph (DFG). Then the CFG and DFG are analyzed to identify the loop regions and collect information for our evaluating model and identify loop regions. In our work, a code will have another type of regions, non-loop regions, except loop regions. Indeed, a region is an union of basic blocks as a unit that our DPS can manipulate. The evaluation model has two goals: one is to partition the remaining part of the loop regions into non-loop regions and the other is enable each region to enter a suitable pipeline mode. The details of partitioning scheme is described in Section 2.2. To activate DPS, for each region we will insert the DPS-enable function in its entrance at compile time so that it can be executed in proper pipeline mode to save energy based on its IPC. Since the switching between pipeline modes best is very hard to decide, we propose an evaluation model to decide the timing during execution. The evaluation model is presented in detail in Section 2.3. In this way, the code will be switched between different pipeline modes at run time. The experiment is performed on the Wattch simulator [1] with DSPstone and Mediabench benchmark suites.

2.2   Evaluation Model for Switching Pipeline Modes

As mentioned in Section 1, the IPC of deep pipeline mode is not always larger than that of shallow mode. As a consequence, the shallow mode may have better power saving than the deep mode according to Equation 2. Thus to reduce power reduction, each region can enter deep mode or shallow mode based on the IPC during execution. To achieve this objective, we conduct an evaluation model to decide the switching timing between deep mode and shallow mode at compiler time. Since the calculation of IPC closely relates to the size of a region, how to partition a code into regions becomes very important to our work. On one hand, if the region size is too large, we may lose the chances to take advantage of switching pipeline modes to save energy. On the other hand, although the small region size can allow us to apply the DPS optimization to a code, it possibly generates severe switching overheads. However what the size of a region is the optimal solution for our approach is very hard to decide, thus we attempt to seek for a principle to guide our selection in this section. With our observations, since the loops usually dominate execution time and power consumption of a code, they are the key to our decision.

To use the loops to partition a code, they must be identified first and then be referred to divided the remaining part into non-loop regions. Below we present our partitioning approach and evaluation model. Given a code $G = (V, E)$, it is divided into two types of regions, Γ_1 and Γ_2, where $\Gamma_1, \Gamma_2 \subseteq V \times E, G = \Gamma_1 \bigcup \Gamma_2$, and $\Gamma_1 \bigcap \Gamma_2 = \emptyset$. Note $\emptyset$ represents the empty set; that is Γ_1 and Γ_2 are disjoint. We first define Γ_1 in case A and then use Γ_1 to define Γ_2 in case B.

Case A: Γ_1 is the set of regions, which are composed of loops. In other words, each region in Γ_1 only contains a loop.

After defining the loop regions, to classify the non-loop regions, we must present our evaluation model first. Then the evaluation model is applied to loop regions for categorizing the non-loop regions. Below we first give some assumptions and then present how to use them to divided the non-loop part of a code into non-loop regions and finally formalize our evaluation model.

For a code $G = (V, E)$, where $V = \{R_1, R_2, ..., R_n\}$ is the set of regions in G and $E = \{(u, v) \mid u, v \in V \text{and } u \neq v\}$. That is, E is the set of edges between regions. For each region R_i, we assume:

- N_{R_i}: the number of instructions in region R_i, for $i = 1, 2, 3, \cdots$
- N_{b_i}: the number of branches in R_i
- B_{ij}: the jth branch in R_i
- $P_{B_{ij}}$: the probability that B_{ij} is taken
- C_i: the number of clock cycles that does not result from any branch in R_i
- CB_{ij}: the number of clock cycles that results from that B_{ij} is taken
- $\overline{CB}_{ij}$: the number of clock cycles that results from that B_{ij} is untaken

According to Equation 1 and Equation 2, IPC predominate the determination of switching pipeline modes during execution. With the information collected previously and the above terminologies, we can present our evaluating model as follows.

$$\Omega_{R_i} = \sum_{j=1}^{N_{b_i}} [P_{B_{ij}} \times CB_{ij} + (1 - P_{B_{ij}}) \times \overline{CB}_{ij}] + C_i \qquad (3)$$

$$\Theta_{R_i} = N_{R_i} / \Omega_{R_i} \qquad (4)$$

Equation 3 estimates the clock cycles required for each region of the target code, which is also applied to classify the non-loop regions. Equation 4 calculates the IPC for each region and is the guideline to enable DPS. Since the loop regions very likely dominates power dissipation of a code, we use the following parameter Λ with the aid of the evaluation function of regions in Γ_1 to partition the non-loop part of a code. Λ is defined as the maximum of all Ω_{R_i} in Γ_1. Formally, it can be described as follows.

$$\Lambda = \{\Omega_R \mid \exists R \in \Gamma_1 \text{ and } \Omega_R \geq \Omega_{R_i}, \text{for } i = 1, \cdots, n\} \qquad (5)$$

Although the loops usually consume the majority of power dissipation for an application, using λ to partition the non-loop part can be furthermore improved. Instead, we adapt Λ as the new parameter by timing a α to it, where α is a real number. Thus Γ_2 can be defined on the basis of $\alpha\Lambda$ in the following case B.

The followings demonstrate how our partitioning approach works using the code segment selected from Matrix of DSPstone benchmark suite as an example. At first, there are two loops existing in this code; thus they are identified as two loop regions and $\Gamma_1 = \{R_{a_1}, R_{a_2}\}$. Then we compare $\Omega_{R_{a_1}}$ and $\Omega_{R_{a_2}}$ and define $\Lambda = \Omega_{R_{a_2}}$ since $\Omega_{R_{a_2}} > \Omega_{R_{a_1}}$. Afterwards, we let $\alpha = 1$ and thus $\alpha\Lambda = \Lambda$ is applied to categorize non-loop parts into regions of the second type. The evaluation value of the first non-loop code segment is smaller than Λ, so it is immediately identified as the first non-loop region, R_{b_1}. Similarly, the evaluation value of the second non-loop code segment is equivalent to Λ, so it is identified as another non-loop region, R_{b_2}. Finally, the evaluation value of the third non-loop code segment is larger than Λ, thus it is further classified into two regions of the second type R_{b_3} and R_{b_4} so that $\Omega_{R_{b_3}} = \Lambda$ and $\Omega_{R_{b_4}} < \Lambda$. As a consequence $\Gamma_2 = \{R_{b_1}, R_{b_2}, R_{b_3}, R_{b_4}\}$.

2.3   DPS Enabling

After the code partitioning has been done, to enable DPS, we insert a function
DPS_enable () into its head of each region to make it executed in deep mode or shallow
mode. The DPS_enable () is implemented as follows.

$$DPS\_enable() \begin{cases} Lda \quad \#SYSCALL\_DEEP2SHAW \\ Call\_Pal \quad \#131 \\ Lda \quad \#SYSCALL\_SHAW2DEEP \\ Call\_Pal \quad \#131 \end{cases}$$

DPS_enable() provides two functionalities to switch between pipeline modes with the
system call of Alpha 21264 Call_Pal #131. #SYSCALL_DEEP2SHAW switches the
pipeline from deep mode to shallow mode and #SYSCALL_SHAW2DEEP switches
the pipeline from shallow mode to deep mode. In this way, we are able to determine
the timing to switch pipeline modes. The Ω value of each region calculated by Equa-
tion 4 is used for DPS_enable() when the code is compiled by our system. Thus the
code will dynamically enter the deep mode or shallow mode during execution after the
DPS_enable() is inserted into it. Finally the optimized DPSed program is performed on
the modified Wattch simulator.

3   Experimental Results

In Section 3.1, we introduce the system configuration of our work and present the ex-
perimental results in Section 3.2.

3.1   System Configuration

The underlying hardware is the Alpha 21264 processor, which contains one fetch buffer,
four integer ALUs, two floating-point ALUs, one integer multiplier/divider, and one
floating-point multiplier/divider, etc. In instruction window, RUU indicates register up-
date unit and LSQ comprises load queue (LQ) and store queue (SQ). Its main features
are summarized in Table 1. To perform our proposed approach, we extend the pipeline
mode from one mode to two modes. We assume that the original pipelining mode is
shallow mode and the new mode is deep mode by constructed by adding extra four
stages to shallow pipeline. It is designed to dynamically disable one of each pair of
stages by making the latches between pipeline stages transparent so that the processor
can switch between these two pipeline modes. The software configuration is shown in
Table 2. The SUIF compiler infrastructure is the front end of our system and generate
CFG and branch information. The operating system is Tru64 UNIX for 64-bit instruc-
tion set architecture. The Wattch simulator is an architectural simulator that provides
cycle-by-cycle simulation and detailed out-of-order issue with multi-level memory sys-
tem. For keeping consistence with our DPS approach, it has been modified to support
shallow pipeline mode and deep pipeline mode.

3.2   Experimental Results

In our experiment, the deep mode is the default pipeline mode and the shallow mode
is chosen during execution if necessary. The energy reduction benefits by the switch-
ing between deep mode and shallow mode depending on the IPC of a region, which is

Table 1. Hardware configuration

Processor Core	
Pipeline length	4 cycles (shallow mode)
	8 cycles (deep mode)
Fetch buffer	8 entries
Functional units	4 Int ALU, 2 FP ALU, 1 Int mult/div,
	1 FP mult/div, 2 mem ports
Instruction win-dow	RUU=80, LSQ=40
Issue width	6 instructions per cycle: 4 Int, 2 FP
Memory Hierarchy	
L1 D-cache size	64KB, 2-way, 32B blocks,
L1 I-cache size	64KB, 2-way, 32B blocks,
L1 latency	1 cycle
L2	Unified, 2MB, 4-way LRU
	32B blocks, 11-cycle latency
Memory latency	100 cycles
TLB size	128-entry, fully-associative,
	30-cycle miss

Table 2. Software Configuration

OS and Software Configuration	
Profiler	SUIF
Compiler	MachSUIF
OS	Tru64 UNIX
Simulator	Wattch v.1.02 with DPS

calculated by equation 4. The experiment is performed on the Wattch simulator with DSPstone. For each program, the baseline is its original energy dissipation and the optimized energy is measured by performing our DPS approach. This benchmark is compiled by the Alpha compiler with default settings and linked with the intrinsic library on Tru64 UNIX operating system.

Figure 4 shows the energy reduction by comparing the baseline energy and the optimized energy for DSPstone. In this experiment, we let $\alpha = 0.25, 0.5, 1.0, \text{and}, 2.0$ to measure the effects of various partitioning sizes of non-loop regions. The energy saving ranges from 2% to 35%, with a mean of reduction 17.8%. As the partitioning size of non-loop region λ becomes larger, the energy saving decreases slowly. With our observations, the large-size region eliminates some chances to switch the pipeline modes based on IPC and thus slightly increases energy consumption. Nine of these programs including adpcm, complex_multiply, complex_update, dot_product, fft, iir_biquad_one_sections, matrix, real_update, and startup, have better energy saving about from 22% to 35%. The reason is that there are many branches in them and thus our DPS approach can take advantage of them to save energy depending upon the contribution of branch penalty to IPC. With our experiences, our approach works better

Fig. 4. Energy reduction of various λ of profile-based DPS for DSPstone


Fig. 5. Chang in IPC for three cases using DSPstone as a benchmark

for the codes with many loops and larger loops. Note that many programs have large outer loops such as event-driven programs or programs with GUI, which may include almost the entire programs. In this experiment, the typical example for above discussion is matrix testbench, and it has the best energy saving about 35%.

Figure 5 demonstrate the effect on IPC for three cases using DSPstone as the metric. Deep and Shallow represent deep and shallow pipeline modes; DPS indicates that applying our DPS approach to these benchmarks. They are still measured for various partitioning sizes of non-loop regions $0.25\lambda, 0.5\lambda, \lambda,$ and 2λ. For DSPstone, the average IPCs of Deep case and Shallow case are 0.4 and 0.52. In DPS case, the average IPC is 0.45, which is between those of deep mode and shallow mode. In DSPstone, the IPCs of adpcm, fft, fir2dim, and matrix are larger. This is because since they are loop-intensive applications and the loops in them contribute a lot to the increase of IPC.

Figure 6 show the effect of our DPS approach on performance for DSPstone. The latency between pipelining stages is designed to be equivalent to increase performance and achieve resource sharing at each clock. In theory, the performance is in direct ratio to the number of pipelining stages and thus the longer pipeline will lead to the performance. Thus, the processor will result in slowdown when executing in shallow mode. The performance will be degraded if the pipelining stages are merged into shallow mode. In reality, the performance may be degraded due to many factors such as pipelining hazards, branch penalty, or switching overhead between pipeline modes. For each benchmark, the performance of deep mode with λ is the baseline to compare those of deep mode and our DPS method for various λs. For the performance of DSPstone in Figure 6, on average, our DPS approach leads to 6.23% degradation in performance.

Fig. 6. Relative performance of various Λs for DSPstone

By contrast, the performance degradation of shallow pipeline mode is 61.62%, which is almost ten times larger than the above one. Although the DPS switches the pipelining modes based on the IPC to save energy, the switching slows down the processor compared to the high-speed execution in deep mode. In addition, larger λ has a better performance than smaller ones since it causes the pipeline to enter the shallow mode more infrequently.

4   Conclusions

DVS has been proven be very effective in low power optimizations, but it cannot further save energy when the voltage reaches its lower bound. Fortunately, DPS can overcome this limitation by adjusting pipeline modes based on IPC. Previous work resolved this issue with hardware techniques and thus increased hardware cost and design complexity. In this paper, we present a DPS technique to reduce power dissipation by proposing an evaluating model so that they can decide the timing of entering the proper pipeline mode. In contrast, our work can eliminate hardware overhead and reduce energy consumption according to the code behavior at compiler time. To investigate the effect of our approach, we perform the experiment with various criteria for DSPstone and Medieabench. In summary, the results show that smaller partitioning sizes of non-loop regions can create optimization space and loop-intensive applications provide more chances to optimize code to save energy.

References

1. Brooks, D., Tiwari, V., Martonosi, M.: Wattch: A framework for architectural level power analysis and optimizations. In: International Symposium on Computer Architecture (2000)
2. Burd, T., Brodersen, R.: Design issues for dynamic voltage scaling. In: International Symposium on Low Power Electronics and Design (2000)
3. Efthymiou, A., Garside, J.D.: Adaptive pipeline depth control for processor power-management. In: IEEE International Conference on Computer Design: VLSI in Computers and Processors (2002)
4. Ernst, D., Kim, N.S., Das, S., Pant, S., Rao, R., Pham, T., Ziesler, C., Blaauw, D., Austin, T., Flautner, K., Mudge, T.: Razor: A low-power pipeline based on circuit-level timing speculation. In: The 36th International Symposium on Microarchitecture (2003)
5. Gerndt, M.: Automatic parallelization for distributed-memory multiprocessing systems. Phd Thesis (1989)

6. Hartstein, A., Puzak, T.R.: The optimum pipeline depth for a microprocessor. In: ACM/IEEE International Symposium on Computer Architecture (2002)
7. Hiraki, M., Bajwa, R.S., Kojima, H., Corny, D.J., Nitta, K., Shridhar, A., Sasaki, K., Seki, K.: Stage-skip pipeline: A low power processor architecture using a decoded instruction buffer. In: International Symposium on Low Power Electronics and Design (1996)
8. Hsu, C.H., Kremer, U.: The design, implementation, and evaluation of a com-piler algorithm for cpu power reduction. In: The ACM SIGPLAN Conference on Programming Languages Design and Implementation (2003)
9. Kessler, R.E., McLellan, E.J., Webb, D.A.: The alpha 21264 microprocessor architecture. In: Intl Conf. Computer Design (1998)
10. Koppanalil, J., Ramrakhyani, P., Desai, S., Vaidyanathan, A., Rotenberg, E.: A case for dynamic pipeline scaling. In: International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (2002)
11. Kornerup, J.: Mapping powerlists onto hypercubes. Masters Thesis (1994)
12. Krishna, C.M., Lee, Y.-H.: Voltage-clock-scaling adaptive scheduling techniques for low power in hard real-time systems. In: The 6th Real Time Technology and Applications Symposium (2000)
13. Lilia, D.J.: Reducing the branch penalty in pipelined processors. IEEE Computer 21(7), 47–55 (1988)
14. Manne, S., Grunwald, D., Klauser, A.: Pipeline gating: Speculation control for power reduction. In: ACM/IEEE International Symposium on Computer Architecture (1998)
15. Parikh, D., Skadron, K., Zhang, Y., Stan, M.: Power-aware branch prediction: Characterization and design. In: The IEEE Transactions on Computers (2004)
16. G. SUIF. Stanford University Intermediate Format, http://suif.stanford.edu