



FAST: Frequency-Aware Static Timing Analysis

KIRAN SETH

Qualcomm

and

ARAVINDH ANANTARAMAN, FRANK MUELLER, and ERIC ROTENBERG

North Carolina State University

Energy is a valuable resource in embedded systems as the lifetime of many such systems is constrained by their battery capacity. Recent advances in processor design have added support for dynamic frequency/voltage scaling (DVS) for saving energy. Recent work on real-time scheduling focuses on saving energy in static as well as dynamic scheduling environments by exploiting idle time and slack because of early task completion for DVS of subsequent tasks. These scheduling algorithms rely on *a priori* knowledge of worst-case execution times (WCET) for each task. They assume that DVS has no effect on the worst-case execution cycles (WCEC) of a task and scale the WCET according to the processor frequency. However, for systems with memory hierarchies, the WCEC typically *does* change under DVS because of requey modulation. Hence, current assumptions used by DVS schemes result in a highly exaggerated WCET. This paper contributes novel techniques for tight and flexible static timing analysis, particularly well-suited for dynamic scheduling schemes. The technical contributions are as follows: (1) We assess the problem of changing execution cycles owing to scaling techniques. (2) We propose a parametric approach toward bounding the WCET statically with respect to the frequency. Using a parametric model, we can capture the effect of changes in frequency on the WCEC and, thus, accurately model the WCET over any frequency range. (3) We discuss design and implementation of the frequency-aware static timing analysis (FAST) tool based on our prior experience with static timing analysis. (4) We demonstrate in experiments that our FAST tool provides safe upper bounds on the WCET, which are tight. The FAST tool allows us to capture the WCET of six benchmarks using equations that overestimate the WCET by less than 1%. FAST equations can also be used to improve existing DVS scheduling schemes to ensure that the effect of frequency scaling on WCET is considered and that the WCET used is *not* exaggerated. (5) We leverage three DVS scheduling schemes by incorporating FAST into them and by showing that the energy consumption further decreases. (6) We compare experimental results using two different energy models to demonstrate or verify the validity of

This work was supported in part by NSF grants CCR-0208581, CCR-0310860 and CCR-0312695. A preliminary version of this paper appeared in the *IEEE Real-Time Systems Symposium*, 2003 [Seth et al. 2003]. This paper does not necessarily reflect or represent the views of Qualcomm, Inc. Authors' addresses: K. Seth, Qualcomm, Inc., 2000 Center Green Way, Cary, NC 27513; A. Anantaraman, F. Mueller and E. Rotenberg, Departments of CS/ECE and Center for Embedded Systems Research, North Carolina State University, Raleigh, NC 27695-7534; email: mueller@cs.ncsu.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2006 ACM 1539-9087/06/0200-0200 \$5.00

simulation methods. To the best of our knowledge, this study of DVS effects on timing analysis is unprecedented.

Categories and Subject Descriptors: D.4.1 [**Operating Systems**]: Process Management—*Scheduling*; D.4.7 [**Operating Systems**]: Organization and Design—*Real-time systems and embedded systems*

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: Real-time systems, scheduling, dynamic voltage scaling, worst-case execution time analysis

1. INTRODUCTION

Power is an important constraint for mobile, battery-powered embedded devices. Limitations on the lifetime of embedded devices have resulted in advances in embedded architectures to extend the lifetime of devices. Microprocessor designs ranging from low-end 8-bit up to high-end 32-bit embedded architectures (e.g., the Atmel Atmega AVR family on the low end and the Intel XScale on the high end, just to name two extremes) support dynamic adjustment of processing speed to prolong battery life. Generally, two techniques are employed in unison. On one side, dynamic frequency scaling allows the speed of instruction execution to change during the operation of a device. On the other side, dynamic voltage scaling modulates the level of the supply voltage upon demand. Both schemes, referred to as DVS in the following, work hand in hand: When the frequency is lowered by a certain degree, the voltage can be also be reduced to a lower level. Furthermore, both scaling techniques impact the power consumption of a device: power scales linearly with the frequency and quadratically with the voltage. Hence, considerable energy savings may result in a concerted approach of dynamic frequency and voltage scaling [Chandrakasan et al. 1992].

Real-time systems are particularly well-suited to profit from DVS. Because of periodic task execution, it is generally not feasible to utilize the range of sleeping modes that modern processors offer. Tasks are frequently invoked (on a periodic basis in the order of a few milliseconds). The time to enter a sleep mode (and the later wakeup time) is in the order of tens of milliseconds, which generally matches the order of magnitude of a real-time task's period. Hence, suspension in sleep modes is not a viable option for real-time systems. Real-time systems, however, often have task sets that underutilize the processor. Hence, reducing the frequency of execution, while still meeting deadlines through DVS, is a viable option resulting in considerable energy reduction.

Recently, a number of hard real-time DVS, scheduling schemes have been studied, ranging from compiler support [Mosse et al. 2000] over numerous static scheduling approaches [Gruian 2001; Pillai and Shin 2001] to dynamic methods [Pillai and Shin 2001; Aydin et al. 2001; Dudani et al. 2002]. All of these approaches have their own merits in that they provide a solution suitable to certain systems, depending on scheduling methods, and utilization bounds of the task sets and architectural properties, such as scaling overhead.

Any DVS scheduling scheme is subject to the same constraints as other hard real-time systems: The worst-case execution time (WCET) of a task has to be

known *a priori*, i.e., safe bounds on a task's execution time have to be obtained. Prior work on static timing analysis provides the means to derive relatively tight WCET bounds for simple embedded architectures, which are provably safe. A number of research groups have addressed various issues in the area of bounding the WCET of a real-time task. Conventional methods for static analysis have been extended from unoptimized programs on simple CISC processors to optimized programs on pipelined RISC processors and from uncached architectures to instruction and data caches [Park 1993; Lim et al. 1994; Healy et al. 1995; Mueller 2000; White et al. 1999; Li et al. 1996]. The challenge of static timing analysis is to provide not only safe but also tight bounds on the WCET in order to impose a high enough processor utilization. These analysis approaches result in tight bounds for deterministic microarchitectures with simple components.

In the context of DVS, static timing analysis is generally assumed to remain valid with frequency scaling. The conjecture is that reducing a processor's frequency still results in the *same* number of cycles of execution for a task. Hence, considering the processor frequency should suffice to derive safe WCET bounds. However, this simplistic view generally *does not hold* for any realistic architectures. Consider the impact of memory references. Any instruction or data reference that is resolved through a main memory access operates at external bus frequency. However, bus frequencies generally diverge from internal processor frequencies and they do *not* scale at the same rate as does the DVS scaling. For example, the first generation Compaq Ipaq has a StrongArm microprocessor (SA-1110) that scales at eight frequencies, but only supports two different external bus frequencies [Corp.].

In short, when static timing analysis is applied in the context of DVS, tightness and safety assumptions may no longer hold: WCET bounds may either not be tight (considerable overestimation upon fast memory operations for lower processor frequencies) or are no longer safe (underestimation potentially leading to missed deadlines upon a reduced data bus frequency). As a result, the memory latency also has to be adjusted to discrete values according to dynamic settings for execution frequencies and memory latencies. Instead of obtaining one discrete WCET through static timing analysis, different values for each processor frequency/bus frequency pair would have to be obtained. While this may still be a feasible approach for a static schedule and for a small number of such frequency pairs, it becomes infeasible for dynamic scheduling paradigms or a large number of frequency pairs. For certain scheduling, approaches that exhibit intratask DVS, such a static approach, becomes impossible if tight bounds for the WCET are to be determined because the point of frequency changes during task execution is typically unknown at static time, e.g., because of dynamic scheduling, preemption, and early completion.

The contribution of this paper is to remedy this problem by promoting a new methodology for frequency-aware static timing analysis (FAST). Instead of obtaining a WCET bound for each frequency pair, FAST takes static timing analysis to a novel level suitable for dynamic scheduling. FAST expresses WCET bounds as a parametric term whose components are frequency-sensitive parameters. On the one side, cycles are interpreted in terms of the processor

frequency; on the other hand, memory accesses are expressed in terms of the memory latency overhead because of the external bus speed. This parametric expression of the WCET allows one to determine on-the-fly the WCET for a given frequency pair. This is particularly appealing when scheduling decisions occur dynamically and when the number of frequency pairs becomes large, such as is the case with state-of-the-art processors with fine-grained frequency settings [Intel 2000].

Another contribution of this paper is its methodology to evaluate benefits of energy conservation. Instead of using a single simulation methodology, as done in most prior work, two different analytical approaches are employed. A commonly used power estimation model on one side is compared to a more detailed power model that separately considers architectural components. The former is based on estimating power via its proportional relation to processor frequency and the square of the voltage, while the latter, known as the Wattch model [Brooks et al. 2000], considers power consumption for the register file, functional units, branch prediction, etc., based on their dynamic utilization in conjunction with frequency and voltage levels. The comparison shows a considerable difference in estimated absolute energy consumption, which indicates that absolute values from simulations can be controversial. Both models loosely agree in that they show an overall reduction in energy consumption because of our approach, which validates our claims about the potential of FAST.

In the following, we detail the technical innovations necessitated by DVS to ensure that safe and flexible WCET predictions may be obtained. We provide motivating examples, discuss the design of our FAST analysis tool, and we show the feasibility of our approach in a set of experiments that demonstrate flexibility and competitiveness while still providing tight bounds on the WCET. Related, as well, are future work and a summary, which conclude our contributions.

2. EFFECTS OF FREQUENCY SCALING ON WCET

In this section, we motivate the need for a parametric frequency model and assess the challenges of supporting this novel model in a static timing analysis tool. We also describe the parametric frequency model in detail and we illustrate the key features in examples.

2.1 Motivation

Real-time systems that use DVS-based scheduling, scale the WCET, assuming that the number of worst-case execution cycles (WCEC) remains constant, even with a change in the frequency. This assumption holds for systems where the memory latency can scale with processor frequency (systems with on-chip memory). In contrast, for a system where the memory latency does not scale with processor frequency (systems with dynamic memory and memory hierarchies), the WCEC of a task *does not* remain constant when the frequency is scaled, since an increase in the frequency typically increases the number of cycles required to access memory. This behavior is caused by a constant access latency for memory references, regardless of changing processor frequencies.

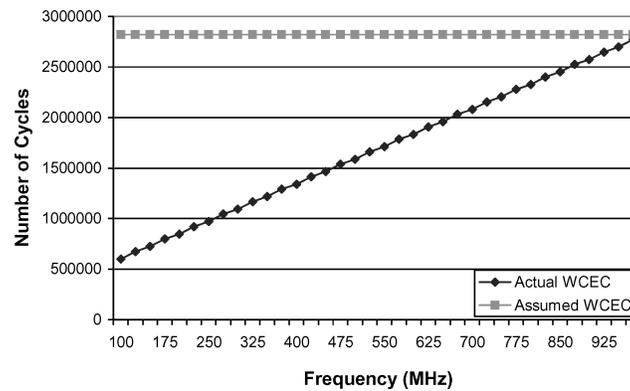


Fig. 1. Actual vs. assumed WCEC for *fft*.

Notice that the memory access time depends on the front-side bus (FSB) instead of the processor frequency. Either the FSB has a constant frequency or it does not provide scaling at the same rate as a processor, i.e., FSB frequencies typically are constrained by a considerably smaller range. Let us assume a constant FSB frequency, which is most common.

By assuming that the WCEC remains constant, one ignores the fact that the WCEC *reduces* with frequency, which results in overestimations of the WCET. Figure 1 depicts results for the C-lab real-time benchmark *fft*, where the actual WCEC for a system with a memory hierarchy is compared to a constant WCEC. The WCEC for the benchmark was calculated for a simple in-order pipeline with instruction and data caches. In this example, it is assumed that the memory access latency is constant. Figure 1 illustrates that the number of WCEC increases proportionally with the processor frequency. This results from an increasing number of wait cycles for a constant time memory latency as the frequency increases. The slope of the actual WCEC depends on the number of accesses to main memory (and the latency to frequency ratio). Hence, the slope depends on the number of misses in the instruction and data caches combined. Therefore, the accuracy of paradigms that measure the worst-case behavior of the instruction and data caches not only control the accuracy of the WCEC, but also affect the accuracy by which the WCEC can be scaled with frequency. Figure 2 depicts the equivalent WCET to the two WCEC curves in Figure 1. The actual WCET depicted indicates the assumption of a constant WCEC, independent of frequency modulations, result in considerable overestimations of the WCET.

The objective of the work described in this paper is to accurately model the actual WCEC and, thereby, the actual WCET of real-time tasks. We derive a parametric frequency model for this purpose. The model provides WCET bounds that remain tight and accurate throughout any frequency range. The parametric model complements real-time systems employing a DVS-base scheduling scheme and it is paramount to achieving higher energy savings. Ignoring the change in WCEC with frequency results in considerably smaller energy savings.

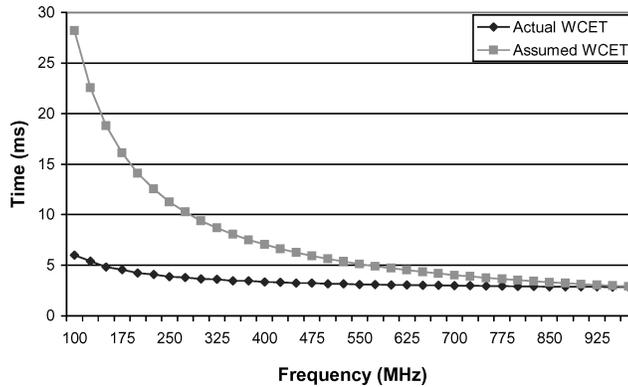


Fig. 2. Actual vs. assumed WCET for *fft*.

2.2 Parametric Frequency Model

Our parametric frequency model can be used for timing analysis with any simple in-order single-issue pipeline. The model is applicable to systems with or without a memory hierarchy. We consider the model in a system with a memory hierarchy in the following and we contribute solutions to the technical challenges posed. We assume that the system is equipped with an on-chip instruction and data cache and that the main external memory has a constant access latency. For now, let us assume that a static timing analyzer has detected a worst-case path, which is an assumption that is lifted in Section 3.2. To accurately model the WCET in systems with memory hierarchies, we propose a parametric frequency model that accurately captures the effect of frequency scaling by splitting the WCEC of a task into two components. The first component, i , captures the ideal number of cycles required to execute the task assuming perfect caches. In other words, i does not scale with frequency. The second component, m , counts the total number of instruction and data cache misses for the task. m is the part of the WCEC that scales with frequency and depends on the memory access latency. If a system without caches is considered, i would count the total number of cycles used for nonmemory operations, while m would count the total number of memory references. Thus, the WCEC is expressed as follows:

$$WCEC = i + mN \quad (1)$$

where N is the number of cycles required to access the memory, which depends on the latency of the memory and the frequency of the processor. For a uniform memory latency, the WCEC can be easily converted into the WCET by dividing by the frequency. This frequency model can accurately model the actual WCET, because it separates the WCEC into components, one that scales and one that does not scale with processor frequency.

The following examples are presented to show that the parametric model can capture the effects of different sequences of instructions in a task. Only sequences that contain data or instruction cache misses are of concern since they are affected during frequency scaling. A sequence of instructions without

A: add R2, R1, R3
 B: load R4, [M1]
 C: add R2, R1, R4
 D: add R2, R1, R5

Fig. 3. Sample instruction sequence.

Cycles	1	2	3	4	5	6	-	-	12	13	14	15	16	17	18	19
IF	A	B	B	B	B	B	-	-	B	C	D					
ID		A								B	C	D				
IS			A								B	C	D			
EX				A								B	C	D		
MEM					A								B	C	D	
WB						A								B	C	D

Fig. 4. Example 1: instruction cache miss.

any cache misses can be captured exclusively by the i component and represents a trivial example of our parametric model. For the following examples, let $N = 10$, as shown in the figures below. We assume separate instruction and data caches and frequency scaling under our model with an arbitrary simple in-order pipeline.

Consider a sequence of four instructions, as shown in the Figure 3. This instruction sequence is executed in a processor with a simple six-stage in-order pipeline. The pipeline stages are fetch (IF), decode (ID), issue (IS), execute (EX), memory access (MEM), and write-back (WB).

1. In Figure 4, we observe the effects of an instruction cache. Consider instruction B resulting in a miss. While instruction B misses in the instruction cache, all other cache accesses result in hits. Since instructions are stalled until the miss on B is resolved, the number of cycles involved can be separated into two components. With $i = 9$ and $m = 1$ in Eq. (1), the WCEC is accurately captured by our model as $WCEC = 9 + 1N$. Hence, the WCEC is accurately modeled for any value of N resulting in an accurate WCET, regardless of frequencies.
2. In Figure 5, we observe the effects of a data cache miss. Instruction B misses in the data cache while all other cache accesses are hits. With $i = 9$ and $m = 1$, the WCEC is again calculated as $9 + 1N$. Since the data miss stalls subsequent instructions, one can separate the number of cycles required for the memory access. However, had the Instruction C or any other stalled instruction performed any useful work instead of being stalled, a potential for overestimation would occur for the model, e.g., for multicycle floating-point operations, and branch mispredictions. Any such overestimation results from the overlap of useful cycles with the memory stall. In our model, the i component counts these useful cycles while the m component counts data miss. Overlap would not be considered by the model. For example, if instruction C took an extra cycle to execute, the new WCEC would become $10 + 1N$. The model does not consider the overlap between the data miss and the extra cycle used by instruction C. A similar problem is also observed

Cycles	1	2	3	4	5	6	-	-	16	17	18	19
IF	A	B	C	D								
ID		A	B	C	D							
IS			A	B	C	D	-	-	D			
EX				A	B	C	-	-	C	D		
MEM					A	B	-	-	B	C	D	
WB						A				B	C	D

Fig. 5. Example 2: data cache miss.

Cycles	1	2	3	4	5	6	-	-	13	14	15	16	-	24	25	26	27
IF	A	B	C	C	C	C	-	-	C	D							
ID		A	B							C	D						
IS			A	B								C	D	-	D		
EX				A	B							C	-	C	D		
MEM					A	B	-	-	B	B	B	B	-	B	C	D	
WB					A										B	C	D

Fig. 6. Example 3: instruction + data cache miss.

in example 1, if the instruction miss overlaps with a high execution latency instruction.

The potential for overestimations implies that the obtained WCET obtained still provides an upper bound on the execution time, albeit not necessarily a tight one. Removing overestimations because of instructions with high execution latencies, however, is non-trivial because instructions may have different execution latencies. Subsequent experiments show that these design choices have a diminishing affect on the tightness of WCET bounds.

3. In Figure 6, we observe the effects of a simultaneous instruction and data cache misses. Instruction B results in a data cache miss while the instruction C results in an instruction cache miss. All other cache accesses are hits. With $i = 9$ and $m = 2$, the $WCEC = 9 + 2N$. The instruction and the data cache misses cannot be serviced together. Hence, instruction B is stalled until instruction C's cache miss is serviced. The model captures all sequences of instructions, where a cache miss stalls yet another cache miss. Notice that the two misses in question need not result from consecutive instructions. We observe some overestimation because of overlapping of some work with the miss cycles.

In the above examples, different combinations of cache misses were considered, which can occur in a simple pipeline. In the presence of these misses, the parametric model accurately captures the worst-case timing behavior for any sequence of instructions. Overestimation is expected when a high execution latency operation overlaps with a miss or when an I-cache miss overlaps with a D-cache miss.

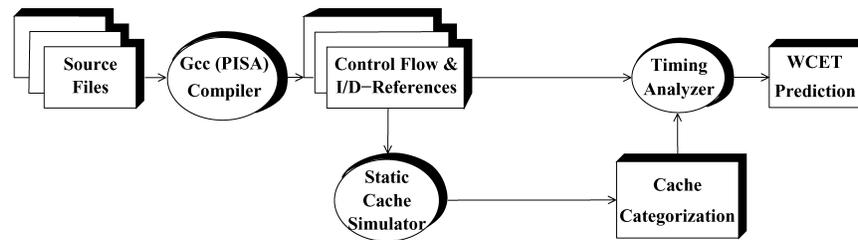


Fig. 7. Obtaining safe WCET bounds.

3. TIMING ANALYSIS

In this section, we describe conventional static timing analysis and briefly contrast the approach to dynamic timing analysis methods. We specify the novel enhancements necessitated by DVS to adapt conventional static timing analysis to a frequency-aware static timing analysis (FAST) tool.

3.1 Static Timing Analysis

Schedulability analysis for hard real-time systems requires that WCET be safely bounded in order to ensure feasibility of scheduling a taskset for a given scheduling policy, such as rate-monotone and earliest-deadline-first scheduling [Liu and Layland 1973]. If the execution time of a task were obtained through dynamic timing analysis based on experimental or trace-driven approaches, these values would not provide a safe bound of the WCET [Wegener and Mueller 2001]. On the one side, it is difficult to determine the worst-case input set even for moderately complex tasks that would exhibit the WCET and to perform exhaustive testing over the entire input space is infeasible, except for trivial cases. On the other side, even if the worst-case input set was known, the interaction between the software and hardware might cause the task to exhibit its WCET for a different input set. The cause of this behavior is architectural complexity, such as complex pipelines and caching mechanisms.

Static timing analysis is a viable alternative to dynamic timing analysis and, while various static approaches have been studied, we will constrain ourselves to one such toolset without loss of generality [Healy et al. 1999; Mueller 2000; White et al. 1999]. The WCET bounds obtained by static timing analysis provide a guaranteed upper bound on the computation time of a task. Static timing analysis performs the equivalent of a traversal over all execution paths to determine timing information independent of a program trace and without tracking values or program variables. Loop bodies only require a few traversals to determine the worst-case behavior of the entire loop because of an efficient fixed-point approach. As the execution paths are traversed, the behavior of the architectural components along the execution paths is captured. The paths are composed to form loops, functions, and, ultimately, the entire application to calculate both WCEC and WCET.

Figure 7 depicts an overview of the organization of this timing analysis toolset. An optimizing compiler has been modified to produce control flow and branch constraint information as a side effect of the compilation of a source file. The original research compiler VPCC/VPO [Benitez and Davidson 1988] was

Table I. Instruction Categories for WCET

Cache Category	Definition
Always miss	Instruction may not be in cache when referenced.
Always hit	Instruction will be in cache when referenced.
First miss	Instruction may not be in cache on 1st reference for each loop execution, but is in cache on subsequent references.
First hit	Instruction is in cache on 1st reference for each loop execution, but may not be in cache on subsequent references.

replaced by GCC with a Portable Instruction Set Architecture (PISA) backend that interfaces with SimpleScalar. Real-time applications are compiled into assembly code using the GCC PISA compiler. The control-flow graph and instruction, as well as data references, are extracted from the assembly code. Upper bounds on the number of iterations performed by loops are provided—a prerequisite for performing static timing analysis. A static instruction cache simulator uses the control-flow information to construct a control-flow graph of the program that consists of the call graph and the control flow of each function. The program’s control-flow graph is then analyzed and a caching categorization for each instruction and data reference in the program is produced. Separate categorizations are provided for each loop level in which the instructions and data references are contained. The categorizations for instruction references are described in Table I. Next, the timing analyzer uses the control flow and constraint information, caching categorizations, and machine-dependent information (e.g., pipeline characteristics) to calculate bounds on the WCET.

The approach in this paper differs from our prior toolset as follows. Our tool separates static I-cache and D-cache (instruction/data cache) analysis. The D-cache analysis currently lacks sufficiently detailed information about references for the GCC compilation phase and D-cache analysis does not fully match the SimpleScalar model. The focus of this paper is on enhancing the timing analyzer with respect to the FAST model and PISA instruction set. However, since we use our SimpleScalar-based architectural simulation environment [Anantaraman et al. 2003] to validate our approach, we have to make simplifying assumptions about data caches. Specifically, we assume a constant number of data cache accesses to be misses for each application to model compulsory misses. The remaining references are considered to be hits, which models a sufficiently large cache. This simplifying assumption does not affect the design of FAST, i.e., our model supports a more precise static data cache analysis as well.

The timing analyzer uses the control-flow information and loop bounds, caching categorizations, and pipeline description to derive WCET bounds. The pipeline simulator considers the effect of structural hazards (an instruction occupying the universal function unit for multiple cycles), data hazards (a load-dependent instruction stalls for at least one cycle if it immediately follows the load), branch prediction (backward–taken/forward–not taken), and cache misses (derived from caching categorizations) for alternative execution paths through a loop body or a function. Static branch prediction is easily accommodated by worst-case analysis: the misprediction penalty is added to the nonpredicted path (not taken path for backward branches and taken path for

forward branches). Path analysis (see below) selects the longest execution path as usual. Once timings for alternate paths in a loop are obtained, a fixed-point algorithm (quickly converging in practice), is employed to safely bound the time of the loop, based on the its body's cycle counts.

The fixed-point approach generally requires path analysis for only a few iterations. Given the longest path for the first iteration, the next-longest path is determined for the second iteration, which may differ from the original path because of caching effects. The lengths of these paths are monotonically decreasing because of cache effects, and once we reach a fixed-point, subsequent loop iterations can be safely approximated by this fixed-point timing value. When the longest paths of consecutive iterations are combined, we account for the pipeline overlap between the tail of the earlier path and the head of the path that follows. The alternative—no overlap—is tantamount to draining the pipeline between iterations. Using this fixed-point approach, the timing analyzer ultimately derives WCET bounds, first for each path, then for loops, and, finally, for functions within the program. A timing analysis tree is constructed, where each node of the tree corresponds to a loop or function. Nodes in the tree are processed in a bottom-up manner. In other words, the WCET for an outer loop/caller is not calculated until the times for all of its inner loops/callees are known. This means that the timing analyzer predicts the WCET for programs by first analyzing the innermost loops and functions before proceeding to higher-level loops and functions, eventually reaching the tree's root (e.g., `main()`). For our purposes, the timing analysis tree provides a convenient method for obtaining WCET for a specific scope, in particular, for subtasks. From the description in this section, it becomes evident that static timing analysis is nontrivial, even for simple pipelines.

3.2 Frequency-Aware Static Timing Analysis

The static timing analysis tool calculates the WCEC for a particular task. However, static timing analysis has to be performed whenever the processor frequency is changed. Reassessing the WCET bound is paramount to temporal safety, since a change in the processor frequency causes a change in the number of cycles required to access the memory, since front-side bus frequencies do not scale at all (or at least not at the same rate). Because of the change in memory latency, the WCEC information for different paths changes, which may result in a different worst-case path than before. Our frequency model can be elegantly incorporated into static timing analysis such that it calculates the number of cycles for *each possible worst-case path* in the program. The following technical innovations to the static timing analysis framework support such flexible calculations.

Instead of using the memory access cycles to simulate the sequence of instructions in the pipeline, the ideal number of cycles is calculated, assuming all cache accesses to be hits. The instruction and data cache misses are accumulated as a side effect to compose a first order polynomial equation describing the WCEC.

Static timing analysis requires different paths through the same node (loop or function) to be compared. The path with the worst WCEC is used as the

WCEC for the node. After integrating the frequency model into the framework, one has to compare two equations to determine which one was to result in a larger number of execution cycles. The challenge here is posed by having to consider both equations: One of them (e.g., for path one) has greater WCEC for some range of frequencies, while the other (for path two) has greater WCEC for the rest of the frequency range. Remember that the frequency model is a first-order polynomial. Consider the case where two equations intersect, i.e., both polynomials have a common solution. We propose three approaches to address this problem.

1. One can maintain an ordered list of equations and the ranges where subsequent polynomials represent a larger WCEC than previous ones. Since the frequency model is a first-order polynomial with different slopes, there exists an intersection point constraining the range for each equation.
2. Alternatively a curve-fitting equation could capture the effects of both equations. This obviates the need for maintaining large numbers of equations, but increases the complexity of the parametric equation. A higher-order polynomial with strict upper bounds on each base polynomial would provide a relatively close fit. The resulting curve would not be as tight as in case (1), but may suffice if the slopes of the original polynomials do not diverge significantly. This would impose more overhead on dynamic scheduling schemes that have to perform additional arithmetic to evaluate the equation upon any scheduling action.
3. Another, easier solution is to declare a valid range of frequencies for the processor. If two equations intersect outside the given range, we simply have to choose the equation that provides the higher WCEC within the valid range. If two equations intersect within this specified range, we use a simple curve-fitting technique through a first-order polynomial that provides a WCEC greater or equal to the values of either of the original equations.

By using one of the above techniques, we ensure that a FAST equation obtained always provides an upper bound on the WCEC of the task, regardless of the chosen frequency. For our FAST framework, we have used the third, which is the easiest technique to bound FAST equations.

4. FAST-DVS SCHEMES

Most DVS scheduling algorithms use the assumption that the WCEC is constant with frequency when scaling the WCET. By not considering the effect on WCEC during frequency modulation, DVS schemes assume a considerably overestimated WCET. Thus, DVS schemes fail to completely utilize available slack, because the scaled WCET is not a tight bound. We have implemented our parametric frequency model as the FAST framework. Parametric equations obtained by FAST can be used in DVS scheduling schemes to ensure that the scaled WCET remains an accurate and tight bound of the execution time for a task. Thus, we can increase the efficiency of DVS schemes and further reduce the energy consumption of the system.

DVS schemes can execute a task set at a lower frequency provided that a schedulability test deems the task set feasible and tasks do not exceed their WCET. For DVS schemes based on earliest-deadline-first (EDF) scheduling, the schedulability test expressed in Eq. 2 must be satisfied by the task set to ensure feasibility. Eq. 2 represents the original Liu and Layland utilization test of the system without considering frequency scaling [Liu and Layland 1973].

$$\frac{C_1}{P_1} + \frac{C_2}{P_2} + \dots + \frac{C_n}{P_n} \leq 1 \quad (2)$$

C_1, C_2, \dots, C_n represent the WCET for each of the n tasks. P_1, P_2, \dots, P_n represent the respective periods of the tasks. As is common in base EDF, tasks' deadlines are assumed to be equal to their periods. Let us now consider a scaling factor α that identifies the actual (scaled) frequency, such that $\alpha = f_c/f_m$, where f_c is the scaled frequency, and f_m is the maximum processor frequency.

Next, let us express Eq. (1) in time instead of cycles where the number of cycles, N , is expressed in terms of the actual frequency, f_c , and the memory latency, L , using the relation $N = L \times f_c$, and f_c is then substituted by $f_m \times \alpha$ by definition of α .

$$C = \frac{WCEC}{f_c} = \frac{i + mLf_c}{f_c} = \frac{i + mLf_m\alpha}{f_m\alpha} \quad (3)$$

Recall that Eq. (2) does not consider the effect of frequency scaling on WCET. By combining Eq. (3) with Eq. (2), we yield a more accurate scaling factor by taking the effects of frequency scaling on WCET into account, as seen in Eq. (4).

$$\frac{i_1 + \alpha m_1 L f_m}{P_1 f_m \alpha} + \dots + \frac{i_n + \alpha m_n L f_m}{P_n f_m \alpha} \leq 1 \quad (4)$$

By solving for α , we get:

$$\begin{aligned} \sum_{j=1}^n \frac{i_j + \alpha m_j L f_m}{P_j f_m} &\leq \alpha \\ \sum_{j=1}^n \frac{i_j}{P_j f_m} + \sum_{j=1}^n \frac{\alpha m_j L}{P_j} &\leq \alpha \\ \sum_{j=1}^n \frac{i_j}{P_j f_m} &\leq \alpha - \alpha \sum_{j=1}^n \frac{m_j L}{P_j} \\ \sum_{j=1}^n \frac{i_j}{P_j f_m} &\leq \alpha \left(1 - \sum_{j=1}^n \frac{m_j L}{P_j} \right) \\ \frac{\sum_{j=1}^n \frac{i_j}{P_j f_m}}{\left(1 - L \sum_{j=1}^n \frac{m_j}{P_j} \right)} &\leq \alpha \\ \frac{\sum_{j=1}^n \frac{i_j}{P_j}}{f_m \left(1 - L \sum_{j=1}^n \frac{m_j}{P_j} \right)} &\leq \alpha \end{aligned} \quad (5)$$

```

EDF-test( $\alpha$ ):
    if  $\frac{\sum_{j=1}^n i_j/P_j}{f_m(1 - L \sum_{j=1}^n m_j/P_j)} \leq \alpha$  return true ;
    else return false;
select-frequency:
    use lowest frequency
     $f_k \in \{f_1, \dots, f_m \mid f_1 < \dots < f_{max}\}$ 
    such that EDF-test( $f_k/f_{max}$ ) is true;

```

Fig. 8. FAST-static voltage scaling for EDF.

The scaling factor in Eq. (5) results in a much lower frequency f_c . The WCET used is not exaggerated, and slack is exploited efficiently.

In our implementation work, we integrated FAST equations into DVS-EDF scheduling as proposed by Pillai and Shin through (a) static voltage scaling, (b) cycle-conserving RT-DVS, and (c) look-ahead RT-DVS [Pillai and Shin 2001]. With only minimal changes to the original algorithms, we integrated the FAST equations into the respective DVS schemes, thereby improving energy savings obtained.

4.1 FAST-Static Voltage Scaling

The static voltage-scaling scheme introduced by Pillai and Shin [2001] uses the modified EDF test shown in Eq. (2) to calculate the scaling factor α . This algorithm uses all static slack in the system. The processor frequency for the entire task set is set statically. Dynamic slack produced during runtime because of early completion of tasks is not considered for frequency scaling. The FAST equations for the WCET can be integrated into the static voltage scheme, as shown in Figure 8. Equation (1) represents the WCET of all tasks and the scaling factor is calculated using Eq. (5). The FAST-static voltage-scaling algorithm performs better than the original static voltage scheme, because it considers the portion of WCET that scales with frequency.

4.2 FAST-Cycle-Conserving RT-DVS

The cycle-conserving RT-DVS by Pillai and Shin [2001] calculates the utilization for a task set at every task release and task completion. Upon task release, the utilization is calculated based on the WCET. Upon task completion, the utilization is calculated by considering the actual execution time of the completed task instead of the WCET. This algorithm uses the static slack available in the system as well as the dynamic slack generated due to early task completions. Figure 9 shows the necessary modifications to the original algorithm to incorporate the FAST equations.

The FAST-cycle-conserving DVS scheme outperforms the original scheme, since it takes the actual execution times as well the scaling levels of previous tasks into account. The scheme derives the current system utilization after task completion by considering the actual execution time. In FAST-cycle-conserving RT-DVS, the total number of cycles and the total number of misses experienced

```

select-frequency():
  use lowest frequency
   $f_k \in \{f_1, \dots, f_{max} | f_1 < \dots < f_{max}\}$ 

  such that  $\frac{\sum_{j=1}^n i_j/P_j}{f_m(1 - L \sum_{j=1}^n m_j/P_j)} \leq f_k/f_{max}$  ;

upon task-release( $T_j$ ):
  set  $i_j = i_{WCET}$  and  $m_j = m_{WCET}$  ;
  select frequency();
upon task-completion( $T_j$ ):
  set  $i_j = i_{actual}$  and  $m_j = m_{actual}$  ;
  /* $m_{actual}$  are the actual number of misses
   for this invocation,
    $i_{actual}$  are the ideal number of cycles for
   this invocation not counting the miss cycles*/
  select frequency();

```

Fig. 9. FAST-cycle-conserving DVS for EDF.

by a task are determined during executing, e.g., by hardware counters, which have become quite common for modern architectures. The actual execution time is also converted into a FAST equation to consider its scaling with frequency. The system utilization and the scaling factor are calculated through Eq. (4) and (5).

4.3 FAST-Look-Ahead RT-DVS

The look-ahead RT-DVS schemes by Pillai and Shin [2001] finds the minimum amount of work that may be performed between now and the next scheduling event without missing any deadlines. All work is deferred till the last possible moment, also referred to as last-chance scheduling [Chetto and Chetto 1989]. As a side effect, the frequency may be increased as execution approaches a deadline. In practice, most tasks complete execution early, i.e., prior to their WCET. Hence, the frequency rarely has to be raised to complete by a deadline. This algorithm also uses all the static slack (idle) as well as most of the dynamic slack. Figure 10 depicts the modified original algorithm to integrate the FAST equations into the DVS scheme. Figure 10 also shows a modification to the look-ahead RT-DVS algorithm for *task completion* by setting $c_left_j = C_j$ (see appendix). The FAST-look-ahead scheme also takes advantage of FAST equations to lower energy consumption of the algorithm. The terms i_left and m_left describe the computation left in the form of a FAST equation. Hardware counters are employed to track total cycles completed and total misses inflicted, while a task is executing. The s component shown in Figure 10 cannot be converted into a FAST equation unless considerable changes are made to the algorithm. Doing so would make the algorithm more aggressive, leading to lower frequencies. To avoid excessive modifications, only the next scheduled task is expressed in the form of a FAST equation. The experiments show that the performance of the algorithm is improved even with minimal modifications to the algorithms.

```

select-frequency(x):    use lowest frequency
                        $f_k \in \{f_1, \dots, f_{max} | f_1 < \dots < f_{max}\}$ 
                       such that  $x \leq f_k / f_{max}$ ;
upon task-release( $T_j$ ):
  set  $c\_left_j = C_j$ ,
   $i\_left_j = i\_wcet_j$  and  $m\_left_j = m\_wcet_j$ ;
  defer();
upon task-completion( $T_j$ ):
  set  $c\_left_j = C_j$ ,
   $i\_left_j = i\_wcet_j$  and  $m\_left_j = m\_wcet_j$ ;
  defer();
during task-execution( $T_j$ ):
  decrement  $c\_left_j$ ,  $i\_left_j$  and  $m\_left_j$ ;
defer():
  set  $U = C_1/P_1 + \dots + C_n/P_n$ ;
  set  $s = 0$ ;
  for  $j = 1$  to  $n$ ,  $T_j \in T_1, \dots, T_n | D_1 \geq \dots \geq D_n$ 
    /*Note: reverse EDF order of tasks*/
    set  $U = U - C_j/P_j$ ;
    set  $x_j = \max(0, c\_left_j - (1 - U)(D_j - D_n))$ ;
    set  $U = U + (c\_left_j - x_j)/(D_j - D_n)$ ;
    set  $s = s + x_j$ ;
   $s = s - x_n$ 
   $t = D_n - \text{current\_time}$ 
  
$$x = \frac{(i\_left_n + s)/t}{f_m(1 - L \times m\_left_n/t)}$$
;
select-frequency(x);

```

Fig. 10. FAST-look-ahead DVS for EDF.

5. EXPERIMENTAL FRAMEWORK

The experimental framework is divided into two sections. The first section is devoted to comparing the WCEC calculated using FAST equations, obtained from the FAST framework, to the WCEC, obtained from the traditional static timing analysis tool. The second section tests and compares FAST-DVS algorithms with the original DVS algorithms proposed by Pillai and Shin [2001].

We assess the energy consumption using two different models for each case: the classical model, based on $E \sim V^2 f$, and an architectural resource model Wattch [Brooks et al. 2000]. The former is widely used in early general-purpose DVS work and in real-time systems to evaluate DVS-scheduling algorithms. The latter has become popular in the architectural community since it integrates with SimpleScalar [Burger et al. 1996]. To provide a proper comparison between the two, the $V^2 f$ model was also integrated into our SimpleScalar-based simulator [Anantaraman et al. 2003]. Notice that the results reported in Section 6 differ from our preliminary paper [Seth et al. 2003], which reported $V^2 f$ -based energy readings obtained from a scheduler simulator. Our new results consistently utilize the SimpleScalar architectural simulator, which, besides the Wattch model, we have enhanced by a real-time scheduler and an

implementation of three DVS scheduling schemes based on EDF, as proposed by Pillai and Shin [2001]. Hence, the DVS scheduling and task dispatch overheads are considered in our framework. The overhead of voltage/frequency switching itself may be considered as part of these overheads.

5.1 Testing the FAST Framework

We redesigned our static timing analyzer [Healy et al. 1999] to create the FAST framework. The FAST tool, like its predecessor [Anantaraman et al. 2003], is based on the portable ISA (PISA) used by the SimpleScalar tool set. All instruction execution latencies are based on the MIPS R10K latencies. Specifically, a constant memory latency of 100 ns is used. We use a 8 KB direct-mapped instruction cache and a 8 KB direct-mapped data cache. For the instruction cache categorizations, the static cache simulator of our existing tool set is used. To obtain data cache categorizations distinguishing hits and misses, we use a scheme that assumes a constant number of data accesses as misses and the remaining references as cache hits. During pipeline simulation, a static branch-prediction scheme using the Ball–Larus heuristic is modeled [Ball and Larus 1993]. Both the static-timing analysis tool and the FAST tool model a simple in-order six-stage pipeline.

When incorporating the frequency model into the static-timing analyzer, two paths with FAST equations that result in intersecting first-order polynomials may be encountered. In this case, we resort to the third method, introduced in Section 3.2, to choose the equation resulting in the worst-case behavior. First, we try to determine if one equation is always greater than the other for the valid range of frequencies (100 MHz-1 GHz). Otherwise, we approximate the two equations by an equation providing a safe upper bound. This may result in slight overestimations but, overall, still provides sufficiently tight bound of the WCEC, as will be seen. We also remove the branch misprediction penalty from the FAST equation if branch misprediction overlaps with a data miss stall. The overestimation caused by instructions with execution latencies higher than one are not removed from the equation as they contribute insignificant savings.

We studied six real-time benchmarks from the C-lab real-time benchmark suite [C-Lab], commonly utilized for WCET experiments. Three floating point benchmarks, *adpcm*, *lms*, and *fft*, as well as three integer benchmarks, *cnt*, *srt*, and *mm*, are analyzed. These benchmarks were compiled by the PISA GCC compiler integrated with our SimpleScalar-based tool set. From the compilation of these benchmarks, the control-flow graphs and instruction layouts were obtained, which are taken as inputs to the FAST analyzer and the static cache analyzer. The FAST output is the WCEC in the form of a parametric equation conforming with our parametric frequency model. The same benchmarks were also exposed to the original static timing analysis tool set for comparison. The original static timing analyzer must be run separately for each frequency under consideration to account for changed memory latency for a given processor frequency. In contrast, the FAST framework captures the same effect in an equation (derived from a single analysis step).

5.2 Testing FAST-DVS Schemes

To test the FAST-DVS schemes, we implemented the algorithms and compiled that into PISA object code to simulate the scheduling overhead, along with each task's execution, within our SimpleScalar-based simulator. Implementation features include generic static voltage-scaling support and scheduling algorithms, ranging from base EDF, cycle-conserving RT-DVS, look-ahead RT-DVS, FAST-static voltage-scaling, FAST-cycle-conserving RT-DVS to FAST-look-ahead RT-DVS. All the scheduling algorithms can choose a frequency between 100 MHz to 1 GHz for the next scheduled task. The base EDF algorithm runs all tasks at 1 GHz. All algorithms switch the processor frequency to 100 MHz during idle times in the schedule, the lowest available frequency, since it is not realistic to put a processor into sleep mode (with millisecond overheads) for frequent task releases (in the order of milliseconds).

A combination of task sets resulting from application workloads of six real-time benchmarks, namely *srt*, *fft*, *mm*, *lms*, *adpcm*, and *cnt*, were studied. The task sets were exposed to the simulator and energy consumption was calculated for all scheduling algorithms. The execution times were derived from exposing the benchmarks to a cycle-accurate pipeline model implemented in our SimpleScalar-based simulator [Anantaraman et al. 2003]. By exploiting a cycle-accurate architectural simulator, we can obtain the total number of cache misses as well as the total number of cycles executed. The execution times obtained from the architectural simulator are scaled with frequency using the same assumption used while formulating the FAST parametric model. Namely, we assume that the total number of execution cycles does not remain constant with frequency. The same execution time scaling method is used for all the voltage-scaling algorithms.

Energy consumption is determined based on the $V^2 f$ and the Wattch models. To evaluate the different FAST-DVS and DVS schemes, we formed several tasksets using the *cnt*, *srt*, *mm*, *adpcm*, *fft*, and *lms* benchmarks. Three groups were formed as follows: G1, *cnt*, *srt*, and *mm* (all integer); G2, *adpcm*, *fft*, and *lms* (all floating point); and G3, *cnt*, *mm*, *fft*, and *lms* (mixed). The periods were chosen for each benchmark and from each group two tasksets are created—one with high utilization and one with low utilization. The high utilization tasksets have a utilization of approximately 0.9, while the low utilization tasksets have a utilization of approximately 0.5.

The frequency/voltage settings used for the scheduling simulator are loosely based on Intel Xscale, which is reported to have five settings ranging from 150 MHz/0.76 V to 1 GHz/1.8 V [Intel 2000]. From the Xscale, we extrapolated 37 settings ranging from 100 MHz/0.70 V to 1 GHz/1.8 V in 25 MHz/0.03 V increments. We calculate energy per cycle at a particular frequency by integrating power over a fixed period of time (e.g., over the hyperperiod) using the relation $Power \sim Voltage^2 \times frequency$.

6. RESULTS FOR FAST FRAMEWORK

The FAST equations for the WCEC for the six benchmarks obtained from the static timing analysis tool and the FAST tool are compiled in Table II and in

Table II. WCEC of FAST vs. Traditional

Benchmarks	Equations		WCET: Static-Timing Analysis/ FAST (WCEC)			
	i	m	100MHZ	400MHZ	700MHZ	1000MHZ
fft	355933	24658	600628/ 602675	1340578/ 1342625	2079876/ 2081993	2820478/ 2822525
adpcm	3026370	544104	8433905/ 8467410	24749525/ 24790530	41065145/ 41113650	57380765/ 57436770
lms	167890	29905	466438/ 466940	1363598/ 1364090	2260748/ 2261240	3157898/ 3158390
cnt	71221	6066	131880/ 131881	313860/ 313861	495840/ 495841	677820/ 677821
mm	2038538	59134	2629877/ 2629878	4403897/ 4403898	6177917/ 6177918	7951937/ 7951938
srt	3509420	102145	4530868/ 4530870	7595218/ 7595220	10659568/ 10659570	13723918/ 13723920

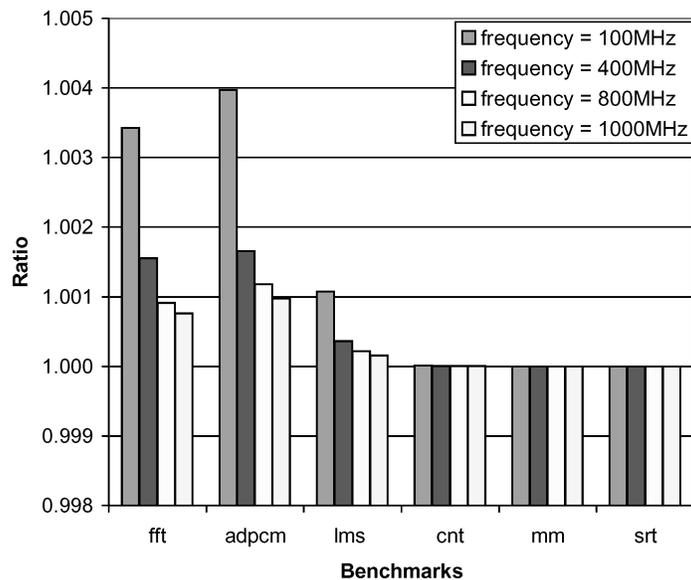


Fig. 11. FAST vs. traditional WCEC.

Figure 11. The FAST scheme differs from conventional static-timing analysis without parametric expressions of frequencies by less than one half a percent. Hence, we conclude that the FAST equations accurately model the WCEC obtained from the static analysis tool. Since the effects of scaling on WCEC are accurately modeled by the FAST equations, the scaling of the WCET can also be accurately captured.

Table II shows the WCEC for all six benchmarks calculated for four different frequencies using the FAST equations and compared with the corresponding WCEC obtained from the static timing analysis tool. Figure 11 plots the ratio of the WCET for the FAST tool and the static timing analysis tool. Table II and Figure 11 show that the FAST bounds on WCET match the bounds obtained by the static timing analyzer exactly for cnt, mm, and srt. For fft, adpcm, and

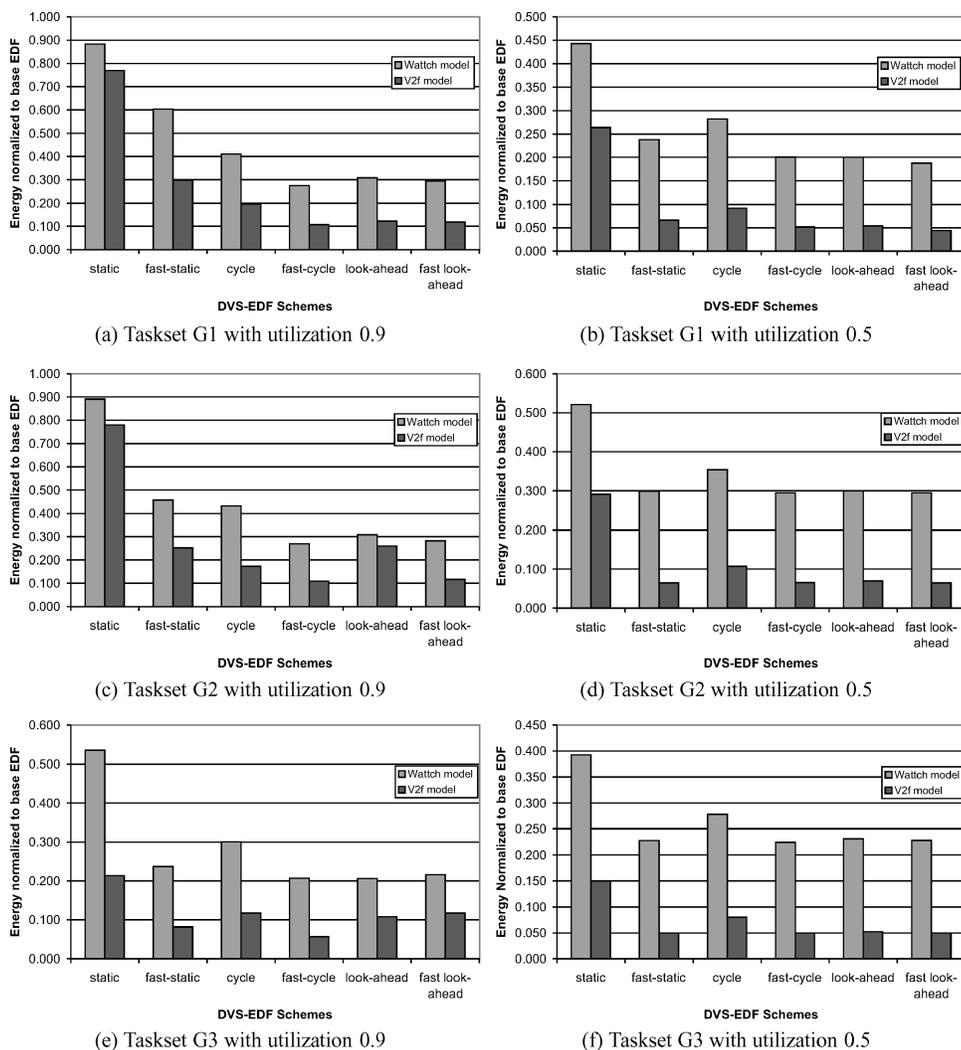


Fig. 12. Energy normalized to base EDF for various task sets.

lms, the FAST bounds on WCET are very close to the bounds obtained by the static timing analyzer. The overestimation in these benchmarks is because of the presence of floating-point operations that have overlapping execution latencies with memory stalls (see Section 2.2, Figure 5). Thus, the FAST tool can accurately model the WCEC of tasks with a negligible error ($< 1\%$) by using our parametric frequency model.

7. RESULTS FOR FAST-DVS SCHEMES

Figures 12a to f depict the energy consumption for both the V^2f and the Watch model of all the DVS schemes normalized to the base EDF scheme for all six tasksets. For each DVS scheme, two bars are presented, the left bar showing

the energy consumption according to the Wattch model and the right bar that of the $V^2 f$ model, each relative to normalized base EDF under the corresponding power model.

The figures show a decrease in energy consumption for all the FAST-DVS schemes when compared to the original RT-DVS schemes. The first, third, and fifth bars in the graphs show the energy consumption for the original RT-DVS schemes. The second, fourth, and sixth bars in the graphs show the improved energy consumption for the FAST-DVS schemes.

For the integer taskset G1, the Wattch model indicates savings of about 30% on energy between static and cycle-conserving RT-DVS and the corresponding FAST variants (Figures 12a and b). For the $V^2 f$ model, savings are even more considerable (in excess of 50%) for these two scheduling schemes. Lower system utilization results in slightly higher energy savings, which can be attributed to exploiting the additional static slack. The look-ahead scheme shows none or only marginal savings under FAST for high and lower utilizations, respectively, regardless of the power model. This is caused by fact that the FAST-look-ahead scheme runs the taskset at a lower frequency and has to recover by raising the frequency more often than the original look-ahead scheme.

The results are also sensitive to the taskset, as a comparison with the floating-point taskset shown by G2. Figures 12c and d indicate that G2 still experiences considerable savings for high utilizations—and slightly lower ones for lower utilizations—under the corresponding FAST scheme. In case of G2, savings for the static and cycle-conserving schemes are even higher than in G1. A comparison between the power models again confirms that the $V^2 f$ model results in higher savings than the Wattch model reports. The results for the integer/floating point mix of G3 in Figures 12e and f show savings at levels between the G1 and G2 tasksets for static and cycle-conserving schemes. The look-ahead version of FAST results in less significant savings, mostly because very aggressive savings are already owing to the original look-ahead scheme.

The differences observed for the $V^2 f$ versus Wattch models indicate that the absolute energy savings obtained by simulation depend on the power model used. Both models show savings relative to base EDF, which validates the FAST approach. However, even relative savings differ by 20%. We believe that the more detailed, architectural Wattch model comes closer to realistically estimating energy savings. The main reason for the inadequacy of the $V^2 f$ model is in its lack of capturing power dependencies following different curves, such as seen in caches and similar architectural structures. In cachelike components, power no longer follows a $V^2 f$ relationship [Zyuban and Kogge 1998]. This explains the lower energy readings for the Wattch model and also indicates that differences between the models depend on the size of caches and similar structures. Hence, the $V^2 f$ model, while suitable as a coarse indicator, may be inaccurate at a more detailed level since it does not take into account overheads of different architectural components.

All results depend on the FAST equation for the benchmarks. The scalability of the WCET depends on the number of misses counted during timing analysis. Because of a worst-case analysis, the number of misses are usually highly exaggerated, especially for data caches. This means that the original schemes

are penalized heavily because of their assumptions about scaling the WCET. Using the FAST equations, the DVS schemes can improve the tightness of the WCET, which is already highly exaggerated, thereby improving energy consumption.

In summary, FAST equations with the RT-DVS schemes are more greedy and result in lower frequencies. The relative energy benefits are highest in the static RT-DVS scheme because it has the most scope for improvement. The cycle-conserving and the look-ahead RT-DVS schemes are dynamic and already scale the frequency aggressively. The addition of the FAST equations to these aggressive schemes enables them to scale the frequency even more aggressively, showing lower energy consumption. However, these dynamic schemes also require higher scheduling overhead with a complexity of $O(n)$, where n denotes the number of tasks. FAST allows simpler, lower complexity DVS schemes, such as the $O(1)$ static RT-DVS variant, to yield results close to their dynamic counterparts. For complex dynamic scheduling schemes, a simpler static scheme in conjunction with FAST may sometimes be the better choice. Overall, benefits for FAST are being observed in all cases.

8. RELATED WORK

Recently, a number of research groups have addressed various issues in the area of predicting the WCET of real-time programs. Conventional methods for static analysis have been extended from unoptimized programs on simple CISC processors to optimized programs on pipelined RISC processors, and from uncached architectures to instruction and data caches [Park 1993; Lim et al. 1994; Healy et al. 1995; Mueller 2000; White et al. 1999; Li et al. 1996]. All these methods obtain discrete values to bound the WCET in a nonparametric fashion.

Vivancos et al. [2001] describe techniques for addressing static timing analysis for variable loop bounds. The so-called parametric timing analysis allows dynamic schedulers to reassess the WCET based on dynamically determined loop bounds during program execution. Chapman et al. [1996] used path expressions to combine a source-oriented parametric approach of WCET analysis with timing annotations, verifying the latter through the former. Bernat and Burns [2000] also proposed using algebraic expressions to represent the WCET of subprograms, where the algebraic expression is parameterized by some of the subprogram's parameters. These approaches differ in that they address fundamental problems in static timing analysis. Our FAST approach, in contrast, aims at isolating execution effects as a function of the processor frequency, a unique, unprecedented approach complementing existing work on static timing analysis.

9. FUTURE WORK

The Fast-Look ahead DVS algorithm in Figure 10 can be improved by considering partial execution of preempted tasks in terms of their instruction (*i_Left*) and memory (*m_Left*) components instead of a more general counter of remaining cycles (*m_Left*). Consider a task k preempted by a release of another

task j in task-release of the algorithm. Currently, the preempted task k is only considered in terms of its $c_left(k)$, not its $i_left(k)$ and $m_left(k)$. Upon calling `defer()`, $i_left(n)$ and $m_left(n)$ be considered only for task n . By not considering the instruction and memory components of task k , a higher frequencies than necessary may be chosen, which is still correct but presents a missed opportunity to further reduce power consumption. As stated in Section 4.3, the s component shown in Figure 10 cannot be directly converted into a FAST equation, since the calculation of c_left is based on i_left and m_left .

To further reduce power, one could normalize the c_left component to the maximum frequency, f_max . By doing so, we assume that the number of misses on the paths taken so far are not exceeding the number of misses on the worst-case paths up to this point, which is valid. Hence, we can calculate

$$c_left(k) = i_left(k) + L * m_left(k)$$

for a memory latency L and the preempted task k upon a task release, i.e., within task-release. This scaled c_left value can then be used in subsequent `defer()` calculations to more tightly bound the required remaining execution time of preempted tasks. Hence, lower frequencies may be chosen so that additional power can be saved.

10. CONCLUSION

In this work, novel techniques for tight and flexible static timing analysis were developed most suitable—but not restricted to—dynamic scheduling schemes. The essence of our approach lies in providing frequency-aware bounds on the WCET through static-timing analysis. Using a frequency-sensitive parametric model, we can capture the effect of combined DFS/DVS on the WCEC and, thus, accurately model the WCET over any frequency range. These techniques are implemented in a FAST tool leveraging prior expertise on static timing analysis. Experiments show the capability of FAST to derive safe upper bounds on the WCET, which are almost as tight (within 1%) as conventional, nonparametric timing analysis. FAST equations can also be used to improve existing DVS scheduling schemes to ensure that the effect of frequency scaling on WCET is considered and that the WCET used is not exaggerated. This is demonstrated by incorporating FAST into three DVS scheduling schemes. Results indicate significant energy savings over the base DVS schedulers because of FAST for two different power models. To the best of our knowledge, this study of DVS effects on timing analysis is unprecedented.

APPENDIX

Modified Look-ahead DVS–EDF

A number of DVS schemes were proposed by Pillai and Shin for scheduling hard real-time systems [Pillai and Shin 2001]. A simple, *static* scaling version uniformly scales the frequency for all tasks based on utilization tests for schedulability, both for rate-monotone and EDF scheduling. *Cycle-conserving* EDF lowers utilization upon task completion temporarily to the proportion of

the actual execution time. *Look-ahead* EDF is an extension to these schemes, that capitalize on early task completion by deferring work for future tasks in favor of scaling the current task. Scaling of the current task occurs based on a modified utilization test that benefits from both idle slots and early task completion. At any completion (both early and on time), the utilization is effectively reduced for the completing task (up until its next release time).

Specifically, upon task completion, $cc_i = c_left_1 = 0$ according to cycle-conserving EDF and look-ahead EDF, respectively. The *defer* calculations of look-ahead EDF then reassesses the utilization based on future and past deadlines for released and completed tasks, respectively.

We modified the look-ahead EDF by setting $c_left_i = C_i$ at task completion instead of assigning a zero value. In addition, we reassess the utilization *strictly* based on the next deadline in the future, regardless of whether tasks are already released. This allows us to look ahead even further in the schedule and, thereby, potentially save additional energy by lowering frequencies more aggressively, and it retains the safety of the schedule by adhering to the EDF utilization test. If the WCET is not fully utilized, then other tasks may still benefit from early completion up to the threshold given by the idle times left in the schedule. This modified look-ahead EDF scheme was implemented in our comparison and is shown to result in up to 34% lower energy consumption than the original scheme. On the average, the modified scheme saves an additional 5–11% of energy for utilizations between 25% and 100%. At high utilizations, our modification occasionally requires between 0.5–8% more energy, which is due to considering an actual time of $cc_i = 0$ in the original scheme up to the next release of a task. Hence, it would be possible to switch between the two schemes based on a utilization threshold as a trigger. Additional savings over the modified scheme because of early completion can only be obtained by considering the density of a schedule at some instance in time, such as given by the maximal schedule utilized in our feedback EDF scheme.

ACKNOWLEDGMENTS

The improvements of Look-ahead DVS–EDF were designed together with Harini Ramaprasad and Sibin Mohan. Comments from the anonymous reviewers helped improve the presentation of derived formulae.

REFERENCES

- ANANTARAMAN, A., SETH, K., PATIL, K., ROTENBERG, E., AND MUELLER, F. 2003. Virtual simple architecture (VISA): Exceeding the complexity limit in safe real-time systems. In *International Symposium on Computer Architecture*. 250–261.
- AYDIN, H., MELHEM, R., MOSSE, D., AND MEJIA-ALVAREZ, P. 2001. Dynamic and aggressive scheduling techniques for power-aware real-time systems. In *IEEE Real-Time Systems Symposium*.
- BALL, T. AND LARUS, J. R. 1993. Branch prediction for free. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*. 300–313.
- BENITEZ, M. E. AND DAVIDSON, J. W. 1988. A portable global optimizer and linker. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*. 329–338.
- BERNAT, G. AND BURNS, A. 2000. An approach to symbolic worst-case execution time analysis. In *25th IFAC Workshop on Real-Time Programming*.
- BROOKS, D., TIWARI, V., AND MARTONOSI, M. 2000. Wattch: A framework for architectural-level power analysis and optimizations. In *Proceedings of the 27th Annual International Symposium*

- on *Computer Architecture*. IEEE Computer Society and ACM SIGARCH, Vancouver, British Columbia, 83–94.
- BURGER, D., AUSTIN, T. M., AND BENNETT, S. 1996. Evaluating future microprocessors: The simple scalar tool set. Technical Report CS-TR-1996-1308, (July) University of Wisconsin, Madison. C-LAB. Wcet benchmarks. Available from <http://www.c-lab.de/home/en/download.html>.
- CHANDRAKASAN, A., SHENG, S., AND BRODERSEN, R. W. April, 1992. Low-power cmos digital design. In *IEEE Journal of Solid-State Circuits*, vol. 27, 473–484.
- CHAPMAN, R., BURNS, A., AND WELLINGS, A. 1996. Combining static worst-case timing analysis and program proof. *Real-Time Systems* 11, 2, 145–171.
- CHETTO, H. AND CHETTO, M. 1989. Some results of the earliest deadline scheduling algorithm. *IEEE Transactions on Software Engineering* 15, 10 (Oct.), 1261–1269.
- CORP., I. Intel StrongARM processors. <http://www.intel.com/design/strong>.
- DUDANI, A., MUELLER, F., AND ZHU, Y. 2002. Energy-conserving feedback edf scheduling for embedded systems with real-time constraints. In *ACM SIGPLAN Joint Conference Languages, Compilers, and Tools for Embedded Systems (LCTES'02) and Software and Compilers for Embedded Systems (SCOPES'02)*. 213–222.
- GRUIAN, F. 2001. Hard real-time scheduling for low energy using stochastic data and dvs processors. In *Proceedings of the International Symposium on Low-Power Electronics and Design ISLPED'01*.
- HEALY, C. A., WHALLEY, D. B., AND HARMON, M. G. 1995. Integrating the timing analysis of pipelining and instruction caching. In *IEEE Real-Time Systems Symposium*. 288–297.
- HEALY, C. A., ARNOLD, R. D., MUELLER, F., WHALLEY, D., AND HARMON, M. G. 1999. Bounding pipeline and instruction cache performance. *IEEE Transactions on Computers* 48, 1 (Jan.), 53–70.
- INTEL. 2000. *Intel XScale Microarchitecture Technical Summary*.
- LI, Y.-T. S., MALIK, S., AND WOLFE, A. 1996. Cache modeling for real-time software: Beyond direct mapped instruction caches. In *IEEE Real-Time Systems Symposium*. 254–263.
- LIM, S.-S., BAE, Y. H., JANG, G. T., RHEE, B.-D., MIN, S. L., PARK, C. Y., SHIN, H., AND KIM, C. S. 1994. An accurate worst case timing analysis for RISC processors. In *IEEE Real-Time Systems Symposium*. 97–108.
- LIU, C. AND LAYLAND, J. 1973. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. of the Association for Computing Machinery* 20, 1 (Jan.), 46–61.
- MOSSE, D., AYDIN, H., CHILDERS, B., AND MELHEM, R. 2000. Compiler-assisted dynamic power-aware scheduling for real-time applications. In *Workshop on Compilers and Operating Systems for Low Power*.
- MUELLER, F. 2000. Timing analysis for instruction caches. *Real-Time Systems* 18, 2/3 (May), 209–239.
- PARK, C. Y. 1993. Predicting program execution times by analyzing static and dynamic program paths. *Real-Time Systems* 5, 1 (Mar.), 31–61.
- PILLAI, P. AND SHIN, K. 2001. Real-time dynamic voltage scaling for low-power embedded operating systems. In *Symposium on Operating Systems Principles*.
- SETH, K., ANANTARAMAN, A., MUELLER, F., AND ROTENBERG, E. 2003. Fast: Frequency-aware static timing analysis. In *IEEE Real-Time Systems Symposium*. 40–51.
- VIVANCOS, E., HEALY, C., MUELLER, F., AND WHALLEY, D. 2001. Parametric timing analysis. In *ACM SIGPLAN Workshop on Language, Compiler, and Tool Support for Embedded Systems*. ACM SIGPLAN Notices, vol. 36, 88–93.
- WEGENER, J. AND MUELLER, F. 2001. A comparison of static analysis and evolutionary testing for the verification of timing constraints. *Real-Time Systems* 21, 3 (Nov.), 241–268.
- WHITE, R. T., MUELLER, F., HEALY, C., WHALLEY, D., AND HARMON, M. G. 1999. Timing analysis for data and wrap-around fill caches. *Real-Time Systems* 17, 2/3 (Nov.), 209–233.
- ZYUBAN, V. AND KOGGE, P. 1998. The energy complexity of register files. In *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED-98)*. ACM Press, New York. 305–310.

Received April 2004; revised November 2004; accepted July 2005