

A Dynamic Voltage Scaling Algorithm for Sporadic Tasks*

Ala' Qadi Steve Goddard
Computer Science & Engineering
University of Nebraska—Lincoln
Lincoln, NE 68588-0115
 {aqadi,goddard}@cse.unl.edu

Shane Farritor
Mechanical Engineering
University of Nebraska - Lincoln
Lincoln, NE 68588-0656
 {sfarritor}@unl.edu

Abstract

Dynamic voltage scaling (DVS) algorithms save energy by scaling down the processor frequency when the processor is not fully loaded. Many algorithms have been proposed for periodic and aperiodic task models but none support the canonical sporadic task model. A DVS algorithm, called DVSSST, is presented that can be used with sporadic tasks in conjunction with preemptive EDF scheduling. The algorithm is proven to guarantee each task meets its deadline while saving the maximum amount of energy possible with processor frequency scaling.

DVSSST was implemented in the μ C/OS-II real-time operating system for embedded systems and its overhead was measured using a stand-alone Rabbit 2000 test board. Though theoretically optimal, the actual power savings realized with DVSSST is a function of the sporadic task set and the processor's DVS support. It is shown that the DVSSST algorithm achieves 83% of the theoretical power savings for a Robotic Highway Safety Marker real-time application. The difference between the theoretical power savings and the actual power savings is due to the limited number of frequency levels the Rabbit 2000 processor supports.

1 Introduction

Many embedded real-time systems consist of a battery operated microprocessor system with a limited battery life. Some of these systems use rechargeable batteries (like cellular phones and robots) while others use dry batteries. In both cases it is very important to maximize the battery life. Dynamic Voltage Scaling (DVS) aims at reducing the power consumption of the system by operating the processor at a lower frequency and thus on a lower voltage.

*Supported, in part, by grants from the National Science Foundation (grant CCR-0208619) and the National Academy of Sciences Transportation Research Board-NCHRP IDEA Program (Project #90).

In CMOS circuits the power consumed by a CMOS gate is proportional to the square of the voltage applied to the circuit, as shown by Equation (1) where C_L is the gate load capacitance (output capacitance), V_{DD} is the supply voltage and f is the clock frequency [29]. The circuit delay t_d is given by Equation (2) where k is a constant depending on the output gate size and the output capacitance and V_T is the threshold voltage [29]. The clock frequency is inversely proportional to the circuit delay; it is expressed using t_d and the logic depth of a critical path as in Equation (3) where L_d is the depth of the critical path [29].

$$P_{CMOS} = C_L V_{DD}^2 f \quad (1)$$

$$t_d = k \frac{V_{DD}}{(V_{DD} - V_T)^2} \quad (2)$$

$$f = \frac{1}{L_d \cdot t_d} \quad (3)$$

It is clear from Equation (1) that reducing the supply voltage will reduce the power consumption. However it also reduces the clock frequency, as shown by Equations (2) and (3), which slows down the processor, meaning that jobs will execute at a slower rate. Thus, the challenge in applying DVS algorithms to real-time systems is to save maximum power while still meeting all temporal requirements of the system.

In recent years significant research has been done in the area of DVS (e.g., [1, 6, 7, 9, 10, 11, 14, 2, 16, 17, 19, 20, 22, 24, 25, 26, 27, 28, 30, 31]). These efforts have resulted in a number of DVS algorithms supporting various task models for embedded and real-time systems. Successful DVS implementations in commercial processors include Intel's Xscale processor [8], Transmeta's Crusoe processor [4] and Rabbit Semiconductors' Rabbit processor [23].

DVS algorithms in [1, 6, 7, 9, 11, 14, 2, 16, 19, 20, 24, 25, 26, 27, 31] support variations of the Liu and Layland periodic task model [15] under RM scheduling or EDF scheduling. Algorithms presented in [16, 17, 22] considered task

models that also support aperiodic requests with soft deadlines or non-periodic tasks with hard deadlines in which job release times were known a priori.

To date, however, no DVS algorithms support the canonical sporadic task model defined by Mok [18] in which tasks have a minimal inter-execution time rather than a fixed period. In this work, a DVS algorithm is presented and evaluated that supports the canonical sporadic task model executed under EDF scheduling. Each task in a sporadic task set $T = \{T_1, T_2, \dots, T_n\}$ has three associated parameters, p , e , and d :

p is the minimum separation period between the release of two consecutive jobs of a task;

e is the worst-case execution time of the job;

d is the relative deadline for the job.

It is assumed in this work that $d = p$ for all tasks. Thus, each task can be described using the tuple (p, e) .

The remainder of this paper is organized as follows. Section 2 describes related work. Section 3 presents our DVS algorithm. Section 4 proves the optimality of the algorithm in terms of schedulability and theoretical power savings. Section 5 presents the implementation and evaluation of the algorithm in a stand-alone environment and in an embedded real-time system. We conclude with a discussion of results in Section 6.

2. Related Work

The algorithms in [1, 19, 25, 26] assume the periodic task model and rely on the principles of intra-task DVS. That is, they adjust the processor voltage level, and hence the processor speed, based on the execution path a task takes and commonly rely on compiler support rather than operating system support to conserve power.

The algorithms in [6, 7, 9, 10, 11, 14, 2, 20, 22, 26, 27, 31] also assume the periodic task model, but rely on an alternative approach to intra-task DVS, called inter-task DVS. In general, inter-task DVS algorithms determine the processor voltage on a task-by-task basis. That is, they adjust the supply voltage at a task level such that idle time is removed from the schedule while guaranteeing that all tasks meet their respective deadlines.

The approach used in this work falls into the category of inter-task DVS. Of the periodic inter-task DVS work identified, the Static Voltage Scaling algorithm developed by Pillai and Shin in [20] is the most closely related to this work. Their Static Voltage Scaling algorithm is an offline algorithm that scales the processor voltage by a factor equal to α where α is the minimum utilization required for the task to remain schedulable under EDF or RM scheduling. This technique is also used in this work to remove deterministic idle time from the schedule, as computed using worst-case execution times (WCET) for each task, but in a slightly different way. The other two on-line algorithms in [20], Cycle

Preserving and Look Ahead, conserve energy by first using Static Voltage Scaling to set the base processor frequency and then further reduce the voltage level when a job executes for less than its WCET or by deferring task execution as much as possible. Our algorithm would give the same result as the Static Voltage Scaling algorithm if all the tasks in the sporadic task set are executed at a periodic rate.

The algorithm presented by Shin and Choi in [26, 27] also sets the initial voltage level using Static Voltage Scaling. They then lower the voltage level further whenever a single task is eligible for execution. Lee et al. [2] developed their DVS algorithms using only two voltage levels and distributing the tasks into two sets, each corresponding to one of the voltage levels: *High* and *Low*. Their work was based on the results of Ishihara and Yasuura [9] who formulated the processor energy optimization problem as a discrete optimization problem that could be solved using linear integer programming techniques. They showed that, in theory, if the processor has a finite number of discrete voltage levels then it is enough to have at most two voltage levels to minimize the energy consumption of the processor.

Kawaguchi et al. [10] presented an approach to schedule a periodic task set by means of task slicing and queues for fixed priority preemptive scheduling, which mainly makes use of the fact that tasks often do not execute with their WCET.

Hong et al. [6, 7] proposed a synthesis technique for variable voltage core based systems containing a set of independent, asynchronous periodic tasks with arbitrary start times (phases) that were scheduled with non-preemptive fixed priority scheduling. Zhang and Chanson present three algorithms in [31] that apply DVS to a periodic task model with non-preemptable sections. This work assumes all tasks are independent and fully preemptive.

The algorithms in [16, 17, 22] consider variations of the periodic task model that support aperiodic requests with soft deadlines or non-periodic tasks with hard deadlines in which job release times are known a priori. Luo and Jha [16] presented an algorithm to schedule periodic tasks, soft aperiodic tasks and hard aperiodic tasks with precedence constraints using task graphs, cyclic scheduling and slack stealing.

Manzak and Chankrabarti [17] used the method of Lagrange multipliers in an iterative way to determine the minimum optimal voltage at which to execute the task set subject to both time and minimum energy constraints. Their method can be applied to a periodic task set or to what they referred to as *Aperiodic Single Tasks*, which is a system of tasks where each task releases a job once in a certain period of time called *Ttotal*. They considered both EDF and RM and presented exact and approximate algorithms to find the minimum voltage, with the exact algorithm having higher complexity. Quan et al. [22] presented an algorithm that can be applied to a periodic task system or to a type of sporadic

task system in which all timing parameters, including task release times, are known a priori.

None of the existing algorithms support the canonical sporadic task model defined by Mok [18] and summarized in Section 1. The next section presents a DVS algorithm for Mok's sporadic task model with preemptive EDF scheduling, which assumes deadlines are equal to periods.

3. The Algorithm

The Dynamic Voltage Scaling for Sporadic Tasks (DVSST) algorithm presented here is classified as an inter-task DVS algorithm for sporadic tasks. That is, it adjusts the processor voltage on a job-by-job basis, where a job represents the release of a sporadic task. Recall from Equations (2) and (3) that the processor frequency is proportional to the voltage level. As with most DVS algorithms, DVSST is defined in terms of processor frequency, rather than voltage levels, since the relationship between the processor frequency and task execution times can be expressed directly.

The DVSST maintains a frequency-scaling factor, α , that represents the percent of the maximum processor frequency. Rather than using the Static Voltage Scaling algorithm of [20] to set the initial frequency level, DVSST starts with a minimum possible frequency-scaling factor, which can be theoretically zero, and scales the processor frequency up and down depending when jobs are released. The scaling factor α is increased by an amount of e_i/p_i when task T_i is first released. Let r_i be the last release time of task T_i . DVSST reduces α at time $r_i + p_i$ by the amount of e_i/p_i if the next job if task T_i was not yet released. When task T_i later releases the job, α is increased by the same amount. The algorithm is explained in detail after we introduce a few definitions.

Definition 1: *The frequency-scaling factor, α , is defined as the ratio between the new processor frequency and the maximum processor frequency:*

$$\alpha = \frac{f_{new}}{f_{max}} \quad (4)$$

Corollary 1: $\alpha \leq 1$.

Proof: The maximum value that we can scale the frequency to is f_{max} . Therefore

$$\alpha \leq \alpha_{max} = \frac{f_{max}}{f_{max}} = 1. \quad \square$$

Definition 2: *The idle-state scaling factor, α_{idle} , is the minimum scaling factor possible when there is no job to execute.*

Theoretically $\alpha_{idle} = 0$, but in many systems α_{idle} must be greater than zero to support platform requirements, or to interact with external devices that trigger the release of a

sporadic task. In this section it is assumed that $\alpha_{idle} = 0$. This assumption is relaxed in the next section when we describe the implementation of the algorithm on a real system.

Definition 3: *The delayed release task set, TD , is a subset of the task set $T = \{T_1, T_2, \dots, T_n\}$ whose last release was at least p_i time units ago. That is, at any time t , $TD = \{T_i | T_i \in T \wedge (t \geq r_i + p_i)\}$ where r_i is the last release time of task T_i or $-\infty$ if task T_i has not yet released a job.*

The DVSST algorithm is shown in Figure 1. Initially, $TD = T$ and the processor frequency is set to α_{idle} . When a task T_i releases a job, the algorithm immediately increases the scaling factor α by an amount equal to $\frac{e_i}{p_i}$ and removes task T_i from the set TD . If any task T_i does not release a job at the end of its minimum separation period p_i , the algorithm reduces the scaling factor α by an amount equal to $\frac{e_i}{p_i}$ and task T_i is added to the set TD . If the algorithm detects that no job is currently executing, then it sets α to the minimum possible value α_{idle} , or in other words, it sets the processor to the idle or sleep mode.

DVSST():
 set $\alpha = \alpha_{idle}$ and $TD = T$ // set initial conditions
 while(true) {
 sleep until ($\exists T_i : (T_i \text{ releases a job and } T_i \in TD)$ or
 ($T_i \notin TD$ and $\text{current_time} \geq r_i + p_i$))
 or (no task is executing)
 if T_i released a job and $T_i \in TD$ then
 // scale up the processor frequency
 set $\alpha = \alpha + \frac{e_i}{p_i}$ and $TD = TD - \{T_i\}$
 else if $T_i \notin TD$ and $\text{current_time} \geq r_i + p_i$
 // scale down the processor frequency
 set $\alpha = \alpha - \frac{e_i}{p_i}$ and $TD = TD + \{T_i\}$
 else // set processor to idle mode
 set $\alpha = \alpha_{idle}$ and $TD = T$ }

Figure 1. The DVSST Algorithm.

The value of α may depend on the previous value of α and since α changes with time, we use α_n to represent the n^{th} change to α at time t . Equation (5) shows how, if at all, α_n is changed at time t .

$$\alpha_n = \begin{cases} \alpha_{idle}, & t = 0 \text{ or no task is executing} \\ \alpha_{n-1} - \frac{e_i}{p_i}, & t \geq r_i + p_i = 0 \text{ and } T_i \notin TD \\ \alpha_{n-1} + \frac{e_i}{p_i}, & T_i \text{ is released at time } t \text{ and } T_i \in TD \\ \text{no change otherwise} & \end{cases} \quad (5)$$

The following example illustrates how the DVSST algorithm scales the processor frequency (voltage) under EDF scheduling.

Example 1: Let us consider the sporadic task set $T_1 = (1,4)$, $T_2 = (1,5)$, $T_3 = (3,10)$. The (un-scaled) system utilization is $U = 0.75$. Let us consider scheduling the jobs that were released in the interval $[0, 20)$ under preemptive

EDF while using the DVSST algorithm to scale the processor frequency (voltage). Assume that the tasks released jobs as follows: T_1 at times 0, 4, 10, and 17; T_2 at times 0, 6, and 11; and T_3 at times 8 and 18. Let job $J_{i,j}$ represent the j^{th} release of task T_i . Figure 2 illustrates the execution of these jobs without DVSST, and Figure 3 illustrates the same jobs executed with DVSST. The specific job attributes for both executions are listed in Table 1.




Figure 2. Executing the example task set under EDF without DVSST.




Figure 3. Executing the example task set under EDF with DVSST. The x-axis represents time and the y-axis represents the frequency scaling factor α , which is set at the start of each execution interval.

Notice that in Figure 2 the processor is idle in the intervals [2,4), [5,6), and [13,17) under EDF scheduling without DVSST. For this set of release times, the DVSST algorithm resulted in an execution in which the processor was never idle during the observed period shown in Figure 3. However, no task missed its deadline—a fact proven in the next section for all feasible task sets.

4. Theoretical Validation

This section addresses the temporal correctness and energy savings possible when sporadic task sets are executed under EDF with DVSST. Section 4.1 presents the temporal correctness and optimality of EDF with DVSST. Section 4.2 quantifies the power savings possible when both the processor voltage and frequency can be scaled, as well as when only the processor frequency can be scaled. It is shown that DVSST is optimal with respect to power savings when only the frequency can be scaled and all tasks execute with their WCET. While this may seem like a strict constraint on the

Job	$r_{i,j}$	$D_{i,j}$	EDF without DVSST			EDF with DVSST		
			Exec Interval	α	% of Job Executed	Execution Interval	α	% of Job Executed
$J_{1,1}$	0	4	[0,1)	1	100%	[0,2.22)	.45	100%
$J_{2,1}$	0	5	[1,2)	1	100%	[2.22,4.44)	.45	100%
$J_{1,2}$	4	8	[4,5)	1	100%	[4.44,5)	.45	24.3%
						[5,6)	.25	25%
						[6,7.13)	.45	50.7%
$J_{2,2}$	6	11	[6,7)	1	100%	[7.13,8)	.45	39.33%
				1		[8,9.21)	.5	60.67%
$J_{3,1}$	8	18	[8,10)	1	66.67%	[9.21,10)	.5	13.13%
			[12,13)	1	33.33%	[12.66,15)	.75	58.34%
						[15,16)	.5	16.66%
						[16,17)	.3	10%
						[17,17.08)	.75	1.87%
$J_{1,3}$	10	14	[10,11)	1	100%	[10,11.33)	.75	100%
$J_{2,3}$	11	16	[11,12)	1	100%	[11.33,12.66)	.75	100%
$J_{1,4}$	17	21	[17,18)	1	100%	[17.08,18.77)	.55	100%
$J_{3,2}$	18	28	[18,21)	1	100%	[18.77,24.22)	.55	100%

Table 1. Job attributes of the example task set when executed under EDF with and without DVSST. The columns labeled $r_{i,j}$ and $D_{i,j}$ represent the release time and absolute deadlines of job $J_{i,j}$. The scaling factor α is set at the start of each execution interval.

processor and task set, Section 5 presents an implementation on a Rabbit 2000 processor for which these constraints hold.

4.1 Temporal Correctness

A voltage (frequency) scaling scheduling algorithm for real-time systems is correct if it guarantees that all jobs meet their deadlines under a specified scheduling algorithm. Example 1, in Section 3, demonstrated that scaling the processor frequency results in new task execution times that are proportional to the frequency-scaling factor. Theorem 1 states that under DVSST, these scaled task execution times results in a scaled processor utilization of one if the sporadic tasks execute at their maximum rate. This theorem provides an intuitive understanding of Theorem 2, which states that (un-scaled) processor utilization less than or equal to one is a necessary and sufficient feasibility condition for sporadic task sets. Before presenting these theorems, however, new definitions are required.

Definition 4: *Scaled-mode execution time*, e_s , is the execution time needed to execute a job under a frequency-scaling factor.

Over any time interval where the scaling factor α is constant, e_s can be calculated by Equation (6) where e_{s_i} is the scaled execution time of task T_i , e_i is the normal execution time of T_i , and α_c is the current scaling factor.

$$e_{si} = \frac{e_i}{\alpha_c} \quad (6)$$

Definition 5: Scaled Mode Utilization, U_s , is the processor utilization while executing at a scaled frequency.

U_s can be calculated over any time interval τ by Equation (7) where α is constant over τ .

$$U_{s\tau} = \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_{si}}{p_i} \quad (7)$$

Definition 6: Scaling Factor Change Interval, τ_{Si} , is the time interval between two consecutive scaling factor changes α_i and α_{i+1} .

Theorem 1: If $\sum_{i=1}^n \frac{e_i}{p_i} \leq 1$ and the DVSST algorithm is used to scale the processor frequency then the processor scaled-mode utilization is always equal to 1 if the tasks are released at their maximum rate.

Proof: Let τ_S be the set of all scaling-factor change intervals between the end of any idle interval and the beginning of the next idle interval where $\tau_S = \{\tau_{S0}, \tau_{S1}, \dots, \tau_{Sm}\}$. If there are no idle intervals, then let τ_S be equal to the length of the hyperperiod of the task set and begin at a hyperperiod boundary. Let $\alpha_S = \{\alpha_{S1}, \alpha_{S2}, \dots, \alpha_{Sm}\}$ be the set of all the scaling factors corresponding to the time intervals in τ_S . From Equation (7), the scaled utilization over any scaling-factor change interval τ_j is

$$U_{s\tau_j} = \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_{si}}{p_i}$$

with $e_{si} = \frac{e_i}{\alpha_{sj}}$ by Equation (6).

Substituting for e_{si} we have

$$U_{s\tau_j} = \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_{si}}{p_i} = \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_i}{p_i \cdot \alpha_{sj}} = \frac{1}{\alpha_{sj}} \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_i}{p_i}$$

but $\alpha_{sj} = \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_i}{p_i}$ since TD contains all the tasks that

did not release a job at its minimum time separation before the start of the interval τ_{Sj} .

Substituting for α_{sj} we have

$$U_{s\tau_j} = \frac{1}{\sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_i}{p_i}} \sum_{\substack{i=1 \\ T_i \notin TD}}^n \frac{e_i}{p_i} = 1.$$

The scaled-mode utilization over the whole time interval τ_S can be calculated as a sum of the products of the utilizations over subintervals times the ratio of the interval to the sum of all intervals, as expressed in the following equation.

$$U_{S\tau} = \sum_{j=1}^m \frac{\tau_{Sj}}{\sum_{k=1}^m \tau_{Sk}} \cdot U_{S\tau_j}$$

since $\sum_{k=1}^m \tau_{Sk}$ is a constant with respect to the

outer summation

$$U_{S\tau} = \frac{1}{\sum_{k=1}^m \tau_{Sk}} \cdot \sum_{j=1}^m \tau_{Sj} \cdot U_{S\tau_j}$$

since $U_{S\tau_j} = 1$

$$U_{S\tau} = \frac{1}{\sum_{k=1}^m \tau_{Sk}} \cdot \sum_{j=1}^m \tau_{Sj} = 1$$

□

From Theorem 1, one might suspect that $U \leq 1$ is a necessary and sufficient feasibility condition for preemptive EDF scheduling with the DVSST algorithm. Theorem 2 states that this is indeed the case, but it cannot be derived directly from Theorem 1. Theorem 1 provides some intuitive insight into why $U \leq 1$ is a necessary and sufficient feasibility condition for preemptive EDF scheduling with the DVSST algorithm, but it does not account for processor demand that spans scaling-factor change intervals. Accounting for this demand is tedious but straightforward. Thus, in the interest of space, the proof of Theorem 2 is contained in [21].

Theorem 2: Let $T = \{T_1, T_2, \dots, T_n\}$ be a sporadic task set with $d_i = p_i$. Preemptive EDF with DVSST will succeed in scheduling T if and only if $\sum_{i=1}^n \frac{e_i}{p_i} \leq 1$.

Proof: See [21]. □

4.2 Power Savings

The amount of power that can be saved depends on whether both frequency and voltage are scaled or frequency alone is scaled. Some processors, such as the Crusoe processor [4], have a feed back loop to scale voltage when the frequency is scaled. Other processors, such as the Rabbit processor [23], can operate on multiple voltage levels but cannot scale the voltage with frequency changes.

Equation (1) shows that power is linearly proportional to the frequency and quadratically proportional to the voltage.

If the processor automatically scales the voltage when the frequency is scaled, then there will be a voltage level corresponding to each frequency level. Let α be the frequency-scaling factor and β be the voltage-scaling factor corresponding to α . From Equation (2) it is clear that the frequency and voltage are related, but the relation between α and β depends on the gate threshold voltage V_T and the voltage itself, V . Equation (8) shows the relation between α and β .

$$\alpha = \frac{(\beta V - V_T)^2}{\beta(V - V_T)^2} \quad (8)$$

Let us compute the power savings of the DVSST algorithm in both cases. First consider the case where only the frequency is scaled, which is the case for the Rabbit processor used in the application described in Section 5. Over any interval in time τ , the normalized power savings will be given by

$$Power\ Savings = \frac{P_{\max} - P_{DVSST}}{P_{\max}} \quad (9)$$

where P_{\max} is the average power consumed by the processor operating at frequency f_{\max} and P_{DVSST} is the average power consumed by the processor operating under the DVSST algorithm. Let τ_S be the set of all scaling-factor change intervals in τ where $\tau_S = \{\tau_0, \tau_1, \dots, \tau_n\}$. Let $\alpha_S = \{\alpha_1, \alpha_2, \dots, \alpha_n\}$ be the set of all scaling factors corresponding to the intervals in τ_S . Since the power over any scaling-factor change interval τ_i is constant, the average power consumed under DVSST in interval τ is calculated as

$$P_{DVSST} = \frac{1}{\tau} \sum_{i=0}^n P_i \cdot \tau_i$$

where P_i is the average power consumed during interval τ_i . Therefore the normalized power savings is

$$Power\ Savings = \frac{P_{\max} - \frac{1}{\tau} \sum_{i=0}^n P_i \cdot \tau_i}{P_{\max}}$$

Recall that from Equation (1) that $P = C_L V_{DD}^2 f$. For P_{\max} , $f = f_{\max}$ and $V = V_{\max}$. For P_{DVSST} , $f = f_i$ because the frequency changes every τ_i and $V = V_{\max}$ because we do not scale the voltage. Therefore the normalized

power savings can be calculated as

$$Power\ Savings = \frac{C f_{\max} V_{\max}^2 - \frac{1}{\tau} \sum_{i=0}^n C f_i V_{\max}^2 \cdot \tau_i}{C f_{\max} V_{\max}^2}$$

but $f_i = \alpha_i f_{\max}$ therefore

$$Power\ Savings = \frac{C f_{\max} V_{\max}^2 - \frac{1}{\tau} \sum_{i=0}^n C \alpha_i f_{\max} V_{\max}^2 \cdot \tau_i}{C f_{\max} V_{\max}^2}$$

$$= \frac{C f_{\max} V_{\max}^2 - \frac{C f_{\max} V_{\max}^2}{\tau} \sum_{i=0}^n \alpha_i \cdot \tau_i}{C f_{\max} V_{\max}^2}$$

$$= 1 - \frac{1}{\tau} \sum_{i=0}^n \alpha_i \cdot \tau_i$$

$$\text{Recall that } \alpha_i = \sum_{\substack{j=1 \\ T_i \notin TD}}^n \frac{e_j}{p_j} = U_{\tau_i}$$

$$Power\ Savings = 1 - \frac{1}{\tau} \cdot \sum_{i=1}^n U_{\tau_i} \cdot \tau_i \\ = 1 - U_{\tau}$$

Now consider the case where both frequency and voltage are scaled. In this case let us keep all the previous assumptions but add $\beta_S = \{\beta_1, \beta_2, \dots, \beta_n\}$ where β_S is the set of voltage-scaling factors corresponding to the intervals in τ_S . The normalized power savings is then computed as follows.

$$Power\ Savings = \frac{P_{\max} - \frac{1}{\tau} \sum_{i=0}^n P_i \cdot \tau_i}{P_{\max}}$$

$$= \frac{C f_{\max} V_{\max}^2 - \frac{1}{\tau} \sum_{i=0}^n C f_i V_i^2 \cdot \tau_i}{C f_{\max} V_{\max}^2}$$

but $f_i = \alpha_i f_{\max}$ and $V_i = \beta_i V_{\max}$

$$= \frac{C f_{\max} V_{\max}^2 - \frac{1}{\tau} \sum_{i=0}^n C \alpha_i \beta_i^2 f_{\max} V_{\max}^2 \cdot \tau_i}{C f_{\max} V_{\max}^2}$$

$$= \frac{C f_{\max} V_{\max}^2 - \frac{C f_{\max} V_{\max}^2}{\tau} \sum_{i=0}^n \alpha_i \beta_i^2 \tau_i}{C f_{\max} V_{\max}^2}$$

$$= 1 - \frac{1}{\tau} \sum_{i=0}^n \alpha_i \beta_i^2 \tau_i$$

We note that in this case the power savings is not equal to $1 - U_{\tau}$ because of the voltage-scaling factor β_i . However, the maximum power that can be saved is still achieved by operating the processor at a frequency equal to the processor utilization.

Theorem 3: If only the frequency can be scaled and the task set is feasibly scheduled, then the processor will save the maximum possible amount of power under DVSSST when all tasks execute with their WCET.

Proof: If the processor only scales the frequency, then the minimum average power consumed in any feasible time interval occurs when the processor is run at a frequency equal to the system utilization over that time interval, assuming WCET is realized. Equation (10) shows the power consumed in this case.

$$P = C \cdot f \cdot V^2 \text{ but } f = \frac{1}{\tau} \sum_{i=1}^n \alpha_i \cdot \tau_i \cdot f_{\max}$$

$$P = \frac{C \cdot V^2 \cdot f_{\max}}{\tau} \sum_{i=1}^n \alpha_i \cdot \tau_i \quad (10)$$

If the processor is running the DVSSST algorithm then average power consumed can be computed using Equation (11).

$$P = \frac{1}{\tau} \sum_{i=1}^n P_i \cdot \tau_i$$

$$P = \frac{1}{\tau} \sum_{i=1}^n C \cdot V^2 \cdot f_i \cdot \tau_i \text{ but } f_i = \alpha_i \cdot f_{\max}$$

$$P = \frac{1}{\tau} \sum_{i=1}^n C \cdot V^2 \cdot \alpha_i \cdot f_{\max} \cdot \tau_i$$

$$P = \frac{C \cdot V^2 \cdot f_{\max}}{\tau} \sum_{i=1}^n \alpha_i \cdot \tau_i \quad (11)$$

Equations (10) and (11) are equal. This proves that if the processor runs the DVSSST algorithm, it will consume the same amount of power as if it was running on a single frequency equal to the task utilization over the whole time interval τ . Thus, DVSSST achieves maximum power savings when only the processor frequency can be dynamically scaled. \square

Theorem 2 states that $U \leq 1$ is a necessary and sufficient condition for schedulability under EDF with DVSSST. Thus, DVSSST does not affect the optimality of EDF scheduling for sporadic task sets. Theorem 3 shows that, in theory, DVSSST is optimal with respect to power savings when only the frequency can be scaled and all tasks execute with their WCET. However, in practice, it is much harder to achieve optimal power savings due to algorithm overhead and limited frequency levels supported by many processors. The next section discusses these implementation issues.

5. Implementation

The DVSSST algorithm was implemented in a modified version of Jean Labrosse's $\mu\text{C}/\text{OS-II}$ (MicroC/OS-II) real

time operating system [12]. The original version of $\mu\text{C}/\text{OS-II}$ uses the RM algorithm to preemptively schedule up to 64 tasks. The modified version used in this study supports EDF scheduling of up to 64K tasks [13]. Algorithm overhead was measured using a stand-alone Rabbit 2000 test board [23]. The actual power savings realized with DVSSST is a function of the sporadic task set and the processor. Rather than create random task sets, we measured the power savings produced by a specific application, the Robotic Highway Safety Marker.

Section 5.1 describes frequency scaling in the Rabbit 2000. Section 5.2 presents slight modifications to the DVSSST algorithm required in practice since currently available embedded processors have a limited number of frequency scaling levels. The overhead created by DVSSST under EDF scheduling on the Rabbit 2000 is reported in Section 5.3. Section 5.4 describes the Robotic Highway Safety Marker and power savings realized for that application.

5.1 Frequency Scaling in the Rabbit 2000

There are two crystal oscillators built into the Rabbit 2000. The main oscillator accepts crystals up to a frequency of 29.4912 MHz and is used to derive the clock for the processor and peripherals. The low power clock oscillator requires a 32.768 kHz crystal, and is used to clock the watchdog timer, a battery backed time/date clock, and a periodic interrupt. The main oscillator can be shut down in a special low-power mode of operation, and the 32.768 kHz oscillator is then used to clock all the things normally clocked by the main oscillator.

The main oscillator can be doubled in frequency and/or divided by 8. If both doubling and dividing are enabled, then there will be a net frequency division by 4. Our model of the Rabbit 2000 has an 18.532 MHz main oscillator. Thus, there are four frequency levels available from the main oscillator: 18.532MHz, 9.266MHz, 4.633MHz and 2.3165MHz—which correspond to 100%, 50%, 25% and 12.5% of the maximum frequency. Since the maximum frequency at which we can operate the processor is 18.532 MHz and the low power mode frequency is 32.768 kHz, the idle-state scaling factor used by DVSSST is $\alpha_{idle} = \frac{32.768kHz}{18.532MHz} = .00176$. In practice, the value of α_{idle} can be close to zero but never zero as assumed in the theoretical presentation of DVSSST.

The Rabbit 2000 processor can operate at different voltages but it does not change the voltage level dynamically when the frequency level is changed. Thus, only the processor frequency will be scaled dynamically, which will result in a linear savings in average power as explained in Section 4.2.

5.2 Modifying DVSST for the Rabbit Processor

There are four non-idle scaling levels available on the Rabbit 2000, rather than the infinite number of levels often assumed in theory. Fortunately, the algorithm can be modified slightly to allow scaling the frequency to a discrete number of levels by rounding the value of α to the next upper scaling level. For example, if we have a processor with scaling levels 0.25, 0.5, 0.75, and 1.0 and the value of α_n at some point in time t as calculated by DVSST is 0.58, then the next upper scaling level to which we set α_n is 0.75.

Another challenge in implementing DVSST on the Rabbit 2000 is that serial communication baud rates cannot be derived from the low-power oscillator. Thus, $\alpha_{idle} = 0.00176$ cannot be used with any application that requires serial communication. Since the wireless transceiver used in the Robotic Highway Safety Marker uses a serial interface to the processor, we use $\alpha_{idle} = \alpha_{min} = 0.125$ so that the application will not lose communication with the other robots.

5.3 Algorithm Overhead

There are two primary sources of overhead created by DVSST: changing frequency levels and detecting when the frequency can be scaled. Changing the processor frequency from one level to another is (approximately) constant, and was measured on the Rabbit 2000 processor to be $120 \mu s$ per frequency change with the main oscillator.

The second source of overhead is largely dependent on how the algorithm detects when it is possible to scale the processor frequency. When a task is released, a check is made to see if the frequency needs to be increased (i.e., if the task $\in TD$). A timer list is used to detect when it is possible to scale down the processor frequency. A timer is set when the task is released and canceled if the task is released again before the timer expires. The processor frequency is scaled down by e_i/p_i whenever a timer expires for task T_i .

The timer list is implemented as a sorted linked list with no effort made to optimize list insertion since most applications that use the Rabbit 2000 have very few tasks; our application has only six tasks and the version of $\mu C/OS-II$ that comes with the board only supports 64 tasks. Thus, insertion into a list of size n has cost $O(n)$. The worst case occurs when an entry needs to be inserted at the end of the list. The list insertion time was measured for up to 512 tasks with random deadlines. For each list length from one to 512, the test was repeated a number of times equal to the list length with random timer values to be inserted. The insertion time was measured for each insertion and the average time of these values for each list length was recorded. The graph shown in Figure 4 plots the average timer list insertion time versus the number of tasks from 20 such experiments. Time is mea-

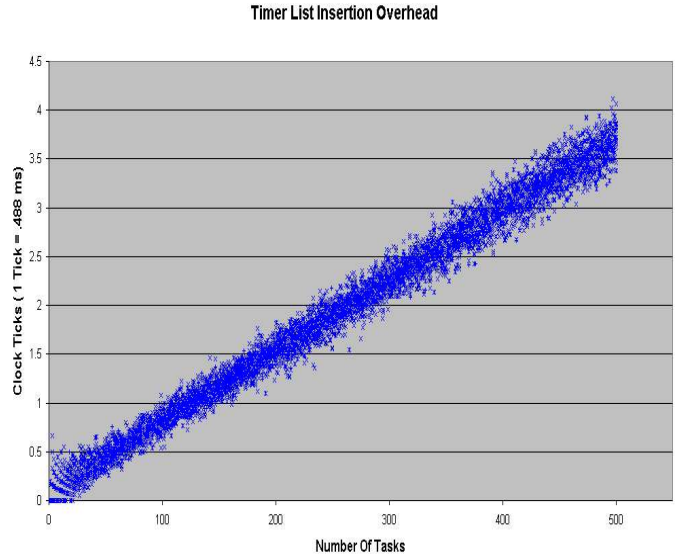


Figure 4. Timer list insertion overhead

sured in terms of periodic clock ticks on the Rabbit 2000, which occur at a rate of 2kHz or one clock tick every $488 \mu s$.

The average insertion time is less than 1 clock tick for a list with less than 125 entries, as shown in Figure 4. The insertion time is about 4 clock ticks (2 ms) for 512 entries. Clearly a more efficient implementation of the timer list should be used for large task sets.

5.4 Power Savings for a Robotic Highway Safety Marker

The Robotic Highway Safety Marker (RSM) is an automated safety device designed to improve road construction work-zone design and safety. A RSM is a semi-autonomous mobile robot that carries a highway safety marker, commonly called a barrel. The RSMs operate in groups that consist of a single lead robot—called the foreman—and worker robots. To date, one foreman and six worker prototype RSMs have been developed. Each worker RSM has a Rabbit 2000 processor running our modified $\mu C/OS-II$. The prototype foreman is more sophisticated than the worker RSMs.

Control of the RSM group is hierarchical and broken into two levels—global and local control—to reduce the per-robot cost. The foreman robot performs global control. To move the robots, the foreman locates each RSM, plans its path, communicates destination points (global waypoints), and monitors performance. Local control is distributed to individual RSMs, which do not have knowledge of other robots and only perform local tasks.

The code for the RSM is implemented as a sporadic task set. The task set only executes after it receives a new waypoint from the foreman. A path from the initial position of the RSM to the new waypoint is computed as a parabola

decomposed into multiple local waypoints. The number of local waypoints depends on the length of the path. The following six sporadic tasks comprise the RSM task set.

- **Serial Task:** reads commands from the foreman via a RF transceiver, converts the command to target destinations, and stores the destinations in a shared queue data structure.
- **Length Task:** calculates the path length, number of iterations, and other values for each target destination.
- **Waypoint Task:** calculates the desired wheel angles for each iteration of a PID control loop.
- **PID Task:** does the PID control for each iteration.
- **Encoder Task:** reads the current wheel angles.
- **Motor Task:** sends commands to each motor.

An abstract processing graph for this task set is shown in Figure 5.4. The precedence relations shown in Figure 5.4 represent the logical precedence constraints on the data processing and do not reflect actual release patterns. For example, to reduce latency in the processing graph, the last four nodes in the processing graph can be released simultaneously with deadline ties broken in favor of producer nodes, as described in [3]. The Serial task is released when data is available on the serial port. When data arrives, the Serial task converts it to a target destination, places it in a shared data structure and releases the Length task. Semaphores are not needed to synchronize access to the data structure, which results in a fully preemptable task set. The Length task calculates the first two local waypoints before the robot begins to move. As the robot moves to waypoint i , waypoint $i+2$ is computed. The design ensures that waypoint $i+2$ is computed before waypoint i is reached.




Figure 5. RSM processing graph.

This task set is modeled as a sporadic task set because the serial task receives commands with a minimum separation of 7.8125ms. The length task is executed the same number of times the serial task is executed. The number of times that Waypoint, PID, Encoder and Motor are executed depends on the number of local waypoints that need to be computed to reach the next global waypoint, which is dependent on the path length. Thus, for each execution of the serial task there may be a different number of executions for the Waypoint,

PID, Encoder and Motor tasks. However, each task has a minimum separation period, as shown in Table 2.

The execution time for these tasks is very deterministic for two reasons. First the Rabbit 2000 has no cache memory, which eliminates memory-caching effects on execution time. Second the tasks repeat almost the same operation each time, with the exception of system initialization where some of the tasks execute a few more lines. Therefore the execution time of these tasks is usually very close to their WCET. The task execution times, shown in Table 2, were determined using an oscilloscope and free I/O pins on the processor.

Task	Period	Execution Time	e_i / p_i
Serial	7.8125ms	100 μ s	.0128
Length	7.8125ms	1ms	.128
Way Point	3 * 7.8125ms	2.5ms	.1066
Encoder	3 * 7.8125ms	350 μ s	.0149
PID	3 * 7.8125ms	1.06ms	.04522
Motor	3 * 7.8125ms	250 μ s	.0106

Table 2. RSM sporadic task set parameters.

The maximum utilization for the task set is $U = 0.31812$, which occurs when all of the tasks execute in a periodic mode for an extended interval of time. If we have no idle periods over an extended interval of time, the lower bound on utilization is when we have only one execution of the serial and length task followed by a very large number of executions of the other tasks. This will result in a processor utilization slightly greater than

$$\frac{e_{Waypoint}}{p_{Waypoint}} + \frac{e_{Encoder}}{p_{Encoder}} + \frac{e_{PID}}{p_{PID}} + \frac{e_{Motor}}{p_{Motor}} = .17732$$

Depending on when commands arrive and the length of the path to be computed, a wide range of utilization values is possible. For any case, the theoretical maximum power savings will be $1 - U_\tau$ (as shown in Section 4.2), where U_τ is the utilization over the time interval τ . The actual power savings achieved is less because we cannot scale the frequency to the desired value; instead we scale it to the nearest upper level of frequency available on the Rabbit 2000, as described in Section 5.2.

As mentioned in Section 5.1, the Rabbit 2000 provides frequency scaling but does not directly adjust the voltage with the frequency. Thus, power savings can be linearly proportional to frequency scaling at best. However, since the Rabbit 2000 provides only a limited number of levels, rather than the unlimited number assumed in theory, there will be a difference between the actual savings and the theoretical power savings.

Figures 6 and 7 show the difference between the actual and the theoretical power savings. The normalized average power savings is plotted against relative utilization values,

where the relative utilization is the ratio of a possible task utilization value to the maximum task utilization (0.31812). Figure 6 shows the normalized theoretical and actual power savings for the task set versus the relative utilization when there are no idle periods. That is, the robot is constantly moving but with destination commands of varying distance. In this case, the minimum relative utilization is 0.55739. Figure 7 shows the normalized theoretical and actual power savings when we have idle periods. That is, when the robot stops for intervals of time.




Figure 6. Power savings with the robot constantly moving.




Figure 7. Power savings with the robot not constantly moving.

Note that the actual power savings deviate from a linear

pattern even though only the processor frequency is scaled and the voltage remains constant. This is because when the frequency is scaled on the Rabbit 2000, it draws less current and the rate at which the current increases or decreases with each frequency level is not exactly linear.

The average ratio of the actual savings to the theoretical savings in both cases is about 83%. This means that DVSSST achieved 83% of the theoretical power savings on the Rabbit 2000 for this application.

If the task set were executed at a periodic rate, the DVSSST would run the processor at a frequency equal to the task utilization, which is the same as the Static Voltage Scaling algorithm of [20], in this case DVSSST will give the same power savings as the Static Voltage Scaling but with more overhead. Other DVS algorithms from the literature are unlikely to improve power savings much, even if the task set executes periodically, because they try to take advantage of the case when tasks do not execute with their WCET. In this application, however, task execution time is very deterministic and there is very little difference between average execution time and WCET.

6. Conclusion

A dynamic voltage-scaling algorithm called DVSSST was presented for sporadic task sets executed under EDF scheduling. It was shown that $U \leq 1$ is a necessary and sufficient schedulability condition for fully preemptive task sets. DVSSST is an inter-task DVS algorithm and the only attempt to save power when jobs execute for less than their WCET is to scale the processor to a minimum frequency level whenever no jobs are pending. DVSSST assumes that resources are not shared between tasks; DVS for resource-sharing sporadic tasks remains an open problem.

DVSSST has been implemented in a modified version of $\mu C/OS-II$ that supports EDF scheduling, and tested with a real-time embedded application, the Robotic Highway Safety Marker (RSM). Though DVSSST is theoretically optimal, results shows that DVSSST saves an average of 83% of the maximum possible theoretical power savings for that application on the Rabbit 2000 processor. Differences between theoretical and actual savings are due to the limited number of frequency levels supported by the Rabbit 2000 processor.

References

- [1] A. Azevedo, I. Issenin, R. Cornea, R. Gupta, N. Dutt. A. Veidenbaum, and A. Nicolau. Profile-Based Dynamic Voltage Scheduling Using Program Checkpoints in the COPPER Framework. *Proceeding of Design, Automation and Test in Europe Conference (DATE)*, March 2002.
- [2] Y. Doh and C. M. Krishna. EDF Scheduling Using Two-Mode Voltage-Clock-Scaling for Hard Real-time Systems. *Proc. of CASES 2001*, pp. 221-228, 2001.

- [3] S. Farritor, M. Rentchler. Robotic Highway Safety Markers. *Proceedings of ASME International Mechanical Engineering Congress*, New Orleans, Louisiana, November 17-22, 2002.
- [4] M. Fleischmann. Crusoe Processor Products and Technology, LongRun Power Management - Dynamic Power Management for Crusoe Processors. http://www.transmeta.com/pdf/white_papers/paper_mfleischmann.17jan01.pdf, Transmeta Inc., January 17, 2001.
- [5] S. Goddard and K. Jeffay. Analyzing the Real-Time Properties of a Data flow Execution Paradigm using a Synthetic Aperture Radar Application. *Proc. 3rd IEEE Real-Time Technology & Applications Symp.*, Montreal, Canada, pp. 60-71, June 1997.
- [6] I. Hong, D. Kirovski, G. Qu, M. Potkonjak, and M. B. Srivastava. Power Optimization of Variable-Voltage Core-Based Systems. *IEEE Trans. Computer-Aided Design*, vol. 18, no. 12, pp. 1702-1714, Dec. 1999.
- [7] I. Hong, G. Qu, M. Potkonjak, and M. B. Srivastava. Synthesis Techniques for Low-Power Hard Real-Time Systems on Variable Voltage Processors. *Proceedings of the IEEE Real-Time Systems Symposium*, pp. 178-187, December 1998.
- [8] Intel XScale microarchitecture, <http://developer.intel.com/design/intelxscale>.
- [9] T. Ishihara and H. Yasuura. Voltage Scheduling Problem for Dynamically Variable Voltage Processors. *Proc. of ISLPED*, pp. 197-202, Aug. 1998.
- [10] H. Kawaguchi, Y. Shin, and T. Sakurai. Experimental Evaluation of Cooperative Voltage Scaling (CVS): A Case Study. *Proceedings of IEEE Workshop on Power Management for Real-Time and Embedded Systems*, pp. 17-23, May 2001.
- [11] W. Kim, J. Kim and S.-L. Min. A Dynamic Voltage Scaling Algorithm for Dynamic-Priority Hard Real-Time Systems Using Slack Time Analysis. *Proceedings of Design Automation and Test in Europe (DATE'02)*, Paris, France, March 2002.
- [12] J. Labrosse. *The Real Time Kernel MicroC/OS-II*, CMP Books, May 2002.
- [13] C-M. Lee, Implementing Rate-Based Execution in MicroC/OS-II. Mater's Project, Dept. of CSE, University of Nebraska-Lincoln, November 27, 2002.
- [14] Y.-H. Lee and C. M. Krishna. Voltage-Clock Scaling for Low Energy Consumption in Real-Time Embedded Systems. *Proceedings of the Sixth Int'l Conf. on Real Time Computing Systems and Applications*, pp. 272-279, 1999.
- [15] C.L.Liu and J. Layland. Scheduling Algorithms for Multiprogramming in a Hard Real-Time Environment. *Journal of the ACM*, Vol.20, pp.46-61, 1973.
- [16] J. Luo and N. K. Jha. Power-conscious Joint Scheduling of Periodic Task Graphs and Aperiodic Tasks in Distributed Real-time Embedded Systems. *Proceedings of ICCAD*, pages 357-364, Nov 2000.
- [17] A. Manzak and C. Chakrabarti. Variable Voltage Task Scheduling for Minimizing Energy or Minimizing Power. *Proceedings IEEE Int. Conf. on Acoustic, Speech, and Signal Processing (ICASSP'00)*, pp. 3239-3242, June 2000.
- [18] A.K.-L. Mok. Fundamental Design Problems of Distributed Systems for the Hard Real Time Environment. Ph.D. Thesis, MIT, Dept. of EE and CS, MIT/LCS/TR-297, May 1983.
- [19] D. Mosse, H. Aydin, B. Childers and R. Melhem, Compiler-Assisted Dynamic Power-Aware Scheduling for Real-Time Applications. *Workshop on Compilers and Operating Systems for Low-Power (COLP'00)*, Philadelphia, PA, Oct. 2000.
- [20] P. Pillai and K. G. Shin. Real-Time Dynamic Voltage Scaling for Low-Power Embedded Operating Systems. *Proc. of the 18th ACM Symp. on Operating Systems Principles*, 2001.
- [21] A. Qadi, S. Goddard, and S. Farritor. DVSST: A Dynamic Voltage Scaling Algorithm for Sporadic Tasks, University of Nebraska - Lincoln, Dept. of CSE, TR-CSE-UNL-2003-2, May 2003. Available via the Web: <http://cse.unl.edu/~goddard/Papers/TR-CSE-UNL-2003-2.pdf>.
- [22] G. Quan and X. Hu. Energy Efficient Fixed-Priority Scheduling for Real-Time Systems on Variable Voltage Processors. *Proceedings of DAC'01: IEEE/ACM Design Automation Conference*, pp. 828-833, June 2001.
- [23] Rabbit Semiconductors. Rabbit 2000 Microprocessor User's Manual, <http://www.rabbitsemiconductor.com/documentation/docs/manuals/Rabbit2000/UsersManual>.
- [24] D. Shin and J. Kim. A Profile-Based Energy-Efficient Intra-Task Voltage Scheduling Algorithm for Hard Real-Time Applications. *Proc. of ISLPED*, 2001.
- [25] D. Shin, J. Kim, and S. Lee. Intra-Task Voltage Scheduling for Low-Energy Hard Real-Time Applications. *IEEE Design and Test Computers*, 18(2): 20-30, 2001.
- [26] Y. Shin and K. Choi. Power Conscious Fixed Priority Scheduling for Hard Real-Time Systems. *Proceedings of the Design Automation Conference*, pp. 134-139, June 1999.
- [27] Y. Shin, K. Choi, and T. Sakurai. Power Optimization of Real-Time Embedded Systems on Variable Speed Processors. *Proceedings of the International Conference on Computer-Aided Design*, pp. 365-368, November 2000.
- [28] M. Weiser, B. Welch, A. Demers, and S. Shenker. Scheduling for Reduced CPU Energy. *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI)*, pp. 13-23, November 1994.
- [29] W. Wolf. *Modern VLSI Design*, Prentice Hall Modern Semiconductor Design Series, Third Edition 2002.
- [30] F. Yao, A. Demers, and S. Shenker. A Scheduling Model for Reduced CPU Energy. *IEEE Symposium on Foundations Computer Science*, pp. 374-382, Oct. 1995.
- [31] F. Zhang and S. T. Chanson, Processor Voltage Scheduling for Real-Time Tasks with Non-Preemptable Sections. *Proceedings of the 23rd IEEE International Real-Time Systems Symposium*, Austin, Texas, pp. 235-245, Dec. 2002.