

On Energy-Optimal Voltage Scheduling for Fixed-Priority Hard Real-Time Systems

HAN-SAEM YUN and JIHONG KIM
Seoul National University

We address the problem of energy-optimal voltage scheduling for fixed-priority hard real-time systems, on which we present a complete treatment both theoretically and practically. Although most practical real-time systems are based on fixed-priority scheduling, there have been few research results known on the energy-optimal fixed-priority scheduling problem. First, we prove that the problem is NP-hard. Then, we present a fully polynomial time approximation scheme (FPTAS) for the problem. For any $\epsilon > 0$, the proposed approximation scheme computes a voltage schedule whose energy consumption is at most $(1 + \epsilon)$ times that of the optimal voltage schedule. Furthermore, the running time of the proposed approximation scheme is bounded by a polynomial function of the number of input jobs and $1/\epsilon$. Given the NP-hardness of the problem, the proposed approximation scheme is practically the best solution because it can compute a near-optimal voltage schedule (i.e., provably arbitrarily close to the optimal schedule) in polynomial time. Experimental results show that the approximation scheme finds more efficient (almost optimal) voltage schedules faster than the best existing heuristic.

Categories and Subject Descriptors: C.3 [Special-Purpose and Application-Based Systems]: Real-Time and Embedded Systems; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*sequencing and scheduling*

General Terms: Algorithms

Additional Key Words and Phrases: Fixed-priority scheduling, real-time systems, approximation algorithms, fully polynomial time approximation scheme, variable voltage processor, dynamic voltage scaling

1. INTRODUCTION

Energy consumption is one of the most important design constraints in designing battery-operated embedded systems such as personal digital assistants, digital cellular phones, and mobile videophones. For such systems, the energy consumption is a critical design factor because the battery operation time is a primary performance measure. The dynamic energy consumption E , which dominates the total energy consumption of CMOS circuits, is given by

This work was supported by grant R01-2001-00360 from the Korea Science and Engineering Foundation. RIACT at Seoul National University provides research facilities for the study.

Authors' address: H.-S. Yun and J. Kim (corresponding author), School of Computer Science and Engineering, Seoul National University, Shilim-dong, Kwanak-ku, Seoul, 151-742, Korea; email: {hsyun, jihong}@davinci.snu.ac.kr.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2003 ACM 1539-9087/03/0800-0393 \$5.00

$E \propto C_L \cdot N_{\text{cycle}} \cdot V_{\text{DD}}^2$, where C_L is the load capacitance, N_{cycle} is the number of executed cycles, and V_{DD} is the supply voltage. Because the dynamic energy consumption E is quadratically dependent on the supply voltage V_{DD} , lowering V_{DD} is an effective technique in reducing the energy consumption. However, lowering the supply voltage also decreases the clock speed, because the circuit delay T_D of CMOS circuits is given by $T_D \propto V_{\text{DD}} / (V_{\text{DD}} - V_T)^\alpha$ [Sakurai and Newton 1990], where V_T is the threshold voltage and α is a technology-dependent constant.

When a given job does not require the maximum performance of a VLSI system, the clock speed (and its corresponding supply voltage) can be dynamically adjusted to the lowest possible level that still satisfies the job's required performance. This is the key principle of the dynamic voltage scaling (DVS) technique. With a recent explosive growth of the portable embedded system market, several commercial variable-voltage processors were developed (e.g., Intel's *Xscale*, AMD's *K6-2+*, and Transmeta's *Crusoe* processors). Targeting these processors, various DVS algorithms [Aydin et al. 2001; Gruian 2001; Hong et al. 1998; Kim et al. 2002; Pillai and Shin 2001; Quan and Hu 2001, 2002; Shin and Choi 1999; Shin et al. 2000; Yao et al. 1995] have been proposed, especially for embedded hard real-time systems.

For hard real-time systems, the goal of voltage scheduling algorithms is to find an *energy-efficient* voltage schedule with all the stringent timing constraints satisfied. A voltage schedule is a function that associates each time unit with a voltage level (i.e., a clock frequency).¹ In this paper, we consider *fixed-priority* real-time jobs running on variable voltage processors.

1.1 Previous Work

Previous investigations on the voltage scheduling problem have focused mainly on real-time jobs running under dynamic-priority scheduling algorithms such as the EDF (earliest-deadline-first) algorithm [Aydin et al. 2001; Hong et al. 1998; Kim et al. 2002; Pillai and Shin 2001]. For example, the problem of energy-optimal EDF scheduling has been well understood. For EDF job sets, the algorithm by Yao et al. [1995] computes the energy-optimal voltage schedules in polynomial time. Although the EDF scheduling policy makes the voltage scheduling problem easier to solve, fixed-priority scheduling algorithms such as the RM (rate monotonic) algorithm are more commonly used in practical real-time systems due to their low overhead and predictability [Liu 2000].

Although there exist several voltage scheduling techniques proposed for fixed-priority real-time tasks (e.g., online scheduling algorithms [Gruian 2001; Pillai and Shin 2001; Shin and Choi 1999] and offline scheduling algorithms [Gruian 2001; Quan and Hu 2001, 2002; Shin et al. 2000]), there have been few research results on the *optimal* voltage scheduling problem for fixed-priority hard real-time systems; neither a polynomial-time optimal voltage scheduling algorithm nor the computational complexity of the problem is known.

¹Throughout the remainder of the paper, we use the term voltage scheduling instead of DVS.

Up to now, the only significant research result on the optimality issue of fixed-priority voltage scheduling is the one presented by Quan and Hu [2002], where energy-optimal voltage schedules for fixed-priority jobs are found by an *exhaustive* algorithm. However, Quan and Hu did not justify their exhaustive approach. If they had presented the computational complexity of the voltage scheduling problem, their result would have been much more significant. Since the worst-case complexity of Quan's algorithm is of higher order than $O(N!)$, where N is the number of jobs, the algorithm is practically unusable for most real-time applications.

Quan and Hu [2001] also proposed a polynomial-time voltage scheduling algorithm for fixed-priority hard real-time systems, which is the best known polynomial-time heuristic for the problem. Although efficient, being a heuristic, this algorithm cannot guarantee the quality of the voltage schedule computed.

1.2 Contributions

In this paper, we give a complete treatment on the optimal voltage scheduling problem for fixed-priority hard real-time systems. As with the work of Quan and Hu [2001, 2002], we assume that the timing parameters of each job is known a priori. Our problem is identical to the one solved by Yao et al. [1995], except that the priority assignment is changed from the dynamic EDF assignment to the fixed assignment. As illustrated by Quan and Hu [2001], the voltage scheduling problem for fixed-priority tasks is more difficult to solve because the preemption relationship among the tasks is much more complex to analyze.

First, we prove that the optimal voltage scheduling problem is NP-hard, which implies that no optimal polynomial-time algorithm is likely to exist. Second, we present a *fully polynomial time approximation scheme* for the problem. A fully polynomial time approximation scheme (FPTAS) is an approximation algorithm that takes any ε (>0) as an additional input and returns a solution whose cost is at most a factor of $(1 + \varepsilon)$ away from the cost of the optimal solution, with the running time bounded by a polynomial both in the size of the input instance and in $1/\varepsilon$ [Woeginger 1999]. Given the NP-hardness of the problem, the proposed approximation scheme is practically the best solution. The proposed approximation scheme computes a near-optimal voltage schedule in polynomial time. By changing ε , the approximation scheme can find a voltage schedule that is provably arbitrarily close to the optimal solution.

The rest of the paper is organized as follows. In Section 2, we formulate the problem and characterize feasible voltage schedules. We describe important properties of an energy-optimal voltage schedule in Section 3, which provide a basis of later proofs. In Section 4, we present the intractability result of the problem including its NP-hardness. The FPTAS for the problem is presented in Section 5. Experimental results are given in Section 6, and we conclude with a summary and directions for future work in Section 7.

2. PROBLEM FORMULATION

We consider a set $\mathcal{J} = \{J_1, J_2, \dots, J_{|\mathcal{J}|}\}$ of priority-ordered jobs with J_1 being the job with the highest priority. A job $J \in \mathcal{J}$ is associated with the following

timing parameters, which are assumed to be known offline:

— r_J : the release time of J .

— d_J : the deadline of J .

— c_J : the number of execution cycles required for J .

We use p_J to denote the priority of the job J . We assume that J has a higher priority than J' if $p_J < p_{J'}$. In the rest of the paper, we use i instead of J_i as a subscript of timing parameters when no confusion arises (e.g., r_i , d_i , and c_i stand for, r_{J_i} , d_{J_i} , and c_{J_i} respectively). Note that our job model can be directly applicable to a periodic real-time system by considering all the task instances within a hyperperiod of periodic tasks.

Since there is a one-to-one correspondence between the processor speed and the supply voltage, we use $S(t)$, the processor speed, to denote the voltage schedule in the rest of the paper. Given a voltage schedule, the job executed at time t can be uniquely determined and is denoted by $job(\mathcal{J}, S, t)$. A voltage schedule $S(t)$ is said to be *feasible* if $S(t)$ gives each job the required number of cycles between its release time and deadline. (An exact characterization of a feasible voltage schedule is given in Section 2.1.)

As with other related work [Quan and Hu 2001, 2002; Yao et al. 1995], we assume that the processor speed can be varied continuously with a negligible overhead both in time and power. Furthermore, we model that the power P , energy consumed per unit time, is a convex function of the processor speed; given a voltage schedule $S(t)$, the power can be written as a function of time by $P(S(t))$. For simplicity, we assume that all the jobs have the same switching activity and that P is dependent only on the processor speed.

The goal of the voltage scheduling problem is, therefore, to find a feasible schedule $S(t)$ that minimizes

$$E(S) = \int_{t_s}^{t_f} P(S(t)) dt \quad (1)$$

where t_s and t_f are the lower and upper limits of release times and deadlines of the jobs in \mathcal{J} , respectively. For the rest of this paper, the energy-optimal voltage schedule of a job set \mathcal{J} is denoted by $S_{\text{opt}}^{\mathcal{J}}$.

2.1 Feasibility Analysis

In this section, we derive a necessary and sufficient condition for a voltage schedule to be feasible, which will provide a basis for the proofs in Section 3. We first introduce some useful notations and definitions.

$W(S, [t_1, t_2])$ is used to denote the number of cycles executed under a voltage schedule $S(t)$ from t_1 to t_2 , that is, $W(S, [t_1, t_2]) = \int_{t_1}^{t_2} S(t) dt$. Among $W(S, [t_1, t_2])$ cycles, $W_i(S, [t_1, t_2])$ denotes the number of cycles between t_1 and t_2 used for executing a set of jobs J_1, J_2, \dots, J_i whose priorities are higher than or equal to p_{J_i} . $R_{\mathcal{J}}$ and $D_{\mathcal{J}}$ represent the sets of release times and deadlines of the jobs in \mathcal{J} , respectively, that is, $R_{\mathcal{J}} = \{r_J | J \in \mathcal{J}\}$ and $D_{\mathcal{J}} = \{d_J | J \in \mathcal{J}\}$. $T_{\mathcal{J}}$ denotes the union of $R_{\mathcal{J}}$ and $D_{\mathcal{J}}$, that is, $T_{\mathcal{J}} = R_{\mathcal{J}} \cup D_{\mathcal{J}}$. Given a job set $\mathcal{J}' \subseteq \mathcal{J}$, $C(\mathcal{J}')$ represents the total workload of jobs in \mathcal{J}' , that is, $C(\mathcal{J}') = \sum_{J \in \mathcal{J}'} c_J$. Furthermore, $\mathbf{I}_{\mathcal{J}'}$

represents the minimum interval that includes the execution intervals of jobs in \mathcal{J}' , that is, $\mathbf{I}_{\mathcal{J}'} = [\min R_{\mathcal{J}'}, \max D_{\mathcal{J}'}]$. $\mathcal{T}^{\mathcal{J}}$ represents the Cartesian product of $[r_{J_i}, d_{J_i}]$, for $1 \leq i \leq |\mathcal{J}|$, that is, $\mathcal{T}^{\mathcal{J}} = [r_{J_1}, d_{J_1}] \times [r_{J_2}, d_{J_2}] \times \cdots \times [r_{J_{|\mathcal{J}|}}, d_{J_{|\mathcal{J}|}}]$. Given voltage schedules $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n$ such that

$$\mathcal{S}_i(t) = 0 \text{ for all } t \notin [\alpha_i, \beta_i] \text{ for all } 1 \leq i \leq n \quad \text{and} \quad \beta_i \leq \alpha_{i+1} \text{ for all } 1 \leq i < n,$$

the concatenation of $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_n$ is

$$\oplus_{i=1}^n \mathcal{S}_i = \mathcal{S}_1 \oplus \mathcal{S}_2 \oplus \cdots \oplus \mathcal{S}_n \stackrel{\text{def}}{=} \sum_{i=1}^n \mathcal{S}_i(t).$$

Since jobs should be released before they can be processed, we assume that a voltage schedule \mathcal{S} always satisfies the constraint that for any $t > 0$, $W(\mathcal{S}, [0, t]) \leq C(\{J | r_J < t\})$.

The condition for a voltage schedule $\mathcal{S}(t)$ to be feasible can be expressed as follows:

Condition I (Feasibility Condition).

There exists a $|\mathcal{J}|$ -tuple $(f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}) \in \mathcal{T}^{\mathcal{J}}$ such that

$$\forall 1 \leq i \leq |\mathcal{J}| \quad \forall r \in \{t | t \in R_{\mathcal{J}} \wedge t < f_{J_i}\}$$

$$W(\mathcal{S}, [r, f_{J_i}]) \geq C(\{J | p_J \leq p_{J_i} \wedge r_J \in [r, f_{J_i}]\}). \quad (2)$$

For a $|\mathcal{J}|$ -tuple $(f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}) \in \mathcal{T}^{\mathcal{J}}$, f_{J_i} can be considered as a modified deadline of J_i , which is equal to or precedes the original deadline d_{J_i} . (The meaning of the $|\mathcal{J}|$ -tuple is further clarified in Section 3.) If $\mathcal{S}(t)$ satisfies Condition I for a given $|\mathcal{J}|$ -tuple $(f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}) \in \mathcal{T}^{\mathcal{J}}$, J_i completes its execution by f_{J_i} for all $1 \leq i \leq |\mathcal{J}|$. Such $|\mathcal{J}|$ -tuples are said to be *valid* with respect to $(\mathcal{J}, \mathcal{S}(t))$. Theorem 2.1 gives a proof for the feasibility condition.

THEOREM 2.1. *Condition I is a necessary and sufficient condition for $\mathcal{S}(t)$ to be feasible.*

PROOF. For the necessary part, suppose that $\mathcal{S}(t)$ is feasible, that is, J_i completes its execution at $f_{J_i} \in (r_{J_i}, d_{J_i}]$ for all $1 \leq i \leq |\mathcal{J}|$. Then, for any $r \in R_{\mathcal{J}}$ such that $r < f_{J_i}$, all the higher priority jobs whose release times are within $[r, f_{J_i}]$ complete their executions by f_{J_i} . So the total amount of work that should be done within $[r, f_{J_i}]$ must be greater than or equal to the sum of workload of the jobs. Thus, we have for all $1 \leq i \leq |\mathcal{J}|$:

$$W(\mathcal{S}, [r, f_{J_i}]) \geq C(\{J | p_J \leq p_{J_i} \wedge r_J \in [r, f_{J_i}]\}).$$

For the sufficient part, assume that Condition I is satisfied for a $|\mathcal{J}|$ -tuple $(f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}})$. By induction on i , we prove that J_i is given its required execution cycles c_{J_i} within $[r_{J_i}, f_{J_i}]$ for all $1 \leq i \leq |\mathcal{J}|$. The base case holds trivially.

For the induction step, assume that the proposition holds for all $k = 1, 2, \dots, i - 1$. Let $r < r_{J_i}$ be the earliest time point in $R_{\mathcal{J}}$ such that no lower priority jobs (i.e., J_k for $k > i$) are executed within $[r, r_{J_i}]$, that is, $W(\mathcal{S}, [r, r_{J_i}]) = W_{i-1}(\mathcal{S}, [r, r_{J_i}])$. If such r does not exist, r is set to r_{J_i} . Then, a

higher priority job J' (i.e., J_l for $l < i$) released before r (i.e., $r_{J'} < r$) must complete its execution before r ; otherwise, since any lower priority jobs cannot be executed within $[r_{J'}, r]$, we have

$$\begin{aligned} W(\mathcal{S}, [r_{J'}, r_{J_i}]) &= W(\mathcal{S}, [r_{J'}, r]) + W(\mathcal{S}, [r, r_{J_i}]) \\ &= W_{i-1}(\mathcal{S}, [r_{J'}, r]) + W_{i-1}(\mathcal{S}, [r, r_{J_i}]) = W_{i-1}(\mathcal{S}, [r_{J'}, r_{J_i}]), \end{aligned}$$

which contradicts the definition of r . Since only higher priority jobs (i.e., J_l for $l < i$) are executed within $[r, r_{J_i}]$, the amount of remaining workload of the higher priority jobs (which are released within $[r, r_{J_i}]$) at time r_{J_i} is $C(\{J_k | 1 \leq k < i \wedge r_{J_k} \in [r, r_{J_i}]\}) - W(\mathcal{S}, [r, r_{J_i}])$. So we have

$$\begin{aligned} W_{i-1}(\mathcal{S}, [r_{J_i}, f_{J_i}]) &\leq C(\{J_k | 1 \leq k < i \wedge r_{J_k} \in [r, r_{J_i}]\}) - W(\mathcal{S}, [r, r_{J_i}]) \\ &\quad + C(\{J_k | 1 \leq k < i \wedge r_{J_k} \in [r_{J_i}, f_{J_i}]\}) \\ &= C(\{J_k | 1 \leq k < i \wedge r_{J_k} \in [r, f_{J_i}]\}) - W(\mathcal{S}, [r, r_{J_i}]). \quad (3) \end{aligned}$$

To complete the induction, we only need to show that $W(\mathcal{S}, [r_{J_i}, f_{J_i}]) - W_{i-1}(\mathcal{S}, [r_{J_i}, f_{J_i}])$ is not smaller than c_{J_i} . (Note that J_i preempts any lower priority jobs.) From (3) and the assumption that Condition I is satisfied, we have

$$\begin{aligned} &W(\mathcal{S}, [r_{J_i}, f_{J_i}]) - W_{i-1}(\mathcal{S}, [r_{J_i}, f_{J_i}]) \\ &\geq W(\mathcal{S}, [r, f_{J_i}]) - C(\{J_k | 1 \leq k < i \wedge r_{J_k} \in [r, f_{J_i}]\}) \quad (\text{from (3.)}) \\ &\geq C(\{J_k | 1 \leq k \leq i \wedge r_{J_k} \in [r, f_{J_i}]\}) \\ &\quad - C(\{J_k | 1 \leq k < i \wedge r_{J_k} \in [r, f_{J_i}]\}) \quad (\text{from (2.)}) \\ &= C(\{J_i\}) = c_{J_i}. \quad \square \end{aligned}$$

A job set \mathcal{J} is said to be an EDF job set if for any $J, J' \in \mathcal{J}$ (where $p_J < p_{J'}$), $d_J \leq d_{J'}$, or $d_{J'} \leq r_J$. When the priority assignment follows the EDF policy, we can prove that Condition I is simplified as follows:

Condition II (EDF Feasibility Condition).

For any $r \in R_{\mathcal{J}}$ and $d \in D_{\mathcal{J}}$ (where $r < d$),
 $W(\mathcal{S}, [r, d]) \geq C(\{J | r_J, d_J \subseteq [r, d]\})$.

LEMMA 2.2. *Given an EDF job set \mathcal{J} , a voltage schedule $S(t)$ of \mathcal{J} is feasible if and only if Condition II is satisfied.*

PROOF. Consider a new job set $\mathcal{J}' = \{J'_1, J'_2, \dots, J'_{|\mathcal{J}|}\}$, where $r_{J'_i} = W(\mathcal{S}, [0, r_{J_i}])$, $d_{J'_i} = W(\mathcal{S}, [0, d_{J_i}])$, $c_{J'_i} = c_{J_i}$, and $p_{J'_i} = p_{J_i}$ for all $1 \leq i \leq |\mathcal{J}|$. Because $W(\mathcal{S}, [0, t])$ is a monotonically increasing function of t , \mathcal{J}' is also an EDF job set (i.e., for any $J'_i, J'_k \in \mathcal{J}'$, where $i < k$, $d_{J'_i} \leq d_{J'_k}$, or $d_{J'_k} \leq r_{J'_i}$). Let $S'(t) = 1$ ($\forall t > 0$) be the voltage schedule of \mathcal{J}' . Then, we can easily verify that the index of the job $job(\mathcal{J}, S, t)$ is the same as that of $job(\mathcal{J}', S', W(\mathcal{S}, [0, t]))$. Therefore, $J_i \in \mathcal{J}$ finishes its execution by its deadline d_{J_i} under $S(t)$ if and only if its corresponding job $J'_i \in \mathcal{J}'$ finishes its execution by $d_{J'_i}$ ($= W(\mathcal{S}, [0, d_{J_i}])$) under S' .

It is well known that all the jobs in an EDF job set meet their deadlines under a constant speed if and only if the utilization ratio for any time interval

is less than or equal to 1 [Liu 2000]. That is, S' is a feasible voltage schedule of \mathcal{J}' if and only if the following is satisfied:

$$\begin{aligned} &\text{For any } r' \in R_{\mathcal{J}'} \text{ and } d' \in D_{\mathcal{J}'} \text{ (where } r' < d'), \\ &C(\{J|J \in \mathcal{J}' \wedge [r_J, d_J] \subseteq [r', d']\}) \leq d' - r'. \end{aligned} \quad (4)$$

Since (4) is equivalent to Condition II, Condition II is a necessary and sufficient condition for $S(t)$ to be a feasible voltage schedule of \mathcal{J} . \square

As shown in Conditions I and II, the complexity of fixed-priority voltage scheduling mainly comes from the inherent exhaustiveness in finding a valid $|\mathcal{J}|$ -tuple. In the EDF scheduling algorithm, it is sufficient for a single $|\mathcal{J}|$ -tuple of the original deadlines to be checked if it satisfies Condition II.

3. SOME PROPERTIES OF OPTIMAL SCHEDULES

In this section, we explain several properties for a feasible voltage schedule to be an energy-optimal schedule. These properties provide a key insight in devising a fast approximation algorithm described in Section 5. The first property, which was proven by Quan and Hu [2001], is that an energy-optimal voltage schedule should be a piecewise-constant function.

The existing optimal voltage scheduling algorithm by Quan and Hu is based on an observation that if a given job set satisfies the requirement of an EDF job set, the optimal voltage schedule can be easily computed by Yao's "peak-power-greedy" algorithm [Yao et al. 1995]. Simply applying Yao's algorithm to a fixed-priority job set may cause some jobs to miss their deadlines. However, if the deadlines of the jobs are appropriately modified before scheduling, Yao's algorithm can yield a feasible optimal schedule as shown in Quan and Hu [2002]. The efficiency of an optimal voltage scheduling algorithm is, therefore, dependent on how efficiently the job set is modified to be an EDF job set. To give a better insight into our approach for solving the voltage scheduling problem, we derive an equivalent result to Quan and Hu [2002] using Conditions I and II.

3.1 Properties on $|\mathcal{J}|$ -Tuples

Given a $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}) \in \mathcal{T}^{\mathcal{J}}$, $\mathcal{J}^{\mathbf{f}}$ represents the job set $\{J'_1, J'_2, \dots, J'_{|\mathcal{J}|}\}$, where $p_{J'_i} = p_{J_i}$, $c_{J'_i} = c_{J_i}$, $r_{J'_i} = r_{J_i}$, and $d_{J'_i} = f_{J_i}$ for all $1 \leq i \leq |\mathcal{J}|$. We say that a $|\mathcal{J}|$ -tuple \mathbf{f} is *EDF ordered* if $\mathcal{J}^{\mathbf{f}}$ follows the EDF priority. Furthermore, $\mathcal{J}^{\mathbf{f}}$ is said to be *EDF-equivalent* to \mathcal{J} . We first establish a link between Conditions I and II.

LEMMA 3.1. *If Condition I is satisfied for a job set \mathcal{J} by a voltage schedule S and an EDF-ordered $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}})$, Condition II is satisfied for a job set $\mathcal{J}^{\mathbf{f}}$ by S .*

PROOF. For any $r \in R_{\mathcal{J}^{\mathbf{f}}}$ and $d \in D_{\mathcal{J}^{\mathbf{f}}}$ ($r < d$), we have

$$\begin{aligned} r &\in \{t | t \in R_{\mathcal{J}} (= R_{\mathcal{J}^{\mathbf{f}}}) \wedge t < d\} \quad \text{and} \\ d &= f_{J_i} \text{ for } \exists f_{J_i} \in D_{\mathcal{J}^{\mathbf{f}}} (= \{f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}\}). \end{aligned}$$

Furthermore, since \mathbf{f} is EDF-ordered, we have

$$\begin{aligned} \forall J'_k \in \mathcal{J}^{\mathbf{f}} \text{ s.t. } r_{J'_k} (= r_{J_k}) \in [r, d (= f_{J_i})), \\ d_{J'_k} = f_{J_k} \leq f_{J_i} = d \quad \text{if } p_{J'_k} \leq p_{J'_i} (= p_{J_i}) \\ d_{J'_k} = f_{J_k} > f_{J_i} = d \quad \text{otherwise.} \end{aligned}$$

Thus, we have for all $J'_k \in \mathcal{J}^{\mathbf{f}}$:

$$p_{J'_k} \leq p_{J'_i} \wedge r_{J'_k} \in [r, d) \Leftrightarrow [r_{J'_k}, d_{J'_k}] \subseteq [r, d]. \quad (5)$$

Finally, by substituting d for f_{J_i} in (2), we have

$$\begin{aligned} W(\mathcal{S}, [r, d]) &\geq C(\{J \in \mathcal{J} \mid p_J \leq p_{J_i} \wedge r_J \in [r, d)\}) \\ &= C(\{J'_k \in \mathcal{J}^{\mathbf{f}} \mid p_{J'_k} \leq p_{J'_i} (= p_{J_i}) \wedge r_{J'_k} (= r_{J_k}) \in [r, d)\}) \\ &= C(\{J' \in \mathcal{J}^{\mathbf{f}} \mid [r_{J'}, d_{J'}] \subseteq [r, d]\}). \quad (\text{from (5).}) \quad \square \end{aligned}$$

LEMMA 3.2. *If Condition II is satisfied for a job set $\mathcal{J}^{\mathbf{f}}$ by a voltage schedule \mathcal{S} where $\mathbf{f} = (f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}})$ is an EDF-ordered $|\mathcal{J}|$ -tuple, Condition I is satisfied for a job set \mathcal{J} by \mathcal{S} .*

PROOF. Let $r \in \{t \mid t \in R_{\mathcal{J}} \wedge t < f_{J_i}\}$. Then, we have

$$r \in R_{\mathcal{J}^{\mathbf{f}}} (= R_{\mathcal{J}}), \quad f_{J_i} \in D_{\mathcal{J}^{\mathbf{f}}} (= \{f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}\}) \text{ and } r < f_{J_i}$$

and substituting f_{J_i} for d in Condition II gives

$$W(\mathcal{S}, [r, f_{J_i}]) \geq C(\{J' \in \mathcal{J}^{\mathbf{f}} \mid [r_{J'}, d_{J'}] \subseteq [r, f_{J_i}]\}).$$

Since \mathbf{f} is EDF-ordered, we have for all $J'_k \in \mathcal{J}^{\mathbf{f}}$ (refer to the proof of Lemma 3.1.):

$$p_{J'_k} \leq p_{J'_i} \wedge r_{J'_k} \in [r, f_{J_i}) \iff [r_{J'_k}, d_{J'_k}] \subseteq [r, f_{J_i}]. \quad (6)$$

Therefore, we have

$$\begin{aligned} W(\mathcal{S}, [r, f_{J_i}]) &\geq C(\{J' \in \mathcal{J}^{\mathbf{f}} \mid [r_{J'}, d_{J'}] \subseteq [r, f_{J_i}]\}) \\ &= C(\{J'_k \in \mathcal{J}^{\mathbf{f}} \mid p_{J'_k} \leq p_{J'_i} (= p_{J_i}) \wedge r_{J'_k} (= r_{J_k}) \in [r, f_{J_i})\}) \\ &= C(\{J \in \mathcal{J} \mid p_J \leq p_{J_i} \wedge r_J \in [r, f_{J_i})\}). \quad \square \end{aligned}$$

From Lemmas 3.1 and 3.2, we can derive the following useful theorem that states how a feasible voltage schedule of a job set can be obtained from its EDF-equivalent job sets.

THEOREM 3.3. *Given a job set \mathcal{J} , let $\mathcal{F}_{\mathcal{J}}$ be the set of all feasible voltage schedules for \mathcal{J} . Then, $\mathcal{F}_{\mathcal{J}} = \cup_{\mathbf{f} \in \mathcal{T}_{\text{EDF}}} \mathcal{F}_{\mathcal{J}^{\mathbf{f}}}$, where \mathcal{T}_{EDF} is the set of all EDF-ordered $|\mathcal{J}|$ -tuples for \mathcal{J} .*

PROOF. To show that $\mathcal{S} \in \mathcal{F}_{\mathcal{J}} \Rightarrow \mathcal{S} \in \cup_{\mathbf{f} \in \mathcal{T}_{\text{EDF}}} \mathcal{F}_{\mathcal{J}^{\mathbf{f}}}$, assume that J_i completes its execution at f_{J_i} ($\leq d_{J_i}$) for all $1 \leq i \leq |\mathcal{J}|$ under $\mathcal{S} \in \mathcal{F}_{\mathcal{J}}$. Let $\mathbf{f} = (f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}})$. Then, $\mathcal{J}^{\mathbf{f}}$ is an EDF job set. If not, we have for some $J'_k, J'_l \in \mathcal{J}^{\mathbf{f}}$ (where $p_{J'_k} < p_{J'_l}$)

$$r_{J'_k} < d_{J'_l} (= f_{J_l}) < d_{J'_k} (= f_{J_k}),$$

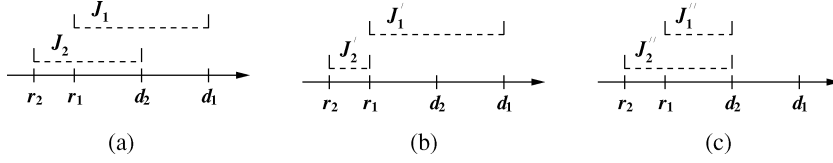


Fig. 1. An example of EDF-equivalent job sets.

which contradicts a fact that once a higher priority job (i.e., J_k) is released during the execution of a lower priority job (i.e., J_l), the higher priority job completes earlier than the lower priority job (i.e., $f_{J_k} < f_{J_l}$). Furthermore, from Lemma 3.1, $S(t)$ is a feasible schedule for the EDF job set \mathcal{J}^f . Thus, we have $S \in \cup_{\mathbf{f} \in \mathcal{T}_{\text{EDF}}} \mathcal{F}_{\mathcal{J}^f}$.

Conversely, given an EDF-ordered $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}})$, let $S \in \mathcal{F}_{\mathcal{J}^f}$ be a feasible schedule for the EDF-equivalent job set \mathcal{J}^f . Then, from Lemma 3.2, S satisfies Condition I for \mathcal{J} . Thus, we have $S \in \mathcal{F}_{\mathcal{J}}$. \square

COROLLARY 3.4. *Given a job set \mathcal{J} , $E(S_{\text{opt}}^{\mathcal{J}}) \leq E(S_{\text{opt}}^{\mathcal{J}^f})$ for any EDF-equivalent job set \mathcal{J}^f . Furthermore, there exists an EDF-equivalent job set \mathcal{J}^f such that $S_{\text{opt}}^{\mathcal{J}} \equiv S_{\text{opt}}^{\mathcal{J}^f}$.*

From Theorem 3.3, there is a one-to-one correspondence between feasible schedules of a fixed-priority job set \mathcal{J} and feasible schedules of \mathcal{J} 's EDF-equivalent job sets. Since the energy-optimal schedule $S_{\text{opt}}^{\mathcal{J}^f}$ for an EDF-equivalent job set \mathcal{J}^f can be directly computed (in polynomial time) by Yao's algorithm [Yao et al. 1995], the problem of finding an energy-optimal (feasible) voltage schedule of \mathcal{J} is reduced to the problem of finding an EDF-equivalent job set \mathcal{J}^f (or to selecting an EDF-ordered $|\mathcal{J}|$ -tuple \mathbf{f}) that minimizes $E(S_{\text{opt}}^{\mathcal{J}^f})$.

Figure 1 shows an example of EDF-equivalent job sets and EDF-ordered $|\mathcal{J}|$ -tuples. Figure 1(a) shows the original job set $\mathcal{J} = \{J_1, J_2\}$. In this example, J_2 has a lower priority but earlier deadline than J_1 , so \mathcal{J} is not an EDF job set. (So Yao's algorithm cannot be directly applied to \mathcal{J} .) In Figures 1(b) and (c), two job sets are shown, which are EDF-equivalent to \mathcal{J} . The job sets $\{J'_1, J'_2\}$ and $\{J''_1, J''_2\}$ are obtained by choosing (r_{J_1}, d_{J_1}) and (d_{J_2}, d_{J_2}) as EDF-ordered $|\mathcal{J}|$ -tuples, respectively. Both job sets follow the EDF priority assignment,² and the optimal voltage schedule for each job set can be computed by Yao's algorithm. (As will be explained below, the energy-optimal voltage schedule of \mathcal{J} is equal to $S_{\text{opt}}^{\{J'_1, J'_2\}}$ or $S_{\text{opt}}^{\{J''_1, J''_2\}}$ depending on the workload of J_1 and J_2 .)

Now, we are to restrict the search space of EDF-ordered $|\mathcal{J}|$ -tuples (equivalently, EDF-equivalent job sets). First, an EDF-ordered $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_1, f_2, \dots, f_{|\mathcal{J}|})$ does not need to be considered if for another EDF-ordered $|\mathcal{J}|$ -tuple $\mathbf{f}' = (f'_1, f'_2, \dots, f'_{|\mathcal{J}|})$ ($\neq \mathbf{f}$), $f_i \leq f'_i$ for all $1 \leq i \leq |\mathcal{J}|$. This is because for any voltage schedule $S(t)$ feasible under \mathbf{f} , $S(t)$ is also feasible under \mathbf{f}' . We define that an EDF-ordered $|\mathcal{J}|$ -tuple \mathbf{f} (or \mathcal{J}^f) is *essential* if such \mathbf{f}' does not exist. (The term “essential” is equivalent to the term “NAP” in Quan and Hu [2002].)

²In Figure 1(c), J''_1 need not have an earlier deadline than J''_2 for the job set to be an EDF job set; $d_{J''_1} = d_{J''_2}$ is sufficient for the job set to be optimally scheduled by Yao's algorithm [Yao et al. 1995].

Quan's optimal algorithm [Quan and Hu 2002] finds an optimal voltage schedule by *exhaustively* enumerating all the essential (or NAP) job sets and then applying Yao's algorithm for each essential job set. Our fast algorithm avoids the exhaustiveness by carefully enumerating the essential job sets.

3.2 $|\mathcal{J}|$ -Permutations

It is easy to check if a $|\mathcal{J}|$ -tuple is EDF-ordered (or essential). On the contrary, it is not obvious how such $|\mathcal{J}|$ -tuples can be enumerated. In this section, we describe how to construct EDF-ordered $|\mathcal{J}|$ -tuples efficiently using a permutation-based analysis.

Given a $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_1, f_2, \dots, f_{|\mathcal{J}|})$, let $\sigma_{\mathbf{f}} : \{1, 2, \dots, |\mathcal{J}|\} \Rightarrow \{1, 2, \dots, |\mathcal{J}|\}$ be a permutation that maps a new tuple index when the tuple elements are sorted in a nondecreasing order, that is, $f_{\sigma_{\mathbf{f}}^{-1}(1)} \leq f_{\sigma_{\mathbf{f}}^{-1}(2)} \leq \dots \leq f_{\sigma_{\mathbf{f}}^{-1}(|\mathcal{J}|)}$. Ties are broken by the priority, that is, if $f_i = f_j$ where $i < j$, $\sigma_{\mathbf{f}}(i) < \sigma_{\mathbf{f}}(j)$. (From now on, we call such σ a $|\mathcal{J}|$ -permutation.) For example, let $\mathbf{f} = (f_1, f_2, f_3, f_4) = (4, 10, 2, 10)$. Then, since $f_3 \leq f_1 \leq f_2 = f_4$, we have $\sigma(3) = 1$, $\sigma(1) = 2$, and (from the tie-breaking rule) $(\sigma(2), \sigma(4)) = (3, 4)$. (Equivalently, we have $(\sigma^{-1}(1), \sigma^{-1}(2), \sigma^{-1}(3), \sigma^{-1}(4)) = (3, 1, 2, 4)$.) Note that $\sigma^{-1}(i)$ denotes the index of the i th smallest element in \mathbf{f} , that is, $f_{\sigma^{-1}(i)}$ is the i th smallest element in \mathbf{f} .

The following lemma states that there cannot exist more than one essential $|\mathcal{J}|$ -tuples whose $|\mathcal{J}|$ -permutations are the same, that is, each essential $|\mathcal{J}|$ -tuple can be uniquely addressed by its corresponding $|\mathcal{J}|$ -permutation (and, obviously, vice versa).

LEMMA 3.5. *For any two essential $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_1, f_2, \dots, f_{|\mathcal{J}|})$ and $\mathbf{f}' = (f'_1, f'_2, \dots, f'_{|\mathcal{J}|})$ ($\mathbf{f} \neq \mathbf{f}'$), $\sigma_{\mathbf{f}} \neq \sigma_{\mathbf{f}'}$.*

PROOF. Suppose $\sigma_{\mathbf{f}} \equiv \sigma_{\mathbf{f}'}$ and let i ($1 \leq i \leq |\mathcal{J}|$) be the largest integer such that $f_{\sigma_{\mathbf{f}}^{-1}(i)} \neq f'_{\sigma_{\mathbf{f}'}^{-1}(i)}$, that is,

$$f_{\sigma_{\mathbf{f}}^{-1}(k)} = f'_{\sigma_{\mathbf{f}'}^{-1}(k)} \quad (= f'_{\sigma_{\mathbf{f}}^{-1}(k)}) \quad \text{for all } i < k < |\mathcal{J}|. \quad (7)$$

Without loss of generality, we can assume $f_{\sigma_{\mathbf{f}}^{-1}(i)} < f'_{\sigma_{\mathbf{f}'}^{-1}(i)}$. Let us consider a new $|\mathcal{J}|$ -tuple $\mathbf{f}'' = (f''_1, f''_2, \dots, f''_{|\mathcal{J}|})$, where

$$f''_k = \begin{cases} f'_k & k = \sigma_{\mathbf{f}}^{-1}(i), \\ f_k & \text{otherwise.} \end{cases}$$

From the definition of \mathbf{f}'' , it can be easily seen that $\sigma_{\mathbf{f}''} \equiv \sigma_{\mathbf{f}} \equiv \sigma_{\mathbf{f}'}$. (We omit the subscripts in the rest of the proof.) We are now to prove that \mathbf{f}'' is EDF-ordered, that is, for any $1 \leq j < k \leq |\mathcal{J}|$,

$$f''_j \leq f''_k \text{ or } f''_k \leq r_{J_j}. \quad (8)$$

Since \mathbf{f} is EDF-ordered, (8) holds for all $1 \leq j < k \leq |\mathcal{J}|$ except for $j = \sigma^{-1}(i)$ or $k = \sigma^{-1}(i)$. So it remains to show that (8) holds for all $1 \leq j < \sigma^{-1}(i) \leq |\mathcal{J}|$ and $1 \leq \sigma^{-1}(i) < k \leq |\mathcal{J}|$.

```

1:    $f_{\sigma^{-1}(|\mathcal{J}|)} := d_{J_{\sigma^{-1}(|\mathcal{J}|)}}$ 
2:   for ( $i := |\mathcal{J}| - 1$  to 1)
3:     let  $\mathcal{J}^H$  be  $\{J_{\sigma^{-1}(k)} \mid i < k \leq |\mathcal{J}| \wedge \sigma^{-1}(k) < \sigma^{-1}(i)\}$ 
4:     if ( $r_{J_{\sigma^{-1}(i)}} \geq \min(\{r_J \mid J \in \mathcal{J}^H\} \cup \{f_{\sigma^{-1}(i+1)}\})$ ) return FALSE
5:     else  $f_{\sigma^{-1}(i)} := \min(\{f_{\sigma^{-1}(i+1)}, d_{J_{\sigma^{-1}(i)}}\} \cup \{r_J \mid J \in \mathcal{J}^H\})$ 
6:     end if
7:   end for

```

Fig. 2. The algorithm to build a $|\mathcal{J}|$ -tuple from a $|\mathcal{J}|$ -permutation.

Case (a). $1 \leq j < \sigma^{-1}(i) \leq |\mathcal{J}|$ (when J_j has a higher priority than $J_{\sigma^{-1}(i)}$.)

If $f_j'' \leq f_{\sigma^{-1}(i)}''$, (8) trivially holds. So we only consider j such that $f_j'' > f_{\sigma^{-1}(i)}''$, that is, $f_j (= f_{\sigma^{-1}(\sigma(j))}) > f_{\sigma^{-1}(i)}' (> f_{\sigma^{-1}(i)})$. From the definition of σ , we have $\sigma(j) > i$. Thus, by substituting $\sigma(j)$ for k in Eq. (7), we have $f_j (= f_j'') = f_j'$. From the assumption, \mathbf{f} is EDF ordered, but we have $f_j' = f_j > f_{\sigma^{-1}(i)}'$. So it must be the case that $f_{\sigma^{-1}(i)}' \leq r_{J_j}$. Therefore, we have

$$f_{\sigma^{-1}(i)}'' = f_{\sigma^{-1}(i)}' \leq r_{J_j}.$$

Case (b). $1 \leq \sigma^{-1}(i) < k \leq |\mathcal{J}|$ (when J_k has a lower priority than $J_{\sigma^{-1}(i)}$.)

First, we can exclude the case when $f_k = f_{\sigma^{-1}(i)}$. Otherwise, we have $\sigma(k) > \sigma(\sigma^{-1}(i)) = i$. (Recall the tie-breaking rule.) But, by the definition of σ , $f_{\sigma^{-1}(\sigma(k))}' (= f_k') \geq f_{\sigma^{-1}(i)}'$ and we finally have

$$f_k' \geq f_{\sigma^{-1}(i)}' > f_{\sigma^{-1}(i)} = f_k,$$

which contradicts Eq. (7).

Second, consider k such that $f_k < f_{\sigma^{-1}(i)}$. \mathbf{f} is EDF-ordered, but we have $f_{\sigma^{-1}(i)} > f_k$. So it must be the case that $f_k \leq r_{J_{\sigma^{-1}(i)}}$. Therefore, we have

$$f_k'' = f_k \leq r_{J_{\sigma^{-1}(i)}}.$$

Finally, for k such that $f_k > f_{\sigma^{-1}(i)}$, we have

$$f_{\sigma^{-1}(i)}'' = f_{\sigma^{-1}(i)}' \leq f_k' = f_k = f_k''.$$

Thus, \mathbf{f}'' is EDF-ordered. However, since we have

$$f_{\sigma^{-1}(i)} < f_{\sigma^{-1}(i)}' = f_{\sigma^{-1}(i)}'' \quad \text{and} \quad f_k = f_k'' \quad \text{for all } 1 \leq k \neq \sigma^{-1}(i) \leq |\mathcal{J}|,$$

\mathbf{f} is not essential, a contradiction. Therefore, $\sigma_{\mathbf{f}} \neq \sigma_{\mathbf{f}''}$. \square

The proof of Lemma 3.5 also implies how to build a unique essential job set for σ .

LEMMA 3.6. *Given a $|\mathcal{J}|$ -permutation σ , the algorithm in Figure 2 finds a unique essential $|\mathcal{J}|$ -tuple for σ if such a $|\mathcal{J}|$ -tuple exists. Otherwise, it returns FALSE.*

PROOF. First, suppose that the essential $|\mathcal{J}|$ -tuple for σ exists and denote it by $\mathbf{f}' = (f_1', f_2', \dots, f_{|\mathcal{J}|}')$. (Note that $f_{\sigma^{-1}(1)}' \leq f_{\sigma^{-1}(2)}' \leq \dots \leq f_{\sigma^{-1}(|\mathcal{J}|)}'$.) We are to prove that $f_{\sigma^{-1}(i)}' = f_{\sigma^{-1}(i)}$, and the algorithm does not abort in line 4 for all $i = |\mathcal{J}|, |\mathcal{J}| - 1, \dots, 1$ by induction on i . The base case holds trivially, that

is, $f'_{\sigma^{-1}(|\mathcal{J}|)} = d_{J_{\sigma^{-1}(|\mathcal{J}|)}} = f_{\sigma^{-1}(|\mathcal{J}|)}$. For the induction step, assume that the proposition holds for all $k = |\mathcal{J}|, |\mathcal{J}| - 1, \dots, i + 1$. Let $\mathcal{J}^H = \{J_{\sigma^{-1}(k)} \mid i < k \leq |\mathcal{J}| \wedge \sigma^{-1}(k) < \sigma^{-1}(i)\}$ (as in line 3 of the algorithm). Note that any job in \mathcal{J}^H has the higher priority than $J_{\sigma^{-1}(i)}$ and that $f'_{\sigma^{-1}(i)} \leq d_{J_{\sigma^{-1}(i)}}$ and $f'_{\sigma^{-1}(i)} \leq f'_{\sigma^{-1}(i+1)}$.

Case (a). $\mathcal{J}^H = \emptyset$.

Suppose that $f'_{\sigma^{-1}(i)} < d_{J_{\sigma^{-1}(i)}}$ and $f'_{\sigma^{-1}(i)} < f'_{\sigma^{-1}(i+1)}$, that is,

$$f'_{\sigma^{-1}(1)} \leq \dots \leq f'_{\sigma^{-1}(i)} < \min\{d_{J_{\sigma^{-1}(i)}}, f'_{\sigma^{-1}(i+1)}\} \leq f'_{\sigma^{-1}(i+1)} \leq \dots \leq f'_{\sigma^{-1}(|\mathcal{J}|)}.$$

Let $\mathbf{f}' = (f'_1, \dots, f'_{\sigma^{-1}(i)-1}, \min\{d_{J_{\sigma^{-1}(i)}}, f'_{\sigma^{-1}(i+1)}\}, f'_{\sigma^{-1}(i+1)}, \dots, f'_{|\mathcal{J}|})$. Then, \mathbf{f}' is EDF ordered, and \mathbf{f}' is not essential, a contradiction. Therefore, we have

$$f'_{\sigma^{-1}(i)} = \min\{d_{J_{\sigma^{-1}(i)}}, f'_{\sigma^{-1}(i+1)}\} = \min\{d_{J_{\sigma^{-1}(i)}}, f_{\sigma^{-1}(i+1)}\} = f_{\sigma^{-1}(i)}.$$

Case (b). $\mathcal{J}^H \neq \emptyset$.

For all $J_{\sigma^{-1}(k)} \in \mathcal{J}^H$, we have $f'_{\sigma^{-1}(i)} < f'_{\sigma^{-1}(k)}$ from the definition of σ (Recall the tie-breaking rule.), and $f'_{\sigma^{-1}(i)} \leq r_{J_{\sigma^{-1}(k)}}$ since \mathbf{f}' is EDF-ordered. Suppose that $f'_{\sigma^{-1}(i)} < \min\{r_{J \mid J \in \mathcal{J}^H}\}$, $f'_{\sigma^{-1}(i)} < d_{J_{\sigma^{-1}(i)}}$, and $f'_{\sigma^{-1}(i)} < f'_{\sigma^{-1}(i+1)}$, that is,

$$\begin{aligned} f'_{\sigma^{-1}(1)} &\leq \dots \leq f'_{\sigma^{-1}(i)} < \min(\{d_{J_{\sigma^{-1}(i)}}, f'_{\sigma^{-1}(i+1)}\} \cup \{r_{J \mid J \in \mathcal{J}^H}\}) \\ &\leq f'_{\sigma^{-1}(i+1)} \leq \dots \leq f'_{\sigma^{-1}(|\mathcal{J}|)}. \end{aligned}$$

Let $\mathbf{f}' = (f'_1, \dots, f'_{\sigma^{-1}(i)-1}, \min(\{d_{J_{\sigma^{-1}(i)}}, f'_{\sigma^{-1}(i+1)}\} \cup \{r_{J \mid J \in \mathcal{J}^H}\}), f'_{\sigma^{-1}(i+1)}, \dots, f'_{|\mathcal{J}|})$. Then, it can be easily shown that \mathbf{f}' is EDF-ordered. Thus, \mathbf{f}' is not essential, a contradiction. Therefore, we have

$$\begin{aligned} f'_{\sigma^{-1}(i)} &= \min(\{d_{J_{\sigma^{-1}(i)}}, f'_{\sigma^{-1}(i+1)}\} \cup \{r_{J \mid J \in \mathcal{J}^H}\}) \\ &= \min(\{d_{J_{\sigma^{-1}(i)}}, f_{\sigma^{-1}(i+1)}\} \cup \{r_{J \mid J \in \mathcal{J}^H}\}) = f_{\sigma^{-1}(i)}. \end{aligned}$$

Furthermore, we have for both cases

$$r_{J_{\sigma^{-1}(i)}} < f'_{\sigma^{-1}(i)} \leq \min(\{r_{J \mid J \in \mathcal{J}^H}\} \cup \{f'_{\sigma^{-1}(i+1)}\}) = \min(\{r_{J \mid J \in \mathcal{J}^H}\} \cup \{f_{\sigma^{-1}(i+1)}\}),$$

and the algorithm does not abort in line 4 at iteration i , which completes the induction.

If the algorithm does not abort, the $|\mathcal{J}|$ -tuple built by the algorithm is always a correct EDF-ordered $|\mathcal{J}|$ -tuple, implying the existence of such $|\mathcal{J}|$ -tuple for σ . Therefore, if such $|\mathcal{J}|$ -tuple does not exist, the algorithm eventually returns FALSE. \square

If a $|\mathcal{J}|$ -permutation σ has the corresponding EDF-ordered $|\mathcal{J}|$ -tuple \mathbf{f} , it is said to be *valid*. Furthermore, if \mathbf{f} is essential, σ is said to be *essential*. From the above argument, we can establish one-to-one correspondences between EDF-ordered $|\mathcal{J}|$ -tuples and valid $|\mathcal{J}|$ -permutations, and between essential $|\mathcal{J}|$ -tuples and essential $|\mathcal{J}|$ -permutations. Figure 3(a) shows a job set with three jobs, and Figures 3(b)–(d) show its EDF equivalent job sets with their $|\mathcal{J}|$ -permutations. Among $3! (= 6)$ possible $|\mathcal{J}|$ -permutations, only three permutations are valid (and essential).

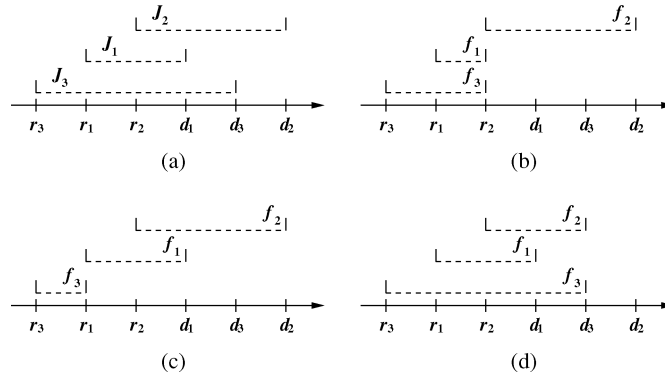


Fig. 3. An example of $|\mathcal{J}|$ -permutations. (a) A job set and its EDF-equivalent job sets for which $(\sigma^{-1}(3), \sigma^{-1}(2), \sigma^{-1}(1)) =$ (b) $(2, 3, 1)$, (c) $(2, 1, 3)$, and (d) $(3, 2, 1)$, respectively. ($(\sigma^{-1}(3), \sigma^{-1}(2), \sigma^{-1}(1)) = (1, 2, 3), (1, 3, 2)$, and $(3, 1, 2)$ are not valid \mathcal{J} -permutations.)

```

1:   $\mathcal{J}' := \{\}, D := \{\}$ 
2:  foreach ( $d_{J_i} \in D_{\mathcal{J}}$  s.t.  $\zeta(d_{J_i}) = 1$ )
3:     $f_i := d_{J_i}$ ,  $\mathcal{J}' := \mathcal{J}' \cup \{J_i\}$ ,  $D := D \cup \{d_{J_i}\}$ 
4:  end foreach /* return FALSE here if  $\mathcal{J}'$  does not follow the EDF priority. */
5:  foreach ( $r_{J_i} \in R_{\mathcal{J}}$  s.t.  $\zeta(r_{J_i}) = 1$  in a decreasing order)
6:     $f_i := \max\{d \in D \mid \mathcal{J}' \cup \{J_i\} \text{ follows the EDF priority where } p_{J_i} = p_{J_i}, r_{J_i} = r_{J_i}, d_{J_i} = d\}$ 
7:    /* return FALSE here if such  $f_i$  does not exist. */
8:     $\mathcal{J}' := \mathcal{J}' \cup \{J_i\}$ ,  $D := D \cup \{r_{J_i}\}$ 
9:  end foreach
10: foreach ( $J_i$  s.t.  $f_i$  is not determined (in any order))
11:    $f_i := \max\{d \in D \mid \mathcal{J}' \cup \{J_i\} \text{ follows the EDF priority where } p_{J_i} = p_{J_i}, r_{J_i} = r_{J_i}, d_{J_i} = d\}$ 
12:   /* return FALSE here if such  $f_i$  does not exist. */
13:    $\mathcal{J}' := \mathcal{J}' \cup \{J_i\}$ 
14: end foreach

```

Fig. 4. The algorithm to build a $|\mathcal{J}|$ -tuple from a bit-vector.

Based on the algorithm in Figure 2, we describe another way to enumerate $|\mathcal{J}|$ -tuples. In the following, r_{J_i} and d_{J_i} are interpreted as symbolic values, not as real numbers. Then, $R_{\mathcal{J}} \cup D_{\mathcal{J}}$ has $2 \cdot |\mathcal{J}|$ distinct symbolic values. Furthermore, the algorithm in Figure 2 is assumed to assign symbolic values to elements of a $|\mathcal{J}|$ -tuple with the following tie-breaking rule in line 5:

(a) $r_{J_i} = r_{J_j} (i < j) : r_{J_i} < r_{J_j}$, (b) $d_{J_i} = d_{J_j} (i < j) : r_{J_i} < r_{J_j}$, (c) $r_{J_i} = d_{J_j} : r_{J_i} < d_{J_j}$.

Given a $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_1, f_2, \dots, f_{|\mathcal{J}|})$, let $\zeta_{\mathbf{f}} : R_{\mathcal{J}} \cup D_{\mathcal{J}} \Rightarrow \{0, 1\}$ be a bit-vector of length $2 \cdot |\mathcal{J}|$ such that

$$\zeta_{\mathbf{f}}(t) = \begin{cases} 1 & t = f_k \text{ for some } 1 \leq k \leq |\mathcal{J}|, \\ 0 & \text{otherwise.} \end{cases}$$

The algorithm in Figure 4 constructs a $|\mathcal{J}|$ -tuple from an arbitrary bit-vector $\zeta : R_{\mathcal{J}} \cup D_{\mathcal{J}} \Rightarrow \{0, 1\}$. The correctness of the algorithm can be proved in a similar manner as the algorithm in Figure 2.

3.3 An Alternative Formulation

The problem formulation given in Section 2 is based on the voltage schedule $S(t)$. In this section, we describe an alternative formulation, based on the following intuitive property, which states that each job runs at the same constant speed if the voltage schedule is an optimal one.

LEMMA 3.7. *For an energy-optimal voltage schedule $S(t)$, $S(t_1) = S(t_2)$ for any t_1 and t_2 such that $job(\mathcal{J}, S, t_1) = job(\mathcal{J}, S, t_2)$.*

PROOF. Given an optimal schedule $S(t)$, suppose that $S(t_1) \neq S(t_2)$ for some t_1 and t_2 such that $job(\mathcal{J}, S, t_1) = job(\mathcal{J}, S, t_2)$. Given that $S(t)$ is optimal, there exist t'_1, t'_2, S_1, S_2 , and Δt such that $S(t) = S_1$ for $t'_1 \leq t \leq t'_1 + \Delta t$, $S(t) = S_2$ for $t'_2 \leq t \leq t'_2 + \Delta t$, and $S_1 \neq S_2$. Let $S(t')$ be defined by

$$S(t') = \begin{cases} \frac{S_1 + S_2}{2} & t'_1 \leq t \leq t'_1 + \Delta t, t'_2 \leq t \leq t'_2 + \Delta t, \\ S(t) & \text{otherwise.} \end{cases}$$

Then, it is obvious that $S(t')$ is feasible, and $E(S') < E(S)$, a contradiction. \square

From Lemma 3.7, it can be shown that the voltage scheduling problem is equivalent to determining the allowed execution time a_i allocated to each J_i . Given a feasible voltage schedule S , the corresponding tuple of the allowed execution times $(a_1, a_2, \dots, a_{|\mathcal{J}|})$, called a *time-allocation tuple*, can be uniquely determined. Conversely, given a time-allocation tuple $\mathbf{A} = (a_1, a_2, \dots, a_{|\mathcal{J}|})$, the corresponding voltage schedule $S_{\mathbf{A}}$ can be uniquely constructed by assigning the constant execution speed c_i/a_i to J_i . \mathbf{A} is said to be *feasible* if the corresponding voltage schedule $S_{\mathbf{A}}$ is feasible.

Let us now consider the exact condition for a time-allocation tuple $\mathbf{A} = (a_1, a_2, \dots, a_{|\mathcal{J}|})$ to be feasible by rewriting Condition I in Section 2 in terms of \mathbf{A} .

Condition III (Feasibility Condition for Time-Allocation Tuples).

There exists a $|\mathcal{J}|$ -tuple $(f_{J_1}, f_{J_2}, \dots, f_{J_{|\mathcal{J}|}}) \in \mathcal{T}^{\mathcal{J}}$ such that

$$\forall 1 \leq i \leq |\mathcal{J}| \quad \forall r \in \{t \mid t \in R_{\mathcal{J}} \wedge t < f_{J_i}\}$$

$$\sum_{J_k / p_{J_k} \leq p_{J_i} \wedge r_{J_k} \in [r, f_{J_i})} a_k \leq f_{J_i} - r. \quad (9)$$

LEMMA 3.8. *Condition III is a necessary and sufficient condition for \mathbf{A} to be feasible.*

PROOF. Given a job set $\mathcal{J} = \{J_1, J_2, \dots, J_{|\mathcal{J}|}\}$ and a time-allocation tuple $\mathbf{A} = (a_1, a_2, \dots, a_{|\mathcal{J}|})$ for \mathcal{J} , consider a new job set $\mathcal{J}' = \{J'_1, J'_2, \dots, J'_{|\mathcal{J}|}\}$, where $c_{J'_i} = a_i$, $r_{J'_i} = r_{J_i}$, $d_{J'_i} = d_{J_i}$, and $p_{J'_i} = p_{J_i}$ for all $1 \leq i \leq |\mathcal{J}|$, that is, \mathcal{J}' is identical to \mathcal{J} except for the workload.

Let $S'(t) = 1$ ($\forall t > 0$) be the voltage schedule of \mathcal{J}' . Then, it is obvious that the response time of J_i under $S_{\mathbf{A}}$ is the same as that of J'_i under S' . Thus, \mathbf{A} is feasible if and only if S' is a feasible voltage schedule for \mathcal{J}' .

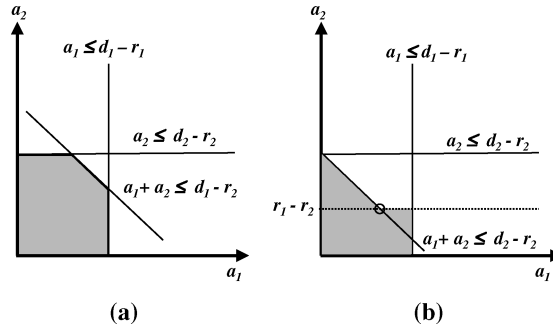


Fig. 5. Solution spaces for (a) an EDF job set and (b) a fixed-priority job set.

After replacing \mathcal{S} and c_{J_i} in Condition I by \mathcal{S}' and a_i , respectively, we have Condition III. \square

By applying the same argument to Condition II, we have the following condition for EDF job sets.

Condition IV (EDF Feasibility Condition for Time-Allocation Tuples).

For any $r \in R_{\mathcal{J}}$ and $d \in D_{\mathcal{J}}$ (where $r < d$),

$$\sum_{J/[r_J, d_J] \subseteq [r, d]} a_i \leq d - r.$$

Now the voltage scheduling problem can be reformulated as follows:

Find a time-allocation tuple $\mathbf{A} = (a_1, a_2, \dots, a_{|\mathcal{J}|})$ such that $E(\mathcal{S}_{\mathbf{A}})$ is minimized subject to Condition III (or Condition IV for an EDF job set).

The energy consumption of the voltage schedule $\mathcal{S}_{\mathbf{A}}$ can be computed directly:

$$E(\mathcal{S}_{\mathbf{A}}) = \sum_{i=1}^{|\mathcal{J}|} a_i \cdot P(c_i/a_i). \quad (10)$$

The set of feasible time-allocation tuples represents the solution space for the voltage scheduling problem stated in terms of time-allocation tuples. For an EDF job set, the solution space is specified by a conjunction of linear inequalities that can be directly obtained from Condition IV. However, this is not the case for a fixed-priority job set; the existential quantifier in Condition III is not always removable. Consequently, the solution space for an EDF job set is a convex set while the solution space for an arbitrary fixed-priority job set may not be a convex set.

Before we present an intractability result for the voltage scheduling problem in the next section, we illustrate the inherent complexity of fixed-priority voltage scheduling based on the results explained in this section. Figures 5(a) and (b) show the solution spaces for an example EDF job set and an example

fixed-priority job set, respectively. As a fixed-priority job set, we use the job set $\{J_1, J_2\}$ of Figure 1. As an EDF job set, we use the same job set $\{J_1, J_2\}$ in Figure 1 with the same timing parameters, but the priority assignment is changed such that it follows the EDF priority assignment, that is, $p_{J_2} < p_{J_1}$. For the EDF job set, we have the following constraint:

$$a_1 \leq d_{J_1} - r_{J_1} \wedge a_2 \leq d_{J_2} - r_{J_2} \wedge a_1 + a_2 \leq d_{J_1} - r_{J_2}.$$

Similarly, we have the following constraint for the fixed-priority job set:

$$\begin{aligned} a_1 \leq d_{J_1} - r_{J_1} \wedge a_2 \leq r_{J_1} - r_{J_2} & \quad (\text{Figure 1(b)}) \vee \\ a_1 \leq d_{J_2} - r_{J_1} \wedge a_1 + a_2 \leq d_{J_2} - r_{J_2} & \quad (\text{Figure 1(c)}). \end{aligned}$$

In Figures 5(a) and (b), the solution spaces for the EDF job set and the fixed-priority job set are depicted as a convex region and a concave region, respectively. (Each point in the shaded regions represents a feasible schedule.) In general, the solution space of any EDF job set with N jobs is represented by a convex set in \mathbf{R}^N , whereas the solution space of a fixed-priority job set is represented by a concave set. Note that for EDF job sets, the objective function, the total energy consumption, can be efficiently minimized by an optimization technique for a convex set (as in Yao's algorithm). However, optimization problems defined on a concave set are generally intractable.

4. INTRACTABILITY RESULT

In this section, we present some observations related to the complexity issue of the optimal fixed-priority scheduling problem. We first show that the decision version of the problem is NP-hard.

THEOREM 4.1. *Given a job set \mathcal{J} and a positive number K , the problem of deciding if there is a feasible voltage schedule $S(t)$ for \mathcal{J} such that $E(S) \leq K$ is NP-hard.*

PROOF. Without loss of generality, we assume that the energy consumption (per CPU cycle) is quadratically dependent on the processor speed. That is, the instantaneous power consumption (per time) is cubically dependent on the processor speed, that is, $P(t) = S(t)^3$. (The reduction can be easily modified for other power functions.) We prove the theorem by reduction from the subset-sum problem, which is NP-complete [Garey and Johnson 1979]:

SUBSET-SUM

INSTANCE: A finite set \mathbf{U} , a size $s : \mathbf{U} \Rightarrow \mathbf{Z}^+$, and a positive integer B .

Question: Is there a subset $\mathbf{U}' \subseteq \mathbf{U}$ such that $\sum_{u \in \mathbf{U}'} s(u) = B$?

Given an instance $\langle \mathbf{U} (= \{u_1, \dots, u_{|\mathbf{U}|}) \rangle, s, B \rangle$ of the subset-sum problem, we construct a job set \mathcal{J} and a positive number K such that there is a voltage schedule $S(t)$ of \mathcal{J} with $E(S) \leq K$ if and only if $\exists \mathbf{U}' \subseteq \mathbf{U}, \sum_{u \in \mathbf{U}'} s(u) = B$. The corresponding job set \mathcal{J} consists of $2 \cdot |\mathbf{U}| + 1$ jobs as follows:

$$\mathcal{J} = \{J_1, J_2, \dots, J_{2 \cdot |\mathbf{U}| + 1}\}$$

where

$$\begin{aligned}
p_{J_i} &= i \quad \text{for all } 1 \leq i \leq 2 \cdot |\mathbf{U}| + 1, \\
r_{J_{2i+1}} &= s(u_{i+1}) + \sum_{j=1}^i 3 \cdot s(u_j), \quad r_{J_{2i+2}} = \sum_{j=1}^i 3 \cdot s(u_j), \\
d_{J_{2i+1}} &= \sum_{j=1}^{i+1} 3 \cdot s(u_j), \quad d_{J_{2i+2}} = 2 \cdot s(u_{i+1}) + \sum_{j=1}^i 3 \cdot s(u_j), \\
c_{J_{2i+1}} &= 8 \cdot \gamma \cdot s(u_{i+1}), \quad c_{J_{2i+2}} = 8 \cdot s(u_{i+1}) \quad \text{for all } 0 \leq i \leq |\mathbf{U}| - 1,
\end{aligned}$$

and

$$r_{J_{2 \cdot |\mathbf{U}| + 1}} = 0, d_{J_{2 \cdot |\mathbf{U}| + 1}} = B + \sum_{j=1}^i 3 \cdot s(u_j), c_{J_{2 \cdot |\mathbf{U}| + 1}} = \sqrt[3]{4} \cdot B,$$

where γ is the unique positive solution of the following quadratic equation:

$$\gamma^2 + \gamma = 1 + \frac{4}{3 \cdot 8^3} \quad \left(\Rightarrow \frac{1}{2} < \gamma < 1 \right).$$

Furthermore, K is set to be

$$K = \left(8^3 + \frac{\gamma^3}{4} \cdot 8^3 \right) \cdot \sum_{i=1}^{|\mathbf{U}|} s(u_i) + 2 \cdot B.$$

From the construction of \mathcal{J} , we have

$$\begin{aligned}
r_{J_{2i+2}} &< r_{J_{2i+1}} (= r_{J_{2i+2}} + s(u_{i+1})) < d_{J_{2i+2}} (= r_{J_{2i+1}} + s(u_{i+1})) \\
&< d_{J_{2i+1}} (= d_{J_{2i+2}} + s(u_{i+1})), \\
[r_{J_{2i+2}}, d_{J_{2i+1}}] &\subset [r_{J_{2 \cdot |\mathbf{U}| + 1}}, d_{J_{2 \cdot |\mathbf{U}| + 1}}] \quad \text{for all } 0 \leq i \leq |\mathbf{U}| - 1
\end{aligned}$$

and

$$[r_{J_{2i+2}}, d_{J_{2i+1}}] \cap [r_{J_{2i'+2}}, d_{J_{2i'+1}}] = \emptyset \quad \text{for all } 0 \leq i \neq i' \leq |\mathbf{U}| - 1.$$

Let $\kappa : \{0, 1\}^{|\mathbf{U}|} \Rightarrow \mathcal{T}^{\mathcal{J}}$ be a function defined by

$$\kappa((b_1, b_2, \dots, b_{|\mathbf{U}|})) = (f_1, f_2, \dots, f_{|\mathcal{J}|})$$

where

$$\begin{aligned}
f_{2i+1} &= d_{J_{2i+1}}, \quad f_{2i+2} = r_{J_{2i+1}} \quad \text{if } b_{i+1} = 0, \\
f_{2i+1} &= f_{2i+2} = d_{J_{2i+2}} \quad \text{if } b_{i+1} = 1 \quad \text{for all } 0 \leq i \leq |\mathbf{U}| - 1,
\end{aligned}$$

and

$$f_{2 \cdot |\mathbf{U}| + 1} = d_{J_{2 \cdot |\mathbf{U}| + 1}}.$$

Then, the set of essential job sets of \mathcal{J} is given by

$$\{\mathcal{J}^{\mathbf{f}} \mid \mathbf{f} = \kappa(\mathbf{b}), \mathbf{b} \in \{0, 1\}^{|\mathbf{U}|}\}.$$

To compute the energy consumption of an essential job set by Yao's algorithm [Yao et al. 1995], we first compare the intensity of each interval. Let

$$I_1 = \frac{c_{J_{2i+2}}}{r_{J_{2i+1}} - r_{J_{2i+2}}}, \quad I_2 = \frac{c_{J_{2i+1}}}{d_{J_{2i+1}} - r_{J_{2i+1}}},$$

$$I_3 = \frac{c_{J_{2i+1}}}{d_{J_{2i+2}} - r_{J_{2i+1}}}, \quad I_4 = \frac{c_{J_{2i+1}} + c_{J_{2i+2}}}{d_{J_{2i+2}} - r_{J_{2i+2}}}$$

and

$$I_5 = \frac{c_{J_{2|\mathbf{U}|+1}}}{B + \delta}.$$

Then, we have

$$I_1 = \frac{8 \cdot s(u_{i+1})}{s(u_{i+1})} = 8 > I_2 = \frac{8 \cdot \gamma \cdot s(u_{i+1})}{2 \cdot s(u_{i+1})} = 4 \cdot \gamma > 2 > \sqrt[3]{4} > I_5 = \frac{\sqrt[3]{4} \cdot B}{B + \delta}$$

and

$$I_4 = \frac{8 \cdot (1 + \gamma) \cdot s(u_{i+1})}{2 \cdot s(u_{i+1})} = 4 + 4 \cdot \gamma > I_3 = \frac{8 \cdot \gamma \cdot s(u_{i+1})}{s(u_{i+1})} = 8 \cdot \gamma > I_5.$$

So the energy consumption of $\mathcal{S}_{\text{opt}}^{\mathcal{J}^f}$ for $\mathbf{f} = \kappa((b_1, b_2, \dots, b_{|\mathbf{U}|}))$ can be computed as follows:

$$E(\mathcal{S}_{\text{opt}}^{\mathcal{J}^f}) = \sum_{i=1}^{|\mathbf{U}|} E_i + E_L$$

where

$$E_i = \begin{cases} \left(8^3 + \frac{\gamma^3}{4} \cdot 8^3\right) \cdot s(u_i) & \left(= \frac{(8 \cdot s(u_i))^3}{s(u_i)^2} + \frac{(8 \cdot \gamma \cdot s(u_i))^3}{(2 \cdot s(u_i))^2}\right) & b_i = 0, \\ \frac{(1 + \gamma)^3}{4} \cdot 8^3 \cdot s(u_i) & \left(= \frac{(8 \cdot (1 + \gamma) \cdot s(u_i))^3}{(2 \cdot s(u_i))^2}\right) & b_i = 1 \end{cases}$$

and

$$E_L = \frac{4 \cdot B^3}{(B + \sum_{i=1}^{|\mathbf{U}|} b_i \cdot s(u_i))^2} \left(= \frac{c_{J_{2|\mathbf{U}|+1}}^3}{(B + \sum_{i=1}^{|\mathbf{U}|} b_i \cdot (d_{J_{2i-1}} - d_{J_{2i}}))^2}\right).$$

Since we have

$$\begin{aligned} \frac{(1 + \gamma)^3}{4} \cdot 8^3 \cdot s(u_i) &= \frac{1 + 3 \cdot \gamma + 3 \cdot \gamma^2 + \gamma^3}{4} \cdot 8^3 \cdot s(u_i) \\ &= \frac{1 + 3 \cdot (1 + 4/(3 \cdot 8^3)) + \gamma^3}{4} \cdot 8^3 \cdot s(u_i) \\ &= \left(8^3 + \frac{\gamma^3}{4} \cdot 8^3\right) \cdot s(u_i) + s(u_i), \end{aligned}$$

we can rewrite $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}^f})$ as follows:

$$E(\mathcal{S}_{\text{opt}}^{\mathcal{J}^f}) = \left(8^3 + \frac{\gamma^3}{4} \cdot 8^3\right) \cdot \sum_{i=1}^{|\mathbf{U}|} s(u_i) + x + \frac{4 \cdot B^3}{(B + x)^2}$$

where

$$x = \sum_{i=1}^{|\mathbf{U}|} b_i \cdot s(u_i).$$

It can be easily shown that $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}^f})$ has the minimum $(8^3 + \frac{\gamma^3}{4} \cdot 8^3) \cdot \sum_{i=1}^{|\mathbf{U}|} s(u_i) + 2B$ ($= K$) at $x = B$. That is, $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}^f}) \leq K$ if and only if

$$\exists (b_1, b_2, \dots, b_{|\mathbf{U}|}) \in \{0, 1\}^{|\mathbf{U}|}, \sum_{i=1}^{|\mathbf{U}|} b_i \cdot s(u_i) = B,$$

which is equivalent to

$$\exists \mathbf{U}' \in \mathbf{U}, \sum_{u \in \mathbf{U}'} s(u) = B.$$

It is obvious that the transformation can be done in polynomial time. Therefore, the problem is NP-hard. \square

From the NP-hardness proof, the problem seems unlikely to have polynomial time algorithms that compute optimal solutions. The NP-hardness of the problem strongly depends on the fact that extremely large input numbers are allowed, as with some other NP-hard problems (e.g., the subset-sum problem and the knapsack problem [Garey and Johnson 1979]). The NP-hardness in the ordinary (but not strong) sense does not rule out possibility of existence of a pseudopolynomial time algorithm or an FPTAS. Since our problem is an optimization problem that handles real numbers, we focus our attention on the FPTAS in the next section.

5. A FAST APPROXIMATION SCHEME

In this section, we present a fully polynomial time approximation scheme (FPTAS) for the problem. We first consider a dynamic programming formulation that always finds the optimal solution, but may run in exponential time. Then, the dynamic programming formulation is transformed into an FPTAS by using a standard technique, the *rounding-the-input-data* technique [Woeginger 1999]. The technique brings the running time of the dynamic program down to polynomial by rounding the input data so that sufficiently close input data are treated by a representative data [Sahni 1976]. The relative error of an approximation scheme depends on how we define the closeness; the smaller the threshold value for the closeness, the smaller the relative error. For a smaller error bound, however, the computation time becomes longer.

5.1 Algorithm for Optimal Solutions

We first present an exponential-time optimal algorithm based on the properties of optimal voltage schedules described in Section 3. The exponential-time algorithm essentially enumerates all the essential job sets. However, unlike Quan's exhaustive algorithm [Quan and Hu 2002], it enumerates the essential job sets intelligently without actually enumerating all of them. Furthermore, it is based

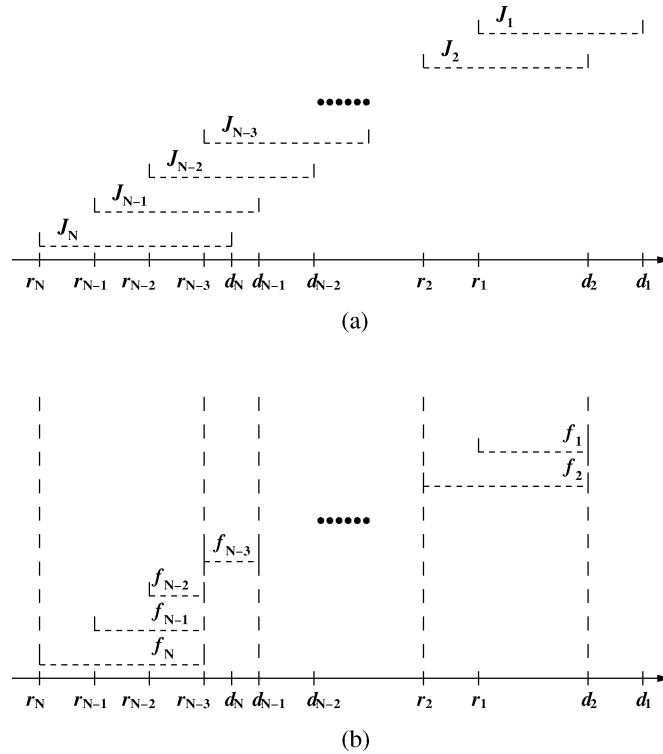


Fig. 6. An example illustrating the optimal algorithm. (a) An original job set; and (b) an essential job set defined by a $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_1, f_2, \dots, f_{N-3}, f_{N-2}, f_{N-1}, f_N) = (d_2, d_2, \dots, d_{N-1}, r_{N-3}, r_{N-3}, r_{N-3})$. Jobs in each subinterval between the thick dashed lines follow the EDF priority assignment and can be optimally scheduled by Yao's algorithm.

on dynamic programming formulation so that it can be easily transformed into an FPTAS by the standard technique.

In formulating the problem by dynamic programming, we first identify appropriate “overlapping” (or reusable) subproblems to which dynamic programming can be applied iteratively. We note that the “optimal substructure” of our problem is naturally reflected by *blocking tuples*, which are just sequences of time points in $\mathcal{T}_{\mathcal{J}}$ in strictly increasing order. (We formally define the blocking tuples later in this section.) That is, the optimal solution of the original problem can be built by just merging the optimal schedules of the subintervals defined by a blocking tuple. Figure 6 shows an example job set and its corresponding EDF-equivalent job set whose time interval is partitioned by a blocking tuple $(r_N, r_{N-3}, d_{N-1}, \dots, r_2, d_2)$, which is depicted by a set of the dashed thick lines in Figure 6(b). Note that jobs in each subinterval follow the EDF-priority assignment.

The original problem is partitioned into subproblems by partitioning the overall time interval into subintervals such that jobs in each subinterval follow the EDF priority assignment. If a job is released within a subinterval with its deadline outside the subinterval, the deadline can be modified to the end of the

subinterval. Each partitioned interval can be optimally scheduled in polynomial time by Yao's algorithm [Yao et al. 1995]. The challenge is how to find the set of subintervals whose optimal subschedules build an energy-optimal voltage schedule.

5.1.1 Basic Idea: The First Example. We now explain the basic idea of the optimal algorithm by describing the optimal algorithm on a simple but illustrative job set $\mathcal{J} = \{J_1, J_2, \dots, J_N\}$ in Figure 6(a), where $r_{i+1} < r_i < d_{i+1} < d_i$ for $1 \leq i < N$. (Note that if the priorities of jobs are reversed, the job set follows the EDF priority.) For this job set, an essential job set \mathcal{J}^e (such as one in Figure 6(b)) is partitioned into $\mathcal{J}_1^e, \mathcal{J}_2^e, \dots, \mathcal{J}_k^e$ such that each \mathcal{J}_i^e ($1 \leq i \leq k$) follows the EDF priority assignment and the union I_i of execution intervals of jobs in \mathcal{J}_i^e (i.e., $I_i = \cup_{J \in \mathcal{J}_i^e} [r_J, d_J]$) does not overlap with I_j ($= \cup_{J' \in \mathcal{J}_j^e} [r_{J'}, d_{J'}]$) for all $1 \leq i \neq j \leq k$. To be more concrete,

$$\text{for all } 1 \leq i < j < k, \quad \forall J \in \mathcal{J}_i^e, J' \in \mathcal{J}_j^e, d_J \leq r_{J'}.$$

Therefore, the optimal voltage schedule $\mathcal{S}_{\text{opt}}^{\mathcal{J}^e}$ of \mathcal{J}^e is equal to the concatenation of the optimal voltage schedules of \mathcal{J}_i^e , that is,

$$\mathcal{S}_{\text{opt}}^{\mathcal{J}^e}(t) \equiv \oplus_{i=1}^k \mathcal{S}_{\text{opt}}^{\mathcal{J}_i^e}(t).$$

Note that $\mathcal{S}_{\text{opt}}^{\mathcal{J}_i^e}$ can be directly computed by Yao's algorithm [Yao et al. 1995] since \mathcal{J}_i^e follows the EDF priority assignment. Therefore, the energy-optimal fixed-priority voltage scheduling problem is further reduced to the problem of finding a partition that gives the energy-optimal voltage schedule for the whole time interval.

In defining a partition, we use a blocking tuple. For example, assume that f_N is selected as r_{N-3} as in Figure 6(b). Then, both f_{N-1} and f_{N-2} should be selected as r_{N-3} , so that the job set becomes EDF-equivalent and, furthermore, essential. As shown in Figure 6(b), these three jobs are separated from the other jobs by a thick vertical line at time r_{N-3} . These jobs constitute the first partitioned job set \mathcal{J}_1^e . The remaining job sets $\mathcal{J}_2^e, \dots, \mathcal{J}_k^e$ can be constructed by applying the same argument. In this way, any essential job set can be partitioned and represented by a blocking tuple.

Let $\mathbf{b} = (b_1, b_2, \dots, b_l) (b_1 < b_2 < \dots < b_l, b_j \in T_{\mathcal{J}})$ be a blocking tuple where

$$\forall 1 \leq j < l, \exists J_i \text{ s.t. } b_j = r_i \wedge b_{j+1} \leq d_i.$$

Then, the corresponding EDF-ordered $|\mathcal{J}|$ -tuple $\mathbf{f} = (f_1, f_2, \dots, f_N)$ is given by

$$f_k = b_j \text{ s.t. } r_k \in [b_{j-1}, b_j) \quad \text{for all } 1 \leq k \leq N.$$

We call such $[b_{j-1}, b_j]$ an *atomic interval*. For example, the intervals $[r_N, r_{N-3}]$ and $[r_N, d_N]$ in Figure 6(a) are atomic, but the interval $[r_N, d_{N-1}]$ is not atomic. (Later, we will formally define the term atomic interval in arbitrary job sets other than this example.) Let t_h be the h th earliest time point in $T_{\mathcal{J}}$, and let $\mathcal{S}_{h,g}$ represent the energy-optimal voltage schedule defined within $[t_h, t_g]$ for the job set $\mathcal{J}_{h,g}$ defined by

$$\mathcal{J}_{h,g} = \{J'_i \mid r_{J'_i} \in [t_h, t_g)\}$$

where

$$r_{J'_i} = r_{J_i}, c_{J'_i} = c_{J_i}, p_{J'_i} = p_{J_i}, \text{ and } d_{J'_i} = \min\{d_{J_i}, t_g\}.$$

Then, we have

$$E(\mathcal{S}_{\text{opt}}^{\mathcal{J}}) = E(\mathcal{S}_{1,|\mathcal{J}|}) = \min \left\{ \sum_{j=1}^{k-1} E(\mathcal{S}_{h_j, h_{j+1}}) \mid 1 = h_1 < h_2 < \dots < h_k = |\mathcal{J}| \right\}$$

and

$$[t_{h_j}, t_{h_{j+1}}] \text{ is atomic for all } j = 1, \dots, k - 1\}.$$

Given an atomic interval $[t_{h_j}, t_{h_{j+1}}]$, $\mathcal{S}_{h_j, h_{j+1}}$ can be directly computed by Yao's algorithm. In this way, the optimal voltage scheduling problem is reduced to a variant of the subset-sum problem. That is, for such job sets as in Figure 6, our problem can be formulated as follows:

Select a tuple (h_1, h_2, \dots, h_k) ($1 = h_1 < \dots < h_k = |\mathcal{J}|$) of integers such that the sum

$$q_{h_1, h_2} + q_{h_2, h_3} + \dots + q_{h_{k-1}, h_k}$$

is minimized subject to $[t_{h_i}, t_{h_{i+1}}]$ is atomic for all $1 \leq i < k$, where $q_{h_j, h_{j+1}}$ denotes $E(\mathcal{S}_{h_j, h_{j+1}})$ (which can be directly computed by Yao's algorithm).

5.1.2 Basic Idea: The Second Example. The example job set in Figure 6 is illustrative in showing how our problem can be formulated by dynamic programming. However, the easily partitionable structure comes from the fact the job set follows the “reverse” EDF priority. For example, in Figure 6, since f_N is set to be r_{N-3} , which is within the execution intervals of J_{N-1} and J_{N-2} , f_{N-1} and f_{N-2} cannot be larger than f_N (or r_{N-3}) so that the modified job set should be EDF-equivalent. Furthermore, f_{N-1} and f_{N-2} are set to be the maximum possible value, f_N , for the modified job set to be essential.

If the priority pattern is not the same as the example job set in Figure 6, the partitioning becomes difficult. For example, the essential job sets in Figures 3(c) and (d) cannot be obtained by the partitioning procedure just explained. In Figure 7(a), J_4 has the lowest priority and the latest deadline, which makes f_4 to be d_4 for all essential job sets (Figures 7(a)–(c)). Therefore, any atomic interval (e.g., $[r_3, r_1]$, $[r_1, d_1]$, or $[r_3, d_3]$) contains partial workload of J_4 , which we call a *background* workload. In the following, we first explain how to extend the dynamic programming formulation to handle the *background* workload. Then, we describe how to explore essential job sets of a given arbitrary job set (as in Figure 3) by dynamic programming.

From Lemma 3.7, the job J_4 in Figure 7 runs at the same speed if the voltage schedule is an optimal one. For the time being, let us assume that the constant speed is among $S_C = \{s_1, s_2, s_3\}$. (For now, S_C is set to be the set of all the possible constant speeds in the optimal voltage schedule. In Section 5.2, we explain how the set S_C is selected such that the size of S_C is bounded by a polynomial function.) For each $s_i \in S_C$, we first compute the amount of background workload

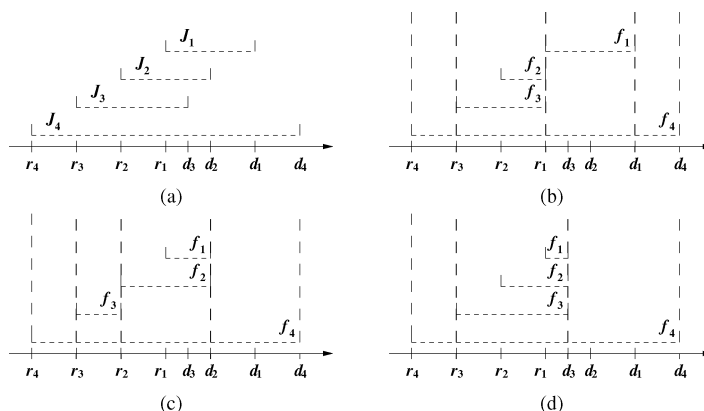


Fig. 7. An example of background workload.

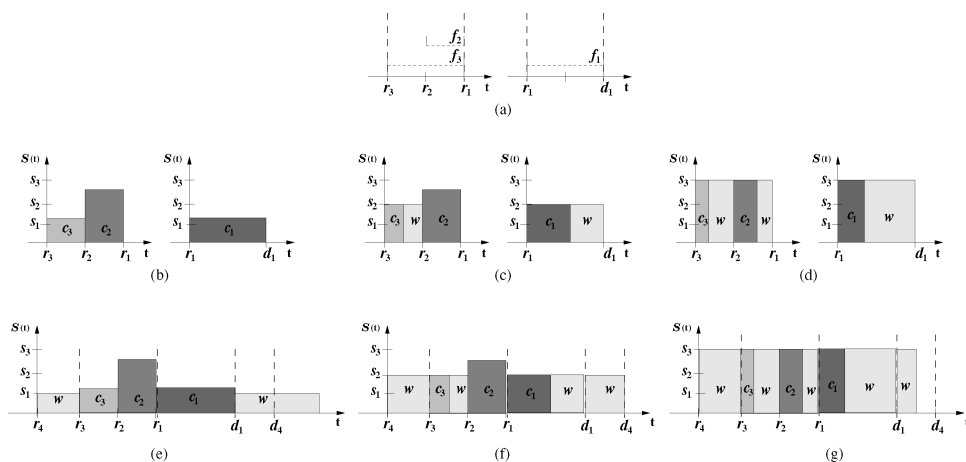


Fig. 8. An example illustrating the algorithm on a job set with background workload; (a) atomic intervals (obtained from the job set in Figure 7(b)). The optimal schedules for two atomic intervals where the speeds of background workload of J_4 are (b) s_1 , (c) s_2 and (d) s_3 , respectively. The voltage schedules for overall time intervals where the speeds of J_4 are (e) s_1 , (f) s_2 , and (g) s_3 , respectively.

of J_4 for each atomic interval, and then find the minimum-energy essential job set (among those in Figures 7(b)–(d)) by using the similar procedure to the previous case in Figure 6. However, unlike the previous case, we discard any job set for which the sum of background workloads executed in overall time interval is less than the total workload of J_4 .

Figure 8(a) shows the atomic intervals $[r_3, r_1]$ and $[r_1, d_1]$, which are obtained from the essential job set in Figure 7(b). Figures 8(b)–(d) show the optimal voltage schedules for the atomic intervals, where J_4 runs at the speed s_1 , s_2 , and s_3 , respectively. The workloads of jobs J_1 , J_2 , and J_3 are denoted by c_1 , c_2 , and c_3 , respectively, and the background workloads are denoted by w . The amount of the background workload (and the resultant optimal voltage schedule) for each atomic interval and speed can be easily computed by a slightly modified

```

1:  $f_{\sigma^{-1}(|\mathcal{J}|)} := d_{J_{\sigma^{-1}(|\mathcal{J}|)}}$ 
2:  $\mathbf{b}_\sigma := (d_{J_{\sigma^{-1}(|\mathcal{J}|)}}$ 
3: for ( $i := |\mathcal{J}| - 1$  to 1)
4:   let  $\mathcal{J}^H$  be  $\{J_{\sigma^{-1}(k)} \mid i < k \leq |\mathcal{J}| \wedge \sigma^{-1}(k) < \sigma^{-1}(i)\}$ 
5:   if ( $r_{J_{\sigma^{-1}(i)}} \geq \min(\{r_J \mid J \in \mathcal{J}^H\} \cup \{f_{\sigma^{-1}(i+1)}\})$ ) return FALSE
6:   else  $f_{\sigma^{-1}(i)} := \min(\{f_{\sigma^{-1}(i+1)}, d_{J_{\sigma^{-1}(i)}}\} \cup \{r_J \mid J \in \mathcal{J}^H\})$ 
7:   end if
8:   if ( $f_{\sigma^{-1}(i)} \leq \min(\{r_{J_{\sigma^{-1}(k)}} \mid i < k \leq |\mathcal{J}|\})$ ) append  $f_{\sigma^{-1}(i)}$  onto the head of  $\mathbf{b}_\sigma$ 
9:   end if
10: end for
11: append  $\min R_J$  onto the head of  $\mathbf{b}_\sigma$ 

```

Fig. 9. The algorithm to build a strongly blocking tuple from a $|\mathcal{J}|$ -permutation.

version of Yao's algorithm [Yao et al. 1995]. That is, when the critical interval is selected, if the speed to be assigned (by the intensity of the critical interval) is less than or equal to the speed of the background workload, we assign the speed of the background workload to all the unscheduled time intervals (including the critical interval). Then, the amount of background workload can be directly computed as in Figures 8(b)–(d).

Once the background workload and the optimal voltage schedule are computed for each atomic interval, we apply the same procedure as in the job set in Figure 6 to find the minimum-energy essential job set and the energy-optimal voltage schedule. In exploring the solution space, we should discard any infeasible schedules. Figure 8(e) shows an infeasible schedule, where J_4 runs at s_1 and cannot complete its execution until its deadline. The voltage schedule in Figure 8(g) is feasible, but not an optimal one. Thus, only the schedule in Figure 8(f) is not removed in the pruning procedure and is compared with another schedules obtained from the essential job sets in Figures 8(c) and (d).

5.1.3 Putting It Altogether. We now describe the optimal algorithm for arbitrary job sets based on the observations from the example job sets. First, we formally define the terms *strongly atomic interval* and *strongly blocking tuple*. Given a valid $|\mathcal{J}|$ -permutation σ , the algorithm in Figure 9 builds the corresponding strongly blocking tuple $\mathbf{b}_\sigma = (b_1, b_2, \dots, b_k)$, where $b_1 < b_2 < \dots < b_k$ and $b_i \in T_{\mathcal{J}}$ for all $1 \leq i \leq k$. The algorithm is identical to the algorithm in Figure 2 except for lines 2, 8, 9, and 11. In line 8, $f_{\sigma^{-1}(i)}$ is selected as an element of a strongly blocking tuple if it partitions the execution interval.

Definition 5.1. Given a valid $|\mathcal{J}|$ -permutation σ , the tuple \mathbf{b}_σ built by the algorithm in Figure 9 is called a *strongly blocking tuple*. An interval $[t, t']$ is *strongly atomic* if there is a strongly blocking tuple $\mathbf{b} = (b_1, b_2, \dots, b_k)$ such that $[t, t'] = [b_i, b_{i+1}]$ for some $1 \leq i < k$. Furthermore, the job set $\mathcal{J}_{[t, t']}$ is defined by

$$\mathcal{J}_{[t, t']} = \{J' \mid J \in \mathcal{J}, r_J \in [t, t']\}$$

```

/*  $T_j = \{t_1, t_2, \dots, t_N\}$  */
1: foreach (strongly-atomic interval  $[t_i, t_j]$ )
2:    $g_{i,j} := E(\mathcal{S}_{\text{opt}}^{j_{[t_i, t_j]}})$ 
3: end foreach
4:  $\mathbf{V} := \{v_1, v_2, \dots, v_N\}$ 
5:  $\mathbf{E} := \{(v_i, v_j) \mid [t_i, t_j] \text{ is strongly-atomic}\}$ 
6: foreach  $((v_i, v_j) \in \mathbf{E})$   $w((v_i, v_j)) := g_{i,j}$  end foreach /* weight of edges */
7: Find the shortest path from  $v_1$  to  $v_N$  in  $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ . /* Note that  $\mathbf{G}$  is acyclic. */
/* The shortest path =  $\langle v_{q_1}, v_{q_2}, \dots, v_{q_k} \rangle$  ( $v_{q_1} = v_1, v_{q_k} = v_N$ ) */
8: return  $\oplus_{j=1}^{k-1} \mathcal{S}_{\text{opt}}^{j_{[v_{q_j}, v_{q_{j+1}}]}}$ 

```

Fig. 10. An exponential-time optimal algorithm based on strongly atomic intervals.

```

1:  $f_{\sigma^{-1}(\mathcal{J})} := d_{J_{\sigma^{-1}(\mathcal{J})}}$ 
2:  $\mathbf{b}_{\sigma}^w := (d_{J_{\sigma^{-1}(\mathcal{J})}})$ 
3: for  $(i := |\mathcal{J}| - 1 \text{ to } 1)$ 
4:   let  $\mathcal{J}^H$  be  $\{J_{\sigma^{-1}(k)} \mid i < k \leq |\mathcal{J}| \wedge \sigma^{-1}(k) < \sigma^{-1}(i)\}$ 
5:   if  $(r_{J_{\sigma^{-1}(i)}} \geq \min(\{r_J \mid J \in \mathcal{J}^H\} \cup \{f_{\sigma^{-1}(i+1)}\}))$  return FALSE
6:   else  $f_{\sigma^{-1}(i)} := \min(\{f_{\sigma^{-1}(i+1)}, d_{J_{\sigma^{-1}(i)}}\} \cup \{r_J \mid J \in \mathcal{J}^H\})$ 
7:   end if
8:   if  $(f_{\sigma^{-1}(i)} \leq \min\{r_J \mid r \in \mathcal{J}^H\})$  append  $f_{\sigma^{-1}(i)}$  onto the head of  $\mathbf{b}_{\sigma}^w$ 
9:   end if
10: end for
11: append  $\min R_J$  onto the head of  $\mathbf{b}_{\sigma}^w$ 

```

Fig. 11. The algorithm to build a weakly blocking tuple from a $|\mathcal{J}|$ -permutation.

where

$$r_{J'} = r_J, c_{J'} = c_J, p_{J'} = p_J, \text{ and } d_{J'} = \min\{d_J, t'\}$$

is said to be *induced* by an interval $[t, t']$.³

For the job set in Figure 3, not only $[r_3, r_2]$, $[r_2, d_2]$ (Figure 3(b)) and $[r_3, r_1]$ (Figure 3(c)) but also $[r_1, d_2]$ (Figure 3(c)) and $[r_3, d_3]$ (Figure 3(d)) are strongly atomic. Note that the intervals $[r_1, d_2]$ and $[r_3, d_3]$ are not covered by the previous definition in Section 5.1.1. Furthermore, (r_3, r_2, d_2) (Figure 3(b)), (r_3, r_1, d_2) (Figure 3(c)), and (r_3, d_3) (Figure 3(d)) are strongly blocking tuples.

Figure 10 shows an optimal algorithm that is based on strongly atomic intervals. The correctness of the algorithm is proved in Appendix A.1. The algorithm may work efficiently for some job sets (e.g., the job set in Figure 6). But the running time may not be bounded by a polynomial function; for the job set in Figure 7, there is only one strongly atomic interval $[r_4, d_4]$ and the algorithm cannot but enumerate all the essential job sets. Furthermore, the algorithm does not have a structure suitable to be transformed into an FPTAS. So we consider another optimal algorithm based on *weakly atomic* intervals, *weakly bounding* tuples, and the background workload. First, we formally define the terms based on the algorithm in Figure 11, which is

³ $[t, t']$ is not required to be strongly atomic.

identical to the algorithm in Figure 9 except for the boxed code segment (line 8).

Definition 5.2. Given a valid $|\mathcal{J}|$ -permutation σ , the tuple \mathbf{b}_σ^w built by the algorithm in Figure 11 is called a *weakly blocking tuple*. An interval $[t, t']$ is *weakly atomic* if there is a weakly blocking tuple $\mathbf{b}^w = (b_1, b_2, \dots, b_k)$ such that $[t, t'] = [b_i, b_{i+1}]$ for some $1 \leq i < k$. Furthermore, the job set $\mathcal{J}_{[t, t']^w}$ is defined by

$$\mathcal{J}_{[t, t']^w} = \{J' \mid J \in \mathcal{J}, r_J \in [t, t'] \wedge (\exists J_h \in \mathcal{J}, p_{J_h} < p_J \wedge r_{J_h} = t' \wedge d_J \in [r_{J_h}, d_{J_h}))\}$$

where

$$r_{J'} = r_J, c_{J'} = c_J, p_{J'} = p_J, \text{ and } d_{J'} = \min\{d_J, t'\}$$

is said to be *weakly induced* by an interval $[t, t']$.

Furthermore, given a weakly blocking tuple $\mathbf{b}^w = (b_1, b_2, \dots, b_k)$ and the corresponding EDF-equivalent job set \mathcal{J}' , any job in $\mathcal{J}' - \cup_{j=1}^{k-1} \mathcal{J}_{[b_j, b_{j+1}]^w}$ is called a *background job* with respect to the weakly blocking tuple \mathbf{b}^w . The workload of background jobs is called *background workload*.

Based on weakly atomic interval, we construct another optimal voltage scheduling algorithm. Figure 12 shows the optimal algorithm that is based on the dynamic programming formulated by weakly atomic intervals. The algorithm identifies weakly atomic intervals and computes the optimal schedule for the weakly atomic interval. (Note that jobs in a weakly atomic interval follow the EDF priority assignment.) In computing the optimal schedule for a weakly atomic interval, we consider the background workload, that is, the algorithm computes the optimal schedule for each candidate background speed in S_C . Given a job set \mathcal{J} , the algorithm first computes the set S_C of candidates for the speed of background workload. For the optimal algorithm, the set S_C is set to be

$$S_C = \left\{ \frac{C(\mathcal{J}')}{\sum_{i=0}^{k-1} (t_{p_{2i+2}} - t_{p_{2i+1}})} \mid \mathcal{J}' \subset \mathcal{J}, t_1 < t_2 < \dots < t_{p_{2k}}, t_j \in T_{\mathcal{J}} \right\}.$$

It is obvious that the speed of the background workload in an optimal voltage schedule is included in S_C . (In the FPTAS presented in Section 5.2, the set S_C is selected such that the size of S_C is bounded by a polynomial function.) Given the optimal schedules of weakly atomic intervals, the algorithm searches the minimum sum of the energy values of the weakly atomic intervals. The correctness of the algorithm is proved in Appendix A.2. The worst-case running time of the algorithm is not bounded by a polynomial function, but it can be easily transformed into an FPTAS.

5.2 Approximation Algorithm

First, we prove a miscellaneous property that is useful in bounding the error of our approximation algorithm.

```

procedure OPTIMAL_VOLTAGE_SCHEDULE ( $\mathcal{J}$ )
  /*  $T_{\mathcal{J}} = \{t_1, t_2, \dots, t_N\}$ ,  $S_C := \{s_1, s_2, \dots, s_n\}$  */
  1: foreach ( $s \in S_C$ )
  2:    $\mathbf{V} := \{v_1, v_2, \dots, v_N\}$ 
  3:    $\mathbf{E} := \{(v_i, v_j) \mid [t_i, t_j] \text{ is weakly-atomic}\}$ 
  4:   foreach ( $(v_i, v_j) \in \mathbf{E}$ )
  5:      $w((v_i, v_j)) := W(\max\{\mathcal{S}_{\text{opt}}^{\mathcal{J}[t_i, t_j]}(t), s\}, [t_i, t_j]) - W(\mathcal{S}_{\text{opt}}^{\mathcal{J}[t_i, t_j]}(t), [t_i, t_j])$  /* weight of edges */
  6:   end foreach
  7:   Find longest paths between all pairs of vertices in  $\mathbf{V}$ . /* Note that  $G$  is acyclic. */
  8:   foreach ( $1 \leq i < j \leq N$  s.t.  $[t_i, t_j]$  is a concatenation of weakly-atomic intervals)
  9:     /* The longest path from  $v_i$  to  $v_j = \langle v_{q_1}, v_{q_2}, \dots, v_{q_l} \rangle$ 
  10:     $c :=$  the weight of the longest path from  $v_i$  to  $v_j$ .
  11:     $E_{i,j}[c] := E(\oplus_{h=1}^{l-1} \max\{\mathcal{S}_{\text{opt}}^{\mathcal{J}[v_{q_j}, v_{q_{j+1}}]}(t), s\}, [t_i, t_j])$ 
  12:   end foreach
  13:   for ( $i := 1$  to  $N-1$ )
  14:     for ( $j := 1$  to  $N-i$ )
  15:        $E_{j,j+i} := \infty^+$ 
  16:       for ( $k := j+1$  to  $j+i$ )
  17:          $c_{j,j+i,k} := C(\{J \in \mathcal{J}^B \mid r_J \in [t_j, t_k] \wedge d_J \in [t_k, t_{j+i}]\})$ 
  18:          $E_{j,j+i,k} := E_{j,k}[c_{j,j+i,k}] + E_{k,j+i}$ 
  19:         if ( $E_{j,j+i} > E_{j,j+i,k}$  and  $\mathcal{S}_{\text{opt}}^{\mathcal{J}[v_j, v_k]^w [c_{j,j+i,k}]}$  is feasible for  $\mathcal{J}_{[t_i, t_j]^w} \cup \{J \in \mathcal{J}^B \mid [r_J, d_J] \subseteq [t_i, t_j]\}$ )
  20:            $E_{j,j+i} := E_{j,j+i,k}$ ,  $h := k$ 
  21:         end if
  22:       end for
  23:        $\mathbf{b}_{j,j+i} := \{t_h\} \cup \mathbf{b}_{j,h} \cup \mathbf{b}_{h,j+i}$ 
  24:     end for
  25:   end for
  /*  $E_{1,N} = E(\mathcal{S}_{\text{opt}}^{\mathcal{J}})$  and  $\mathcal{S}_{\text{opt}}^{\mathcal{J}} \equiv \mathcal{S}_{\text{opt}}^{\cup_{h=1}^{N-1} \mathcal{J}_{[b_h, b_{h+1}]^w} \cup \mathcal{J}^B}$  where  $\mathbf{b}_{1,N} = (b_1, b_2, \dots, b_l)$  */
  26:    $\mathcal{J}_{\text{opt}} := \cup_{h=1}^{N-1} \mathcal{J}_{[b_h, b_{h+1}]^w} \cup \mathcal{J}^B$  where  $\mathbf{b}_{1,N} = (b_1, b_2, \dots, b_l)$ 
  27:   return  $\mathcal{S}_{\text{opt}}^{\mathcal{J}_{\text{opt}}}$  /*  $\mathcal{J}_{\text{opt}}$  is an EDF job set. So,  $\mathcal{S}_{\text{opt}}^{\mathcal{J}_{\text{opt}}}$  can be directly computed by Yao's algorithm */
end procedure

```

Fig. 12. An exponential-time optimal algorithm based on weakly atomic intervals.

LEMMA 5.3. *Given a function $P : \mathbf{R}^+ \Rightarrow \mathbf{R}^+$ and a constant $0 < \varepsilon < 1$, if*

$$0 < x_1 < x_2 < \left(1 + \frac{\varepsilon \cdot \log 2}{\max\{\eta(x) \mid x > 0\}}\right) \cdot x_1,$$

where

$$\eta(x) = \frac{P'(x)}{P(x)} \cdot x,$$

then $P(x_2) < (1 + \varepsilon) \cdot P(x_1)$.

PROOF. From the condition, we have

$$\log x_2 - \log x_1 < \log \left(1 + \frac{\varepsilon \cdot \log 2}{\max\{\eta(x) \mid x > 0\}}\right) < \frac{\varepsilon \cdot \log 2}{\max\{\eta(x) \mid x > 0\}}. \quad (11)$$

Let $y_1 = \log x_1$ and $y_2 = \log x_2$. Then we have

$$\begin{aligned} \log P(x_2) - \log P(x_1) &= \log P(e^{y_2}) - \log P(e^{y_1}) \\ &\leq (y_2 - y_1) \cdot \max \left\{ \frac{d(\log P(e^y))}{dy} \right\}. \end{aligned}$$

From (11) and

$$\frac{d(\log P(e^y))}{dy} = \frac{P'(e^y)}{P(e^y)} \cdot e^y = \eta(e^y),$$

we have

$$\log P(x_2) - \log P(x_1) < \frac{\varepsilon \cdot \log 2}{\max\{\eta(x)|x > 0\}} \cdot \max\{\eta(x)|x > 0\} = \varepsilon \cdot \log 2.$$

It follows that

$$P(x_2) < e^{\varepsilon \cdot \log 2} \cdot P(x_1) < e^{\log(1+\varepsilon)} \cdot P(x_1) = (1 + \varepsilon) \cdot P(x_1). \quad \square$$

For a power function $P(s) = \alpha \cdot s^n$, we have $\eta(s) = n$. In the following, we use ρ_P to denote $\log 2 / \max\{\eta(x)|x > 0\}$. From Lemma 5.3, we can construct an FPTAS as in Figure 13. The FPTAS is slightly different from the algorithm in Figure 12. To bring the running time down to polynomial, we use S'_C instead of S_C :

$$S'_C = \{ \min\{S_C\} \cdot (1 + \varepsilon \cdot \rho_P)^k \mid k = 0, 1, \dots, l \}$$

where

$$\min\{S_C\} \cdot (1 + \varepsilon \cdot \rho_P)^{l-1} < \max\{S_C\} \leq \min\{S_C\} \cdot (1 + \varepsilon \cdot \rho_P)^l.$$

THEOREM 5.4. *APPROX_VOLTAGE_SCHEDULE is a fully polynomial time approximation scheme for the voltage scheduling problem.*

PROOF. Let s_1 and s_2 be elements of S'_C such that $s_2 = s_1 \cdot (1 + \varepsilon \cdot \rho_P)$. Given a weakly atomic interval $[t_i, t_j]$, we have for $t_i \leq t \leq t_j$:

$$\max \left\{ \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}}(t), s_2 \right\} \leq (1 + \varepsilon \cdot \rho_P) \cdot \max \left\{ \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}}(t), s_1 \right\}.$$

Thus, from Lemma 5.3, we have for $t_i \leq t \leq t_j$

$$P \left(\max \left\{ \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}}(t), s_2 \right\} \right) \leq (1 + \varepsilon) \cdot P \left(\max \left\{ \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}}(t), s_1 \right\} \right),$$

which implies

$$E \left(\max \left\{ \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}}(t), s_2 \right\}, [t_i, t_j] \right) \leq (1 + \varepsilon) \cdot E \left(\max \left\{ \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}}(t), s_1 \right\}, [t_i, t_j] \right).$$

Let us compare $E_{j,k}[c']$ in line 21 of APPROX_VOLTAGE_SCHEDULE and $E_{j,k}[c_{j,j+i,k}]$ in line 18 of OPTIMAL_VOLTAGE_SCHEDULE. Let s' and s be the corresponding elements in S'_C and S_C , respectively. Then, from the definition of S'_C , we have $s' < (1 + \varepsilon \cdot \rho_P) \cdot s$, which implies $E_{j,k}[c'] < (1 + \varepsilon) \cdot E_{j,k}[c_{j,j+i,k}]$. Therefore, $E_{1,N} < (1 + \varepsilon) \cdot E(S_{\text{opt}}^{\mathcal{J}})$.

```

procedure APPROX_VOLTAGE_SCHEDULE ( $\mathcal{J}, \varepsilon$ )
  /*  $T_j = \{t_1, t_2, \dots, t_N\}$  */
  /*  $S'_C = \{\min\{S_C\} \cdot (1 + \delta)^k \mid k = 0, 1, \dots, \lceil \log_{1+\delta}(\max\{S_C\} / \min\{S_C\}) \rceil\}$  where  $\delta = \varepsilon \cdot \rho_P$  */
  1: Initialize  $C_{i,j} := \{\}$  for  $1 \leq i < j \leq N$ .
  2: foreach ( $s \in S'_C$ )
  3:    $\mathbf{V} := \{v_1, v_2, \dots, v_N\}$ 
  4:    $\mathbf{E} := \{(v_i, v_j) \mid [t_i, t_j] \text{ is weakly-atomic}\}$ 
  5:   foreach ( $(v_i, v_j) \in \mathbf{E}$ )
  6:      $w((v_i, v_j)) := W(\max\{S_{\text{opt}}^{j_{[t_i, t_j]}}(t), s\}, [t_i, t_j]) - W(S_{\text{opt}}^{j_{[t_i, t_j]}}(t), [t_i, t_j])$  /* weight of edges */
  7:   end foreach
  8:   Find longest paths between all pairs of vertices in  $\mathbf{V}$ . /* Note that  $G$  is acyclic. */
  9:   foreach ( $1 \leq i < j \leq N$  s.t.  $[t_i, t_j]$  is a concatenation of weakly-atomic intervals)
  10:    /* The longest path from  $v_i$  to  $v_j = \langle v_{q_1}, v_{q_2}, \dots, v_{q_l} \rangle$ 
  11:     $c :=$  the weight of the longest path from  $v_i$  to  $v_j$ .
  12:     $E_{i,j}[c] := E(\oplus_{h=1}^{l-1} \max\{S_{\text{opt}}^{j_{[v_{q_h}, v_{q_{h+1}]}}}(t), s\}, [t_i, t_j])$ 
  13:     $C_{i,j} := C_{i,j} \cup \{c\}$ 
  14:   end foreach
  15: for ( $i := 1$  to  $N - 1$ )
  16:   for ( $j := 1$  to  $N - i$ )
  17:     $E_{j,j+i} := \infty^+$ 
  18:    for ( $k := j + 1$  to  $j + i$ )
  19:      $c_{j,j+i,k} := C(\{J \in \mathcal{J}^B \mid r_J \in [t_j, t_k] \wedge d_J \in [t_k, t_{j+i}]\})$ 
  20:      $c' := \min\{c \in C_{j,k} \mid c \geq c_{j,j+i,k}\}$ 
  21:      $E_{j,j+i,k} := E_{j,k}[c'] + E_{k,j+i}$ 
  22:     if ( $E_{j,j+i} > E_{j,j+i,k}$  and
  23:        $S_{\text{opt}}^{j_{[t_j, t_k]}^w [c_{j,j+i,k}]}$  is feasible for  $\mathcal{J}_{[t_i, t_j]}^w \cup \{J \in \mathcal{J}^B \mid [r_J, d_J] \subseteq [t_i, t_j]\}$ )
  24:        $E_{j,j+i} := E_{j,j+i,k}$ ,  $h := k$ 
  25:     end if
  26:   end for
  27:    $\mathbf{b}_{j,j+i} := \{t_h\} \cup \mathbf{b}_{j,h} \cup \mathbf{b}_{h,j+i}$ 
  28: end for
  29: /*  $E_{1,N} < (1 + \varepsilon) \cdot E(S_{\text{opt}}^j)$  */
  30:  $\mathcal{J}_\varepsilon := \cup_{h=1}^{l-1} \mathcal{J}_{[b_h, b_{h+1}]}^w \cup \mathcal{J}^B$  where  $\mathbf{b}_{1,N} = (b_1, b_2, \dots, b_l)$ 
  31: return  $S_{\text{opt}}^{\mathcal{J}_\varepsilon}$  /*  $\mathcal{J}_\varepsilon$  is an EDF job set. So,  $S_{\text{opt}}^{\mathcal{J}_\varepsilon}$  can be directly computed by Yao's algorithm */
end procedure

```

Fig. 13. The fully polynomial time approximation scheme.

Finally, since we have

$$|S'_C| = 1 + \lceil \log_{1+\varepsilon \cdot \rho_P}(\max\{S_C\} / \min\{S_C\}) \rceil \quad (12)$$

$$< 2 + \frac{\log(\max\{S_C\} / \min\{S_C\})}{\varepsilon \cdot \log(1 + \rho_P)}, \quad (13)$$

the running time is bounded a polynomial function of $|\mathcal{J}|$ and $1/\varepsilon$. \square

6. EXPERIMENTAL RESULTS

In order to evaluate how the proposed FPTAS performs, we have performed several experiments using the FPTAS described in Figure 13. For a comparison, we also implemented Quan's heuristic [Quan and Hu 2001], which is currently

Table I. Experimental Results for Three Real-World Real-Time Applications

Applications	Normalized Energy			CPU Time(s)	
	MPEG4	CNC	Avionics	CNC	Avionics
No. of jobs	22	289	1372	289	1372
FPTAS					
$\varepsilon = 0.1\%$	1	1	1	44.71	4506.63
$\varepsilon = 0.5\%$	1.003	1.004	1.003	11.67	1021.48
$\varepsilon = 1.0\%$	1.006	1.008	1.007	6.12	631.15
$\varepsilon = 1.5\%$	1.012	1.013	1.011	5.16	512.32
$\varepsilon = 2.0\%$	1.017	1.018	1.018	3.81	313.15
Quan [Quan and Hu 2001]	1.041	1.062	1.059	4.76	580.32

the best polynomial-time voltage scheduling algorithm for fixed-priority real-time tasks. We compared the energy efficiency and computation time between two algorithms.⁴

In our experiments, we assumed that the energy consumption is quadratically dependent on the supply voltage. For a given supply voltage V , the corresponding clock frequency f is proportional to $(V_{DD} - V_{TH})^\alpha / V_{DD}$, where V_{TH} and α are assumed to be 0.5 V and 1.3 [Sakurai and Newton 1990].

We constructed test job sets from periodic task sets of three real-world applications: MPEG4 Videophone [Shin et al. 2001], CNC [Kim et al. 1996], and Avionics [Locke et al. 1991]. Table I summarizes the experimental results for these job sets. In each experiment, the execution time of each job (i.e., task instance) was randomly drawn from a Gaussian distribution⁵ within the range of $[WCET/10, WCET]$ of each task. Results were normalized over the energy consumption of each application scheduled by the proposed FPTAS with $\varepsilon = 0.1\%$. As shown in Table I the FPTAS outperforms Quan's algorithm, spending reasonable CPU times. In the experiments, actual errors were always less than given ε 's. (We omit CPU times for MPEG4 Videophone because they are less than 0.1.)

We also performed experiments using synthesized job sets with the varying number of jobs from 50 to 1600. We conjectured that one of the key parameters affecting the performance of Quan's algorithm is the degree of interference among jobs. Since the degree of interference is mainly dependent on the lengths of the execution intervals of the jobs, we generated three classes of job sets as follows: for the first class of job sets (Class 1), the release time and the length of the execution interval of a job are selected under the uniform distribution within $[0, 1000]$ and $[50, 100]$, respectively. The workload of each job was randomly selected from a uniform distribution within $[0.2, 1.0]$. (Note that it is sufficient to consider only the relative values of workloads, since the maximum processor speed can be always appropriately adjusted.) For the second class of jobs (Class 2) and the third class of jobs (Class 3), we used $[100, 300]$ and $[300, 500]$ (instead of $[50, 100]$) for the length of the execution intervals,

⁴We have implemented the exhaustive optimal algorithm by Quan and Hu [2002] as well for experiments. This algorithm, however, takes an excessive amount of time. For example, it took more than a day when $N = 25$. Therefore, we cannot include the experimental results for this algorithm.

⁵With the mean $m = \frac{WCET/10 + WCET}{2}$ and the standard deviation $\sigma = \frac{WCET - WCET/10}{6}$.

Table II. Experimental Results for Synthesized Jobs (Class 1)

Job sets	Normalized Energy					
	\mathcal{J}_1	\mathcal{J}_2	\mathcal{J}_3	\mathcal{J}_4	\mathcal{J}_5	\mathcal{J}_6
No. of jobs	50	100	200	400	800	1600
FPTAS						
$\varepsilon = 0.1\%$	1	1	1	1	1	1
$\varepsilon = 0.5\%$	1.003	1.003	1.004	1.004	1.003	1.003
$\varepsilon = 1.0\%$	1.008	1.007	1.009	1.009	1.008	1.009
$\varepsilon = 1.5\%$	1.013	1.012	1.012	1.014	1.014	1.014
$\varepsilon = 2.0\%$	1.016	1.016	1.019	1.018	1.019	1.019
Quan [Quan and Hu 2001]	1.044	1.047	1.051	1.054	1.052	1.071

Table III. Experimental Results for Synthesized Jobs (Class 2)

Job sets	Normalized Energy					
	\mathcal{J}_1	\mathcal{J}_2	\mathcal{J}_3	\mathcal{J}_4	\mathcal{J}_5	\mathcal{J}_6
No. of jobs	50	100	200	400	800	1600
FPTAS						
$\varepsilon = 0.1\%$	1	1	1	1	1	1
$\varepsilon = 0.5\%$	1.004	1.004	1.003	1.004	1.003	1.004
$\varepsilon = 1.0\%$	1.009	1.007	1.007	1.008	1.009	1.009
$\varepsilon = 1.5\%$	1.013	1.012	1.014	1.014	1.013	1.014
$\varepsilon = 2.0\%$	1.018	1.016	1.018	1.018	1.019	1.019
Quan [Quan and Hu 2001]	1.055	1.062	1.070	1.079	1.103	1.127

Table IV. Experimental Results for Synthesized Jobs (Class 3)

Job sets	Normalized Energy					
	\mathcal{J}_1	\mathcal{J}_2	\mathcal{J}_3	\mathcal{J}_4	\mathcal{J}_5	\mathcal{J}_6
No. of jobs	50	100	200	400	800	1600
FPTAS						
$\varepsilon = 0.1\%$	1	1	1	1	1	1
$\varepsilon = 0.5\%$	1.004	1.004	1.004	1.003	1.004	1.004
$\varepsilon = 1.0\%$	1.009	1.007	1.007	1.009	1.008	1.009
$\varepsilon = 1.5\%$	1.014	1.013	1.014	1.013	1.014	1.014
$\varepsilon = 2.0\%$	1.018	1.017	1.019	1.018	1.019	1.019
Quan [Quan and Hu 2001]	1.094	1.114	1.121	1.134	1.142	1.137

respectively. Note that Class 1, Class 2, and Class 3 correspond to job sets with low, medium, and high degrees of interference among the jobs. Tables II, III, and IV show the experimental results for Class 1, Class 2, and Class 3. As shown in tables, in general, the higher the degree of interferences becomes, the larger the improvement of our algorithm over Quan's algorithm.

7. CONCLUSIONS

We investigated the problem of energy-optimal voltage scheduling for fixed-priority real-time systems implemented on a variable voltage processor. First,

we proved the NP-hardness of the problem. Our complexity analysis provided an important new insight into the problem.

Knowing the NP-hardness of the problem as the best practical solution, we described a fully polynomial time approximation scheme for the problem. That is, for any $\varepsilon > 0$, the proposed approximation scheme computes a voltage schedule whose energy consumption is bounded by $(1 + \varepsilon)$ times that of the optimal voltage schedule. Furthermore, the running time of the proposed approximation scheme is bounded as well by a polynomial function of the number of input jobs and $1/\varepsilon$. Experimental results show that the proposed approximation scheme runs sufficiently fast even for a small error bound (i.e., 0.5%).

While the proposed approximation scheme is efficient for general fixed-priority job sets, the proposed scheme can be further extended in several directions. For example, we are interested in devising more efficient algorithms for more specialized job sets such as job sets from periodic task sets. In addition, we plan to modify the proposed approximation scheme to work under a more realistic processor model with a limited number of voltage levels and voltage transition overheads.

APPENDIX: PROOFS

A.1 Proof of the Correctness of the Algorithm in Figure 10

We first prove some properties on strongly blocking tuples and strongly atomic intervals. Note that for an interval $[t, t']$, $\mathbf{I}_{\mathcal{J}_{[t, t']}} \subseteq [t, t']$ since $t \leq r_J < d_J \leq t'$ for all $J \in \mathcal{J}_{[t, t']}$. Therefore, for a strongly blocking tuple $\mathbf{b} = (b_1, b_2, \dots, b_k)$, $\mathbf{I}_{\mathcal{J}_{[b_1, b_2]}}$, $\mathbf{I}_{\mathcal{J}_{[b_2, b_3]}}$, \dots , $\mathbf{I}_{\mathcal{J}_{[b_{k-1}, b_k]}}$ are disjoint. Now, we prove that a job set can be partitioned by strongly blocking tuples as with the job set in Figure 6, so that the formulation described in Section 5.1.1 can be extended to cover arbitrary job sets.

LEMMA A.1. *Given a job set \mathcal{J} and an essential $|\mathcal{J}|$ -tuple \mathbf{f} , $\mathcal{J}^{\mathbf{f}} \equiv \bigcup_{j=1}^{k-1} \mathcal{J}_j$, where $\mathbf{b}_{\sigma_j} = (b_1, b_2, \dots, b_k)$ and \mathcal{J}_j is an EDF-equivalent job set of $\mathcal{J}_{[b_j, b_{j+1}]}$ for all $1 \leq j < k$.*

PROOF. Let $\mathcal{J}^{\mathbf{f}} = \{J'_1, J'_2, \dots, J'_{|\mathcal{J}|}\}$ and let $\mathcal{J}_j = \{J'_l \in \mathcal{J}^{\mathbf{f}} | r_{J'_l} (= r_{J_l}) \in [b_j, b_{j+1}]\}$. Then, $\{\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_{k-1}\}$ forms a partition of $\mathcal{J}^{\mathbf{f}}$, that is,

$$\mathcal{J}^{\mathbf{f}} \equiv \bigcup_{j=1}^{k-1} \mathcal{J}_j \quad \text{and} \quad \mathcal{J}_j \cap \mathcal{J}_{j'} = \emptyset \quad \text{for all } 1 \leq j \neq j' < k.$$

Thus, it suffices to show that \mathcal{J}_j is an EDF-equivalent job set of $\mathcal{J}_{[b_j, b_{j+1}]}$ for all $1 \leq j < k$. Let $i_j = \max\{i | f_{\sigma^{-1}(i)} = b_j\}$ for all $1 \leq j \leq k$, and suppose that $d_{J'_l} > b_{j+1}$ for a job $J'_l \in \mathcal{J}_j$. Then, we have $\sigma(l) > i_{j+1}$, since

$$f_{\sigma^{-1}(\sigma(l))} = f_l = d_{J'_l} > b_{j+1} = f_{\sigma^{-1}(i_{j+1})}.$$

From line 8 of the algorithm in Figure 9, we have

$$b_{j+1} = f_{\sigma^{-1}(i_{j+1})} \leq \min\{r_{J_{\sigma^{-1}(k)}} | i_{j+1} < k \leq |\mathcal{J}|\} \leq r_{J_{\sigma^{-1}(k)}} \Big|_{k=\sigma(l) (> i_{j+1})} = r_{J'_l},$$

which contradicts $r_{J'_l} (= r_{J_l}) \in [b_j, b_{j+1})$. Therefore, $d_{J'_l} \in [b_j, b_{j+1}]$ for all $J'_l \in \mathcal{J}_j$. Furthermore, \mathcal{J}_j follows the EDF priority, since it is a subset of the EDF job set \mathcal{J}^f .

It remains to show that $|\mathcal{J}_j| = |\mathcal{J}_{[b_j, b_{j+1}]}|$ and there is a bijective function $\alpha : \mathcal{J}_{[b_j, b_{j+1}]} \Rightarrow \mathcal{J}_j$ such that

$$\forall J' \in \mathcal{J}_{[b_j, b_{j+1}]}, p_{J'} = p_{\alpha(J')}, c_{J'} = c_{\alpha(J')} \quad \text{and} \quad r_{J'} = r_{\alpha(J')}. \quad (14)$$

For the former, we have

$$|\mathcal{J}_j| = |\{J' \in \mathcal{J}^f | r_{J'} \in [b_j, b_{j+1})\}| = |\{J \in \mathcal{J} | r_J \in [b_j, b_{j+1})\}| = |\mathcal{J}_{[b_j, b_{j+1}]}|.$$

For the latter, we define α such that $\alpha(J') = J''$ if and only if $p_{J'} = p_{J''}$. Then, it is clear that α is a bijective function and (14) holds. \square

LEMMA A.2. *Let $S(t) = \bigoplus_{j=1}^{h-1} \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_j, t_{j+1}]}}$ for $\min R_{\mathcal{J}} = t_1 < t_2 < \dots < t_h = \max D_{\mathcal{J}}$ ($t_j \in T_{\mathcal{J}}$). Then, S is a feasible voltage schedule of \mathcal{J} . Furthermore,*

$$E(S) = \sum_{j=1}^{h-1} E\left(\mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_j, t_{j+1}]}}\right) \geq E\left(\mathcal{S}_{\text{opt}}^{\mathcal{J}}\right).$$

PROOF. Let $u_{[t_0, t'_0]}(t)$ be defined by

$$u_{[t_0, t'_0]}(t) = \begin{cases} 1 & t_0 \leq t \leq t'_0, \\ 0 & \text{otherwise,} \end{cases}$$

Since $\mathbf{I}_{[t_j, t_{j+1}]} \subseteq [t_j, t_{j+1}]$, S is feasible if $S(t) \cdot u_{[t_j, t_{j+1}]}(t)$ is a feasible schedule of $\mathcal{J}_{[t_j, t_{j+1}]}$ for all $1 \leq j < h$. By definition, $S(t) \cdot u_{[t_j, t_{j+1}]}(t) = \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_j, t_{j+1}]}}$ is a feasible schedule of $\mathcal{J}_{[t_j, t_{j+1}]}$ for all $1 \leq j < h$. $E(S) = \sum_{j=1}^{h-1} E(\mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_j, t_{j+1}]}})$ holds trivially from $\mathbf{I}_{[t_j, t_{j+1}]} \subseteq [t_j, t_{j+1}]$. Finally, since S is feasible, $E(S) \geq E(\mathcal{S}_{\text{opt}}^{\mathcal{J}})$. \square

The following lemma implies how an energy-optimal voltage scheduling problem can be partitioned into subproblems.

LEMMA A.3. *Let*

$$E_1 = \min \left\{ \sum_{j=1}^{k-1} E\left(\mathcal{S}_{\text{opt}}^{\mathcal{J}_{[b_j, b_{j+1}]}}\right) \mid (b_1, b_2, \dots, b_k) \text{ is a strongly blocking tuple.} \right\},$$

$$E_2 = \min \left\{ \sum_{j=1}^{h-1} E\left(\mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_j, t_{j+1}]}}\right) \mid \min R_{\mathcal{J}} = t_1 < t_2 < \dots < t_h = \max D_{\mathcal{J}}, t_j \in T_{\mathcal{J}} \right\},$$

and

$$E_3 = \min \left\{ \sum_{j=1}^{h-1} E\left(\mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_j, t_{j+1}]}}\right) \mid [t_j, t_{j+1}] \text{ is a subinterval of a strongly atomic interval for all } 1 \leq j < h \right\}.$$

Then, $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}}) = E_1 = E_2 = E_3$.

PROOF. Let

$$\mathbf{S}_1 = \left\{ \bigoplus_{j=1}^{k-1} \mathcal{S}_{\text{opt}}^{\mathcal{J}^{[b_j, b_{j+1}]}} \mid (b_1, b_2, \dots, b_k) \text{ is a strongly blocking tuple.} \right\}$$

and define \mathbf{S}_2 and \mathbf{S}_3 similarly. Then, from Lemma A.2, $E_i = \min\{E(S) \mid S \in \mathbf{S}_i\}$ for $i = 1, 2, 3$. By definition, $\mathbf{S}_1 \subseteq \mathbf{S}_3 \subseteq \mathbf{S}_2$ and consequently $E_2 \leq E_3 \leq E_1$. Furthermore, $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}}) \leq E_2$ from Lemma A.2. From Theorem 3.3 and Lemma A.1, $\mathcal{S}_{\text{opt}}^{\mathcal{J}} \in \mathbf{S}_1$. Thus, we have $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}}) \geq E_1$, which implies $E(\mathcal{S}_{\text{opt}}^{\mathcal{J}}) = E_1 = E_2 = E_3$. \square

From Lemma A.3, it is obvious that the algorithm in Figure 10 always computes an optimal voltage schedule.

A.2 Proof of the Correctness of the Algorithm in Figure 12

We start with some lemmas to prove the correctness of the algorithm. First, note that for an interval $[t, t']$, $\mathbf{I}_{\mathcal{J}_{[t, t']^w}} \subseteq \mathbf{I}_{\mathcal{J}_{[t, t']}} \subseteq [t, t']$ since $\mathcal{J}_{[t, t']^w} \subseteq \mathcal{J}_{[t, t']}$. Therefore, for a weakly blocking tuple $\mathbf{b}^w = (b_1, b_2, \dots, b_k)$, $\mathbf{I}_{\mathcal{J}_{[b_1, b_2]^w}}, \mathbf{I}_{\mathcal{J}_{[b_2, b_3]^w}}, \dots, \mathbf{I}_{\mathcal{J}_{[b_{k-1}, b_k]^w}}$ are disjoint.

LEMMA A.4. *Given a weakly blocking tuple \mathbf{b}^w , let $\mathcal{J}_{\mathbf{b}^w}^B$ represent the set of background jobs with respect to \mathbf{b}^w . Then, $\mathcal{J}_{\mathbf{b}_1^w}^B \equiv \mathcal{J}_{\mathbf{b}_2^w}^B$ for any weakly blocking tuples \mathbf{b}_1^w and \mathbf{b}_2^w .*

PROOF. Let $\mathbf{b}_1^w = (b_1, b_2, \dots, b_k)$ and $\mathbf{b}_2^w = (b'_1, b'_2, \dots, b'_k)$. Assume that $J \in \mathcal{J}_{\mathbf{b}_1^w}^B$ and $r_J \in [b_j, b_{j+1})$. From the definition of a background job, we have

$$\exists k > j + 1, d_J \geq b_{j+1}, \text{ and } p_J > \max \left\{ p_{J'} \mid J' \in \bigcup_{l=j+1}^{k-1} \mathcal{J}_{[b_l, b_{l+1}]^w} \right\}. \quad (15)$$

Suppose that $J \notin \mathcal{J}_{\mathbf{b}_2^w}^B$. From (15), we have

$$[b_{j+1}, b_{j+2}] \subseteq (b'_{j'}, b'_{j'+1}] \text{ for } r_J \in [b'_{j'}, b'_{j'+1});$$

a contradiction. So $\mathcal{J}_{\mathbf{b}_1^w}^B \subseteq \mathcal{J}_{\mathbf{b}_2^w}^B$. Similarly, we have $\mathcal{J}_{\mathbf{b}_2^w}^B \subseteq \mathcal{J}_{\mathbf{b}_1^w}^B$. \square

Lemma A.4 states that we can specify background jobs irrespective of weakly blocking tuples. For the rest of this paper, we use \mathcal{J}^B to represent the set of background jobs.

LEMMA A.5. *Given a job set \mathcal{J} and an essential $|\mathcal{J}|$ -tuple \mathbf{f} , let $\mathbf{b}_{\text{opt}}^w = (b_1, b_2, \dots, b_k)$. Then, for any weakly atomic interval $[b_j, b_{j+1}]$ ($1 \leq j < k$) and a background job J , we have the following, assuming jobs are executed under $\mathcal{S}_{\text{opt}}^{\mathcal{J}}$.*

- (a) $d_J \in [b_j, b_{j+1})$: J completes its execution by b_j .
- (b) $r_J \in [b_j, b_{j+1})$: J completes its execution by b_{j+1} .
- (c) $[b_j, b_{j+1}] \subseteq [r_J, d_J]$ executes its partial workload at constant speed.

Furthermore, for any interval $[t, t'] \subseteq [b_j, b_{j+1}]$, $\mathcal{J}_{[t, t']^w}$ is an EDF job set.

PROOF. Case (a) and Case (b) are obvious from the construction of the weakly blocking tuple $\mathbf{b}_{\text{opt}}^w$. Case (c) follows from Lemma 3.7. Finally, suppose that $\mathcal{J}_{[t, t']^w}$ is not an EDF job set. Then, we have

$$\exists J, J' \in \mathcal{J}_{[t, t']^w} \text{ s.t. } p_J > p_{J'}, d_J \in (r_{J'}, d_{J'}),$$

and the algorithm in Figure 11 selects $r_{J'}$ ($\in (b_j, b_{j+1})$) as an element of $\mathbf{b}_{\sigma_f}^w$; a contradiction. \square

From Lemma A.5, we characterize the optimal schedule in terms of weakly atomic intervals, weakly blocking tuples, and background workload.

LEMMA A.6. *Given a job set \mathcal{J} and an essential $|\mathcal{J}|$ -tuple \mathbf{f} ,*

$$\mathcal{S}_{\text{opt}}^{\mathcal{J}} \equiv \bigoplus_{j=1}^{k-1} \mathcal{S}_{\text{opt}}^{\mathcal{J}_j} \quad (16)$$

where $\mathbf{b}_{\sigma_f}^w = (b_1, b_2, \dots, b_k)$ and $\mathcal{J}_j = \mathcal{J}_{[b_j, b_{j+1}]} \cup \{J_j^b\}$ such that

$$\begin{aligned} r_{J_j^b} &= b_j, d_{J_j^b} = b_{j+1}, p_{J_j^b} = \max\{p_J | J \in \mathcal{J}_{[b_j, b_{j+1}]}\} + 1, \text{ and} \\ c_{J_j^b} &= c_j^b \text{ for some } c_j^b \geq 0. \end{aligned}$$

PROOF. From Lemma A.5, we have

$$\{\text{job}(\mathcal{J}, \mathcal{S}_{\text{opt}}^{\mathcal{J}}(t), t) | t \in [b_j, b_{j+1}]\} \equiv \mathcal{J}_{[b_j, b_{j+1}]^w} \cup \mathcal{J}' \cup \mathcal{J}''$$

where

$$\mathcal{J}' = \{J' \in \mathcal{J}^B | r_{J'} \in [b_j, b_{j+1}]\}$$

and

$$\mathcal{J}'' = \{J' \in \mathcal{J}^B | [b_j, b_{j+1}] \subseteq [r_{J'}, d_{J'}]\}.$$

From Case (b) of Lemma A.5, $\mathcal{J}_{[b_j, b_{j+1}]^w} \cup \mathcal{J}' \equiv \mathcal{J}_{[b_j, b_{j+1}]}$, and from Case (c), $\mathcal{J}'' = \{J_b\}$. So, we have

$$\mathcal{S}_{\text{opt}}^{\mathcal{J}}(t) \cdot u_{[b_j, b_{j+1}]}(t) \equiv \mathcal{S}_{\text{opt}}^{\mathcal{J}_j} \quad \text{for all } 1 \leq j < k,$$

which is equivalent to (16). \square

From Lemma A.6, the voltage scheduling problem is reduced to the problem of finding a weakly blocking tuple $\mathbf{b}^w = (b_1, b_2, \dots, b_k)$ and the amount of background workload $c_{[b_j, b_{j+1}]}^B$ for each weakly atomic interval $[b_j, b_{j+1}]$. To find the *background speed* $s_{[b_j, b_{j+1}]}^B$ instead of the amount of background workload makes it possible to exploit Lemma 3.7.

LEMMA A.7. *Given a weakly atomic interval $[t_1, t_2]$, let $\mathcal{J}' = \mathcal{J}_{[t_1, t_2]} \cup \{J^b\}$, where*

$$r_{J^b} = t_1, d_{J^b} = t_2, p_{J^b} = \max\{p_J | J \in \mathcal{J}_{[t_1, t_2]}\} + 1, \text{ and } c_{J^b} = c_{[t_1, t_2]}^B (> 0),$$

and let $s_{[t_1, t_2]}^B$ be the constant speed of J^b under $\mathcal{S}_{\text{opt}}^{\mathcal{J}'}$. Then,

$$\mathcal{S}_{\text{opt}}^{\mathcal{J}'}(t) = \begin{cases} \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_1, t_2]}}(t) & t \text{ s.t. } \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_1, t_2]}}(t) > s_{[t_1, t_2]}^B, \\ s_{[t_1, t_2]}^B & t \text{ s.t. } \mathcal{S}_{\text{opt}}^{\mathcal{J}_{[t_1, t_2]}}(t) \leq s_{[t_1, t_2]}^B. \end{cases}$$

Furthermore, $s_{[t_1, t_2]}^B$ strictly increases as $c_{[t_1, t_2]}^B$ increases, and vice versa.

PROOF. From Lemmas A.5 and A.6, both $\mathcal{J}_{[t_1, t_2]}$ and \mathcal{J}' follow the EDF priority and their optimal voltage schedules are obtained by Yao's algorithm

[Yao et al. 1995]. For an interval $[t'_1, t'_2] \subset [t_1, t_2]$ such that $S_{\text{opt}}^{\mathcal{J}_{[t'_1, t'_2]}}(t) > s_{[t'_1, t'_2]}^{\text{B}}$, Yao's algorithm selects the same speed for $S_{\text{opt}}^{\mathcal{J}'}(t)$. For the other intervals, $S_{\text{opt}}^{\mathcal{J}'}(t) = s_{J^{\text{b}}}$ since $[t_1, t_2] \subseteq [r_{J^{\text{b}}}, d_{J^{\text{b}}}]$.

Because $W(S_{\text{opt}}^{\mathcal{J}'}, [t_1, t_2])$ strictly increases as $s_{[t_1, t_2]}^{\text{B}}$ increases, and $c_{[t_1, t_2]}^{\text{B}} = W(S_{\text{opt}}^{\mathcal{J}'}, [t_1, t_2]) - W(S_{\text{opt}}^{\mathcal{J}_{[t_1, t_2]}}, [t_1, t_2])$, $c_{[t_1, t_2]}^{\text{B}}$ increases as $s_{[t_1, t_2]}^{\text{B}}$ increases. Hence, it follows that $s_{[t_1, t_2]}^{\text{B}}$ increases as $c_{[t_1, t_2]}^{\text{B}}$ increases (and vice versa). \square

Definition A.8. Given a job set \mathcal{J} and background workload c , the job set \mathcal{J} with background workload c is defined as

$$\begin{aligned} \mathcal{J}[c] &\stackrel{\text{def}}{=} \mathcal{J} \cup \{J^{\text{b}}\} \quad \text{where } r_{J^{\text{b}}} = R_{\mathcal{J}}, d_{J^{\text{b}}} = D_{\mathcal{J}}, \\ p_{J^{\text{b}}} &= \max\{p_J | J \in \mathcal{J}\} + 1, \text{ and } c_{J^{\text{b}}} = c. \end{aligned}$$

Furthermore, given a job set $\mathcal{J}[c]$, the constant speed of background workload under $S_{\text{opt}}^{\mathcal{J}[c]}$ is called a *background speed* of $\mathcal{J}[c]$ and is denoted by $BS(\mathcal{J}, c)$.

The following lemma is an extension of Lemma A.7 for arbitrary intervals.

LEMMA A.9. *Given a job set $\mathcal{J}[c]$*

$$\begin{aligned} S_{\text{opt}}^{\mathcal{J}[c]} &\equiv \bigoplus_{j=1}^{k-1} S_{\text{opt}}^{\mathcal{J}_{[b_j, b_{j+1}]}, c_j} \quad \text{for } b_1, \dots, b_k \in \mathcal{T}_{\mathcal{J}}, b_1 < \dots < b_k \text{ such that} \\ c &= \sum_{j=1}^{k-1} c_j \text{ and } BS(\mathcal{J}_{[b_j, b_{j+1}]}, c_j) = BS(\mathcal{J}_{[b_{j'}, b_{j'+1}]}, c_{j'}) \text{ for all } 1 \leq j \neq j' < c_j. \end{aligned}$$

PROOF. Directly from Lemmas A.6 and 3.7. \square

Along with Lemma A.9, the following lemma implies how the problem can be reduced to a dynamic programming formulation.

LEMMA A.10. *Given $t_i, t_j, t_m \in \mathcal{T}_{\mathcal{J}}$, where $t_i < t_m \leq t_j$, let*

$$\mathcal{J}_{[t_i, t_j]}^{\text{B}w} = \mathcal{J}_{[t_i, t_j]}^w \cup \{J \in \mathcal{J}^{\text{B}} | [r_J, d_J] \subseteq [t_i, t_j]\}$$

and

$$c_{[t_i, t_m]}^{\text{B}} = C(\{J \in \mathcal{J}^{\text{B}} | [r_J, d_J] \subseteq [t_i, t_m] \} \wedge d_J \in [t_m, t_j]),$$

and let $S_{\text{opt}}^{[t_i, t_j]}$ represent $S_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}^{\text{B}w}}$. Then,

$$S_{\text{opt}}^{[t_i, t_j]} \in \left\{ S_{\text{opt}}^{\mathcal{J}_{[t_i, t_m]}^w, c_{[t_i, t_m]}^{\text{B}}} \oplus S_{\text{opt}}^{[t_m, t_j]} \mid S_{\text{opt}}^{\mathcal{J}_{[t_i, t_m]}^w, c_{[t_i, t_m]}^{\text{B}}} \text{ is feasible for } \mathcal{J}_{[t_i, t_m]}^{\text{B}w} \right\}.$$

PROOF. If all the jobs in $\{J \in \mathcal{J}^{\text{B}} | [r_J, d_J] \subseteq [t_i, t_j]\}$ run at the same speed under $S_{\text{opt}}^{[t_i, t_j]}$, $S_{\text{opt}}^{[t_i, t_j]} \equiv S_{\text{opt}}^{\mathcal{J}_{[t_i, t_j]}^w, c_{[t_i, t_j]}^{\text{B}}}$. Otherwise, there must exist $t_m \in \mathcal{T}_{\{J \in \mathcal{J}^{\text{B}} | [r_J, d_J] \subseteq [t_i, t_j]\}} (\subseteq \mathcal{T}_{\mathcal{J}})$ such that all the jobs in $\{J \in \mathcal{J}^{\text{B}} | [r_J, d_J] \subseteq [t_i, t_m] \wedge d_J \in [t_m, t_j]\}$ finish their executions by t_m with the same constant speed and all the jobs in $\{J \in \mathcal{J}^{\text{B}} | [r_J, d_J] \subseteq [t_m, t_j]\}$ are not executed before t_m under $S_{\text{opt}}^{[t_i, t_j]}$.

Therefore, we have $S_{\text{opt}}^{[t_i, t_j]} \equiv S_{\text{opt}}^{\mathcal{J}_{[t_i, t_m]}^w, c_{[t_i, t_m]}^{\text{B}}} \oplus S_{\text{opt}}^{[t_m, t_j]}$, where $S_{\text{opt}}^{\mathcal{J}_{[t_i, t_m]}^w, c_{[t_i, t_m]}^{\text{B}}}$ is feasible for $\mathcal{J}_{[t_i, t_m]}^{\text{B}w}$. \square

COROLLARY A.11. Let $E_{\text{opt}}^{[t_i, t_j]}$ denote $E(S_{\text{opt}}^{[t_i, t_j]})$ where $S_{\text{opt}}^{[t_i, t_j]}$ is defined as in Lemma A.10. Then,

$$E_{\text{opt}}^{[t_i, t_j]} = \min \left(\left\{ E \left(S_{\text{opt}}^{\mathcal{J}_{[t_i, t_m]^w} [c_{[t_i, t_m]}^B]} \right) + E_{\text{opt}}^{[t_m, t_j]} \mid t_m \in \mathcal{T}_{\mathcal{J}}, t_i < t_m < t_k, \right. \right. \\ \left. \left. S_{\text{opt}}^{\mathcal{J}_{[t_i, t_m]^w} [c_{[t_i, t_m]}^B]} \text{ is feasible for } \mathcal{J}_{[t_i, t_m]^w}^B \right\} \right).$$

The correctness of the algorithm in Figure 12 directly follows from Lemma A.10 and Corollary A.11.

ACKNOWLEDGMENTS

We would like to thank Gang Quan and the anonymous referees for their helpful comments and suggestions.

REFERENCES

- AYDIN, H., MELHEM, R., MOSSE, D., AND ALVAREZ, P. M. 2001. Dynamic and aggressive scheduling techniques for power-aware real-time systems. In *Proceedings of Real-Time Systems Symposium*.
- GAREY, M. AND JOHNSON, D. 1979. *Computers and Intractability*. W.H. Freeman and Company: San Francisco, CA.
- GRUIAN, F. 2001. Hard real-time scheduling for low-energy using stochastic data and DVS processors. In *Proceedings of International Symposium on Low Power Electronics and Design*, 46–51.
- HONG, I., QU, G., POTKONJAK, M., AND SRIVASTAVA, M. B. 1998. Synthesis techniques for low-power hard real-time systems on variable voltage processors. In *Proceedings of Real-Time Systems Symposium*, 178–187.
- KIM, W., KIM, J., AND MIN, S. L. 2002. A dynamic voltage scaling algorithm for dynamic-priority hard real-time systems using slack time analysis. In *Proceedings of Design, Automation and Test in Europe*.
- KIM, N., RYU, M., HONG, S., SAKSENA, M., CHOI, C., AND SHIN, H. 1996. Visual assessment of a real-time system design: A case study on a CNC controller. In *Proceedings of Real-Time Systems Symposium*, 300–310.
- LIU, W.-S. 2000. *Real-Time Systems*. Prentice Hall, Englewood, Cliffs, NJ.
- LOCKE, C., VOGEL, D., AND MESLER, T. 1991. Building a predictable avionics platform in Ada: A case study. In *Proceedings of Real-Time Systems Symposium*.
- PILLAI, P. AND SHIN, K. G. 2001. Real-time dynamic voltage scaling for low-power embedded operating systems. In *Proceedings of ACM Symposium on Operating Systems Principles*.
- QUAN, G. AND HU, X. 2001. Energy efficient fixed-priority scheduling for real-time systems on variable voltage processors. In *Proceedings of Design Automation Conference*, 828–833.
- QUAN, G. AND HU, X. 2002. An optimal voltage schedule for real-time systems on a variable voltage processor. In *Proceedings of Design, Automation and Test in Europe*.
- SAHNI, S. 1976. Algorithms for scheduling independent tasks. *J. ACM* 23, 116–127.
- SAKURAI, T. AND NEWTON, A. 1990. Alpha-power law MOSFET model and its application to CMOS inverter delay and other formulas. *IEEE Journal of Solid State Circuits* 25, 2, 584–594.
- SHIN, D., KIM, J., AND LEE, S. 2001. Low-energy intra-task voltage scheduling using static timing analysis. In *Proceedings of the 38th Design Automation Conference*.
- SHIN, Y. AND CHOI, K. 1999. Power conscious fixed priority scheduling for hard real-time systems. In *Proceedings of Design Automation Conference*, 134–139.
- SHIN, Y., CHOI, K., AND SAKURAI, T. 2000. Power optimization of real-time embedded systems on variable speed processors. In *Proceedings of International Conference on Computer-Aided Design*, 365–368.

WÖEGINGER, G. J. 1999. When does a dynamic programming formulation guarantee the existence of an FPTAS? In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, 820–829.

YAO, F., DEMERS, A., AND SHENKER, S. 1995. A scheduling model for reduced CPU energy. In *Proceedings of IEEE Annual Foundations of Computer Science*, 374–382.

Received March 2002; accepted July 2002