

## Heuristics for Static Voltage Scheduling Algorithms on Battery-Powered DVS Systems

Tetsuo Yokoyama, Gang Zeng, Hiroyuki Tomiyama, and Hiroaki Takada  
 Graduate School of Information Science, Nagoya University  
 Furo-cho, Chikusa-ku, Nagoya 464-8603, Japan  
 {yokoyama,sogo,tomiyama,hiro}@ertl.jp

### Abstract

*The principles for good design of battery-aware voltage scheduling algorithms for both aperiodic and periodic task sets on dynamic voltage scaling (DVS) systems are presented. The proposed algorithms are based on greedy heuristics suggested by several battery characteristics and Lagrange multipliers. To construct the proposed algorithms, we use the battery characteristics in the early stage of scheduling more properly. As a consequence, the proposed algorithms show superior results on synthetic examples of periodic and aperiodic tasks from the task sets which are excerpted from the comparative work, on uni-processor platforms. Especially, for some large task sets, the proposed algorithms enable previously unschedulable task sets due to battery exhaustion to be schedulable.*

### 1. Introduction

Increase in battery capacity and improvements of battery utilization have proved to be indispensable due to higher demand for functionality resulting in drastic increase of energy consumption. On the other hand, the energy density of the battery, though gradually improved, already reached a half or one third of the theoretical limits [1]. It is therefore of great importance to reduce the energy consumption by controlling software in the battery-powered embedded systems. The challenge is to appropriately handle the non-trivial battery characteristics, such as *recovery* and *non-linearity* of battery capacity which are dependent on the current load history [1]–[3]. A real battery recovers its charge when it is idle. The current load affecting the total capacity of batteries is not uniform but rather depends on the load history. Once the voltage of the battery reaches its threshold, a battery becomes *exhausted*. Then, its current is unavailable and

never recovers without an external power supply. To use battery capacity to the greatest extent, it is impossible to apply existing energy optimization technics as they are. Therefore, the aforementioned battery characteristics must be reflected.

Battery characteristics have been widely studied [4]. Although computationally expensive low-level electrochemical models (*e.g.*, Dualfoil [5]) are accurate, high-level battery models provide reasonably accurate approximation gained by lightweight computation [6]–[11]. Based on those models, a number of battery-aware voltage scheduling algorithms on dynamic voltage scaling (DVS) systems are proposed [2], [3], [12].

Specifically, incorporating battery properties, such as recovery effect and non-linearity, Rakhmatov and Vrudhula [6], [7] developed a high-level analytical battery model with only two configuration parameters for each battery instance. Accuracy of their model was confirmed by the low-level electrochemical simulation Dualfoil [5] within approximately 5 % error according to their report. Rakhmatov and Vrudhula [2] specified various important properties of their mathematically formulated cost function, and several efficient static battery-aware voltage scheduling algorithms. Chowdhury and Chakrabarti [3] identified several insightful battery properties and extended the work of Rakhmatov, Vrudhula, and Chakrabarti's [2], [13]. They proposed battery-aware voltage scheduling algorithms not only for periodic tasks on uniprocessor platform but also for both aperiodic and periodic tasks on both uni- and multi-processor platforms.

Their improvement relies on heuristics derived from battery characteristics, such as the steepest profile and non-increasing ordering in the early stage of scheduling. In this paper, we investigate the implications caused by battery properties, which are identified in the aforementioned previous work [2], [3]. As the result of the proper investigation of those properties, we propose new static voltage scheduling algorithms for the battery-powered embedded

systems, based on greedy heuristics suggested by several battery properties and Lagrange multipliers. We target uniprocessor systems. Our method shows better results in the aperiodic task sets of time varying load used in the periodic task sets on uniprocessor systems. For periodic tasks, we need a special care, not necessarily considered in energy minimization; the optimal profiles are not always the same among different hyperperiods.

This paper is organized as follows. First, we present the non-ideal battery characteristic using a motivating example (Section 2). After representing the system and battery models we rely on (Section 3), we propose our voltage scheduling algorithms (Section 4). The proposed algorithms are evaluated through comprehensive experiments (Section 5). Finally, we conclude with a few remarks (Section 6).

## 2. Motivation

Fig. 1 represents four profiles of two identical tasks (current 250 mA, duration 2 min in the highest voltage setting, deadline 6 min) on a battery-powered uniprocessor system, in which processor speed and power change continuously.

Fig. 1(a) shows the increase current profile. Since the idle time increases the nominal residual charge in batteries, the later we measure the nominal battery capacity, the higher value we obtain. The consumed capacity is measured by the objective function  $\sigma_B$  (see Section 3 in detail). Intuitively, the higher  $\sigma_B$  stands for the smaller residual charge available in batteries at observation time  $B$ . Once  $\sigma_B$  reaches some threshold (denoted as  $\alpha$  in this paper), the batteries become exhausted; The batteries are inactive and no longer recoverable without external power supply. The objective function at time 6 min ( $\sigma_6$ ) is 1566 mA·min and it decreases to 1115 mA·min at time 12 min ( $\sigma_{12}$ ). Each of two objective functions is the worst in the four cases.

For battery-unaware scheduling, it is well-known that only a single processor speed is sufficient to obtain the optimal profile [14]. The time between the end of task execution and deadline is called *slack time*. If we distribute the slack time equally among two identical tasks, we obtain level current profile (Fig. 1(b), cf. [12]), which is an optimal profile in terms of the amount of energy dissipated from batteries. However, due to non-linearity of the objective function, it is not always an optimal profile as far as optimization of battery residual charge is concerned. As we will see later on, this profile is optimal in case of ideal battery.

The decrease current profile results in better battery performance compared to increase current profile. Fig. 1(c) is an optimal profile minimizing the objective function  $\sigma_{12}$ .

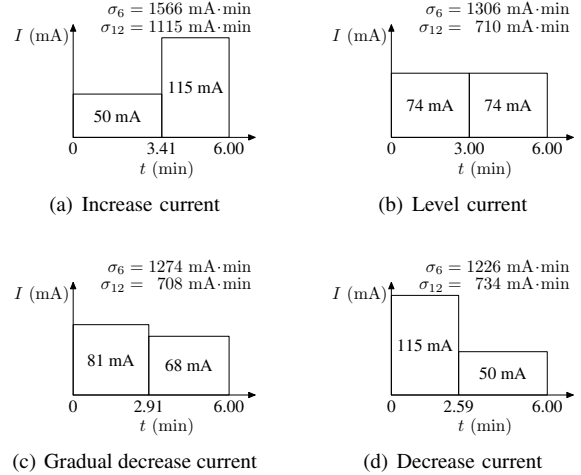


Figure 1. Motivating example.

INPUT:  $S_V, S_\phi, S_I, S_r, S_\Delta, S_d, G, B, \alpha, \beta, p$

OUTPUT:  $S_V^*, S_t$

OBJECTIVE:

$$\text{minimize } \sigma_B \text{ s.t. } \sigma_B = \sum_{k=0}^{n-1} I_k^* F(B, t_k, t_k + \Delta_k^*, \beta)$$

CONSTRAINTS:

- 1)  $r_k \leq t_k$
- 2)  $\forall t. \#(t \in [t_k, t_k + \Delta_k^*]) \leq p$
- 3)  $k'$  depends on  $k$  in  $G \Rightarrow t_k + \Delta_k^* \leq t_{k'}$
- 4)  $\forall k. t_k + \Delta_k^* \leq d_k$  and  $\forall k. t_k + \Delta_k^* \leq B$
- 5)  $\forall t \leq B. \sum_{k=0}^{n-1} I_k^* F(t, \min\{t, t_k\}, \min\{t, t_k + \Delta_k^*\}, \beta) \leq \alpha$

Figure 2. Battery-aware voltage scheduling problem.

Despite higher energy consumption in the level current approach (Fig. 1(b)), the battery charges  $\sigma_6, \sigma_{12}$  are reduced.

Fig. 1(d) is an optimal profile minimizing  $\sigma_6$ . This decrease current is obtained by swapping the two tasks in the increase current profile in Fig. 1(a). Those two profiles consume exactly the same amount of energy, but the residual available charges differ (34.2% improvement with respect to  $\sigma_{12}$ ). This implies the importance of the task order (*i.e.*, the *history of charge*) in the battery-aware voltage scheduling. The gradient of this task load is steeper than in the gradual decrease current profile (Fig. 1(c)), and this results in the worse effect on  $\sigma_{12}$ . Namely, the steepest non-increasing load current profile is *not* always optimal. Hence, heuristics for choosing the steepest non-increasing load current profile [3] is not always the best choice.

This example becomes an intratask voltage scheduling problem [15], if we regard two identical tasks as one

piece of task and have opportunity to change the speed when a half of the task is finished. In the battery-aware voltage scheduling, even when each task consumes power uniformly, there is still room to optimize the total available capacity in the battery by switching voltages during the task execution. The difference of  $\sigma_6$  between Fig. 1(b) and Fig. 1(d) shows not negligible improvement (6.1 %).

Energy minimization is *not* exactly equivalent to battery optimization. In summary, when considering the battery-aware voltage scheduling, we need to pay attention to the recovery effect and the history of charge, and we should not rely too much on the steepest non-increasing profile heuristics to obtain an efficient profile.

The importance of the research on the better battery utilization is also reinforced by the fact that it is somewhat independent of the energy optimization of the other parts of the systems. For example, if a subsystem, which consumes 10 % of the total energy, reduces the energy consumption by 50 %, only 5 % is reduced in total. However, if energy, which remained in a battery after battery failure, is used in the system it is exactly equal to the increase of the available energy for performing system.

### 3. Preliminaries

This section explains the assumptions used in the following sections, in order to make this paper self-contained.

#### 3.1. System Configuration

We assume DVS-enabled uniprocessors. For the sake of simplicity, the DC-DC conversion efficiency is assumed to be 100 %. The ratio of the initial task duration  $\Delta$  and the new task duration  $\Delta^*$  after scaling the task voltage  $V_{dd}$  down by factor  $s$  (*i.e.*,  $V_{dd}/s$ ) is

$$\frac{\Delta^*}{\Delta} = s \left( 1 + \frac{2(s-1)V_{th}}{V_{dd} - V_{th}} \right) \quad (1)$$

with  $V_{th}$  a threshold voltage. The battery current  $I_{batt}$  scales by  $s^3$ , *i.e.*,

$$\frac{I_{batt}^*}{I_{batt}} = \frac{1}{s^3} \quad (2)$$

with  $I_{batt}^*$  a battery current after scaling the task voltage.

Our target processor is compatible with the StrongARM SA-1100 microprocessor [16]. The operating voltage ranges over set  $\{3.3, 3.0, 2.7, 2.5, 2.0\}$ . The threshold voltage  $V_{th}$  is 0.4 V. For the sake of simplicity, we assume no time and energy overhead due to the change of voltage setting.

In most modern embedded systems, dynamic energy consumption is dominant. Therefore, in this paper, static energy consumption is ignored, although it is one of major

concerns in the near future. All the assumptions above are consistent with the work on battery-conscious voltage scheduling [3] for comparison purpose.

#### 3.2. Battery Model

There has been no single model capturing the behavior of both charge and voltage at the same time. We focus on charge sensitive model and ignore voltage sensitive models [5], [17]. As a model for variable load, we use the high-level analytical model of Rakhmatov and Vrudhula's [6], [7]. The load on battery  $i(t)$  and battery life time  $L$  are related by equation

$$\alpha = \int_0^L i(t)dt + \int_0^L i(t) \left( 2 \sum_{m=1}^{\infty} e^{-\beta^2 m^2 (L-t)} \right) dt \quad (3)$$

with  $\alpha$  (mA·min) and  $\beta$  ( $\text{min}^{-1/2}$ ) being constants, uniquely determined for each battery. Intuitively,  $\alpha$  represents the battery's theoretical capacity and  $\beta$  the recovery rate.

The first term on the right hand side sums up the discharged capacity from the battery from time 0 to  $L$ . The second term represents the residual charge in the battery, unavailable at time  $L$ . It should be noted that once battery voltage becomes lower than some threshold, the residual charge in the battery cannot be used any more. If the second term on the right hand side is negligible (*e.g.*,  $\beta \rightarrow \infty$ ), the battery behavior is nearly optimal. If  $\alpha$  is significantly large, we do not have to consider the battery failures (exhaustion). By means of insertion of the period of no load, *i.e.*,  $i = 0$ , or low load, the battery life  $L$  increases; the battery at rest recovers its charge. Let  $\tau_k$  ( $k$  in short) be tasks executed during the period of  $\Delta_k$  starting at time  $t_k$ . For brevity, we estimate the load of each task  $k$  on the battery to be a constant  $I_k$  in the highest voltage setting. If the tasks are sequentially computed and the battery life ends at task  $u$ , the battery capacity equation computed with continuous current function (3) becomes the discrete equation [2], [6]

$$\alpha = \sum_{k \in S_u} I_k F(L, t_k, t_k + \Delta_k, \beta) + I_u F(L, t_u, L, \beta) \quad (4)$$

with a set  $S_u$  consisting of the tasks executed before task  $u$  and auxiliary function

$$F(T, s, f, \beta) = f - s + 2 \sum_{m=1}^{m_{\max}} \frac{e^{-\beta^2 m^2 (T-f)} - e^{-\beta^2 m^2 (T-s)}}{\beta^2 m^2} \quad (5)$$

in which  $s$  represents the start time,  $f$  the finishing time, and  $T$  the observation time. If  $t \notin [t_k, t_k + \Delta_k)$  for all  $\tau_k \in S_u + \{\tau_u\}$  in the uniprocessor systems, then  $t$  is in the idle time. Since in theory  $m_{\max}$  is infinite, the number of terms in the sum of infinite series provides a tradeoff

between the accuracy and the amount of computation. In [6] and [7], a graph ranging  $1 \leq \beta^2 L \leq 10^2$  shows that the sums of the first 10 and 100 000 terms create negligible difference. This fact implies that the first 10 terms are a good approximation.

We justify this observation analytically. Since the formula (summands) under  $\Sigma$  notation in  $F$  is monotonically decreasing with  $m$  and  $T$ , by using the solution of Basel problem,<sup>1</sup> we obtain, for any  $m_{\max}$

$$\sum_{m=1}^{m_{\max}} \frac{e^{-\beta^2 m^2 (T-f)} - e^{-\beta^2 m^2 (T-s)}}{m^2} < e^{-\beta^2 (T-f)} \frac{\pi^2}{6}. \quad (6)$$

Given  $m_{\max}$  and the upper bound of  $f-s$ , the upper bound of error is obtained. For example, assuming  $f-s \geq 1$  min,  $m_{\max} = 10$ ,  $\beta = 0.637 \text{ min}^{-1/2}$ ,  $T-f = 10$  min, error is bounded by  $1.03 \times 10^{-3}$ . The sum of the first 10 terms has approximately at most 0.2 % error. Therefore, we use  $m_{\max} = 10$  in the formula of  $F$  (5).

We focus on the task sets of the middle duration range (0.5 min to 20 min). This is because firstly the load frequencies higher than 1 Hz can be filtered owing to the late response of the battery device [1] and battery-charge optimization is not effective for very fine-grained ( $< 10$  ms) tasks [18]. Secondly, for very coarse-grained ( $> 30$  min) tasks, battery-aware voltage scheduling is not much superior to energy optimal scheduling [18]; Battery optimization is almost equivalent to energy minimization in this range. Not a few tasks of PDA, such as playing music and movie, and controllers, reside in this middle time range.

In this paper, we do not consider self-discharge mechanisms, aging caused by discharge/recharge repetition, and dependence on temperature, since this model does not take them into account

### 3.3. Cost Function

A profile of  $n$  tasks consists of a set of the current  $I_k$ , the starting time  $t_k$ , and the duration  $\Delta_k$ . For a given profile of  $n$  tasks and the observation time  $B$ , the battery-aware cost function [7]

$$\sigma_B = \sum_{k=0}^{n-1} I_k F(B, t_k, t_k + \Delta_k, \beta) \quad (7)$$

is to be minimized. The subscript  $B$  of  $\sigma$  is omitted if it is insignificant or obvious from its context. Intuitively,  $\sigma_B$  models a measure of the residual capacity in the battery available to use at time  $B$ . When battery parameter  $\beta$  or the observing time  $B$  becomes large ( $\beta \rightarrow \infty$ ,  $B \rightarrow \infty$ ),

<sup>1</sup> Euler solved the Basel problem and obtained the formula  $\sum_{m=1}^{\infty} 1/m^2 = \pi^2/6$ .

the third term on the right hand side in the formula of  $F$  (5) disappears, and the battery becomes ideal ( $\lim_{\beta \rightarrow \infty} \sigma_B = \sigma_{\infty}$ ).

To make the value  $\sigma$  meaningful, two conditions must be satisfied. Firstly, all tasks must terminate before observation time  $B$ :

$$\forall k. t_k + \Delta_k \leq B. \quad (8)$$

Secondly, the battery must satisfy the *endurance constraint*; the battery must not be exhausted before  $B$ :

$$\forall t \leq B. \alpha \geq \sum_{k=0}^{n-1} I_k F(t, \min\{t, t_k\}, \min\{t, t_k + \Delta_k\}, \beta). \quad (9)$$

It should be noted that if the values of the second and third arguments of  $F$  are equal,  $F$  becomes zero.

For comparison purpose, we use the same battery parameters in all profiles (including motivating examples in Section 2) as in [2], [3], [7] ( $\alpha = 40\ 375 \text{ mA}\cdot\text{min}$ ,  $\beta = 0.273 \text{ min}^{-1/2}$ ).

We will use the largest deadline of given tasks for the observation time  $B$ , instead of the end time of each profile, which was used in the previous work [2], [3]. As a result, even if all tasks finished earlier than the deadline, it is not always disadvantageous because of the usage of the remaining time for recovery.

## 4. Battery-Aware Voltage Scheduling

A voltage scheduling problem for the battery-powered DVS system can be formulated as a non-linear optimization problem (Fig. 2). In this paper, we do not consider task preemption. The input consists of six ordered sets of voltages  $S_V$  on which a system operates, frequency  $S_\phi$ , current  $S_I$ , release time  $S_r$ , duration  $S_\Delta$ , and deadline  $S_d$ , task dependency graph  $G$ , observation time  $B$ , battery parameters  $\alpha$ ,  $\beta$ , specified in Section 3.3, and the number of processors  $p$ . The output consists of two ordered sets representing scheduled voltage  $S_{V^*}$  and start time  $S_t$  for each task.  $S_{V^*}$  is uniquely determined by scheduled current  $I^*$ s and/or scheduled duration  $\Delta^*$ s.

The objective function to be minimized is  $\sigma_B$  in the battery-aware cost function (7). Five constraints are imposed: 1) release time, 2) number of processor, 3) dependency, 4) deadline, and 5) current load endurance. First, release time of each task must be smaller or equal to each starting time. Second, the number of processes running simultaneously must be smaller or equal to the number of processors. Third, the task dependency represented by  $G$  must be preserved. Fourth, the deadline  $d_k$  for each task  $k$  should be met, and the length of profile must be smaller than observation time  $B$ . Fifth, the battery must survive all the tasks without battery exhaustion until observation time  $B$ .

A voltage scheduling problem is NP-hard, even if tasks have the fixed-priorities [19]. Therefore, the efficient and effective heuristics is needed. The energy minimization problem is in general an instance of the battery-aware energy optimization problem, since the latter becomes the former when  $\alpha, \beta \rightarrow \infty$ .

We will add the following assumptions and make the problem simpler. For simplicity we do not consider intra-task voltage scheduling [15], *i.e.*, the power and performance are uniform within a single task. The arrival time, the deadline, the current, and the dependence relation are known in advance before execution. These conditions are not too restrictive in the embedded systems.

Our major idea is simple greedy heuristics. The objective function  $\sigma$  reaches its minimum if the equation

$$\frac{\partial \sigma}{\partial \Delta_1} = \frac{\partial \sigma}{\partial \Delta_2} = \dots = \frac{\partial \sigma}{\partial \Delta_n} \quad (10)$$

is satisfied. This equation is obtained by the Lagrange multiplier method. Here, we assume there is no slack time between tasks. The performance of the task will be decreased, so that the energy increase is the most effective. The effectiveness of the energy increase is measured by the decreased cost per the decreased duration  $\partial \sigma / \partial \Delta_k$ . Since we focus on discrete DVS and configurations, the effectiveness is measured by

$$\frac{\sigma^* - \sigma}{\Delta_k^* - \Delta_k}. \quad (11)$$

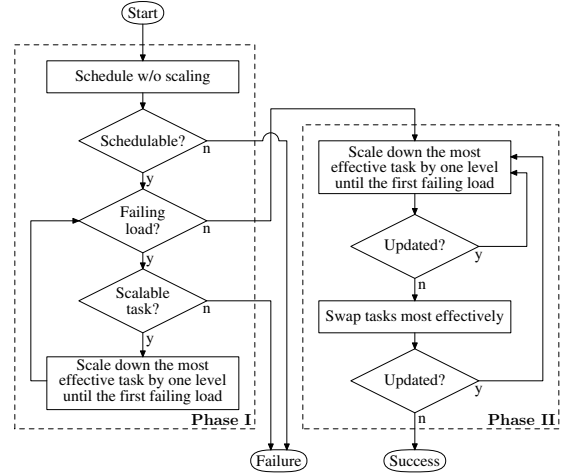
It should be noted that this expression does not assume any specific function  $\sigma$ .

#### 4.1. Voltage Scheduling for Uniprocessor Systems

The proposed algorithm consists of two phases (Fig. 3): I) obtain a feasible solution and II) distribute slack time.

In Phase I, battery-unaware scheduling algorithm without voltage scaling is used. In this paper, we use the earliest deadline first (EDF) algorithm, but this choice is not essential. The power is scaled down starting from the highest power initial solution, which hopefully satisfies the deadline constraints. When it does not satisfy the endurance constraint, our scheduling algorithm returns “Failure”.

To repair the battery failure, we repeatedly scale down the speed of the failed tasks or the tasks appeared before them, in such a way that reduction of their speed by one level results in the greatest value using the discrete voltage downscaling effectiveness measure (11) within timing constraints (greedy choice). It should be noted that the downscaled task is not always the failed task unlike in case of the scheduling algorithms in [3], where even the failed task is not assigned the lowest voltage. If the



**Figure 3. Battery-aware uniprocessor voltage scheduling algorithm.**

task is not a failing task, the voltage downscaling is not guaranteed to be superior to the insertion of the battery idle time. But such slack time tends to be relatively larger compared to our time range. Therefore, as in [3] and unlike in [2], [13], we do not consider the insertion of the idle time. If the effectiveness of DVS evaluated by the discrete voltage downscaling effectiveness measure (11) is negative, the voltage of the task is not scaled down. This case never occurs if assumption

$$V \geq V' \Rightarrow \Delta(V) \leq \Delta(V') \text{ and } I(V)\Delta(V) \geq I(V')\Delta(V') \quad (12)$$

holds at any moment. Strictly speaking in our assumption, due to duration scaling equation (1), it does not always hold, but the effect is limited and can be ignored when designing scheduling algorithm.

In Phase II, we repeatedly use the available slack time by scaling down speeds of tasks to achieve the most effective decrease of the cost with respect to the discrete voltage downscaling effectiveness measure (11). Next, we swap the task order if the result has lower cost without violating the deadline constraints. We repeat this phase until no tasks can be swapped. Consideration of batteries is a necessary condition for this new method of swapping tasks to optimize the energy efficiency. It should be noted that, if all tasks have the same release time and deadline, the scheduled subprofiles are always placed in the non-increasing order [2], [3].

Generally, the energy efficiency is achieved only if  $\sigma$  has an asymptotic lower bound being the function proportional to an exponential function of  $\Delta$ , in which exponent  $x$  is greater than one, *i.e.*,  $\sigma > \Delta^x (x > 1)$ . If  $\sigma$  is monotonic with respect to  $\Delta$ , scaling down the voltage to the next level is more effective than to any other levels.

This observation justifies the reduction of the speed of the most effective task by one level at each time.

The proposed algorithm depends on heuristics and thus it is not guaranteed to return the optimal solution, as we will see in the next example (Fig. 4). Nevertheless experiments in Section 5 show that for the same task set proposed algorithm results in better schedules when comparing to previous works [3].

We do not assume continuity and differentiability of the objective function, and we can apply our method to the practical objective functions, most of which are neither continuous nor differentiable, especially when we consider memory accesses and peripherals.

Table 1 describes an initial task set of aperiodic tasks. Fig. 4(a) shows the load profile given by EDF scheduling algorithm in the highest voltage setting. The value of objective function at infinite time  $\sigma_\infty$  is 28.0 % smaller than its value at time 38 min, *i.e.*,  $\sigma_{38}$ . The difference shows the theoretical maximum bound of recovery. This profile returns the worst result among four profiles given in Figure 4. It should be noted that to schedule the tasks we used the objective function at time 38 min.

Fig. 4(b) shows the load profile achieved by two existing approaches, *i.e.*, non-increasing [3] and exclusive down scaling [13]. These two accidentally result in the same profile. Both algorithms assign all tasks to the highest available voltage, schedule tasks by the battery-unaware algorithm, recover the failure if any occurs by downscaling the task, and repeatedly distribute the remaining slack to preceding tasks whose voltage can be lowered (as the result, the profile becomes steep) without violating deadline constraints.

Fig. 4(c) shows the load profile achieved by the proposed algorithm. Unlike two previously presented algorithms, our algorithm does not always return the non-increasing profile, as this figure shows. However, both objective functions  $\sigma_{38}$  and  $\sigma_\infty$  show the improved values.

Fig. 4(d) shows the load profile by the exhaustive search, resulting in the optimal solution at time 38 min. While the proposed algorithm is not optimal, even when there are no task dependencies, objective function obtained in our approach differs from the optimal solution by at most 4 %. Note that the optimal solution was obtained by brute-force search. The optimal solution does not always result in a non-increasing order. The simultaneous reduction of the objective functions at different observation times is by no means inevitable. In fact, the value of  $\sigma_\infty$  in Fig. 4(c) is more optimal than one in Fig. 4(d).

It should be noted that, for the task set in Table 1, proposed algorithm returns more efficient profile than the previously proposed algorithms. This result implies that, on the contrary to the intuition described in [3], the non-increasing scheduling algorithm is not optimal for the case

**Table 1. Initial task specifications**

| Task # | Duration (min) | Deadline (min) | Current (mA) |
|--------|----------------|----------------|--------------|
| 1      | 7              | 18             | 650          |
| 2      | 5              | 10             | 800          |
| 3      | 8              | 26             | 400          |
| 4      | 10             | 38             | 380          |

**Table 2. Initial task specifications of periodic tasks**

| Task # | Duration (min) | Deadline (min) | Current (mA) |
|--------|----------------|----------------|--------------|
| 1      | 0.5            | 2              | 250          |
| 2      | 0.2            | 4              | 100          |
| 3      | 1.0            | 6              | 500          |

when there is no task dependency.

## 4.2. Scheduling Periodic Tasks

The least common multiple (LCM) of the periods of all periodic tasks is called a *hyperperiod*. In a conventional voltage scheduling problem, the simple repetition of the optimal profile for hyperperiod is optimal. However, in a battery-aware voltage scheduling, it is *not* true due to the non-linearity of the objective function (see the experimental result in Section 5). For different hyperperiods, the optimal profile is not always the same.

Fig. 5 illustrates the comparison of slack utilization by three algorithms, for the periodic tasks specified in Table 2. Profile shown in Fig. 5(a) is generated by the level current algorithm [12], which aims to reduce the peak power. Fig. 5(b) is substantially improved due to the non-increasing task ordering achieved by battery-aware periodic task voltage scheduling algorithm in [3], even though the peak current is significantly greater than in case of the level current profile. The objective function  $\sigma_{12}$ , the value after execution of the first hyperperiod, shows significant improvement (45.6 %). All sequential subprofiles are non-increasing, but the profile in Fig. 5(b) still has a lot of unused slack time.

The profile produced by the proposed algorithm (Fig. 5(c)), in which swapping task ordering plays an important role, further distributes the slack time. The objective function  $\sigma_{12}$  shows the further considerable improvement (33.7 % from the non-increasing profile, 63.3 % from the level current profile). The resulting profile of the proposed algorithm does not always have a non-increasing task ordering. In fact Fig. 5(c) contains a non-increasing subprofile, while the overall tendency is non-increasing. Better slack time utilization, however, shows better battery efficiency.

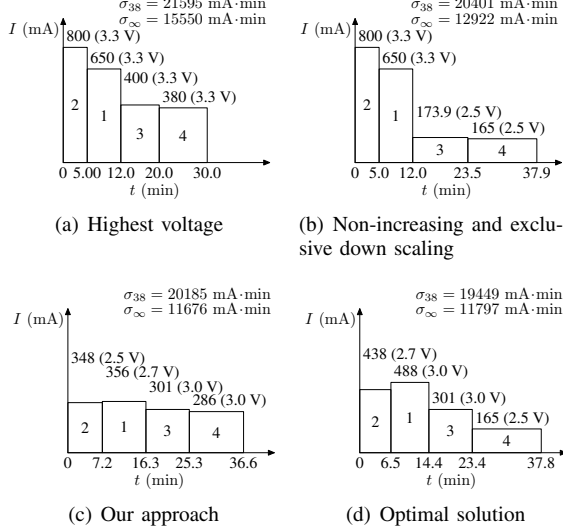


Figure 4. Load profiles for tasks in Table 1.

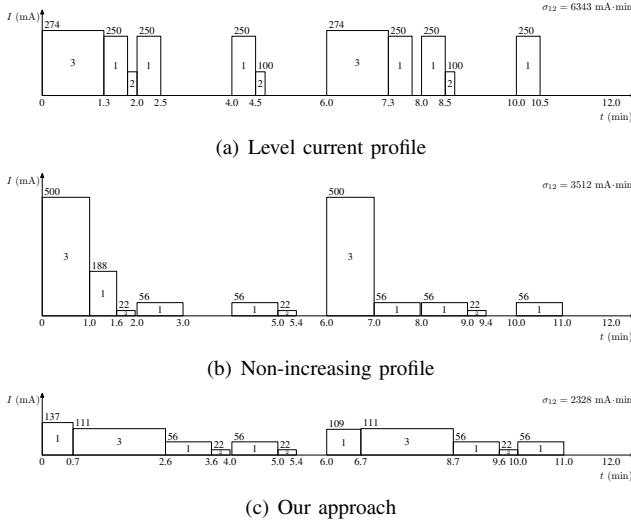


Figure 5. Profiles of the first hyperperiod.

## 5. Experimental Results

For comparison purposes, throughout this paper, we used the same task sets as in the existing work [3]. Table 3 shows the current of the real-time applications running on ITSy [3]. The arrival time of all tasks was assumed to be time 0 to use the greatest freedom of schedulability.

**Aperiodic tasks on uniprocessor systems.** Four task graphs are presented in Fig. 6: the strictly increasing profile (Case I), the independent tasks (Case II), and randomly generated task sets (Case III and IV). The number in the circle denotes the number of tasks specified in Table 3. Each task has a pair of duration (min) and deadline (min).

Table 3. Task specifications

| Task # | Name        | Current (mA) |
|--------|-------------|--------------|
| 1      | MPEG        | 208.92       |
| 2      | DICT206high | 186.53       |
| 3      | TTS206      | 102.89       |
| 4      | TTS74high   | 98.77        |
| 5      | TTS74low    | 98.77        |
| 6      | WAV206high  | 79.28        |
| 7      | WAV59       | 70.71        |

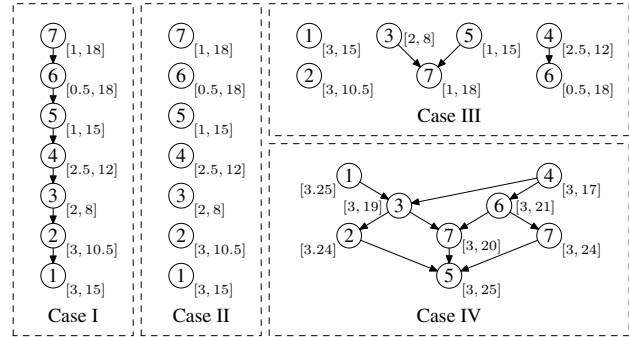


Figure 6. Aperiodic task graphs.


For comparison purpose, we used exclusive down scaling algorithm (*ExclusivDownScaling2*( $\cdot$ ) in [2, p.304]), which showed the best result among algorithms described in [2] for their example task sets, and the profile produced by non-increasing algorithm [3]. Table 4 shows that all the resulting profiles of the proposed algorithm were not worse than exclusive down scaling profiles and non-increasing profiles. The proposed algorithm was approximately 2 % superior on average and 3 % at the maximum.

Non-increasing algorithm did not return the result for the long profiles, e.g., 160 identical tasks of task #1 (duration 1 min in the highest voltage, deadline 240 min). Non-increasing algorithm cannot recover the battery exhaustion, since it does not take into account downscaling of tasks preceding failing task. The proposed algorithm, however, successfully returned the feasible profile.

**Periodic tasks on uniprocessor systems.** Fig. 7 presents task sets of the same period dependent tasks (Case I) and the different period independent tasks (Case II). Each periodic task set was executed 20 times and the results are presented in Table 5. The first column of each case describes the battery charge imposed by the first period and the second column by the 20th period. In all cases, our profiles are superior to non-increasing profiles, in the way that the improvement is up to 19.0 %, with an average of 15.5 %. Interestingly, in our result, the different periods had the different voltage profiles, while in non-increasing profiles those were the same.

**Table 4. Charges given by the aperiodic task voltage scheduling for uniprocessor systems for tasks in Fig. 6**

| Algorithm\ task sets   | I    | II   | III  | IV   |
|------------------------|------|------|------|------|
| Exclusive down scaling | 3258 | 3039 | 3039 | 5265 |
| Non-increasing         | 3258 | 2834 | 2928 | 5807 |
| Proposed algorithm     | 3254 | 2799 | 2800 | 5265 |



**Figure 7. Periodic task graphs.**

## 6. Concluding Remarks

We have proposed static voltage scheduling algorithms for battery-powered DVS systems based on the studied battery characteristics and our analysis. The proposed algorithms are extensions to Chowdhury and Chakrabarti's [3], and are designed by using greedy heuristics. Periodic and aperiodic voltage scheduling on uniprocessor platforms outperformed those in the comparative work [6], [7].

**Acknowledgment.** The authors wish to thank Takuya Azumi, Krzysztof Jozwik, Seiya Shibata, and Hideki Takase for providing useful comments on a draft of this paper. Part of this work was supported by CREST, JST.

## References

- [1] T. L. Martin, "Balancing batteries, power, and performance: system issues in CPU speed-setting for mobile computing," Ph.D. dissertation, Carnegie Mellon University, 1999.
- [2] D. Rakhmatov and S. Vrudhula, "Energy management for battery-powered embedded systems," *ACM Transactions on Embedded Computing Systems*, vol. 2, no. 3, pp. 277–324, 2003.
- [3] P. Chowdhury and C. Chakrabarti, "Static task-scheduling algorithms for battery-powered DVS systems," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 13, no. 2, pp. 226–237, 2005.
- [4] R. Rao, S. Vrudhula, and D. N. Rakhmatov, "Battery modeling for energy-aware system design," *IEEE Computer*, vol. 36, no. 12, pp. 77–87, 2003.
- [5] T. F. Fuller, M. Doyle, and J. Newman, "Simulation and optimization of the dual lithium ion insertion cell," *Journal of the Electrochemical Society*, vol. 141, pp. 1–10, 1994.
- [6] D. N. Rakhmatov and S. B. K. Vrudhula, "An analytical high-level battery model for use in energy management of portable electronic systems," *Proc. International Confer-*

**Table 5. Charges given by the periodic task voltage scheduling at the end of specified hyperperiod for uniprocessor systems for tasks in Fig. 7**

| Algorithm\ task sets | I    |      | II   |      |
|----------------------|------|------|------|------|
|                      | 1st  | 20th | 1st  | 20th |
| Non-increasing       | 1534 | 2013 | 1939 | 2198 |
| Proposed algorithm   | 1234 | 1705 | 1663 | 1918 |

*ence on Computer-Aided Design (ICCAD 01)*, IEEE Press, Nov. 2001, pp. 488–493.

- [7] D. Rakhmatov, S. Vrudhula, and D. Wallach, "A model for battery lifetime analysis for organizing applications on a pocket computer," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 11, no. 6, pp. 1019–1030, 2003.
- [8] P. Rong and M. Pedram, "Battery-aware power management based on Markovian decision processes," *Proc. International Conference on Computer-Aided Design (ICCAD 02)*, Nov. 2002, pp. 707–713.
- [9] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino, and R. Scarsi, "Discrete-time battery models for system-level low-power design," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 9, no. 5, pp. 630–640, 2001.
- [10] C.-F. Chiasserini and R. R. Rao, "Energy efficient battery management," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 7, pp. 1235–1245, 2001.
- [11] T. L. Martin and D. P. Siewiorek, "A power metric for mobile systems," *Proc. International Symposium on Low Power Electronics and Design (ISLPED 96)*, IEEE Press, Aug. 1996, pp. 37–42.
- [12] J. Luo and N. K. Jha, "Battery-aware static scheduling for distributed real-time embedded systems," *Proc. Design Automation Conference (DAC 01)*, ACM Press, Jun. 2001, pp. 444–449.
- [13] D. Rakhmatov, S. Vrudhula, and C. Chakrabarti, "Battery-conscious task sequencing for portable devices including voltage/clock scaling," *Proc. Design Automation Conference (DAC 02)*, ACM Press, Jun. 2002, pp. 189–194.
- [14] T. Ishihara and H. Yasuura, "Voltage scheduling problem for dynamically variable voltage processors," *Proc. International Symposium on Low Power Electronics and Design (ISLPED 98)*, Aug. 1998, pp. 197–202.
- [15] D. Shin and J. Kim, "Optimizing intratask voltage scheduling using profile and data-flow information," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 26, no. 2, pp. 369–385, 2007.
- [16] *Intel StrongARM SA-1100 Microprocessor Developer's Manual*, Intel, 1999.
- [17] D. N. Rakhmatov, "Battery voltage prediction for portable systems," in *IEEE International Symposium on Circuits and Systems*. IEEE, 2005, pp. 4098–4101.
- [18] R. Rao and S. Vrudhula, "Battery optimization vs energy optimization: Which to choose and when?" *Proc. International Conference on Computer-Aided Design (ICCAD 05)*, IEEE Computer Society, Nov. 2005, pp. 439–445.
- [19] H.-S. Yun and J. Kim, "On energy-optimal voltage scheduling for fixed-priority hard real-time systems," *ACM Transactions on Embedded Computing Systems*, vol. 2, no. 3, pp. 393–430, 2003.