

A Framework of Concurrent Task Scheduling and Dynamic Voltage and Frequency Scaling in Real-Time Embedded Systems with Energy Harvesting

Xue Lin¹, Yanzhi Wang¹, Siyu Yue¹, Naehyuck Chang² and Massoud Pedram¹

¹University of Southern California, Los Angeles, CA USA

²Seoul National University, Seoul, Korea

¹{xuelin, yanzhiwa, siyuyue, pedram}@usc.edu, ²{naehyuck}@elpl.snu.ac.kr

Abstract

Energy harvesting is a promising technique to overcome the limitation imposed by the finite energy capacity of batteries in conventional battery-powered embedded systems. In particular, the question of how one can achieve full energy autonomy (i.e., perpetual, battery-free operation) of a real-time embedded system with an energy harvesting capability (RTES-EH) by applying a global control strategy is investigated. The energy harvesting module is comprised of a Photovoltaic (PV) panel for harvesting energy and a supercapacitor for storing any excess energy. The global controller performs optimal operating point tracking for the PV panel, state-of-charge management for the supercapacitor, and energy-harvesting-aware real-time task scheduling with dynamic voltage and frequency scaling (DVFS) in the embedded load device. The controller, which accounts for dynamic V-I characteristics of the PV panel, terminal voltage variation and self-leakage of the supercapacitor, and power losses in voltage converters, employs a cascaded feedback control structure with an inner control loop determining the V-I operating point of the PV panel and an outer supervisory control loop performing real-time task scheduling and setting the voltage and frequency level in the embedded load device (to keep the state-of-charge of the supercapacitor in a desirable range). Experimental results show that the proposed global controller lowers the task drop rate in a RTES-EH by up to 60% compared with baseline controller within the same service time.

Keywords

Task scheduling, DVFS, photovoltaic, real-time embedded system, supercapacitor

1. Introduction

Minimizing power consumption remains one of the critical design challenges for battery-powered devices and systems. *Dynamic power management* (DPM) [1][2][3] and *dynamic voltage and frequency scaling* (DVFS) [4][5][6] have been proven to be effective for reducing energy dissipation while meeting performance constraints. However, even elaborated DVFS or DPM cannot make the battery operate forever, and the battery should be recharged or replaced to maintain the system operation.

Additionally, some applications do not allow battery recharging or replacement. An example is sensor nodes that are deployed in radioactive surroundings. *Energy harvesting* to increase the lifespan of such applications has been actively explored [7][8][9]. It is regarded as a prospective method for overcoming the energy limitation in traditional battery-powered embedded systems and achieving energy autonomy. Available ambient energy sources for energy harvesting include sunlight, wind, tidal wave, etc. The harvested energy is limitless, but the amount of harvested energy varies significantly with changes in environmental conditions. Therefore, an energy storage element (e.g., a rechargeable battery or a supercapacitor [7][9]) becomes essential for continuous power supply in an embedded system with energy harvesting.

A number of recent works have been carried out on power minimization algorithms and techniques for *real-time embedded systems with energy harvesting* (RTES-EH) [10][11][12][13]. The major difference between these works and the traditional real-time task scheduling and DVFS works [4][5][6] comes from both the

intermittent nature of harvested energy and limited size of the energy storage element. Enhancement of the *system availability* is the primary purpose of the RTES-EH task scheduling. The lazy scheduling algorithm executes tasks at full speed as late as possible [10], but the task slacks are not exploited for energy saving. Later works have chosen solar power as the energy harvesting source and presented algorithms for improving the deadline miss rate and reducing energy dissipation by using DVFS [11][12][13].

We present a *global control algorithm* for a RTES-EH, which adopts a PV panel as the energy harvesting source, a supercapacitor as the energy storage element, and a real-time sensor node as the embedded load device. The proposed global controller performs simultaneous optimal operating point tracking of the PV panel, state-of-charge (SoC) management of the supercapacitor, as well as effective real-time task scheduling with DVFS in the sensor node.

Key shortcomings of the previous works and contributions of the present work are summarized next.

- 1) Previous works rely on the battery as the energy storage element [10][11][12][13], which prevents the RTES-EH from having a long life because batteries tend to have a rather short cycle life (up to several hundred charge-discharge cycles) after which they have to be replaced. Nowadays supercapacitors constitute a more promising storage element for a RTES-EH due to (i) orders-of-magnitude longer cycle life compared to batteries and (ii) higher power capacity to deal with bursty power demands of load devices. Because the terminal voltage variation of a supercapacitor is much larger than that of a battery [14], it is not possible to apply the previous works (which assume ideal efficiency for the DC-DC converters in the system) to the supercapacitor-based RTES-EH. In this paper, we present a global control algorithm for RTES-EH, which fully accounts for the supercapacitor characteristics.
- 2) There is some power loss in the power converters, connecting the energy source to energy storage elements, and these elements to the load [15]. The converter power losses are significant in the RTES-EH especially for the supercapacitor-based system because the supercapacitor terminal voltage variation can be quite high. This paper considers variations in the converter efficiency based on accurate power converter models in order to optimize the overall system availability.
- 3) Most previous approaches assume ideal component models for both the energy source and the storage element. Accurate models should consider the nonlinear V-I curve of the PV panel, the energy overhead for charging/discharging storage elements, and the self-discharge of supercapacitor. We develop an accurate component model-based global control algorithm for the RTES-EH, thereby, enhancing the overall system availability.

We first identify two fundamental optimization problems i.e., the *supercapacitor charging problem* and the *converter-aware frequency setting problem*. The first problem aims at maximizing the supercapacitor charging current by optimally setting the operating point of the PV panel, similar to the idea proposed in [16]. This problem is different from the traditional *maximum power point tracking* (MPPT) technique [17] in that the latter neglects the power loss in the converter. The second problem focuses on finding the optimal sensor node execution frequency for a single task instance without deadline limit, such that the amount of energy extracted from

the supercapacitor is minimized. The solution is based on the observation in [15] that the minimum execution frequency in the sensor node is no longer always optimal in total energy saving for a single task instance without deadline limit when the converter power loss is considered.

Based on the optimal solutions of the aforesaid two fundamental problems, we further propose a near-optimal global control algorithm based on a cascaded feedback control structure. The outer supervisory control loop maintains the supercapacitor SoC through the whole operation period (from sunrise to sunset) so as to enhance the overall system availability and reduce the task instance drop rate. The inner control loop determines the V-I operating point of the PV panel (and hence the input current of the supercapacitor.) Experimental results demonstrate that the proposed global control algorithm on RTES-EH significantly lowers the task instance drop rate by up to 60% compared with baseline control algorithms.

2. Component Models

The energy source of the RTES-EH is a PV panel. The PV panel output voltage V_{pv} and current I_{pv} satisfy certain V-I characteristics given the solar irradiance G , as shown from the PV model in [18][19].

We employ PWM (pulse width modulation) buck-boost power converter as both the charger and DC-DC converter in the RTES-EH. The charger regulates its output current to the value set by the microcontroller, and the DC-DC converter regulates its output voltage. We adopt the converter power model from [20]. The power conversion efficiency of a power converter is a function of its input and output voltages and currents.

Supercapacitors have a small internal resistance and thus small power loss in charging/discharging operations. They exhibit a higher volumetric power density and a longer cycle life compared with batteries [21]. These features make supercapacitors preferred in the RTES-EH with frequent charging and discharging.

However, a primary disadvantage of the supercapacitor is a full-range terminal voltage variation as function of its SoC. The terminal voltage $V_{cap}(t)$ of a supercapacitor is a linear function of its SoC $Q_{cap}(t)$, given by $V_{cap}(t) = Q_{cap}(t)/C_{cap} = \sqrt{2E_{cap}(t)}/C_{cap}$, where $E_{cap}(t)$ is the energy stored in the supercapacitor, and C_{cap} is the capacitance. The terminal voltage variation incurs a significant power conversion efficiency variation.

In addition, a supercapacitor may lose 40% of its stored energy per day by self-discharge [21]. The voltage decay of a supercapacitor after a time period Δt is given by

$$V_{cap}(t + \Delta t) = V_{cap}(t) \cdot e^{-\frac{\Delta t}{\Gamma}}, \quad (1)$$

when no external power supply or load is connected to it. The parameter Γ is the self-discharge time constant.

3. RTES-EH System

3.1. System architecture




Figure 1: Block diagram of the target RTES-EH.

The RTES-EH architecture is shown in Figure 1, which is comprised of a PV panel as the energy source, a supercapacitor as the energy storage, a real-time DVFS-enabled sensor node as the load device, and power converters. A charger connects the PV panel and the supercapacitor, and a DC-DC converter connects the supercapacitor and the sensor node. The RTES-EH operates from sunrise $T_{sunrise}$ until sunset T_{sunset} . $G(t)$ denotes the solar

irradiance at time t . The $G(t)$ profile is predicted via solar irradiance prediction algorithms. The output voltage and current of the PV panel at time t are $V_{pv}(t)$ and $I_{pv}(t)$, respectively. $V_{pv}(t)$ and $I_{pv}(t)$ satisfy certain V-I characteristics given $G(t)$. We control the PV panel operating point $(V_{pv}(t), I_{pv}(t))$ through controlling the charger's output current $I_{cap,in}(t)$.

We denote the terminal voltage, input current and output current of the supercapacitor at time t by $V_{cap}(t)$, $I_{cap,in}(t)$ and $I_{cap,out}(t)$, respectively. The supercapacitor stored energy is calculated by

$$E_{cap}(t) = E_{cap}(T_{sunrise}) + \int_{T_{sunrise}}^t V_{cap}(\tau) \cdot (I_{cap,in}(\tau) - I_{cap,out}(\tau) - I_{cap,sd}(\tau)) d\tau, \quad (2)$$

where $E_{cap}(T_{sunrise})$ is the initial capacitor energy at $T_{sunrise}$, and $I_{cap,sd}(t)$ is the self-discharge current, which is calculated using (1).

We assume a single-core microprocessor as the sensor node. We denote its supply voltage, input current, and power consumption by $V_{proc}(t)$, $I_{proc}(t)$ and $P_{proc}(t)$, respectively. We denote the sensor execution frequency by $f(t)$. The sensor has N discrete operating frequencies: $f_{min} = f_1 < f_2 < \dots < f_N = f_{max}$; each operating frequency f_n corresponds to a supply voltage $V_{proc}(t) = V_{proc,n}$ and a power consumption level $P_{proc}(t) = P_{proc,n}$, where $1 \leq n \leq N$.

The input voltage, input current, output voltage, and output current of the charger are $V_{pv}(t)$, $I_{pv}(t)$, $V_{cap}(t)$, and $I_{cap,in}(t)$, respectively. The charger power loss $P_{conv,in}(t)$ is a function of $V_{pv}(t)$, $I_{pv}(t)$, and $V_{cap}(t)$ as shown in [20]. Similarly, the power loss $P_{conv,out}(t)$ of the DC-DC converter is a function of its input voltage, output voltage, and output current, i.e., $V_{cap}(t)$, $V_{proc}(t)$, and $I_{proc}(t)$, respectively [20]. The following two functions hold due to the energy conservation law:

$$V_{pv}(t) \cdot I_{pv}(t) = V_{cap}(t) \cdot I_{cap,in}(t) + P_{conv,in}(t), \quad (3)$$

$$V_{cap}(t) \cdot I_{cap,out}(t) = V_{proc}(t) \cdot I_{proc}(t) + P_{conv,out}(t). \quad (4)$$

3.2. Real-time task set

We assume a frame-based preemptive hard real-time system in the sensor node. The frame length is T_{LCM} , which is the hyper-period of all the tasks. The frame is executed repeatedly from $T_{sunrise}$ until T_{sunset} . The number of frames during the whole execution time is $K = \frac{T_{sunset} - T_{sunrise}}{T_{LCM}}$. Each k -th frame ($1 \leq k \leq K$) begins at $T_{sunrise} + (k-1)T_{LCM}$ and ends at $T_{sunrise} + kT_{LCM}$. The duration of a frame is in the order of minutes or seconds, which is much shorter than the whole RTES-EH operating time. The earliest deadline first (EDF) scheduling [4] is incorporated in the sensor node to schedule periodic tasks, which are independent from each other.

We denote the set of M real-time periodic tasks by $\{T_1, T_2, \dots, T_M\}$. Each task i ($1 \leq i \leq M$) has a period P_i and workload W_i . The workload is described as the number of clock cycles to complete an instance of task i . Each task i releases its task instances periodically. The j -th ($1 \leq j \leq T_{LCM}/P_i$) instance of task i in the k -th frame, denoted by $\mathcal{J}_{i,j,k}$, is released at time $T_{sunrise} + (k-1)T_{LCM} + (j-1)P_i$, which is the absolute release time of that task instance. The relative deadline of task i is assumed to be the same as its period P_i . Hence, the absolute deadline of task instance $\mathcal{J}_{i,j,k}$ is equal to $T_{sunrise} + (k-1)T_{LCM} + jP_i$.

4. Problem Formulation

Let $f_i(t)$ denote the sensor execution frequency assigned to task i at time t . Considering real implementation, we assume that $f_i(t)$ can only take discrete values in the set $\{f_{min} = f_1, f_2, \dots, f_{max} = f_N\}$. If task i is not executed at time t , we have $f_i(t) = 0$. The single-core microprocessor assumption ensures that at most one of the $f_i(t)$ values at time t is non-zero. We have $f(t) = \sum_{1 \leq i \leq M} f_i(t)$. We adopt a more general definition of a *schedule* of task instances where task instance dropping or timing constraint violation is allowed to accommodate the intermittent nature in the input energy of the

RTES-EH [11][12]. We define $\delta_{i,j,k}$ such that $\delta_{i,j,k} = 1$ indicates completion of task instance $\mathcal{T}_{i,j,k}$ before its deadline, i.e.,

$$\int_{T_{sunrise}+(k-1)T_{LCM}+jP_i}^{T_{sunrise}+(k-1)T_{LCM}+(j-1)P_i} f_i(t)dt = W_i. \quad (5)$$

In contrast, $\delta_{i,j,k} = 0$ indicates that $\mathcal{T}_{i,j,k}$ has missed the deadline, i.e.,

$$\int_{T_{sunrise}+(k-1)T_{LCM}+jP_i}^{T_{sunrise}+(k-1)T_{LCM}+(j-1)P_i} f_i(t)dt < W_i. \quad (6)$$

We minimize the task instance drop rate, or equivalently, we maximize total number of fully executed task instances, given by

$$Total_Executed = \sum_{1 \leq k \leq K} \sum_{1 \leq i \leq M} \sum_{1 \leq j \leq T_{LCM}/P_i} \delta_{i,j,k}. \quad (7)$$

$P_{proc}(t)$ and $V_{proc}(t)$ depend on $f(t)$ where $f(t) = \sum_{1 \leq i \leq M} f_i(t)$. $I_{cap,out}(t)$ is determined by $P_{proc}(t)$, $P_{conv,out}(t)$ and $V_{cap}(t)$ using (4). Besides the $f_i(t)$ profiles, the PV operating point $(V_{pv}(t), I_{pv}(t))$ is also a set of control variables of the RTES-EH. $I_{cap,in}(t)$ is determined by $(V_{pv}(t), I_{pv}(t))$, $P_{conv,in}(t)$ and $V_{cap}(t)$ using (3). We calculate $E_{cap}(t + \Delta t)$, i.e., the energy stored in the supercapacitor at time $t + \Delta t$, using (2). Hence, we achieve the supercapacitor SoC control (which is equivalent to the control of its stored energy) through task scheduling with DVFS (i.e., finding the $f_i(t)$ profile for each task i) and PV panel operating point tracking.

The formal statement of the global control problem for the RTES-EH is given as follows.

RTES-EH global control problem statement:

Given the solar irradiance profile $G(t)$ and initial supercapacitor stored energy $E_{cap}(T_{sunrise})$,

Find a schedule and frequency allocation of tasks, represented by $f_i(t)$ ($1 \leq i \leq M$), and the PV operating point $(V_{pv}(t), I_{pv}(t))$ for $t \in [T_{sunrise}, T_{sunset}]$,

Maximize the total number of completely executed task instances before deadline, given by (7),

Subject to the energy conservation law (2), (3), (4).

5. Model-Based Optimization

We first define two fundamental optimization problems that are necessary in solving the original global control problem: the *supercapacitor charging* (SC) problem and the *converter-aware frequency setting* (CA-FS) problem.

5.1. Supercapacitor charging problem

The SC problem aims at maximizing $I_{cap,in}(t)$ by optimally setting $(V_{pv}(t), I_{pv}(t))$, given $G(t)$ and $E_{cap}(t)$, for $t \in [T_{sunrise}, T_{sunset}]$. Solving the SC problem optimally at time $t \in [T_{sunrise}, T_{sunset}]$ is required for the optimal solution of the global control problem for RTES-EH.




Figure 2: V-I and V-P characteristics of the source PV panel.

Figure 2 shows the V-I and voltage-power (V-P) characteristics of the PV panel. For a specific $G(t)$, $V_{pv}(t)$ decreases as $I_{pv}(t)$ increases, and the output power $V_{pv}(t) \cdot I_{pv}(t)$ is maximized at the *maximum power point* (MPP) marked by red dots in Figure 2. Conventional *MPP tracking* (MPPT) technique maintains the PV panel operating point at its MPP. However, it cannot guarantee the maximum amount of energy transferred into the supercapacitor due

to the efficiency variation of the charger. The *maximum power transfer tracking* (MPTT) technique takes into account the non-ideal behavior of the charger, thereby maximizing the energy transferred into the supercapacitor [14]. We adopt the MPTT method to solve the SC problem optimally at time $t \in [T_{sunrise}, T_{sunset}]$.

5.2. Converter-aware frequency setting problem

The CA-FS problem aims to find the optimal sensor execution frequency for a specific task instance without deadline constraint, given E_{cap} (or equivalently, V_{cap}), such that the energy extracted from the supercapacitor is minimized. The motivation is that the minimum frequency in the sensor node is no longer always optimal in energy saving for a single task instance without deadline constraint when the converter power loss is considered [15].

We derive the total energy extracted from the supercapacitor during execution of the task instance as a function of the sensor node execution frequency f . Note that f can only take discrete values from $\{f_{min} = f_1, f_2, \dots, f_{max} = f_N\}$. The sensor supply voltage and power consumption during task execution are functions of f , denoted by $V_{proc}(f)$ and $P_{proc}(f)$, respectively. The total execution time of the task instance is $T_{exe}(f) = W/f$, where W is the workload of the task instance. We assume that V_{cap} does not change during a single task instance execution since the execution time is too short to make noticeable SoC change. The DC-DC converter power loss during task execution is a function of its input voltage V_{cap} , output voltage $V_{proc}(f)$ and output current $I_{proc}(f) = P_{proc}(f)/V_{proc}(f)$ as specified in [20]. We denote the DC-DC converter power loss by $P_{conv,out}(f, V_{cap})$. The amount of energy extracted from the supercapacitor during the task instance execution is given by:

$$E_{extract}(f, V_{cap}) = T_{exe}(f)(P_{proc}(f) + P_{conv,out}(f, V_{cap})) \quad (8)$$

which is not a monotonically increasing function of f . This implies that using the lowest feasible sensor node voltage and frequency for task execution does not always minimize the total energy drawn from supercapacitor. Let $f_{opt}(V_{cap})$ denote the optimal sensor execution frequency that minimizes the total energy drawn from supercapacitor. The optimal frequency is a function of V_{cap} and is calculated by

$$f_{opt}(V_{cap}) = \underset{f_{min} \leq f \leq f_{max}}{\operatorname{argmin}} E_{extract}(f, V_{cap}). \quad (9)$$

6. Global Control Algorithm

We present a cascaded feedback control based algorithm, comprised of an outer supervisory control and an inner control loop. The outer supervisory control, which is performed at each decision epoch (to be defined later), maintains the supercapacitor SoC through task scheduling with DVFS, as well as (possibly) selectively dropping some tasks. The inner control loop sets the optimal PV panel operating point. The outer supervisory control loop and the inner control loop are performed at the rate of once every 300 seconds (which is equal to T_{LCM}) and five milliseconds, respectively.

Decision epochs are defined as the start times of each frame of tasks. We denote the k -th ($1 \leq k \leq K$) decision epoch by $T_{s,k}$, and have $T_{s,k} = T_{sunrise} + (k-1)T_{LCM}$. At each decision epoch $T_{s,k}$, the supervisory control algorithm finds a schedule and sets the frequency for each task instance in the k -th frame, i.e., $f_i(t)$ ($1 \leq i \leq M$) for $t \in [T_{s,k}, T_{s,k+1}]$. The supervisory control algorithm is aware of future power harvesting and power consumption in the RTES-EH in order to effectively control the supercapacitor SoC.

6.1. Motivation

We have the following observations on the RTES-EH:

- 1) At the beginning (i.e., in the morning) and at the end (i.e., in the evening) of the RTES-EH operation, the solar irradiance level is low; while the solar irradiance becomes abundant in the middle (i.e., at noon) of system operation in general.
 - 2) At the beginning of RTES-EH operation, the supercapacitor is lack of stored energy due to the self-discharge overnight.
- The above observations lead to the following motivations for the proposed global control algorithm:

- 1) There exists a desirable amount of supercapacitor stored energy (or equivalently, the supercapacitor terminal voltage) when considering efficiency variations of both the charger and the DC-DC converter. The task instance drop rate is minimized when the amount of energy stored in supercapacitor is close to such value.
- 2) Selectively dropping some task instances at the beginning of the RTEs-EH operation may eventually help increase the total number of completely executed task instances since the supercapacitor stored energy may increase (towards the desirable value) along with system operation.
- 3) It is desirable to use up the energy stored in supercapacitor by $T_{\text{sun\textit{set}}}$ since the stored energy will be eventually dissipated due to self-discharge at night.

6.2. Outer supervisory control loop

In order to effectively perform supervisory control at each decision epoch, we first identify two subproblems and provide corresponding solutions: the *converter-aware task scheduling* (CA-TS) problem (without task instance dropping) and the *supercapacitor energy control* (SEC) problem. Then we provide the supervisory control algorithm based on the solutions of the two subproblems as well as the above-described motivations. Figure 3 illustrates the relationship between the control algorithms and the subproblems in this paper in order for better understanding.




Figure 3: The relationship between algorithms and subproblems.

6.2.1. Converter-aware task scheduling problem

Suppose we are at the k -th decision epoch $T_{s,k}$. The supercapacitor stored energy and terminal voltage are $E_{cap}(T_{s,k})$ and $V_{cap}(T_{s,k})$, respectively. We are given a set \mathfrak{I} of task instances. The absolute release time and absolute deadline of all task instances in \mathfrak{I} are within the time period $[T_{s,k}, T_{s,k+1}]$. We are going to find a schedule and frequency allocation of all the task instances in \mathfrak{I} , such that the amount of energy extracted from the supercapacitor is minimized. Task instance dropping is not allowed in this problem.

Finding the optimal schedule and frequency allocation of a set of task instances on a continuously variable frequency processor is optimally solved in [4] based on the EDF scheduling. In the solution of the CA-TS problem, we incorporate the facts: (i) a practical sensor node adopts discretely variable frequency (and voltage) levels, and (ii) when the converter power loss is considered, there exists an optimal sensor execution frequency based on which the amount of energy extracted from the supercapacitor is minimized, as discussed

Algorithm 1: Solution of the CA-TS Problem.

Generate an optimal continuously variable frequency (and voltage) allocation and schedule for all task instances in \mathfrak{I} , using the method in [4]

Find the optimal sensor node execution frequency $f_{opt}(V_{cap}(T_{s,k}))$ by optimally solving the CA-FS problem

For each task instance:

If the allocated frequency is less than $f_{opt}(V_{cap}(T_{s,k}))$:

Allocate the frequency $f_{opt}(V_{cap}(T_{s,k}))$ to that task instance

Use the algorithm proposed in [22] to generate the optimal discretely variable frequency allocation and schedule for all task instances in \mathfrak{I}

Return the frequency allocation and schedule

in Section 5.2. The basic idea is to use (at most) two discrete frequency (and voltage) levels for each task instance. The two frequency levels are the immediate neighbors of the frequency allocated to that task instance by the scheduling algorithm in [4], which is based on a sensor node with continuously variable frequency levels. Algorithm 1 provides the details of the proposed solution of the CA-TS problem. Similar to the solution of the CA-FS problem, we assume that the supercapacitor voltage will not change significantly during the time period $[T_{s,k}, T_{s,k+1}]$.

6.2.2. Supercapacitor energy control problem

Suppose we are at the k -th decision epoch $T_{s,k}$. The SEC problem aims to find the optimal amount of supercapacitor stored energy, denoted by $E_{cap,k}^{tar}$. The overall system availability will be maximized when the amount of energy stored in the supercapacitor is near such a value. In the following, we discuss about the equivalent problem of finding the optimal supercapacitor voltage $V_{cap,k}^{tar}$.

For each possible V_{cap} , we calculate the *estimated system availability* as follows. We assume that the supercapacitor terminal voltage will not change (i.e., $V_{cap}(t) \approx V_{cap}$) for simplicity in estimation. The estimated input energy of the supercapacitor during $[T_{s,k}, T_{\text{sun\textit{set}}}]$, denoted by $\hat{E}_{in,k,total}(V_{cap})$, is calculated using:

$$\hat{E}_{in,k,total}(V_{cap}) = V_{cap} \cdot \int_{T_{s,k}}^{T_{\text{sun\textit{set}}}} \max_{(V_{pv}(t), I_{pv}(t)); V_{cap}(t)=V_{cap}} I_{cap,in}(t) dt, \quad (10)$$

where at time $t \in [T_{s,k}, T_{\text{sun\textit{set}}}]$, we maximize the (estimated) supercapacitor input current $I_{cap,in}(t)$ via finding the optimal $(V_{pv}(t), I_{pv}(t))$ (i.e., optimally solving the SC problem.) Next, we calculate the estimated output energy of the supercapacitor. We solve the CA-TS problem assuming that the supercapacitor terminal voltage is V_{cap} . We obtain the schedule of task instances in the k -th frame, represented by $f_i(t)$ ($1 \leq i \leq M$) for $t \in [T_{s,k}, T_{s,k+1}]$. We calculate the (estimated) supercapacitor discharging current $I_{cap,out}(t)$ based on $f(t) = \sum_{1 \leq i \leq M} f_i(t)$ and V_{cap} in the way described in Section 4. The estimated output energy of the supercapacitor during $[T_{s,k}, T_{s,k+1}]$ as a function of V_{cap} , denoted by $\hat{E}_{out,k}(V_{cap})$, is calculated using:

$$\hat{E}_{out,k}(V_{cap}) = V_{cap} \cdot \int_{T_{s,k}}^{T_{s,k+1}} I_{cap,out}(t) dt. \quad (11)$$

Since the supercapacitor voltage is assumed to be the constant value V_{cap} , the estimated output energy of the supercapacitor over each k' -th frame ($k < k' \leq K$) is also equal to $\hat{E}_{out,k}(V_{cap})$.

The estimated self-discharge energy loss in the supercapacitor during $[T_{s,k}, T_{\text{sun\textit{set}}}]$ as a function of V_{cap} , denoted by $\hat{E}_{sd,k,total}(V_{cap})$, is calculated using:

$$\begin{aligned} \hat{E}_{sd,k,total}(V_{cap}) &= V_{cap} \cdot I_{cap,sd}(V_{cap}) \cdot (T_{\text{sun\textit{set}}} - T_{s,k}) \\ &= (K - k + 1) \cdot C_{cap} \cdot \frac{(V_{cap})^2}{\Gamma} \cdot T_{LCM}, \end{aligned} \quad (12)$$

where $I_{cap,sd}(V_{cap})$ is the supercapacitor self-discharge current as a function of V_{cap} . The above equation is derived from (1).

We define the ratio of $\hat{E}_{in,k,total}(V_{cap}) - \hat{E}_{sd,k,total}(V_{cap})$ to $\hat{E}_{out,k}(V_{cap})$ the estimated system availability. The estimated system availability represents the estimated number of frames that can be executed in the sensor node during $[T_{s,k}, T_{\text{sun\textit{set}}}]$, when keeping the supercapacitor energy unchanged. The optimal (target) supercapacitor terminal voltage $V_{cap,k}^{tar}$ is the optimal V_{cap} that maximizes the estimated system availability. $V_{cap,k}^{tar}$ is calculated by:

$$V_{cap,k}^{tar} = \operatorname{argmax}_{V_{cap}} \frac{\hat{E}_{in,k,total}(V_{cap}) - \hat{E}_{sd,k,total}(V_{cap})}{\hat{E}_{out,k}(V_{cap})}. \quad (13)$$

6.2.3. Supervisory control algorithm

The objective of the supervisory control algorithm is to maintain the amount of supercapacitor stored energy (or equivalently, its SoC)

around or above a desirable value, through effective task scheduling with DVFS and selective task instance dropping. Consider the supervisory control algorithm performed at decision epoch $T_{s,k}$. The supercapacitor energy and terminal voltage are $E_{cap}(T_{s,k})$ and $V_{cap}(T_{s,k})$, respectively. The optimal (target) supercapacitor stored energy satisfies $E_{cap,k}^{tar} = \frac{1}{2} \cdot C_{cap} \cdot (V_{cap,k}^{tar})^2$. We use $e_{cap,k} = E_{cap}(T_{s,k}) - E_{cap,k}^{tar}$ to denote the difference between the current amount of energy stored in the supercapacitor and the target value. The supervisory control algorithm goes as follows:

Case I ($e_{cap,k} \geq 0$): In this case, the amount of energy stored in the supercapacitor is abundant (higher than the target value.) We schedule all the task instances in the k -th frame during time period $[T_{s,k}, T_{s,k+1}]$ using Algorithm 1, without task instance dropping.

Case II ($e_{cap,k} < 0$): In this case, the objective of the supervisory control algorithm is to make the supercapacitor energy at the next decision epoch, $E_{cap}(T_{s,k+1})$, equal to or higher than $E_{cap}(T_{s,k}) + \delta \cdot |e_{cap,k}|$, in which δ is a predefined feedback control parameter. We may have to drop some task instances to achieve this goal. The basic idea is to keep dropping the most energy consuming task instance and performing task re-scheduling using Algorithm 1, until the estimated amount of energy stored in the supercapacitor at time $T_{s,k+1}$ satisfies that requirement. Details are shown in Algorithm 2.

However, the 3rd motivation in Section 6.1 specifies that it is desirable to use up the energy stored in supercapacitor by the time T_{sunset} . In other words, there is no need for keeping the supercapacitor stored energy around or above the target value near the end of the RTES-EH operation period (i.e., in the evening.) We can schedule all the task instances without dropping in this case. We use the following estimation-based procedure to decide at a decision epoch $T_{s,k}$ whether performing supercapacitor energy maintenance is necessary. We first calculate the estimated supercapacitor input energy $\hat{E}_{in,k,total}(V_{cap}(T_{s,k}))$ and the self-discharge energy loss

Algorithm 2: The Supervisory Control Algorithm.

Calculate the values $\hat{E}_{in,k,total}(V_{cap}(T_{s,k}))$, $\hat{E}_{out,k,total}(V_{cap}(T_{s,k}))$ and $\hat{E}_{sd,k,total}(V_{cap}(T_{s,k}))$

Initialize the set \mathfrak{S} to include all task instances in the k -th frame

If $\hat{E}_{in,k,total}(V_{cap}(T_{s,k})) + E_{cap}(T_{s,k}) \geq \hat{E}_{out,k,total}(V_{cap}(T_{s,k})) + \hat{E}_{sd,k,total}(V_{cap}(T_{s,k}))$:

Use Algorithm 1 to generate a schedule of all task instances in \mathfrak{S}

Else:

Calculate $E_{cap,k}^{tar}$ by solving the SEC problem

Calculate $e_{cap,k} = E_{cap}(T_{s,k}) - E_{cap,k}^{tar}$

If $e_{cap,k} \geq 0$:

Use Algorithm 1 to generate a schedule of all task instances in \mathfrak{S}

Else:

Calculate the estimated supercapacitor input energy during time period $[T_{s,k}, T_{s,k+1}]$, given by $\hat{E}_{in,k}(V_{cap}(T_{s,k})) = V_{cap}(T_{s,k}) \cdot \int_{T_{s,k}}^{T_{s,k+1}} \max(V_{pv}(t), I_{pv}(t)); V_{cap}(t) = V_{cap}(T_{s,k}) I_{cap,in}(t) dt$

Calculate the estimated supercapacitor self-discharge energy loss during $[T_{s,k}, T_{s,k+1}]$, given by $\hat{E}_{sd,k}(V_{cap}(T_{s,k})) = C_{cap} \cdot \frac{(V_{cap}(T_{s,k}))^2}{\Gamma} \cdot T_{LCM}$

Use Algorithm 1 to generate a schedule of all task instances in \mathfrak{S}

Calculate the estimated supercapacitor output energy $\hat{E}_{out,k}(V_{cap}(T_{s,k}))$ during $[T_{s,k}, T_{s,k+1}]$ based on the schedule

While $E_{cap}(T_{s,k}) + \hat{E}_{in,k}(V_{cap}(T_{s,k})) - \hat{E}_{out,k}(V_{cap}(T_{s,k})) - \hat{E}_{sd,k}(V_{cap}(T_{s,k})) < E_{cap}(T_{s,k}) + \delta \cdot |e_{cap,k}|$:

Drop the most energy consuming task instance from \mathfrak{S}

Use Algorithm 1 to generate a schedule for the new set \mathfrak{S}

Calculate $\hat{E}_{out,k}(V_{cap}(T_{s,k}))$ based on the new schedule

Return the schedule and the corresponding frequency allocation

$\hat{E}_{sd,k,total}(V_{cap}(T_{s,k}))$ over $[T_{s,k}, T_{sunset}]$ where the supercapacitor voltage is $V_{cap}(T_{s,k})$, using the method described in Section 6.2.2. Next, we calculate the estimated supercapacitor output energy $\hat{E}_{out,k,total}(V_{cap}(T_{s,k}))$ over $[T_{s,k}, T_{sunset}]$ where the supercapacitor voltage is $V_{cap}(T_{s,k})$. $\hat{E}_{out,k,total}(V_{cap}(T_{s,k}))$ is calculated by

$$\hat{E}_{out,k,total}(V_{cap}(T_{s,k})) = (K - k + 1) \hat{E}_{out,k}(V_{cap}(T_{s,k})), \quad (14)$$

where $\hat{E}_{out,k}(V_{cap}(T_{s,k}))$ is calculated as described in Section 6.2.2.

If $\hat{E}_{in,k,total}(V_{cap}(T_{s,k})) + E_{cap}(T_{s,k}) \geq \hat{E}_{out,k,total}(V_{cap}(T_{s,k})) + \hat{E}_{sd,k,total}(V_{cap}(T_{s,k}))$, we conclude that the stored energy plus the future harvested energy is enough for scheduling all the task instances until T_{sunset} , and subsequently, the supervisory control algorithm schedules all the task instances in the k -th frame without task instance dropping. Otherwise, we properly control the supercapacitor stored energy based on selective task instance dropping, as discussed before. Details of the proposed supervisory control algorithm are given in Algorithm 2. We again make the assumption that the supercapacitor terminal voltage will not change significantly during $[T_{s,k}, T_{s,k+1}]$.

6.3. Inner control loop

The inner control loop determines the optimal V-I operating point of the PV panel (and hence the input current of the supercapacitor.) The inner control loop is performed every five milliseconds in this paper. The inner control loop determines the optimal $(V_{pv}(t), I_{pv}(t))$ by optimally solving the SC problem. Besides, the sensor node execution frequency $f(t)$ is given by the task schedule and frequency allocation generated by the supervisory control algorithm. The inner control loop also controls $V_{proc}(t)$ using the DC-DC converter so that the embedded sensor node properly runs at frequency $f(t)$.

7. Experimental Results




Figure 4: Daily solar irradiance profiles.

We compare the proposed global control algorithm with baseline control algorithm on the RTES-EH in terms of task instance drop rate. The solar irradiance profiles used in our experiments are measured at Duffield, VA, in the year of 2007. Figure 4 shows four different daily solar irradiance profiles, in which the solar irradiance is normalized by the irradiance value G_{STC} at standard test condition. The following experiments are based on Profile 1 shown in Figure 4. We use $T_{sunrise} = 6:00$ AM and $T_{sunset} = 6:00$ PM. We use a 500 F supercapacitor as the energy storage in the RTES-EH.

The embedded sensor node exploited in our experiments has six discrete operating frequencies: 500 MHz, 600 MHz, 700 MHz, 800 MHz, 900 MHz and 1000 MHz. Correspondingly, the sensor node also has six discrete supply voltage levels 1.0 V, 1.2 V, 1.4 V, 1.6 V, 1.8 V and 2.0 V, and six power levels: 0.62 W, 0.89 W, 1.21 W, 1.58 W, 2.00 W and 2.47 W. Specifications of the task set in the embedded sensor node are provided in Table 1. Table 1 shows for

each task \mathcal{T}_i its period P_i (which is equal to its relative deadline.) It also provides the execution time of an instance of each task \mathcal{T}_i , given by W_i/f_{max} , when the sensor node runs at its highest frequency f_{max} . The frame length T_{LCM} , which is the hyper-period of all the tasks, is 300 s in the RTES-EH.

Table 1: Task set specifications.

Task	\mathcal{T}_1	\mathcal{T}_2	\mathcal{T}_3	\mathcal{T}_4	\mathcal{T}_5
P_i	50 s	60 s	100 s	150 s	300 s
W_i/f_{max}	10 s	10 s	15 s	15 s	30 s

The baseline system control algorithm adopted in our experiments is based on the same RTES-EH system architecture. It uses the method proposed in [22] to find a discretely variable frequency allocation and schedule of all task instances in the sensor node. The baseline system incorporates a simple procedure to protect the supercapacitor storage from energy depletion. When the sensor node is about to execute a new task instance, the baseline controller checks whether the current supercapacitor terminal voltage is less than a predefined threshold value. If so, that task instance will be dropped.




Figure 5: Comparison in task instance drop rate between the proposed global control algorithm and baseline algorithm.

Figure 5 illustrates the comparison on the task instance drop rate between the proposed global control algorithm and the baseline system control algorithm. The X-axis of Figure 5 is the initial terminal voltage $V_{cap}(T_{sunrise})$ of the supercapacitor, while the Y-axis is the overall task instance drop rate. Figure 5 shows that the proposed global control algorithm consistently outperforms the baseline algorithm. It achieves a significant reduction in task instance drop rate, by up to 60%, compared with the baseline algorithm.




Figure 6: Comparison on the supercapacitor terminal voltage between the proposed global control algorithm and baseline algorithm.

Figure 6 illustrates the comparison on $V_{cap}(t)$ for $t \in [T_{sunrise}, T_{sunset}]$ between the proposed global control algorithm and the baseline control algorithm. In the left subfigure, $V_{cap}(T_{sunrise}) = 1V$; in the right subfigure, $V_{cap}(T_{sunrise}) = 4V$. Figure 6 shows that the baseline control algorithm greedily executes all the task instances at the beginning of system operation, when the solar irradiance level is relatively low. Consequently, the supercapacitor terminal voltage rapidly decreases, thereby resulting in high task instance drop rate afterwards. On the other hand, the proposed global control algorithm properly maintains the supercapacitor terminal voltage no less than a desirable value, thereby resulting in higher overall system availability and lower task instance drop rate.

8. Conclusion

In this paper, we propose an efficient global control algorithm of a real-time embedded system with energy harvesting (RTES-EH) in order to achieve full system energy autonomy. The global controller performs optimal operating point tracking for the PV panel, SoC management of the supercapacitor, and energy-harvesting-aware real-time task scheduling with DVFS in the sensor node. Experimental results demonstrate that the proposed global control algorithm on RTES-EH significantly lowers the task instance drop rate by up to 60% compared with baseline control algorithms.

Acknowledgement

This work is supported in part by the Software and Hardware Foundations program of the NSF's Directorate for Computer & Information Science & Engineering, the Center for Integrated Smart Sensors funded by the Ministry of Science, ICT & Future Planning as Global Frontier Project (CISS-2012054193). The SPORT lab at USC and ICT at SNU provides research facilities for this study.

References

- [1] Y. H. Lu, L. Benini, and G. De Micheli, "Low-power task scheduling for multiple devices," *Proc. of International Workshop on HW/SW Co-Design*, 2000.
- [2] S. Liu, Q. Qiu, and Q. Wu, "Task merging for dynamic power management of cyclic applications in real-time multi-processor systems," in *ICCD*, 2006.
- [3] Y. Wang, Q. Xie, A. Ammari, and M. Pedram, "Deriving a near-optimal power management policy using model-free reinforcement learning and Bayesian classification," in *DAC*, 2011.
- [4] F. Yao, A. Demers, and S. Shenker, "A scheduling model for reduced CPU energy," *IEEE Foundations of Computer Science*, 1995.
- [5] T. Ishihara and H. Yasuura, "Voltage scheduling problem for dynamically variable voltage processors," in *ISLPED*, 1998.
- [6] G. Quan and X. S. Hu, "Minimum energy fixed-priority scheduling for variable voltage processors," in *DATE*, 2002.
- [7] V. Raghunathan, A. Kansal, et al, "Design considerations for solar energy harvesting wireless embedded systems," *Proc. of IPSN*, 2005.
- [8] X. Jiang, J. Polastre, and D. E. Culler, "Perpetual environmentally powered sensor networks," *Proc. of IPSN*, 2005.
- [9] F. Simjee and P. H. Chou, "Everlast: long-life, supercapacitor-operated wireless sensor node," in *ISLPED*, 2006.
- [10] C. Moser, D. Brunelli, L. Thiele, and L. Benini, "Lazy scheduling for energy-harvesting sensor nodes," *Proc. of Working Conference on Distributed and Parallel Embedded Systems*, 2006.
- [11] S. Liu, Q. Qiu, and Q. Wu, "Energy aware dynamic voltage and frequency selection for real-time systems with energy harvesting," in *DATE*, 2008.
- [12] S. Liu, Q. Wu, and Q. Qiu, "An adaptive scheduling and voltage/frequency selection algorithm for real-time energy harvesting systems," in *DAC*, 2009.
- [13] S. Liu, J. Lu, Q. Wu, and Q. Qiu, "Harvesting-aware power management for real-time systems with renewable energy," *IEEE TVLSI*, 2012.
- [14] Y. Kim, N. Chang, Y. Wang, and M. Pedram, "Maximum power transfer tracking for a photovoltaic-supercapacitor energy system," in *ISLPED*, 2010.
- [15] Y. Choi, N. Chang, and T. Kim, "DC-DC converter-aware power management for low-power embedded systems," *IEEE TCAD*, 2007.
- [16] Y. Wang, X. Lin, Y. Kim, N. Chang, and M. Pedram, "Enhancing efficiency and robustness of a photovoltaic power system under partial shading," in *ISQED*, 2012.
- [17] T. ESRAM, J. Kimball, P. Krein, P. Chapman, and P. Midya, "Dynamic maximum power point tracking of photovoltaic arrays using ripple correlation control," *IEEE Trans. on Power Electronics*, 2006.
- [18] W. Lee, Y. Kim, Y. Wang, et al, "Versatile high-fidelity photovoltaic module emulation system," in *ISLPED*, 2011.
- [19] X. Lin, Y. Wang, S. Yue, D. Shin, N. Chang, and M. Pedram, "Near-optimal, dynamic module reconfiguration in a photovoltaic system to combat partial shading effects," in *DAC*, 2012.
- [20] Y. Wang, Y. Kim, Q. Xie, N. Chang, and M. Pedram, "Charge migration efficiency optimization in hybrid electrical energy storage (HEES) systems," in *ISLPED*, 2011.
- [21] M. Pedram, N. Chang, Y. Kim, and Y. Wang, "Hybrid electrical energy storage systems," in *ISLPED*, 2010.
- [22] W. Kwon and T. Kim, "Optimal voltage allocation techniques for dynamically variable voltage processors," in *DAC*, 2003.