Cost-Constrained QoS Optimization for Approximate Computation Real-Time Tasks in Heterogeneous MPSoCs

Tongquan Wei, Junlong Zhou, Kun Cao, Peijin Cong, Mingsong Chen,
Gongxuan Zhang, X. Sharon Hu, and Jianming Yan

Abstract—Internet of Things (IoT) devices such as video-based detectors or road side units are being deployed in emerging applications like sustainable and intelligent transportation systems. Oftentimes, stringent operation and energy cost constraints are exerted on this type of applications, necessitating a hybrid supply of renewable and grid energy. The key issue of a cost-constrained hybrid of renewable and grid power is its uncertainty in energy availability. The characteristic of approximate computation that accepts an approximate result when energy is limited and executes more computations yielding better results if more energy is available, can be exploited to intelligently handle the uncertainty. In this paper, we first propose an energy-adaptive task allocation scheme that optimally assigns real-time approximate-computation tasks to individual processors and subsequently enables a matching of the cost-constrained hybrid supply of energy with the energy demand of the resultant task schedule. We then present a quality of service (QoS)-driven task scheduling scheme that determines the optional execution cycles of tasks on individual processors for optimization of system QoS. A dynamic task scheduling scheme is also designed to adapt at runtime the task execution to the varying amount of the available energy. Simulation results show that our schemes can reduce system energy consumption by up to 29% and improve system QoS by up to 108% as compared to benchmarking algorithms.

Index Terms—Hybrid Energy Systems, Real-Time MPSoCs, Approximate Computation, QoS Optimization.

I. INTRODUCTION

AS a global infrastructure for information society, IoT has enabled various emerging applications [1]–[3]. Sustainable and intelligent transportation is one of such applications where video-based detection systems or road side unit (RSU)-based vehicular ad-hoc networks are utilized for better traffic planning and management. IoT devices such as video detectors and RSUs consume power of up to 10W [4], resulting in an unaffordable energy cost and thus hindering the deployment of intelligent transportation systems in metropolitan cities. This

T. Wei, K. Cao, P. Cong, and J. Yan are with the Department of Computer Science and Technology, East China Normal University, Shanghai 200062, China. M. Chen is with the Shanghai Key Laboratory of Trustworthy Computing, East China Normal University, Shanghai 200062, China. M. Chen is the corresponding author: mschen@sei.ecnu.edu.cn.

J. Zhou and G. Zhang are with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing 210094, China.

X. S. Hu is with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN 46656, USA.

situation can potentially be mitigated from two approaches. One is for these systems to scavenge energy from renewable generations as complement to the energy drained from power grid. The other one is to perform approximate computing for applications like image or video-based object identification such that results of best quality is obtained by consuming a given amount of energy. In this paper, we are interested in designing such systems that are powered by hybrid energy and supposed to finish under budget a mission with best quality.

The key issue of a hybrid energy system under budget of the energy cost is its uncertainty. This uncertainty stems from the intermittence nature of renewable generation and time-varying pricing of electric grid. A system powered by such hybrid energy may fail to execute a task to completion by deadline due to lack of energy, resulting in a timing fault and degraded system performance. Approximate computation approach [5] can minimize the possibility that a task misses its deadline due to the non-stationarity of powering in hybrid energy systems. In the approximate computation approach, a task is decomposed into a mandatory part followed by an optional part. The mandatory part must execute to completion to produce an acceptable result while the optional part refines the generated result. Based on the observation that a timely approximate result is preferable to a precise result too late [5], the approximate computation technique can be used to avoid timing faults. This is achieved by producing an approximate result of acceptable quality by the deadline when the system cannot produce an exact result in time due to lack of energy. In this paper, we explore the design of a hybrid energy real-time system using the approximate computation technique.

The studied hybrid energy real-time system attempts to combine renewable generation and electric grid with energy storage to deliver a cost-constrained power supply. In fact, real-time embedded systems powered by conventional energy and renewable energy have been studied separately in the recent past. With regard to real-time embedded platforms that are either powered by batteries or directly connected to power grid, their energy management in the last two decades has concentrated on minimizing energy consumption for a longer lifetime and clear financial advantages [6]. On the other hand, for real-time embedded platforms that are powered by renewable generations, the focus of research is mainly on improving the efficiency to take advantage of renewable energy, reducing the capacity of energy storage, and minimizing the deadline miss rate of real-time tasks. For instance, Severini et al. [7]

presented an improved version of lazy scheduling algorithm [8] for energy harvesting systems. The improved solution can foresee at runtime the task energy starving, hence allows obtaining a more conservative and efficient management of energy with respect to the original solution. Chetto et al. [9] showed that when the incoming ambient energy or task arrivals have a pure stochastic nature, the EDF scheduling algorithm remains an attractive scheduler owing to easy implementation and no need for estimating the stored or harvested energy. Abdeddaim et al. [10] designed an optimal fixed-priority solution to the real-time scheduling problem that handles both energy and timing constraints for energy harvesting systems.

All the above works on renewable powered systems focus on improving energy efficiency and schedule timeliness for energy harvesting embedded systems, however, the concept of approximate execution is not considered. Stavrinides et al. [11] combined the approximate computation and bin packing strategy to evaluate the impact of input error on the performance of a heterogeneous distributed real-time system, however, the energy design constraint is not taken into account.

It is natural for approximate computing systems to maximize their QoS given a certain cost or energy budget. Cortes et al. [12] designed an approach to maximize rewards for real-time approximate computation systems. The voltage at which each task runs and the number of optional cycles are determined under the timing and energy constraint. Similarly, Yu et al. [13] presented a runtime scheduling algorithm for approximate computation tasks. The algorithm maximizes system QoS under energy constraints by optimally redistributing slack generated at runtime. Since these approaches assume a constant constraint of energy, they are not well suited for cost-constrained hybrid energy systems where the available energy varies due to intermittence of renewable and dynamic pricing of grid power. Considering the uncertainty in availability of energy in embedded harvesting systems, Kooti et al. [14] designed an energy management technique to maximize system QoS. However, the scheme specifically targets applications in which QoS constraints allow drop out of some of real-time tasks, which doesnot comply with the stringent timing requirements of general real-time systems. In addition, the scheme is based on a platform of single processor only powered by harvesting energy, thus are not applicable for the state-of-the-art hybrid energy multi-core processor platform.

In this paper, we propose an approximate computation based task allocation and scheduling scheme for a heterogeneous MPSoC real-time system that is powered by a hybrid of energy harvested from environments and drained from power grid. The major contributions are summarized as follows.

- We present a static energy-efficient task allocation scheme that adaptively assigns approximate computation tasks to individual processors considering the uncertainty in the hybrid of the renewable and grid energy. The energy-adaptive task allocation scheme can effectively enable the matching of the hybrid energy supply with the energy consumption of the resultant task schedule.
- We propose a static scheduling scheme for tasks allocated to individual processors. The proposed scheduling scheme determines the number of optional cycles of


Fig. 1: The diagram of the system architecture.

each task to be executed such that the system QoS is maximized under budget of the cost including energy and battery aging. The static task schedule is also adapted at runtime to the fluctuating amount of available energy.
- We conduct extensive simulation experiments to verify the effectiveness of the proposed schemes in improving system energy efficiency and QoS. Simulation results have demonstrated that the proposed schemes have better performance as compared to benchmarking schemes.

The rest of the paper is organized as follows. Section II introduces the system architecture and models, and Section III defines the problem and gives the overall framework of the proposed scheme. Section IV presents the proposed uncertainty-aware and cost-constrained task allocation, and Section V describes the proposed QoS-driven task scheduling. The effectiveness of the proposed approach is verified in Section VI and concluding remarks are given in Section VII.

II. SYSTEM ARCHITECTURE AND MODELS

The system consists of four major parts: the energy source module, storage module, dissipation module, and power grid. As shown in Fig. 1, the energy source module automatically scavenges renewable energy from environments at the power of $P_{harv}(t)$, and converts the renewable generation into electrical energy. The energy storage module, which is typically in the form of a battery, serves as a buffer against the uncertainty in harvested energy. When the harvested energy is more than the energy consumed by the dissipation module, the extra energy is stored in the storage module. A heterogeneous multiprocessor system-on-chip (MPSoC) is considered as the energy dissipation module, which drains energy at the rate of $P_{cons}(t)$ from the energy source, storage module, and/or power grid that provide a power supply of $P_{sup}(t)$. We are interested in a system that is designed under a cost budget. We assume that the harvested energy is free, thus, the amount of energy drained from the battery and/or grid depends upon the cost of battery aging and the dynamic price of grid electricity.

A. Processor and Task Model

The MPSoC system consists of M processors, denoted by $\Theta = \{\Theta_1, \Theta_2, \cdots, \Theta_M\}$, where each processor Θ_m $(1 \leq m \leq M)$ is a typical DVFS-enabled processor that can operate with a set of discrete supply voltage and frequency pairs $(v_{m,r}, f_{m,r})$ $(1 \leq r \leq x_m)$, where $v_{m,1} < \cdots < v_{m,r} < \cdots < v_{m,x_m}$, $f_{m,1} < \cdots < f_{m,r} < \cdots < f_{m,x_m}$, and x_m is the voltage/frequency level of Θ_m. Real-time tasks

are supposed to schedule and execute on the system. Consider a task set Γ consisting of N independent real-time tasks $\{\tau_1, \tau_2, \cdots, \tau_N\}$. The task set Γ is a frame-based task set, in which all tasks share a common deadline d that is also the frame. Different tasks exhibit different power consumptions on the same processor, even executing at the same operating speed and temperature. This is due to the fact that power consumptions of tasks strongly rely on circuit activities and usage patterns of different functional units [15]. Thus, the activity factor of a task, denoted by μ (ranging in (0,1]), is introduced to capture how intensively functional units are being utilized by the task [16].

We consider approximate computation tasks in this work. Each task τ_i ($1 \leq i \leq N$) is logically decomposed into a mandatory part with execution cycles M_i and an optional part with execution cycles O_i [17]. The mandatory part must execute to completion before the deadline and generate an acceptable result, while the optional part refines and improves the result. The characteristic of an approximate computation modeled task τ_i is therefore described by a quadruple $\tau_i : \{\mu_i, M_i, O_i, d\}$, where μ_i is the activity factor, and d is the common deadline. M_i is the mandatory cycles of τ_i that must be completed before the deadline while O_i is the maximum optional cycles of τ_i. Since optional cycles are partially executed, we introduce a variable to represent the executed optional cycles of τ_i, which is denoted as o_i and holds for $0 \leq o_i \leq O_i$. Then, the actual length l_i of τ_i, measured by the total execution cycles, can be expressed as

$$l_i = M_i + o_i. \quad (1)$$

B. Battery SoH Degradation Model

A practical and widely used SoH degradation model of Li-ion batteries [18]–[20] is adopted to estimate the SoH degradation for cycled charging and discharging of a Li-ion battery cell. Before presenting the estimation of SoH degradation, we formally define the state-of-charge (SoC) of a battery at first, that is,

$$SoC = \frac{C_{bat}}{C_{full}} \times 100\%, \quad (2)$$

where C_{bat} is the amount of charge stored in the battery and C_{full} is the battery full charge capacity. The amount of state-of-health (SoH) degradation, denoted by D_{SoH}, is defined as

$$D_{SoH} = \frac{C_{full}^{nom} - C_{full}}{C_{full}^{nom}} \times 100\%, \quad (3)$$

where C_{full}^{nom} is nominal value of C_{full} for a fresh new battery.

The adopted SoH degradation model estimates the battery SoH degradation in a cycled charging/discharging pattern, where a (charging/discharging) cycle is defined as a charging process of the battery cell from SoC_{low} to SoC_{high} followed by a discharging process from SoC_{high} to SoC_{low}. Thus the average SoC and SoC swing in a cycle are calculated as

$$SoC_{avg} = (SoC_{low} + SoC_{high})/2, \quad (4)$$

$$SoC_{swing} = SoC_{high} - SoC_{low}. \quad (5)$$

The battery SoH degradation $D_{SoH,cycle}$ during one cycle, accounting for SoC_{avg} and SoC_{swing}, is given by

$$D_1 = K_{co} \cdot e^{(SoC_{swing}-1) \cdot \frac{T_{ref}}{K_{ex} \cdot T_B}} + 0.2\frac{\eta}{\eta_{life}}, \quad (6)$$

$$D_2 = D_1 \cdot e^{4K_{SoC} \cdot (SoC_{avg}-0.5)} \cdot (1 - D_{SoH}), \quad (7)$$

$$D_{SoH,cycle} = D_2 \cdot e^{K_T \cdot (T_B - T_{ref}) \cdot \frac{T_{ref}}{T_B}}, \quad (8)$$

where K_{co}, K_{ex}, K_{SoC}, and K_T are battery specific parameters. T_B and T_{ref} are battery's operation temperature and reference temperature, respectively. η is the duration of this charging/discharging cycle and η_{life} is the calendar life of this battery. The total SoH degradation (in reference to a fresh battery) after $\mathcal{W}$ charging and discharging cycles is

$$D_{SoH} = \sum_{w=1}^{\mathcal{W}} D_{SoH,cycle}(w), \quad (9)$$

where $D_{SoH,cycle}(w)$ is the SoH degradation in the w^{th} cycle.

However, the above SoH degradation model can only be applied to the fixed charging/discharging pattern, that is, the battery experiences the cycles with the same SoC_{swing} and SoC_{avg}. In reality, a battery may not follow this pattern. Hence, a cycle-decoupling method was proposed [19] to build an improved SoH degradation model of battery for charging/discharging cycles with arbitrary patterns. The SoH degradation in an arbitrary cycle and the total degradation can thus be derived using Eq. (8) and (9), respectively.

C. Energy Model

The energy of the concerned system is modeled from the perspective of both supply and demand. We first describe the model for energy supply. Let $P_{harv}(t)$ be the harvesting power and $E_{harv}(t_1, t_2)$ be the energy scavenged from environments during time interval $[t_1, t_2]$, then $E_{harv}(t_1, t_2)$ is calculated as

$$E_{harv}(t_1, t_2) = \int_{t_1}^{t_2} P_{harv}(t)dt. \quad (10)$$

As illustrated in Fig. 1, the energy source module, storage module, and power grid can provide the energy to the dissipation module. Let $E_{sup}(t_1, t_2)$ represent the system available supply energy during time interval $[t_1, t_2]$, and $E_{bat}(t_1, t_2)$ and $E_{grid}(t_1, t_2)$ indicate the energy drained from the battery and power grid during $[t_1, t_2]$, respectively, then we have

$$E_{sup}(t_1, t_2) = E_{harv}(t_1, t_2) + E_{bat}(t_1, t_2) + E_{grid}(t_1, t_2). \quad (11)$$

We then describe the model for energy demand. The total energy consumption of an IoT device depends on its processors, memory, disks, cooling system, and wireless components. Although there are many research work [21]–[23] focusing on reducing system overall energy consumption, we only consider processor energy consumption in the energy dissipation model since processor is the dominant source of overall energy consumption of some systems like IoT security authentication system [24]–[26]. The power consumption of a processor can be modeled as the sum of static/leakage power consumption P_{sta} and dynamic power consumption P_{dyn}. The static power

P_{sta} is temperature dependent and consumed by the leakage current required to maintain basic state of circuits. The leakage current changes super linearly with temperature. The dynamic power P_{dyn} is related to processor switching activity and can be estimated by a function of supply voltage v and frequency f, that is, $P_{dyn} \propto v^2 f$. Thus, the overall power consumption of processor Θ_m when executing task τ_i at the supply voltage/frequency $(v_{m,r}, f_{m,r})$ is formulated as

$$
\begin{aligned}
P_{cons}(\tau_i, \Theta_m) =& \omega_1 \cdot v_{m,r} + \omega_2 \cdot v_{m,r} \cdot T_m(t) \\
&+ C_{eff,m} \cdot \mu_i \cdot v_{m,r}^2 \cdot f_{m,r},
\end{aligned} \tag{12}
$$

where ω_1 and ω_2 are both non-negative architecture-dependent constants of Θ_m, $T_m(t)$ is the operating temperature of Θ_m at time instance t, $C_{eff,m}$ is the effective switching capacitance of Θ_m, and μ_i is the activity factor of τ_i.

The static power is always consumed to maintain basic circuits and the dynamic power is only consumed when executing tasks. Thus, based on Eq. (12), the total energy consumed (or demanded) by the system during the scheduling horizon $\mathcal{H}$, denoted by E_{dem}, is calculated as

$$
\begin{aligned}
E_{dem} =& \sum_{m=1}^{M} \left(\omega_1 \cdot v_{m,r} \cdot \mathcal{H} + \int_0^{\mathcal{H}} \omega_2 \cdot v_{m,r} \cdot T_m(t) dt \right) \\
&+ \sum_{m=1}^{M} \sum_{\tau_i \in \Gamma_m} \left(C_{eff,m} \cdot \mu_i \cdot v_{m,r}^2 \cdot l_i \cdot \frac{\mathcal{H}}{d} \right),
\end{aligned} \tag{13}
$$

where Γ_m is the subset of tasks allocated to processor Θ_m. Oftentimes, $\mathcal{H}$ is a multiple of d.

III. PROBLEM DEFINITION AND OVERALL FRAMEWORK

Our goal is to reduce the impact of uncertainties in the energy scavenged from ambient environments, and improve the efficiency of renewable energy and QoS of the system under a certain cost budget $Cost_{bdg,\mathcal{H}}$. The following subsections describe the system cost and QoS function, formulate the QoS optimization problem, and overview the proposed solution.

A. System Cost Function

We adopt a slotted time model that deals with all the system decisions and constraints in discrete time intervals of equal length. More specifically, the whole scheduling horizon $\mathcal{H}$ is divided into $\mathcal{W}$ time slots with equal and constant length of Δt. We assume that the power grid has a dynamic pricing function, and the price of one unit of energy (in kWh) during the w^{th} time slot is denoted by $PX_{grid}[w]$. In our formulation, solar energy is considered free for simplicity, though its cost can be normalized and incorporated in the formulation. Since solar energy is free, the total cost function $Cost_{tot,\mathcal{H}}$ during $\mathcal{H}$ is comprised of two parts. One is the cost $Cost_{aging,\mathcal{H}}$ associated with battery aging and the other is the energy cost $Cost_{grid,\mathcal{H}}$ due to charges from power grid. That is,

$$
Cost_{tot,\mathcal{H}} = Cost_{aging,\mathcal{H}} + Cost_{grid,\mathcal{H}}. \tag{14}
$$

Let $Cost_{bat}$ be the cost to purchase a fresh new battery, $D_{SoH,\mathcal{H}}$ be the amount of SoH degradation during $\mathcal{H}$, and $D_{SoH,end}$ be the amount of SoH degradation indicating the end-of-life of a battery, the cost of battery aging during $\mathcal{H}$ is

$$
Cost_{aging,\mathcal{H}} = Cost_{bat} \frac{D_{SoH,\mathcal{H}}}{D_{SoH,end}}. \tag{15}
$$

Oftentimes, $Cost_{bat}$, the cost to purchase and replace a battery is fixed, and the $D_{SoH,end}$ indicating the end-of-life of a battery takes value of 70% [20]. Since $Cost_{bat}$ and $D_{SoH,end}$ are assumed constant, it can be seen from Eq. (15) that the cost of battery aging can be minimized by optimizing $D_{SoH,\mathcal{H}}$, the amount of SoH degradation during the scheduling horizon $\mathcal{H}$.

For the w^{th} time slot, let $P_{sup}[w]$ be the power supplied to the energy dissipation module, $P_{harv}[w]$ be the harvesting power of energy source module, and $P_{bat}[w]$ be the power drained from the battery. We assume that the grid electricity price $PX_{grid}[w]$ and the battery aging cost in the w^{th} time slot is fixed. Then we define a binary variable β that is set to 1 when the electricity cost in a time slot is greater than the battery aging cost, and 0 otherwise. When the incurred cost of grid electricity is different from the battery aging cost in a time slot, the energy source of lower cost is selected to power the MPSoC. Given these and based on the principle of energy conservation illustrated in Fig. 1, the grid power can then be derived as $P_{sup}[w] - P_{harv}[w] - \beta \cdot P_{bat}[w]$, and its cost is

$$
\begin{aligned}
Cost_{grid,\mathcal{H}} = \sum_{w=1}^{\mathcal{W}} PX_{grid}[w](P_{sup}[w] - P_{harv}[w] \\
- \beta \cdot P_{bat}[w])\Delta t.
\end{aligned} \tag{16}
$$

From the viewpoint of energy dissipation module, $P_{sup}[w]$ can also be written in the form of $P_{dem}[w] + \Delta P[w]$, that is, $P_{sup}[w] = P_{dem}[w] + \Delta P[w]$, where $\Delta P[w]$ is the difference between the power supply and demand of energy dissipation module. Thus, Eq. (16) can be expressed as

$$
\begin{aligned}
Cost_{grid,\mathcal{H}} = \sum_{w=1}^{\mathcal{W}} PX_{grid}[w](P_{dem}[w] \\
+ \Delta P[w] - P_{harv}[w] - \beta \cdot P_{bat}[w])\Delta t.
\end{aligned} \tag{17}
$$

The solar harvesting profile $P_{harv}[w]$ can be estimated using historical data. The power drain of battery $P_{bat}[w]$ can be derived by an SoH degradation-aware battery management policy [19] that minimizes the amount of SoH degradation $D_{SoH,\mathcal{H}}$ and hence $Cost_{aging,\mathcal{H}}$. As the cost of battery aging $Cost_{aging,\mathcal{H}}$ is derived and the cost budget $Cost_{bdg,\mathcal{H}}$ is given, the cost of grid power $Cost_{grid,\mathcal{H}}$ is then determined. In Eq. (17), since $Cost_{grid,\mathcal{H}}$, $P_{harv}[w]$, $P_{bat}[w]$, $PX_{grid}[w]$, Δt are all known, $P_{dem}[w] + \Delta P[w]$ is decided. When $P_{dem}[w]$ is optimized and $\Delta P[w] = 0$, the energy consumed per execution cycle is minimized and the energy supply matches with the energy demand. In other words, more energy can be used to execute task optional cycles for QoS improvement.

B. System QoS Function

It has been shown that the quality of a task can be represented as a linear or concave function of optional cycles of the task [5]. The more cycles the optional part of the task executes, the higher QoS the task generates. Thus, we quantitatively define a simple yet effective QoS function for a system, which is the sum of the executed CPU cycles of optional parts of all the
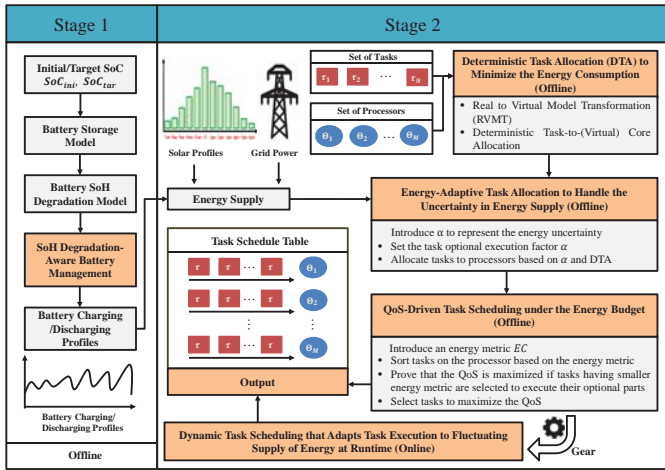
Fig. 2: The main design flow of the proposed scheme.

real-time tasks. It is denoted by Q and is expressed as

$$Q = \sum_{i=1}^{N} o_i. \tag{18}$$

Clearly, the system QoS is a function of executed optional cycles of real-time tasks.

C. Cost-Constrained QoS Optimization Problem

Given a set of approximate-computation real-time tasks and a set of heterogeneous processors powered by renewable and grid energy, design a task-to-processor allocation that adapts to the uncertainty of system available energy, and a task scheduling scheme that arranges assigned tasks on individual processors to maximize the system QoS under the constraint of a cost budget. We assume that the system energy demand cannot exceed the energy supply, and the mandatory parts of all the tasks must be finished before the common deadline d. The problem is formulated into the below form.

$$\text{Maximize:} \quad Q = \sum_{i=1}^{N} o_i$$
$$\text{Subject to:} \quad 0 \leq o_i \leq O_i \tag{19}$$
$$E_{dem} \leq E_{sup} \tag{20}$$
$$\sum_{\tau_i \in \Gamma_m} M_i / f(\tau_i) \leq d \tag{21}$$
$$Cost_{tot,\mathcal{H}} \leq Cost_{bdg,\mathcal{H}}. \tag{22}$$

D. Overview of the Proposed Two-Stage Scheme

The key issue of the studied system is that uncertainties in the energy scavenged from ambient environments result in low energy utilization and/or high deadline miss rate. We address the problem by designing approximate computation-based task allocation and scheduling algorithms that enhance the energy efficiency, ensure system timing constraints, and maximize the system QoS under the constraint of a cost budget. The design flow of the proposed scheme is illustrated in Fig. 2. Given cost budget of a system, the uncertainties in energy supply of energy dissipation modules stem from the intermittence of renewable energy and dynamic pricing of grid energy. We tackle the uncertainties in stage 1 of the scheme by adopting an

SoH degradation-aware battery management policy [19] that minimizes the cost of battery aging, and conducting an energy-adaptive task allocation that minimizes the system energy demand and adapts the energy demand to the availability of energy supply including the grid and renewable energy.

In stage 1 of the proposed scheme, we first adopt an SoH degradation-aware battery management policy [19] to reduce the cost of battery aging. Based on the battery storage and SoH degradation model adopted, the SoH degradation-aware battery management policy [19] takes the battery initial SoC and target SoC as input and determines the battery charging/discharging current profile using convex optimization techniques. The generated charging/discharging current profile is a near-optimal charge management policy that can minimize the battery SoH degradation, thus extend the cycle life of battery. For a given cost budget formulated in Eq. (14), it is preferable to utilize free renewable energy, or battery storage and/or grid power, whichever incurs lower cost during the scheduling horizon $\mathcal{H}$. The cost of battery aging and grid power are given in Eq. (15) and (17), respectively.

We then design an energy adaptive task allocation (ATA) algorithm in stage 2 of the proposed scheme. The algorithm first introduces and initializes a variable α, named the task optional execution factor, to represent the uncertainty in energy supply. It allocates tasks to processors using a deterministic task allocation (DTA) scheme that consists of Real_to_Virtual_Model_Transformation (RVMT) and Deterministic Task-to-(Virtual) Core Allocation [27] based on the α. Then, it adapts system energy demand to renewable and grid supply of energy by using the task optional execution factor as a control knob to maximize the degree of match between system energy demand and supply. Through the deterministic task allocation and task optional execution factor tuning, the system energy consumption is minimized, the energy supply is fully utilized, and the cost of grid energy is reduced.

The stage 2 of the proposed scheme aims at improving the QoS of the system after SoH degradation-aware battery management policy minimizes the cost of battery aging and the energy adaptive task allocation scheme reduces the cost of grid energy. The QoS-driven task scheduling algorithm proposed in this stage first introduces an energy metric that indicates the importance of every task on the processor in terms of improving system QoS, then proves that the system QoS is maximized if tasks having smaller energy metric are selected to execute their optional parts. It finally develops a task selection scheme that chooses the tasks with smaller metric under the constraint of cost budget.

The above QoS-driven task scheduling is designed for the offline scenario under the assumption of a fixed energy supply. However, the actual energy harvested from environments may not be the same as the energy estimated by prediction techniques due to the fluctuating nature of renewable generation. In addition, the cost-constrained energy drained from the battery and power grid may vary due to different initial SoC of the battery and time varying price of grid power. Therefore, an online QoS-driven task scheduling is designed to adapt the offline schedules to fluctuating energy supply at runtime.

IV. TASK ALLOCATION UNDER COST CONSTRAINT

The task allocation scheme first assumes deterministic energy sources and minimizes the energy consumption by intelligently assigning tasks to processors, then iteratively adapts the task allocation to uncertainties in energy availability.

A. DTA to Minimize Energy Consumption

The concerned system consists of heterogeneous multiprocessors, each of which is DVFS-enabled and supports a set of discrete supply voltage and frequency pairs. Hence, operating frequencies of tasks need to be determined during task allocation. To this end, we propose an RVMT method that converts the processor model with multiple voltage and frequency levels to a model of multiple virtual cores, each of which has a fixed supply voltage and frequency level. This method can effectively decrease one dimension of optimization for energy consumption by reducing task-to-(real) processor allocation and frequency selection problem to task-to-(virtual) core allocation problem. Below, we first show the proposed RVMT processor model transformation, analyze the energy optimality of allocating tasks to virtual cores and present a theorem on optimum task allocation, then develop a task-to-(virtual) core allocation heuristic based on the theorem.

1) Real_to_Virtual Model Transformation (RVMT): The system energy consumption during the scheduling horizon $\mathcal{H}$ given in Eq. (13) can be re-written as

$$E_{dem} = \sum_{m=1}^{M}\left(\omega_1 \cdot v_{m,r} \cdot \mathcal{H} + \int_0^{\mathcal{H}} \omega_2 \cdot v_{m,r} \cdot T_m(t)dt\right) + \frac{\mathcal{H}}{d}\sum_{m=1}^{M}\sum_{r=1}^{x_m}\left(C_{eff,m}\cdot v_{m,r}^2 \sum_{\tau_i \in \Gamma_{m,r}}\mu_i \cdot l_i\right), \quad (23)$$

where $\Gamma_{m,r}$ is the subset of tasks allocated to processor Θ_m and operated at frequency $f_{m,r}$. The two items of Eq. (23) are the static and dynamic energy consumption, respectively. The static energy consumption depends on temperature while the dynamic energy consumption depends on task allocation. Thus both thermal-aware task scheduling and energy-efficient task allocation are helpful to minimize the system overall energy consumption. Since thermal-aware task scheduling for minimizing the static energy consumption has been well studied in [27], we focus on energy-efficient task allocation to minimize the dynamic energy consumption in this paper.

Let E_{dyn} represent the dynamic energy consumption, which can be formulated into the product of two vectors, that is,

$$E_{dyn} = \frac{\mathcal{H}}{d}\sum_{m=1}^{M}\sum_{r=1}^{x_m}(C_{eff,m}v_{m,r}^2\sum_{\tau_i \in \Gamma_{m,r}}\mu_i l_i)$$
$$= \frac{\mathcal{H}}{d}\sum_{m=1}^{M}\sum_{r=1}^{x_m}\mathcal{G}_{m,r}\mathcal{Q}_{m,r} = \frac{\mathcal{H}}{d}\mathcal{G}\mathcal{Q}, \quad (24)$$

where $\mathcal{G}_{m,r} = C_{eff,m}v_{m,r}^2$, $\mathcal{Q}_{m,r} = \sum_{\tau_i \in \Gamma_{m,r}}\mu_i l_i$, M is the number of processors, and x_m is the number of frequency levels supported by the m^{th} processor. Vector $\mathcal{G} = [\mathcal{G}_{1,1}, \mathcal{G}_{1,2}, \cdots, \mathcal{G}_{1,x_1}, \cdots\cdots, \mathcal{G}_{M,1}, \mathcal{G}_{M,2}, \cdots, \mathcal{G}_{M,x_M}]^T$ captures processor dependent parameters and vector $\mathcal{Q} = [\mathcal{Q}_{1,1}, \mathcal{Q}_{1,2}, \cdots, \mathcal{Q}_{1,x_1}, \cdots\cdots, \mathcal{Q}_{M,1}, \mathcal{Q}_{M,2}, \cdots, \mathcal{Q}_{M,x_M}]$ captures task related parameters. Here, $\mathcal{G}_{m,r} \in \mathcal{G}$ is referred to as the power dissipation factor of processor Θ_m, $\mathcal{Q}_{m,r} \in \mathcal{Q}$
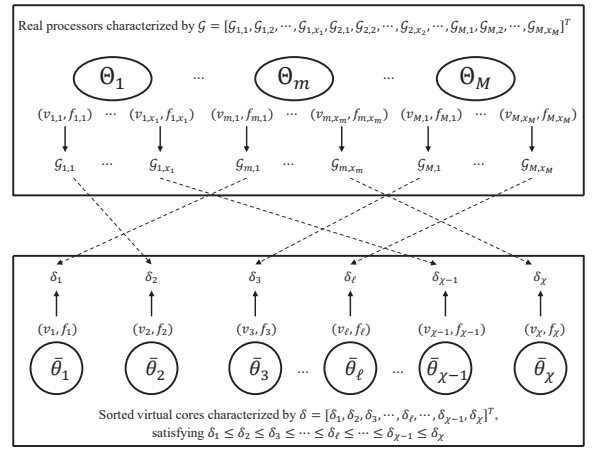

Fig. 3: Illustration of the **RVMT** transformation.

is referred to as the power dissipation factor of subset $\Gamma_{m,r}$ allocated to processor Θ_m at frequency level $f_{m,r}$, and $\mu_i l_i$ is referred to as the power dissipation factor of task τ_i. It is clear that $\mathcal{G}$ is constant since $C_{eff,m}$ and $v_{m,r}$ are known for the given MPSoC system Θ while $\mathcal{Q}$ is not since it depends on the task assignment and frequency selection (e.g., $\Gamma_{m,r}$). In addition, for a given task set Γ, the sum of power dissipation factor of tasks in the set is a constant (denoted by $\mathcal{V}$) and can be expressed as $\sum_{m=1}^{M}\sum_{r=1}^{x_m}\mathcal{Q}_{m,r} = \sum_{m=1}^{M}\sum_{r=1}^{x_m}\sum_{\tau_i \in \Gamma_{m,r}}\mu_i l_i = \sum_{i=1}^{N}\mu_i l_i = \mathcal{V}$.

According to the above formulation, the power dissipation of the MPSoC system can be characterized by vector $\mathcal{G} = [\mathcal{G}_{1,1}, \mathcal{G}_{1,2}, \cdots, \mathcal{G}_{1,x_1}, \cdots\cdots, \mathcal{G}_{M,1}, \mathcal{G}_{M,2}, \cdots, \mathcal{G}_{M,x_M}]^T$. Since the number of power dissipation factors in the vector $\mathcal{G}$ is greater than the number of processors supporting multiple voltage levels, we propose a model transformation method that converts the DVFS-enabled MPSoC system of multiple voltage levels into a virtual core system, each core of which supports only one voltage level. During the transformation, the virtual cores obtained from a certain processor in the MPSoC system are assumed to share the same characteristics as the processor except for the supply voltage and frequency. As a result, each element in the power dissipation vector $\mathcal{G}$ corresponds to a core in the virtual core system. We then sort the cores in the increasing order of power dissipation factors.

Fig. 3 illustrates the proposed RVMT processor model transformation. As demonstrated in the figure, the virtual core system can be represented by $\bar{\Theta} = \{\bar{\theta}_1, \bar{\theta}_2, \cdots, \bar{\theta}_\ell, \cdots, \bar{\theta}_\mathcal{X}\}$, where every core $\bar{\theta}_\ell \in \bar{\Theta}$ ($1 \le \ell \le \mathcal{X}$, and $\mathcal{X} = \sum_{m=1}^{M}x_m$) has a fixed voltage and frequency pair (v_ℓ, f_ℓ). The virtual core system $\bar{\theta}$ is then characterized by vector $\delta = [\delta_1, \delta_2, \cdots, \delta_\ell, \cdots, \delta_\mathcal{X}]^T$ and $\delta_1 \le \delta_2 \le \cdots \le \delta_\ell \le \cdots \le \delta_\mathcal{X}$ holds, where $\delta_\ell = C_{eff,\ell}v_\ell^2$ is referred to as the power dissipation factor of $\bar{\theta}_\ell$. Accordingly, the power dissipation of subsets allocated to virtual cores can be represented by vector $\xi = [\xi_1, \xi_2, \cdots, \xi_\ell, \cdots, \xi_\mathcal{X}]$, where $\xi_\ell = \sum_{\tau_i \in \Gamma_\ell}\mu_i l_i$ is referred to as the power dissipation factor of subset allocated to $\bar{\theta}_\ell$, and $\xi_1 + \cdots + \xi_\ell + \cdots + \xi_\mathcal{X} = \mathcal{V}$ holds. Thus the dynamic energy consumption given in Eq. (24) can be expressed as

$$E_{dyn} = \frac{\mathcal{H}}{d}\delta\xi. \quad (25)$$

Algorithm 1: Deterministic task-to-core assignment

Input: 1) Task set represented by $\Gamma = \{\tau_1, \tau_2, \cdots, \tau_N\}$
    2) MPSoC represented by $\Theta = \{\Theta_1, \Theta_2, \cdots, \Theta_M\}$
Output: task-to-(virtual) allocation $\{\Gamma_1, \Gamma_2, \cdots, \Gamma_\mathcal{X}\}$

1 transform the real processor system Θ to virtual core system $\bar{\Theta}$ by $\bar{\Theta} = \mathrm{RVMT}(\Theta)$, where the $\mathcal{X}$ cores of $\bar{\Theta}$ are in increasing order of their power dissipation factors;
2 **for** $\ell = 1$ to $\mathcal{X}$ **do**
3   initialize the utilization and task subset of virtual core $\bar{\theta}_\ell \in \bar{\Theta}$ by $U(\bar{\theta}_\ell) = 0$ and $\Gamma_\ell = \varnothing$;
4 $\ell = 1$;
5 **while** $\Gamma \neq \varnothing$ and $\ell \leq \mathcal{X}$ **do**
6   sort the tasks $\tau_i \in \Gamma$ in the decreasing order of task power dissipation factors $\mu_i l_i$ using heapsort;
7   **for** $i = 1$ to $size(\Gamma)$ **do**
     /* use First-Fit to group tasks */
8   **if** $\mathrm{RT}(\tau_i, \bar{\theta}_\ell) == \textbf{\textit{true}}$ and $u(\tau_i, \bar{\theta}_\ell) + U(\bar{\theta}_\ell) \leq 1$ **then**
9   assign task τ_i to core $\bar{\theta}_\ell$;
10   update Γ, Γ_ℓ, and $U(\bar{\theta}_\ell)$ by $\Gamma = \Gamma - \tau_i$, $\Gamma_\ell = \Gamma_\ell + \tau_i$, and $U(\bar{\theta}_\ell) = U(\bar{\theta}_\ell) + u(\tau_i, \bar{\theta}_\ell)$;
11    $\ell \leftarrow \ell + 1$;
12 **if** $\Gamma \neq \varnothing$ and $\ell > \mathcal{X}$ **then**
13   **exit(1)**;   /* Exit when infeasible */

Let $\xi^* = [\xi_1^*, \xi_2^*, \cdots, \xi_\ell^*, \cdots, \xi_\mathcal{X}^*]$ be a vector that denotes the optimum power dissipation factors of subsets allocated to individual virtual cores, and E_{dyn}^* be the dynamic energy consumption of the optimum task allocation. We prove below that the dynamic energy consumption E_{dyn} given in Eq. (25) is minimized when the virtual core with smaller power dissipation factor (δ) ends up with the subset of its allocated tasks having a larger power dissipation factor (ξ^*).

Theorem 1: *If the virtual core power dissipation factors $\delta_1 \leq \delta_2 \leq \cdots \leq \delta_\mathcal{X}$ hold for $\delta = [\delta_1, \delta_2, \cdots, \delta_\mathcal{X}]^T$, and the sum of the corresponding task subset power dissipation factors $\xi_1^* + \xi_2^* + \cdots + \xi_\mathcal{X}^*$ is fixed for $\xi^* = [\xi_1^*, \xi_2^*, \cdots, \xi_\mathcal{X}^*]$, then the dynamic energy consumption E_{dyn} is minimized to E_{dyn}^* if $\xi_1^* \geq \xi_2^* \geq \cdots \geq \xi_\ell^* \geq \cdots \geq \xi_\mathcal{X}^*$ holds.*

2) Deterministic Task-to-(Virtual) Core Allocation: We propose a sub-optimal task allocation heuristic which is motivated by the theorem presented in Section IV-A1, that is, allocating the subset having a larger power dissipation factor to the virtual core having a smaller power dissipation factor can minimize the dynamic energy consumption. The heuristic operates as follows. Tasks in the subset with the maximum power dissipation factor is allocated to the virtual core with the minimum power dissipation factor, and tasks in the subset with the next maximum power dissipation factor is allocated to the virtual core with the next minimum power dissipation factor. This process repeats until all subsets of tasks are allocated to cores. In addition, the constraints of task deadline and processor capacity limit are examined during the allocation.

The details of the heuristic are described in Alg. 1. Inputs to the algorithm are task set Γ and processor set Θ. Let $\bar{\Theta}$ be the corresponding virtual core system, and $\bar{\theta}_\ell$ be the ℓ^{th} virtual core of the virtual system model. Let $U(\bar{\theta}_\ell)$ be the utilization of virtual core $\bar{\theta}_\ell$, which is calculated as the sum of utilization of tasks assigned to the core, and $u(\tau_i, \bar{\theta}_\ell)$ be the utilization of task τ_i if it is assigned to virtual core $\bar{\theta}_\ell$, which is calculated as the quotient of task execution time and frame size.

The algorithm first transforms the MPSoC to a virtual core system using $\bar{\Theta} = \mathrm{RVMT}(\Theta)$, in which the cores are arranged in the increasing order of power dissipation factors. The utilization and subset of all virtual cores $\bar{\theta}_\ell$ $(1 \leq \ell \leq \mathcal{X})$ are initialized by $U(\bar{\theta}_\ell) = 0$ and $\Gamma_\ell = \varnothing$, respectively. It then iteratively implements the process of task-to-(virtual) core allocation if the task set is not empty and not all the cores in $\bar{\Theta}$ have been considered. Before starting the iteration, the core index ℓ is set to 1. In each round of iteration, the algorithm first sorts the tasks of set Γ in the decreasing order of task power dissipation factors $\mu_i l_i$, then allocates tasks to cores and construct subset of tasks in a first fit manner under the real-time constraint $\mathrm{RT}(\tau_i, \bar{\theta}_\ell) ==$ **true** and processor capacity limit $u(\tau_i, \bar{\theta}_\ell) + U(\bar{\theta}_\ell) \leq 1$. $\mathrm{RT}(\tau_i, \bar{\theta}_\ell)$ is a procedure that checks the real-time constraint of a task and returns **true** if the task is finished before its deadline, and **false** otherwise. If allocating task τ_i to core $\bar{\theta}_\ell$ can satisfy these constraints, the task is allocated to the core, and task set Γ, subset Γ_ℓ, and utilization $U(\bar{\theta}_\ell)$ are hence updated. The process then moves to the next iteration and considers the allocation of the next task. Otherwise, the task is not assigned and the process directly moves to the next iteration. If there is no feasible schedule for the system under the constraints, the algorithm exits.

B. ATA to Handle Uncertainties in Energy Availability

Based on the deterministic task allocation, this section handles the uncertainty in energy sources by adapting the execution of approximate-computation tasks to the energy availability. We introduce a variable α to denote the ratio of the number of executed optional cycles to the maximum optional cycles of a task. The α is named as task optional execution factor and falls within the range [0, 1]. Using α, the total execution cycles of task τ_i given in Eq. (1) can be re-written as

$$l_i = M_i + \alpha \cdot O_i. \tag{26}$$

The key issue in dealing with the energy uncertainty is to derive the relationship between task optional execution factors and energy demand of the system. In other words, the optional execution factor of a task is expected to serve as a coarse adjustment knob to control the degree of match between system energy demand and supply. To this end, we first show that the energy demand of a system increases when the optional execution factor of a task in the given set increases, as described in Theorem 2.

Theorem 2: *Given task $\tau_i \in \Gamma$, its two different optional execution factors α and α', and the task allocation scheme in Alg. 1, the inequality $E_{dem} > E_{dem}'$ holds if $\alpha > \alpha'$, where E_{dem} and E_{dem}' are the energy demand of the generated task schedule corresponding to α and α' of task τ_i, respectively.*

Since the optional execution factor of a task is positively related to system energy demand, we can determine the task optional execution factor for given energy demand by using a simple yet effective binary search-based approach. We assume that the optional execution factor of every task is the same such that the adaptive task allocation is conducted at a coarse granularity. We refine the optional execution factors of tasks such that each task has its own optional execution factor at the task scheduling stage, as shown in Section V.

Algorithm 2: Iterative task allocation based on task optional execution factors

Input: Task set Γ && MPSoC Θ && cost budget $Cost_{bdg,\mathcal{H}}$ && a sufficiently small positive number ϵ

Output: task to virtual core assignment $\{\Gamma_1, \Gamma_2, \cdots, \Gamma_{\mathcal{X}}\}$

1 derive available energy E_{sup} during $[t, t + \mathcal{H}]$ using Eq. (11);
2 $min = 0, max = 1, \alpha = (min + max)/2$;
3 calculate task execution cycles for all tasks using Eq. (26) based on α, call **Alg. 1** to allocate tasks, and calculate energy demand E_{dem} using Eq. (13);
4 **while** $|E_{sup} - E_{dem}| > \epsilon$ **do**
5   **if** $E_{dem} > E_{sup}$ **then**
6       $max = \alpha, \alpha = (min + max)/2$;
7   **else**
8       $min = \alpha, \alpha = (min + max)/2$;
9   update l_i for all tasks using Eq. (26) based on α;
10   allocate tasks using **Alg. 1** based on l_i;
11   update E_{dem} of the task schedule using Eq. (13);

To handle the uncertainty in energy availability, we adapt system energy demand to fluctuating energy supply by using the task optional execution factor as a coarse-grain control knob to adjust the system energy consumption. The proposed adaptive task allocation algorithm is described in Alg. 2. It takes as input task set Γ, processor set Θ and the given cost budget $Cost_{bdg,\mathcal{H}}$, and outputs $\mathcal{X}$ subsets $\{\Gamma_1, \Gamma_2, \cdots, \Gamma_{\mathcal{X}}\}$. It first derives the budget-constrained energy E_{sup} that is available to support the system operation during the scheduling horizon $\mathcal{H}$ based on Eq. (11). It then utilizes a binary search-based approach to derive a common task optional execution factor for all tasks in set Γ. It initializes the common task optional execution factor α to 0.5. It calculates task execution cycles using Eq. (26), calls Alg. 2 to allocate tasks, and computes energy demand E_{dem} using Eq. (13). Once energy demand E_{dem} of the task allocation is derived, the algorithms iteratively adapt the energy demand to the fluctuating supply of energy E_{sup} according to $|E_{sup} - E_{dem}| > \epsilon$, where ϵ is a sufficiently small positive number. Alg. 2 stops when energy demand of the generated schedule closely matches to the energy supply, and returns the allocation solution.

V. QoS-Driven Task Scheduling

This section maximizes the system QoS by determining tasks whose optional parts should be executed. Once these tasks are picked out, a runtime scheme needs to be designed to adapt task schedule to the fluctuating energy availability.

A. Static Task Selection for QoS optimization

As given in Eq. (18), the QoS of a system is defined as the sum of executed CPU cycles of optional parts of all tasks in the given set. To maximize the QoS of a system under the energy constraint, we propose to develop a task selection scheme that chooses certain tasks and executes the optional parts of these tasks. Let E_ℓ denote the energy consumed by tasks in subset Γ_ℓ allocated to the virtual core $\bar{\theta}_\ell$, then E_ℓ is given by

$$E_\ell = \sum_{\tau_i \in \Gamma_\ell} P_{cons}(\tau_i, \bar{\theta}_\ell) \cdot \frac{o_i + M_i}{f_\ell} = \sum_{\tau_i \in \Gamma_\ell}$$
$$\frac{P_{cons}(\tau_i, \bar{\theta}_\ell)}{f_\ell} \cdot o_i + \sum_{\tau_i \in \Gamma_\ell} \frac{P_{cons}(\tau_i, \bar{\theta}_\ell)}{f_\ell} \cdot M_i, \quad (27)$$

where $P_{cons}(\tau_i, \bar{\theta}_\ell)$ is the power consumption of task τ_i on core $\bar{\theta}_\ell$, f_ℓ is the operating frequency of core $\bar{\theta}_\ell$, o_i is the actual optional execution cycles of task τ_i, and M_i is the mandatory part of task τ_i. The energy demand or consumption of the system can be expressed as the sum of energy E_ℓ consumed by individual cores. It approximates closely to the energy supply of the system by a parameter $\epsilon > 0$ since tasks are assigned to individual cores using the energy adaptive algorithm (i.e., Alg. 2) presented in Section IV-B. Thus, we have

$$E_{sup} \approx E_{dem} = \sum_{\ell=1}^{\mathcal{X}} E_\ell = \underbrace{\sum_{\ell=1}^{\mathcal{X}} \sum_{\tau_i \in \Gamma_\ell} \frac{P_{cons}(\tau_i, \bar{\theta}_\ell)}{f_\ell} \cdot o_i}_{E_{optl}}$$
$$+ \underbrace{\sum_{\ell=1}^{\mathcal{X}} \sum_{\tau_i \in \Gamma_\ell} \frac{P_{cons}(\tau_i, \bar{\theta}_\ell)}{f_\ell} \cdot M_i}_{E_{mand}}, \quad (28)$$

where the first item E_{optl} is the energy consumed by executing optional parts while the second item E_{mand} is the energy consumed by executing mandatory parts of real-time tasks. When a scheduling decision is made at the start of the scheduling horizon $\mathcal{H}$, the estimated energy supply E_{sup} is a fixed value. In addition, the E_{mand} shows that the energy consumed by all allocated mandatory parts during the scheduling horizon $\mathcal{H}$ is invariable. As a result, the total energy consumption of the optional parts given by the E_{optl} is fixed.

The problem of selecting tasks for QoS optimization is hence refined into finding optimal optional cycles for all allocated tasks under the energy budget constraint. Let

$$EC_{\ell,i} = \frac{P_{cons}(\tau_i, \bar{\theta}_\ell)}{f_\ell} \quad (29)$$

be an energy metric that gives the energy consumption of task τ_i on core $\bar{\theta}_\ell$ per execution cycle. Since every task is only allowed to bind with one core, we use EC_i rather than $EC_{\ell,i}$ to denote the energy metric of τ_i on $\bar{\theta}_\ell$ after task assignment is determined. In addition, we sort all tasks allocated to cores in ascending order of EC_i for the sake of easy presentation. Thus, the energy consumed by optional parts of all tasks is

$$E_{optl} = EC_1 \cdot o_1 + EC_2 \cdot o_2 + \cdots + EC_N \cdot o_N, \quad (30)$$

and $EC_1 < EC_2 < \cdots < EC_N$ holds. In Eq. (30), the metric EC_i and task executed optional cycles $o_i = \alpha \cdot O_i$ are determined once tasks are bound to individual cores. However, the α is derived in the task allocation for maximizing energy efficiency. Now we aim to maximize system QoS in the task scheduling under the energy budget by re-determining the executed optional cycles o_i of tasks using an individual optional execution factor α_i.

Clearly, EC_i is determined once tasks are bound to individual cores. Moreover, the LHS of Eq. (30) is determined after task allocation. Thus, our goal of optimization for system QoS, which is given in Eq. (18) as the sum of executed optional cycles of all tasks, becomes maximizing the sum of o_i for $1 \leq i \leq N$ under the constraint of Eq. (30). This can be achieved by developing a task selection scheme that wisely chooses certain tasks and executes their optional cycles. Given

---

Algorithm 3: Determine task optional execution cycles

Input: Subsets of assigned tasks produced by Alg. 2
Output: Tasks with optional execution cycles updated

1 **for** $\ell = 1$ to $\mathcal{X}$ **do**
2   **for** $i = 1$ to $size(\Gamma_\ell)$ **do**
3   calculate the energy metric EC_i;
4   sort the tasks allocated to core $\bar{\theta}_\ell$ in ascending order of EC_i using heapsort;
5 $E_{optl} = E_{sup} - E_{mand}$;
6 **for** $\ell = 1$ to $\mathcal{X}$ **do**
7   **for** $i = 1$ to $size(\Gamma_\ell)$ **do**
8   **if** $E_{optl} > 0$ **then**
9   **if** $EC_i \times O_i \leq E_{optl}$ **then**
10                  $o_i = O_i, \alpha_i = 1$;
11   **else**
12                  $o_i = \frac{E_{optl}}{EC_i}, \alpha_i = \frac{E_{optl}}{EC_i \times O_i}$;
13              $E_{optl} -= EC_i \times o_i$;
14   **else**
15              $o_i = 0, \alpha_i = 0$;
16          $i \leftarrow i + 1$;

Algorithm 4: Adapt task execution to fluctuating energy at runtime

Input: Optional execution cycles $\{o_1, o_2, \cdots, o_N\}$ determined by Alg. 3 && subset Γ_ℓ of tasks on core $\bar{\theta}_\ell$

1 sort tasks in Γ_ℓ in ascending order of metric EC using heapsort;
2 $h_\ell = \frac{\sum_{i=1}^{size(\Gamma_\ell)} EC_i \cdot (M_i + o_i)}{f^{-1}(Cost_{bdg})}$;
3 $i \leftarrow 1$;
4 **while** $i \leq size(\Gamma_\ell)$ **do**
5      $l_i = (o_i + M_i)$;
6   **while** $EC_i \cdot l_i > h_\ell \cdot f^{-1}(Cost_{bdg}) \cdot \frac{l_i}{f_\ell \mathcal{H}}$ && $l_i \geq M_i$ **do**
7          $l_i = \sigma \cdot l_i$;   /* scale the length l_i to reduce energy demand of task τ_i */
8   **while** $h_\ell \cdot f^{-1}(Cost_{bdg}) \cdot \frac{l_i}{f_\ell \mathcal{H}} - EC_i \cdot l_i > \epsilon$ && $l_i \geq M_i$ **do**
9          $l_i = (1 + \sigma) \cdot l_i$;   /* scale the length l_i to increase energy demand of task τ_i */
10   **if** $h_\ell \cdot f^{-1}(Cost_{bdg}) \cdot \frac{l_i}{f_\ell \mathcal{H}} - EC_i \cdot l_i \leq \epsilon$ && $l_i \geq M_i$ **then**
11   execute task τ_i with execution cycles of l_i;
12   **else**
13   drop task τ_i and $l_i = 0$;   /* drop τ_i when the energy supply is insufficient */
14   update battery storage;
15      $i \leftarrow i + 1$;

---

these, we present and prove a theorem below, which shows the QoS is maximized if the tasks having smaller energy metric are selected to execute their optional parts.

Theorem 3: *Given the energy budget E_{optl} of optional parts of the task set Γ, the QoS of the system, as defined in Eq. (18), is maximized if the tasks having smaller energy metric are selected to execute their optional parts.*

Based on the theorem, we conclude that executing optional parts of tasks with smaller energy metric can improve the system QoS. A heuristic is thus designed to determine the execution cycles of optional parts of all tasks, which is shown in Alg. 3. The algorithm describes the procedure to re-determine optional execution cycles of tasks in subsets for maximizing the system QoS. The algorithm first calculates EC, the energy metric that gives the energy consumption per execution cycle of a task and sorts the tasks allocated to individual cores in ascending order of the metric, then calculates the energy budget E_{optl} for optional parts of all tasks. It derives optional execution cycles for each task based on the analysis that executing optional parts of tasks with smaller energy metric can improve the system QoS. Given energy metric EC_i and maximum optional cycles O_i of τ_i, the energy consumed by the task can be expressed as $EC_i \times O_i$. If energy budget E_{optl} for optional parts is large enough, maximum optional cycles of τ_i can be executed by $o_i = O_i$; otherwise, only $o_i = E_{optl}/EC_i$ optional cycles can be executed. No optional execution cycles will be executed if the energy budget is exhausted. After τ_i is examined, the energy budget is updated accordingly and the process moves to the next task. The algorithm outputs all tasks once their optional execution cycles o_i are updated.

B. Dynamic Energy Adaptation for QoS Optimization

The proposed scheme generates a cost-constrained task schedule at design time by mapping tasks to processors and executing optional cycles of those tasks that can maximize the system QoS. However, due to the fact that the amount of renewable and grid energy available under a cost budget fluctuates, the static task schedule needs to be adjusted at runtime. Therefore, we provide a dynamic task scheduling algorithm that adapts task execution to fluctuating energy available at runtime. The dynamic algorithm is developed based on the QoS-aware static algorithm described in Alg. 4.

Consider core $\bar{\theta}_\ell$ with allocated tasks in subset Γ_ℓ. Energy available for core $\bar{\theta}_\ell$ to execute tasks includes renewable energy harvested from environments, and energy drained from battery and grid. Due to the intermittence of renewable energy and time varying price of grid power, the amount of energy available to support the operation of the system under a given cost budget fluctuates. Assume that the actual energy available will be allocated to individual cores in proportion to the optimal energy value derived using Alg. 4. The actual energy available is then weighted by a factor h_ℓ, which is defined as the ratio of the energy consumed by $size(\Gamma_\ell)$ tasks on core $\bar{\theta}_\ell$ to the energy available for the system under the cost budget of $Cost_{bdg}$. Let $f^{-1}(Cost_{bdg})$ denote the energy available for the given cost budget $Cost_{bdg}$, then the factor h_ℓ is given by

$$h_\ell = \frac{\sum_{i=1}^{size(\Gamma_\ell)} EC_i \cdot (M_i + o_i)}{f^{-1}(Cost_{bdg})}. \tag{31}$$

The dynamic algorithm running on core $\bar{\theta}_\ell$ is given in Alg. 4. It takes as input the N task optional execution cycles determined by Alg. 4 and the subset Γ_ℓ of tasks allocated to core $\bar{\theta}_\ell$. To achieve a higher system QoS, the algorithm sorts tasks on the core in ascending order of energy metric EC so that it first executes tasks with smaller energy consumption per execution cycle. It calculates the factor h_ℓ for $\bar{\theta}_\ell$ using Eq. (31). For task τ_i in subset Γ_ℓ, the algorithm calculates the execution cycles l_i of the task. The energy demand of the task is expressed as $EC_i \cdot l_i$ and the energy available on $\bar{\theta}_\ell$ during time interval $\frac{l_i}{f_\ell}$ is expressed as $h_\ell \cdot f^{-1}(Cost_{bdg}) \cdot \frac{l_i}{f_\ell \mathcal{H}}$. The dynamic algorithm aims to adapt the static task schedule to the variation in energy supply. That is, it tries to reduce the runtime

mismatch between the energy supply and demand before a task is executed. On one hand, if energy demand of τ_i is greater than energy supply of the task, then execution cycles of the task and in turn the energy demand of the task are iteratively reduced by an empirical factor $\sigma \in [0, 1]$. The iteration stops when the energy demand is less than or equal to the energy supply. On the other hand, if the energy supply of τ_i is greater than the energy demand of the task by a small positive number $\epsilon > 0$, the execution cycles of the task and in turn the energy demand of the task are iteratively increased by an empirical factor $\sigma \in [0, 1]$. After execution cycles and energy demand of τ_i are updated in accordance with the fluctuation of energy supply, the task is executed if its execution cycles l_i is greater than its mandatory cycles M_i; otherwise, τ_i is dropped and l_i is reset to 0. The energy available for core $\bar{\theta}_\ell$ is updated and the procedure moves to the next task.

VI. Evaluation

A. Experimental Settings

We validate the proposed cost-constrained task allocation and scheduling scheme through two sets of simulation experiments. The first set of simulations is based on synthetic real-time tasks while the second set of simulations is based on real-life benchmarks. As introduced in Section IV and V, our scheme is composed of four algorithms. The first one is a deterministic task allocation (DTA, Alg. 1) proposed in [27] that minimizes the energy consumption by utilizing heterogeneities of both processors and tasks. The second one is an adaptive task allocation (ATA, Alg. 2) that uses a common optional execution factor to handle the uncertainty in energy availability and adopts DTA to allocate tasks. The third one is a static task selection (STS, Alg. 3) that selects certain tasks to complete their optional parts for maximizing QoS under the energy budget. The tasks with smaller EC are given higher priority to execute optional parts. The three algorithms are developed at design time. The fourth one is a dynamic task scheduling (DTS, Alg. 4) which adjusts the static schedule to adapt to the fluctuating harvested energy at runtime.

Accordingly, four comparative experiments are carried out in each set of simulations to validate our scheme from different perspectives. First, we compare the energy consumption of our DTA scheme [27] with that of the hybrid worst-fit genetic algorithm (HWGA) [28] to validate the efficiency of our ATA algorithm in saving energy, which is developed based on DTA. We also investigate the degree of match between energy demand and supply by using our ATA to validate the efficiency of our ATA in utilizing the energy supply, which specifically exploits an optimal common task optional execution factor to maximize the degree of match. Second, we compare the QoS of our ATA with that of DTA under a given cost constraint to validate the efficiency of our scheme in improving QoS at task allocation stage. Third, we compare the QoS of our STS with that of baseline method RAND, RVS, and benchmarking algorithm critical-task-first (CTF) [14] to validate the efficiency of our scheme in improving QoS at task scheduling stage. Finally, to validate the effectiveness of our runtime scheme DTS, we compare the QoS of DTS with that


Fig. 4: Normalized energy consumption of 30 task sets using our DTA and benchmarking algorithm HWGA [28].
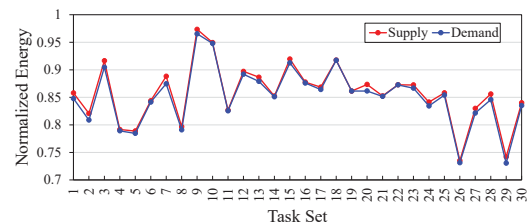

Fig. 5: Normalized energy supply and demand of our proposed ATA algorithm for executing 30 task sets.

of the dynamic algorithm, gradient curve shifting (GCS) [13]. These algorithms are described below.

- **DTA** [27] is an energy-efficient task allocation algorithm designed for systems of deterministic energy supply. It can minimize dynamic energy consumption by assigning the subset of tasks having a large power dissipation factor to the processor having a small power dissipation factor.
- **HWGA** [28] is a state-of-the-art approach that integrates a worst-fit based partitioning heuristic with the genetic algorithm to generate a task allocation that reduces the energy consumption while satisfying all system constraints.
- **RAND** is a method that randomly selects tasks whose optional parts are to be completed under the budget.
- **RVS** is a method that selects the task whose energy metric is larger to execute first, which is exactly the opposite of our proposed task selection algorithm STS.
- **CTF** [14] is a method that assigns QoS-critical jobs higher priorities such that their optional cycles can be completed first. The QoS-critical jobs are defined as tasks with larger maximum optional cycles.
- **GCS** [13] is a dynamic scheduling method that determines the best allocation of slack cycles generated at runtime for maximizing QoS under energy constraint.

B. Simulation for Synthetic Real-Time Tasks

We perform this simulation based on an MPSoC system with 8 heterogeneous processors ($M = 8$). Our processor model is built on 65nm technology and the parameters of processors can be found in [27], [29], [30]. Task activity factors μ are uniformly distributed in the interval $[0.4, 1]$, demonstrating the heterogeneous nature of tasks [16]. The size N of task set Γ is set to 200. The worst case execution cycles (WCEC) of tasks in the set Γ are assumed to be in the range of $[4 \times 10^7, 6 \times 10^8]$. Each task τ_i is instantiated by randomly picking two WCECs from the range, one is for its mandatory part M_i and the other is for its maximum optional part O_i. The common deadline d of tasks is assumed to be

$1.5 \sum_{i=1}^{N} M_i / f_{max}$, where f_{max} is the maximum frequency. Multiple task sets are constructed to validate the proposed scheme in the simulation.

Solar source is selected as the renewable generation, and the trace of harvesting power $P_{harv}(t)$ is generated according to the equation $P_{harv}(t) = \left| \psi \cdot \Psi(t) \cdot \cos(\frac{t}{70\pi}) \cdot \cos(\frac{t}{100\pi}) \right|$ [8], [31], where ψ is a coefficient and $\Psi(t)$ is a normally distributed random variable with variance 1 and mean 0. The ψ is used to generate two different $P_{harv}(t)$, the first of which is taken as predicted harvesting power and the second of which is used as actual harvesting power. The real-time pricing model offered by Ameren Illinois Corporation [32] is adopted as the pricing model of grid electricity. A Li-lion battery bank equipped with 4Ah normal capacity and 15V terminal voltage [19] is served as the energy storage module. The whole scheduling horizon $\mathcal{H}$ is set to 24 hours and divided into 24 time slots with equal and constant length of 1 hour.

1) Evaluate the Effectiveness of Our ATA Scheme in Energy Management: Our ATA scheme is developed based on the energy-efficient DTA scheme [27] that minimizes the system energy consumption by fully exploiting the heterogeneities of both processors and tasks, and an optimal task common optional execution factor that can maximize the degree of match between system energy demand and supply. To evaluate the effectiveness of our ATA in energy management, we first compare our DTA with benchmarking algorithm HWGA [28] in terms of energy consumption, then investigate the degree of match between energy demand and supply using our ATA.

Fig. 4 shows the normalized energy consumption of 30 task sets using our DTA scheme [27] and benchmarking algorithm HWGA [28]. It has been demonstrated in the figure that our DTA consumes less energy as compared to HWGA. To be specific, the energy consumption of DTA is 19.7% lower than that of HWGA on average. Furthermore, the energy savings achieved by DTA over HWGA can be up to 29.0%. For example, the normalized energy consumed by executing task set 1 using HWGA and DTA are 0.854 and 0.607, respectively. In addition to the efficiency of reducing energy consumption achieved by DTA, our scheme ATA can make the most of energy supply by introducing task optional execution factor. Fig. 5 plots the energy supplied by the harvesting system and the energy demanded by our ATA for executing 30 task sets. It can be easily seen from the figure that our scheme achieves a close match between energy supply and demand.

2) Evaluate the Effectiveness of Our ATA and STS Scheme for QoS Improvement: Two comparative experiments are carried out to validate the effectiveness of our proposed static QoS-aware task selection scheme under varying cost budgets. We utilize a normalized cost budget in the comparative study, the maximum of which is set to 1. That is, the normalized cost budget varies in the range of (0, 1]. The first comparative study is conducted from the perspective of task allocation, where two allocation schemes, the DTA without considering energy uncertainty and ATA considering energy uncertainty, are compared with respected to QoS. Fig. 6 presents the normalized QoS obtained when executing the optional cycles of tasks in 10 sets over a scheduling horizon of 24 hours using the proposed DTA and ATA scheme under varying cost


Fig. 6: Normalized QoS of 10 synthetic task sets over 24 hours using the proposed DTA and ATA scheme under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.


Fig. 7: Normalized QoS of 10 synthetic task sets over 24 hours using the proposed STS and baseline method RVS under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.

budgets. As compared to DTA, ATA achieves higher QoS of up to 55.1%, especially when the cost budget is low. For example, the average QoS of ATA are 51.8% and 42.1% higher than that of DTA for $Cost_{bgt} = 0.6$ and 0.7, respectively.

The second comparative experiments investigate the impact of STS on system QoS. To be specific, we compare the QoS of our STS with that of baseline task selection methods RAND, RVS, and benchmarking algorithm CTF [14] to validate the efficiency of our scheme in improving QoS at the task scheduling stage. Fig. 7 shows the normalized QoS obtained when executing optional cycles of tasks in 10 sets of tasks using the proposed STS and the baseline method RVS under the cost budget of $Cost_{bgt} = 0.6$ and 0.7. As compared to RVS, the average QoS of our STS scheme are 97.6% and 72.4% higher in the case of $Cost_{bgt} = 0.6$ and 0.7, respectively. Fig. 8 compares the proposed STS and baseline method RAND in terms of the QoS obtained when executing


Fig. 8: Normalized QoS of 10 synthetic task sets over 24 hours using the proposed STS and baseline method RAND under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.


Fig. 9: Normalized QoS of 10 synthetic task sets over 24 hours using our STS and benchmarking method CTF [14] under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.

optional cycles of tasks under the cost budget of $Cost_{bgt} = 0.6$ and 0.7, respectively. It can be seen from the figure that the average QoS of STS are 75.2% and 51.8% higher than that of RAND for $Cost_{bgt} = 0.6$ and 0.7, respectively. A comparative study of the proposed STS and benchmarking method CTF is demonstrated in Fig. 9. The figure indicates that the average QoS of our STS scheme are 61.2% and 39.7% higher than that of the CTF method for $Cost_{bgt} = 0.6$ and 0.7, respectively.

From the above analyses, we can draw the followings.

- From the perspective of task allocation, the ATA that considers the uncertainty in energy availability outperforms the DTA that doesnot in terms of system QoS.
- As to selecting tasks on individual processors to be executed, the STS outperforms methods RVS, RAND, and CTF in terms of system QoS. This is because our STS utilizes heterogeneities of both tasks and processors for QoS improvement under given cost constraints.
- Our ATA and STS scheme outperform the benchmarking methods regarding to system QoS improvement, especially when the given cost budget is low. This is because when the given cost budget and the corresponding available energy is low, our ATA and STS scheme can efficiently utilize the intermittent available energy so that a larger portion of the optional part of a task is finished.

3) Evaluate the Effectiveness of Our DTS Scheme for QoS Improvement: The effectiveness of our dynamic task scheduling scheme, referred to as DTS, is also verified in the simulation. We compare the proposed dynamic task scheduling scheme DTS with benchmarking algorithm GCS [13] in terms of runtime QoS. Fig. 10 gives the normalized QoS achieved when executing optional cycles of tasks in 10 sets under the cost budget of $Cost_{bgt} = 0.6$ and 0.7, respectively.

We can observe from the figure that the average QoS of our DTS scheme are 17.7% and 15.1% higher than that of the GCS [13] algorithm for $Cost_{bgt} = 0.6$ and 0.7, respectively. This is primarily due to the fact that our proposed DTS utilizes heterogeneities of both processors and tasks to enhance system energy efficiency in a way so that the energy consumed per execution cycle is minimized. Owing to the enhanced energy efficiency achieved by judiciously assigning tasks to processors, DTS could not only guarantee the completion of task mandatory parts, but also execute more task optional parts. In addition to the enhanced energy efficiency that leads to a higher QoS, the DTS could directly improve QoS by intelligently selecting tasks on individual processors to execute their optional parts. On the contrary, the GCS [13] improves system QoS by only using slack claimed in the runtime, which renders relatively small headroom for QoS improvement.

C. Simulation for Real-Life Benchmarks

In this simulation, a heterogeneous MPSoC system [24] that consists of an AMD Athlon processor with three supply voltage/frequency levels and an TI DSP processor with two supply voltage/frequency levels is adopted. The three supply voltage and frequency pairs of AMD Athlon are $(0.89V, 1.8GHz)$, $(1.12V, 2.4GHz)$, $(1.34V, 3GHz)$, and the two supply voltage and frequency pairs of TI DSP are $(0.98V, 2.0GHz)$,


Fig. 10: Normalized QoS of 10 synthetic task sets over 24 hours using our DTS and benchmarking method GCS [13] under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.
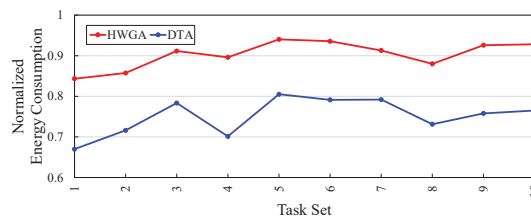

Fig. 11: Normalized energy consumption of 10 real-life task sets using our DTA and benchmarking algorithm HWGA [28].

$(1.42V, 3.0GHz)$ [24], respectively. The settings of solar source, Li-lion battery bank, and real-time pricing are the same as that of simulation for synthetic real-time tasks. The tool MEGA [33] that incorporates approximate computation is utilized in this simulation to generate real-life benchmarks.

1) Evaluate the Effectiveness of Our ATA Scheme in Energy Management: Fig. 11 shows the normalized energy consumption of 10 real-life task sets using our DTA scheme [27] and benchmarking algorithm HWGA [28]. It has been demonstrated in the figure that our DTA consumes less energy as compared to HWGA. To be specific, the energy consumption of DTA is 16.9% lower than that of HWGA on average. Furthermore, the energy savings achieved by DTA over HWGA can be up to 21.7%. For example, the normalized energy consumed by executing the real-life task set 4 using HWGA and DTA are 0.896 and 0.701, respectively. Fig. 12 plots the energy supplied by the harvesting system and the energy demanded by our ATA scheme for executing 10 real-life task sets. It can be easily seen from the figure that our scheme achieves a close match between energy supply and demand.

2) Evaluate the Effectiveness of Our ATA and STS Scheme for QoS Improvement: Fig. 13 presents the normalized QoS obtained when executing the optional cycles of real-life tasks in 10 sets over a scheduling horizon of 24 hours using the proposed DTA and ATA scheme under varying cost budgets. Compared to DTA, ATA achieves a higher QoS of up to 68.8%, especially when the cost budget is low. For example,


Fig. 12: Normalized energy supply and demand of our proposed ATA algorithm for executing 10 real-life task sets.

Fig. 13: Normalized QoS of 10 real-life task sets over 24 hours using our DTA and ATA scheme under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.


Fig. 14: Normalized QoS of 10 real-life task sets over 24 hours using our STS and baseline method RVS under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.

the average QoS of ATA are 59.6% and 48.3% higher than that of DTA for $Cost_{bgt} = 0.6$ and 0.7, respectively. Fig. 14 shows the normalized QoS obtained when executing optional cycles of real-life tasks in 10 sets using the proposed STS algorithm and baseline method RVS under the cost budget of $Cost_{bgt} = 0.6$ and 0.7. As compared to RVS, the average QoS of STS are 108.6% and 63.8% higher for $Cost_{bgt} = 0.6$ and 0.7, respectively. Fig. 15 compares the proposed STS algorithm and baseline method RAND in terms of the QoS obtained when executing optional cycles of real-life tasks in 10 sets under the cost budget of $Cost_{bgt} = 0.6$ and 0.7, respectively. It can be seen from the figure that the average QoS of STS are 78.6% and 40.4% higher than that of RAND for $Cost_{bgt} = 0.6$ and 0.7, respectively. A comparative study of the proposed STS algorithm and benchmarking method CTF is demonstrated in Fig. 16. The figure indicates that the average QoS of STS are 69.9% and 26.7% higher than that of CTF for $Cost_{bgt} = 0.6$ and 0.7, respectively.


Fig. 15: Normalized QoS of 10 real-life task sets over 24 hours using our STS and baseline method RAND under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.


Fig. 16: Normalized QoS of 10 real-life task sets over 24 hours using our STS and baseline method CTF [14] under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.


Fig. 17: Normalized QoS of 10 real-life task sets over 24 hours using our DTS and benchmarking method GCS [13] under varying cost budgets: (a) $Cost_{bgt} = 0.6$, (b) $Cost_{bgt} = 0.7$.

3) Evaluate the Effectiveness of Our DTS Scheme for QoS Improvement: Fig. 17 gives the normalized QoS achieved when executing optional cycles of real-life tasks in 10 sets under the cost budget of $Cost_{bgt} = 0.6$ and 0.7, respectively. We can observe from the figure that the average QoS of our DTS scheme are 20.8% and 13.5% higher than that of the GCS [13] algorithm for $Cost_{bgt} = 0.6$ and 0.7, respectively.

VII. CONCLUSION

We propose to utilize the characteristic of approximate computation to intelligently handle uncertainties in energy availability of an MPSoC real-time system that is powered by a hybrid of energy harvested from environments and drained from power grid. For systems of such power supply, we design a static energy-adaptive task allocation scheme and an QoS-driven task scheduling scheme that not only reduce the energy consumption but also improve the system QoS. We also design a dynamic task scheduling algorithm that adapts task execution to fluctuating energy available at runtime. We conduct extensive simulations to validate the effectiveness of our schemes. Results show that our algorithms can reduce energy consumption by up to 29% and improve system QoS by up to 108% as compared to benchmarking schemes.

VIII. ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their helpful suggestions. The proof of theorems are omitted due to page limit.

REFERENCES

[1] M. Elhebeary, M. Ibrahim, M. Aboudina, and A. Mohieldin, "Dual-source self-start high-efficiency micro-scale smart energy harvesting system for IoT applications," *IEEE TIE*, 2017.
[2] J. Zhou, J. Yan, T. Wei, M. Chen, and X. Hu, "Energy-adaptive scheduling of imprecise computation tasks for QoS optimization in real-time MPSoC systems," *DATE*, pp. 1402-1407, 2017.
[3] K. Huang, B. Hu, L. Chen, A. Knoll, and Z. Wang, "ADAS on COTS with OpenCL: a case study with lane detection," *IEEE TC*, 2017.
[4] Peek Traffic Corporation, Video Detection Products, available at http://peektraffic.com/products_video_detection.php.
[5] J. Liu, K. Lin, W. Shih, A. Yu, J. Chung, and W. Zhao, "Algorithms for scheduling imprecise computations," *Springer US*, 1991.
[6] M. Bambagini, M. Marinoni, H. Aydin, and G. Buttazzo, "Energy-aware scheduling for real-time systems: a survey," *ACM TECS*, vol. 15, no. 1, 2016.
[7] M. Severini, S. Squartini, and F. Piazza, "Energy-aware lazy scheduling algorithm for energy-harvesting sensor nodes," *NCA*, vol. 23, no. 7-8, pp. 1899-1908, 2013.
[8] C. Moser, D. Brunelli, L. Thiele, and L. Benini, "Real-time scheduling for energy harvesting sensor nodes," *RTS*, vol. 37, no. 3, pp. 233-260, 2007.
[9] M. Chetto and A. Queudet, "Clairvoyance and online scheduling in real-time energy harvesting systems," *RTS*, vol. 50, no. 2, pp. 179-184, 2014.
[10] Y. Abdeddaim, Y. Chandarli, and D. Masson, "The optimality of PFPasap algorithm for fixed-priority energy-harvesting real-time systems," *ECRTS*, pp. 47-56, 2013.
[11] G. Stavrinides and H. Karatza, "Scheduling real-time DAGs in heterogeneous clusters by combining imprecise computations and bin packing techniques for the exploitation of schedule holes," *FGCS*, vol. 28, no. 7, pp. 977-988, 2012.

[12] L. Cortes, P. Eles, and Z. Peng, "Quasi-static assignment of voltages and optional cycles in imprecise-computation systems with energy considerations," *IEEE TVLSI*, vol. 14, no. 10, pp. 1117-1129, 2006.

[13] H. Yu, B. Veeravalli, and Y. Ha, "Dynamic scheduling of imprecise-computation tasks in maximizing QoS under energy constraints for embedded systems," *ASPDAC*, pp. 452-455, 2008.

[14] H. Kooti, N. Dang, D. Mishra, and E. Bozorgzadeh, "Energy budget management for energy harvesting embedded systems," *RTCSA*, pp. 320-329, 2012.

[15] Y. Liu, R. Dick, L. Shang, and H. Yang, "Thermal vs energy optimization for DVFS-enabled processors in embedded systems," *ISQED*, pp. 204-209, 2007.

[16] H. Huang, V. Chaturvedi, G. Quan, J. Fan, and M. Qiu, "Throughput maximization for periodic real-time systems under the maximal temperature constraint," *ACM TECS*, vol. 13, no. 2s, 2014.

[17] M. Amirijoo, J. Hansson, and S. Son, "Specification and management of QoS in real-time databases supporting imprecise computations," *IEEE TC*, vol. 55, no. 3, pp. 304-319, 2006.

[18] A. Millner, "Modeling Lithium Ion battery degradation in electric vehicles," *CITRES*, pp. 349-356, 2010.

[19] Y. Wang, X. Lin, Q. Xie, N. Chang, and M. Pedram, "Minimizing state-of-health degradation in hybrid electrical energy storage systems with arbitrary source and load profiles," *DATE*, 2014.

[20] T. Cui, S. Chen, Y. Wang, S. Nazarian, Q. Zhu, and M. Pedram, "Optimal control of PEVs for energy cost minimization and frequency regulation in the smart grid accounting for battery state-of-health degradation," *DAC*, 2015.

[21] C. Perera, D. Talagala, C. Liu, and J. Estrella, "Energy-efficient location and activity-aware on-demand mobile distributed sensing platform for sensing as a service in IoT clouds," *IEEE TCSS*, vol. 2, no. 4, pp. 171-181, 2015.

[22] J. Kwak, Y. Kim, J. Lee, and S. Chong, "DREAM: dynamic resource and task allocation for energy minimization in mobile cloud systems," *IEEE JSAC*, vol. 33, no. 12, pp. 2510-2523, 2015.

[23] S. Zhuravlev, J. Saez, S. Blagodurov, A. Fedorova, and M. Prieto, "Survey of energy-cognizant scheduling techniques," *IEEE TPDS*, vol. 24, no. 7, pp. 1447-1464, 2013.

[24] K. Li, X. Tang, and K. Li, "Energy-efficient stochastic task scheduling on heterogeneous computing systems," *IEEE TPDS*, vol. 25, no. 11, pp. 2867-2876, 2014.

[25] W. Trappe, R. Howard, and R. Moore, "Low-energy security: limits and opportunities in the Internet of Things," *IEEE Security & Privacy Magazine*, vol. 13, no. 1, pp. 14-21, 2015.

[26] P. Schaumont, "Security in the Internet of Things: A challenge of scale," *DATE*, pp. 674-679, 2017.

[27] J. Zhou, T. Wei, M. Chen, J. Yan, X. Sharon Hu, and Y. Ma, "Thermal-aware task scheduling for energy minimization in heterogeneous real-time MPSoC systems," *IEEE TCAD*, vol. 35, no. 8, pp. 1269-1282, 2016.

[28] S. Saha, Y. Lu, and J. Deogun, "Thermal-constrained energy-aware partitioning for heterogeneous multi-core multiprocessor real-time systems," *RTCSA*, pp. 41-50, 2012.

[29] W. Liao, L. He, and K. Lepak, "Temperature and supply voltage aware performance and power modeling at microarchitecture level," *IEEE TCAD*, vol. 24, no. 7, pp. 1042-1053, 2005.

[30] G. Quan and V. Chaturvedi, "Feasibility analysis for temperature constraint hard real-time periodic tasks," *IEEE TII*, vol. 6, no. 3, pp. 329-339, 2010.

[31] J. Chen, T. Wei, and J. Liang, "State-aware dynamic frequency selection scheme for energy-harvesting real-time systems," *IEEE TVLSI*, vol. 22, no. 8, pp. 1679-1692, 2013.

[32] Day Ahead Pricing, Ameren Illinois, avialable at https://www2.ameren.com/RetailEnergy/RealTimePrices.

[33] Y. Tirat-Gefen, D. Silva, and A. Parker, "Incorporating imprecise computation into system-level design of application-specific heterogeneous multiprocessors," *DAC*, pp. 58-63, 1997.

Kun Cao is currently pursuing the Ph.D. degree with the Department of Computer Science and Technology, East China Normal University, Shanghai, China. His current research interests are in the areas of high performance computing, multiprocessor systems-on-chip and cyber physical systems.

Peijin Cong received the B.S. degree from the Department of Computer Science and Technology, East China Normal University, Shanghai, China, in 2016. She is currently pursuing the master degree with the Department of Computer Science and Technology, East China Normal University, Shanghai, China. Her current research interest is in the area of power management in mobile devices.

Mingsong Chen (S'08-M'11) received the B.S. and M.E. degrees from the Department of Computer Science and Technology, Nanjing University, Nanjing, China, in 2003 and 2006, respectively, and the Ph.D. degree in Computer Engineering from the University of Florida, Gainesville, FL, USA, in 2010. He is currently a full Professor with the Department of Embedded Software and Systems, East China Normal University, Shanghai, China. His current research interests include design automation of cyber-physical systems, formal verification techniques, and mobile cloud computing.

Gongxuan Zhang (SM'12) received the B.S. degree in electronic computer from Tianjin University in 1983 and the M.S. and Ph.D. degrees in Computer Application from the Nanjing University of Science and Technology in 1991 and 2005, respectively. He was a Senior Visiting Scholar in Royal Melbourne Institute of Technology from 2001.9 to 2002.3 and University of Notre Dame from 2017.7 to 2017.10. Since 1991, he has been with the Nanjing University of Science and Technology, where he is currently a professor in the School of Computer Science and Engineering. His current research interests include multicore and parallel processing and distributed computing.

Tongquan Wei (M'11) received the Ph.D. degree in electrical engineering from Michigan Technological University, Houghton, MI, USA, in 2009. He is currently an Associate Professor with the Department of Computer Science and Technology, East China Normal University, Shanghai, China. His current research interests include real-time embedded systems, green and reliable computing, parallel and distributed systems, and cloud computing.

Xiaobo Sharon Hu (S'85-M'89-SM'02-F'16) received the B.S. degree from Tianjin University, Tianjin, China, the M.S. degree from the New York University Tandon School of Engineering, Brooklyn, NY, USA, and the Ph.D. degree from Purdue University, West Lafayette, IN, USA. She is currently a Professor with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN, USA. She has authored more than 250 papers in the related areas. Her current research interests include real-time embedded systems, low-power system design, and computing with emerging technologies.

Junlong Zhou (S'15-M'17) received the Ph.D. degree in Computer Science from East China Normal University, Shanghai, China, in 2017. He was a Visiting Scholar with the University of Notre Dame, Notre Dame, IN, USA, during 2014-2015. He is currently an Assistant Professor with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, China. His research interests include real-time embedded systems, cloud computing, and cyber physical systems.

Jianming Yan received the master's degree from the Department of Computer Science and Technology, East China Normal University, Shanghai, China, in 2016. He is currently a senior software engineer with Meituan.com Corporation, Beijing, China. His research interests include task allocation and scheduling techniques in heterogeneous real-time MPSoC systems.