Available online at www.sciencedirect.com
ScienceDirect
Journal of Power Sources 166 (2007) 284–296
JOURNAL OF POWER SOURCES
www.elsevier.com/locate/jpowsour

Coordinated discharge of a collection of batteries

Shivakumar Sastry *, Omer Gimdogmus, Tom T. Hartley, Robert J. Veillette

Department of Electrical and Computer Engineering, The University of Akron, Akron, OH 44325-3904, United States

Received 31 October 2006; received in revised form 22 December 2006; accepted 22 December 2006
Available online 16 January 2007

Abstract

Collections of batteries are used to supply energy to a variety of applications. By utilizing the energy in such a collection efficiently, we can improve the lifetime over which energy can be supplied to the application. We say that the discharge of a collection of batteries is coordinated when, at the end of discharge, the difference in the remaining capacity of individual batteries is small. This paper presents a decision-maker based on a goal-seeking formulation that coordinates the discharge of a collection of batteries. This formulation allows us to use a simple battery model and simple decision-making algorithms. We present results from MATLAB simulations that demonstrate the performance of the decision-maker when energy is drawn out of the collection in three different discharge scenarios. The new decision-maker consistently improves the discharge efficiency obtained using scheduling methods. Our results show that when the discharge is coordinated, the lifetime of the collection is extended.

Keywords: Goal-seeking; Decision-maker; Batteries; Coordinated discharge

1. Introduction

A battery, or electrochemical cell, is an important source of energy that is widely used in applications such as medical devices, hybrid-electric vehicles, communication systems and spacecraft. Because the internal electrochemical processes of a battery are complex and difficult to model, it is difficult to estimate the state of charge (SoC) of a battery and the remaining duration over which it can supply energy to a load [1]. Further, cells or batteries are often used in series or parallel connected collections to satisfy the demands of the applications; it is even more difficult to estimate the remaining duration over which such a collection could supply energy to the load.

We consider the problem of discharging a collection of batteries in an efficient manner when the collection includes a number of spare batteries. Each battery in the collection can be switched in or switched out under the control of a decision-maker. We say that the discharge of a collection of batteries is *coordinated*, when the amount of remaining charge in each of the batteries at the end of discharge is nearly equal. That is, given a collection, $b_1, b_2, \ldots, b_n$, of n batteries, if $\mathrm{rc}(b_i)$ is the remaining charge in battery b_i, the discharge is coordinated when $|\mathrm{rc}(b_i) - \mathrm{rc}(b_j)|$ is small for $1 \le i, j \le n$. When the discharge of a collection is coordinated, the lifetime of the collection [2] is also extended.

We use a goal-seeking formulation as a basis to design a new decision-maker that coordinates the discharge of a collection of batteries. This formulation is based on the goal-seeking paradigm proposed by Mesarovic and Takahara [3] which has been used extensively to study large-scale and complex systems [4–6]. To the best of our knowledge, this formulation has not been used as a basis for battery management. This formulation allows us to use a simple model of a battery [7] and simple decision-making algorithms. We consider the variability in the unique dynamic behavior of individual batteries as being caused by two uncertainties. First, to account for variations introduced by the manufacturing processes, we assume that the initial charge and capacity of a battery is uncertain. Second, to account for differences in the dynamic behavior, we assume that the parameters of the battery models are uncertain. We define a metric called *discharge efficiency* that captures the lifetime over which a collection supplies energy, the energy drawn out of the collection, and the amount of charge that remains in the collection at the end of discharge, i.e., when the collection is no longer able to satisfy the requirements of the load. By using three typical discharge scenarios in which loads draw energy from a collection of batteries, we present results that demonstrate the discharge effi-

* Corresponding author. Tel.: +1 3309727646; fax: +1 3309726487.
E-mail address: ssastry@uakron.edu (S. Sastry).

0378-7753/$ – see front matter
doi:10.1016/j.jpowsour.2006.12.083

ciency of the decision-maker in each scenario using MATLAB simulations.

Consider an example scenario in which an application uses batteries that can nominally provide 1.35 V and have an initial charge of 0.86 Ah. If the application requires a minimum of 14 V, at least 11 batteries must be connected in series to satisfy this demand. Let us say we choose to use 12 batteries in series for this application; if these batteries are discharged by drawing a constant current of about 250 mA, then the collection could supply energy to the load for about 190 min assuming that all the batteries discharge uniformly. If there are 12 spare batteries available, then all the 12 batteries could be replaced at the end of discharge. If, however, there are only eight spare batteries available, then the spares are not useful to the application.

By using a decision-maker that periodically selects 12 out of the 20 batteries during discharge, in a round-robin manner, the lifetime of the collection can be extended to about 320 min. The new decision-maker presented in this paper extends the lifetime to about 430 min.

Section 2 presents the background. Section 3 describes the simulation approach and presents baseline measurements that are used to evaluate the performance of the new decision-maker. Section 4 presents the goal-seeking paradigm and Section 5 presents a goal-seeking formulation of the coordination problem, the decision-maker based on this formulation and simulation results. Finally, Section 6 presents our conclusions.

2. Background

An accurate estimate of the SoC of a battery is critical for a decision-maker that coordinates discharge. SoC is a measure of the remaining charge in a battery; if $Q_{b,\max}$ represents the maximum stored charge or capacity and $Q_b(t)$ represents the charge in a battery b at time t, then

$$\mathrm{SoC}_b(t) = \frac{Q_b(t)}{Q_{b,\max}} \times 100\%.$$

Assuming that the initial charge of a battery is known accurately, we can integrate the current to estimate the SoC of the battery as

$$\mathrm{SoC}_b(t) = \frac{Q_b(0) - \int_0^t i_b(t)\,\mathrm{d}t}{Q_{b,\max}} \times 100\%.$$

This method is usually referred to as ampere–hour integration or counting. However, this method unrealistically assumes that the discharging process is 100% efficient and may yield inaccurate estimates of the SoC. It also assumes an accurate knowledge of the capacity $Q_{b,\max}$ and of the initial charge $Q_b(0)$. Despite its inaccuracies, this method can be combined with other methods to estimate the SoC.

The terminal voltage also gives some information about SoC for many battery chemistries, with Lithium Ion being a notable exception. Because terminal voltage is affected by load conditions, a more accurate estimate of the SoC can be obtained from an open-circuit voltage measurement [8]. It is necessary to allow the battery to rest for a substantial duration of time to allow the electrochemical processes to reach equilibrium in order to get an accurate open-circuit voltage measurement.

2.1. Estimating battery lifetime

The duration of time over which a battery could supply energy to a load is called the *lifetime* of the battery. When an application uses a collection of batteries or cells, the idea of lifetime can be extended to the collection of batteries in the obvious manner. The SoC of a battery is a critical factor that influences the remaining lifetime of a battery. Because of the inherent difficulties in estimating SoC, it is also difficult to estimate the remaining lifetime of a battery. Other factors that affect the lifetime of a battery are depth-of-discharge, rate of discharge, current, temperature, shock and vibrations. In addition, when one battery in a collection is supplying a disproportionate amount of energy to compensate for some other weak battery, the lifetimes of the good batteries are decreased.

Fig. 1 shows the voltage profile of a NiMH battery through its lifetime when it is being discharged at a constant current. This figure shows that the voltage exhibits a sharp “knee” at the end of discharge where the voltage drops quickly. By detecting the knee a decision-maker could determine when a battery is about to become completely discharged or needs to rest for recovery.

2.2. Charge equalization

When a collection of batteries is used in an application, the unique behavior of the individual batteries results in an unbalanced SoC across the collection. Several techniques have been described in the literature to equalize the charge that is accepted by each battery in a collection when rechargable batteries are being charged [9–13] where the objective is to maintain the SoC of each battery as close as possible to the SoC of the other batteries in the collection. Resistive shunts, flying capacitor and energy converters have been used in the literature to equalize charge during discharge.

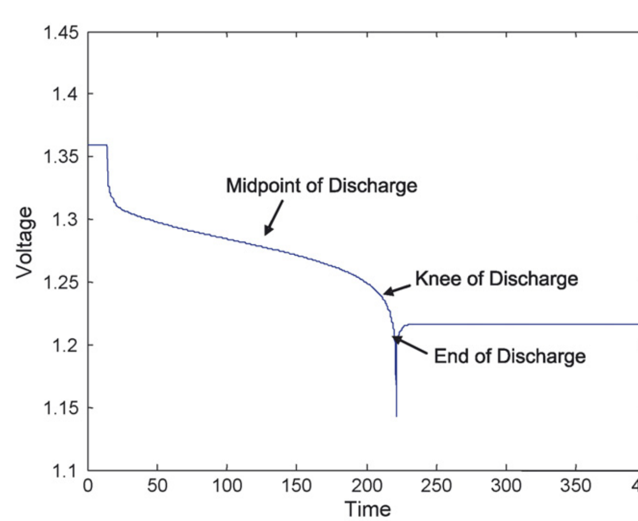

Fig. 1. Discharge profile of NiMH battery.

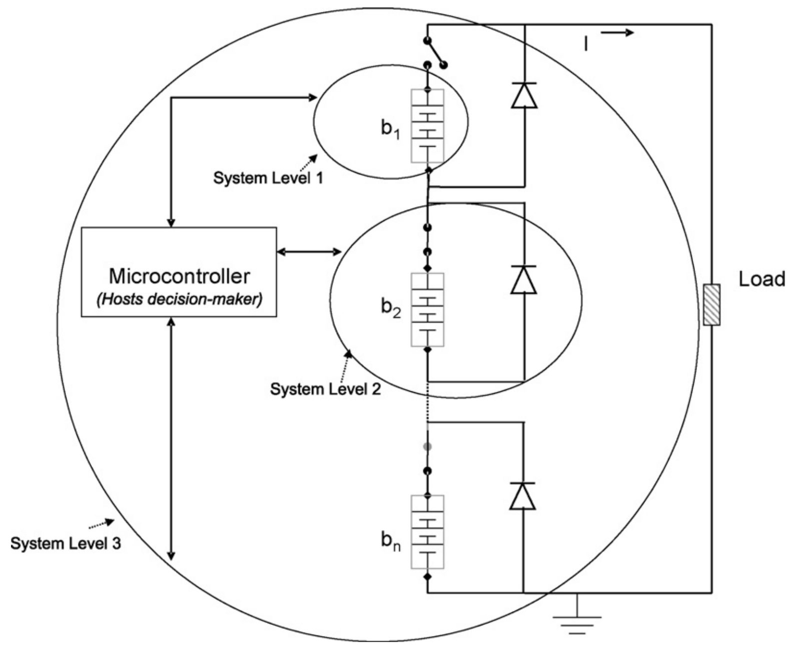

Fig. 2. System views of a collection of batteries.

Consider a collection of batteries connected in series in which there is a resistive shunt in parallel with every battery [9]. The amount of current through the shunt resistor is proportional to the voltage of its associated battery. Batteries that have higher SoC will tend to dissipate more power through the resistive shunts. While this approach is useful in a charging scenario, it has little use in a discharging scenario.

Instead of dissipating power through resistive shunts the use of a flying capacitor is a non-dissipative method in which the objective is to balance the voltage by moving energy from one battery to another by using active voltage or current converter elements [12]. A capacitor is connected to the collection of batteries and this capacitor can be charged selectively from one of the batteries in the collection that has a higher voltage. To select the source battery, a switch associated with the battery can be closed. Once the capacitor is charged, it is connected in parallel with a battery that has a lower voltage. While little power is dissipated in this method of equalization, it requires a large number of switches. The capacitor must be large and the switching currents must be high. Charge equalization is only effective to the extent that voltage is correlated with SoC.

2.3. Battery scheduling

When a collection includes spare batteries, charge equalization effects may be achieved by using battery scheduling techniques. Battery scheduling is a recent approach that has been used to extend the lifetime of a collection of batteries when the collection includes a number of spare batteries [2]. A decision-maker, which is hosted on a microcontroller [14,15], selects batteries that must rest and discharge using strategies discussed in Section 2.5. To discuss such strategies and our new decision-maker, we first present a systems view.

2.4. Systems view

As shown in Fig. 2, the collection of batteries and a microcontroller that hosts the decision-maker can be viewed to be at three levels. System Level 1 comprises a single battery.[1] At System Level 2, each battery is associated with a switch and a diode that is useful to either connect a battery for discharge or disconnect a battery for rest. At System Level 3, we consider the microcontroller that hosts a decision-maker and all the interfacing hardware and software necessary to select batteries and make measurements at all the three system levels.

A decision-maker observes the operational condition of the collection of discharging batteries at discrete sampling instants (or time ticks), $t_1, t_2, \ldots, t_f$. At some time-tick, t_f, the collection of batteries would no longer be able to satisfy the demands of the load. We refer to the interval of time between two successive time ticks as a *time-slice*, $T_s = |t_i - t_{i+1}|$ for all $1 \le i \le f-1$. For convenience, we assume that the first instant of time is t_1.

Given n batteries, in each time-slice the decision-maker must identify the m batteries that discharge and the $k = n - m$ batteries that rest. The requirements of the application and the terminal voltages of the batteries guide the choice of the number, m, of batteries that must be discharged in each time-slice. We assume a hypothetical application in which the load requires a minimum of 14 V. A collection of $n = 20$ NiMH batteries, each with a nominal capacity of 1 Ah and nominal voltage of 1.3 V are assumed to be available. When the terminal voltage of a battery is less than 1.25 V, that battery is defined to be completely discharged. The discharge process is terminated when the collection of batteries can no longer supply a minimum voltage of 4 V to the load.

[1] We refer to a battery pack, i.e., a fixed number of batteries packaged in a single module, as a battery for convenience.

2.5. Selection strategies for battery scheduling

We consider three selection strategies for scheduling batteries. These are queued selection, sliding window selection and random selection.

In the queued selection strategy, the decision-maker, initially, selects batteries $b_1, b_2, \ldots, b_m$ randomly for discharge. In some time-slice (t_i, t_{i+1}), if one of the selected batteries is completely discharged, it is replaced by one of the k spare batteries, by disconnecting the discharged battery and connecting the spare batteries in some pre-defined order. This strategy is simple and is commonly used in portable devices [2]. Because batteries display considerable variation in lifetime, it is not possible to predict when a spare battery will need to be introduced. It is necessary to estimate the state of discharge of each battery by monitoring the voltage or counting the ampere–hours. Even if we assume that the individual batteries will discharge uniformly, this approach will not extend the lifetime of the collection when $k<m$, and more than k batteries are completely discharged. Such a state of termination is not desirable because there would be at least k batteries that are not fully discharged.

In the sliding window selection strategy the decision-maker selects batteries $b_i, b_{i+1}, \ldots, b_{i+(m-1)}$ to discharge in time-slice $(t_i, t_{i+\mathrm{i}})$ while the remaining batteries rest. The discharge process terminates when the collection can no longer supply minimum voltage to the load. This strategy ensures that all the n batteries in the collection are used at about the same level. However, there is no means to compensate for variability in the individual batteries in this strategy.

In the random selection strategy, the decision-maker randomly selects m batteries, in each time-slice, to discharge while the remaining k batteries rest. This strategy typically utilizes the available batteries better than the queued selection strategy; however, the amount of charge remaining in the collection of n batteries at the end of discharge cannot be determined a priori because of the random selections.

3. Simulation approach

This section presents the battery model, methods used to estimate charge, representations that capture the unique behavior of individual batteries, discharge scenarios, and the discharge efficiency metric. We then present an application scenario that is the basis for the results reported in this paper. Using the battery model described here, we establish baseline measures to compare the performance of the decision-maker presented in Section 5.

3.1. Battery model

We selected a commonly used AA NiMH battery as a representative example of batteries used in applications. A model for such a battery was reported in [7]. This model uses a linear-in-the-parameters approximation to the electrode equation. The terminal voltage is expressed as

$$v(t) = k_1 - k_2 \times i(t) + k_3 \times Q_\mathrm{d}(t) + k_4 \times \log\left(\frac{Q(t)}{Q_\mathrm{max}}\right) - k_5 \times \log\left(\frac{Q_\mathrm{max} - Q(t)}{Q_\mathrm{max}}\right), \quad (1)$$

where t represents time (min), $v(t)$ the terminal voltage (V) at time t, $i(t)$ the current (A) at time t, Q_max the maximum capacity (Ah), $Q(t)$ the stored charge (Ah), $Q_\mathrm{d}(t)$ the diffusing charge (Ah), and $k_1, \ldots, k_5$ are the parameters that are obtained from a battery discharge curve.

Using a 1 Ah rechargeable NiMH cell with a degraded capacity of 0.86 Ah, the nominal values for these parameters are $k_1 = 1.33$ V, $k_2 = -0.12$ V A^{-1}, $k_3 = 20$ V Ah^{-1}, $k_4 = 0.015$ V and $k_5 = -0.013$ V.

Using the model in [7], the stored charge and diffusing charge are calculated as

$$\frac{\mathrm{d}Q(t)}{\mathrm{d}t} = \frac{1}{60} \times i(t) - 4.83 \times 10^{-6} \times Q(t),$$

and

$$\frac{\mathrm{d}Q_\mathrm{d}(t)}{\mathrm{d}t} = -0.35 \times Q_\mathrm{d}(t) - 0.001126 \times i(t).$$

3.2. Unique behavior of batteries

The initial SoC and the dynamic behavior of each battery are unique and difficult to characterize. The initial SoC of each battery varies significantly and critically affects the behavior of the battery. To capture this phenomenon in the simulation, we randomly vary the initial SoC in the battery models within 32% of 0.86 Ah. Thus, the initial SoC of each battery in the n batteries we consider are randomly chosen to be between 0.58 and 1.135 Ah.

To capture the differences in the dynamic behavior of each battery in the collection, we randomly vary the values of the parameters k_i within 10% of the nominal values presented in the preceding section.

3.3. Estimating charge

We use a method that combines open-circuit measurement and ampere–hour counting to estimate SoC. Using the NiMH battery model, we simulated the discharge of a single battery. We approximated the relationship between V_oc (in V) and SoC as

$$\mathrm{SoC} = (1.5265 - 3.4482 \times V_\mathrm{oc} + 2.5949 \times V_\mathrm{oc}^2 - 0.6506 \times V_\mathrm{oc}^3) \times 10^6\%. \quad (2)$$

Eq. (2) is used to obtain the starting point $\hat{Q}_\mathrm{max}$ for the ampere–hour counting method for estimating SoC. In addition, we used a table based approach to estimate the SoC where we approximated the SoC based on the terminal voltage of the battery.

The decision-maker acquires two open-circuit voltage measurements, $V_\text{oc-1}$ and $V_\text{oc-2}$, for each battery in the collection.

For each battery, $V_{oc\text{-}1}$ is measured before the discharge process begins and $V_{oc\text{-}2}$ is measured during the first time the battery rests. The SoC corresponding to these two measurements is obtained from Eq. (2) as SoC_1 and SoC_2. The ampere–hours drawn out of these batteries between the times at which these measurements, Q_{out}, was estimated using the ampere–hour counting method. These measurements allowed us to estimate the maximum capacity as

$$\hat{Q}_{max} = \frac{Q_{out}}{SoC_1 - SoC_2} \times 100\%. \tag{3}$$

3.4. Discharge scenarios and efficiency

Because of the diversity in battery applications, we use three discharge scenarios, namely constant current discharge (CCD), constant power discharge (CPD), and randomly varying current discharge (RVCD), to validate the decision-makers. In CCD, constant current is drawn out of the collection of batteries when supplying energy to the load. In CPD, power delivered to the load is maintained constant; this is necessary for example in an application that uses switching voltage regulators. When the voltage supplied by the batteries drops, current is increased to maintain constant power. In RVCD, we assumed that the current drawn out of the collection varies at random within 10% of the nominal value used for CPD—to represent applications where the load may change at random. For example, in a space mission, certain experiments can be turned on or off optionally. In all these discharge scenarios, the batteries supply energy to the load until they can no longer satisfy the requirements of the load.

Clearly, the remaining lifetime, t_{rem}, of a colletion of batteries, depends on the discharge scenario. In CCD, the remaining time is approximately

$$t_{rem} = \frac{Q_{rem}}{I}$$

where

$$Q_{rem} = Q_{max} - \int_0^t i(t)\,dt.$$

In CPD,

$$t_{rem} = \frac{\int_0^{Q_{rem}} v(Q)\,dQ}{P}$$

where P is the power maintained at the load. It is difficult to estimate the remaining time in RVCD.

Any decision-maker discharging a collection of batteries must consider the lifetime of the collection, the energy drawn out of the collection, and the utilized fraction of the total charge capacity of the collection. Because the importance of these efficiency metrics differs from one application to another, we define discharge efficiency (DE) as

$$DE = c_1 \times \frac{\text{Lifetime}}{\text{MPL}} + c_2 \times \frac{\text{EnergyOut}}{\text{MPE}} + c_3 \times \frac{\text{UsedCapacity}}{Q_{max}}, \tag{4}$$

where Lifetime is the lifetime over which the collection supplies energy to the load, EnergyOut the energy that is drawn out of the collection, UsedCapacity the amount of available capacity that is used, MPL the maximum possible lifetime that could be achieved, and MPE is the maximum possible energy that could be drawn out.

$c_i > 0$ and $\sum_i c_i = 1$. In this paper, we use the values $c_1 = c_2 = 0.25$ and $c_3 = 0.5$. MPL, MPE and Q_{max} are obtained from manufacturer specifications.

3.5. Baseline measurements

We now use the battery model to simulate a collection of batteries in a variety of scenarios. These results are established as a baseline to evaluate the performance of the decision-maker presented in Section 5.

Each simulation is executed with $T_s = 1$ min. For CCD, the discharging current is chosen to be 250 mA. In CPD, the fixed power maintained is 3.75 W. To maintain constant power, the current through the battery is increased as the battery discharges and the voltage decreases. For RVCD, we select the current to vary within 10% of a nominal 238 mA load current.

All the results presented in the figures are based on a single execution of the simulation and are intended to be illustrative. The results shown in the accompanying tables represent average values obtained from 10 executions of each simulation.

We first simulated a collection of 12 batteries in the three discharge scenarios. Fig. 3 shows the stored charge versus time of the collection of batteries in CCD. In this simulation the collection supplies the voltage required by the load for about 190 min. A total of 718.97 Wh energy was drawn out of the collection and the charge remaining in the collection is 2.8 Ah. Fig. 4 shows the effect of changing the discharge current in this baseline CCD scenario. When the discharge current is 185 mA, the collection supplies the voltage required by the load for 270 min; 779.83 Wh energy was drawn out and the remaining charge is 2.2 Ah. When the same 12 batteries were discharged at 500 mA, the collection supplied the voltage required by the application for

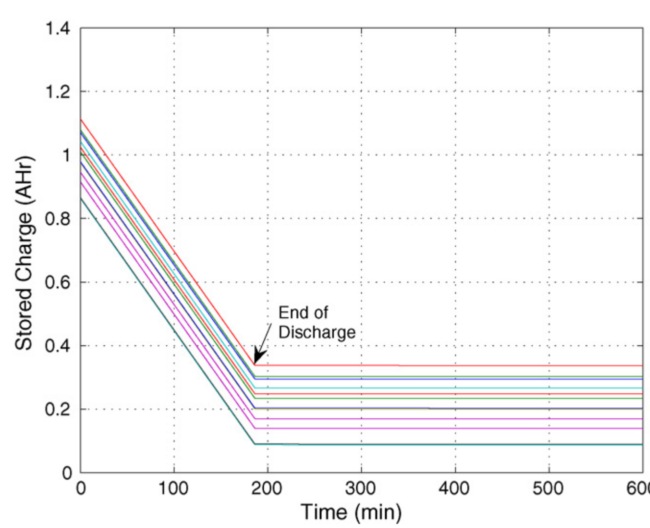

Fig. 3. Stored charge vs. time for baseline in CCD 250 mA.

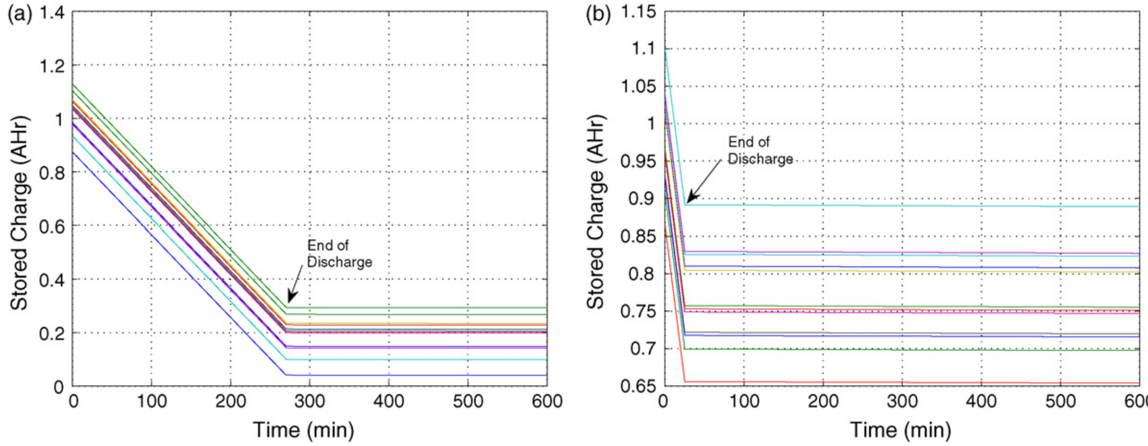

Fig. 4. Effects of discharge current on baseline CCD. (a) 185 mA and (b) 500 mA.

Table 1
Baseline measurements for 12 batteries

Scenario	Lifetime (min)	Energy out (Wh)	Remaining charge (Ah)
CCD	189.3	12.191	2.461
CPD	195.9	12.216	2.457
RVCD	214.8	12.621	2.316

25 min; the energy drawn out is 189.48 Wh and 9.24 Ah of charge remained in the collection. These variations are consistent with the expectation that discharge current impacts the performance of a collection of batteries.

Table 1 shows the lifetime, energy drawn out, and the remaining charge in the collection of 12 batteries in all the three discharge scenarios. As noted earlier, the values in this table are average values obtained after 10 executions of the simulation in each scenario when the discharge current is 250 mA.

Fig. 5 shows the stored charge versus time of the collection of batteries in CCD when discharged using the queued selection strategy. For this simulation, $T_s = 1$ min and the discharging process was terminated when the collection could no longer supply the minimum voltage required by the load. There are several batteries in the collection that have remaining charge, and yet, this charge is not sufficient to supply the minimum voltage required by the load. Despite this approach being the most commonly used one in portable electronic devices, note that the time is not extended significantly consistent with the observation by Benini et al. [2].

Fig. 6 shows the stored charge versus time of the collection of batteries in CCD when the batteries are selected at random in each time-slice. Notice that the lifetime is extended to about 311 min. Unlike the queued selection strategy, all the available batteries are used during discharge. The difference in the remaining charge of individual batteries is, however, large.

Fig. 7 shows the stored charge versus time the collection of batteries in CCD. Noticed that the lifetime is extended considerably to about 330 min. The difference in the remaining charge of individual batteries, however, remains wide. Fig. 8 shows that when the discharge current is reduced to 185 mA, the lifetime is increased to about 460 min. When the discharge current is increased to 500 mA, the lifetime reduces to about 160 min.

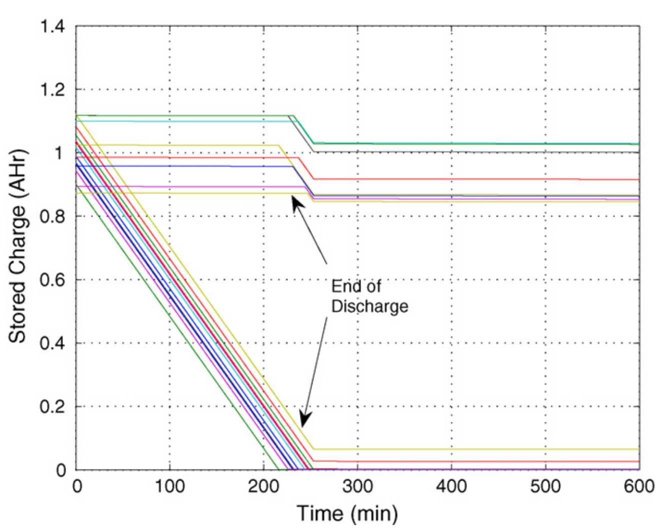

Fig. 5. Stored charge vs. time for queued selection in CCD 250 mA.

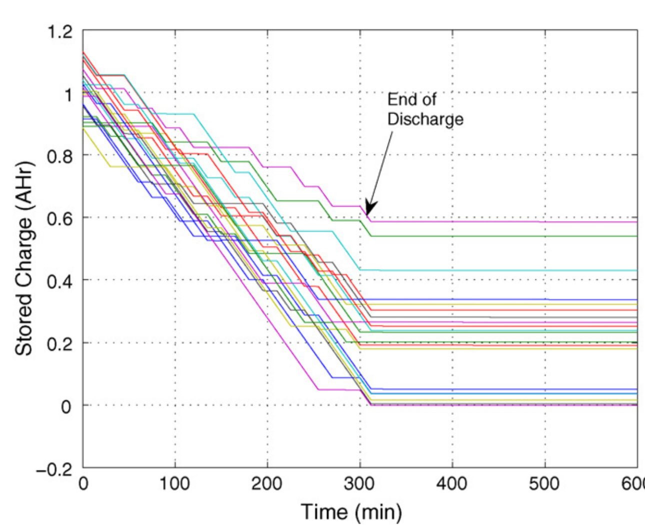

Fig. 6. Stored charge vs. time for random selection in CCD 250 mA.

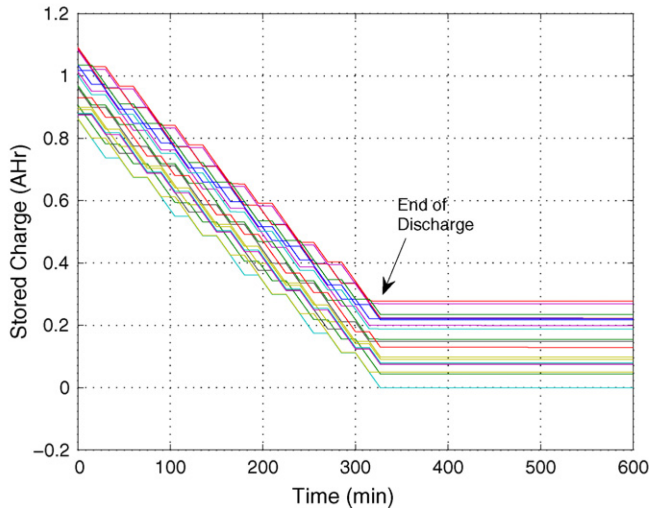

Fig. 7. Stored charge vs. time for sliding window selection in CCD 250 mA.

By comparing these results to the baseline values shown in Figs. 3 and 4, we note that the addition of spares improves the performance of the collection of batteries. Because the sliding window selection strategy performs better than queued selection and random selection, we present how the sliding window selection strategy performs when the discharge current changes in Table 2.

4. Goal-seeking paradigm

The goal-seeking paradigm [3] is an approach to modeling and describing systems that provides an alternative to the well-known state-transition paradigm. The state-transition paradigm is based on the assumption that the states of a system, Z, are precisely describable. The dynamics of the system are described by a *State-Transition* function

$$S_1: \quad Z \otimes X \rightarrow Z$$

where X is a set of control inputs or disturbances that affect the system in any given time-slice. The outputs produced by the

Table 2
Effects of discharge currents—sliding window selection

Scenario	Lifetime (min)	Energy (Wh)	Remaining capacity (Ah)
CCD			
185 mA	466.37	22.419	2.508
250 mA	345.83	22.26	2.493
500 mA	168.1	20.896	3.164
CPD			
185 mA	354.42	22.143	2.783
250 mA	353.46	22.085	2.515
500 mA	353.9	22.152	2.718
RVCD			
185 mA	528.56	22.911	2.295
250 mA	384.16	22.332	2.538
500 mA	184.16	20.749	3.374

system at any time-tick are determined by the mapping

$$S_2: \quad Z \rightarrow \Psi,$$

where Ψ is the set of system outputs.

In the goal-seeking paradigm, there is no attempt to describe the system states and hence the system model is necessarily simplified relative to what one may expect when using a state-transition paradigm. Instead, the decision-making process is formulated using the following sets and functions. There is a set of *Alternate Actions*, Π, from which the decision-maker can select actions. Anticipated system perturbations and disturbances are represented as a set of *Uncertainties*, Δ. If a given perturbation $\delta_i \in \Delta$ occurs, it would impact the success of a selected action. Consequences are outputs that are produced by the system; the set of *Consequences*, Ψ, includes all outcomes that may result from the implementation of some action. The decision-maker uses a function called *Reflection*, Ξ: $\Pi \otimes \Delta \rightarrow \Psi$, as its view of the environment. Suppose that the decision-maker selects an action $\pi_1 \in \Pi$; the decision-maker uses Ξ to estimate the consequence, $\psi_1 \in \Psi$ that π_1 would produce if a given perturbation occurs. An *Evaluation Set*, Λ, represents a *Performance Scale* that is used to compare the results of alternate actions according to an *Evaluation Mapping*, Ω: $\Psi \otimes \Pi \rightarrow \Lambda$. That is, if the decision-maker has the

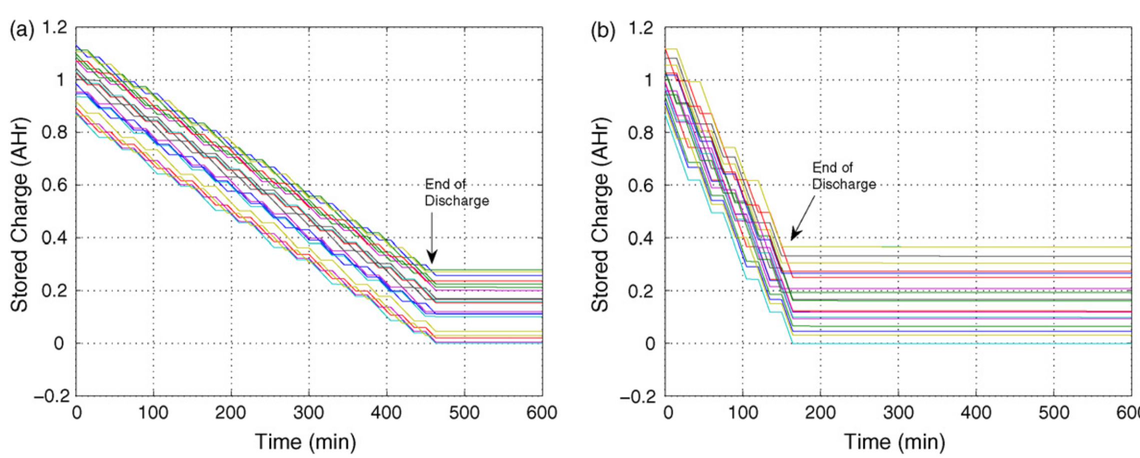

Fig. 8. Effects of discharge current on sliding window CCD. (a) 185 mA and (b) 500 mA.

option to select one of two actions $\pi_1, \pi_2 \in \Pi$, and these actions are expected to result in consequences $\psi_1, \psi_2 \in \Psi$, respectively, then the decision-maker uses values of Λ as the metric to determine whether one of the two actions is preferred over the other. Ω is also used to evaluate the actual measured output of the system. A *Tolerance Function*, Γ: $\Pi \otimes \Psi \rightarrow \Lambda$ provides a bound on how much the performance can vary before being considered as unsatisfactory. Using these artifacts and transformations, the task of the decision-maker may be stated as

> Continue to select an action $\pi \in \Pi$ as long as the outcome is within tolerance limits, i.e., $\Omega(\pi, \psi) > \Gamma(\pi, \psi)$, for any possible perturbation $\delta \in \Delta$.

The goal-seeking approach aims to find a satisfying solution that is within an acceptable tolerance limit. Such an approach is useful when it is not possible, or desirable, to construct a precise model of a system. Consequently, in this paradigm, the control of a complex system does not require a complex decision-maker. The next section presents an application of this approach to coordinate the discharge of a collection of batteries.

5. Discharge coordination

A simple view of the interaction between the decision-maker and the collection of batteries is depicted in Fig. 9. The decision-maker maintains the performance of the collection of batteries within established limits by selecting alternate actions from Π, while considering only its simple view of the environment, Ξ. When the batteries are discharging, the terminal voltage of the collection of batteries, V_t, is used as a measure of the performance and evaluated using the evaluation mapping, Ω. As long as the V_t is within an acceptable range, as indicated by Γ, the decision-maker does not select new alternate actions. Whenever the performance is outside the acceptable range, the decision-maker selects new alternate actions to maintain the performance within the acceptable range. The discharge process is terminated when the batteries can no longer supply the minimum voltage required by the load.

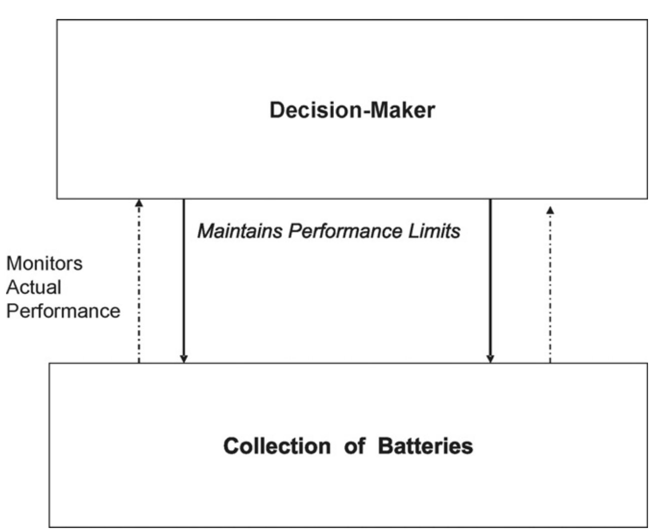

Fig. 9. Simple view of goal-seeking interactions.

In the remainder of this section, we present a goal-seeking formulation of the problem of coordinating battery discharge. The results demonstrate the improved performance of the decision-maker relative to the performance of scheduling methods discussed in Section 2.5.

5.1. Specifying goal-seeking artifacts

The principal objectives for the decision-maker are to extend the lifetime of the collection of batteries and improve the discharge efficiency. The decision-maker must select actions that enable the collection of batteries to maintain the terminal voltage of the collection at System Level 3 within a tolerance limit specified for the load.

5.1.1. Alternate actions, Π

Alternate actions represent the choices that are available to the decision-maker. The decision-maker must, in each time-slice, determine whether a given battery b_i must discharge or rest in the following time-slice. Given that the collection has a total of n batteries, it suffices to use a binary vector, bv, of n elements, where $\mathrm{bv}(i) = 1$ if battery b_i is selected for discharge and $\mathrm{bv}(i) = 0$ otherwise. If we denote the set of all possible binary vectors of n elements as the set Π, then the decision-maker must select one member of this set in each time-slice so that the selected batteries would discharge in the following time-slice. The selection of an element in Π is determined by other several higher-level decisions discussed in the following paragraphs.

The number of batteries that are discharging, m, is directly related to the terminal voltage V_t that can be delivered by the collection of batteries. The decision-maker selects m and k in each time-slice from the set

$$\Pi_1 = \{(m, k) : m \leq n, k \leq n, m + k = n\}.$$

The choice of $\mathrm{bv} \in \Pi$ is restricted by the choice of $(m, k) \in \Pi_1$.

We represent the choice of battery selection strategies available to the decision-maker as

$$\Pi_2 = \{\mathrm{QS}, \mathrm{RS}, \mathrm{SW}, \mathrm{GS}\}$$

where QS represents queued selection, RS represents random selection, SW represents sliding window and GS represents SoC selection. The first three strategies have been discussed in Section 2.5 and SoC selection is discussed in Section 5.2. The decision-maker can change the selection strategy from one member of Π_2 to another, at a time-tick that is chosen dynamically by the decision-maker based on its observations. Such a time-tick can be selected from the set

$$\Pi_3 = \{1, 2, \ldots, f\}.$$

The set

$$\Pi_4 = \{1, 10, 15\}$$

represents the choices available for the duration of the time-slice T_s (min). Finally, the set

$$\Pi_5 = \{\mathrm{AhM}, \mathrm{OCVM}\}$$

where AhM represents ampere–hour measurement and OCVM represents open-circuit voltage measurement. These are two methods for estimating SoC that the decision-maker uses.

Thus, the decision-maker's choice of bv $\in \Pi$ is determined by the member in the set

$$\Pi_1 \otimes \Pi_2 \otimes \Pi_3 \otimes \Pi_4 \otimes \Pi_5$$

that is selected by the decision-maker. In our simulation, we selected a time-tick from Π_3 only once. The choice of a member in Π_2 was also made once at the beginning of the discharge and changed once during discharge. In a general setting, these choices can be made by a decision-maker more than once.

5.1.2. Uncertainties, Δ

We represent the set of uncertainties as

$$\Delta = \Delta_1 \otimes \Delta_2.$$

Δ_1 represents the set of perturbations that could occur to cause the unique behavior of individual batteries. We represent such perturbations by varying the parameters, $k_1, \ldots, k_5$, of the battery model shown in Eq. (1) randomly within 10% of the nominal values. Thus, each member $\delta_1 \in \Delta$ represents a particular set of values for the parameters $k_1, \ldots, k_5$. Such variation in the parameters affect the stored charge, life time and voltage of each battery. In our simulation, the decision-maker does not consider members of Δ_1 before making decisions—the effects of these variations manifest as different dynamic behaviors of individual batteries.

Δ_2 represents the perturbations that occur during battery manufacturing. We aggregate these variations as a difference in the initial charge of the battery. To represent such variations, we vary the initial charge of an individual battery randomly by 32% of 0.86 Ah.

5.1.3. Consequences, Ψ

The consequences represent the terminal voltage of the collection of batteries. When selecting actions for the next time-slice, the decision-maker computes an expected consequence using the reflection function Ξ that is defined in Section 5.1.4. When the batteries in the collection are discharging, the decision-maker measures the terminal voltage V_t at System Level 3 (Fig. 2). Clearly, if an action is selected and some perturbation occurs, the measured voltage of the collection of batteries would reflect the effect of the perturbation.

5.1.4. Reflection, Ξ

The reflection function, Ξ: $\Pi \otimes \Delta \rightarrow \Psi$, represents the decision-maker's view of the environment. To emphasize that the consequence identified by the reflection is an estimated value, we use $\hat{\psi}$ to denote an estimated consequence.[2] Whenever the decision-maker has to select among a set of actions, $\pi \in \Pi$, it uses Ξ to estimate the consequence if a given perturbation, $\delta \in \Delta$, occurs.

[2] The actual consequence ψ, i.e., the measured value V_t, is not necessarily the same as the estimated value, $\hat{\psi}$.

The expected consequence is computed based on the decision-maker accounting for some uncertainty, i.e., $\hat{\psi} = \Xi(\pi, \hat{\delta})$. If V_{avg} represents the average voltage that can be delivered by a selected battery, then a simple function that can be used for Ξ is

$$\hat{\psi}_{m,k}(t_i) = V_{avg}(t_{i-1}) \times m.$$

To account for uncertainties, we adjust V_{avg} by the estimated SoC in the selected batteries as

$$\hat{\psi}_{m,k}(t_i) = \sum_{j=1}^{n} V_{avg}(t_{i-1}) \times \frac{\mathrm{SoC}_j^{t_i} \times \mathrm{bv}(j)}{\mathrm{SoC}_{avg}^{t_{i-1}}}$$

where bv(j) represents the jth element of the choice Π_1; bv(j) = 1 if the corresponding battery is discharging and 0 otherwise. $\mathrm{SoC}_j^{t_i}$ is the estimated SoC for battery b_j at time-tick t_i. In our simulation, this is estimated by varying parameters k_i of the battery model randomly by 10%; and the open-circuit voltage is computed using the battery model (Eq. (1)). If voltage from the model is less than $V_{EoD} = 1.25$ V, $\mathrm{SoC}_j^{t_i}$ for b_j is set to 0 to reflect that the battery is completely discharged.

5.1.5. Evaluation Set, Λ

The Evaluation Set (or performance scale) is a metric used to compare outcomes of selected actions. This scale helps the decision-maker to determine which alternate action is preferable over other choices. We used the closed interval [0, 1] on the real line as the scale—with 0 representing an undesirable choice and 1 representing the most desirable choice.

5.1.6. Evaluation mapping, Ω

The evaluation mapping is a function that maps a selected action and its consequence (estimated or measured) to a value, $\lambda_i \in \Lambda$, on the performance scale, i.e., Ω: $\Pi \otimes \Psi \rightarrow \Lambda$. The decision-maker uses this mapping both to select alternate actions using the estimated consequences and to evaluate the consequence of selected actions based on measured values.

When selecting batteries for discharge, the expected consequence should be close to the minimum voltage $V_{min} = 14$ V, otherwise, there is needless dissipation of energy. Similarly, we would like the measured consequence also to be close to V_{min}. We captured these requirements as

$$\Omega(\pi, \tilde{\psi}_{m,k}) = \begin{cases} \exp^{-(\tilde{\psi}_{m,k} - V_{min})} \times \dfrac{1}{m} \times \displaystyle\sum_{j=1}^{n} \mathrm{SoC}_j^{t_i} \times \mathrm{bv}(j), & \text{when } \tilde{\psi}_{m,k} \geq V_{min} \\ 0, & \text{when } \tilde{\psi}_{m,k} < V_{min} \end{cases} \quad (5)$$

$\tilde{\psi}_{m,k}$ represents the estimated consequence $\hat{\psi}_{m,k}$ when Ω is being used by the decision-maker to select actions. When evaluating the measured consequence, $\tilde{\psi}_{m,k}$ represents V_t. In addition to representing the requirement that the selected batteries can supply the voltage desired, the evaluation mapping also ensures that only the smallest necessary number of batteries is selected.

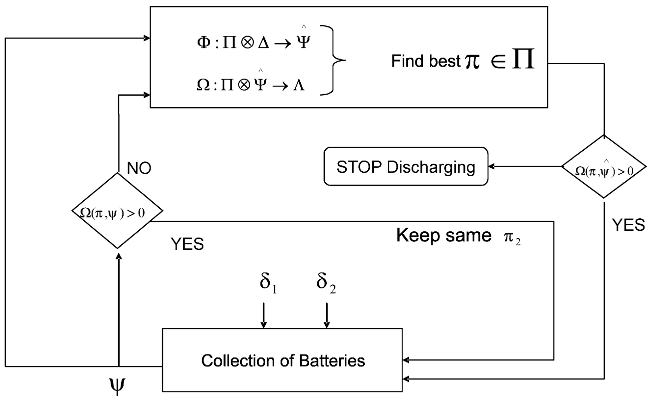

Fig. 10. Detailed view of goal-seeking interactions.

5.1.7. *Tolerance function, Γ*

The tolerance function, Γ: $\Pi \otimes \Psi \to \Lambda$, is a bound on the amount of variation that can be tolerated in performance before a solution is considered as unsatisfactory. During discharge, we required $V_t > 14$ V, i.e., we required $\Omega(\pi, \tilde{\psi}_{m,k}) > 0$ With $V_{min} = 14$ V.

5.2. *Execution of the decision-maker*

Fig. 10 shows the interactions between the collection of batteries and a decision-maker that coordinates its discharge.

In our simulations, the discharge process is initiated by selecting batteries using the sliding window strategy (in Π_3) and a time-slice duration of $T_s = 10$ min (in Π_5). When the batteries are discharging, the decision-maker measures the terminal voltage of the collection at System Level 3. This measurement is used to calculate an estimated average terminal voltage of each battery that is discharging; when this average value drops below 1.275 V, the decision-maker changes the selection strategy at a time-tick (in Π_4) to SoC selection with $T_s = 1$ min.

Since the estimated SoC is a critical factor of coordinating the discharge of a collection of batteries, we explored four different methods to estimate SoC. To support the SoC selection strategy, the terminal voltage is measured at System Level 1 for each battery that is discharging. We first assumed that the rate of change of terminal voltage at System Level 1 is correlated to the SoC of the battery and used this rate as a guide to select batteries with smaller rates of change. Second, we used a pre-constructed estimation table that mapped an open-circuit voltage measurement to SoC. Third, we used both a linear relationship between V_{oc} and SoC and the cubic polynomial relationship (Eq. (2)). Finally, we combined the method of estimating charge discussed in Section 3.3 and the method of ampere–hour counting to estimate the SoC. One can argue that because ampere–hour counting is used, this is approach uses the concept of “state”. This use of “state” is limited to the purpose of estimating whether or not a particular battery can be discharged in the next time-tick. In this approach, we have not attempted to construct a state-transition model for the collection of batteries.

The collection of batteries were discharged in all the three scenarios discussed for selection strategies in Section 2.5. We observed that in CCD and CPD, it was most effective to estimate SoC using a combination of charge estimation and ampere–hour counting where V_{oc} was mapped to SoC using the cubic polynomial in Eq. (2). For RVCD, we observed that the SoC estimation based on a pre-constructed table was more effective than an estimation using the cubic polynomial.

As in Section 2.5, the results presented in the graphs in this section are based on a single execution of the simulation. Results presented in Tables 1 and 3 are the average values obtained from 10 executions of the simulation.

Table 3
Performance of discharging methods at 250 mA

Scenario	Lifetime (min)	Energy (Wh)	Remaining capacity (Ah)
Queued selection			
CCD	231.9	14.909	8.335
CPD	240.5	15.025	8.257
RVCD	259.3	15.196	8.114
Random selection			
CCD	337.53	21.801	3.103
CPD	352.29	22.012	2.938
RVCD	375.5	22.128	3.303
Sliding window selection			
CCD	345.83	22.261	2.428
CPD	353.46	22.09	2.514
RVCD	384.16	22.332	2.538
Goal-seeking			
CCD	425.7	25.618	0.199
CPD	411.0	25.683	0.169
RVCD	472.4	25.685	0.561

5.3. *Simulation results*

Fig. 11 shows the stored charge versus time of the collection of batteries in CCD. Note that the lifetime of the collection of batteries has been significantly extended over the lifetime observed under the sliding window selection strategy in Fig. 7. In addition, note that the remaining charge in the collection is lower than the remaining charge observed in sliding window selection. Fig. 12 shows that the goal-seeking approach per-

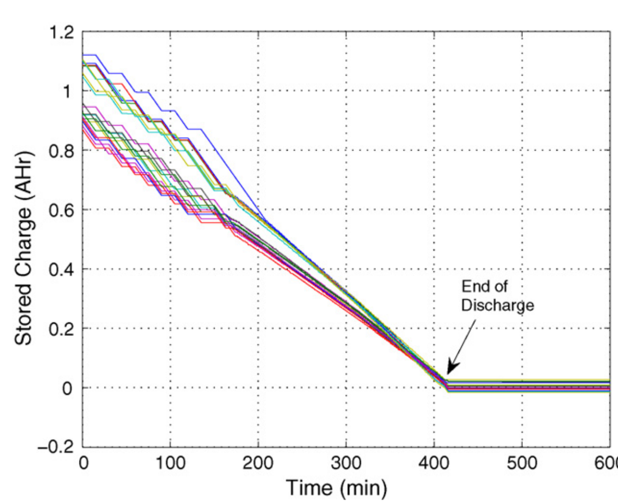

Fig. 11. Stored charge vs. time in CCD 250 mA.

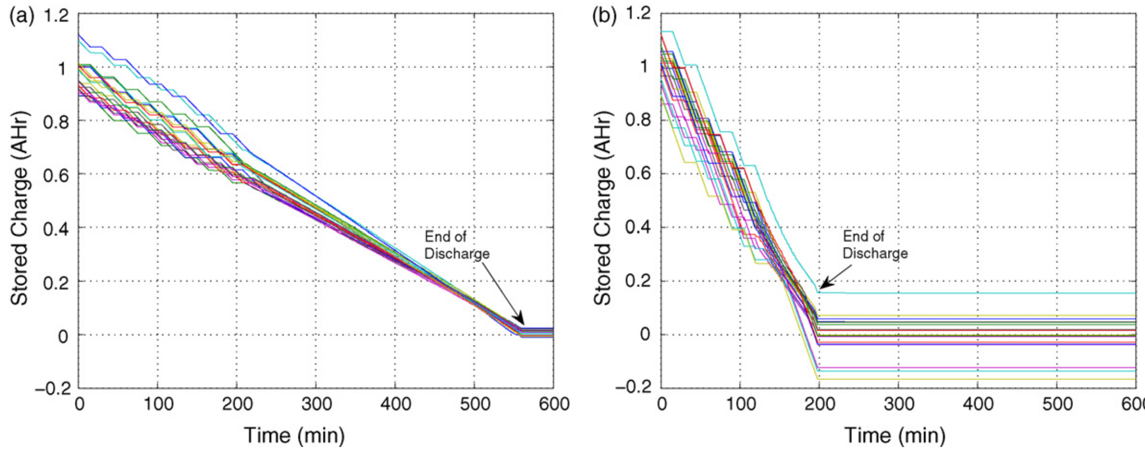

Fig. 12. Effects of discharge current in CCD. (a) 185 mA and (b) 500 mA.

forms as expected when the discharge current is changed to 185 or 500 mA.

Fig. 13 shows stored charge versus time of the collection of batteries in CPD. Once again, the decision-maker extends the lifetime and reduces remaining charge when compared to the performance of selection strategies that is shown in Table 3.

Fig. 14 shows stored charge versus time in RVCD when the SoC was estimated using the cubic polynomial. Fig. 15 shows a similar relationship when the decision-maker used an estimation table to estimate SoC. It may be noticed that the remaining charge at the end of discharge is reduced considerably when an estimation table is used in RVCD.

5.4. Discussion

Table 1 shows the baseline performance for a collection of 12 batteries in which there are no spares. This baseline allows us to recognize that when we discharge a collection of 20 batteries in which 8 batteries are available as spare batteries, the performance of the collection is significantly improved. Table 3 shows the performance of all the methods in the three discharge scenarios. It can be noticed that the decision-maker based on

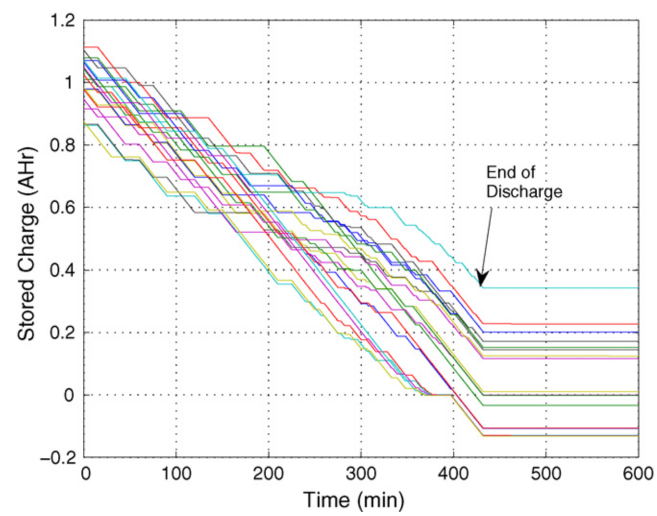

Fig. 14. Stored charge vs. time in RVCD using cubic polynomial 250 mA.

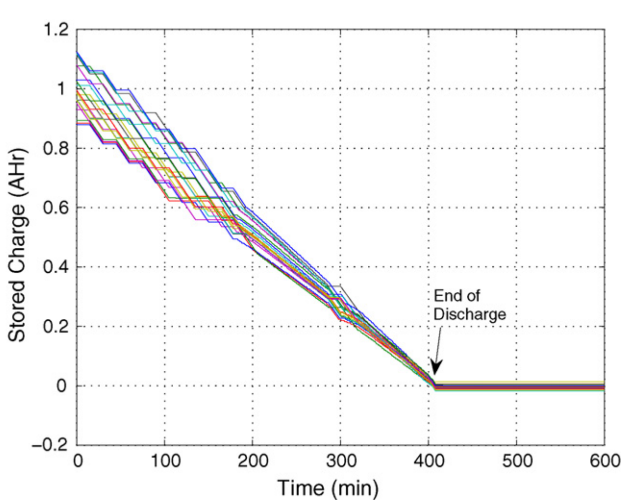

Fig. 13. Stored charge vs. time in CPD 250 mA.

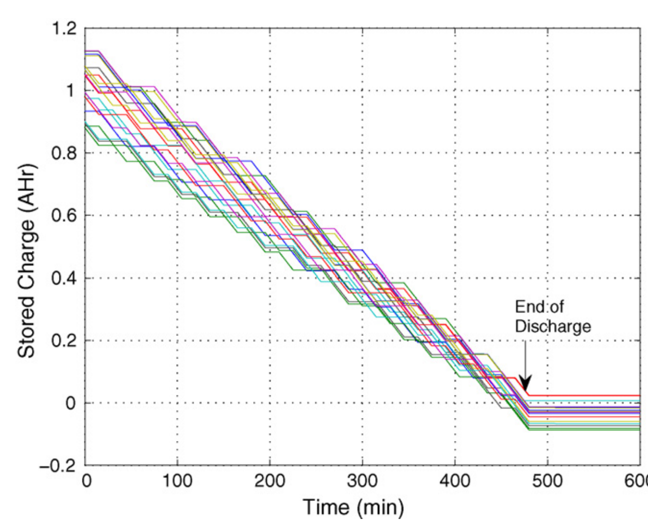

Fig. 15. Stored charge vs. time in RVCD using estimation tables 250 mA.

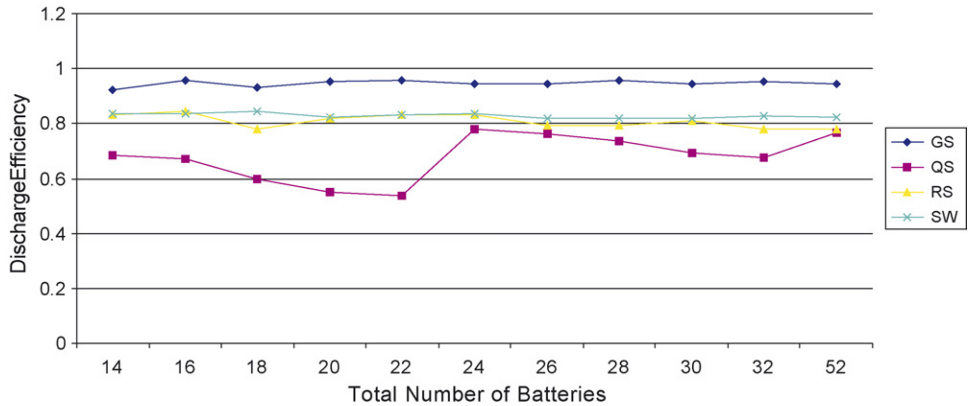

Fig. 16. Discharge efficiency.

Table 4
Effects of discharge currents—goal-seeking

Scenario	Lifetime (min)	Energy (Wh)	Remaining capacity (Ah)
CCE			
185 mA	572.8	22.485	0.194
250 mA	425.7	25.618	0.199
500 mA	193.2	23.805	0.764
CPD			
185 mA	408.9	25.508	0.145
250 mA	411.0	25.683	0.169
500 mA	407.5	25.475	0.167
RVCD—polynomial			
185 mA	554.31	22.415	2.997
250 mA	410.0	22.499	3.065
500 mA	201.5	22.290	2.592
RVCD—estimation table			
185 mA	632.90	22.272	0.708
250 mA	472.40	25.685	0.561
500 mA	220.40	24.185	1.890

the goal-seeking paradigm performs better than all the other methods in all the discharge scenarios.

Table 4 shows the performance of the new decision-maker when the discharge current varies. Comparing the values in this table with the ones in Table 2 we can see that the new decision-maker based on the goal-seeking paradigm consistently performs better than the sliding window selection strategy.

We defined the discharge efficiency metric to capture how effectively the decision-makers could utilize the available energy while extending lifetime. Table 5 presents the discharge efficiency of all the methods discussed in this paper when the coefficients were selected to emphasize that it is important to minimize the remaining capacity. It may be noted that the decision-maker based on the goal-seeking paradigm improves the efficiency consistently.

Table 5
Discharge efficiency

Selection method	CCD	CPD	RVCD
Queued selection	0.577	0.576	0.582
Random selection	0.805	0.812	0.805
Sliding window	0.832	0.834	0.830
Goal-seeking	0.943	0.931	0.905

The principle of using redundancy to extend lifetime and efficiency is not new in engineering. However, since the battery scheduling methods [2] and the decision-maker based on the goal-seeking formulation are both based on this principle, it is important to compare how well the two approaches utilize the available redundancy.

Fig. 16 shows how the discharge efficiency varies when the number of spare batteries available is changed. It is seen that the discharge efficiency when using queued selection (QS) increases only when certain number of spare batteries are available. This reflects the fact that this strategy is effective only when the number of spare batteries is at least as many as the minimum number of batteries that are required to meet the voltage demands of the load. The discharge efficiency of sliding window (SW) selection is consistent and steady; in contrast for the random selection (RS) we see that the efficiency varies. This variation is because of wide variance of the remaining charge in the collection at the end of discharge. The new decision-maker that is based on the goal-seeking paradigm (GS) performs consistently better than all the other methods when the utilization of available energy in the collection is important; i.e., $c_1 = c_2 = 0.25$ and $c_3 = 0.5$. The discharge efficiency of the decision-maker based on the goal-seeking paradigm increases from 0.93 to 0.96 when the number of spare batteries changes from two to four. Thus, our decision-maker is able to better utilize the available spare batteries, especially when the number of spare batteries is low.

6. Conclusions

We developed a goal-seeking formulation for the problem of coordinating the discharge of a collection of batteries motivated by the need to cope with the uncertainties that are inherent in batteries. We used this formulation to design a new decision-maker and evaluated the performance of the decision-maker while discharging a collection of batteries in constant current, constant power and randomly varying current discharge scenarios. To evaluate the performance of the decision-maker, we used lifetime, energy supplied and remaining capacity at the end of discharge as factors and defined a metric called discharge efficiency that combined all these factors. In all three discharge scenarios, the new decision-maker, by making use of the individual battery voltage measurements, consistently achieves better discharge efficiency than the scheduling methods reported in the literature.

Based on these results, we are encouraged to consider designing new decision-makers that optimally charge a collection of batteries in the presence of uncertainties.

Acknowledgements

The first author was supported in part by a College of Engineering, Research Initiation Award during this investigation. We thank the anonymous referees for the comments that helped us improve the presentation of this paper.

References

[1] H.J. Bergveld, W.S. Kruijt, P.H.L. Notten, Battery Management Systems, Kluwer Academic Publishers, 2002.
[2] L. Benini, A. Macii, E. Macii, M. Poncino, R. Scarsi, IEEE Trans. Very Large Scale Integration (VLSI) Syst. 11 (6) (2003) 1136–1143.
[3] M.D. Mesarovic, Y. Takahara, Mathematical Theory of General Systems, Academic Press, 1974.
[4] M.D. Mesarovic, D.L. McGuiness, D.A. West, Cybernetics of Global Change: Human Dimension and Managing of Complexity. MOST Policy Paper 3, UNESCO, 1996.
[5] S. Sastry, Assembly Autom. 24 (2) (2004) 201–209.
[6] Y. Takahara, M.D. Mesarovic, Organizational Structure, Academic Press, 2003.
[7] T. Hartley, C. Lorenzo, NASA Aerospace Battery Workshop, Huntsville, 2002.
[8] P.T. Krein, R. Balog, 24th International Telecommunications Energy Conference, 2002, pp. 515–523.
[9] J. Chatzakis, K. Kalaitzakis, N.C. Voulgaris, S.N. Manias, IEEE Trans. Ind. Electron. 50 (5) (2003) 990–999.
[10] N.H. Kutkut, H.L.N. Wiegman, D.M. Divan, D.W. Novotny, IEEE Trans. Aerospace Electron. Syst. 34 (1) (1998) 235–246.
[11] N.H. Kutkut, H.L.N. Wiegman, D.M. Divan, D.W. Novotny, IEEE Trans. Ind. Appl. 35 (1) (1999) 28–35.
[12] Y.-S. Lee, M.W. Chen, K.L. Hsu, J.Y. Du, C.F. Chuang, Conference on Computers, Communications, Control and Power Engineering, vol. 3, 2002, pp. 2042–2045.
[13] C.S. Moo, Y.C. Hsieh, I.S. Tsai, IEEE Trans. Aerospace Electron. Syst. 39 (2) (2003) 704–710.
[14] A. Jossen, V. Spath, H. Doring, J. Garche, Telecommunications Energy Conference, 1999, pp. 1–6.
[15] F.C. Lee, J.D. van Wyk, D. Boroyevich, G.-Q. Lu, Z. Liang, P. Barbosa, IEEE Control Syst. Mag. (2002) 4–23.