

On Dynamic Reconfiguration of A Large-Scale Battery System

Hahnsang Kim and Kang G. Shin
Real-Time Computing Laboratory

Department of Electrical Engineering and Computer Science
The University of Michigan

Ann Arbor, MI 48109-2121, U.S.A.

Email: {hahnsang, kgshin}@eecs.umich.edu

Abstract—Electric vehicles powered with large-scale battery packs are gaining popularity as gasoline price soars. Large-scale battery packs usually consist of an estimated 12,000 battery cells connected in series and parallel, which are susceptible to battery-cell failures. Unfortunately, current battery-management systems are unable to handle the inevitable battery-cell failures very well. To address this problem, we propose a dynamic reconfiguration framework that monitors, reconfigures, and controls large-scale battery packs online. The framework is built upon a syntactic bypassing mechanism that provides a set of rules for changing the battery-pack configuration, and a semantic bypassing mechanism by which the battery-cell connectivity is reconfigured to recover from a battery-cell failure. In particular, the semantic bypassing mechanism is dictated by constant-voltage-keeping and dynamic-voltage-allowing policies. The former policy is effective in preventing unavoidable voltage drops during the battery discharge, while the latter policy is effective in supplying different amounts of power to meet a wide-range of application requirements. Our experimental evaluation has shown the proposed framework to enable the battery packs to be 9 times as fault-tolerant as a legacy scheme.

Key Words: Programmability, self-reconfiguration, self-healing, dynamic supply voltage

I. INTRODUCTION

Demand for electric vehicles with hybrid drive has soared worldwide due mainly to a recent sharp increase in fuel prices. According to a recent survey [7], in 2008 alone, 36.0% motorists worldwide want to buy a car with hybrid drive while 45.8% of them are interested in buying full-electric cars. Electric cars are powered entirely with electrical energy from tens of thousands of battery cells. These battery cells are grouped and assembled as a set of battery packs. Individual cells in a pack, which are exposed to, and must operate in a harsh environment, have different operating characteristics due to difference in their manufacturing tolerances, uneven temperature conditions across the pack, or non-uniform ageing patterns. These, in turn, have crucial effects on the charge/discharge of battery cells. In a series chain of battery cells, a *weak* battery cell with low capacity reaches its full charge state well before the rest of the battery cells in the chain, hence overcharging and overheating itself. On the other hand, when the weak cell cannot reach its full charge owing to a high *self-discharge* [2] and/or a short-circuited cell, good battery cells could become overcharged. In a series chain

of battery cells, an open-circuited cell causes the others in the chain to become open-circuited as well. All of these phenomena eventually lead to a *battery-cell failure*, which is inevitable especially in large-scale battery packs.

The most commonly-used method for managing a large-scale battery system is module-based, where battery cells are grouped into smaller modules of battery cells, each of which is monitored, controlled, and balanced by the corresponding *local* controller while a group of modules are managed by a *central* controller [21]. In such a modular battery-management system, individual electronic control units (ECUs) collect information—such as cell voltage and current, temperature, etc.—on their serially-connected battery packs via an equalizer connected to each battery cell, and then process and report the collected information to the central ECU responsible for making the local ECUs work properly. To handle the case of any battery cell becoming open-circuited, the battery pack can be made configurable [8] by adding additional controllable switches around each battery cell, detouring any faulty battery cell [1, 6, 22]. These solutions are, however, dedicated to micro-scale batteries based on *static configuration* and hence limited to physical processes, requiring more interactive computations online to cope with a large-scale battery-management system.

There are two main challenges in developing a dynamic reconfiguration framework for large-scale battery-management systems. First, the framework should be able to reconfigure battery connectivity online, upon detection of a battery-cell failure. Healthy battery cells should also be kept in use, possibly in the form of two hierarchical layers of connectivity: battery cells in each pack (cell-level) and packs in the entire battery system (pack-level). Second, unlike battery-powered portable devices, a large-scale battery-management system, especially for electric vehicles, requires multiple output terminals of the power source (from the battery packs), supplying different voltages for different applications and/or devices. Physical separation of battery packs is, however, rarely an option mainly for cost reasons.

In this paper we propose a novel dynamic reconfiguration framework to monitor, reconfigure, and control a large set of battery packs online, achieving high resilience to battery-cell failures and thus extending the battery packs' *operation-*

time. By operation-time, we mean cumulative time of drawing current from a battery cell until it no longer delivers the current to the load. In other words, when the terminal voltage of the battery cell falls below the *cutoff voltage*—a threshold by which to determine whether the battery cell has been fully discharged, the operation-time ends. Our framework is built upon *syntactic* and *semantic bypassing mechanisms*. The syntactic bypassing mechanism provides a set of rules for changing configurations. The semantic bypassing mechanism reconfigures the battery connectivity in case of a battery-cell failure with *voltage balancing* across the series chains of the battery cells/packs under consideration, thereby keeping the supply voltage constant, called a *constant-voltage-keeping policy*. By contrast, the *dynamic-voltage-allowing policy* changes the supply voltage to the underlying applications and/or devices online. These two policies are applied throughout the entire battery operation-time, thus extending the operation of the battery packs.

The main contributions of this paper are two-fold as follows, by addressing the above-mentioned two challenges. First, our dynamic reconfiguration framework, to our best knowledge, is the first comprehensive way of coping with battery-cell failures in large-scale battery systems. In particular, the connectivity of battery cells in and out of the battery packs can simultaneously be configured, thus effectively providing fine-grained supply voltage and improving flexibility and scalability in reconfiguration. Second, our framework “customizes” supply voltage online according to the needs of the underlying applications and devices. Also, the framework can accommodate multiple applications/devices at the same time by virtually partitioning and re-allocating the power source drawn from the battery packs to meet the applications’ requirements. This customizability enhances the applicability and deployability of our reconfiguration framework.

The rest of the paper is organized as follows. Section II provides background information on dynamics of batteries, a generic battery-management system, and the motivations of this work. Section III describes the design of a dynamic reconfiguration framework including battery monitoring and the two bypassing mechanisms: constant-voltage-keeping and dynamic-voltage-allowing policies. Section IV evaluates the performance of our reconfiguration framework. We discuss the related work in Section V and conclude the paper in Section VI.

II. BATTERY, BATTERY-MANAGEMENT SYSTEM, AND MOTIVATIONS

We first provide a physical insight into rechargeable batteries, then describe a generic battery-management system for battery-powered devices which is followed by the description of motivations behind our work.

A. Battery Dynamics

A rechargeable battery cell (e.g., NiCd, NiMH, and Lithium-ion) is capable of converting chemical energy to electrical energy, and vice versa, via electrochemical oxidation and

reduction reactions [3, 11]. These reactions involve the exchange of electrons through the load between electro-active species in two electrodes inside the battery cell, generating a flow of electric current. Ideally, the total number of current units, or Coulomb, from a battery cell will always be the same throughout its entire life cycle. In reality, however, the characteristics of a battery cell are nowhere close to being ideal due to the uncertainty of reaction kinetics and diffusion processes and/or active material dissolution [11] in the battery cell over time.

We describe below the characteristics of a real-life battery for completeness (see [2, 19] for details). First, the battery terminal voltage is not constant during its discharge; voltage drops non-linearly with a discharge rate. The higher the discharge rate, the steeper does the voltage drop. For this reason, a DC-DC convertor is required to shift and stabilize the supply voltage [9]. Second, battery capacity varies with the discharge rate; the higher the discharge rate, the lower the battery capacity. Third, batteries have limited charge recovery effects at a high discharge rate; a high load current for a short period of time causes a higher concentration gradient among electro-active species, making the unused charge unavailable due to the lag between reaction and diffusion rates [11]. Thus, when the battery is allowed to rest for some time at a low (or zero) charge rate, the voltage that dropped temporarily goes back up. Last, temperature also affects internal resistance and full charge capacity; the lower the temperature, the higher the internal resistance, thereby reducing full charge capacity. On the other hand, the high temperature leads to *self-discharge* [2], reducing the actual capacity to be delivered [3, 11]. In addition to these characteristics, some batteries, e.g., NiCd batteries, are known to have *memory effect* [3], while Lithium-ion batteries do not.

Apart from temporary changes in battery capacity, as listed above, the batteries lose their capacity to some extent due to unwanted side reactions including electrolyte decomposition, active material dissolution, and passive film formation [11, 19], thereby increasing internal resistance and ultimately causing a battery-cell failure. Several possible failure modes [13] exist, making it difficult to predict. First, an *open circuit* can be a fail-safe mode for other battery cells in the series chain including an open-circuited battery cell, because it limits further damage to the other battery cells. However, this is not useful to the applications because all the battery cells in the series chain become open-circuited and unusable. Second, a *short circuit* that has an abnormal low electrical resistance incurs almost no voltage drop, so that the rest of the battery cells in the chain could be slightly overloaded while the whole battery pack (i.e., a set of the battery cells) remains functional. Last, a possible *explosion* is avoided via a protection circuit that detects and stops an extremely high current.

B. Battery-Management System

A generic battery-management system offers functions to optimally use battery cells for applications. Examples of the functions [3, 11] are handling proper charge and discharge

of battery cells, and protecting battery cells from misuse. A battery-management system is in general composed of (1) a battery cell or a battery pack; (2) a controller that estimates the battery pack condition, including the *state of charge* and the *state of health*; (3) a control unit that turns on/off the charge/discharge, depending on control signal from the controller; (4) a thermistor that measures the battery temperature inside the pack; (5) a temperature fuse that cuts off the current if the control unit experiences abnormal heating; (6) a power module that is capable of charging the battery pack, monitoring the charge for preventing overcharge from inflicting damage to the battery pack, and often powering the load directly; and (7) a DC-DC converter that converts the electrical energy from the battery pack/power module into magnetic energy and back into electrical energy with the lowest possible supply voltage that is suitable for operations of the load. It is also worth mentioning that DC-DC converters are most efficient when their input voltage is closest to their output voltage [12].

Fig. 1 illustrates a simple battery-management system. During the charge, the current is supplied to both the battery and the load, while it is supplied by the battery pack during the discharge. In this paper, we consider the relation between one controller and its multiple control units each of which monitors and controls a single battery cell, and then extend to the relation between the controllers, which will be detailed in Section III.

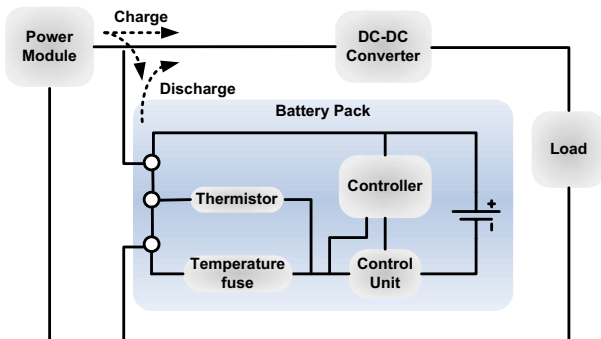


Fig. 1. A schematic diagram of a generic battery-management system based on a model [3]

C. Motivations

As mentioned earlier, we are interested in a large-scale battery-management system for complicated applications, such as electric vehicles, that can cope with tens of thousands of battery cells connected in series and parallel. Brand-new electric cars that General Motors manufactures are powered by a set of battery packs that consist of an estimated 12000 battery cells in total each of which is monitored and controlled by an ECU, with approximately 600 volts of supply voltage provided.¹

There are two main motivations behind this work. First, in a series chains of many battery cells, failure of a single battery

¹This information was obtained from our private communication with specialists in powertrain and battery-management systems divisions at General Motors in September 2008.

cell can prevent the entire battery system from functioning properly. Simply replacing the faulty cell with a new one is rarely an option because the new battery cell may become short-circuited. Second, re-allocating the power source drawn from the battery cells online to meet the needs of underlying applications has great potential for cost efficiency. For instance, the power source from the battery cells diverges via several output terminals. In case of electric cars, one provided in 600-volt current and the others in 12-volt current each flow into the engine and the underlying applications, such as power windows, windshield wipers, or heating, ventilation, and air conditioning systems. Some of these applications can require more power as their demand scales up. Meanwhile, not all applications may be active simultaneously. In such a case, a portion of the power source that is temporarily left unused by some of the applications can be re-allocated to the others needing more power, thus satisfying their requirements efficiently.

These two aspects lead us to the development of a re-configuration framework for fault-tolerant large-scale battery-management systems.

III. A DYNAMIC RECONFIGURATION FRAMEWORK

This section describes the architecture, the battery-management model, and two bypassing mechanisms, i.e., syntactic bypassing and semantic bypassing, based on which battery-cell connectivity is adaptively self-reconfigured. Also, the battery-management model is extended to support multiple input/output terminals, followed by the comparison of capacity versus operation-time of battery packs.

A. The Architecture

The design of a dynamic reconfiguration framework for large-scale battery-management systems is guided by a simple principle: one should be able to bypass *any* battery cell. At the same time, the architecture includes as few switches to be placed around the battery cells as possible. Fig. 2 illustrates the architecture of the dynamic reconfiguration framework. The controller enforces the two policies by orchestrating its control units. The enforcement of the policies is in the form of turning relevant switches on/off via the control unit. The control unit controls two sensors to monitor the battery condition, bypass and series switches (i.e., $S_B - (2)$ and $S_S - (3)$) by which to either skip itself or connect to the next battery cell, input and parallel switches (i.e., $S_I - (1)$ and $S_P - (4)$) by which to connect the battery cells in series or parallel, input and output terminal switches (i.e., S_{IT} and S_{OT}) by which to allow the battery pack to provide multiple terminals. For instance, when all the battery cells are to be connected in parallel, each control unit i turns $(S_{i,I}, On)$ and $(S_{i,P}, On)$, while it turns $(S_{i,S}, On)$. In addition, when $(S_{i,IT}, On)$ and $(S_{i,OT}, On)$ stay for all i , the interface of the battery pack has single input and single put terminals. When the selector switch on the control unit is turned off to the feedback, the corresponding battery cell remains open.

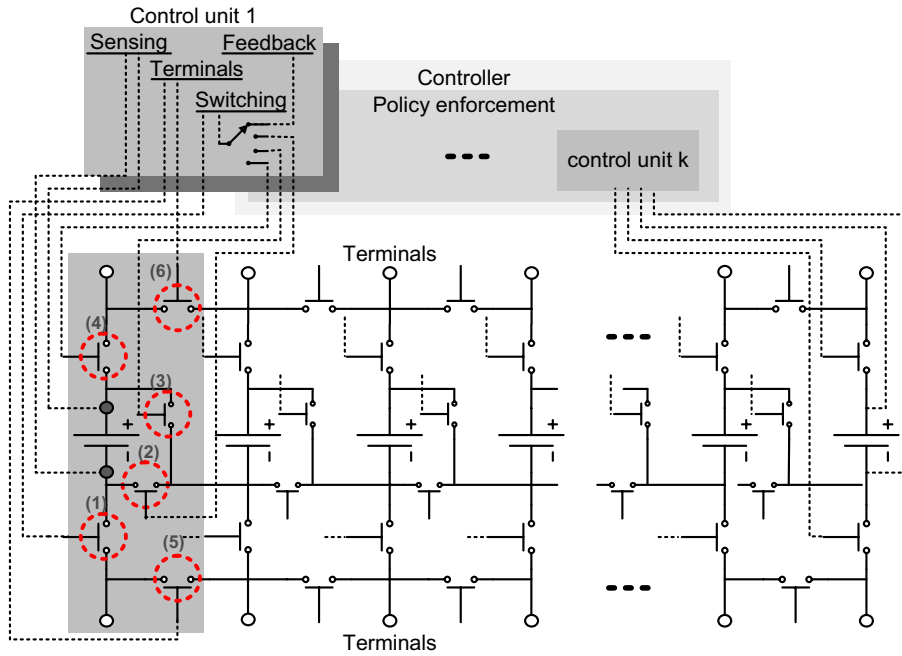


Fig. 2. Schematic diagram of the dynamic reconfiguration framework for a large-scale battery-management system that consists of circuit logics and a controller that encompasses k control units. Each control unit is responsible for 6 switches, i.e., input switch ($S_{1,I}$ -(1)), bypass switch ($S_{1,B}$ -(2)), series switch ($S_{1,S}$ -(3)), parallel switch ($S_{1,P}$ -(4)), input-terminal switch ($S_{1,IT}$ -(5)), and output-terminal switch ($S_{1,OT}$ -(6)). In addition to these switches, the control unit monitors the condition of its battery cell via two sensors attached to the ends of the battery cell.

The architecture of the dynamic reconfiguration framework is represented as $\Psi = (E, F, S, D)$, where E is an array of sensors, $\{E_1, \dots, E_i, \dots, E_k\}$ each of which reads the voltage and the current of a corresponding battery cell. F denotes an array of feedback switches, $\{F_1, \dots, F_i, \dots, F_k\}$, that the controller maintains to determine which cell to be bypassed. When a battery-cell failure in device i is detected, (F_i, On) is turned. S denotes an array of the switches, $\{S_1, \dots, S_i, \dots, S_k\}$, where S_i is composed of $S_{i,I}$, $S_{i,O}$, $S_{i,B}$, $S_{i,S}$, $S_{i,P}$, $S_{i,IT}$, and $S_{i,OT}$. D is a set of battery devices, $\{D_1, \dots, D_i, \dots, D_k\}$. The connectivity of these devices is thought of as an $n_s \times n_p$ matrix:

$$\begin{pmatrix} D_{1,1} & \cdots & D_{1,n_p} \\ \vdots & \ddots & \vdots \\ D_{n_s,1} & \cdots & D_{n_s,n_p} \end{pmatrix}, \quad (1)$$

where n_s is the number of battery cells connected in a series chain and n_p is the number of the series chains connected in parallel. Throughout the paper, we use V_d and V_a to denote the voltage demand and the average voltage of battery cells (or a set of battery packs), respectively. Similarly, we define

$$f_N(\Psi) = \sum_{i=1}^k 1(F_i), \quad (2)$$

where $1(F_i)$ is an indication function, i.e., if (F_i, Off) holds then the function returns 1 else it returns 0.

B. Monitoring

Control units monitor changes in the state of charge (SOC) and voltage of its battery cells. The SOC of a battery cell can

be estimated by measuring and integrating the current flowing into and out of the battery cell over time, called Coulomb count [3]. In practice, voltage and temperature are also figured in as battery variables. Thus, function $f_{V,T}(\int Idt)$ [3] that is based on the content of the coulomb count returns SOC. On the other hand, in general, direct voltage measurement is not accurate enough to be used as an indicator because of its dependency on the discharge rate and temperature; voltage can be estimated by applying Kalman filter [14–16]. The definition and derivation of these estimations are, however, out of the scope of this paper. Instead, we assume that an integrated recursive function, $f_{V,I,T}(SOC, \int Idt)$, is given and returns $[V, SOC]$.

At every Δt , the controller checks the SOC of each battery cell via the corresponding control unit and triggers a *rotation* event if

$$\frac{\min(SOC_1, \dots, SOC_k)}{\max(SOC_1, \dots, SOC_k)} > \delta \quad (3)$$

holds, where δ denotes a threshold that bounds the maximum variation of SOC. Clearly, the larger the δ , the higher the battery cells become unbalanced. Furthermore, the variation needs to be adjusted with δ , in conjunction with Δt , because the larger the Δt , the larger the variation. In particular, Δt is inversely proportional to the discharge rate.

A faulty cell is generally regarded as a battery cell that can be charged as low as 80% nominal capacity and/or that has voltage as low as the cut-off voltage in a fully charged state. Thus, when battery cell i is determined faulty, (F_i, On) is turned in control unit i .

At every Δt , the controller also checks the average voltage

and triggers a *reconfiguration* event if

$$V_d \leq V_a \cdot n_s < V_d + \alpha \quad (4)$$

holds, where α specifies an upper bound of voltage unbalancing. Obviously, α is tuned based on the granularity in supply voltage.

C. Syntactic-Bypassing Mechanism

The syntactic-bypassing mechanism is characterized by three rules that describe how to bypass a battery cell (D_k) in different connection topologies as shown in Fig. 3. When the event of a battery-cell failure is triggered, i.e., (F_i, On) , the controller checks which switches are currently in control unit i . In the case of $(S_{i,S}, On)$, rule 1 is applied: control unit i turns off its current on-switch and on its bypass switch, i.e., $(S_{i,B}, On)$. If $(S_{i,S}, Off)$ holds, it is toggled. In rule 2, control unit i turns off its input switch. Then, control unit $i + 1$ that has no faulty battery cell turns on its input switch. In rule 3, if the parallel switch is on, control unit i turns it off and then control unit $i - 1$ that has no faulty battery cell turns on its parallel switch and off its series switch. These three rules are applied mutually exclusively.

Rule 1:

if $(S_{i,S}, On) \ \& \ (F_i, On)$
 Switch $(S_{i,S}, Off)$ and $(S_{i,B}, On)$;
else
 Switch $(S_{i,S}, On)$;
end if

Rule 2:

if $(S_{i,I}, On) \ \& \ (F_i, On)$
 Switch $(S_{i,I}, Off)$;
 Switch $(S_{i+1,I}, On)$, where (F_{i+1}, Off) ;
else
 Switch $(S_{i,I}, On)$;
end if

Rule 3:

if $(S_{i,P}, On) \ \& \ (F_i, On)$
 Switch $(S_{i,P}, Off)$;
 Switch $(S_{i-1,P}, On)$ and $(S_{i-1,S}, Off)$, where (F_{i-1}, Off) ;
else
 Switch $(S_{i,P}, On)$;
end if

Fig. 3. Syntactic bypassing rules

Although the syntactic bypass mechanism is fundamental in reconfiguring battery-cell connectivity, care should be taken in the case of multiple series chains connected in parallel so that the parallel groups may agree on the same voltage, called *voltage balancing* unless different output terminals are supported (to be discussed in Section III-E). Otherwise, inverse charge occurs from a higher voltage series chain to a lower

voltage series one, causing the average voltage between the two chains to drop.

D. Semantic-Bypassing Mechanism

The semantic-bypassing mechanism is applied to supply a wide range of voltages while abiding by voltage balancing across the parallel group of the series chains. To achieve this, two policies are defined. First, a *constant-voltage-keeping policy* is specified to keep the supply voltage as constant over the battery operation-time as possible in spite of the battery-cell failure. To this end, the series chain containing the faulty battery cell is bypassed. However, it is possible that the voltages of both used and unused healthy battery cells in the series chain may drift apart over time, resulting in unbalanced voltages between the battery cells within the series chain. For this reason, the rotation event is triggered during the monitoring, reconfiguring the battery-cell connectivity. For connectivity reconfiguration, battery cells at the lowest level of their SOC are singled out first. Fig. 4 illustrates the reconfiguration of battery cells in a series chain under the constant-voltage-keeping policy.

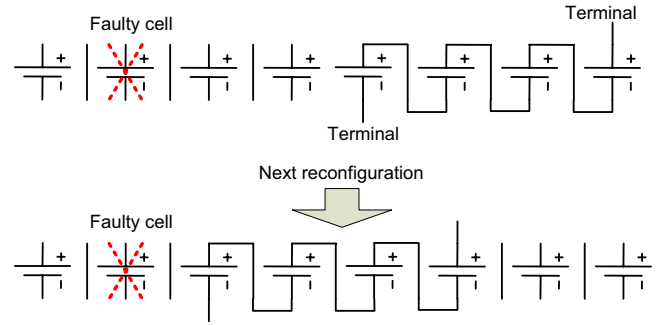


Fig. 4. Reconfiguration of battery-cell connectivity under the constant-voltage-keeping-policy: An instance of periodic reconfiguration appears in the lower part.

The number of battery cells to be bypassed is calculated as follows. Given V_d , n_s is first calculated by $\lceil \frac{V_d}{V_a} \rceil$; use of V_a offsets the nonlinear voltage drop during their operation-time. n_p is then derived from $\lfloor \frac{f_N(\Psi)}{n_s} \rfloor$, resulting in $(f_N(\Psi) - n_s \cdot n_p)$ healthy battery cells to be bypassed. This procedure repeats during the next Δt .

Alternatively, the *dynamic-voltage-allowing policy* is defined to improve the maximum deliverable power, given available battery cells, at the expense of a voltage drop that corresponds to a single battery-cell's voltage. Under the dynamic-voltage-allowing policy, healthy battery cells in each series chain are singled out as shown in Fig. 5; during Δt , n_p is fixed and n_s is then calculated by $\lfloor \frac{f_N(\Psi)}{n_p} \rfloor$, resulting in $(f_N(\Psi) - n_s \cdot n_p)$ healthy battery cells to be bypassed. As with the constant-voltage-keeping policy, the battery cells are singled out based on their SOC. To achieve the maximum deliverable power, $n_s \cdot n_p$ determines which policy to be selected; this will be evaluated in Section IV.

Once n_s and n_p are determined, connecting battery cells is straightforward based on the syntactic bypassing rules, as




Fig. 5. Reconfiguration of battery-cell connection under the dynamic-voltage-allowing-policy

shown in Fig. 6. At first, rule 2 is applied to open a new series chain and then rule 1 is followed to connect battery cells in series. Finally, rule 3 is applied to close the series chain. This procedure repeats until n_p parallel groups are created. As a consequence, the supply voltage of $V_d \cdot n_s$ is provided. The voltage, however, decays during the battery operation-time due to voltage degradation or battery-cell failure. For this reason, the reconfiguration event is triggered, reconfiguring the battery-cell connectivity with V_d fixed (varied) under the constant-voltage-keeping (dynamic-voltage-allowing) policy.

Battery-Cell Connection (Ψ, n_s, n_p) :

/* Rules 1, 2, and 3 are presented in Fig. 3 */

```

for  $i \leftarrow 1 : n_p$ 
  Apply Rule 2;
  repeat
    Apply Rule 1;
  until  $(n_s - 2)$  cells are connected
  Apply Rule 3;
end for

```

Fig. 6. Battery cells are connected in series and parallel

E. Extension to Multiple Battery Packs

The extension to the dynamic reconfiguration framework is represented as $\Upsilon = (E, F, S, \Psi)$, where $\Psi = \{\Psi_1, \dots, \Psi_2, \dots, \Psi_k\}$. Each Ψ_i is reconfigured via a ‘local’ controller and Υ via a ‘central’ controller. Obviously, the two policies described earlier can be applied in the central controller. Similarly to the local controller, the central controller determines Ψ_i as failure when (F_i, On) is detected; some battery cells in Ψ_i may be reused, so it is not a good idea to bypass Ψ_i . In other words, healthy battery cells in Ψ_i are still kept in use in such a way that bypassing Ψ_i is allowed only when change in the maximum deliverable power in Ψ is minimum. To this purpose, the smallest number of healthy battery cells in Ψ , n_m , and the sum of the numbers of healthy battery cells to be bypassed in each Ψ_k , n_b , are compared.

Thus, Ψ_i continues to be in use until $n_m > n_b$. This decision is systematically made via a decision function described in Fig. 7. For instance, when Ψ_1 , Ψ_2 , and Ψ_3 each have 6 battery cells, suppose Ψ_1 has one battery cell failed (i.e., $f_N(\Psi_1) = 5$) and Ψ_2 has two failed (i.e., $f_N(\Psi_2) = 4$), resulting in $n_m = 4$. In this case, Ψ_2 can be used by bypassing one and two healthy battery cells in Ψ_1 and Ψ_3 , respectively, since $n_b = 3$. On the other hand, in the case where Ψ_2 has an additional battery cell failed, Ψ_2 itself is bypassed, since $n_b = 5$.

Battery-Pack Bypassing Decision (Υ) :

$$n_m \leftarrow \min(f_N(\Psi_1), \dots, f_N(\Psi_k));$$

$$n_b \leftarrow \sum_{i=1}^n (f_N(\Psi_i) - n_m);$$

if $n_m \leq n_b$

return $\Upsilon - \{\Psi_i \mid n_m = f_N(\Psi_i)\}$;

end if

return Υ .

Fig. 7. Battery pack usability

A central controller reconfigures all the battery cells in Υ in conjunction with local controllers, generating a wide range of supply voltages for the load. Given V_d , the central controller calculates the number of battery cells to be connected in series for Ψ_k and Υ , i.e.,

$$n_s \cdot N_s = \lceil \frac{V_d}{V_a} \rceil, \quad (5)$$

where $n_s \leq f_N(\Psi_k)$ is the number of battery cells in a series chain in Ψ_k and $N_s \leq f_N(\Upsilon)$ is the number of battery cells in a series chain in Υ . This equation holds on the condition that if $f_N(\Psi_k) \leq f_N(\Upsilon)$, then $n_s \leq N_s$, or if $f_N(\Upsilon) \leq f_N(\Psi_k)$, then $N_s \leq n_s$. After n_s and N_s are resolved subject to the condition, n_p in Ψ_k is calculated by $\lfloor \frac{f_N(\Psi_k)}{n_s} \rfloor$. Likewise, N_p (the number of series chains connected in parallel) in Υ is calculated by $\lfloor \frac{f_N(\Upsilon)}{N_s} \rfloor$. As a consequence, the local controllers and the central controller apply the battery-cell connecting procedure in Fig. 6 with arguments of (Ψ, n_s, n_p) and (Υ, N_s, N_p) , respectively, thereby resulting in all the battery cells in and out of the battery packs configured in tandem.

F. Multi-Terminals

To provide multi-terminals with different supply voltages for various applications, a simple allocation policy is defined as shown in Fig. 8. The central controller attempts to resolve the number of parallel groups for all applications. Each application is characterized by its V_d . The V_d of application k determines $N_{s,k}$. The sum of $N_{p,1}, N_{p,2}, \dots, N_{p,k}$, given the total number of healthy battery cells, $N(\Upsilon)$, leads to q parallel groups for all the applications. At this point, if battery cells are left available, they are first distributed to high-priority applications, i.e., those with a high demand voltage. This distribution continues until the remaining cells are not enough to be distributed. Thus, the central controller allocates the power source of $N_{s,k} \times N_{p,k}$ to each application k .

Multi-Terminal-Based Grouping :

*/** $\{N_{s,1}, \dots, N_{s,i}, \dots, N_{s,k} \mid N_{s,1} > \dots > N_{s,i} > \dots > N_{s,k}\}$ **/*

$[N_{p,1}, \dots, N_{p,i}, \dots, N_{p,k}] \leftarrow 1;$

$q \leftarrow \lfloor \frac{N(\Upsilon)}{N_{s,1} + \dots + N_{s,i} + \dots + N_{s,k}} \rfloor;$

for $j \leftarrow 1 : k$

if $(f_N(\Upsilon) - q \sum_{i=1}^k N_{s,i}) < \sum_{i=1}^j N_{p,i} \cdot N_{s,i}$

$N_{p,j} \leftarrow 0;$

end if

end for

$[N_{p,1}, \dots, N_{p,i}, \dots, N_{p,k}] \leftarrow [N_{p,1}, \dots, N_{p,i}, \dots, N_{p,k}] + q;$

return $[N_{p,1}, \dots, N_{p,i}, \dots, N_{p,k}]$

Fig. 8. Different terminals provide different supply voltages

G. Capacity vs. Operation-time

Large-scale battery cells, e.g., for EVs and HEVs, are packed in such a way that n_s battery cells are connected in series, providing the required supply voltage, and n_p parallel groups are connected in parallel, determining flows of the current (I), resulting in the required capacity. The capacity, because of the nonlinearities of batteries, cannot be derived simply by the ideal battery capacity equation:

$$C = T \cdot I, \quad (6)$$

where T is the discharge time (the battery operation-time). Instead, empirical Peukert's relation [11] models nonlinearities for the case of a constant current load by introducing an empirical parameter as: $C = T \cdot I^\alpha$, where $\alpha > 1$ is called Peukert's value, which typically ranges between 1.2 and 1.4 [11].

For our purpose, we model this nonlinearity using discretization of a flow of the current, similar to the approach followed by Rakhmotov and Vrudhula in [17]. That is, real-world systems are characterized by loads that are variable over time. Such variable loads are approximated by piece-wise constant loads, represented by a set of M current levels (i_1, \dots, i_M), in which M is used to characterize the load and is determined by the quantization interval, $\Delta t (= t_i - t_{i+1})$ which is a fraction of the total operation time, T . That is, $I_i(t) = \sum_{i=1}^M i_i \cdot 1_{[t_{i-1}, t_i)}(t)$, where $1_A(t)$ is an indicator function. So, the smaller the Δt , the higher the accuracy in the characterization of the load. In the case where $\Delta t = T$, the load is constant. The patterns of the load can be obtained via empirical measurements, resulting in a discharge profile for a battery cell or a pack of battery cells. Thus, the model of Eq. (6) generalizes to

$$C = T \cdot I_i(t). \quad (7)$$

The total load is the sum of the current that is loaded from individual parallel groups, i.e., $I = I_1 + \dots + I_i + \dots + I_{n_p}$ and it is uniformly distributed at some point in time within a certain acceptable threshold in discrepancy, leading to $I = n_p \cdot I_i$. This

results in

$$C = T \cdot \frac{I(t)}{n_p}. \quad (8)$$

When cell failures occur, the number of available parallel groups equals $n_p - N(t)$, where $N(t)$ is the total number of failures occurred in the battery-cell array by time t . In the ABS, the number of available parallel groups is defined as

$$n_A = n_p - \lceil \frac{N(t)}{n_s} \rceil, \quad 0 \leq N(t) \leq n_p \cdot n_s. \quad (9)$$

Since the numbers of these failures that occur in disjoint time intervals are independent, $N(t)$ is Poisson distributed with a battery-cell failure rate, λ . So, the average total number of cell failures that occur by time t is proportional to t , resulting in $\lambda \cdot t$. This equation is applied to Eq. (9), yielding:

$$C = T \cdot \frac{I(t)}{n_A} = T \cdot \frac{I(t)}{n_p - \lceil \frac{\lambda \cdot t}{n_s} \rceil}, \quad (10)$$

where $t < \frac{n_p \cdot n_s}{\lambda} \leq T$.

On the other hand, in a legacy scheme, the load for a series chain of operational battery cells increases in proportion to the total number of cell failures across the n_p parallel groups as:

$$n_L = \begin{cases} n_p - N(t), & 0 \leq N(t) \leq n_p \\ 0, & n_p < N(t) \leq n_p \cdot n_s. \end{cases} \quad (11)$$

Clearly, the linear increase in the load is due to the fact that it fails to reuse any healthy battery cells in the series chain containing a faulty cell. So, the available capacity in following the legacy scheme is calculated by

$$C = T \cdot \frac{I(t)}{n_L} = T \cdot \frac{I(t)}{n_p - \lambda \cdot t}, \quad (12)$$

where $t < \frac{n_p}{\lambda} \leq T$. Therefore, the higher λ , the more operation-time gain over the legacy scheme; it is also inversely proportional to the number of battery cells in series, n_s .

Either of the two policies described above is applied, based on the configuration of the battery-cell connectivity. To maximize battery-cell utilization, the capacity of power that the entire battery cells deliver is selected as a criterion to compare the two policies. If an $m \times n$ matrix represents a combination of n_s battery cells in a series chain and there are n_p parallel groups, any element of battery cells in the matrix is assumed to become faulty independently of others. For instance, when one battery cell fails, $(n_s - 1) \cdot n_p$ of power is provided, based on the dynamic-voltage allowing policy—for simplicity, it is assumed the element of each battery cell is capable of 1 volt and 1 ampere, while $n_s \cdot (n_p - 1)$ of power based on the constant-voltage keeping policy. So, the break-even point in selecting the policy is found when $n_s = n_p$. When more than one battery cell fail, the number of battery cells left unused due to the faulty-cell detouring reflects a measure of the capacity. In other words, a ratio (r) of the number of columns (c) to the number of rows (w) counted on faulty cells in the matrix can be a factor in the decision to make, comparing with the total size of the matrix. So, the break-

even point is determined by

$$r = \frac{c}{w} = \frac{n_p}{n_s}, \quad (13)$$

and hence, when $r > \frac{n_p}{n_s}$, the dynamic-voltage allowing policy is chosen, providing more capacity of power than the constant-voltage keeping policy.

IV. EVALUATION

This section first describes our evaluation methodology and then evaluate the performance of our framework in comparison with a legacy scheme that cannot configure the battery-cell connectivity online.

A. Evaluation Methodology

The metrics used for evaluation of battery performance include the battery operation-time and the supply voltage. The operation-time is proportional to the total capacity of the battery cells/packs, while the supply voltage determines the deliverable power.

We simulate the battery dynamics using Dualfoil [10] which is widely used for designing multiple battery systems [6, 18]. Using Dualfoil is sufficient to demonstrate the way the battery connectivity is dynamically reconfigured.

B. Evaluation Result

The reconfiguration framework effectively “masks” the effects of a battery-cell failure, thus extending the battery operation-time, while the legacy scheme significantly suffers battery-capacity loss and hence reduces the operation-time. The battery operation-time is computed with the maximum deliverable power and the amount of current constantly drawn from the battery pack. Obviously, the more the battery-cell failures, the higher the reduction in the battery operation-time. Fig. 9 shows the results of comparing the battery operation-times. Clearly, the legacy scheme loses a significant amount of span as the number of faulty battery cells increases. The reason for this is that the failure of one battery cell results in the loss of the series chain including the faulty battery cell. By contrast, the reconfiguration framework reuses the remaining healthy battery cells in the series chain as backup cells. So, despite additional battery-cell failure in other chains, they are replaced with surviving healthy battery cells. Fig. 9-(a) shows the fault-tolerance capability of the proposed reconfiguration framework. For instance, when $\lambda \cdot t = 6$ through 9 and $\lambda \cdot t = 12$ through 15, the battery-pack’s operation-time remains constant irrespective of an increase in number of battery-cell failures. The difference in operation-time between the two mechanisms gets larger as the frequency of battery-cell failures gets higher. As can be seen in Fig. 9-(b), the operation-time gain achieved by the reconfiguration framework grows substantially with an increase in number of battery cells in a series chain (n_s) in each parallel group, thus enhancing the availability of backup battery cells. This is effective even for the case of connecting two battery cells in series (i.e., $n_s = 2$), achieving a factor of 5 gain. Clearly, the more the battery cells in series, the larger the gain.

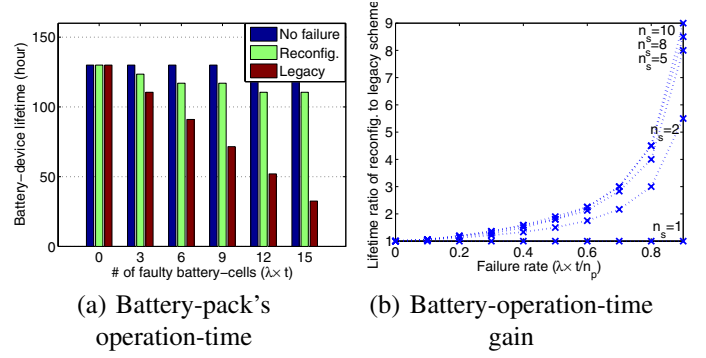


Fig. 9. Comparison of battery-pack’s operation-time and gain with battery-cell failures: The battery pack is formed, unless specified otherwise, by 5 battery cells in a series chain ($n_s = 5$) with 20-battery parallel groups ($n_p = 20$). The constant current of 200 mA is assumed drawn from the battery pack, and the capacity of each battery cell is 1.3 AH.

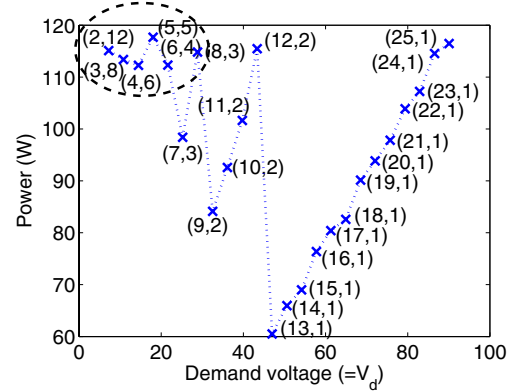


Fig. 10. Changes in demand voltage and presentation of the corresponding power: A pair corresponds to (n_s , n_p).

The dynamic-voltage-allowing policy aims to meet the demand of wide-ranging supply voltages from different applications while keeping deliverable power maximum. Fig. 10 shows changes in the demand voltage and the corresponding maximum deliverable power resulting from a 25-battery-cell pack that is based on the configuration of setting actual supply voltage and capacity of each battery cell to 3.6 Volts and 1.3 AH, respectively, with jitter of 2.5% allowed. So, maximum deliverable power is bounded by between an estimated 114 W and 120 W. This power can be delivered in a combination of 5 parallel groups and 5 battery cells of a series chain in each group (i.e., (5, 5)), or one parallel group with 25 battery cells in series (i.e., (25, 1)). Interestingly, a good range of supply voltages, corresponding to the group circled in Fig. 10, is provided while keeping maximum deliverable power reasonably constant. This implies that appropriately tuning the battery connection can improve the utilization of battery cells while meeting the demand of the underlying applications. In the meantime, the connectivity of (9, 2) or (13, 1) appears inefficient with respect to the utilization of battery cells. However, failure of any battery cell or a voltage drop can be resolved by virtually replacing them with backup battery cells, thereby maintaining the required voltage level.

The dynamic-voltage-allowing and constant-voltage-keeping policies are devised for different purposes: the former aims to meet the demand of wide-ranging supply




Fig. 11. Comparison of dynamic-voltage-allowing and constant-voltage-keeping policies with maximum deliverable power

voltages, while the latter is to sustain an acceptable range of supply voltage against battery failures or a possible voltage drop during the battery operation-time, both with the deliverable power kept maximum. So, the two policies can be compared with respect to the deliverable power. Fig. 11 shows the distribution of power magnitudes between the constant-voltage-keeping and the dynamic-voltage-allowing policies. In battery connectivity, when $n_s > n_p$, the dynamic-voltage-allowing policy is effective in supplying the maximum deliverable power, while when $n_p > n_s$, the constant-voltage-keeping policy is a better choice. The reason for this lies in the utilization of unused battery cells/packs. Obviously, the break-even point occurs when $n_s = n_p$.

As mentioned earlier, since the voltage drop is unavoidable, the constant-voltage-keeping policy is applied to keep the supply voltage above or equal to the demand voltage while the supply voltage is being monitored. The monitoring interval (Δt) is directly associated with a degree to which the system may suffer due to the voltage drop below the demand. Obviously, the higher the frequency of monitoring, the shorter the time an application suffers, but the higher the overhead of monitoring. Fig. 12-(a) shows changes in supply voltage with two different discharge rates during the operation-time of a 700-battery-cell pack. It is assumed that each battery cell is discharged independently, following the distribution of discharging a Lithium-ion battery that is simulated with the configuration of providing output voltage of 4.3 volts and nominal capacity of 1.3 AH. Demand voltage (V_d) for an application is assumed to be 600 volts. In the case where the battery pack is discharged at C rate (in Fig. 12-(b)), when the battery pack is monitored every $\Delta t (= 10)$, it is detected at the 10-th time interval when the supply voltage drops below V_d , reconfiguring the battery pack connectivity into 4 parallel groups with 143 battery cells in a series chain, i.e., (143, 4), providing an estimated 604 volts. In the case of C2 rate (in Fig. 12-(c)), the underlying application suffers 5 times more battery-capacity loss than at the normal discharge rate. In particular, the more steeply does the supply voltage drop, the larger the difference between the supply and demand voltages. This case can be improved by reducing the monitoring interval ($\Delta t = 10$). As can be seen in Fig. 12-(d), with the monitoring interval halved ($\Delta t = 5$), on-time detection of the voltage drop is improved by 67%.

V. RELATED WORK

A synergetic battery pack (SBP) [8] is a simple battery charger designed to charge 4 battery cells connected in series, with accompanying control circuitry. In SBP, two switches at each battery cell are set boundary, making it connected or disconnected to power buses. Like this scheme, a battery switch array system [1, 22] presents a flexible and configurable architecture for arranging micro-scale battery cell, allowing faulty battery cells to be detoured. Unlike our proposed framework, this system, similar to a reconfigurable multi-cell battery [6], fails to realize in practice the dynamic selection and switching of battery cells, especially based on battery characteristics such as state-of-charge (SOC), state-of-health, and load requirements.

Accurate SOC estimation is important; the voltage or current may not be an exclusive indicator due to the unpredictability of both battery and user behaviors. Most existing adaptive SOC systems are based on well-known estimation techniques, such as Kalman filter [14–16], a fuzzy logic model [20] and artificial neural network model [4]. As for estimation, a book-keeping system [3], based on coulomb counting, takes into account the self-discharge rate, temperature, capacity efficiency, and cycle life to improve SOC accuracy.

Appropriate steering of discharge current helps extend battery operation-time. Benini *et al.* [2] found that virtual parallel discharge is more effective for maximizing battery operation-time than sequential discharge. Chiasserini and Rao [5] considered the recovery effect in discharge scheduling.

VI. CONCLUSION

We now discuss some of the remaining and make concluding remarks.

A. Discussion

There are several remaining issues. First, in the architecture of our framework, one controller monitors and controls battery cells via its own control units. However, we do not exclude the possibility of the controller failure which has as critical an impact on the battery-management system as the battery-cell failure. Thus, the quantity of battery cells for which a single controller is responsible could be a criterion of assessing design risk. In addition to the controller, other units, such as controllable switches, battery-state-monitoring sensors, wires and others, are also fault-prone. Clearly, the higher the quality of these units, the more fault-tolerant the management system, but the higher manufacturing cost.

Second, energy dissipation across circuit units is inevitable in practice. In particular, DC-DC converters (i.e., down-conversion) are a major part. The converter's efficiency is, in general, represented by a ratio of power delivered by the converter to power delivered to the converter. Despite a great deal of engineering effort made thus far, the closer to the output voltage level the input voltage gets, the higher the efficiency, assuming the same amount of current flow in and out of the converter. In this sense, the dynamic-voltage-allowing policy supports overall high energy efficiency. This is part of our future inquiry.

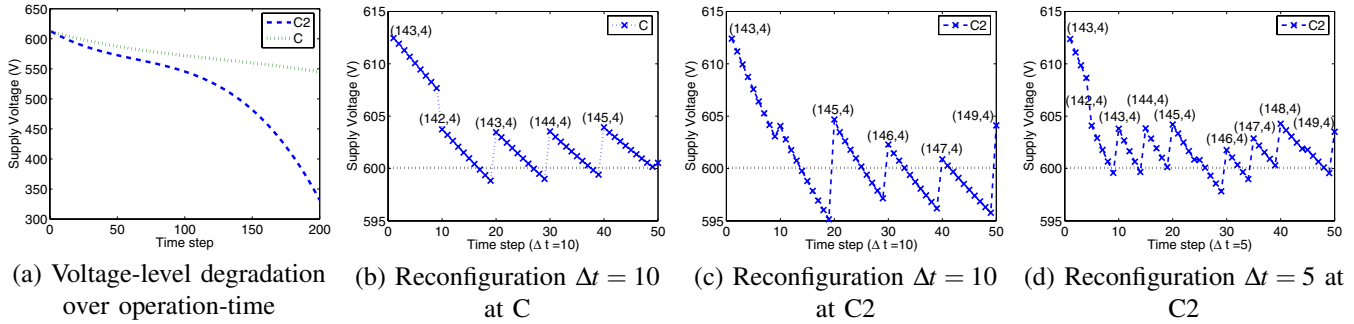


Fig. 12. Dynamic reconfiguration subject to a demand voltage ($V_d = 600$) with respect to different discharge rates: A pair corresponds to (n_s, n_p) .

Last, efficient charge and discharge of large-scale battery packs are equally important in extending their operation-time, along with the dynamic reconfiguration framework; capacity efficiency depends on the degree to which batteries are discharged. Developing efficient scheduling algorithms on charge/discharge [2, 5] is an important area of research.

B. Conclusive Remarks

Large-scale battery-management systems must cope with the battery-cell failure, especially for electric vehicles. In this paper, we have presented a novel dynamic reconfiguration framework to enhance the dependability of battery-management systems as well as the flexibility in reconfiguration of battery cell connectivity. We began by designing the architecture of the dynamic reconfiguration framework and proposed two bypassing mechanisms. We then analyzed the performance of our framework, in comparison with a legacy scheme in terms of the battery operation-time and the maximum deliverable power. The experimental results show that the proposed framework allows large-scale battery packs to be more fault-tolerant by a factor of 9. We addressed two challenges: (1) The online reconfiguration of both cell-level and pack-level battery-cell connectivity works in tandem, thereby achieving the capability of providing fine-grained supply voltage and the high scalability in extending to larger battery-management systems; (2) The power source from the battery packs is customized online to meet wide-ranging demands from multiple applications simultaneously with their requirements incrementally met. In summary, our framework offers a unique solution to the vulnerability of large-scale battery-management systems to the battery-cell failure, and greatly improves the scalability to the extension and customizability of the power source.

REFERENCES

- [1] Mahmoud Alahmad, Herbert Hess, Mohammad Mojarradi, William West, and Jay Whitacre. Battery switch array system with application for jpl's rechargeable micro-scale batteries. *J. of Power Sources*, 177(2):566–578, 2008.
- [2] Luca Benini, Davide Bruni, Alberto Macii, Enrico Macii, and Massimo Poncino. Discharge current steering for battery lifetime optimization. *Trans. on Com.*, 52(8):985–995, Aug. 2003.
- [3] Henk Jan Bergveld, Wanda S. Kruijt, and Peter H.L. Notten. *Battery Management Systems: Design by modelling*. ISBN 1-4020-0832-5. Kluwer Academic Publishers, 2002.
- [4] C.C. Chan, E.W.C. Lo, and Shen Weixiang. The available capacity computation model based on artificial neural network for lead-acid batteries in electric vehicles. *J. of Power Sources*, 87(1-2):201–204, 2000.
- [5] Carla-Fabiana Chiasserini and Ramesh R. Rao. Energy efficient battery management. *JSAC*, 19(7):1235–1245, July 2001.
- [6] Song Ci, Jiucui Zhang, Hamid Sharif, and Mahmoud Alahmad. A novel design of adaptive reconfigurable multicell battery for power-aware embedded network sensing systems. In *Globecom*, pages 1043–1047, Washington, DC, USA, 2007. IEEE.
- [7] Green Car Congress. International study shows global gains in consideration of hybrid and electric vehicles. <http://www.greencarcongress.com/2008/06/international-s.html>.
- [8] Amanda Davis, Ziyad M. Salameh, and Stephen S. Eaves. Evaluation of lithium-ion synergetic battery pack as battery charger. *Trans. on Energy Conversion*, 14(3):830–835, 1999.
- [9] Power Designer. Dc-dc converter basics. <http://www.powerdesigners.com>.
- [10] Marc Doyle, Thomas F. Fuller, and John Newman. Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell. *J. of Power Sources*, 140(6):1526–1533, 2003.
- [11] D. Linden and T.B. Reddy. *Handbook of Batteries*. ISBN 978-0-07-135978-8. McGraw-Hill, 3rd edition, 2002.
- [12] Maxim. Source resistance: The efficiency killer in DC-DC converter circuits. http://www.maxim-ic.com/appnotes.cfm/an_pk/3166, 2004.
- [13] mPower. Why batteries fail. http://www.mpoweruk.com/failure_modes.htm.
- [14] Gregory L. Plett. Extended kalman filtering for battery management systems of lipb-based hev battery packs part 1. background. *J. of Power Sources*, 134(2):252–261, 2004.
- [15] Gregory L. Plett. Extended kalman filtering for battery management systems of lipb-based hev battery packs part 2. modeling and identification. *J. of Power Sources*, 134(2):262–276, 2004.
- [16] Gregory L. Plett. Extended kalman filtering for battery management systems of lipb-based hev battery packs part 3. state and parameter estimation. *J. of Power Sources*, 134(2):277–292, 2004.
- [17] Daler N. Rakhmatov and Sarma B. K. Vrudhula. An analytical high-level battery model for use in energy management of portable electronic systems. In *ICCAD*, pages 488–493, Piscataway, NJ, USA, 2001. IEEE.
- [18] Ravishankar Rao, Sarma Vrudhula, and Daler Rakhmatov. Analysis of discharge techniques for multiple battery systems. In *ISLPED*, pages 44–47. ACM, Aug. 2003.
- [19] Ravishankar Rao, Sarma Vrudhula, and Daler N. Rakhmatov. Battery modeling for energy-aware system design. *Computer*, 36(12):77–87, 2003.
- [20] Pritpal Singh, Craig Fennie Jr., and David Reisner. Fuzzy logic modelling of state-of-charge and available capacity of nickel/metal hydride batteries. *J. of Power Sources*, 136(2):322–333, 2004.
- [21] Thomas Stuart, Fang Fang, Xiaopeng Wang, Cyrus Ashitiani, and Ahmad Pesaran. A modular battery management system for HEVs. In *SAE*, pages 1–9. SAE International, June 2002.
- [22] Vinesh Sukumar, Mahmoud Alahmad, Kevin Buck, Herbert Hess, Harry Li, Dave Cox, Fadi Nessir Zghoul, Jeremy Jackson, Stephen Terry, Ben Blalock, M.M. Mojarradi, W.C. West, and J.F. Whitacre. Switch array system for thin film lithium microbatteries. *J. of Power Sources*, 136(2):401–407, 2004.