Battery-Driven System Design: A New Frontier in Low Power Design*

†Kanishka Lahiri ‡Anand Raghunathan †Sujit Dey †Debashis Panigrahi

†Dept. of ECE, UC San Diego, La Jolla, CA
‡C & C Research Labs, NEC USA, Princeton, NJ

Abstract— **As an increasing number of electronic systems are powered by batteries, battery life becomes a primary design consideration. Maximizing battery life requires system designers to develop an understanding of the capabilities and limitations of the batteries that power such systems, and to incorporate battery considerations into the system design process. Recent research has shown that, the amount of energy that can be supplied by a given battery varies significantly, depending on how the energy is drawn. Consequently, researchers are attempting to develop new battery-driven approaches to system design, which deliver battery life improvements over and beyond what can be achieved through conventional low-power design techniques. This paper presents an introduction to this emerging area, surveys promising technologies that have been developed for battery modeling and battery-efficient system design, and outlines emerging industry standards for smart battery systems.**

I. INTRODUCTION

Battery powered electronic systems, and the integrated circuits within them, account for a large and rapidly growing revenue segment for the computer, electronics, and semiconductor industries. For instance, the revenue from wireless voice/data handsets is expected to exceed that from PCs in the near future, and the use of wireless Internet access is expected to overtake fixed Internet access in the next few years. For battery powered systems, the battery life directly impacts the system's utility, and the duration and extent of its mobility. The battery life of a system is determined by the capacity of the energy source (*i.e.*, battery), and the energy drawn by the rest of the system.

Improvements in semiconductor fabrication and wireless communication technologies promise to enable advances in ubiquitous information access and manipulation (anytime, anywhere computing and communications). Unfortunately, projected improvements in the capacity of batteries (5-10% CAGR) are much slower than what is needed to support the increasing complexity, functionality and performance of the systems they power. Figure 1 illustrates a widening "battery gap", between trends in processor power consumption [1], and improvements in battery capacity [2]. Bridging this gap is a challenge that system designers must face for the foreseeable future.

The need to improve battery life has, in large part, driven the research and development of low power design techniques for electronic circuits and systems [3], [4], [5], [6]. Low power design techniques are successful in reducing the energy that is drawn from the battery, and hence improve battery life to some extent. However, truly maximizing battery life requires an understanding of both the source of energy and the system that consumes it. It has been shown that, the amount of energy that can be supplied by a given battery varies significantly, depending on how the energy is drawn [7]— [23].

Battery-driven system design, which refers to the process of designing a system with careful consideration of the battery and its characteristics, promises to provide further improvements in

* This work was supported in part by the National Science Foundation, under NSF grant 99-12414

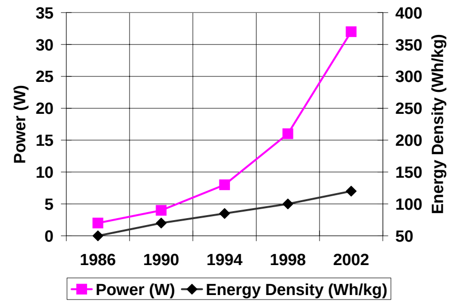

Fig. 1. A widening "battery gap", due to rapidly increasing power requirements and slowly improving battery technology.

battery life beyond what can be achieved by conventional low power design technologies. This paper aims to introduce and summarize the relatively new field of battery-driven system design. We present a brief review of contemporary battery technologies, and highlight the characteristics of batteries that are relevant to designers of battery powered electronic systems. We present an overview of techniques that have been developed for battery modeling and analysis, as well as techniques for the design of battery friendly system architectures. We also summmarize emerging industry standards for implementing smart battery systems.

II. BACKGROUND

In this section, we first provide an overview of battery technologies that are commonly used to power portable electronic appliances. Next, we briefly describe the principles of operation of a battery, and go on to explain the important characteristics of batteries that need to be considered for the design of battery-efficient systems.

A. Overview of Battery Technologies

In this sub-section, we describe battery technologies that have been developed over the last two decades to meet the increasing demand for smaller, lighter, higher capacity re-chargeable batteries for portable appliances. When comparing different battery technologies, several considerations arise. These include energy density (charge stored per unit weight of the battery), cycle life (the number of discharge/charge cycles prior to battery disposal), environmental impact, safety, cost, available supply voltage, and charge/discharge characteristics. Figure 2 illustrates the development of re-chargeable battery technologies, and compares them in terms of typical energy density, based on data from [2], [24]. The most popular re-chargeable battery technologies for portable electronic appliances include:

- **Nickel Cadmium:** This is a mature technology, and has been successfully used for several decades to develop re-chargeable batteries for portable electronic devices. Its advantages include low cost, and high discharge rates. While Ni-Cd technology has been losing ground in recent years owing to its low energy density and toxicity, it is still used in low cost applications like portable radios, CD/tape players, *etc*.

- **Nickel Metal-Hydride:** These batteries have been in widespread use in the recent years for powering laptop computers. They have roughly twice the energy density of Ni-Cd batteries. However, they have shorter cycle life, are more expensive, and are inefficient at high rates of discharge.
- **Lithium Ion:** This is the fastest growing battery technology today, with significantly higher energy densities, and cycle life about twice that of Ni-MH batteries. Lithium ion batteries are more sensitive to characteristics of the discharge current, are more expensive than Ni-MH batteries, and can be unsafe when improperly used. However, longer lifetimes have made them the most popular battery choice for notebooks, PDAs, and cellular phones.
- **Reusable Alkaline:** While disposable alkaline batteries have been used for many years, reusable alkaline manganese technology has developed as a low cost alternative in which energy density and cycle life are compromised. While the initial energy density of reusable alkaline batteries is higher than Ni-Cd, it has been found to decrease rapidly with cycle life. For instance, after 10 cycles, a 50% reduction, and after 50 cycles, a 75% reduction in energy density is commonly observed [25].
- **Lithium Polymer:** This emerging technology enables ultra thin batteries (less than 1 mm thickness), and is expected to suit the needs of light-weight next-generation portable computing and communication devices. Additionally, they are expected to improve over current lithium ion technology in terms of energy density and safety. However, these batteries are currently expensive to manufacture, and face challenges in internal thermal management [17].

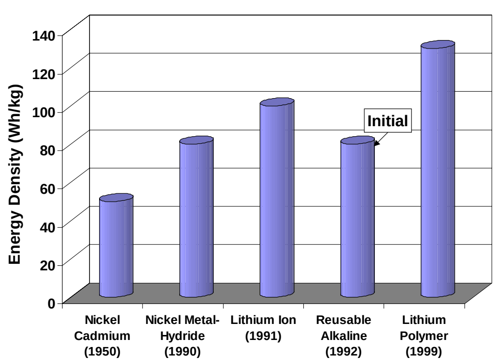

Fig. 2. Energy density and year of first commercial deployment for different battery technologies.

Next, we briefly describe some concepts and terminology associated with battery discharge.

B. Principles of Battery Discharge

The basic components of a battery are shown in Figure 3, through the example of a lithium/thionyl chloride cell. The battery consists of an anode (lithium), a cathode (carbon), and an electrolyte. The electrolyte separates the two electrodes and provides a mechanism for the transfer of charge between them. During battery discharge, oxidation at the anode (Li) results in the generation of (i) electrons, which flow through the external circuit, and (ii) positively charged ions (Li^+), which by diffusion, move through the electrolyte ($SOCl_2$) towards the cathode. Reduction reactions occur at available reaction sites in the cathode, generating negatively charged ions (Cl^-), which combine with the positive ions (Li^+) to generate an insoluble compound ($LiCl$) that gets deposited on the cathode. Sites where the compound is deposited become inactive, making them unavailable for further use. As discharge proceeds, more and more reaction sites are made unavailable, eventually leading to a state of complete discharge.

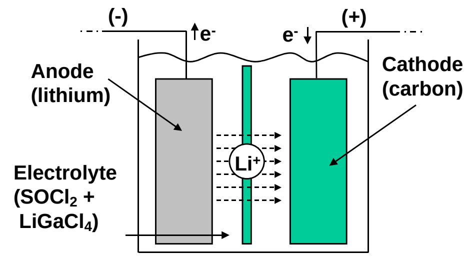

Fig. 3. Basic structure of a lithium/thionyl chloride battery.

A battery is characterized by the open-circuit voltage (V_{OC}), *i.e.*, the initial potential of a fully charged battery under no-load conditions, and the cut-off voltage (V_{cut}) at which the battery is considered discharged. Battery capacity can be expressed in three ways. The *theoretical capacity* of a battery is based on the amount of energy stored in the battery, and is an upper bound on the total energy that can be extracted in practice. The *standard capacity* of a battery is the energy that can be extracted when it is discharged under standard load conditions, which are specified by the manufacturer. For example, a typical lithium ion battery may have standard capacity $500\ mAh$ when discharged at a constant current of $125\ mA$, at $25°C$. The *actual capacity* of a battery is the amount of energy that the battery delivers under a given load, and is usually used (along with battery life) as a metric to judge the battery efficiency of the load system. A battery-efficient system is one where the discharge profile characteristics result in improved actual capacity. While the actual capacity may exceed the standard capacity, it cannot exceed the theoretical capacity of the battery. In this paper, we use the term battery capacity to refer to the actual capacity of the battery.

Next, we present some characteristics of batteries whose understanding is crucial for battery-driven system design.

C. Impact of Discharge Characteristics on Battery Capacity

Two important effects that make battery performance sensitive to the profile of the discharge current are (i) *rate capacity effects*, which are due to a dependency between the actual capacity of a battery and the magnitude of the discharge current, and (ii) *recovery effects*, which are due to recovery of charge during idle periods (when no charge is drawn). Both these phenomena can significantly affect the battery capacity and the lifetime of a battery. We next provide a short description of the electrochemical phenomena responsible for these effects.

C.1 Rate Capacity Effects

The lifetime of a battery depends on the availability and reachability of active reaction sites in the cathode. During periods of discharge where the rate of discharge is low (current drawn is small), the distribution of inactive reaction sites throughout the material of the cathode is more or less uniform. However, if the rate of discharge is high (large current is drawn), reductions occur only at the outer surface of the cathode. This results in the surface of the cathode being coated with an insoluble compound, preventing access to many active internal reaction sites. Consequently, the battery is declared discharged even though many active cathode sites remain un-utilized, effectively decreasing the total capacity of the battery. The effect of this phenomenon is a dependency between battery capacity and the rate at which it is discharged [18].

C.2 Recovery Effects

Besides availability of active reaction sites in the cathode, the availability of charged ions (Li^+) is also a factor that determines the amount of energy that can be delivered by a battery [19]. When no current is drawn, the concentration of positively charged ions (Li^+) is uniform at the electrode-electrolyte interface. When current is drawn from the battery, positively charged ions are consumed at the cathode-electrolyte interface, and replaced by new ions that diffuse from the anode through the electrolyte. When the current drawn is sufficiently large, the rate of diffusion fails to keep up with the rate at which ions are consumed at the cathode. As a result, the concentration of positively charged ions decreases near the cathode and increases near the anode, degrading the battery's output voltage. However, if the battery is allowed to idle for a period of time, the concentration gradient decreases (due to diffusion), and charge recovery takes place at the cathode. As a result, the capacity and lifetime of the battery increase.

To summarize, the amount of energy that can be extracted from a battery, and consequently its lifetime, are sensitive to (i) the magnitude of the discharge current, and (ii) the presence of idle times in the discharge profile. We next turn our attention to battery models that have been developed to capture such battery characteristics.

III. BATTERY MODELING

Battery models capture the characteristics of real-life batteries, and can be used to predict their behavior under various conditions of charge/discharge. Battery models are useful tools for a battery-driven system design approach, because they enable analysis of the discharge behavior of the battery under different design choices (*e.g.,* system architectures, power management policies), without resorting to time consuming (and expensive) prototyping and measurement for each alternative [26].

In the following discussion, we consider several different battery models. We classify these approaches into (a) analytical models, (b) electrical circuit models (c) stochastic models, and (d) electro-chemical models. Additionally, there are several other important factors that characterize a battery model, each of which may also serve as a basis for classification. For example, battery models could be classified on the basis of (i) the types of load current supported (*e.g.,* constant load *vs.* variable load), (ii) supported battery technologies (Lithium-ion, Ni-Cd *etc*), (iii) the set of battery effects captured by the model (*e.g.,* rate capacity effects, recovery effects, thermal effects), (iv) computational efficiency, and (v) accuracy in predicting real-life battery behavior.

A. Analytical models

Several battery models have been developed where analytical expressions are formulated to calculate actual battery capacity and lifetime using discharge current values, operating environment characteristics, and physical properties of the battery as parameters [7], [8], [9]. One of the simplest and earliest analytical models is Peukert's formula, which expresses the non-linear relationship between the battery capacity and the rate of discharge [27]. Peukert's formula states that the actual capacity Q of a battery is given by:

$$Q = \frac{k}{I^\alpha}$$

where I is the load current (assumed to be constant), k is a constant capturing electro-chemical properties, physical construction, and operating environment of the battery, and α is a constant that captures the rate capacity effect.

Recently, sophisticated analytical models have been proposed, where the capacity is expressed in terms of more complex functions of discharge current, temperature, and several constants. Various techniques are used to determine the exact form of the function. For example, regression analysis is used in [7] to estimate coefficients in an expression based on a Weibull failure model. This is a constant load model that captures rate capacity and thermal effects. The model described in [8] is a variable load model that uses empirical equations to capture the rate-capacity effect. It can be used to predict variation in capacity for current discharge profiles that follow different probability density functions. A variable load model described in [9] uses laws of chemical kinetics to derive the mathematical expression, and uses statistical techniques to estimate the parameters.

In summary, analytical battery models can include both constant-load and variable-load models. All of these models capture rate-capacity effects, some (*e.g.,* [7]) capture thermal effects, though none address recovery due to relaxation during idle periods. These models are flexible, and can be easily configured for specific batteries. They are computationally efficient, requiring evaluation of simple analytical expressions.

B. Electrical Circuit Models

We next consider a class of techniques that model battery discharge using an equivalent electrical circuit [10], [11], [12], [13]. Most of these approaches are based on constructing a SPICE model of a coupled network to represent the battery. In [11], the charge stored in the battery is modeled using a capacitor, while voltage across the capacitor is used to represent the output voltage. The discharge process is modeled by continuously applying corrections to both the charge stored in the capacitor, and the output voltage. For example, to model the rate-capacity effect, a voltage source is used to subtract from the voltage across the capacitor. The discharge current indexes into a lookup table to control this voltage source. Further corrections account for (i) the cycle life of the battery (the number of times the battery has been charged and discharged), (ii) internal resistance (which may vary with time), and (iii) thermal characteristics.

The first application of such circuits for modeling alkaline, zinc-air, and Ni-Cd batteries was described in [10]. Modeling of Ni-MH batteries using similar techniques is described in [12], while lithium-ion batteries were modeled using this approach in [11]. A discrete-time model of a lithium-ion battery is presented in [13], which approximates the analog circuit model of [11] while improving computational efficiency.

Electrical circuit models are variable-load models that are capable of modeling rate-capacity and thermal effects. However, none of the known circuit models account for recovery effects. This class of battery models lends itself to easy simulation using existing circuit simulators, and has been used in practice to analyze many common battery technologies.

C. Stochastic models

A stochastic model of a battery is described in [14], [15], where the battery is represented by a finite number of charge units, and the discharge behavior of the battery is modeled using a discrete-time transient stochastic process. As the stochastic process evolves over time (which is divided into a sequence of equally sized slots), the state of the battery is tracked by the number of remaining units of charge. In each time slot, the average discharge current is measured and used to determine the number of charge units consumed. If this average is non-zero, the number of charged units drained is obtained from a look-up table (or a plot) that contains rate capacity data. This lookup is similar to the approach described in [16]. However, if the slot is an idle slot, (*i.e.*, no current is drawn), then the battery recovers a certain number of charge units. The exact number of charge units recovered is modeled using a decreasing exponential probability density function, which is based on (i) the state-of-charge of the battery, and (ii) coefficients, which depend on the specific battery as well as the discharge characteristics. Over time, the battery steps through a sequence of states, from a state of full charge to either (i) a state where the cut-off voltage is reached, or (ii) a state where the theoretical capacity is exhausted.

The stochastic model described above can account for variable loads, and currently has been used to model lithium-ion batteries. It accounts for both rate-capacity and recovery effects. However, it does not take into account thermal effects. Its accuracy stems from the fact that accurate battery specifications are used to construct rate capacity data, and from the detailed analysis performed for the exact formulation of the probability density functions. Moreover, its computation requirements are modest, enabling it to be employed in system-level simulations.

D. Electro-chemical Models

This includes the class of battery models that directly consider the electro-chemical processes, thermodynamic processes, and the physical construction, when modeling discharge of the battery [17]— [23]. These models are significantly more detailed (and hence more accurate) than any of the previously described approaches. For example, [18] and [19] describe a model where concentrated solution theory and porous electrode theory are used to construct a set of differential equations for a specific battery. The model has a large number of parameters, including electrode geometries, concentration of the electrolyte, diffusion coefficients, transfer coefficients, reaction rate constants, *etc*. Numerical techniques are used to solve the equations to predict battery capacity under different load conditions. Battery models that belong to this class are very closely tied to specific batteries. For example, [21] describes separate models for neutral electrolyte based and acidic electrolyte based lithium/thionyl chloride batteries. The model presented in [18] applies to lithium-ion batteries with a solid polymer electrolyte and composite insertion cathode.

In summary, electro-chemical models are capable of analyzing many discharge effects under variable loads, including rate-capacity effects, thermal effects, and recovery effects. They are considerably more detailed than the previously described models, and hence are the most accurate among those considered here. Unfortunately, they are also the most computationally intensive.

E. Summary

Based on the above discussion, we draw the following conclusions regarding the various battery models: In terms of flexibility, the electro-chemical models are the least flexible, *i.e.*, it is difficult to use them for modeling any given battery. On the other hand, configuring the circuit level models, analytical models and the stochastic models for different types of batteries is relatively easy. Additionally, the electro-chemical models being principally targeted to designers and manufacturers of batteries rather than systems, make use of many proprietary parameters which are typically unavailable to a system designer. In terms of accuracy and efficiency, the electro-chemical models are the most accurate, and also the most computation intensive. Analytical models are at the other extreme, being computationally efficient, but limited in the discharge effects that they model. Electrical circuit models can be simulated with high efficiency, ignoring the effects of recovery of charge during idle periods. The stochastic model can be efficiently used in a simulation framework and is capable of modeling rate capacity as well as recovery effects.

IV. BATTERY-DRIVEN SYSTEM DESIGN

In this section, we describe various approaches to the design of battery-efficient systems. The first class of techniques that we consider are based on optimizing the system architecture itself. These include battery-driven policies for frequency scaling, task scheduling, voltage scaling, and power management. The second class of techniques we consider addresses multi-battery systems through battery scheduling techniques aimed at optimizing the discharge of constituent batteries. Finally, we describe techniques that enable battery-efficient operation of portable appliances in a wireless network by making use of new network protocols and traffic shaping techniques.

A. Battery-efficient System Architectures

In this sub-section, we present approaches to battery-efficient design of HW/SW system architectures.

A.1 Frequency Scaling

While frequency scaling is a common approach to reducing the average power consumption, CPU frequency scaling for battery powered computers is examined in [28] in terms of its impact on battery life, system performance, and power consumption. A commonly used history-based policy (not optimized for batteries) for CPU frequency scaling is presented in [29]. This policy dynamically calculates the CPU frequency for the next time interval based on (i) run- and idle-time percentages of the previous time interval, and (ii) work left over, in case the CPU frequency was too slow in the previous time interval. A larger percentage of idle-time in the previous time interval results in decreasing the CPU clock frequency by a small constant for the next time interval, while not going below a lower bound. A larger percentage of run-time results in increasing CPU frequency by a small constant, while not going beyond the CPU's maximum clock frequency.

It is shown in [28] that by using the above history based frequency scaling policy, important factors like non-ideal battery behavior are neglected. Hence, the lower bound on CPU clock frequency is redefined to maximize a a metric combining battery capacity, performance and power. In [28], both analytical and table driven techniques are described for the computation of this lower bound.

In summary, frequency scaling approaches use information from a battery model to vary the clock frequency of system components dynamically at run time. Since they also use workload characteristics (run- and idle-time percentages), and models of system power and performance, these approaches can be used to ensure efficient use of the battery without significantly compromising system performance.

A.2 Battery-aware Task Scheduling

Battery-aware static scheduling algorithms have been developed for real-time systems that aim at improving the system current discharge profile [30]. In this approach, a static schedule is generated from a given task graph using standard scheduling algorithms. The schedule is then subjected to a series of post-processing stages, wherein transformations are applied to optimize the corresponding discharge current profile for the battery. The transformations are applied to the schedule in a two step manner. First, the initial schedule is transformed using "global shifting", which reduces peak current consumption, and increases flexibility in the schedule. This serves as a high-quality initial solution for the second step, which is a set of local transformations to the schedule. These local transformations attempt to change the position of scheduled events (*e.g.*, interchange, shift forward, shift backward) in order to minimize a cost function that captures delay and the actual energy drawn from the battery, based on the model described in [8].

This approach explicitly tailors the current discharge profile of the system to meet battery characteristics. However, its applicability is limited to systems that can be statically scheduled.

A.3 Supply Voltage Scaling

We next describe recent approaches that apply battery-driven techniques towards scaling the supply voltage of the system. In [8], an analytical technique is presented to optimally select V_{dd} (the supply voltage) to find the best trade off between battery capacity and performance. The objective is to find a value of V_{dd} that minimizes the "battery discharge-delay" product, which is given by the product of the actual charge drained from the battery and the delay of the circuit for a given task. To achieve this, the battery discharge-delay product is mathematically expressed in terms of V_{dd} and other known (or measurable) parameters. The analysis uses (i) an analytical model of the battery (described in Section III and in [8]), (ii) the distribution of the discharge current profile, and (iii) a model for CMOS circuit delay in terms of V_{dd} and V_{th}. Note that this approach does not attempt to modify the shape of the current discharge profile, either statically or dynamically. Rather, if the current discharge profile (or distribution) can be statically determined, this technique can be used to select a constant supply voltage to jointly optimize circuit delay and battery life.

The approach presented in [30] applies to statically scheduled real-time systems, and is based on re-allocating slack times for tasks to enable supply voltage scaling. In this technique, an initial static schedule, which satisfies all performance constraints, is subjected to a set of global and local transformations to (i) shape the current discharge profile with a knowledge of the battery rate capacity characteristics, and (ii) facilitate re-allocation of the amount of slack assigned to each task. Slack available for various tasks is then exploited by voltage scaling, with the objective of flattening the discharge profile while meeting all real time constraints.

Though the approaches of [8] and [30] are both static approaches, the important difference between the two is that the latter results in an explicit modification to the shape of the discharge current profile, since the voltages of different tasks may be scaled differently.

A.4 Dynamic Power Management

While system-level power management is in itself a well-researched area ([6], [31]), recent research has proposed new power management schemes that specifically target battery-efficiency rather than average power.

The battery-driven system-level power management technique described in [32] is based on a policy that controls the operation state of the system according to the state-of-charge of the battery. As a case study, a battery powered digital audio recorder is considered. The aim of the power management policy is to provide a graceful degradation in audio quality as the battery approaches a state of complete discharge, along with an extension in battery life. To achieve this objective, the dynamic power management policy used is as follows. As long as the battery output voltage is above a certain threshold, the recorder plays high quality music, which results in a high rate of battery discharge (low battery efficiency). When the battery output voltage falls below a threshold voltage, the policy forces the player to degrade the quality of music, causing the system to discharge the battery at a lower rate (higher battery efficiency).

By including feedback from the battery, this work extends standard dynamic power management techniques (*e.g.*, time outs), which depend only on the energy consuming components of the system and their workload. These power management policies exploit rate capacity effects to improve battery-efficiency of the system when the battery nears a state of complete discharge. This approach can be described as being "reactive", since the battery-driven policies come into play only when the state-of-charge of the battery falls below a certain threshold.

B. Battery Scheduling and Management

With increasing energy requirements in portable appliances, systems with multiple batteries are no longer uncommon. A typical approach taken in case of these appliances is to discharge each battery sequentially and completely. However, recent work has shown this policy may not be the best for battery life. In this sub-section, we discuss approaches towards efficient management of multi-battery systems. In particular, we describe contributions made towards battery scheduling: how to distribute the current demand of a system among a set of batteries. We consider three classes of battery scheduling techniques. The first class includes static scheduling techniques (which do not use any run-time information), while the second and third classes are dynamic scheduling techniques that use measurements of battery terminal voltage and discharge current, respectively.

B.1 Static Battery Scheduling

Static battery scheduling approaches do not make use of any run-time information from either (i) the system discharge profile, or (ii) the state-of-charge of the various batteries. These techniques include:

- Serial Scheduling [33]: In this approach, all the batteries are discharged in sequence — an approach currently adopted in most products that are powered by multiple batteries.
- Random Scheduling [34]: In this approach, at every discharge interval, a battery is chosen at random from the array of batteries to power the system for a fixed discharge interval.
- Round-Robin Scheduling [32], [33], [34]: In this approach, for each discharge interval, the battery to be used is chosen in a round-robin fashion.

Both random and round-robin scheduling result in an improvement in the lifetime of the system as compared to serial scheduling. That is because these approaches provide individual batteries the opportunity to recover lost capacity during idle periods (at times when other batteries are being used). Experiments reported in [34] indicate that the round-robin scheme performs better than the random scheme (when the current demand follows a Poisson distribution).

Both these approaches require a selector circuit to switch between batteries at some frequency F_{sw}. The value of F_{sw} directly relates to the discharge interval, the duration for which each battery is continuously discharged without idling. When $F_{sw} = 0$, the batteries are discharged in sequence. However, improvements in battery capacity are observed with increasing F_{sw}, till such time diminishing returns are obtained, due to (i) large time constants internal to the batteries, and (ii) energy consumed by the switching circuit [32].

B.2 Terminal Voltage Based Battery Scheduling

Several scheduling techniques have been suggested that make use of information regarding the state-of-charge of the battery, estimated from the terminal voltage of the battery. Modified round-robin approaches are described in [32], [34] that take into account the output voltage. In these approaches, a fixed round-robin scheme is used till the voltage of one or more batteries falls below some threshold voltage V_{th}. Batteries with $V_{out} < V_{th}$ are disconnected from the load (made "inactive"), which gives them a chance to recover some charge. Meanwhile, round-robin scheduling is applied among the remaining "active" batteries in the pack. Inactive batteries may recover enough charge to re-enter the set of active batteries. Finally, when no active batteries are left (*i.e.*, with $V_{out} > V_{th}$), the batteries are again discharged in round-robin fashion.

In [33], the discharge interval is not fixed (as in the previously described approaches), but is adapted to the state-of-charge. The policy used is as follows: the battery chosen for discharge is used continuously till the output voltage falls below a certain threshold, at which point, it is disconnected from the load. The next battery to be used is chosen based on the output voltages and the past idle times of the various candidate batteries.

B.3 Discharge Current Based Battery Scheduling

In [35], the authors present the design of a two-battery power supply based on heterogeneous batteries with different rate capacity characteristics. Consider a two-battery system with batteries named A and B. At low rates of discharge, battery A has higher capacity per unit weight than battery B. However, at high rates of discharge, battery B has higher capacity per unit weight than A. In other words, battery B lasts longer under high rates of discharge, while A lasts longer under low rates of discharge.

The following scheduling policy is used to maximize the lifetime of this two-battery supply: if the rate of discharge is less than a threshold current, battery A is connected to the system, else battery B is connected. The authors provide an analysis that allows a designer to select the size for each battery (by optimally partitioning a fixed weight) for a given distribution of the system current profile. Experiments (using the battery model described in [11]) show that using the two-battery supply with discharge current based scheduling significantly improves the lifetime of the system, compared to the case where a single large battery of type A or B is used.

In this sub-section, we described various scheduling techniques that have been proposed for extracting longer lifetimes from multi battery systems by taking advantage of the charge recovery phenomenon and rate capacity effects.

C. Battery-efficient Traffic Shaping and Routing

In a network of battery operated devices, the architecture of each constituent node is not the only factor that determines battery lifetime. In this case, an important role is played by the network protocols and communication traffic patterns in determining battery efficiency and lifetime. In this sub-section, we highlight some recent work in networking protocols and traffic shaping algorithms for wireless networks that aim at improving battery life of the mobile nodes.

A traffic shaping algorithm for battery-powered portable devices is described in [36]. The idea is to forcibly interrupt the discharge of the battery whenever the battery charge falls below a certain level, and allow it to recover lost capacity. During these interruptions, incoming packets awaiting processing are buffered. Discharge of the battery resumes when (i) the battery recovers sufficient charge to process at least one request, and (ii) there are one or more pending requests.

A battery-efficient routing protocol for ad hoc wireless networks is presented in [37]. In this protocol, the metric for choosing the best route between a source and destination is based on (i) the state-of-charge of the battery-powered nodes along the route, (ii) the transmission energy of each node on the route, and (iii) penalties due to rate capacity considerations for nodes that transmit at high power levels. The protocol selects routes with low impact on the battery capacity of the constituent nodes, hence allows recovery of charge to occur on inactive nodes.

V. SMART BATTERY SYSTEMS

In this section, we briefly describe the smart battery system (SBS) standard, which is an emerging industry standard relating to smart batteries and the systems that deploy them. Initially developed by Intel and Duracell [38], the SBS specifications are currently being developed by an industry consortium called the SBS Implementers Forum [39]. The SBS specification is currently supported by over 80 companies that include battery, semiconductor, software, and system manufacturers, and suppliers of battery electronics [39].

The aim of the SBS forum is to create open standards that enable systems to be aware of and better communicate with the batteries that power them, improve battery efficiency, and promote inter-operability between products from battery, software, semiconductor, and system vendors.

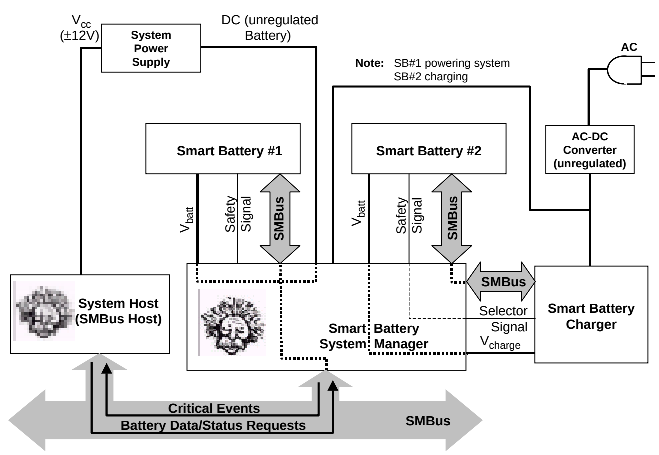

Fig. 4. A typical implementation of a smart battery system [39].

The SBS specifications include the following components

- **System Management Bus** (SMBus), which defines the protocols for the battery to communicate with other system components. Since the SMBus can also be used as a control bus for other low-speed system communications, the SMBus specifications are defined by an independent forum [40]
- **Battery Data Set**, which defines the information that is provided by a smart battery to the system host, smart charger, smart battery system manager, and other system components. Accuracy of the battery data set, which determines the accuracy of the battery information presented by the system to the user, is also covered by the specifications.
- **Smart Battery Charger**, which enables the charging characteristics to be controlled by the batteries themselves, in contrast to conventional chargers that have fixed charging characteristics hardwired for specific battery chemistries

and configurations. Smart battery chargers result in improvements in system safety, usable energy, charging time, and cycle life.

- **Smart Battery System Selector**, which is used in multi-battery systems to select the battery that will actually supply power to the load system. It is also responsible for reporting any changes in the selector state to the system's power management software.
- **Smart Battery System Manager**, which manages the usage of all the smart batteries in a system. The SBS manager is an alternative to the smart battery selector that provides for added functionality, including the possibility for multiple batteries to simultaneously power the system. It also schedules and controls the charging of multiple batteries, and reports the characteristics of batteries powering the system to the management software.

A typical example of a smart battery system is shown in Figure 4, that includes two smart batteries, a smart battery charger, and a SBS manager, which communicate among themselves and with the host system via the SMBus. In the figure, the SBS manager has configured smart battery #1 to power the system, while smart battery #2 is charging. In addition, the flow of information relating to critical events (*e.g.*, battery low alarm) and battery data/status requests is shown in the figure.

Smart battery systems promise several advantages: (i) Better and more accurate sensing of battery data by the system will lead to battery aware power management policies and hence improved battery life, (ii) Systems will be able to use batteries of any chemistry type; new technologies can be directly utilized by existing SBS compatible systems, (iii) Smart charging will improve the safety and cycle life, while reducing charging times, and (iv) Battery data can be used not only by the host system but also by systems across the network, eventually supporting remote diagnosis and repair.

VI. CONCLUSIONS

Battery-driven system design is a rapidly evolving area that has the potential to significantly improve the lifetimes of battery-powered systems, over and beyond conventional low power design methodologies. In this paper, we surveyed this important emerging area of research, describing several promising new technologies in battery modeling, battery-efficient system architecture design, smart battery systems, and wireless networking protocols for battery operated devices. We believe significant opportunities exist for developing battery-driven system design methodologies and tools, as well as battery-friendly system architectures. We envision active research in this area in the near future, which will help reduce the growing battery-gap, by enabling the design of truly battery-efficient systems.

References

[1] S. H. Gunther, F. Binns, D. M. Carmean, and J. C. Hall, "Managing the Impact of Increasing Microprocessor Power Consumption." Intel Technology Journal, First Quarter, 2001, available online at http://developer.intel.com/technology/itj.

[2] I. Buchmann, "Batteries in a Portable World." http://www.cadex.com.

[3] A. R. Chandrakasan and R. W. Brodersen, *Low Power Digital CMOS Design*. Kluwer Academic Publishers, Norwell, MA, 1995.

[4] J. Rabaey and M. Pedram (Editors), *Low Power Design Methodologies*. Kluwer Academic Publishers, Norwell, MA, 1996.

[5] A. Raghunathan, N. K. Jha, and S. Dey, *High-level Power Analysis and Optimization*. Kluwer Academic Publishers, Norwell, MA, 1998.

[6] L. Benini and G. De Micheli, *Dynamic Power Management: Design Techniques and CAD Tools*. Kluwer, 1998.

[7] K. C. Syracuse and W. D. K. Clark, "A statistical approach to domain performance modeling for oxyhalide primary lithium batteries," in *Proc. Annual Battery Conference on Applications and Advances*, 1997.

[8] M. Pedram and Q. Wu, "Design considerations for battery-powered electronics," in *Proc. Design Automation Conf.*, pp. 861–866, June 1999.

[9] D. Rakhmatov and S. B. K. Vrudhula, "Time to failure estimation for batteries in portable systems," in *Proc. Int. Symp. Low Power Electronics & Design*, pp. 88–91, Aug. 2001.

[10] S. Hageman, "Simple PSPICE models let you simulate common battery types," *Electronic Design News*, vol. 38, pp. 117–132, Oct. 1993.

[11] S. Gold, "A PSPICE macromodel for lithium-ion batteries," in *Proc. Annual Battery Conference on Applications and Advances*, pp. 9–15, 1997.

[12] M. Glass, "Battery electrochemical non-linear dynamic SPICE model," in *Energy Conversion Engineering Conference*, pp. 292–297, 1996.

[13] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino, and R. Scarsi, "Discrete time battery model for high level power estimation," in *Proc. Design Automation & Test Europe (DATE) Conf.*, pp. 35–39, Mar. 2000.

[14] C. F. Chiasserini and R. R. Rao, "Importance of a pulsed battery discharge in portable radio devices," in *Proc. of MOBICOMM*, Aug. 1999.

[15] D. Panigrahi, C. F. Chiasserini, S. Dey, R. R. Rao, A. Raghunathan, and K. Lahiri, "Battery life estimation for mobile embedded systems," in *Proc. Int. Conf. VLSI Design*, pp. 55–63, Jan. 2001.

[16] T. Simunic, L. Benini, and G. D. Micheli, "Energy-efficient design of battery-powered embedded systems," *IEEE Trans. VLSI Systems*, vol. 9, pp. 15–28, Feb. 2000.

[17] S. Li and J. W. Evans, "Electrochemical-thermal model of lithium polymer batteries," *J. Electrochem. Soc*, vol. 147, pp. 2086–2095, June 2000.

[18] M. Doyle, T. F. Fuller, and J. S. Newman, "Modeling of galvanostatic charge and discharge of lithium/polymer/insertion cell," *J. Electrochem. Soc*, vol. 140, pp. 1526–1533, June 1993.

[19] T. F. Fuller, M. Doyle, and J. S. Newman, "Relaxation phenomena in lithium-ion insertion cells," *J. Electrochem. Soc*, vol. 141, pp. 982–990, Apr. 1994.

[20] T. I. Evans, T. V. Nguyen, and R. E. White, "A mathematical model of a lithium/thionyl chloride primary cell," *J. Electrochem. Soc*, vol. 136, pp. 328–329, Feb. 1989.

[21] K. Tsaur and R. Pollard, "Mathematical modeling of the lithium/thionyl chloride static cell," *J. Electrochem. Soc*, vol. 131, pp. 975–984, May 1984.

[22] W. B. Gu and C. Y. Wang, "Thermal-electrochemical modeling of battery systems," *J. Electrochem. Soc*, vol. 147, no. 8, pp. 2910–2922, 2000.

[23] W. B. Gu, C. Y. Wang, J. W. Weidner, R. G. Jungst, and G. Nagasubramanium, "Computational fluid dynamic modeling of a lithium/thionyl chloride battery with electrolyte flow," *J. Electrochem. Soc*, vol. 147, no. 2, pp. 427–434, 2000.

[24] N. Kularatna, *Power Electronics Design Handbook*, ch. 5, pp. 137–174. EDN Series for Design Engineers, 1998.

[25] "Rayovac corporation: Rechargeable alkaline product guide." http://www.rayovac.com/busoem/oem/specs/ren6c.shtml.

[26] S. Park, A. Savvides, and M. B. Srivastava, "Battery capacity measurement and analysis using lithium coin cell battery," in *Proc. Int. Symp. Low Power Electronics & Design*, pp. 382–387, Aug. 2001.

[27] D. Linden and T. Reddy, *Handbook of Batteries*. McGraw-Hill, NY, 2001.

[28] T. L. Martin, *Balancing Batteries, Power and Performance: System Issues in CPU Speed-Setting for Mobile Computing*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, 1999. Available on-line at http://www.eb.uah.edu/ tlmartin/papers.html.

[29] W. Weiser, B. Welch, A. Demers, and S. Shenker, "Scheduling for reduced CPU energy," in *Proc. USENIX Symposium on Operating Systems Design*, pp. 13–23, Nov. 1994.

[30] J. Luo and N. K. Jha, "Battery-aware static scheduling for distributed real-time embedded systems," in *Proc. Design Automation Conf.*, pp. 444–449, June 2001.

[31] L. Benini, A. Bogliolo, and G. D. Michelli, "A survey of design techniques for system level dynamic power management," *IEEE Trans. VLSI Systems*, vol. 8, pp. 299–316, June 2000.

[32] L. Benini, G. Castelli, A. Macii, and R. Scarsi, "Battery-driven dynamic power management," *IEEE Design & Test of Computers*, vol. 18, pp. 53–60, Apr. 2001.

[33] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino, and R. Scarsi, "Extending lifetime of portable systems by battery scheduling," in *Proc. Design Automation & Test Europe (DATE) Conf.*, pp. 197–201, Mar. 2001.

[34] C. F. Chiasserini and R. R. Rao, "Energy Efficient Battery Management," *IEEE J. on Selected Areas in Communications*, vol. 19, pp. 1235–1245, July 2001.

[35] Q. Wu, Q. Qiu, and M. Pedram, "An interleaved dual-battery power supply for battery powered electronics," in *Proc. Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 387–390, Jan. 2000.

[36] C. F. Chiasserini and R. R. Rao, "Improving battery performance by using traffic shaping techniques," *IEEE J. on Selected Areas in Communications*, vol. 19, pp. 1385–1394, July 2001.

[37] C. F. Chiasserini and R. R. Rao, "Routing protocols to maximize battery efficiency," in *Proc. of Milcom*, pp. 496–500, Oct. 2000.

[38] "Intel and Duracell create SBS (http://www.intel.com/ebusiness/products/mobile/initiatives/ar424.htm).".

[39] "Smart Battery System Implementers Forum (http://www.sbs-forum.org).".

[40] "System Management Bus (http://www.smbus.org).".