Battery Modeling for Energy-Aware System Design

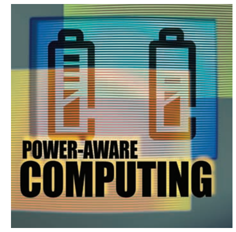

Computationally feasible mathematical models are now available that capture battery discharge characteristics in sufficient detail to let designers develop an optimization strategy that extracts maximum charge.

Ravishankar Rao
Sarma Vrudhula
University of Arizona, Tucson

Daler N. Rakhmatov
University of Victoria

Many features of modern portable electronic devices—such as high-speed processors, colorful displays, optical/magnetic storage drives, and wireless network interfaces—carry a significant energy cost. However, advances in battery technology have not kept pace with rapidly growing energy demands.[1,2]

Most laptops, handheld PCs, and cell phones use rechargeable electrochemical batteries—typically, lithium-ion batteries—as their portable energy source. These batteries take anywhere from 1.5 to 4 hours to fully charge, but they can run on this charge for only a few hours or, in the case of some newer pocket PCs, up to about 14 hours.

The battery has thus emerged as a key parameter to control in the energy management of portables.[3-8] To meet the stringent power budget of these devices, researchers have explored various architectural, hardware, software, and system-level optimizations to minimize the energy consumed per useful computation.

Maximizing the number of useful computations is effectively a problem of maximizing battery lifetime subject to system performance constraints. Given a load applied to a battery over a certain period, information about when the battery fails as well as its *state of charge*, or remaining capacity, at any time can be used to trade off system performance for battery lifetime at both the design stage and runtime, possibly with the user's active participation. For example, an energy-aware picture phone could let a user trade off image quality with talk time and the number of photos the phone could take using the remaining battery capacity.

Incorporating battery-state information into a lifetime optimization strategy requires a mathematical model that captures battery nonlinearities. Accurate low-level models[9-11] based on the differential equations that describe the complex phenomena occurring in an electrochemical cell have been around for about a decade, but solving these equations can take days. In recent years, however, researchers have developed high-level battery models[4,5,7,12-15] that reduce simulation time while predicting relevant variables with acceptable accuracy.

BATTERY DISCHARGE BEHAVIOR

Because the energy drawn from a battery is not always equivalent to the energy consumed in device circuits, understanding discharge behavior is essential for optimal system design.

Batteries consist of cells arranged in series, parallel, or a combination of both. Two electrodes—an anode and a cathode, separated by an electrolyte—constitute each cell's active material. When the cell is connected to a load, a reduction-oxidation reaction transfers electrons from the anode to the cathode. This transfer converts the chemical energy stored in the active material to electrical energy, which flows as a current in the external circuit.[16] As the battery discharges, its voltage drops; when this voltage falls below a certain cutoff, the battery disconnects from the load.

We define capacity in terms of charge units rather than energy.[17] *Full charge capacity* is the remaining

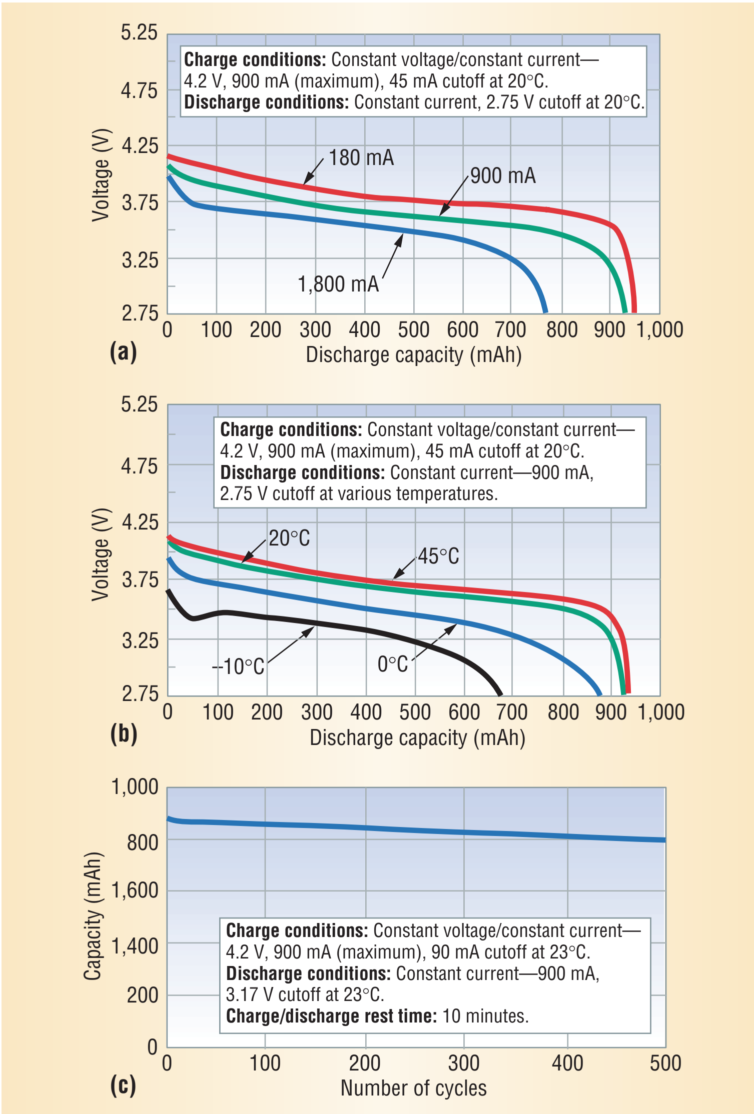

Figure 1. Lithium-ion battery discharge characteristics: (a) rate-dependent capacity, (b) temperature effect, and (c) capacity fading. Reprinted with permission, Matsushita Electric Industrial Co., Ltd.

capacity of a fully charged battery at the beginning of a discharge cycle, and *full design capacity* is the remaining capacity of a newly manufactured battery. Further, *theoretical capacity* is the maximum amount of charge that can be extracted from a battery based on the amount of active material it contains, *standard capacity* is the amount of charge that can be extracted from a battery when discharged under standard load and temperature conditions, and *actual capacity* is the amount of charge a battery delivers under given load and temperature conditions.[3]

Like other electrochemical systems, the laws of thermodynamics, electrode kinetics, and transport phenomena determine the complex set of equations that govern battery behavior.[18] Thus, as Figure 1[19] shows, battery discharge behavior is sensitive to numerous factors including the discharge rate, temperature, and the number of charge-recharge cycles. Consequently, battery discharge behavior deviates significantly from the behavior of an ideal energy source.

Rate-dependent capacity

Battery capacity decreases as the discharge rate increases. To illustrate this phenomenon, Figure 2 shows a simplified symmetric electrochemical cell in which similar processes occur at both electrodes.

In a fully charged cell (Figure 2a), the electrode surface contains the maximum concentration of active species. When the cell is connected to a load, a current flows through the external circuit; active species are consumed at the electrode surface and replenished by diffusion from the bulk of the electrolyte. However, this diffusion process cannot keep up with the reaction process, and a concentration gradient builds up across the electrolyte (Figure 2b).

A higher load current results in a higher concentration gradient[9] and thus a lower concentration of active species at the electrode surface. When this concentration falls below a certain threshold, which corresponds to the voltage cutoff, the electrochemical reaction can no longer be sustained at the electrode surface. At this point, the charge that was unavailable at the electrode surface due to the gradient remains unusable (Figure 2d) and is responsible for the reduction in capacity.

However, the unused charge is not physically "lost," but simply unavailable due to the lag between reaction and diffusion rates. Decreasing the discharge rate effectively reduces this lag as well as the concentration gradient. If the battery's load goes to zero, the concentration gradient flattens out after a sufficiently long time, reaching equilibrium again (Figure 2c). The concentration of active species near the electrode surface following this rest period makes some unused charge available for extraction.

System designers can exploit this *charge recovery effect* to control the discharge rate to maximize battery lifetime under performance constraints. However, at sufficiently low discharge rates, the battery will behave like an ideal energy source. For example, in Figure 1a, battery capacity will not significantly differ from that of the 180-mAh curve for constant currents below 900 mA.

Temperature effect

Temperature also strongly affects battery discharge behavior. Below room temperature (around 25°C), chemical activity in the cell decreases and internal resistance increases, reducing full charge capacity and increasing the slope of the discharge curve. At much higher temperatures, a decrease in internal resistance increases the full charge capacity and voltage. However, the higher rate of chemical activity, or *self-discharge*, can reduce the actual capacity delivered.[16] Unlike the discharge rate, temperature is not an easily controllable variable in energy-aware system design.

Capacity fading

Because of their high energy density and capacity, lithium-ion batteries are the popular choice for many portable applications. However, these batteries lose a portion of their capacity with each discharge-charge cycle. This *capacity fading* results from unwanted side reactions including electrolyte decomposition, active material dissolution, and passive film formation.[16] These irreversible reactions increase cell internal resistance, ultimately causing battery failure.

To deal with this problem, system users can attempt to control the depth of discharge before recharging. Typically, a battery subjected to shallow discharges—that is, voltage is still relatively high when recharging occurs—will be good for more cycles than a battery subjected to deep discharges—for example, until the cutoff voltage is reached.

BATTERY MODELS

Researchers have developed numerous computationally feasible mathematical models that capture battery behavior in sufficient detail. *Physical* models provide a detailed description of the physical processes occurring in the battery. *Empirical* models consist of ad hoc equations describing battery behavior with parameters fitted to match experimental data. *Abstract* models represent a battery as electrical circuits, discrete-time equivalents, stochastic process models, and so on. *Mixed* models offer a simplified view of the physical processes with empirically fitted parameters.

Models in each category can be evaluated according to four basic criteria:

- *Accuracy.* How closely do the predicted values of the battery variables of interest—lifetime, voltage, and so on—match experimental data? Can the model handle a general case of time-varying loads? Does it account for the temperature effect and capacity fading?
- *Computational complexity.* How long do the simulations take?
- *Configuration effort.* How many parameters can the model estimate? Does the model require in-depth knowledge of battery chemistry?
- *Analytical insight.* Do the equations describing the model provide some qualitative understanding of battery behavior? Is such insight useful in exploring ways to trade off lifetime and performance?

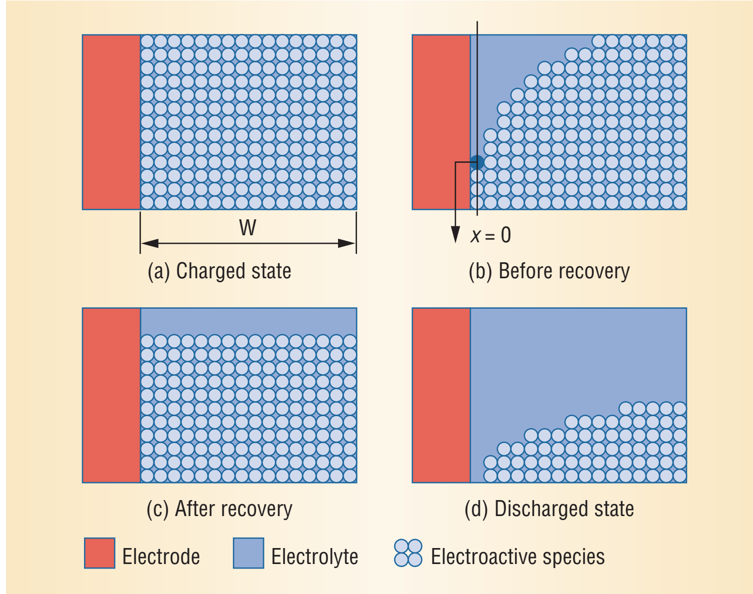

Figure 2. Battery operation. In a symmetric electrochemical cell, similar processes occur at both electrodes: (a) fully charged state, (b) before recovery, (c) after recovery, and (d) discharged state.

Table 1 summarizes a number of representative battery models with respect to these criteria and describes some of their applications.

Physical models

Physical models are the most accurate and have great utility for battery designers as a tool to optimize a battery's physical parameters. However, they are also the slowest to produce predictions and the hardest to configure, providing limited analytical insight for system designers.

Marc Doyle, Thomas F. Fuller, and John Newman[9,10] developed an isothermal electrochemical model that describes the charge and discharge of a lithium (anode)/polymer (electrolyte)/insertion (cathode) cell for a single cycle. This model uses concentrated solution theory to derive a set of differential equations that, when solved, provide cell potential values as a function of time.[18]

Dualfoil[20] is a Fortran program that uses this model to simulate lithium-ion batteries. The program reads the load profile as a sequence of constant current steps, and the battery lifetime is obtained from the output by reading off the time at which the cell potential drops below the cutoff voltage. Researchers have used Dualfoil to evaluate other battery models, and have extended the lithium/polymer/insertion cell model to include additional factors such as energy balance and capacity fading.[11]

Nevertheless, simulating a given lithium-ion battery can require specifying more than 50 parameters based on knowledge of the structure, chemical composition, capacity, temperature, and other characteristics. In addition, solving the model's interdependent partial differential equations requires using complex numerical techniques. As a result, simulating each load profile can take several hours or even days.

Table 1. Battery models and applications.

Model	Temperature effect	Capacity fading	Accuracy	Computational complexity	Configuration effort	Analytical insight	Applications
Physical							
Lithium-polymer-insertion cell (Doyle et al.)	Yes	Yes; support for Arrhenius temperature dependence and cycle aging added by Rong and Pedram	Very high	High	Very high (> 50 parameters)	Low	
Empirical							
Peukert's law	Yes; needs recalibration for each temperature	No	Medium (14% average error for constant load, 8% average error for interrupted and variable loads)	Low	Low (2 parameters)	Low	
Battery efficiency (Pedram and Wu)	Yes; needs recalibration for each temperature	No	Medium	Low	Low (2 parameters)	Low	Design of interleaved dual-battery power supply; load splitting for maximum lifetime of multibattery systems
Weibull fit (Syracuse and Clark)	Yes	No	Medium	Low	Low (3 parameters)	Low	
Abstract							
Electrical-circuit (Gold)	Yes	Yes	Medium (12% error predicting cell voltage and thermal characteristics, 5% error predicting cycle aging)	Medium	Medium (> 15 parameters)	Medium	
Electrical-circuit (Bergveld et al.)	Yes	No	Medium	Medium	High (> 30 parameters)	Medium	Thermostatic charge method: high charging efficiency
Discrete-time (Benini et al.)	Yes	No	Medium (1% compared to Hspice continuous-time model)	Medium	Medium (>15 parameters)	Medium	Dynamic Power Management; multibattery discharge
Stochastic (Chiasserini and Rao)	No	No	High (1%)	Low	Low (2 parameters)	Medium (stochastic model of load pattern assumed)	Shaping load pattern to exploit charge recovery
Mixed							
Analytical high-level (Rakhmatov et al.)	No	No	High (5%)	Medium	Low (2 parameters)	High	Task scheduling by sequencing and *V*/*f* scaling; analysis of discharge methods for multibattery systems
Analytical high-level (Rong and Pedram)	Yes	Yes	High (3.5%)	Medium	Medium (> 15 parameters)	High	

Empirical models

Empirical models are the easiest to configure, and they quickly produce predictions, but they generally are the least accurate. Although they work well in certain special cases, the constants used have no physical significance, which seriously limits their analytical insight.

Peukert's law. Some models attempt to capture nonideal discharge behavior using relatively simple equations in which the parameters match empirical data. While an ideal battery with capacity C discharged at a constant current would be expected to have a lifetime L given by $C = LI$, Peukert's law[16] expresses this as a power law relationship, $C = LI^{\alpha}$. The exponent provides a simple way to account for rate dependence. However, the α values for different temperatures must be obtained empirically, and the fit is not always accurate.[21]

Though easy to configure and use, Peukert's law does not account for time-varying loads. Most batteries in portable devices experience widely varying loads—for example, a pocket PC user may run a movie player application followed by a notes editor, which yields a profile with two very different loads for the battery.

Battery efficiency model. Massoud Pedram and Qing Wu[5] model battery *efficiency*—the ratio of actual capacity to theoretical capacity—as a linear-quadratic function of the load current. They derive bounds on the actual power consumed for different current distributions with the same average current and show that these bounds depend on the current's maximum and minimum values. Among all distributions with the same mean, a constant current (least variance) would give the longest battery lifetime, and a uniformly distributed current (highest variance) would give the shortest.

This model accounts for rate dependence and can handle variable loads. Researchers have used it, with slight modifications, to maximize the lifetime of multibattery systems,[6] to minimize the discharge-delay product in an interleaved dual-battery system design,[22] and in static task scheduling for real-time embedded systems.[23]

Weibull fit model. K.C. Syracuse and W.D.K. Clark[13] used statistical methods to model the discharge behavior of lithium-oxyhalide cells. For a fixed load and temperature, they noted battery voltage values at various stages of discharge. They then fit a Weibull model with three coefficients to these values to express voltage as a function of *delivered capacity,* or charge lost. Syracuse and Clark estimated the coefficients for different load/temperature combinations similarly, and modeled the coefficients' variation as a quadratic surface. They used a similar method to predict battery lifetime as a function of load and temperature.

Abstract models

Abstract models attempt to provide an equivalent representation of a battery.

Instead of modeling discharge behavior either by describing the electrochemical processes in the cell or by empirical approximation, abstract models attempt to provide an equivalent representation of a battery. Although the number of parameters is not large, such models also employ lookup tables that require considerable effort to configure. In addition, despite acceptable accuracy and computational complexity, these models have limited utility for automated design space exploration because they lack analytical expressions for many variables of interest.

Electrical-circuit and discrete-time models are particularly useful when compatible models of other system components—circuit models or VHSIC Hardware Description Language (VHDL) models—are available to simulate the entire system in a single continuous-time or discrete-time environment.

Electrical-circuit models. Steven C. Hageman[24] and Sean Gold[12] have each proposed PSpice circuits consisting of linear passive elements, voltage sources, and lookup tables to model nickel-metal-hydride and lithium-ion batteries, respectively. Henk Jan Bergveld, Wanda S. Kruijt, and Peter H.L. Notten[14] likewise devised an electrical-circuit model of a nickel-cadmium battery by grouping the mathematical equations describing the battery processes.

In Gold's approach, capacity fading is modeled by a capacitor C_{CAP} whose capacitance decreases linearly with the number of cycles. The load current I minus a rate-dependence offset flows through this capacitance. The voltage across C_{CAP} represents the ratio of delivered capacity to full charge capacity. This normalized state of charge is then converted via a lookup table into a voltage V_{COMP}.

The temperature effect is modeled as a resistor-capacitor circuit with two temperature-dependent sources, $V_{AMBIENT} \propto T$ and $E_{RISE} \propto I^2R_{cell}$, where T is the ambient temperature and R_{cell} is cell internal resistance. The main loop computes the cell voltage by superposing the effect of the state of charge, temperature, and cell internal resistance.

Electrical-circuit models are inherently continuous-time and, while their simulation times are faster than those of physical models, they are still time-consuming. For example, while the number of circuit

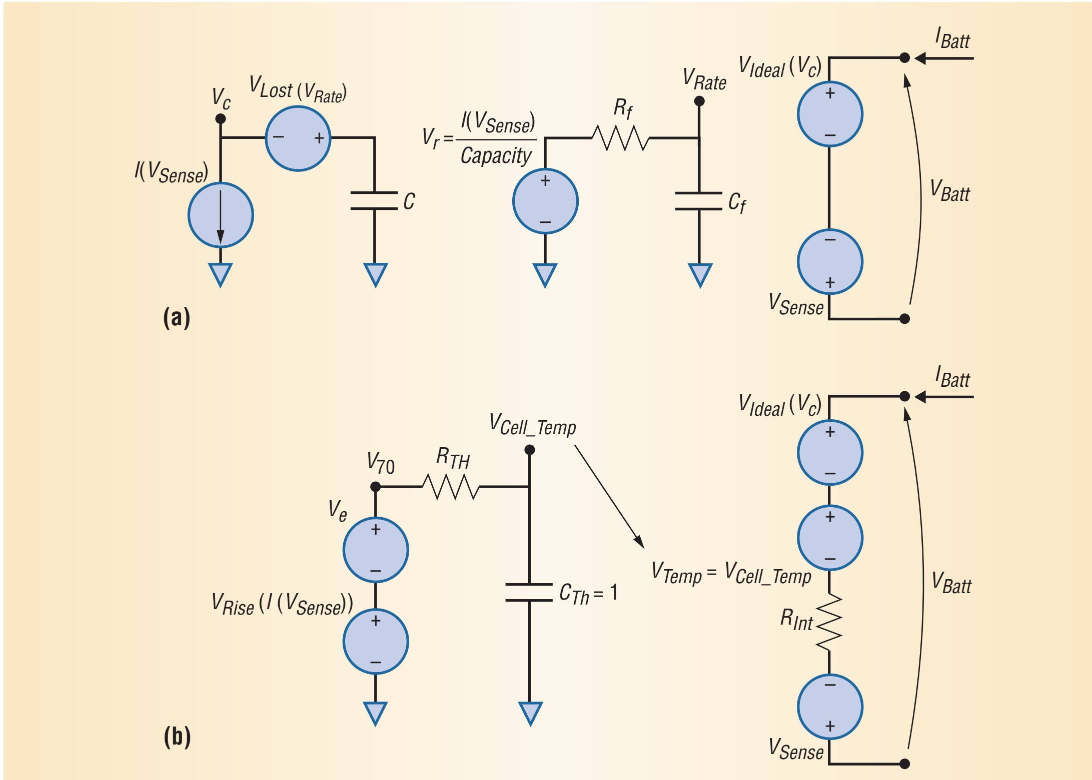

Figure 3. Continuous-time model used by Benini et al. in their discrete-time approximation. The model incorporates battery voltage dependence on (a) first-order effects and (b) second-order effects.

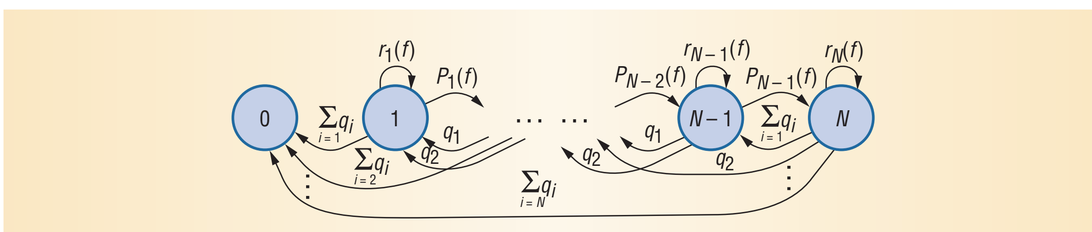

Figure 4. Stochastic cell model. Given a cell state of charge 0, 1, ..., N, each discharge demand of i units causes a transition to i states lower, while rest periods cause state transitions to successively higher states.

parameters in Gold's model is not large, configuring the lookup tables requires substantial effort.

Discrete-time model. Using VHDL, Luca Benini and colleagues[15] approximated the continuous-time model shown in Figure 3 to a discrete-time model. Their approach incorporates battery voltage dependence on first-order effects—charge state, discharge rate, and discharge frequency—and the second-order effects of temperature and internal resistance. A lookup table models DC-DC converter characteristics. For constant and time-varying loads, this model predicts lifetime values for different battery types that are similar to those of the continuous-time model on which it was based. Researchers have used the discrete-time model to compare different Dynamic Power Management[15,25] and multibattery discharge techniques.[26]

Stochastic model. Carla-Fabiana Chiasserini and Ramesh R. Rao[7] developed a battery model, shown in Figure 4, that represents charge recovery as a decreasing exponential function of the state of charge and discharged capacity. Assuming each cell's load to be a pulsed discharge, this model represents discharge and recovery as a transient stochastic process.

Each discharge demand of i units causes a transition to i states lower, while rest periods cause state transitions to successively higher states. Capacity gain is expressed as $G = A_{cu}/N$, where A_{cu} represents the average number of charge units and N is the *nominal capacity*—the charge extractable by a constant load. Using the Dualfoil simulator, the researchers obtained curves for G as a function of the discharge rate for different values of load current density and fitted two of the model parameters to match these curves.

This model is useful for representing pulsed discharge, as it can obtain capacity gain for different types of stochastic loads analytically without simulation; Chiasserini and Rao reported several analytical results related to distributing the load between two cells of a battery package. However, because it concentrates only on charge recovery, their model does not account for other battery nonlinearities. Debashis Panigrahi and colleagues[27] added a lookup table to incorporate rate dependence, resulting in an abstract model that is both fast and capable of producing predictions closely matching Dualfoil predictions.

Mixed models

Some models combine a high-level representation of a battery for which experimental data determines the parameters with analytical expressions based on physical laws. For example, Daler N. Rakhmatov and Sarma Vrudhula[4] developed a high-level analytical model that characterizes a battery using two constants, α and β, derived from the lifetime values for a series of constant load tests. The α parameter is a measure of the battery's theoretical capacity, while β models the rate at which the active charge carriers are replenished at the electrode surface.

Starting with Faraday's law for electrochemical reaction and Fick's laws[28] for concentration behavior during one-dimensional diffusion in an electrochemical cell, these researchers obtained the following expression relating the load i, battery lifetime L, and battery parameters:

$$\alpha = \int_0^L i(\tau)\ d\tau + \lim_{\epsilon \to 0^+} 2 \sum_{m=1}^{\infty} \int_0^{L-\epsilon} i(\tau)\ e^{-\beta^2 m^2 (L -}$$

The first term represents the charge the load consumed over the period $[0, L)$, while the second term represents the charge that was "unavailable" at the electrode surface at the time of failure L. The unavailable charge models the effect of the concentration gradient that builds up as the flow of active species through the electrolyte falls behind the rate at which they discharge at the electrode surface.

Battery lifetime predictions using this model closely match both Dualfoil simulation results and experimental measurements.[29,30] The simulation time is moderate, and the authors point out that it is possible to trade off accuracy with speed by reducing the number of terms in the summation and approximating the continuous-time load waveform $i(t)$ to an N-step staircase (in the extreme case of a constant load approximation, $N = 1$). However, the model does not account for the effect of temperature and capacity fading on the discharge characteristics. Compared to the stochastic model, it has higher computational complexity but requires less configuration effort and offers more analytical insight.

Peng Rong and Pedram[17] recently proposed a high-level battery model to estimate remaining capacity that considers both the temperature effect and capacity fading with successive cycles but assumes a constant current load. They derived an expression for cell terminal voltage as a function of time and, using the Arrhenius dependence on temperature of cell kinetics and transport phenomena, obtained an expression for the bulk properties of the active material as a function of the temperature. They also derived an expression for film thickness as a function of the temperature, discharge rate, and number of cycles. Using these quantities, they define *state of charge* as remaining capacity/full charge capacity and *state of health* as full charge capacity/full design capacity.

These capacity ratios match well with Dualfoil simulations, and the model effectively captures the effect of temperature and cycle aging on the battery state of charge. However, the expressions for remaining capacity are more involved than those in Rakhmatov and Vrudhula's model, requiring configuration of more than 15 different parameters to set up the equivalent battery. In addition, the constant-load assumption limits the model's application for optimizing portable systems with highly variable loads.

Using a model to understand battery behavior can help system designers devise optimal battery management algorithms and policies.

APPLICATIONS

Using one of these models to understand battery behavior can help system designers devise optimal battery management algorithms and policies. Examples of such management include shaping the discharge current profile under performance constraints, developing optimal charging procedures, and customizing batteries for a given application under volume and weight constraints.

Battery-aware power supply design

The digital circuits in most modern electronic devices are designed using complementary metal oxide semiconductor logic. The supply voltage V_{dd} and the threshold voltage V_{th} characteristic of CMOS transistors affect the power consumed during switching in these circuits.

To minimize the product of battery discharge and delay, Wu, Qinru Qiu, and Pedram[22] use the battery efficiency model[5] to compute the optimal V_{dd}. They define battery discharge as the ratio between the actual energy drawn from the battery and the total energy stored in a new battery. For CMOS circuits, the delay is proportional to $V_{dd}/(V_{dd} - V_{th})^{\alpha}$, where $1 \le \alpha \le 2$.

These researchers also propose an interleaved power supply system, shown in Figure 5, to discharge a pair of batteries with different current-capacity characteristics. For a total energy

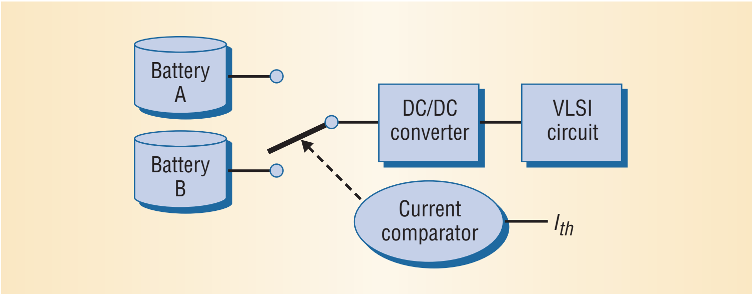

Figure 5. Interleaved power supply system. A dual-battery system offers a 25 percent improvement in power supply over a single optimal battery.

constraint, they find the distribution of active material weight between the two batteries that maximizes system lifetime and then compare the load value to a threshold to choose the more efficient battery. Hspice simulations using random current distributions show that a dual-battery system offers a 25 percent improvement in power supply over a single optimal battery.[22]

Static task scheduling for real-time embedded systems

Drawing on previous work,[5,8] Jiong Luo and Niraj K. Jha[23] proposed a battery-aware scheduling algorithm for real-time embedded systems that support variable voltages. The algorithm seeks to reduce the mean value of the discharge current and shape the discharge profile to maximize battery lifetime. The actual power drawn from the battery is the cost function to be minimized.

After obtaining an initial feasible schedule with a list-scheduling algorithm, Luo and Jha used a global shifting transformation to reduce peak power consumption. They then applied local transformations involving iteratively sequencing and shifting tasks, starting at points along the hyperperiod with highest power consumption, to reduce the cost function. Next, they performed voltage-clock scaling for processing elements that support variable voltages by distributing the total available slack time among all the tasks. They chose speed and voltage reduction ratios for each task to minimize total energy consumption.

Rakhmatov, Vrudhula, and Chaitali Chakrabarti[21,31] used an analytical model of a battery to develop a cost function $\sigma(t)$ of a battery as a function of the time-varying load $i(t)$. The cost function is the sum of the actual charge lost to the load $l(t)$ and the temporarily unavailable charge $u(t)$. The task-scheduling problem involved assigning start times t_k and voltage-frequency combinations V_k and ϕ_k for each of a set of N tasks to minimize the cost function of the chosen schedule, subject to the following constraints:

- the scheduling maintains task dependencies,
- the time by which all tasks complete does not exceed a deadline B, and
- the battery does not fail before completing all tasks.

Minimizing the charge lost to the load—or effectively, the energy consumed—after completing all tasks was the objective of several early approaches to task scheduling, but the authors point out that the charge lost is actually a lower bound on σ. Given the difficulty of deriving an exact solution to the task-scheduling problem, they proposed heuristics for the general case starting from initial solutions corresponding to the minimum-charge, lowest-power, or highest-power load profile. These heuristics are based on the provable properties of the cost function. The researchers subsequently improved the load profile by inserting rest periods, voltage up/down scaling, and task sequencing.

Load-profile shaping for multibattery systems

More portable devices such as laptops employ multiple batteries. Because researchers have found the traditional method of discharging batteries in sequence to be suboptimal, there is increasing interest in developing new discharge methods at both the experimental and analytical level.

Experimental work. Benini and colleagues[26] used a discrete-time model[15] to simulate three different techniques:

- *sequentially* discharging each battery until it fails;
- *static switching*—discharging each battery for a fixed duration and in round-robin schedule; and
- *dynamic switching*—scheduling the healthiest battery for discharge at any instant dynamically while the other batteries rest.

For comparison, the authors also simulated a monolithic equivalent of the multibattery system. They generally found the lifetimes to follow the relation monolithic ≥ dynamic switching ≥ static switching ≥ sequential discharge. They also observed that, as frequency increased in the static switching case, the resulting lifetime approached that of the monolithic battery.

Another effort led by Benini[6] incorporates fast switching between batteries to achieve a "virtual parallel" discharging of multiple batteries. Modeling rate dependence using an approach similar to that of Pedram and Wu,[5] the researchers performed nonlinear optimization to split the load current over a set of multiple batteries to maximize system lifetime. Their proportional current-allocation scheme was a moderate improvement compared to equally dividing the load among all batteries.

Davide Bruni and colleagues[32] implemented the virtual parallel scheme along with one in which multiple batteries are connected in series and the combined voltage down-converted and demonstrated good improvement in system lifetime for high current loads.

Analytical work. Chiasserini and Rao[7] applied results from load balancing in computer systems to distributing the load between two cells of a battery package. They first considered a *delay-free approach* that provides charge units to the load as soon as they are required, then a *delayed approach* that introduces some delay so that the discharge profile can be shaped to maximize battery lifetime. They used a stochastic cell model to analytically show that a *best of two* approach is better than the round-robin and random scheduling approaches. The delayed approach is similar to dynamic switching[26] but buffers requests if no cell is active. The goal is to let each cell recover as much charge as required to maximize the charge it delivers.

The charge delivered using the delayed approach hypothetically equals the battery's theoretical capacity, at the cost of delay. However, assuming an infinite buffer to hold the load's requests for charge units is unrealistic for most portable applications. Also, many applications, such as the display, cause a constant drain on the battery; such background discharge can be significant and yet cannot be modeled stochastically.

We used a high-level battery model[4] to obtain an upper bound on the lifetime of a multibattery system for a given load.[33] This study showed that the lifetime of multiple batteries discharged

- sequentially is no greater than that of an equivalent monolithic battery discharged by the same load;
- simultaneously (in parallel) is equal to that of an equivalent monolithic battery discharged by the same load; and
- by switching at a fixed frequency approaches that of an equivalent monolithic battery at high frequencies, when both are discharged by the same constant load.

Our results also demonstrate that parallel discharge performs as well as a monolithic battery, while switching techniques achieve this performance only asymptotically. As technology supporting simultaneous discharge of multiple batteries is available,[34] we conclude that parallel discharge is preferable to more complex switching techniques.

Battery-aware Dynamic Power Management

DPM policies attempt to minimize a system's average power consumption by shifting to low-power modes such as standby, sleep, and off if the system remains idle after a certain time-out period. These periods are based on the overhead due to mode transitions and the energy savings resulting from the transition. However, DPM policies do not consider the battery's state of charge in determining when to change modes.

Benini and colleagues[25] proposed closed-loop DPM policies that exploit battery-state information from a discrete-time battery model[15] to change the system state. They implemented a simple scheme to switch between a "fine" and a lower-power "raw" play mode on an MP3 player, based on whether the battery voltage was above or below a certain threshold. The researchers showed significant improvements in lifetime with a small performance penalty.

Research in battery-aware optimization is now moving from stand-alone devices to ad hoc and distributed-sensor networks.

The accurate mathematical modeling of batteries is now a mature field, and researchers have applied such models fairly successfully in optimizing system behavior to achieve maximum lifetime. Because many of these models are independent of the battery chemistry, they should remain relevant as technology advances. The task schedulers of portable device operating systems ultimately must incorporate algorithms that dynamically adapt system behavior based on the battery's state of charge. Implementation efforts are already under way—for example, the advanced configuration and power interface specification implemented in most modern laptops offers power-saving options.

Research in battery-aware optimization is now moving from stand-alone devices to networks of wireless devices—specifically, ad hoc and distributed-sensor networks. The collaborative nature of these networks provides ample ground for using battery-state information to improve the nodes' efficiency. Battery life is especially important in such networks because they are often deployed in potentially hazardous or unreachable conditions to sense data for reconnaissance, environmental-monitoring, or health-monitoring purposes. Developing efficient routing protocols, medium-access protocols, and discharge-shaping techniques to maximize battery life are active areas of research in this field. ■

Acknowledgment

This work was carried out at the National Science Foundation's State/Industry/University Cooperative Research Centers' Center for Low Power Electronics, which is supported by the NSF under grant EEC-9523338, the Department of Commerce of the State of Arizona, and an industrial consortium.

References

1. I. Buchmann, *Batteries in a Portable World: A Handbook on Rechargeable Batteries for Non-Engineers,* 2nd ed., Cadex Electronics, 2001.
2. S.H. Gunther et al., "Managing the Impact of Increasing Microprocessor Power Consumption," *Intel Technology J.,* Q1 2001; www.intel.com/technology/itj/q12001/articles/art_4.htm.
3. K. Lahiri et al., "Battery-Driven System Design: A New Frontier in Low Power Design," *Proc. Joint 15th Int'l Conf. VLSI Design/7th Asia and South Pacific Design Automation Conf.,* IEEE CS Press, 2002, pp. 261-267.
4. D.N. Rakhmatov and S.B.K. Vrudhula, "An Analytical High-Level Battery Model for Use in Energy Management of Portable Electronic Systems," *Proc. 2001 IEEE/ACM Int'l Conf. Computer-Aided Design,* IEEE Press, 2001, pp. 488-493.
5. M. Pedram and Q. Wu, "Design Considerations for Battery-Powered Electronics," *Proc. 36th ACM/IEEE Design Automation Conf.,* ACM Press, 1999, pp. 861-866.
6. L. Benini et al., "Discharge Current Steering for Battery Lifetime Optimization," *Proc. 2002 Int'l Symp. Low-Power Electronics and Design,* ACM Press, 2002, pp. 118-123.
7. C.F. Chiasserini and R.R. Rao, "Energy Efficient Battery Management," *IEEE J. Selected Areas in Comm.,* vol. 19, no. 7, 2001, pp. 1235-1245.
8. T.L. Martin, *Balancing Batteries, Power and Performance: System Issues in CPU Speed-Setting for Mobile Computing,* doctoral dissertation, Dept. Electrical and Computer Eng., Carnegie Mellon Univ., 1999.
9. M. Doyle, T.F. Fuller, and J. Newman, "Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell," *J. Electrochemical Soc.*, vol. 140, no. 6, 1993, pp. 1526-1533.
10. T.F. Fuller, M. Doyle, and J. Newman, "Simulation and Optimization of the Dual Lithium Ion Insertion Cell," *J. Electrochemical Soc.,* vol. 141, no. 1, 1994, pp. 1-10.
11. G.G. Botte, V.R. Subramanian, and R.E. White, "Mathematical Modeling of Secondary Lithium Batteries," *Electrochimica Acta,* vol. 45, nos. 15-16, 2000, pp. 2595-2609.
12. S. Gold, "A PSPICE Macromodel for Lithium-Ion Batteries," *Proc. 12th Ann. Battery Conf. Applications and Advances,* IEEE Press, 1997, pp. 215-222.
13. K.C. Syracuse and W.D.K. Clark, "A Statistical Approach to Domain Performance Modeling for Oxyhalide Primary Lithium Batteries," *Proc. 12th Ann. Battery Conf. Applications and Advances,* IEEE Press, 1997, pp. 163-170.
14. H.J. Bergveld, W.S. Kruijt, and P.H.L. Notten, "Electronic-Network Modeling of Rechargeable NiCd Cells and Its Application to the Design of Battery Management Systems," *J. Power Sources,* vol. 77, no. 2, 1999, pp. 143-158.
15. L. Benini et al., "Discrete-Time Battery Models for System-Level Low-Power Design," *IEEE Trans. VLSI Systems,* vol. 9, no. 5, 2001, pp. 630-640.
16. D. Linden and T. Reddy, *Handbook of Batteries,* 3rd ed., McGraw-Hill, 2001.
17. P. Rong and M. Pedram, "An Analytical Model for Predicting the Remaining Battery Capacity of Lithium-Ion Batteries," *Proc. 2003 Design, Automation and Test in Europe Conf. and Exposition,* IEEE CS Press, 2003, pp. 1148-1149.
18. J.S. Newman, *Electrochemical Systems,* 2nd ed., Prentice Hall, 1991.
19. Matsushita Electronic Industrial Catalogue, lightweight prismatic lithium ion (CGA series) batteries, CGA523450A; http://industrial.panasonic.com/www-data/pdf2/ACA4000/ACA4000CE190.pdf. This information is generally descriptive only and is not intended to make or imply any representation guarantee or warranty with respect to any cells and batteries.
20. J.S. Newman, "FORTRAN Programs for Simulation of Electrochemical Systems," Dualfoil.f Program for Lithium Battery Simulation; www.cchem.berkeley.edu/~jsngrp/fortran.html.
21. D. Rakhmatov, S. Vrudhula, and C. Chakrabarti, "Battery-Conscious Task Sequencing for Portable Devices Including Voltage/Clock Scaling," *Proc. 39th Design Automation Conf.,* ACM Press, 2002, pp. 189-194.
22. Q. Wu, Q. Qiu, and M. Pedram, "An Interleaved Dual-Battery Power Supply for Battery-Operated Electronics," *Proc. 2000 Conf. Asia and South Pacific Design Automation,* IEEE Press, 2000, pp. 387-390.
23. J. Luo and N.K. Jha, "Battery-Aware Static Scheduling for Distributed Real-Time Embedded Systems," *Proc. 38th Design Automation Conf.,* ACM Press, 2001, pp. 444-449.
24. S.C. Hageman, "PSpice Models Nickel-Metal-Hydride Cells," *EDN Access,* 2 Feb. 1995; www.

reed-electronics.com/ednmag/archives/1995/020295/03di1.htm.

25. L. Benini et al., "Battery-Driven Dynamic Power Management," *IEEE Design and Test of Computers,* vol. 18, no. 2, 2001, pp. 53-60.
26. L. Benini et al., "Extending Lifetime of Portable Systems by Battery Scheduling," *Proc. 2001 Conf. Design, Automation and Test in Europe,* IEEE Press, 2001, pp. 197-203.
27. D. Panigrahi et al., "Battery Life Estimation of Mobile Embedded Systems," *Proc. 14th Int'l Conf. VLSI Design,* IEEE CS Press, 2001, pp. 57-63.
28. A.J. Bard and L.R. Faulkner, *Electrochemical Methods: Fundamentals and Applications,* 2nd ed., John Wiley & Sons, 2000.
29. D. Rakhmatov, S. Vrudhula, and D.A. Wallach, "Battery Lifetime Prediction for Energy-Aware Computing," *Proc. 2002 Int'l Symp. Low Power Electronics and Design,* ACM Press, 2002, pp. 154-159.
30. D. Rakhmatov, S. Vrudhula, and D.A. Wallach, "A Model for Battery Lifetime Analysis for Organizing Applications on a Pocket Computer," to appear in *IEEE Trans. VLSI Systems,* vol. 11, no. 6, 2003.
31. D. Rakhmatov and S. Vrudhula, "Energy Management for Battery-Powered Embedded Systems," *ACM Trans. Embedded Computing Systems,* vol. 2, no. 3, 2003, pp. 277-324.
32. D. Bruni, L. Benini, and B. Riccó, "System Lifetime Extension by Battery Management: An Experimental Work," *Proc. 2002 Int'l Conf. Compilers, Architecture, and Synthesis for Embedded Systems,* ACM Press, 2002, pp. 232-237.
33. R. Rao, S. Vrudhula, and D. Rakhmatov, "Analysis of Discharge Techniques for Multiple Battery Systems," *Proc. 2003 Int'l Symp. Low Power Electronics and Design,* ACM Press, 2003, pp. 44-47.
34. M. Gurries, "Dual Battery Power Manager Increases Run Time by 12% and Cuts Charge Time in Half," Linear Technology Design Note 277; www.linear.com/pdf/dn277f.pdf.

Ravishankar Rao *is a graduate student in the Electrical and Computer Engineering Department at the University of Arizona, Tucson. His research interests include energy management for portable battery-powered systems. Rao received a BE in electronics and communication engineering from Regional Engineering College, Surathkal, India. He is a student member of the IEEE. Contact him at rao@ece.arizona.edu.*

Sarma Vrudhula *is a professor in the Electrical and Computer Engineering Department at the University of Arizona, Tucson, and director of the NSF Center for Low Power Electronics. His research interests are in design automation and VLSI computer-aided design. Vrudhula received a PhD in electrical engineering from the University of Southern California. He is a senior member of the IEEE. Contact him at sarma@ece.arizona.edu.*

Daler N. Rakhmatov *is an assistant professor in the Department of Electrical and Computer Engineering at the University of Victoria, British Columbia, Canada. His research interests include modeling and optimization of embedded systems, with an emphasis on hardware-software support for energy-efficient reconfigurable computing. Rakhmatov received a PhD in electrical and computer engineering from the University of Arizona, Tucson. He is a member of the IEEE. Contact him at daler@ece.uvic.ca.*