

Non-ideal Battery Properties and Low Power Operation in Wearable Computing

Thomas L. Martin, Daniel P. Siewiorek

Institute for Complex Engineered Systems, Carnegie Mellon University, Pittsburgh, PA 15213
[tlm,dps]@cs.cmu.edu

Abstract

This paper describes non-ideal properties of batteries and how these properties may impact power-performance trade-offs in wearable computing. The first part of the paper details the characteristics of an ideal battery and how these characteristics are used in sizing batteries and estimating discharge times. Typical non-ideal characteristics and the regions of operation where they occur are described. The paper then covers results from a first-principles, variable-load battery model, showing likely areas for exploiting battery behavior in mobile computing. The major result is that when battery behavior is non-ideal, lowering the average power or the energy per operation may not increase the amount of computation that can be completed in a battery life.

1. Introduction

Two of the major constraints on mobile and wearable computing are size and weight, of which the battery is a large portion. Reducing the battery is thus a key to reducing the overall system bulk. The usual approach to achieving this is to decrease the power consumption of the hardware, either by power management, putting unused systems into low power modes; or by a power-performance trade-off, completing a computation at a slower speed for less power. While power management and power-performance trade-offs are important for all mobile computers, they are more so for wearable computers because of their performance intensive user interfaces and their tighter constraints on size and weight. Previous work in power-performance trade-offs attempted to minimize the energy-delay product [4] or the energy per operation [9]. When the system is battery-powered, however, minimizing either measure may not maximize the computations per battery life. Non-ideal battery properties may come into play, as will be shown using both simulation results from a first-principles battery model. These properties must be considered during wearable computer design and for software control of power management and power-performance trade-offs.

The paper begins by describing the ideal characteristics of batteries in Section 2. Then Section 3 covers the non-ideal characteristics of batteries and the regions of operation where the characteristics are exhibited. Section 4 presents simulation results from a first-principles, variable-load model of Li-ion cells. Finally, Section 5 summarizes ways in which the non-ideal battery properties can be exploited.

2. Ideal battery properties and discharge time estimates

The two most important properties of batteries from the viewpoint of someone using them are voltage and capacity. An ideal battery has a constant voltage throughout a discharge, which drops instantaneously to zero when the battery is fully discharged, and has constant capacity for all values of the load, as shown in Figure 1.

For sizing batteries, the battery voltage should be in the allowable range of the power supply of the device in question. The battery voltage is considered to be the rated voltage of the battery, e.g. 1.2V per cell for NiCd batteries and 3.6V per cell for most Li-ion batteries. The charge capacity of the battery is typically given in terms of Amp-hours or milliAmp-hours and is called the battery's "C" rating. The C rating is used in the battery industry to normalize the load current to the battery's capacity [6]. The advantage of C ratings is that it allows battery manufacturers to present one graph of discharge curves for batteries of similar construction but different capacities. Loads are then measured relative to the C rating, e.g. a 10 mA load on a battery with a rated capacity of 100 mAh is a load of 0.1C.

For mobile systems, the discharge time T is usually estimated to be the battery's rated voltage V multiplied by the charge capacity C , divided by the average power P of the system, or $T = (C \times V) / P$. The rated voltage multiplied by the charge capacity is the battery's nominal energy capacity, typically given in Watt-hours (1 Wh = 3600 J). As Section 4 will show, this method will overestimate the battery life if the load has a large peak value.




Figure 1. Characteristics of an ideal battery: Constant voltage and constant capacity

3. Non-ideal battery properties

While ideally a battery has constant voltage and capacity, in practice both vary widely. Figure 2a shows the battery voltage as a function of discharge time for two different loads. Load 1 is smaller than load 2. Because of resistance and other losses, the voltage throughout the discharge is lower for load 2 than load 1. The voltage for each load also drops over the course of the discharge due to changes in the battery's active materials and reactant concentrations.

The capacity also varies with the value of the load. The two major ways in which it varies are loss of capacity with increasing load, and an effect called recovery where an intermittent load may have a larger capacity than a continuous load. Figure 2b shows the loss of capacity with increasing load current for a typical NiCd battery. The capacity decreases by about 40% over a range of discharge rates of 0.1C to 10C. (Note that the capacity in Figure 2b exceeds 100% at low rates because the C rating is specified as the capacity for a given time of discharge. The capacity in Figure 2b was measured at the 2 hour rate, since 100% capacity occurs at 0.5C. If the capacity had been measured at the 10 hour rate, 100% would have occurred at 0.1C.)

The second non-ideal capacity property, recovery, is shown in Figure 2c. A reduction of the load for periods of time results in an increase in battery capacity. The voltage rises while the load is reduced, and the overall time of discharge increases. This phenomena occurs because, during the time when load is reduced, reactants in the battery diffuse to the reaction location, allowing more of them to be used during the life of the battery. The degree to which the battery recovers depends on the discharge rate and the length of time the load is reduced, as well as the details of

the battery construction.

It is widely known that the battery voltage varies during discharge. For example, power supplies are usually rated over a range of input voltages. When a power supply is used with a battery, it is necessary to ensure that the range of the supply's input voltage includes the range of the battery voltage during discharge. Since the voltage variation is widely known, this paper will not focus on it. The non-ideal capacity properties, on the other hand, are not widely known, and so will be the main subject of the remainder of this work. Given that a battery's discharge time is typically estimated using ideal values of voltage and capacity, the loss of capacity can lead to an overestimate of the discharge time for large loads. While a chart such as Figure 2b or a model such as Peukert's equation [6] allows one to account for the loss of capacity for loads that are constant and continuously on, in general loads are intermittent and variable. If recovery occurs, then the duration of the off times of the load must be considered in addition to the duration of its on times and its value while on. Models that account for both capacity loss and recovery are needed to determine if recovery occurs for the loads encountered in mobile computing, and if so, to properly estimate battery lives for intermittent loads.

4. Results with Doyle's variable load model

The typical load of a mobile computer system is not constant, but variable. A model is needed, then, to estimate the discharge time with variable loads. A variable-load model would ideally possess the following characteristics:

- Accurate relative capacity information (i.e. if several loads are simulated, then the model should correctly predict the relative difference in discharge times,

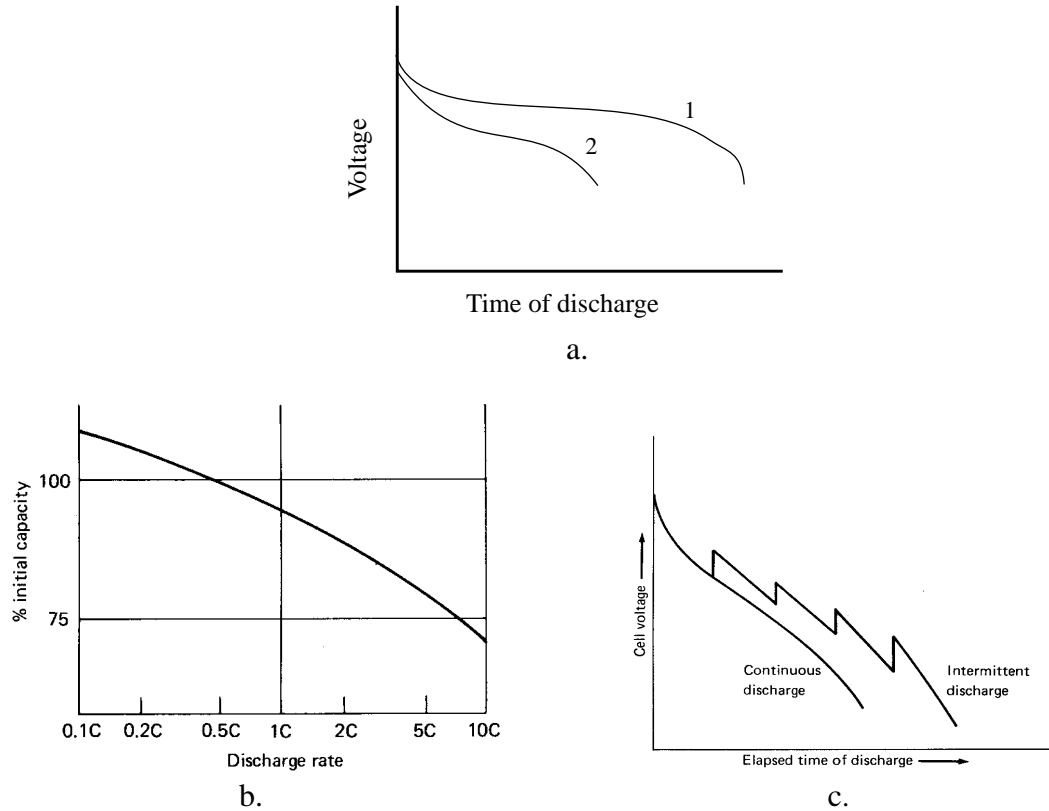


Figure 2. Non-ideal battery properties: (a) voltage change, (b) loss of capacity, and (c) recovery (after [6])

even if the actual differences are inaccurate.)

- Applicable to a variety of battery types
- Intuitive parameters and behavior
- Ease of correlation to actual cells

Of these four criteria, the first is the most important for this paper. The last three will become more important when battery models are more widely used in mobile system design. A number of battery models were investigated [1][3][5][7], but Doyle's model inspired the most confidence due to its having been created solely for Li-ion cells and due to its use in industry [2]. The other models had not been created for use with Li-ion cells and hence results with them would have required lengthy correlation with actual cells before their predictions could have been trusted.

Doyle's model was used to study the effect of intermittent discharges on the capacity. It was found that peak power predicts battery capacity better than average power.

Figure 3 shows the model results for battery capacity versus average power for continuous discharges over a range of loads, and for intermittent discharges for several combinations of peak power and duty cycle. The intermittent discharges were square waves with an off power of 0 W/kg. The two major features of the results are that the capacity decreases as the load power increases for continuous loads, and that there is a range where the peak power of an intermittent load rather than the average power is a stronger indicator of the battery's capacity. For example, the 300 W/kg continuous load results in a battery capacity of 90 Wh/kg (point A in the figure) and the 75 W/kg continuous load results in a battery capacity of 140 Wh/kg (point B), while the intermittent load with a peak power of 300 W/kg and duty cycle of 25% (i.e. an average power of 75 W/kg, point C) results in a capacity of approximately 100 Wh/kg. Thus using the average power of this intermittent load would over-estimate the battery capacity by about 40% (i.e., point B's 140 Wh/kg would be expected), while using the peak power would under-estimate it by only about 10% (i.e., point A's 90 Wh/kg

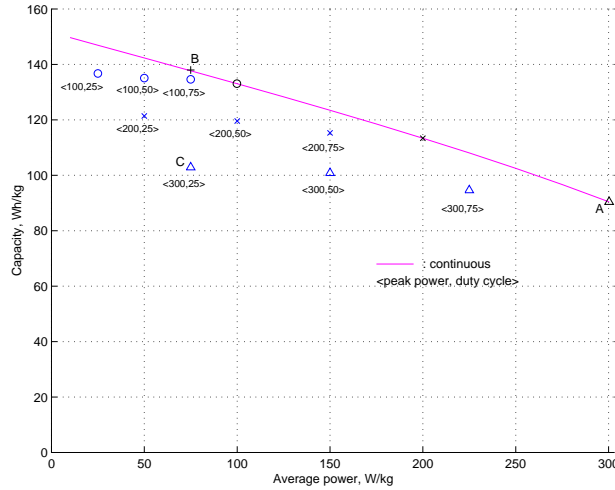


Figure 3. Doyle's Li-ion model results for capacity versus average power, showing difference between continuous and intermittent loads of same average value.

would be expected). To put these results in more common terms, the 75 W/kg continuous load B would have a battery life of about 1.9 hours, while the intermittent load C, with the same average power, 75 W/kg, would have a battery life of about 1.3 hours. Only when the peak power is below about 50 W/kg (about a 3 hour discharge when continuously on) would the peak and average power give about the same estimate of battery life.

The characteristics displayed in Figure 3 mean that minimizing energy per operation may not maximize computations per battery life. For example, suppose a mobile system has a dynamic power profile that is cyclic, having periods of activity with a high peak power followed by idle periods of low power. If one has a choice between a 20% reduction in the energy per cycle by reducing the idle power and a 20% reduction in the energy per cycle by reducing the active power, the average power is reduced by 20% in both cases. If the battery capacity were constant as is commonly assumed, one would expect that the battery life would increase by a factor of $1/(1-20\%) = 1.25$ for both cases. But because the capacity is determined by the peak power, the battery life will be increased more by reducing the active power than by reducing idle power. Not only will the average power be reduced but the capacity available will be increased. Hence, once all the subsystems that can be put into idle mode are put into idle mode, one should focus on reducing the power during the active time rather than focus on reducing the power during the idle time.

A second example is if the designer has a choice between reducing the active time and the active power by some factor. Both will result in the same decrease in the

average power. But again, reducing the active power will result in a bigger increase in battery life when the active power is large. This means that the focus should be on reducing peak power rather than reducing duty cycle.

For a more concrete example of each method of reducing average power, consider the dynamic power profile as shown in Figure 4. The average power, P_{ave} , is equal to $(P_{active} \times t_{active} + P_{idle} \times t_{idle})/t_{cycle}$. To reduce the average power P_{ave} , the active power can be reduced (A), the idle power can be reduced (B), or the active duty cycle can be reduced (C). Table 1 shows the results from Doyle's model for the waveform of Figure 4. The waveform was simulated for three different values of initial average power, and the desired reduction in average power for each case was 20%. As expected, reducing active power (A) results in the greatest increase in battery life when the peak power is large. Reducing idle power (B) always results in the least increase in battery life. Reducing the duty cycle (C) always does better than reducing the idle power and does as well as reducing peak power only for the lowest value of peak power. But when the peak power is larger, reducing the duty cycle does not increase the battery life by as much as reducing the active power.

The column labeled “% difference from expected” refers to difference between the simulated battery life of the modification and what would be expected given the initial battery life and the factor by which the power was reduced. For example, the initial battery life of the waveform with the 300 W/kg peak power is 51 minutes. Because the average power for each of the modifications is 80% of the initial waveform, one would expect the battery

Table 1. Doyle's model results for waveform of Figure 4.

Waveform modification	Duty Cycle, $t_{active}/t_{cycle}, \%$	Peak power, W/kg	Idle power, W/kg	Average power, W/kg	Battery life, minutes	% difference from expected
none	20	300	75	120	51	--
A	20	180	75	96	83	+30
B	20	300	45	96	67	+5
C	9.3	300	75	96	68	+7
none	20	200	50	80	87	--
A	20	120	50	64	132	+21
B	20	200	30	64	117	+8
C	9.3	200	50	64	118	+9
none	20	100	25	40	202	--
A	20	60	25	32	268	+6
B	20	100	15	32	253	0
C	9.3	100	25	32	268	+6

life for them to be $51/0.8 = 64$ minutes. But this ignores the non-ideal capacity behavior. As the results show, the non-ideal capacity behavior can cause two loads with the same average power to have greatly different battery lives. The dynamic power must be considered as well as the average power.

The results of Figure 3 may explain why the advertised

battery life of the typical notebook computer is greater than what users realize in practice: Suppose the notebook manufacture is advertising an estimated battery life rather than a measured one. If the manufacturer estimates the battery life by using the battery's rated capacity and the average power of the system, then the estimate will be too large because of the loss of capacity of the battery at higher rates. While the notebook computer manufacturers reap an advantage by




Figure 4. Dynamic power profile example. Modifications A, B, and C reduce the average power.

advertising a longer battery life than is achievable in practice, obviously a motive to be considered, they may simply be using the rated battery capacity rather than the capacity available at the notebook's peak power.

The results shown in Figure 3 and from other simulations with Doyle's model (not shown due to space limitations) show that recovery is a much smaller effect than loss of capacity for loads that would be typical of mobile computing. This has two consequences. First, models of battery behavior under continuous loads can be used to estimate battery life. Second, continuous discharges are sufficient for measuring the effect of a change to lower power so long as the energy consumed while the system is idle is accounted for.

5. Conclusions

Because the battery is a key factor in the overall system weight and volume, its characteristics must be carefully considered. Non-ideal battery properties can lead to misestimates of battery life. Models that capture the non-ideal behavior are necessary both for wearable computer design and for software control of power management and power-performance trade-offs [8].

In summary, this paper has shown the following:

- Battery capacity will vary with load power.
- Peak power is a better indicator of battery capacity than average power. Estimating battery life using average power can be overly optimistic if peak power is large.
- Total system power must be considered. Power-performance trade-offs made by examining a subsystem in isolation may not lead to an increase in the computations per battery life because total peak power is ignored.
- Peak power should be reduced wherever possible, which means background operations should be performed serially rather than concurrently. Serial operation is better than concurrent operation when each consumes roughly the same energy.
- Reducing active energy is more important than reducing idle energy.
- Continuous behavior can be used to estimate intermittent behavior.

Because of non-ideal battery behavior, reducing average power or energy per operation may not increase the amount of computation completed in a battery life. Battery behavior must be considered to properly make decisions about

low power operation in wearable computing.

Acknowledgments

The authors are affiliated with the Wearable Computing Laboratory at the Institute for Complex Engineered Systems at Carnegie Mellon University. The Wearable Computing Lab receives funding from the National Science Foundation and DARPA. The authors are grateful to Marc Doyle for allowing access to his battery model.

References

- [1] M. Doyle and J. Newman, "The Use of Mathematical Modeling in the Design of Lithium/Polymer Battery Systems," *Electrochimica Acta*, vol. 40, no. 13-14; 1995; pp. 2191-2196.
- [2] M. Doyle, J. Newman, and J. Reimers, "A quick method for measuring the capacity versus the discharge rate for a dual lithium-ion insertion cell undergoing cycling," *Journal of Power Sources*; vol. 52, no. 2; December 1994; pp. 211-216.
- [3] S. Hageman, "Simple PSpice models let you simulate common battery types," *EDN*, October 28, 1993; pp. 117-132.
- [4] M. Horowitz, T. Indermaur, and R. Gonzalalez, "Low-Power Digital Design," *Proceedings of the 1994 Symposium on Low Power Electronics*, October 1994, pp. 8-11.
- [5] E. Hyman, W. C. Spindeler, and J. F. Fatula, "Phenomenological Discharge Voltage Model for Lead-Acid Batteries," *Proceedings of the AIChE Meeting*, November 1986, pp. 78-86.
- [6] D. Linden, *Handbook of Batteries*, 2nd ed. New York: McGraw-Hill, Inc. 1995.
- [7] J. Manwell and J. McGowan, "Lead Acid Battery Storage Model for Hybrid Energy Systems," *Solar Energy* Vol. 50, No. 5, pp. 399-405.
- [8] T. Martin and D. Siewiorek, "The Impact of Battery Capacity and Memory Bandwidth on CPU Speed-Setting: A Case Study," *Proceedings of the 1999 International Symposium on Low Power Electronics and Design*, to appear.
- [9] M. Weiser, B. Welch, A. Demers, and S. Shenker, "Scheduling for Reduced CPU Energy," *Proceedings of the 1st USENIX Symposium on Operating Systems Design and Implementation*, November 1994, pp. 13-23.