# Extending Lifetime of Portable Systems by Battery Scheduling

L. Benini[◇] G. Castelli[‡] A. Macii[‡] E. Macii[‡] M. Poncino[‡] R. Scarsi[‡]

[◇] Università di Bologna
Bologna, ITALY 40136

[‡] Politecnico di Torino
Torino, ITALY 10129

**Abstract**

*Multi-battery power supplies are becoming popular in electronic appliances of the latest generations, due to economical and manufacturing constraints. Unfortunately, a partitioned battery subsystem is not able to deliver the same amount of charge as a monolithic battery with the same total capacity. In this paper, we define the concept of battery scheduling, we investigate policies for solving the problem of optimal charge delivery, and we study the relationship of such policies with different configurations of the battery subsystem. Results, obtained for different workloads, demonstrate that the choice of the proper scheduling can make, in the best case, system lifetime as close as 1% of that guaranteed by a monolithic battery of equal capacity.*

## 1 Introduction

Many portable electronic systems of the latest generations utilize power supplies made of multiple batteries. This is done, primarily, to increase user's flexibility in exploiting the existing trade-off between weight and capacity of the battery subsystem. As an example, it is common practice replacing the floppy-disk drive of a note-book computer with a second battery pack before starting a long, intercontinental flight, while only a single battery is used when a short, domestic flight is taken.

Flexibility in modifying the configuration of the battery subsystem according to the user's specific needs is not the only motivation for adopting a partitioned power supply. Fabrication constraints for both the battery cells (custom-designed battery packs with ad-hoc size and capacity are used only in high-end electronic appliances, low-cost products normally utilize batteries with standard shape and capacity) and the application (battery packs need to adapt to shape and size of the case) are equally important factors that may hamper the choice of a monolithic battery cell.

Supporting multi-battery power supplies is thus becoming a major requirement for current electronic products, and significant advances in the technologies that enable the realization of multi-battery devices have been achieved very recently [1]. Among others, smart batteries (i.e., batteries equipped with specialized hardware that provides state-of-charge information under software control), dedicated battery-to-bus interfaces (e.g., Intel SMBUS [2], an ACPI-compliant [3], two-wire interface similar to Philips $I^2C$ [4] that enables the exchange of information between smart batteries and the system), and battery selector circuits (e.g., the *Intelligent Charger* by $O_2$Micro [5] and the *General-Purpose Smart Battery Charger/Selector IC* by Mitsubishi Semiconductors [6]) play an increasingly important role in the multi-battery electronics scenario.

Unfortunately, the choice of distributing the total amount of capacity to different battery packs has some penalizing effects on the total charge that can be extracted from the battery subsystem. This is due to two reasons: First, a selector circuit is required to allow a smooth, transparent transition of the system from an almost discharged battery to the next one. Although low-power techniques are normally used to design this type of circuits, some overhead in consumption has to be expected. Second, more important, under a fixed system workload, the lifetime guaranteed by a battery does not scale linearly with the capacity. For example, if a monolithic battery of capacity $C$ guarantees a lifetime of $LT$ time units, two batteries of capacity $C/2$ which are discharged one after the other ensure a lifetime smaller than $LT$ [7]. As additional side-effect, several small batteries tend to weigh more than a monolithic battery with the same total capacity, due to the additional packaging overhead and replication of the battery control circuitry which is aboard each pack.

Weight increase can only be faced by battery manufacturers; on the other hand, there are ways to address the problem of decreased usable battery capacity by adopting smart policies for charge extraction from the battery subsystem.

In [8], a set of battery-driven dynamic power management schemes, some of which are well-suited to multi-battery power supplies, have been presented. They help in maximizing battery usage by properly modifying the current load, and are all based on the fact that the amount of charge that can be extracted from a battery depends on how the current is drawn over time [9, 10, 11, 12, 13]. The idea was to utilize the various battery packs in an alternate fashion, instead of following a strict, sequential scheme, as it is done in most appliances now available on the market. The rationale behind this choice was that electro-chemical cells can recover some amount of deliverable charge if they are allowed to rest after a period of high-current discharge. Results have shown that, for some of the policies we have experimented with, battery lifetime increased by 10-15% with respect to the case of sequential battery discharge.

In this paper, the idea of multi-battery management is extended. In particular, we introduce the concept of *battery scheduling*, and we explore the impact of different scheduling policies on the achievable lifetime increase.

More formally, we solve the following problem: Given a partitioned battery subsystem (i.e., $N$ battery packs, each of which is characterized by its own nominal capacity), and fixed the system workload, find the optimal scheduling policy that maximizes the total lifetime for the given workload.

We have investigated three classes of scheduling algorithms with increasing generality, as well as complexity of the required selector circuitry, and we have performed extensive experimentation on power supply subsystems with different configurations and containing up to four battery packs of varying capacities. This choice was made in accordance with the SMBUS specification [2], which states that partitioning a battery into more than four packs may not payoff, due to the intrinsic management overhead this solution does introduce. Three current waveforms with different characteristics and profiles have been used as system workloads. The results we have obtained demonstrate that scheduling schemes allow to fully recover the loss in deliverable charge caused by the adoption of a partitioned battery supply subsystem. In fact, the difference in lifetime with respect to the monolithic case can be as small as 1%.
The remainder of this article is organized as follows. In Section 2 we review two of the macroscopic phenomena that differentiate a battery from an ideal power supply and that are at the basis of the scheduling algorithms we propose to optimize charge extraction from the partitioned battery power supply. Section 3 describes in details the three classes of scheduling algorithms we have devised. Experimental results are collected and discussed in Section 4. Finally, Section 5 concludes the paper.

## 2 Battery Properties

From the system designer's point of view, two are the physical properties of interest in a battery: Output voltage and battery capacity. In an ideal battery, the voltage is constant over a complete discharge cycle, and it drops to zero when the battery is fully discharged. In practice, however, voltage decreases as the time of discharge increases. As a matter of fact, a battery is considered exhausted when its output voltage falls below a given voltage threshold (e.g., 80% of the nominal voltage). This behavior motivates the adoption of DC-DC converters for voltage stabilization when batteries are used to power up digital systems.
Beside this intuitive and very well known fact, there are two additional factors that differentiate real batteries from ideal power supplies and that are at the basis of the battery scheduling policies we discuss in this paper:

- *The capacity of a battery depends on the discharge current.* At higher currents, a battery is less efficient in converting its chemically stored energy into available electrical energy. For increasing load currents, the battery capacity progressively deviates from the nominal value.
- *A battery can recover some of its deliverable charge when it is given some rest.* Due to electro-chemical phenomena, battery cells can deliver a larger amount of charge if periods of discharge are interleaved with periods in which no current is drawn.

The second property was exploited in [8] to develop power management policies for a two-battery system in which the two batteries alternate in providing current to the load. In this way, the battery temporarily disconnected from the load can rest, while the other one powers the system.

## 3 Battery Scheduling

Appliances powered by partitioned battery subsystems will never be able to entirely exploit the available amount of charge, as if the battery cell would be monolithic. At the same time, data point out that a smart choice of the policy adopted to extract the charge can help in reducing the efficiency gap between monolithic and partitioned power supplies [8].
The problem we address in this paper is that of deciding, given a multi-battery system, which battery pack should be connected to the load, at any point in time during system operation, in order to make the overall lifetime approaching the best case (i.e., monolithic battery). In other words, we investigate solutions to the *battery scheduling* problem.
Since factors like configuration of the battery subsystem (i.e., number and capacity of the battery packs in the power supply) and profile of the current drawn by the load (i.e., the workload) may have a substantial impact on the performance of the scheduling strategy, we introduce three classes of algorithms, characterized by increasing generality and complexity of the supporting circuitry. The quality of the results they can provide will be assessed experimentally in Section 4.

### 3.1 Serial Scheduling

The first class, that we call *serial scheduling*, simply discharges all the batteries one after the other. In other words, a battery is disconnected from the load only when it is empty. This policy offers a single degree of freedom in its implementation: The order in which the batteries are discharged. If the battery subsystem consists of $N$ packs, then there is a total on $N!$ possible serial schedulings. Clearly, changing the order in which the batteries are discharged matters only if all batteries do not have the same capacity and the workload is not constant over time.
Although it is known that serial scheduling is not very effective, we consider it for our experiments because it is the policy which is currently adopted by the majority of electronic products powered by multi-battery systems; therefore, it provides a good term of comparison for the performance of the other types of scheduling we introduce next.

### 3.2 Static Scheduling

Policies of the second class, called *static scheduling*, discharge the batteries following a round-robin scheme. Each battery stays connected to the load for a fixed amount of time, then it is disconnected. It will be used again after all other batteries have been discharged for their assigned period of time. Policies of this class are thus characterized by two parameters: The order in which batteries are discharged and the time period (or time slice) during which each battery is connected to the load. Assuming $N$ battery packs and $M$ time slices each pack can be assigned, there is a total of $N! \cdot M^N$ possible static schedulings that should be explored. Some pruning of the search space must then be adopted, the simplest being the choice of a single time slice for all batteries. In this case, the number of schedulings to be explored reduces to $N!$.

### 3.3 Dynamic Scheduling

The last set of policies that we have investigated uses specific battery information to drive battery selection. This information can be related to the physical state of the batteries, and can be expressed by the output voltage, the state of charge (SOC), or the elapsed discharge time. Clearly, the chosen representative quantity should be observable at the battery interface to allow the practical implementation of the policy. We classify this class of policies as *dynamic scheduling*. Dynamic scheduling no longer assumes a unique time slice being assigned to each battery pack. Rather, each battery is discharged for a different amount of time depending on the quantity that is observed.

We note that, in principle, there is no guarantee that dynamic scheduling outperforms the static approach. However, the former holds the advantage of adapting the rest time which is given to each battery pack to the actual conditions of the various batteries. Therefore, it is likely to provide a finer tuning of the times in which charge recovery can take place, and may then be more effective for battery subsystems containing packs of different capacities.

## 4 Results

### 4.1 Experimental Setup

To run our experiments with the scheduling policies of Section 3, we have defined three pairs of workloads with different profiles.

The first two pairs (*CC* and *SW*) are artificial workloads; the current profiles have the same shape, but they differ for the value of the average current and/or for the absolute current levels:

- *CC*: Constant current load of $0.5A$ for $CC_a$ and of $1.0A$ for $CC_b$.
- *SW*: Square wave with 50% duty-cycle, average current value of $0.5A$ and current levels of $(0.3A, 0.7A)$ for $SW_a$, average current value of $1.0A$ and current levels of $(0.8A, 1.2A)$ for $SW_b$.

The third pair of workloads ($RL$) refers to a real-life example; current profiles have been obtained by monitoring the activity of a Personal Digital Assistant (PDA) over a few hours of operation [13]. For this example, the temporal waveforms of the current drawn by the system from the power supply do not have the same shape; the mean value of workload named $RL_a$ is lower than that of $RL_b$.

Regarding the battery subsystem, we have chosen to experiment with four different configurations (denoted, in the sequel, by *BS1*–*BS4*). All of them respect the constraints set by the SMBUS specification [2] (e.g., at most four battery packs), and correspond to organizations that are found, in practice, in most electronic products.

The sum of the capacity of all packs in the partitioned power supply is chosen to be the same for all configurations ($C_{\mathrm{TOT}} = 1.35Ahr$). This is also the capacity of the monolithic battery we use as the baseline for comparison of all the schedulings.

All the battery packs in the power supply have the same nominal output voltage of $4.1V$, independently of their capacities. In the experiments, a battery pack is regarded as exhausted, therefore no longer selectable by the scheduler, after its output voltage drops below $3.3V$.

The four battery systems we consider are the following:

- *BS1*: Two battery packs of capacity $C_{\mathrm{TOT}}/2$ each.
- *BS2*: Four battery packs of capacity $C_{\mathrm{TOT}}/4$ each.
- *BS3*: Four battery packs, one of which has *back-up* purposes (thus has a very small capacity, i.e., $C_{\mathrm{backup}} = C_{\mathrm{TOT}}/20$), and the remaining three have the same capacity of $(C_{\mathrm{TOT}} - C_{\mathrm{backup}})/3$.
- *BS4*: Four battery packs, all with different capacities: $C_{\mathrm{TOT}}/3$, $C_{\mathrm{TOT}}/4$, $C_{\mathrm{TOT}}/5$, and $13 \cdot C_{\mathrm{TOT}}/60$.

The model used for battery lifetime estimation is the event-driven model described in [13].

### 4.2 Scheduling Experiments

#### 4.2.1 Serial Scheduling

In serial scheduling there is a total of $N!$ experiments to run. Since the chosen battery subsystems have at most four batteries, exhaustive exploration is feasible.

Table 1 collects the results regarding system lifetime, which is expressed in seconds. Row *Mono* refers to the monolithic power supply, while rows *BS1*–*BS4* refer to the partitioned cases. Finally, columns indicate the workload. Notice that, for space reasons, for each combination of power supply configuration and workload, only the best result is provided (out of the $N!$ that have been explored).

| | $CC_a$ | $CC_b$ | $SW_a$ | $SW_b$ | $RL_a$ | $RL_b$ |
|---|---|---|---|---|---|---|
| Mono | 8744 | 3878 | 8617 | 3878 | 10993 | 10589 |
| BS1 | 8129 | 3352 | 7895 | 3240 | 10308 | 9945 |
| BS2 | 6839 | 2794 | 6399 | 2675 | 8793 | 8578 |
| BS3 | 6835 | 2792 | 6474 | 2659 | 8817 | 8584 |
| BS4 | 6820 | 2785 | 6478 | 2653 | 8793 | 8567 |

Table 1: Results: Serial Scheduling.

The data clearly show the efficiency decrease that a partitioned battery determines on system lifetime, regardless of the workload, if serial scheduling is used. The effect is more visible for power supplies with more than two battery packs (lifetime reduction is in the range [18.9-31.6]%), while it is more under control for *BS1* (the range here is [6.08 - 16.45]%).

#### 4.2.2 Static Scheduling

The search space for static scheduling is much wider than in the case of serial scheduling, since it has two dimensions: The order in which battery packs are discharged and the duration of the time slice assigned to each pack. Exhaustive exploration is thus infeasible, and some pruning criteria must be introduced. We have chosen those described in Section 3, that assume each battery pack is not given any choice in the selection of its own time slice. More specifically, we have first tried the solution where a time slice of 10 seconds is used for all battery packs.

Table 2 displays the results. As in the case of serial scheduling, for each combination of power supply configuration and workload, only the best result is shown (out of the $N!$ that have been explored). Data show sensible improvements over serial scheduling. The gap with respect to the lifetime provided by the monolithic battery is now in the range [2.8-13.9]% for *BS1* and [10.3-29.5]% for the four-battery systems.

| | $CC_a$ | $CC_b$ | $SW_a$ | $SW_b$ | $RL_a$ | $RL_b$ |
|---|---|---|---|---|---|---|
| Mono | 8744 | 3878 | 8617 | 3878 | 10993 | 10589 |
| BS1 | 8468 | 3669 | 7978 | 3339 | 10582 | 10289 |
| BS2 | 7758 | 2958 | 6878 | 2735 | 9683 | 9499 |
| BS3 | 7709 | 2938 | 6958 | 2798 | 9706 | 9410 |
| BS4 | 7474 | 2859 | 7268 | 2790 | 9443 | 9262 |

Table 2: Results: Static Scheduling ($10s$ Time Slice).

The choice of a time slice of 10 seconds for the previous experiment was totally arbitrary. Therefore, in order to acquire some further insight on how the choice of the frequency at which the various battery packs must be connected/disconnected from the load, we have run the same experiment with a shorter time slice, namely 2 seconds. The results of this exploration are collected in Table 3, and they show an even more dramatic improvement over serial scheduling than static scheduling with a 10 seconds time slice. In fact, the distance from the monolithic case is confined in the range [1.1-7.8]% for *BS1* and [4.3-23.1]% for the four-battery power supplies.

| | $CC_a$ | $CC_b$ | $SW_a$ | $SW_b$ | $RL_a$ | $RL_b$ |
|---|---|---|---|---|---|---|
| Mono | 8744 | 3878 | 8617 | 3878 | 10993 | 10589 |
| BS1 | 8626 | 3834 | 8177 | 3576 | 10853 | 10438 |
| BS2 | 8366 | 3554 | 7760 | 3233 | 10477 | 10119 |
| BS3 | 8256 | 3616 | 7739 | 3213 | 10383 | 9664 |
| BS4 | 7850 | 3124 | 7579 | 2980 | 9967 | 9700 |

Table 3: Results: Static Scheduling ($2s$ Time Slices).

This consistent lifetime increase over all battery systems and load profiles suggests that further decreasing the time slice value may improve lifetime accordingly. We have therefore completed this analysis by determining the lifetimes for the two battery configuration as a function of the inverse of the time slice (i.e., the switching frequency, $f_{sw}$). The plot of Figure 1 shows the corresponding results for the case of a constant current of $0.5A$ (i.e., workload $CC_a$). When $f_{sw}$ is very low, the batteries are discharged in sequence, and lifetime is minimum. This corresponds to serial discharging. As $f_{sw}$ increases, lifetime increases as well, and it asymptotically tends to the lifetime provided by a monolithic battery with double capacity (i.e., $1.35Ahr$ in the plot). This behavior is consistently observed also for all the four battery subsystems.
The bottom line of this exploration is that, in principle, a multi-battery system using a static scheduling policy where batteries are switched with an ideally infinite frequency will perform as if an equivalent (i.e., with equal total capacity) monolithic battery is attached to the load. No better solution can be found, no matter what type of current load or composition of the battery subsystem.

However, to support a static scheduling scheme with a high switching frequency, a high-switching battery selector is required. Commercially available battery selectors are used for the serial discharging policies of Section 3.1, that is, they connect a load to a battery with switching periods in the order of the lifetime of a single battery. The usability of these devices in the context of high-frequency static battery scheduling still needs to be investigated and, at this point in time, remains an open issue.

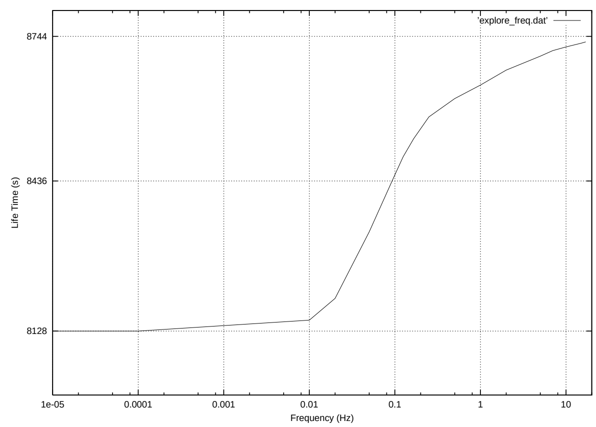


Figure 1: Battery Life-Time vs. $f_{sw}$.

#### 4.2.3 Dynamic Scheduling

The distinctive feature of dynamic scheduling is that it tries to adapt the time slice given to a battery pack to the actual SOC. In other words, the currently active battery is disconnected from the load as soon as its output voltage drops below a predefined threshold. Such threshold is specified as a percentage of the output voltage the battery pack holds at the time it is connected to the load and starts delivering the current.
After a battery pack is disconnected, it must be given some rest time in order to let it recover some of its charge, and then it must be rescheduled. A strategy is then needed for choosing the next battery pack to be connected to the load. We have experimented with the following policies:

- $V_{max}$: Select the battery pack with highest SOC.
- $V_{min}$: Select the battery pack with lowest SOC.
- $T_{max}$: Select the battery pack that has been unused for the longest time.
- $T_{min}$: Select the battery that has been unused for the shortest time.

If more than one battery is eligible for selection, the initial capacity of the pack is used as first-level tie-breaker (i.e., the battery with larger capacity is chosen), while random selection is used if further ties occur.
Independently of the chosen policy, once the next battery pack is selected, a check is performed to make sure that a minimum rest time has elapsed. If this is not the case, the next battery in the schedule is selected, and the one that missed its turn will be reconsidered when a new selection will take place.

We observe that policy $T_{max}$ is actually very similar to static scheduling, with the difference that disconnection from the load of the active battery pack is voltage-driven rather than time-driven.

The choice of the threshold is critical for the performance of dynamic scheduling. Intuitively, it should be very close to 100%. In fact, this would imply that each pack stays connected to the load for a very short time, and thus the switching frequency automatically becomes very high. However, as mentioned for the case of static scheduling, switching too quickly from one battery pack to another may not be advisable. In addition, the selection policies described above automatically prevent the scheduling of a battery if a rest time greater than a predefined value has not elapsed. Dead-lock conditions may thus occur (i.e., none of the batteries can be selected) if switching between batteries takes place too often. For our experiments, a threshold of 95% has been used for all the battery packs. This value was chosen because it was the highest value in the range [90-99]% that avoided the occurrence of dead-lock situations for all the workloads and all the power supply configurations. Table 4 reports all the results.

The data tell that policy $V_{max}$ is generally better than the others, except for two cases where $T_{max}$ is more effective. In general, we have that the gap between lifetime of the monolithic power supply and *BS1* is in the range of [2.5-5.1]%, and in the range of [2.5-20.3] for *BS2*-*BS4*. By comparing these values to those provided by static scheduling with $2s$ time slices, we observe that in some cases the latter are better. This is happens primarily for battery systems in which all packs have the same capacity; in these cases, switching at constant times helps in optimizing charge extraction. On the other hand, dynamic scheduling is more flexible and it is thus more appropriate in providing the best performance in the case of etherogeneous battery subsystems and highly irregular current profiles, as those occurring in power-managed applications (e.g., *BS4*).

We also note that, for *BS1*, all dynamic policies yield exactly the same scheduling. This is due to the fact that a battery pack cannot be scheduled as the next battery of itself, since a minimum rest time is imposed.

| | | $CC_a$ | $CC_b$ | $SW_a$ | $SW_b$ | $RL_a$ | $RL_b$ |
|---|---|---|---|---|---|---|---|
| Mono | | 8744 | 3878 | 8617 | 3878 | 10993 | 10589 |
| BS1 | $V_{max}$ | 8410 | 3781 | 8173 | 3690 | 10532 | 10214 |
| | $V_{min}$ | 8410 | 3781 | 8173 | 3690 | 10532 | 10214 |
| | $T_{max}$ | 8410 | 3781 | 8173 | 3690 | 10532 | 10214 |
| | $T_{min}$ | 8410 | 3781 | 8173 | 3690 | 10532 | 10214 |
| BS2 | $V_{max}$ | 8242 | 3768 | 7820 | 3651 | 10341 | 9962 |
| | $V_{min}$ | 7716 | 3262 | 7669 | 3170 | 9882 | 9502 |
| | $T_{max}$ | 8242 | 3780 | 7939 | 3390 | 10311 | 9962 |
| | $T_{min}$ | 7877 | 3306 | 7695 | 3218 | 9982 | 9581 |
| BS3 | $V_{max}$ | 8224 | 3676 | 7896 | 3550 | 10366 | 9959 |
| | $V_{min}$ | 7564 | 3146 | 7337 | 3050 | 9660 | 9315 |
| | $T_{max}$ | 8082 | 3596 | 7799 | 3459 | 10187 | 9803 |
| | $T_{min}$ | 7853 | 3259 | 7498 | 3216 | 9863 | 9539 |
| BS4 | $V_{max}$ | 8125 | 3677 | 7919 | 3517 | 10287 | 9902 |
| | $V_{min}$ | 7737 | 3148 | 7570 | 3079 | 9891 | 9464 |
| | $T_{max}$ | 8017 | 3383 | 7636 | 3358 | 10135 | 9718 |
| | $T_{min}$ | 7974 | 3113 | 7499 | 3090 | 9891 | 9671 |

Table 4: Results: Dynamic Scheduling.

## 5 Conclusions

Manufacturing and cost constraints force the producers of modern electronic appliances to adopt multi-pack batteries as power supplies. It has been shown that a partitioned battery in which the components are connected serially to the load (i.e., a new pack comes in use only after the active one is exhausted) is not able to guarantee the same system lifetime as a monolithic battery of equal total capacity.

In this paper, we have introduced the concept of battery scheduling (i.e., the policy used to connect and disconnect the packs of a multi-battery power supply to the current load during system operation); we have proposed a number of scheduling algorithms, characterized by increasing efficiency and complexity of the selector circuitry which is in charge of connecting the various batteries to the system, and we have performed an exhaustive experimental exploration.

On average, we have observed that static and dynamic scheduling allow system lifetime to increase significantly with respect to serial scheduling, and to close the gap between partitioned and monolithic power supplies down to 1% in the best case.

## References


[1] Q. Wu, Q. Qiu, M. Pedram, "An Interleaved Dual-Battery Power Supply for Battery Operated Electronics" *ASPDAC-00*, pp. 387-390, Jan. 2000.

[2] Intel, *System Management Bus*, http://www.smbus.org.

[3] Compaq, Intel, Microsoft, Phoenix, Toshiba, *Advanced Configuration and Power Interface*, http://www.teleport.com/ acpi.

[4] Philips, *Inter-IC($I^2C$) Bus*, http://www-us.semiconductors.com/i2c.

[5] $O_2$Micro, *Intelligent Charger*, http://www.o2micro.com/products.

[6] Mitsubishi Semiconductors, *General-Purpose Smart Battery Charger/Selector IC*, http://www.mitsubishichips.com/products/assps/assp.htm.

[7] T. L. Martin, *Balancing Batteries, Power and Performance: System Issues in CPU Speed-Setting for Mobile Computing*, Ph.D. Thesis, CMU, Dept. of ECE, Aug. 1999.

[8] L. Benini, G. Castelli, A. Macii, E. Macii, R. Scarsi, "Battery-Driven Dynamic Power Management of Portable Systems," *ISSS-00*, pp. 25-30, Madrid, Spain, Sep. 2000.

[9] M. Doyle, T. F. Fuller, J. Newmann, "Modeling of Galvanostatic Charge and Discharge of the Lithium/Polymer/Insertion Cell," *Journal of the Electrochemical Society*, Vol. 140, No. 6, pp. 1526-1533, Jun. 1993.

[10] T. Martin, D. Sewiorek, "A Power Metric for Mobile Systems," *ISLPED-96*, pp. 37-42, Monterey, CA, Aug. 1996.

[11] M. Pedram, Q. Wu, "Battery-Powered Digital CMOS Design," *DATE-99*, pp. 72-76, Munich, Germany, Mar. 1999.

[12] M. Pedram, Q. Wu, "Design Considerations for Battery Powered Electronics," *DAC-36*, pp. 861-866, New Orleans, LA, Jun. 1999.

[13] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino, R. Scarsi, "A Discrete-Time Battery Model for High-Level Power Estimation," *DATE-00*, pp. 35-39, Paris, France, Mar. 2000.