

Battery Life Estimation of Mobile Embedded Systems *

Debashis Panigrahi †, Carla Chiasserini ‡, Sujit Dey †, Ramesh Rao †,
Anand Raghunathan § and Kanishka Lahiri †

†Dept. of Electrical and Computer Engg., University of California, San Diego, CA.

‡Dipartimento di Eletttronica, Politecnico di Torino, Italy.

§C & C Research Labs, NEC USA, Princeton, NJ.

Abstract

Since battery life directly impacts the extent and duration of mobility, one of the key considerations in the design of a mobile embedded system should be to maximize the energy delivered by the battery, and hence the battery lifetime. To facilitate exploration of alternative implementations for a mobile embedded system, in this paper we address the issue of developing a fast and accurate battery model, and providing a framework for battery life estimation of Hardware/Software (HW/SW) embedded systems.

We introduce a stochastic model of a battery, which can simultaneously model two key phenomena affecting the battery life and the amount of energy that can be delivered by the battery: the Rate Capacity effect and the Recovery effect. We model the battery behavior mathematically in terms of parameters that can be related to physical characteristics of the electro-chemical cell. We show how this model can be used for battery life estimation of a HW/SW embedded system, by calculating battery discharge demand waveforms using a power co-estimation technique. Based on the discharge demand, the battery model estimates the battery lifetime as well as the delivered energy. Application of the battery life estimation methodology to three system implementations of an example TCP/IP network interface subsystem demonstrate that different system architectures can have significantly different delivered energy and battery lifetimes.

1 Introduction

As the need for mobile computation and communication increases, there is a strong demand for design of Hardware/Software (HW/SW) Embedded Systems for mobile applications. Maximizing the amount of energy that can be delivered by the battery, and hence the battery life, is one of the most important design considerations for a mobile embedded system, since it directly impacts the extent and duration of the system's mobility. To enable exploration of alternative implementations for a mobile system, it is critical to develop fast and accurate battery life estimation techniques for embedded systems. In this paper, we focus on developing such a battery model, and provide a framework for battery-life estimation of HW/SW embedded systems.

Previous research on low power design techniques [1, 2], tries to minimize average power consumption either by reducing the average current drawn by a circuit keeping the supply voltage fixed or by scaling the supply voltage statically or dynamically. However, as shown in this paper, designing to minimize average power consumption does not necessarily lead to optimum battery lifetime. Additionally, the above techniques assume that the battery subsystem is an

ideal source of energy which stores or delivers a fixed amount of energy at a constant output voltage. In reality, it may not be possible to extract the energy stored in the battery to the full extent as the energy delivered by a battery greatly depends on the current discharge profile. Hence, accurate battery models are needed to specifically target the battery life and the amount of energy that can be delivered by a battery in the design of a mobile system.

The lifetime of a battery, and the energy delivered by a battery, for a given embedded system strongly depend on the current discharge profile. If a current of magnitude greater than the rated current of the battery is discharged, then the efficiency of the battery (ratio of the delivered energy and the energy stored in the battery) decreases, in other words, the battery lifetime decreases [3, 10]. This effect is termed as the *Rate Capacity Effect*. Additionally, if a battery is discharged for short time intervals followed by idle periods, significant improvements in the delivered energy seem possible [11, 13]. During the idle periods, also called *Relaxation Times*, the battery can partially recover the capacity lost in previous discharges. We call this effect as the *Recovery Effect*.

An accurate battery model, representing fine-grained electro-chemical phenomenon of cell discharge using Partial Differential Equations (PDE), was presented in [14]. However, it takes prohibitively long (days) to estimate the battery lifetime for a given discharge demand of a system. Hence, the PDE models cannot be used for design space exploration. Some SPICE level models of battery have been developed [6, 7], which are faster than the PDE model. However, the SPICE models can take into account the effect of *Rate Capacity* only. Based on the *Rate Capacity* effect, a system-level battery estimation methodology was proposed in [4, 5]. Recently, a Discrete-Time battery model was proposed for high-level power estimation [8]. Though it is faster than the previous models, it does not consider the *Recovery* effect.

In this paper, we describe a stochastic battery model, taking into account both the *Recovery* effect and the *Rate Capacity* effect. The proposed model is fast as it is based on stochastic simulation. Also, by incorporating both *Recovery* and *Rate Capacity* effects, it represents physical battery phenomena more accurately than the previous fast models. We also show how this model can be used for estimating the battery lifetime and the energy delivered by the battery for a HW/SW system, by calculating battery discharge demand waveforms using a power co-estimation technique [9]. Based on the discharge demand, the battery model estimates the battery lifetime as well as the delivered energy. Finally, we demonstrate how this framework can be used for system level exploration using a TCP/IP network interface subsystem. The results indicate that the energy delivered by the battery and the lifetime of the battery can be significantly in-

*This work was supported by NEC USA Inc., and by the California Micro Program

creased through architectural explorations.

The rest of the paper is organized as follows. Section 2 motivates the need for an accurate battery life estimation methodology by illustrating that the battery life and the energy delivered by the battery can be affected significantly by tradeoffs at the system level. Section 3 provides background on the physical phenomena inside a battery, which affect the the battery lifetime as well as the delivered energy. The proposed battery model is described in section 4. The methodology used to calculate current waveforms is described in section 5. In section 6, we demonstrate how the battery life estimation methodology can be used to evaluate alternate implementations in the design of Battery Efficient Systems. Section 7 concludes the paper and explores future research.

2 Motivation : Exploration For Battery Efficient Architectures

In this section, we present the effect of system architectures on the delivered energy and the lifetime of battery. Our investigations motivate the need for fast and accurate battery life estimation techniques that can be used for system level exploration.

We analyze the performance of an example TCP/IP network interface subsystem with respect to the delivered energy and the lifetime of the battery. The subsystem consists of the part of the TCP/IP protocol stack performing the checksum computation (Figure 1). *Create_Packet* receives a packet, stores it in a shared memory, and enqueues its starting address. *IP_Chk* periodically dequeues packet information, erases specific bits of the packet in memory, and coordinates with *Checksum* to verify the checksum value. Figure 1 shows a candidate architecture for the TCP/IP system where *Create_Packet* and *Packet_Queue* are software tasks mapped to a SPARC processor, while *IP_Chk* and *Checksum* are each mapped to dedicated hardware. Packet bits are stored in a single shared memory accessed through a common system bus.




Figure 1: Architecture of TCP/IP Network Subsystem

We study the effect of alternate ways of packet processing by the system on the battery lifetime as well as the energy delivered by the battery. In the first implementation (Sys_A), the packets are processed sequentially as shown in Figure 2(a). To estimate the battery lifetime and the energy delivered by the battery of the implementation, the current profiles of each component of the system need to be calculated. Figures 2(b) and (c) show the current demands (in mA) for some of the components of the system for Sys_A plotted over time

(in ms), calculated using the methodology described in Section 5. The cumulative current profile for Sys_A is shown in Figure 2(d).

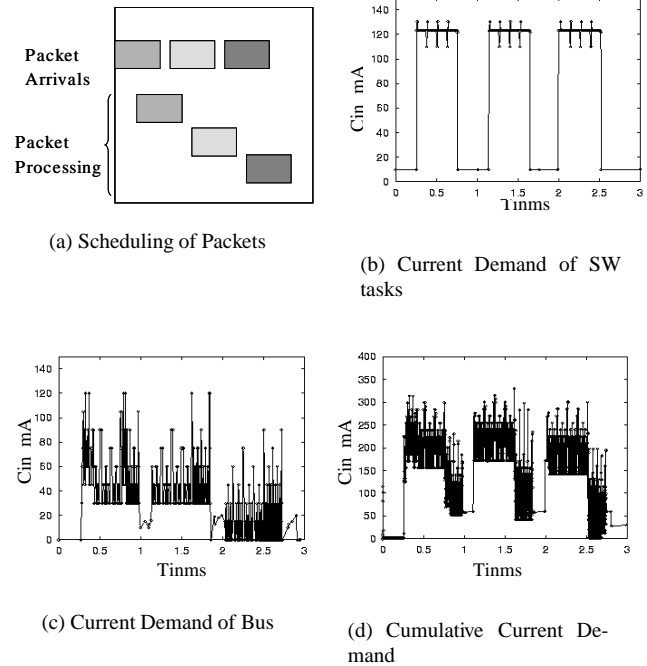


Figure 2: Sys_A : Scheduling of Packets and Current Discharge Demands

Figure 3(a) shows another way of scheduling the packets. In this implementation (Sys_B), instead of sequential processing of packets as in Sys_A, the first two packets are processed in a pipelined manner; for example, while *Checksum* is processing the first packet, *Create_Packet* starts writing the next packet to memory. The cumulative current profile for Sys_B is shown in Figure 3(b).

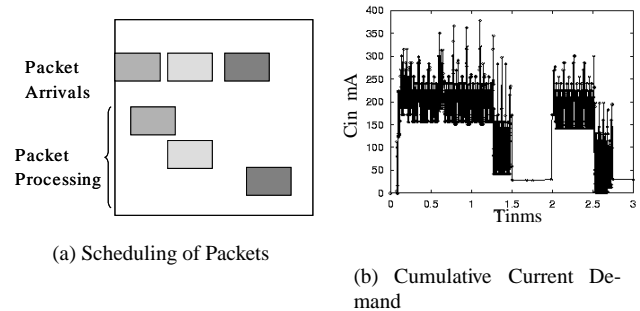


Figure 3: Sys_B : Scheduling of Packets and the Cumulative Current Demand

As shown in Table 1, the average current (and hence average power consumed) is very similar for both the alternative system implementations. Table 1 shows the battery lifetimes and the specific energy delivered by the battery, calculated using an accurate battery model [14] for both implementations, Sys_A and Sys_B. The table also reports the number of packets processed by each system before the battery used by

the systems is completely discharged. Note that, though the two discharge demands have almost equal average current requirement, the delivered specific energy and lifetime of the battery differ significantly.

Table 1: Comparison of Battery Lifetime and Delivered Energy

Discharge Demand	Average Curr. (mA)	Battery Life (ms)	Del. Spec. Energy (Wh/Kg)	Packets Processed
Sys_A	123.8	357053	15.12	119018
Sys_B	124.2	536484	18.58	178828

The results show that the battery life of a mobile embedded system can be improved significantly by system level trade-offs. For instance, as shown in Table 1, Sys_A can process 119018 packets before the battery used for the system is discharged, where as Sys_B can process 178828 packets using the same battery. The results also show that the delivered energy and the lifetime of a battery can be significantly different for current demands with the same average current requirement. Above results motivate us to develop a fast and accurate battery model that can be used in design exploration for battery optimal design of mobile embedded systems.

The next section provides a brief background on the operation of a battery, emphasizing the physical phenomena that affect the battery lifetime and the energy delivered by the battery.

3 Battery Background

A battery cell consists of an anode, a cathode, and electrolyte that separates the two electrodes and allows transfer of electrons as ions between them. During discharge, oxidation of the anode (Li for LiIon battery) produces charged ions (Li^+), which travel through the electrolyte and undergo reduction at the cathode. The reaction sites (parts of the cathode where reductions have occurred) become inactive for future discharge because of the formation of an inactive compound. *Rate Capacity Effect* (dependency of energy delivered by a battery on magnitude of discharge current) and *Recovery Effect* (recovery of charged ions near cathode) are two important phenomena that affect the delivered energy and the lifetime of a battery. A short description of the physical phenomena responsible for these effects follows.

The lifetime of a cell depends on the availability and reachability of active reaction sites in the cathode. When discharge current is low, the inactive sites (made inactive by previous cathode reactions) are distributed uniformly throughout the cathode. But, at higher discharge current, reductions occur at the outer surface of the cathode making the inner active sites inaccessible. Hence, the energy delivered (or the battery lifetime) decreases since many active sites in the cathode remain un-utilized when the battery is declared discharged.

Besides non-availability of active reaction sites in the cathode during discharge, the non-availability of charged ions (lithium ions for lithium insertion cell) can also be a factor determining the amount of energy that can be delivered by a battery [11]. Concentration of the active species (charged ions i.e. Li^+) is uniform at electrode-electrolyte interface at zero current. During discharge, the active species are consumed at the cathode-electrolyte interface, and replaced by new active species that move from electrolyte solution to cathode

through diffusion. However, as the intensity of the current increases, the concentration of active species decreases at the cathode and increases at the anode and the diffusion phenomenon is unable to compensate for the depletion of active materials near the cathode. As a result, the concentration of active species reduces near the cathode decreasing the cell voltage. However, if the cell is allowed to idle in between discharges, concentration gradient decreases because of diffusion, and charge recovery takes place at the electrode. As a result, the energy delivered by the cell, and hence the lifetime, increases. Summarizing, the amount of energy that can be delivered from a cell and the lifetime of a cell, depend on the value of the discharge current and the idle times in the discharge demand. In the next subsection, we define some notations which will be used in the rest of the paper.

3.1 Notations and Definitions

A battery cell is characterized by the open-circuit potential (V_{OC}), i.e., the initial potential of a fully charged cell under no-load conditions, and the cut-off potential (V_{cut}) at which the cell is considered discharged. Two parameters are used to represent the cell capacity: the *theoretical* and the *nominal* capacity. The former is based on the amount of energy stored in the cell and is expressed in terms of ampere-hours. The latter represents the energy that can be obtained from a cell when it is discharged at a specific constant current (called the rated current, C_{rated}). Battery data-sheets typically represent the capacity of the cell in terms of the nominal capacity.

Finally, to measure the cell discharge performance, the following two parameters are considered: *Battery Lifetime* and *Delivered Specific Energy*. *Battery Lifetime* is expressed as seconds elapsed until a fully charged cell reaches the V_{cut} voltage. *Delivered Specific Energy* is the amount of energy delivered by the cell of unit weight, i.e., expressed as watt-hour per kilogram.

In the next section, we describe our basic stochastic battery-model that captures the *Recovery Effect*, and a description of an extension of the model to incorporate the *Rate Capacity Effect*.

4 Battery Models

The fine-grained electro-chemical phenomena underlying the cell discharge are represented by the accurate model, based on PDE (Partial Differential Equations) [14], which involve a large number of parameters depending on the type of cell. The set of results that can be derived through the PDE model is limited since as the discharge current and the cut-off potential decrease, the computation time becomes exceedingly large. Hence, the accurate PDE model cannot be used for system-level exploration of a mobile embedded system. In the following section, we present a more tractable parametric model that captures the essence of the recovery mechanism.

4.1 Stochastic Battery Model

We model the battery behavior mathematically in terms of parameters that can be related to the physical characteristics of an electro-chemical cell [15, 16]. The proposed stochastic model focuses on the *Recovery Effect* that is observed when *Relaxation Times* are allowed in between discharges.

Let us consider a single cell and track the stochastic evolution of the cell from the fully charged state to the completely discharged state. We define the smallest amount of capacity that may be discharged as a *charge unit*. Each fully charged cell is assumed to have a maximum available capacity of T charge units, and a nominal capacity of N charge units. The nominal capacity, N , is much less than T in practice and represents the charge that could be extracted using a constant discharge profile. Both N and T vary for different kinds of cells and values of discharge current.

We represent the cell behavior as a discrete time transient stochastic process, that tracks the cell state of charge. Figure 4 shows a graphical representation of the process. At each time unit, the state of charge decreases from state i to the state $i - n$ if n charge units are demanded from the battery. Otherwise, if no charge units are demanded, the battery may recover from its current state of charge (i) to a higher state (greater than i). The stochastic process starts from the state of full charge ($V = V_{OC}$), denoted by N , and terminates when the absorbing state 0 ($V = V_{cut}$) is reached, or the maximum available capacity T is exhausted. In case of constant current discharge, N successive charge units are drained and the cell state goes from N to 0 in a time period equal to N time units. By allowing idle periods in between discharges, the battery can partially recover its charge during the idle times, and thus we can drain a number of charge units greater than N before reaching the state 0.




Figure 4: Stochastic process representing the cell behavior

In this model, the discharge demand is modeled by a stochastic process. Let us define q_i to be the probability that in one time unit, called slot, i charge units are demanded. Thus, starting from N , at each time slot, with probability q_i ($i > 0$), i charge units are lost and the cell state moves from state z to $z - i$ (see Figure 4). On the other hand, with probability q_0 an idle slot occurs and the cell may recover one charge unit (i.e., the cell state changes from state z to $z + 1$) or remain in the same state.

The recovery effect is represented as a decreasing exponential function of the state of charge of the battery. To more accurately model real cell behavior, the exponential decay coefficient is assumed to take different values as a function of the discharged capacity.

During the discharge process, different phases can be identified according to the recovery capability of the cell. Each phase f ($f=0, \dots, f_{max}$) starts right after d_f charge units have been drained from the cell and ends when the amount of discharged capacity reaches d_{f+1} charge units. The probability of recovering one charge unit in a time slot, conditioned on being in state j ($j=1, \dots, N-1$) and phase f is

$$p_j(f) = \begin{cases} q_0 e^{-g_N(N-j)-g_C(f)} & f = 0 \\ q_0 e^{-g_N(N-j)-g_C(f)d_f} & f = 1, \dots, f_{max} \end{cases} \quad (1)$$

where g_N and g_C are parameters that depend on the recovery capability of the battery, and q_0 is the probability of an idle slot. In particular, a small value of g_N represents a high cell conductivity (i.e., a great recovery capability of the cell), while a large g_N corresponds to a high internal resistance, (i.e., a steep discharge curve for the cell). The value of g_C is related to the cell potential drop during the discharge process, and therefore, to the discharge current. Given the recovery probability, the probability to remain in the same state of charge in an idle time while being in phase f is

$$\begin{aligned} r_j(f) &= q_0 - p_j(f) & j=1, \dots, N-1 \\ r_N(f) &= q_0. \end{aligned} \quad (2)$$

We assume that g_N is a constant, whereas g_C is a piecewise constant function of the number of charge units already drawn off the cell, that changes value in correspondence with d_f ($f = 1, \dots, f_{max}$). We have $d_0=0$ and $d_{f_{max}+1} = T$, while for d_f ($f = 1, \dots, f_{max}$) proper values are chosen according to the configuration of the battery.

One simulation step of the battery cell assuming the input discharge demand as Bernoulli arrival is shown in Figure 5.

```

Simulation_Step
inputs:      Current_State, Recovery_Probability[],
Discharge_Rate
outputs: Next_State
begin
  Generate a random number R between 0 and 1;
  If (R < Discharge_Rate) then
    Next_State := Current_State - 1;
  else if (R < Recovery_Probability[Current_State]) then
    Next_State := Current_State + 1;
  end if
end

```

Figure 5: The Basic Simulation_Step

4.2 Validation of the Stochastic Model

In this section we present a comparison between results obtained through the stochastic model and those derived from the PDE model of a dual lithium ion insertion cell. The discharge demand is assumed to be a Bernoulli process with probability q that one charge unit is required in a time slot. Note that in this discharge process the probability of an idle time (q_0) is $(1 - q)$ where as q_i ($i > 1$) is equal to zero.

The PDE model was numerically solved by using a program developed by Newman *et al.* [17]. Results relate to the first discharge cycle of the cell; thus, discharge always starts from a value of positive open-circuit potential equal to 4.3071 V. We consider that the cut-off potential is equal to 2.8 V and the current impulse duration is equal to 0.5 ms.

Results obtained from the stochastic model are derived under the following assumptions: $f_{max} = 3$, N equal to the number of impulses obtained through the PDE model under con-

stant discharge, and T equal to the number of impulses obtained through the PDE model when $q = 0.1$.

Figure 6 presents the behavior of the delivered capacity normalized to the nominal capacity versus the discharge rate (q) at which the current impulses are drained for three values of current density: $I=90, 100, \text{ and } 110 \text{ A/m}^2$. It can be seen that the curves obtained from the PDE and the stochastic models match closely.




Figure 6: Comparison of results obtained by the Stochastic Model (s) with the PDE model (e)

For these cases, we assume $g_N=0$ and vary the parameters g_C and d_f ($f=1, \dots, f_{max}$) of the stochastic model according to the considered value of current density. Following this procedure, we obtain a maximum error equal to 4% and an average error equal to 1%.

4.3 Enhancement For Deterministic Discharge Profile

To be able to use the model for battery-life estimation of SoC designs, we enhanced the model to accept any deterministic discharge demand, like the ones shown in Figures 2 and 3, instead of a stochastic discharge profile.

At each step of the simulation, we look at the input discharge demand and draw an appropriate number of charge units (or change the state of the battery). Since a battery cannot respond to instantaneous changes in current, we have calculated average current over a time period of τ , which is based on the time constant that characterizes the electrochemical phenomena. Given any current demand waveforms, we calculate the average current drawn over each time period and convert it to an appropriate number of charge units to be drawn in the battery model. We used a time constant $\tau = 0.5$ ms for our experiments in this paper.

4.4 Incorporation of Rate Capacity Effect

As we have described earlier, the efficiency of the battery decreases when the discharge current is more than the rated current (or C_{rated}). If the efficiency of the battery is μ , ($0 < \mu \leq 1$), for current I , we can claim that the actual current drawn is I/μ not I . For example, if there is a demand of 2 charge units and the efficiency of the battery is 60% at that current level, then we would draw 3 charge units instead of 2.

To incorporate this *Rate Capacity* effect, we change the number of charge units to be actually discharged. We calculate the actual number of charge particles to be discharged

by looking up a table, which stores the relationship between the demanded charge units and actual charge units to be discharged, calculated based on the simulation of the PDE model. The basic step of the simulation, incorporating the *Rate Capacity* effect for any deterministic discharge demand, is described in Figure 7.

```

Simulation_Step
inputs: Current_State, Current_Demand,
Recovery_Probability[], Efficiency_Table[]
outputs: Next_State
variables: Actual_Demand
begin
    Generate a random number R between 0 and 1;
    Actual_Demand := Efficiency_Table[Current_Demand];
    If (Current_Demand > 0) then
        Next_State := Current_State - Actual_Demand;
    else if ( R < Recovery_Probability[Current_State]) then
        Next_State := Current_State + 1;
    end if
end

```

Figure 7: The Simulation_Step modeling *Rate Capacity* effect

We present results to validate our enhanced model in Section 6. For estimating the battery lifetime and the energy delivered by the battery of SoC designs, we need current demand waveforms for the system. The next section describes the methodology used to calculate cycle-accurate current demand waveforms for HW/SW mobile systems.

5 Generating System Current Discharge Profiles

In this section, we briefly describe the techniques that we used to generate system-level current discharge profiles. Please note that, while the focus of this paper is accurate and efficient estimation of the lifetime and delivered energy of the battery, system current discharge profiles are required as inputs to the battery models presented earlier, hence we include a description of how they can be generated.

We adapted the system-level power estimation framework presented in [9] for our purpose. It consists of a discrete-event co-simulation environment, power models for various system components (including processors, synthesized hardware, system-level buses, and memories), and speedup techniques that enable efficient power estimation for complex systems. As shown in [9], co-simulation (of the various system components) is necessary in order to obtain accurate power estimates. This is especially important in the context of battery life estimation, since we need accuracy in the current (hence, power) waveforms, not just the average power or total energy. To estimate the power dissipation in a programmable processor component, we use an instruction-set simulator enhanced with an instruction-level power model [18]. For application-specific hardware, we use fast-synthesis and power simulation of the synthesized netlist. For system-level buses, we estimate the switched capacitive power consumed in the bus by using bus line capacitances (calculated from bus length and width parameters provided

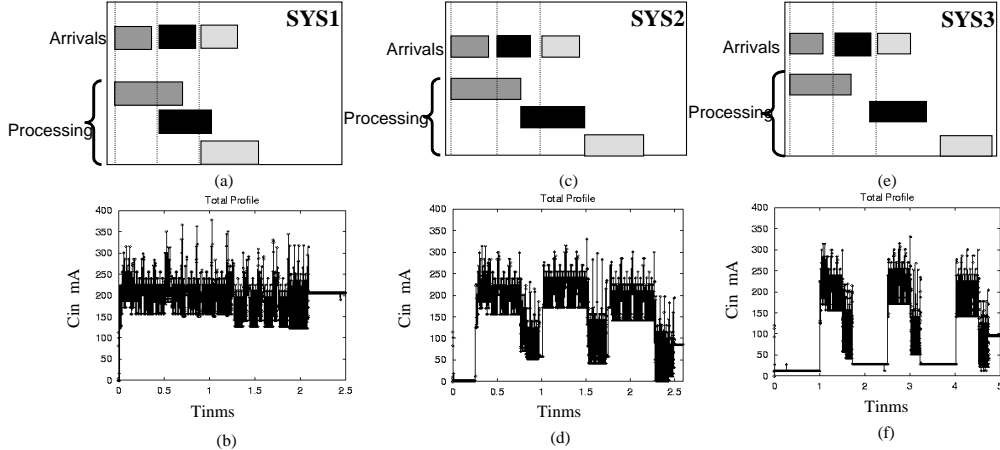


Figure 8: Scheduling Mechanism and Current Profiles of Example Implementations

by the designer), and switching activity (derived from the bus trace provided by the co-simulation tool). For memory cores, we use the access and idle power specifications from the core provider’s data sheets.

6 Experimental Results & Application

In this section, we demonstrate how our proposed battery life estimation technique can be used to evaluate different system implementations and select a battery-optimal implementation. We report results obtained by applying the proposed battery life estimation methodology on three system implementations, SYS1, SYS2 and SYS3, of the TCP/IP network interface subsystem, shown in Figure 1. In the first implementation (SYS1), the incoming packets are processed in a pipelined manner; for example, while *Chksum* is processing one packet, *Create_Packet* starts writing the next packet to memory. Figure 8(a) shows the arrival and processing of packets over time for SYS1. The cumulative current profile for this system is shown in Figure 8(b), with the total current (in mA) plotted over time (in ms). Alternative ways of processing the packets lead to the second (SYS2) and the third (SYS3) implementations, as shown in Figures 8(c) and 8(e) respectively. The corresponding current profiles are shown in Figures 8(d) and 8(f) respectively.

Table 2: Characteristics of Example Discharge Profiles

System	Max_Cur (mA)	Avg_Cur (mA)	% Slots Above Rated Current	% Idle Slots	Latency (ms)
SYS1	380	175	100	0	2.5
SYS2	330	154	60	0	2.5
SYS3	330	85	30	40	5.0

Table 2 shows the characteristics of the discharge profiles, in terms of the maximum current (*Max_Cur*), the average current (*Avg_Cur*), percentage of time slots with current above the rated current (*% Slots Above Rated Current*), percentage of idle slots in the discharge demand (*% Idle Slots*) and finally the latency of processing the packets (*Latency*). For our experiments, we used a Li-ion battery with the following parameters : $V_{OC} = 4.3V$, $V_{cut} = 2.8V$, and $C_{rated} = 125mA$. Table 2 shows that neither SYS1 nor SYS2 have any idle slots in the discharge demand, where as the discharge demand for SYS3 has 40% idle slots. It can be noticed that SYS1 violates the rated current maximum number of times among the three im-

plementations.

We measured the *Delivered Specific Energy* (the energy delivered by a battery of unit weight) and the *Battery Lifetime* for the above implementations using the proposed battery model. The battery model, we used for our simulation, has Nominal Capacity(N) = 650,000 *charge units* and Theoretical Capacity(T) = 1,000,000 *charge units*, obtained from the accurate PDE model of the considered Li-ion battery. The re-charging probabilities are appropriately set to accurately model the battery phenomena.

Table 3 shows the results of estimation using our stochastic battery model for the discharge profiles. In our estimation, we repeat simulation of any given discharge demand till the battery is discharged and report the delivered specific energy and battery lifetime. Each row in the table represents one of the example implementations described earlier. Columns 2 and 3 report the delivered specific energy and the battery lifetime when only the effect of *Rate Capacity* is modeled, as proposed in [5, 6]. Similarly, columns 4 and 5 show estimation results when only the effect of *Recovery* is considered, like the model proposed in [16]. The last three columns report results when both the *Rate Capacity* and *Recovery* effects are modeled, as proposed in this paper. In this case, in column 7, we also report the number of packets processed before the battery is completely discharged.

It can be seen from Table 3 that the different implementations have significantly different delivered energy and battery lifetimes. For example, looking at columns 6, 7 and 8 in Table 3, we see that SYS1 has 1.369 Wh/Kg delivered specific energy and 16875ms of lifetime, and processes 20250 packets before the battery is discharged. SYS2, whose discharge profile has 40% less slots above the rated current than the discharge profile of SYS1 (Table 2), can draw more than twice as much delivered energy from the same battery as SYS1, and shows four times increase in the battery lifetime as well as the number of packets processed. However, SYS3, whose discharge profile also has 40% idle slots, enhances the delivered energy and the battery lifetime even more significantly, by a factor of 1.3 and 2.3 over SYS2, respectively. SYS3 also shows 15% increase in the number of packets processed.

Table 3 also shows that both the *Rate Capacity* effect

Table 3: Estimation of Battery Life and Delivered Energy Using Stochastic Model

System	Rate Capacity Effect		Recovery Effect		Rate Capacity & Recovery Effect		
	Delivered Spec. Energy (Wh/Kg)	Life Time (ms)	Delivered Spec. Energy (Wh/Kg)	Life Time (ms)	Delivered Spec. Energy (Wh/Kg)	Life Time (ms)	Packets Processed
SYS1	1.369	16875	13.357	163650	1.369	16875	20250
SYS2	3.754	67717	15.553	280543	3.754	67717	81260
SYS3	2.858	88383	32.924	1115616	4.974	153817	92290

Table 4: Comparison with PDE model : Speed and Accuracy

System	Delivered Spec. Energy (Wh/Kg)			Life Time (ms)			Computation Time	
	STOC	PDE	% Err	STOC	PDE	% Err	STOC	PDE
SYS1	1.36	1.33	2.25	16785	17264	2.85	18.62 sec	>1 Day
SYS2	3.75	3.79	1.06	67717	65723	2.94	19.52 sec	>1 Day
SYS3	4.97	5.07	2.01	153817	154956	1.00	40.35 sec	>2 Days

and *Recovery* effect need to be considered simultaneously to be able to accurately measure the significant differences. For example, if only the *Rate Capacity* effect is considered (columns 2 and 3), the delivered energy and battery lifetimes obtained for SYS1 and SYS2 are accurate, as the corresponding discharge profiles do not have any idle times. However, the estimates for SYS3, which have significant idle times, are very inaccurate since the *Recovery* effect is not considered. Note that, though the fraction of current above the rated current is less in SYS3 than SYS2 (Table 2), the delivered energy calculated for SYS3 is less than that for SYS2. This is because the amount of offset from the rated current is more in case of SYS3 than in SYS2 (which is due to inaccuracy introduced by calculating average current instead of using instantaneous current). On the other hand, when only the *Recovery* effect is modeled (columns 4 and 5), we get significant improvement in delivered specific energy for SYS3. However, the values of specific energy delivered for all the three cases are significantly more than when only the *Rate Capacity* effect is considered, since these results do not reflect the diminishing battery efficiency due to the *Rate Capacity* effect. As noted earlier, by considering both the effects simultaneously, our stochastic model can estimate the delivered energy and lifetimes accurately, as shown in the last two columns of Table 3.

Finally, we validate our proposed model by comparing it with the accurate PDE model. Table 4 shows the delivered specific energy and the battery lifetimes for the three implementations, and corresponding percentage errors. The last two columns show the CPU time for the two models.

Our experimental results demonstrate that our stochastic model based approach is significantly faster than the PDE model. Since it only takes time in the order of seconds, as opposed to days taken by the PDE model, our approach can be used for design-space exploration to identify the most battery-efficient implementation of a mobile embedded system, which is not possible to do with the existing PDE model. At the same time our proposed approach does not compromise on accuracy, by modeling both the *Rate Capacity* and the *Recovery* effects, as opposed to the existing *Rate Capacity* based estimation techniques [5, 6].

7 Conclusion and Future Work

We have presented a stochastic model of battery and a framework for estimating the battery life as well as the delivered energy for system-level design space exploration of battery-

powered mobile embedded systems. The model proposed is fast enough to enable iterative battery life estimation for system level exploration. At the same time, it is very accurate, as validated using an accurate PDE model. In future, we will use the developed framework for exploring how system architectures can be made battery-efficient and for suggesting the optimum battery configuration for any HW/SW embedded system.

References

- [1] A. R. Chandrakasan and R.W. Broderson, *Low Power Digital CMOS Design*, Kluwer Academic Publishers, Norwell, MA, 1995.
- [2] J. Rabaey and M. Pedram (Editors) *Low Power Design Methodologies*, Kluwer Academic Publishers, Norwell, MA, 1996.
- [3] URL : <http://www.valence-tech.com/products/65krear.html>
- [4] T. Simunic, L. Benini and G. De Micheli, "Cycle-Accurate Simulation of Energy Consumption in Embedded Systems", *Proceedings of DAC, 1999*, New Orleans, USA, 1999
- [5] T. Simunic, L. Benini and G. De Micheli, "Energy-Efficient Design of Battery-Powered Embedded Systems", *Proceedings 1999 ISLPEd*, pp 212-217, 1999
- [6] S. Gold, "A PSPICE Macromodel for Lithium-Ion Batteries", *Proceedings of the 12th Battery Conference*, pp. 9-15, 1997.
- [7] S.C.Hageman, "Simple PSPICE Models Let You Simulate Common Battery Types", *EDN*, pp. 117-132, Oct. 1993.
- [8] L. Benini, G. Castelli, A. Macii, E. Macii, M. Poncino and R. Scarsi, "A Discrete-Time Battery Model for High-Level Power Estimation", *Proceedings of DATE, 2000*, pp 35-39.
- [9] M. Lajolo, A. Raghunathan, S. Dey and L. Lavagno, "Efficient Power Co-estimation Techniques for System-on-Chip Design", *Proceedings of DATE, 2000*, pp 27-34.
- [10] M. Doyle and J.S. Newman, "Analysis of capacity-rate data for lithium batteries using simplified models of the discharge process", *J. Applied Electrochem.*, vol. 27, no. 7, pp. 846-856, July 1997.
- [11] T.F. Fuller, M. Doyle, and J.S. Newman, "Relaxation phenomena in lithium-ion-insertion cells," *J. Electrochem. Soc.*, vol. 141, no. 4, pp. 982-990, Apr. 1994.
- [12] R.M. LaFollette, "Design and performance of high specific power, pulsed discharge, bipolar lead acid batteries," *10th Annual Battery Conference on Applications and Advances*, Long Beach, pp. 43-47, Jan. 1995.
- [13] B. Nelson, R. Rinehart, and S. Varley, "Ultrafast pulse discharge and recharge capabilities of thin-metal film battery technology", *11th IEEE International Pulsed Power Conference*, Baltimore, pp. 636-641, June 1997.
- [14] M. Doyle, T.F. Fuller, J.S. Newman, "Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell," *J. Electrochem. Soc.*, vol. 140, pp. 1526-1533, 1993.
- [15] C.F. Chiasserini, R.R. Rao, "A traffic control scheme to optimize the battery pulsed discharge," *Proc. of Milcom'99*, Atlantic City, NJ, Nov. 1999.
- [16] C.F. Chiasserini, R.R. Rao, "Energy efficient battery management," *Proc. of Infocom 2000*, Tel Aviv, Israel, March 2000.
- [17] J.S. Newman, *FORTAN programs for simulation of electrochemical systems*, <http://www.cchem.berkeley.edu/~jsngrp/>.
- [18] V. Tiwari, S. Malik and A. Wolfe, "Power Analysis of embedded software: A first step towards software power minimization", *IEEE Trans. VLSI Systems*, Vol. 2, No 4, pp. 437-445.