

Lifetime Improvement by Battery Scheduling

Marijn R. Jongerden¹ and Boudewijn R. Haverkort^{1,2}

¹ University of Twente

Centre for Telematics and Information Technology
Design and Analysis of Communication Systems (DACs)
Enschede, The Netherlands

{jongerdenmr, brh}@ewi.utwente.nl

² Embedded Systems Institute

Eindhoven, The Netherlands
boudewijn.haverkort@esi.nl

Abstract. The use of mobile devices is often limited by the lifetime of their batteries. For devices that have multiple batteries or that have the option to connect an extra battery, battery scheduling, thereby exploiting the recovery properties of the batteries, can help to extend the system lifetime. Due to the complexity, work on battery scheduling in literature is limited to either small batteries or to very simple loads. In this paper, we present an approach using the Kinetic Battery Model that combines real size batteries with realistic random loads. The results show that, indeed, battery scheduling results in lifetime improvements compared to the sequential usage of the batteries. The improvements mainly depend on the ratio between the average discharge current and the battery capacity. Our results show that for realistic loads one can achieve up to 20% improvements in system lifetime by applying battery scheduling.

1 Introduction

Many autonomous devices rely on batteries for their power supply. The capacity of the batteries is finite, and the duration with which one can use the device is limited by the battery lifetime. Lifetime, here, is the time of one discharge period of the battery, from full to empty. Although the battery lifetime depends mostly on its capacity and the level of the load applied to it, another important influence is *how the battery is used*, *i.e.*, its usage pattern [3].

When a battery is continuously discharged, a high current will cause it to provide less energy until the end of its lifetime than a lower current. This effect is termed the *rate-capacity effect*. On the other hand, during periods of low or no discharge current, the battery can recover to a certain extent. This is termed the *recovery effect*.

One approach to improve system lifetime is to connect one or more extra batteries. In this case, the batteries mostly are used in sequential order, the next one is used when the previous one has reached the end of its lifetime. Although this clearly prolongs the device lifetime, it is not the most efficient way. By using

the batteries one after each other one does not exploit the rate-capacity and recovery effect. Indeed, by switching regularly between the batteries one will give the batteries time to recover from the applied load. This will lead to longer system lifetimes, as we show in this paper.

Some research has already been done on battery scheduling. However, the approaches that use realistic random loads are limited to very small batteries, *cf.* [2, 17], and the approaches that do have real size batteries, such as [1], use only a limited number of test loads which are mostly very regular. The former leads to an overestimation of the improvement obtained by battery scheduling. For the latter the question remains how battery scheduling will perform under realistic loads.

In this paper, which has been presented earlier at the UK-PEW 2011 [7], we study the impact of battery scheduling when using real size batteries with a variety of realistic (random) loads. Various battery schedulers are modeled, and using the Kinetic Battery Model (KiBaM) [8–10] the overall system lifetime is computed for randomly generated loads. Our results show that for realistic loads one can achieve up to 20% improvements in system lifetime by applying battery scheduling.

The rest of the paper is structured as follows. In Section 2 an overview is given of the related work. Section 3 describes the used Kinetic Battery Model, and gives an expression for the maximum possible lifetime gain according to this model. The results of the simulations are given in Section 4. Finally, we conclude in Section 5.

2 Related work

The scheduling of batteries has attracted quite some attention in the literature. Over the years various kinds of battery models have been developed. An overview of the main modeling approaches is given in [6]. Part of these models have been used to study the problem of battery scheduling. We consider here the main approaches. The most important scheduling schemes that are studied are:

- Sequential scheduling: a next battery is only picked when the current one is empty.
- Round robin scheduling: at fixed moments in time another battery is used. The batteries are used in a fixed order.
- Pick-best scheduling: at fixed moments in time the status of all batteries is checked and the best battery is used. What is the best battery can be determined in several ways, for example the battery with the highest voltage, or the battery that has been used for the shortest period of time.

Benini et al. [1] use an electrical-circuit model to describe the batteries. They consider sequential scheduling, round robin scheduling and various types of pick-best scheduling, where either the output voltage or the time that a battery has not been used determines which battery is to be scheduled. The different scheduling schemes are applied to several battery configurations containing up to four

batteries. The loads that have been used are simple continuous and intermitted loads and two real-life example load profiles. Which scheduler performs best depends on the applied load.

Chiasserini and Rao [2] use a discrete-time Markov battery model to compare three different scheduling schemes in a multiple battery system. In the model, the recovery of the battery is considered as a random process. Also, the workload is stochastic. Next to the commonly used round robin and pick-best scheduler, also a random scheduler is considered. The schedulers are compared for different job arrival rates. The results show that the pick-best scheduler outperforms the other two. However, the complexity of the used models limits the analysis to cases with only small batteries.

In all this work the battery scheduling is limited to simple scheduling schemes. All show that battery scheduling gives longer system lifetime than when the batteries are used sequentially. However, they do not indicate whether longer lifetime could be possible by using even smarter scheduling.

Sarkar and Adamou [17] propose an algorithm for computing an optimal scheduling scheme based on the stochastic battery model of Chiasserini and Rao. To do this, they translate the problem to a stochastic shortest path problem. The optimal solution can only be computed for very small batteries. However, they do show that pick-best scheduling performs close to optimal.

Another optimization approach is taken in [4, 5], in which the batteries are modeled using priced-timed automata. With model checking techniques the schedule that gives the maximum lifetime is computed. The result is compared to the simple sequential, round robin and pick-best schedulers. Although the results show that the round robin and pick-best schedulers are sometimes far from optimal, these schedulers are much better than the sequential scheduler. The model actually shows that sequential scheduling results in the shortest lifetime possible. In [5] a first step towards random loads was taken. However, the priced-timed automata model was limited to very small battery capacities. In this paper we combine the random loads with realistic battery capacities, in order to obtain a better prediction of the potential gain of battery scheduling.

All these studies show that by applying battery scheduling, the system lifetime will be extended. However, the improvement varies a lot between the different modeling approaches. Where Benini et al. [1] predict an average improvement of approximately 11% for a two battery system, Chiasserini and Rao [2] show improvements of more than 100%.

3 Kinetic battery model

3.1 Introduction

The battery model we use is the Kinetic Battery Model (KiBaM) of Manwell and McGowan [8–10]. This model is very intuitive, and the simplest model that includes the two important non-linear battery properties, the rate-capacity effect and the recovery effect [6]. The rate-capacity effect is the effect that less charge

can be drawn from the battery when the discharge current is increased. However, some of the charge left behind in the battery after a period with a high discharge current will be available for usage after a period with no or low current. This is the recovery effect.

In the model the battery charge is distributed over two wells: the available-charge well and the bound-charge well (cf. Figure 1). For the full battery, a

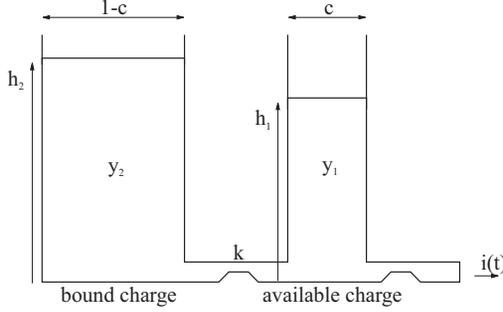


Fig. 1. Two-well-model of the Kinetic Battery Model

fraction c of the total capacity is put in the available charge well, and a fraction $1 - c$ in the bound charge well. The available charge well supplies electrons directly to the load ($i(t)$), whereas the bound-charge well supplies electrons only to the available-charge well. The charge flows from the bound charge well to the available charge well through a “valve” with fixed conductance, k . Next to this parameter, the rate at which charge flows between the wells depends on the height difference between the two wells. The heights of the two wells are given by: $h_1 = \frac{y_1}{c}$ and $h_2 = \frac{y_2}{1-c}$. The change of the charge in both wells is given by the following system of differential equations:

$$\begin{cases} \frac{dy_1}{dt} = -i(t) + k(h_2 - h_1), \\ \frac{dy_2}{dt} = -k(h_2 - h_1), \end{cases} \quad (1)$$

with initial conditions $y_1(0) = c \cdot C$ and $y_2(0) = (1 - c) \cdot C$, where C is the total battery capacity. The battery is considered empty when there is no charge left in the available charge well, $y_1 = 0$.

One can solve the differential equations using Laplace transformations when the load is constant ($i(t) = I$). In this case the evolution of the charge in the two charge wells is given by:

$$\begin{cases} y_1(t) = cC - Ict - \frac{I(1-c)}{k'} (1 - e^{-k't}), \\ y_2(t) = (1-c)C - (1-c)It + \frac{I(1-c)}{k'} (1 - e^{-k't}), \end{cases} \quad (2)$$

where k' is defined as $k' = k/(c(1-c))$.

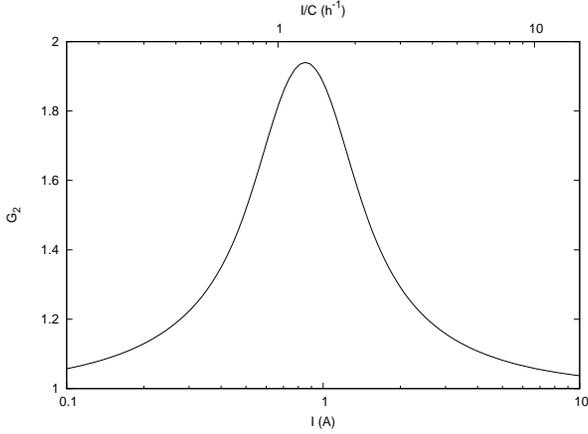


Fig. 2. The maximum lifetime gain for a system with two batteries as a function of the constant discharge current. Note that the current is plotted in a logarithmic scale. The top x-axis gives the current normalized to the capacity of one battery.

From the equation for the available charge one can obtain the battery lifetime (t_s) by setting $y_1 = 0$:

$$t_s = \frac{C}{I} - \frac{1}{k'} \left(\frac{1-c}{c} - W \left(\frac{1-c}{c} e^{-\frac{c k'}{I} + \frac{1-c}{c}} \right) \right), \quad (3)$$

where W denotes the Lambert W function. The Lambert W function is the inverse function of $f(W) = W e^W$ [14].

3.2 Maximum possible lifetime gain

In [5] it is shown that, according to the KiBaM model, in theory the best way to discharge the batteries in a multiple battery system is by using them in parallel. For a system with N identical batteries discharged with a continuous current I the lifetime is given by:

$$t_{p,N} = \frac{NC}{I} - \frac{1}{k'} \left(\frac{1-c}{c} - W \left(\frac{1-c}{c} e^{-\frac{NCk'}{I} + \frac{1-c}{c}} \right) \right). \quad (4)$$

This equation is similar as (3) with $\frac{I}{N}$ substituted for I .

Using Equation (3) and (4) we can compute the maximum possible gain one can obtain by applying battery scheduling in the case of a constant discharge current. The system lifetime when using N batteries sequentially will be $N t_s$, hence the maximum possible gain with N batteries G_N is given:

$$G_N = \frac{t_{p,N}}{N t_s}. \quad (5)$$

In Figure, 2 the gain for a system with 2 batteries (G_2) is given as a function of the discharge current. The batteries that have been used in this computation

are similar to those used in [5], *i.e.*, $c = 0.166$ and $k = 2.815 \cdot 10^{-4} \text{ s}^{-1}$. However, here the capacity is increased to a realistic value, $C = 2400 \text{ As}$, instead of the much smaller capacity of 330 As used in [5]. This type of battery is used in the Itsy pocket computer, which was also simulated by Rakhmatov *et al.* in [15, 16].

The discharge current has been varied between 0.1 A and 10 A . For this system of batteries the highest gain is obtained at a discharge current of approximately 0.85 A , where the gain is more than 1.9 . The peak can be explained as follows. When the discharge current gets too high, the available charge well will be depleted too fast and the slow recovery process will hardly increase the usable capacity, even when scheduling is applied. At low discharge currents the loss of capacity due to the rate capacity effect is low, *i.e.*, the flow of charge from the bound to the available charge well can keep up with the demand, and little charge will be left behind in the bound charge well. Therefore, the gain of allowing batteries to recover by the scheduling is limited. However, at a discharge current of 0.1 A the gain still is approximately 1.05 , and a 5% lifetime extension is still a considerable improvement.

When we look at Equation 4 and Equation 3 we see that the discharge current I always appears in direct relation with the battery capacity C , in the form $\frac{C}{T}$. This implies that when the battery capacity is halved, and the other battery parameters stay the same, the discharge current needs to be halved as well to obtain the same lifetime gain. Using the top x-axis Figure 2 shows how the maximum lifetime gain depends on the current normalized to the capacity of one battery (I/C).

Of course, the shape of curve depicted in Figure 2 and the position of the maximum highly depend on the battery parameters c and k , as well as the number of batteries N . The plots in Figure 3 show how the curve changes when one of the parameters is varied. In all the three subfigures the curve of Figure 2 is given as reference, drawn with a solid line.

In Figure 3(a) the number of batteries (N) is varied. The increase of the number of batteries leads to an increase of the gain (G_N). This can be understood as follows. When more batteries are used, the discharge current per battery will drop. The flow of charge from the bound charge well to the available charge well now can keep up better with discharge current, and more charge will be available for the load.

Figure 3(b) shows that when the fraction of available charge (c) is increased the gain will be lower. As more charge is directly available the lifetime of sequential discharge will increase, and dividing the load will be less beneficial. In the extreme case that $c = 1$, the batteries will behave as ideal batteries, and all charge will always be available for the load. In this case, there will be no difference in lifetime between sequential and parallel discharge, and $G_N = 1$ for all currents.

Finally, Figure 3(c) shows that an increase of k leads to an increase of the current at which the gain is maximal. The increase of k causes the flow of the charge from the bound charge well to the available charge well to be faster. In

this way, the flow will be able to keep up with higher discharge currents, and the gain will be largest at a higher discharge current.

The results shown above indicate that a system with two identical batteries used in parallel will behave as one with double the capacity. However, the possibility of connecting batteries directly in parallel is under debate. Where [12] claims lithium batteries are well suited to connect in parallel, [13] says one should not do this. One problem of connecting batteries in parallel is that even for two batteries of the same type a difference in potential can occur. When this happens a current will flow between the batteries, resulting in a loss of capacity and, even worse, possibly damage to the batteries. Using batteries in parallel requires extra electronic circuitry, which consumes some power and decreases efficiency. Also, in some situations, like the routing problem described in the Section 1, parallel usage is simply impossible. Using a simple scheduling scheme, like round robin scheduling, one can circumvent the problems of parallel usage, and still obtain an improvement in system lifetime.

4 Battery scheduling results

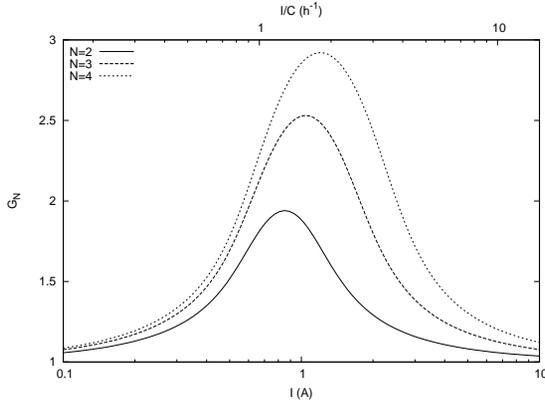
4.1 Simulation set-up

Like in the previous section, the batteries we model are the lithium-ion batteries that are used in the Itsy pocket computer. The modeled battery has a capacity of 2400 As, with $c = 0.166$ and $k = 2.815 \cdot 10^{-4} \text{ s}^{-1}$. The load currents are within the range of the Itsy pocket computer, up to 600 mA [15]. In [6], we have shown that the KiBaM yields good results in lifetime computations for this type of battery and loads. We generate 10000 random load traces, which are subsequently input to the system lifetime computations using the KiBaM equations (1). We introduce randomness into the loads in three steps. First, we introduce random on-times, and keep the discharge current constant. Second, we introduce random discharge currents and keep the periods of discharge constant. Finally, we model full random loads based on a Markov model that mimics a typical usage pattern of a mobile device.

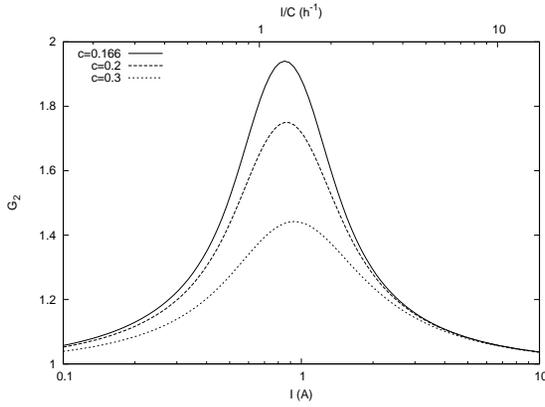
In the analysis we use four basic scheduling schemes:

- *sequential*: a next battery is chosen when the current one is empty.
- *load-round-robin*: the batteries are chosen in a fixed order, a switch between batteries takes place at the moment the discharge current is changed to another positive current.
- *best-of-two*: at the moment the load changes the battery with the most charge in the available-charge well is chosen.
- *time-round-robin*: the batteries are chosen in a fixed order, a switch between batteries takes place after a fixed amount of time has passed.

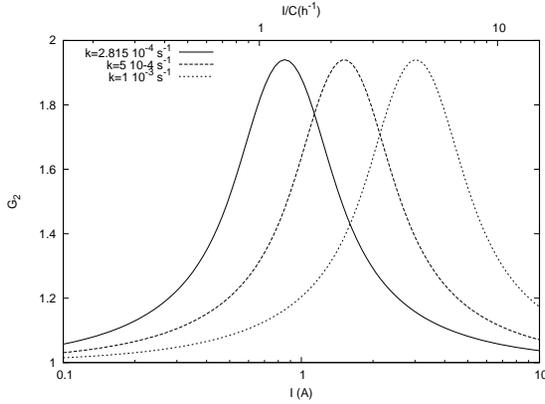
Three of these schedulers were also used in [5] in the setting of priced timed automata: sequential, load-round-robin and best-of-two. These schedulers are used here to see what the effect is of the bigger battery capacity on the lifetime gain. The time-round-robin scheduler is used to approach the maximum



(a) Varying number of batteries (N), $c = 0.166$, $k = 2.815 \cdot 10^{-4} \text{ s}^{-1}$ and $C = 2400 \text{ As}$.



(b) Varying fraction available charge (c), $N = 2$, $k = 2.815 \cdot 10^{-4} \text{ s}^{-1}$ and $C = 2400 \text{ As}$.



(c) Varying conductance (k), $N = 2$, $c = 0.166$ and $C = 2400 \text{ As}$.

Fig. 3. Maximum possible gain (G_N) for various battery parameters.

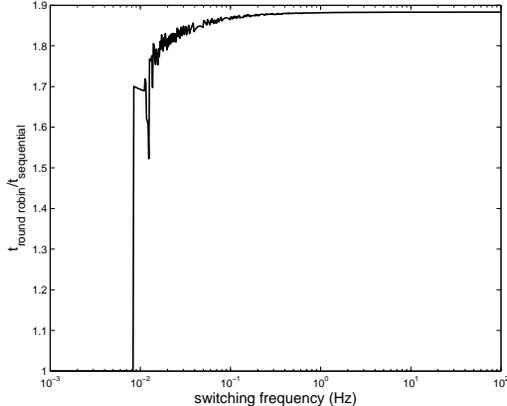


Fig. 4. Gain of using a round robin scheduler compared to sequential usage as a function of the switching frequency.

lifetime. As discussed before, in Section 3.2, parallel discharge, which leads to the maximum lifetime, may not be practically possible. However, the lifetime of parallel discharge can be easily approached by using a fast switching round robin scheduler, as will be shown in the next section.

Of course, the switching between the batteries in all the used schedulers will cost some extra energy, especially for the time-round-robin scheduler which will switch between the batteries the most often. However, the energy needed to switch between batteries will be negligible compared to the actual load. In [11], Matsuura presents a low-power pulse generator which operates with a discharge current of $0.15 \mu\text{A}$ at a voltage of 1.5 V . This current is at least a factor 1000 less than the discharge current the device operates at, which is in the order of mA. Therefore, the cost of switching using the time-round-robin scheduler can be neglected without introducing any significant error to the computed system lifetime.

4.2 Round robin frequency dependence

In order to find what switching frequency is efficient to approach the maximum lifetime, we investigate how the gain in lifetime depends on the switching frequency in case of round robin scheduling. The system of two batteries is discharged with a constant current of 1 A . We compare the system lifetime obtained with the round robin scheduler with that of sequential battery usage.

In Figure 4, we show the ratio of the system lifetime using round robin scheduling to the lifetime with sequential scheduling as a function of the round robin switching frequency. We see that the gain in lifetime of using the scheduler grows to a level of 1.89 when the switching frequency is increased. This level is the gain one would get with parallel discharge, which can be seen as switching

with infinite frequency. The figure shows that already for a switching frequency of 1 Hz the gain is close to optimal, so switching at higher frequencies is not necessary. Therefore we use a switching frequency of 1 Hz for the time-round-robin scheduler in the simulations in the following sections.

On the side of the low frequencies, smaller than 0.1 Hz, the graph fluctuates with clear downward tendency, that is, a small increase of the switching frequency may result in a considerable change in lifetime. This can be explained as follows. At very low frequency, lower than 0.008 Hz, the batteries are emptied in one period and the round robin scheduler results in sequential usage. When the frequency is increased the point will be reached where the first battery will not be emptied completely before the switch takes place. While the second battery is used the first can recover. Due to this recovery time the battery can be used longer, and the system lifetime is increased. This results in the first jump in the graph. Every time the batteries can recover for one period more a next jump in the graph occurs. The size of the jumps decreases as the frequency increases, since the extra recovery time will be shorter at higher frequencies. Between the jumps the system battery lifetime decreases, since the extra recovery time decreases as the switching frequency is increased. Thus, the ratio between on and off time will decrease until the next jump occurs.

4.3 Random times

As first random load, we take an on-off load with 250 mA on-current. The off periods last 1 minute, and the on periods are uniformly distributed over the interval $[\frac{1}{2}, \frac{3}{2}]$ minute. This load has also been used in [5], but there the modelled batteries had a capacity that was approximately 8 times smaller than the real capacity, which is used here.

We compute the system lifetime for 10000 randomly generated loads using the four schedulers mentioned in Section 4.1. In Figure 5 the empirical lifetime distributions, expressed as the frequency count of the lifetimes, is provided for the different schedulers. The bin size for these distributions has been set to 1 minute. For clarity the histograms are plotted using lines. In Table 1 the mean and variance of the computed lifetimes are given for the different schedulers. As can be observed, clear system lifetime improvement is obtained when battery scheduling is applied. On average the load-round-robin and best-of-two scheduler outperform sequential usage by 6% and 6.6% respectively. Also, the two schedulers perform only slightly worse than the time-round-robin scheduler.

When we compare these results with those in [5], in which a gain of 65% was observed, we see that the relative gain in lifetime obtained by battery scheduling is much less than for smaller batteries. This is related to the result in Section 3.2, where the maximum possible gain is given as a function of the discharge current, as follows. The mean of the discharge current of the used loads is 125 mA. This gives a ratio between the load and the battery capacity of $0.125 \text{ A}/0.666 \text{ Ah} = 0.1875 \text{ h}^{-1}$ for the real size battery. For the smaller battery used in [5] the ratio is $0.125 \text{ A}/0.0916 \text{ Ah} = 1.36 \text{ h}^{-1}$. Using the top x-axis in Figure 2 one sees that the ratio of 0.1875 h^{-1} allows for a gain of less

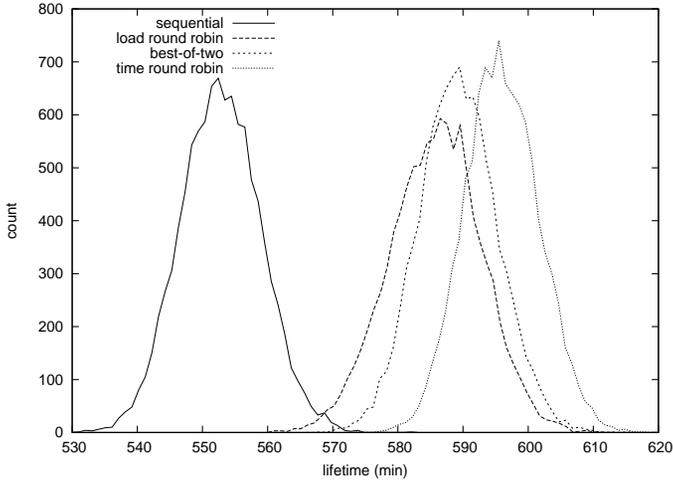


Fig. 5. Empirical lifetime distributions generated with 10000 on-off loads with random on-times.

Table 1. Mean and variance of the lifetimes obtained with the different schedulers for the loads with random on-times.

scheduler	mean lifetime (min)	variance (min^2)
sequential	552.87	39.36
load-round-robin	585.90	50.39
best-of-two	589.33	37.44
time-round-robin	596.01	33.38

than 10%, whereas the ratio of 1.36 h^{-1} is close to the peak value of a maximum possible gain of 90%.

4.4 Random currents

The second set of random loads is also used in [5]. In this set of random loads every minute we uniformly choose the discharge current from the set $\{0, 100, 200, 300, 400, 500\}$ mA. The current will stay constant for one minute until the next current is picked. We use the same schedulers as in the previous section. The load-round-robin and best-of-two scheduler now make a scheduling decision every minute, when the new current is picked.

Again, 10000 loads were generated. The lifetime distributions for these loads are given in Figure 6, and the numbers for the mean and variance of the simulations are given in Table 2. The trend is similar to the previous random load. The best-of-two scheduler performs slightly better than the load-round-robin scheduler, and both perform close to the time-round-robin scheduler. The average improvements relative to the sequential scheduler are 16% and 18% for the load-round-robin and best-of-two, respectively. This is much better than for the

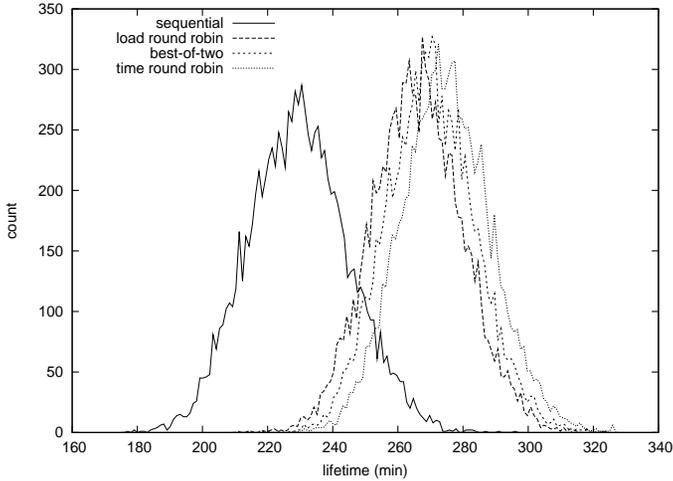


Fig. 6. Empirical lifetime distributions generated with 10000 loads with random discharge current.

Table 2. Mean and variance of the lifetimes obtained with the different schedulers for the loads with random discharge currents.

scheduler	mean lifetime (min)	variance (min^2)
sequential	229.55	237.98
load-round-robin	266.12	206.73
best-of-two	270.10	195.44
time-round-robin	274.84	197.20

loads with random on times due to the higher average discharge current. For the loads with random currents the average discharge current is 250 mA. In Figure 2 we can see that for a discharge current of 250 mA the maximum lifetime gain is just under 20% when the system is discharged with a continuous current of 250 mA. On the other hand, the maximum gain for the random on-times, which have an average discharge current of 125 mA, is approximately 10%.

Due to the higher variance in discharge current, the variance in lifetime is larger for this load, as visible through the “wider” graphs, and the numbers for the variance in Table 1 and 2.

When we compare the results with those of [5], we see that the difference in lifetime gain is not as large as with the previous set of random loads. The timed-automata approach in [5] resulted in a system lifetime that was 26% longer than the sequential schedule. For the smaller batteries used in [5], the ratio between the discharge current and the battery capacity is $0.250 \text{ A}/0.0916 \text{ Ah} = 2.73 \text{ h}^{-1}$. This ratio leads, according to Figure 2, to a maximum possible gain of approximately 26%, which is the obtained gain.

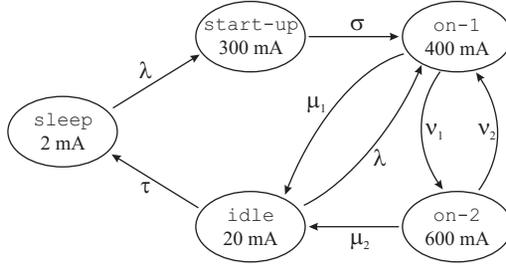


Fig. 7. State transition diagram of the workload model.

Table 3. Transition rates of the Markov model.

transition	λ	σ	μ_1	ν_1	μ_2	ν_2	τ
rate (min^{-1})	$\frac{1}{5}$	2	$\frac{1}{14}$	$\frac{1}{14}$	$\frac{1}{25}$	$\frac{4}{25}$	$\frac{1}{2}$

4.5 Full random load

The final step in introducing randomness into the loads is having both random discharge times and random currents. This is done by using a Markov model that represents a simple workload of a device. The state transition diagram of this Markov model is given in Figure 7. The device has 5 different states: **sleep**, **start-up**, **on-1**, **on-2** and **idle**. In the **sleep** state the device draws a 2 mA current from the battery. From the **sleep** state the device first has to start-up before it can go to the **on-1** state. The start-up takes 30 seconds on average, and during start-up the discharge rate is 300 mA. From the **on-1** state a transition is made either to the **idle** state, or to the **on-2** state, both with probability $\frac{1}{2}$. In the **on-1** and **on-2** state the discharge current is 400 and 600 mA, respectively. The average residence time in the **on-1** and **on-2** state is 7 and 6 minutes, respectively. From the **on-2** with probability $\frac{4}{5}$ it will go back to **on-1**, and with probability $\frac{1}{5}$ go to **idle**. In the **idle** state the current is 20 mA, and the average time it takes to go back to **sleep** is 2 minutes. The used discharge currents are based on the average discharge currents for different modes of the Itsy pocket computer [16]. An overview of the transition rates is given in Table 3.

Again, we use the sequential, load-round-robin, best-of-two scheduler and time-round-robin to compute the system lifetime. The Markov model is used to generate 10000 random loads. For the load-round-robin and best-of-two scheduler the scheduling choices are made at the state changes.

In Figure 8 the empirical lifetime distributions for the Markov workload model are given. In Table 4 the mean and variance of the lifetimes are given. Again, we see the same order in performance of the four schedulers. However, the difference between the time-round-robin scheduler, and the best-of-two and load-round-robin scheduler is a lot larger. Even though the average improvement compared to sequential discharge of the load-round-robin and best-of-two scheduler is 10.3% and 12.6%, respectively, the time-round-robin scheduler leads to

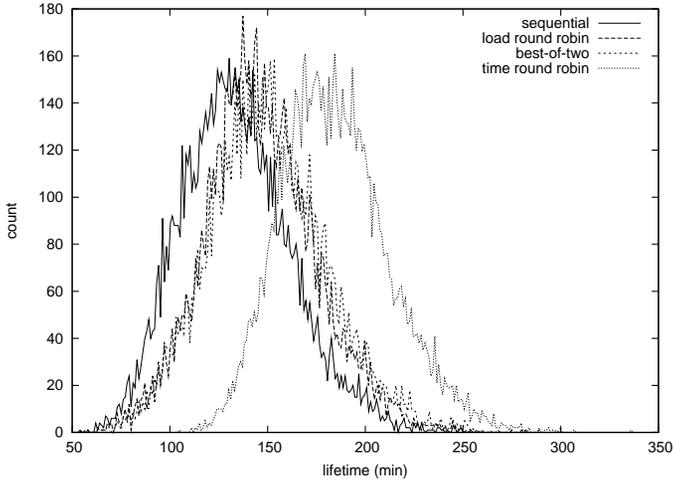


Fig. 8. Empirical lifetime distributions generated with 10000 Markov model loads.

Table 4. Mean and variance of the lifetimes obtained with the different schedulers for the Markov model loads.

scheduler	mean lifetime (min)	variance (min^2)
sequential	133.72	794.12
load-round-robin	145.99	803.83
best-of-two	149.11	892.97
time-round-robin	183.61	775.06

an even longer lifetime. The time-round-robin scheduler outperforms sequential scheduling with 40.6% for this workload. The large difference between the time-round-robin scheduling and the round robin and best-of-two schedulers is caused by the longer average time between scheduling moments.

5 Conclusions

In this paper we have looked at the impact of battery scheduling for real size batteries using several random loads. Our results show that the use of battery scheduling can improve the system lifetime. The analysis of the kinetic battery model shows that parallel discharge leads to the largest lifetime improvement. The actual lifetime gain highly depends on the load-capacity ratio.

Parallel discharge is not always possible, as it may damage the batteries. However, we show that using a high frequency round-robin scheduler one can approach the lifetime obtained by parallel discharge. The simulation results show that also for more complex random loads, battery scheduling helps to improve the system lifetime considerably. The gain in lifetime compared to sequential

discharge of the batteries for the different schedulers varies with the type of load. The average maximum lifetime gain can be well predicted by computing the maximum possible lifetime gain for a continuous discharge current using the average current of the random load.

References

1. Benini, L., Castelli, G., Macii, A., Macii, E., Poncino, M., Scarsi, R.: Extending lifetime of portable systems by battery scheduling. In: Design, Automation and Test in Europe (DATE). pp. 197–203. IEEE Computer Society, Los Alamitos, CA, USA (2001)
2. Chiasserini, C., Rao, R.: Energy efficient battery management. *IEEE Journal on Selected Areas in Communications* 19(7), 1235 – 1245 (2001)
3. Cloth, L., Haverkort, B.R., Jongerden, M.R.: Computing battery lifetime distributions. In: Proc. IEEE DSN 2007. pp. 780–789. IEEE Computer Society Press (2007)
4. Jongerden, M., Haverkort, B., Bohnenkamp, H., Katoen, J.P.: Maximizing system lifetime by battery scheduling. In: Proc. IEEE DSN 2009. pp. 63–72. IEEE Computer Society Press (2009)
5. Jongerden, M., Mereacre, A., Bohnenkamp, H., Haverkort, B., Katoen, J.P.: Computing optimal schedules for battery usage in embedded systems. *IEEE Transactions on Industrial Informatics* 6(3), 276–286 (August 2010)
6. Jongerden, M.R., Haverkort, B.R.: Which battery model to use? *IET Software* 3(6), 445–457 (2009)
7. Jongerden, M., Haverkort, B.: Lifetime improvement by battery scheduling. In: UK Performance Engineering Workshop, July 7-8, 2010, Bradford, UK
8. Manwell, J., McGowan, J.: Lead acid battery storage model for hybrid energy systems. *Solar Energy* 50, 399–405 (1993)
9. Manwell, J., McGowan, J.: Extension of the kinetic battery model for wind/hybrid power systems. In: Proc. of the 5th European Wind Energy Association Conference (EWEC '94). pp. 284–289 (1994)
10. Manwell, J., McGowan, J., Baring-Gould, E., Stein, W., Leotta, A.: Evaluation of battery models for wind/hybrid power system simulation. In: Proc. of the 5th European Wind Energy Association Conference (EWEC '94). pp. 1182–1187 (1994)
11. Matsuura, Y.: Low-power consumption reference pulse generator. United States Patent 4,618,837 (1986)
12. Battery University.com: <http://www.batteryuniversity.com/partone-24.htm> (april 2010)
13. Low Cost Batteries.com: <http://www.lowcostbatteries.com/articles.asp?id=107> (april 2010)
14. Wolfram Mathworld Lambert-W Function: <http://mathworld.wolfram.com/lambertw-function.html> (april 2010)
15. Rakhmatov, D., Vrudhula, S., Wallach, D.A.: Battery lifetime predictions for energy-aware computing. In: Proc. of the 2002 International Symposium on Low Power Electronics and Design (ISLPED '02). pp. 154–159 (2002)
16. Rakhmatov, D., Vrudhula, S., Wallach, D.A.: A model for battery lifetime analysis for organizing applications on a pocket computer. *IEEE Transactions on VLSI Systems* 11(6), 1019–1030 (2003)
17. Sarkar, S., Adamou, M.: A framework for optimal battery management for wireless nodes. In: Proc. IEEE INFOCOM 2002. pp. 179–188 (2002)