Evaluation of Battery Lifetimes using Inhomogeneous Markov Reward Models

Lucia Cloth, Boudewijn R. Haverkort, Marijn Jongerden
University of Twente
[lucia,jongerndenmr,brh]@ewi.utwente.nl

***Abstract*— The usage of mobile devices like cell phones, navigation systems, or laptop computers, is limited by the lifetime of the included batteries. This lifetimes depends naturally on the rate at which energy is consumed, however, it also depends on the usage pattern of the battery. Continuous drawing of a high current results in an excessive drop of residual capacity. However, during intervals with no or very small currents, batteries do recover to a certain extend. We model this complex behaviour with an inhomogeneous Markov reward model. The state-dependent reward rates thereby correspond to the power consumption of the attached device and to the available charge, respectively. We develop new numerical algorithms for the computation of the distribution of the consumed energy and show how different workload patterns influence the overall lifetime of a battery.**

I. Introduction

With the proliferation of cheap wireless access technologies, such as wireless LAN, Bluetooth as well as GSM, the number of wireless devices an average citizen is using has been steadily increasing since a few years. Such devices not only add to the flexibility with which we can do our work, but also add to our reachability and our security. Next to these personal wireless devices, an ever growing number of wireless devices is used for surveillance purposes, most notably in sensor-type networks. A common issue to be dealt with in the design of all of these devices is power consumption. Since all of these devices use batteries of some sort, mostly rechargeable, but sometimes not even that, achieving low power consumption for wireless devices has become a key design issue. This fact is witnessed by many recent publications on this topic, and even a special issue of *IEEE Computer* (November 2005) devoted to it [1].

Low-power design is a very broad area in itself, with so-called "battery-driven system design" a special branch of it, that becomes, due to the reasons mentioned, more and more important. A key issue to be addressed is to find the right tradeoff between battery usage and required performance: how can we design a (wireless) system such that with a given battery, good performance (throughput, reachability, and so on) is obtained, for a long-enough period. Stated differently, how should the processes in the wireless device be organised such that the battery lifetime (which determines the system lifetime) will be as high as possible. Indeed, it has been observed recently that due to the specific physical nature of batteries, achieving the longest battery lifetime is not always achieved by "just" trying to minimise the power consumption. Instead, also the way in which the power is consumed, that is, the employed current levels and the current-extraction patterns do play a role in the battery lifetime.

In order to obtain a better insight in the lifetime of batteries, a wide variety of models has been developed. We will discuss some of these models in the next section. What has not been done, however, is the combination of such power consumption models in a versatile way with performance models for mobile communication systems, thereby taking into account typical physical aspects of battery operation. It is exactly this issue that we address in the current paper.

Our approach will be to describe the operation of a system with an abstract workload model, describing the various states the wireless device can be in, together with the energy consumption rates in those states. Also, the transition possibilities between these states will be represented in the workload model. Such a description can be interpreted as a Markov-reward model in which accumulated reward stands for the amount of energy consumed. The system or battery lifetime would then be equal to the time until a certain level of consumption (the battery its energy) is reached. Determining this time, or better, its distribution, could be done with well-known techniques for performability evaluation. However, such an approach does not well take into account the physical aspect of battery operation. Indeed, studies on batteries reveal that the battery depletion rate in general is non-linear in time, and, moreover, also depends on the amount of energy still in the battery. Furthermore, in periods when a battery is not used, subtle but important battery-restoration effects are in place, that apparently refill the battery. Translating such effects to a Markov reward model context, this would amount to models in which, possibly, the reward and transition rates depend on time and/or on the amount of reward accumulated so far, and in which both positive and negative reward rates are in place. There are no generally applicable algorithms for evaluating completion-time like quantities for such models available.

So, next to the fact that we combine battery models with performance models, the current paper also presents two new algorithms for evaluating performability-like distributions for a new and very general class of time- and accumulated reward-dependent Markov-reward models.

We note that this is not the first work that addresses performability-like measures in a non-homogeneous context. In the 1990's some work has appeared on the computation of transient state probabilities for pure Markovian models; no rewards were addressed there [2], [3], [4]. A recent paper by Telek *et al.* characterises the performability distribution in inhomogeneous MRMs through a coupled system of partial differential equations that is solved through discretisation, and used to derive systems of ordinary differential equations to determine moments of accumulated reward [5]. In [6] we presented a simpler version of the discretisation algorithm (cf. Section V) as a model for non-homogeneous failure processes.

The rest of the paper is organised as follows. We briefly introduce into the world of batteries and battery models in Section II. We then fix some notation for inhomogeneous Markov reward models in Section III and present the Markov reward models for batteries used in the rest of the paper in Section IV. In Section V we describe the algorithms for the computation of the battery lifetime. Section VI discusses the results obtained for the models and in Section VII we conclude this paper.

II. Batteries

The two most important properties of a battery are its voltage (expressed in volts V) and its capacity (mostly expressed in Ampere-hour, Ah); the product of these two quantities gives the energy stored in the battery. For an ideal battery the voltage stays constant over time until the moment it is completely discharged, then the voltage drops to zero. The capacity in the ideal case is the same for every load for the battery. Reality is different, though: the voltage drops during discharge and the capacity is lower under a higher load.

In the ideal case it would be easy to calculate the lifetime of a battery. The lifetime (L) in the case of a constant load is the capacity (C) over the load current (I (Ampere)):

$$L = C/I.$$

Due to various nonlinear effects this relation does not hold for real batteries. A simple approximation for the lifetime under constant load can be made with Peukert's law [7]:

$$L = \frac{a}{I^b},$$

where $a > 0$ and $b > 1$ are constants which depend on the battery. This relation does not hold for a variable load. Following Peukert's law, all load profiles with the same average would have the same lifetime. Experimentally it can be shown that this is not the case. One of the effects playing an important role here is the recovery effect of the battery, as follows.

All batteries are driven by electro-chemical reactions. During the discharge, an oxidation reaction at the anode takes place. In this reaction electrons are produced, which are released into the (connected) circuit. At the cathode a reduction reaction takes place. Here electrons are accepted from the circuit and consumed in the reaction:

$$\begin{aligned} O_1 + ne^- \rightarrow R_1, &\qquad \text{cathode} \\ R_2 \rightarrow O_2 + me^-, &\qquad \text{anode} \end{aligned}$$

As an example of a chemical reaction, this is what happens in the highly-used Lithium-ion batteries [8]:

$$\begin{aligned} Li_{1-x}CoO_2 + xLi^+ + xe^- &\rightarrow LiCoO_2, \\ CLi_x &\rightarrow C + xLi^+ + xe^-. \end{aligned}$$

These are the reactions for discharging the battery. For charging the battery the arrows in the reaction equations are directed to the left.

In a lithium ion battery, the Li^+ ions made at the anode have to diffuse to the cathode when a current is drawn from the battery. When the current is too high the internal diffusion cannot keep up with the rate the ions react at the cathode. As a result, the positive charge at the cathode drops and rises at the anode. This causes a drop in the output voltage of the battery. However, when the battery is less loaded for a while, the ions have time enough to diffuse again and charge recovery takes place.

Another effect that occurs when high currents are drawn is that no reaction sites (molecules) are available in the cathode. At small load (low currents) the reaction sites are uniformly distributed over the cathode. But at high currents the reduction takes place only at the surface of the cathode. Due to this, the reaction sites in the internal of the cathode become unreachable. This also results in a drop of the effective capacity of the battery.

In an attempt to get grip on these physical battery processes, a variety of models has been proposed. The simplest models are purely analytical and similar to Peukert's law. With more detail, so-called equivalent electrical circuit models have been proposed, that can be evaluated (simulated) using a package such as Spice [9]. With even more detail, electro-chemical models have been proposed; although these models can be very accurate for predicting battery lifetime under concrete loads, these models are often too large and complicated to be used as part of high-level system models [10]. Recently, also stochastic models have been proposed,in which the battery charge is discretised and in which probabilistic transitions between charge levels are included to account for the above presented effects [11]. With these, in essence, Markovian models, also the effect of workload variations (around a given mean) have been studied [12].

The approach that we follow in this paper is different. We see the battery capacity as a fixed starting point, and describe the various processes that use energy (the normal workload, including its non-linear effects, restoration, and so on) as stochastic processes (continuous-time Markov chains) with rewards rates, both positive and negative, and possibly time and charge-level dependent, that account for the battery discharge. In doing so, we cast the battery lifetime problem in the well-known framework of performability evaluation (or completion-time evaluation), albeit in an extended model class, namely one in which reward and transition rates may be time- and accumulated reward-dependent.

III. Model and Measure

Homogeneous case. A (homogeneous) Markov reward model (MRM) consists of a finite state space $S = \{1, \ldots, N\}$, the transition rate matrix $\mathbf{Q} \in \mathbb{R}^{N\times N}$ and a reward vector $\underline{r} \in \mathbb{R}^N$.

The matrix $\mathbf{Q}$ is an infinitesimal generator matrix, i.e., the entries $q_{i,j} \geqslant 0$, $j \neq i$, and $q_{i,i} = \sum_{j\in S, j\neq i} q_{i,j}$. The diagonal entry $q_{i,i}$, which is often denoted as q_i, describes the rate at which state i is left. This rate is to be interpreted as the rate of a negative exponential distribution, i.e., the probability that state i is left within s seconds is given as $1-e^{-q_i\cdot s}$. The next state then is j with probability $q_{i,j}/q_i$. The initial distribution of states at time $t = 0$ is denoted as $\underline{\alpha}$. The generator matrix $\mathbf{Q}$ together with $\underline{\alpha}$ determines the CTMC $X(t)$.

When in state i, reward is accumulated with rate r_i. The total reward accumulated when residing in state i from time t_1 until time $t_2 \geqslant t_1$ is denoted $y_i(t_1, t_2)$ and equals

$$y_i(t_1, t_2) = r_i \cdot (t_2 - t_1).$$

Given the state process $X(t)$, the accumulated reward at time t, $Y(t)$, is defined as

$$Y(t) = \int_0^t r_{X(s)} ds.$$

The distribution of $Y(t)$, the so-called performability distribution [13], [14], equals

$$F^Y(t, y) = \Pr\{Y(t) \leqslant y\}.$$

The corresponding density (with respect to y) equals

$$\begin{aligned} f^Y(t, y) &= \frac{\partial F^Y(t, y)}{\partial y} \\ &= \lim_{h\to 0} \Pr\{y \leqslant Y(t) \leqslant y+h\}. \end{aligned}$$

Inhomogeneous case. In the inhomogeneous case, the transition rate matrix $\mathbf{Q}$ and the reward vector $\underline{r}$ can depend on the time t (time-inhomogeneous) and the accumulated reward y (reward-inhomogeneous). We then have $\mathbf{Q}(t, y)$ and $\underline{r}(t, y)$, where y is the current level of accumulated reward. The reward accumulated between time t_1 and $t_2 \geqslant t_1$ when residing completely in state i is described by the following differential equation with initial value $y_i(t_1, 0) = 0$:

$$\frac{dy_i(t_1, t_2)}{dt_2} = r_i(t_2, y_i(t_1, t_2)).$$

The accumulated reward until time t in this case is defined as

$$Y(t) = \int_0^t r_{X(s)}(t, Y(t)) ds.$$

An MRM can easily have more than one reward structure. State i is then equipped with reward rates $r_{i,1}$ through $r_{i,K}$, i.e., we have a reward matrix $\mathbf{R}(t, \underline{y}) \in \mathbb{R}^{N\times K}$ for $\underline{y} \in \mathbb{R}^K$. The accumulated reward is then a vector of random variables $\underline{Y}(t) = (Y_1(t), \ldots, Y_K(t))$ and its distribution is defined as

$$F^{\underline{Y}}(t, (y_1, \ldots, y_K)) = \Pr\{Y_1(t) \leqslant y_1, \ldots, Y_K(t) \leqslant y_K\}.$$

Battery case. For the modelling of a battery we need an MRM that is time-homogeneous but reward-inhomogeneous and has two types of rewards. We therefore denote the generator matrix as $\mathbf{Q}(y_1, y_2)$ and the reward rates as $\mathbf{R}(y_1, y_2) \in \mathbf{R}^{N\times 2}$. The reward accumulated in a state i between time t_1 and time t_2 is described by the following differential equations with initial values $y_{i,1}(t_1, 0) = y_{i,2}(t_1, 0) = 0$:

$$\begin{aligned} \frac{dy_{i,1}(t_1, t_2)}{dt_2} &= r_{i,1}\left(y_{i,1}(t_1, t_2), y_{i,2}(t_1, t_2)\right), \\ \frac{dy_{i,2}(t_1, t_2)}{dt_2} &= r_{i,2}\left(y_{i,1}(t_1, t_2), y_{i,2}(t_1, t_2)\right). \end{aligned}$$

The accumulated reward is then defined as

$$\begin{aligned} \underline{Y}(t) &= (Y_1(t), Y_2(t)) \\ &= \int_0^t \underline{r}_{X(s)}(\underline{Y}(s)) ds \\ &= \int_0^t \left(r_{X(s),1}\left(Y_1(s), Y_2(s)\right), r_{X(s),2}\left(Y_1(s), Y_2(s)\right)\right) ds, \end{aligned}$$

and its distribution equals

$$F^{(Y_1,Y_2)}(t, y_1, y_2) = \Pr\{Y_1(t) \leqslant y_1, Y_2(t) \leqslant y_2\}. \quad (1)$$

We assume that the accumulated rewards have to be non-negative and are bounded by a minimum $\underline{l} = (l_1, l_2)$ and a maximum $\underline{u} = (u_1, u_2)$. This is absolutely reasonable when considering batteries because their charge is always between 0 and a predefined capacity C. We then have

$$\begin{aligned} f^{(Y_1,Y_2)}(t, y_1, y_2) = 0 \qquad &\text{for } y_1 < l_1 \text{ or } y_2 < l_2 \\ &\text{or } y_1 > u_1 \text{ or } y_2 > u_2 \qquad (2) \end{aligned}$$

In the following we often consider the joint distribution of state and accumulated rewards, that is,

$$F_i(t, y_1, y_2) = \Pr\{X(t) = i, Y_1(t) \leqslant y_1, Y_2(t) \leqslant y_2\},$$

with density $f_i(t, y_1, y_2)$. The distribution of the accumulated rewards can then be calculated using

$$F^{(Y_1,Y_2)}(t, y_1, y_2) = \sum_{i\in S} F_i(t, y_1, y_2).$$

IV. Battery models

We model a battery-powered device as a MRM with two reward types. The CTMC states $\{1, \ldots, N\}$ of the MRM reflect the different operating modes of the device. Reward rates of the first type indicate the power consumption in each state i. These reward rates $r_{i,1}$ are non-negative and homogeneous, that is, they are constant over the time and the accumulated reward.

Fixing a capacity C for the battery, a naive evaluation of the battery lifetime is already possible: the battery is empty at time t, if the accumulated reward at time t exceeds C. This does, however, not take into account the more complex behaviour of the battery capacity. This comes into play via the reward rates of the second type. Whenever the power consumption rate $r_{i,1}$ is positive in state i, the maximum battery capacity is going to degrade with a negative rate $r_{i,2}(y_1, y_2)$. This rate depends on the current values of the already consumed power $y_1$1 and on the current maximum capacity y_2 as follows: if the difference between consumed power and capacity $y_2 - y_1$ is still big, we want to have a slow decrease of the capacity. In case the difference is small, that is, the battery is almost empty, we want to have a quick decrease. Additionally, the decrease should be proportional to the current consumption rate. We therefore chose

$$r_{i,2}(y_1, y_2) = -r_{i,1} \cdot \frac{1}{y_2 - y_1}, \quad r_{i,1} > 0, \tag{3}$$

as inhomogeneous reward rate of the second type for state i. If the reward rate $r_{i,1}$ is zero (no power is consumed in state i), the battery capacity increases with a constant rate c, but y_2 can never exceed the residual capacity $C - y_1$:

$$r_{i,2}(y_1, y_2) = \begin{cases} c, & y_2 < C - y_1, \\ 0, & \text{otherwise.} \end{cases} \qquad (r_{i,1} = 0).$$

The interesting question for battery-powered devices is "When does the battery get empty"? For the MRM model, the battery is empty at time t if the consumed power $Y_1(t)$ reaches the current capacity $Y_2(t)$, that is, if $Y_1(t) \geqslant Y_2(t)$. Since the accumulated rewards $Y_1(t)$ and $Y_2(t)$ are random variables, we can only indicate the *probability* that the battery is empty at time t by integrating over the density:

$$\begin{aligned} &\Pr\{\text{battery empty at time } t\} \\ &= \Pr\{Y_1(t) \geqslant Y_2(t)\} \\ &= \int_0^C \int_0^{x_1} f^{(Y_1,Y_2)}(t, x_1, x_2)\, dx_2\, dx_1. \end{aligned} \tag{4}$$

Figure 1 shows how the depletion of a battery takes place: if energy is consumed (increase of the lower curve), the battery capacity degrades (decrease of the upper curve). If there is no energy consumption, the capacity recovers. If the overall amount of consumed energy meets the actual capacity, no further current can be drawn from the battery - the end of its lifetime.

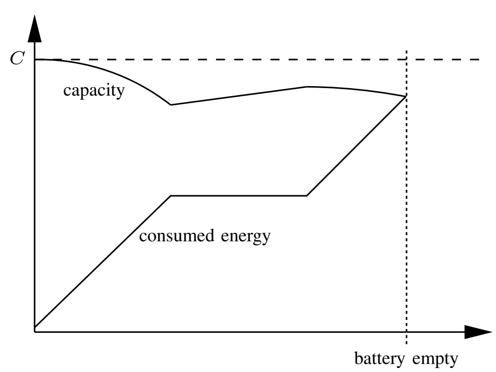

Fig. 1. Example battery lifetime

In the following we consider two workload models of battery powered devices. The first, simple one consists of three states as depicted in Figure 2. At the beginning the model is in `idle` (0) state. With rate $\lambda = 2$ per hour there is the necessity to send data over the wireless interface. If such data is present, the model moves into the `send` (1) state. The sending of data is complete in 10 minutes on average (resulting in a sending rate of $\mu = 6$ per hour). From the idle state the device can also move into an power-saving `sleep` 92) state, this is done – on average – once per hour ($\tau = 1$). The power-consumption rate is low when idling ($r_{0,1} = 7$mA), it is high when sending data ($r_{1,1} = 210$mA) and negligible in the sleep state ($r_{2,1} = 0$mA). With a typical battery capacity $C = 800$mAh (check your cell phone!), this means that theoretically the device can be almost 4 hours in send mode or more than 110 hours in idle mode. We set the recover rate

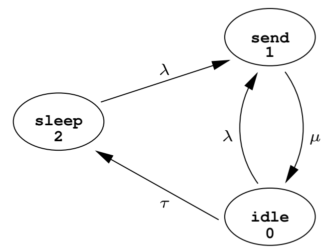

Fig. 2. State transition diagram for the simple model

for the battery capacity if the device is sleeping to $c = 300$.

The battery capacity can only recover if the device is sleeping. To extend the overall battery lifetime it seems to be beneficial to have short periods of high sending activity (bursts) and long periods without sending activity. In the modelled wireless device this could be achieved by accumulating the data to be transmitted and then send all in a row instead of transmitting lower amounts of data more frequently. This can be modelled by switching on and off the flow of arriving data. When the flow is active, data arrives with a very high rate. If the flow is inactive, the device can safely go to sleep. Figure 3 shows a state-transition diagram for such a burst model. It has the same sending rate μ and timeout rate τ as the simple model. Bursts start with rate `switch_on`=1 per hour and stop with rate `switch_off`=6 per hour. To make any results of the two models comparable, we have chosen $\lambda_{\text{burst}} = 182$ such that the steady-state probability to be in `off-send` or `on-send` in the burst model is the same as the probability to be in `send` in the simple model. As could be expected, the steady-state probability to be in `sleep` is higher in the burst model than in the simple model.

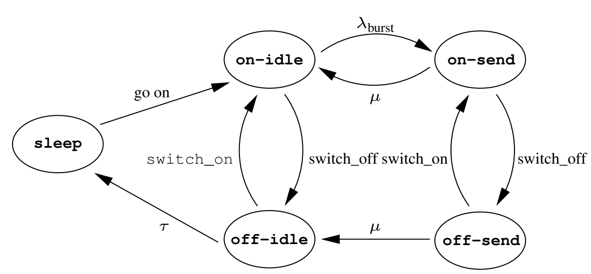

Fig. 3. State transition diagram for the burst model

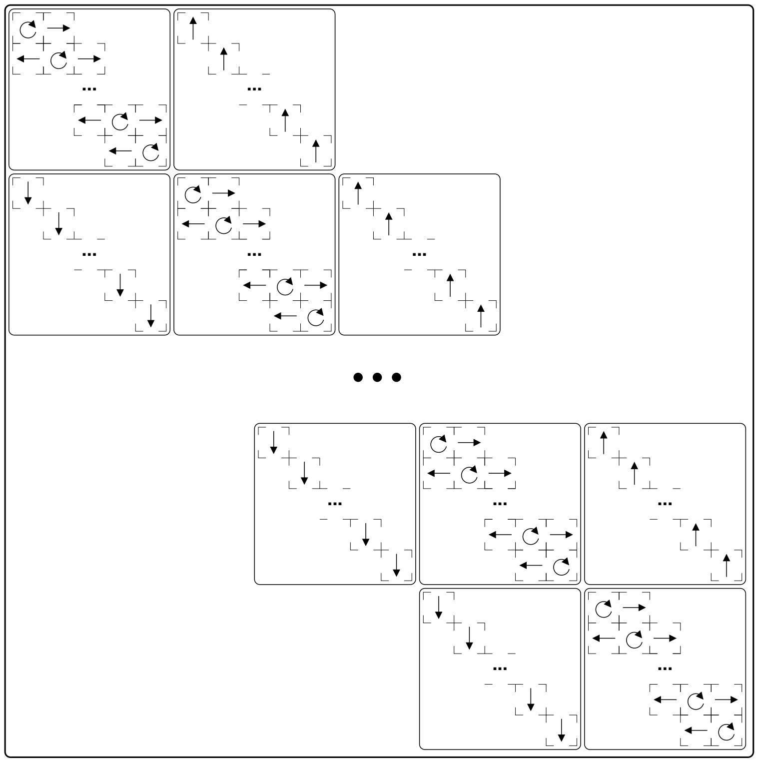

Fig. 4. Structure of the new generator matrix $\mathbf{Q}^*$.

V. Algorithms

In this section we present two algorithms for the computation of the distribution of the accumulated reward in an inhomogeneous Markov reward model. The first algorithm is a Markovian approximation, in which the computation is reduced to the transient solution of a pure CTMC via uniformisation. It was first described for homogeneous MRMs with positive reward rates in the CSRL context [15], [16], then extended to reward-inhomogeneous models with positive reward rates [17]. For battery models we complete the algorithm by allowing multiple reward structures and negative reward rates.

The second algorithm is an extension of the discretisation algorithm of Tijms and Veldman [18]. We adapted their computation scheme to deal with time- and reward-inhomogeneous MRMs including multiple reward structures and negative reward rates [6].

A. Markovian approximation

We approximate the joint distribution of state process and accumulated reward by the transient solution of a derived homogeneous CTMC. It is applicable if the generator matrix and the reward rates depend on the current accumulated reward and not on the current time. This is exactly the case with our battery model and we therefore restrict the presentation to a two dimensional reward structure, even though the approach applies for three or more reward types equally well.

The joint distribution of state and accumulated reward (1) can be rewritten by summing over evenly-sized subintervals of the reward intervals $[l_1, y_1]$ and $[l_2, y_2]$:

$$F_i(t, y_1, y_2) = \sum_{j_1=\frac{l_1}{\Delta}}^{\frac{y_1}{\Delta}-1} \sum_{j_2=\frac{l_2}{\Delta}}^{\frac{y_2}{\Delta}-1} \Pr \left\{ \begin{array}{l} X_t = i, \\ Y_1(t) \in [j_1\Delta, j_1(\Delta+1)], \\ Y_2(t) \in [j_2\Delta, j_2(\Delta+1)] \end{array} \right\}.$$

We want to approximate the terms $\Pr\{X_t = i, Y_1(t) \in [j_1\Delta, j_1(\Delta+1)], Y_2(t) \in [j_2\Delta, j_2(\Delta+1)]\}$ in such a way that the computation is done for a pure CTMC (without rewards). This is accomplished as follows. An MRM modelling a battery can be seen as having an infinite and uncountable state space $S \times \mathbb{R} \times \mathbb{R}$, where state (s, y_1, y_2) indicates that the "CTMC part" of the MRM is in state s and the accumulated reward of the first type is y_1 and of the second type is y_2. For our approximation we break down the uncountable state space to a finite one. Let

$$S^* = S \times \left\{ \frac{l_1}{\Delta}, \dots, \frac{u_1}{\Delta} \right\} \times \left\{ \frac{l_2}{\Delta}, \dots, \frac{u_2}{\Delta} \right\}$$

be the state space of the new CTMC. A state (s, j_1, j_2) then indicates that the MRM is in state s and has accumulated rewards in the intervals $(j_1\Delta, j_1(\Delta+1)]$ and $(j_2\Delta, j_2(\Delta+1)]$, respectively. The initial distribution $\underline{\alpha}^*$ depends on the original initial distribution α and the initial values for the accumulated rewards a_1 and a_2:

$$\alpha^*_{(i,j_1,j_2)} = \begin{cases} \alpha_i, & a_1 \in (j_1\Delta, j_1(\Delta+1)] \text{ and} \\ & a_2 \in (j_2\Delta, j_2(\Delta+1)], \\ 0, & \text{otherwise.} \end{cases}$$

The distribution of the accumulated rewards is then approximated as

$$F^{(Y_1,Y_2)} \approx \sum_{i\in S} \sum_{j_1=\frac{l_1}{\Delta}}^{\frac{y_1}{\Delta}-1} \sum_{j_2=\frac{l_2}{\Delta}}^{\frac{y_2}{\Delta}-1} \pi_{(i,j_1,j_2)}(t),$$

where $\pi_{(i,j_1,j_2)}(t)$ is the transient probability of residing in state (i, j_1, j_2) at time t in the derived CTMC.

For battery models, the probability that the battery is already empty at time t, cf. (4), is approximated as:

$$\Pr\{\text{battery empty at time } t\} \approx \sum_{i\in S} \sum_{j_1=\frac{l_1}{\Delta}}^{\frac{y_1}{\Delta}-1} \sum_{j_2=\frac{l_2}{\Delta}}^{j1} \pi_{(i,j_1,j_2)}(t).$$

Two types of transitions are possible in the new CTMC with generator $\mathbf{Q}^*$: transitions taken from the original CTMC and transitions between different reward levels (for each of the two reward types). An entry in the new generator matrix $\mathbf{Q}^*$ is defined depending on the type of transition it represents. Figure 4 shows the structure of the generator matrix $\mathbf{Q}^*$. Each small block corresponds to a fixed j_1 and j_2 and has dimension $N \times N$, each of the big block corresponds to one j_2.

Transitions from the original generator. If the original CTMC part of two states $(i, j1, j2)$ and $(i', j1, j2)$ are are different $(i \neq i')$ but the reward levels are identical, the entry is taken from the original generator. Since it is a reward-inhomogeneous MRM, the current reward level $(j_1\Delta, j_2\Delta)$ must be taken into account, that is,

$$Q^*_{(i,j_1,j_2),(i',j_1,j_2)} = Q_{i,i'}(j_1\Delta, j_2\Delta).$$

In Figure 4 these entries are found in the blocks :

Transitions indicating a change in the first accumulated reward. If the CTMC states are identical, the levels of the first accumulated reward are different and the levels of the second accumulated reward are again identical, the entry indicates a change in the first accumulated reward. Such a change can only happen between neighbouring levels, hence, either $j_1' = j_1 - 1$ (entries in blocks) or $j_1' = j_1 + 1$ (entries in blocks).

$$Q^*_{(i,j_1,j_2),(i,j_1',j_2)} = \begin{cases} \frac{r_{i,1}(j_1\Delta,j_2\Delta)}{\Delta}, & r_{i,1} > 0 \text{ and} \\ & j_1' = j_1 + 1, \\ \frac{-r_{i,1}(j_1\Delta,j_2\Delta)}{\Delta}, & r_{i,1} < 0 \text{ and} \\ & j_1' = j_1 - 1, \\ 0, & \text{otherwise.} \end{cases}$$

Transitions indicating a change in the second accumulated reward. If $j_2' = j_2 - 1$ or $j_2' = j_2 + 1$ with identical i and j_1, the entries in blocks and are defined as

$$Q^*_{(i,j_1,j_2),(i,j_1,j_2')} = \begin{cases} \frac{r_{i,2}(j_1\Delta,j_2\Delta)}{\Delta}, & r_{i,2} > 0 \text{ and} \\ & j_2' = j_2 + 1, \\ \frac{-r_{i,2}(j_1\Delta,j_2\Delta)}{\Delta}, & r_{i,2} < 0 \text{ and} \\ & j_2' = j_2 - 1, \\ 0, & \text{otherwise.} \end{cases}$$

All other off-diagonal entries of $\mathbf{Q}^*$ are zero, the diagonal entries are defined as the negative row sums.

B. Discretisation

General case. The second algorithm we present approximates the density $f_i(t, \underline{y})$ by discretising both t and $\underline{y}$ in units of size $\Delta > 0$, such that the probability that two ore more state transitions happen within a period of length Δ is negligible. Given Δ, we can approximate the densities $f_i(t, \underline{y})$ as follows (we use $\tilde{f}_i(t, \underline{y})$ to denote the approximation):

$$\begin{aligned} \tilde{f}_i(t, \underline{y}) &= \tilde{f}_i(t-\Delta, \underline{y} - \underline{r}_{i,\cdot}(t, \underline{y})\Delta) \cdot (1 - q_i(t, \underline{y})\Delta) \\ &\quad + \sum_{j\neq i} \tilde{f}_j(t-\Delta, \underline{y} - \underline{r}_{j,\cdot}(t, \underline{y})\Delta) \cdot q_{j,i}(t, \underline{y})\Delta. \end{aligned} \quad (5)$$

The first summand of the right-hand side accounts for the case that no transition took place in the previous Δ time units, meaning that the state at time $t - \Delta$ was also state i and the accumulated rewards were accordingly smaller. The second summand accounts for the case that Δ time units back, the state was $j \neq i$, and hence, the reward over the last Δ time units was accumulated according to the reward rates $\underline{r}_j$. From this equation the need to choose Δ small enough is clear, as in the product with a transition rates $q_{i,j}$ a probability should be the result.

Battery case. In what follows we show how the computation proceeds for the battery models of Section IV. However, the presented algorithm and results generalise to arbitrary numbers of reward rates as well as time-inhomogeneity.

The approximated density function is computed at discretised time steps $\tau = 1, \ldots, t/\Delta$, where discrete time τ corresponds to continuous time instant $\tau\Delta$. The same discretisation is carried out for the accumulated rewards: for the first type we have $j_1 = l_1/\Delta, \ldots, u_1/\Delta$ and for the second type $j_2 = l_2/\Delta, \ldots, u_2/\Delta$. If the initial accumulated rewards at time 0 are (a_1, a_2), the computation is initialised for discrete time $\tau = 1$ by

$$\tilde{f}_i(\Delta, j_1, j_2) = \begin{cases} \alpha_i/\Delta, & j_1 = a_1 + r_{i,1}\Delta \text{ and} \\ & j_2 = a_2 + r_{i,2}\Delta, \\ 0, & \text{otherwise.} \end{cases}$$

In the subsequent iterations for $\tau = 2, \ldots, t/\Delta$, the approximation of the density for all $j_1 = l_1/\Delta, \ldots, u_1/\Delta$ and $j_2 = l_2/\Delta, \ldots, u_2/\Delta$ is computed using (5) and (2).

Figure 5 illustrates the computation of the approximated densities for time instant $\tau + \Delta$ given the densities at time τ. For each state $i \in S$, each x_1 and each x_2 the value of $\tilde{f}_i(\tau + \Delta, x_1, x_2)$ has to be calculated. If the corresponding rates in the generator matrix $\mathbf{Q}(x_1, x_2)$ are non-null, this calculation accesses previously computed values of the density at time τ.

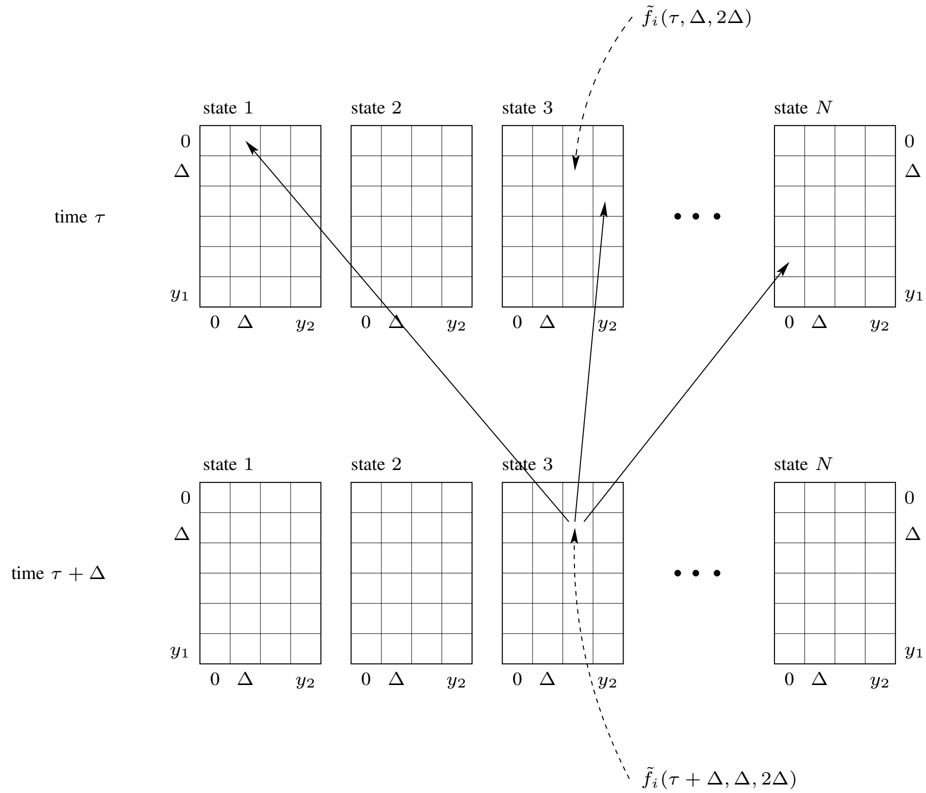

Fig. 5. One iteration in the discretisation algorithm

	discretisation		Markovian approximation	
	space	time	space	time
general model	$\mathcal{O}\left(N \cdot \frac{\prod_{k=1}^{K} y_k}{\Delta^K} \cdot\right)$	$\mathcal{O}\left(N^2 \cdot \frac{t}{\Delta} \cdot \frac{\prod_{k=1}^{K} y_k}{\Delta^K} \cdot\right)$	$\mathcal{O}\left(N \cdot \frac{\prod_{k=1}^{K} y_k}{\Delta^K} \cdot\right)$	$\mathcal{O}\left(N^2 \cdot t \cdot \frac{\prod_{k=1}^{K} y_k}{\Delta^K} \cdot\right)$
battery model	$\mathcal{O}\left(N \cdot \frac{y_1}{\Delta} \cdot \frac{y_2}{\Delta}\right)$	$\mathcal{O}\left(N^2 \cdot \frac{t}{\Delta} \cdot \frac{y_1}{\Delta} \cdot \frac{y_2}{\Delta}\right)$	$\mathcal{O}\left(N \cdot \frac{y_1}{\Delta} \cdot \frac{y_2}{\Delta}\right)$	$\mathcal{O}\left(N^2 \cdot t \cdot \frac{y_1}{\Delta} \cdot \frac{y_2}{\Delta}\right)$

TABLE I

COMPLEXITIES

The scheme only works for *integer rewards*, otherwise one would access values outside the grid of discretised accumulated reward levels. While this is no restriction in theory (one can scale the reward rates accordingly) it substantially increases the size of the grids in practice.

The performability distribution is then approximated as follows

$$F^{(Y_1,Y_2)}(t, y_1, y_2) = \sum_{i \in S} \int_{l_1}^{y_1} \int_{l_2}^{y_2} f_i(t, z_1, z_2)\, dz_2\, dz_1$$
$$\approx \sum_{i \in S} \Delta \cdot \sum_{j_1 = \frac{l_1}{\Delta}+1}^{\frac{y_1}{\Delta}} \sum_{j_2 = \frac{l_2}{\Delta}+1}^{\frac{y_2}{\Delta}} \tilde{f}_i(t, j_1\Delta, j_2\Delta),$$

and the probability that a battery is empty is

$$\Pr\{\text{battery empty at time } t\}$$
$$\approx \sum_{i \in S} \Delta \cdot \sum_{j_1 = \frac{l_1}{\Delta}+1}^{\frac{y_1}{\Delta}} \sum_{j_2 = \frac{l_2}{\Delta}+1}^{j1} \tilde{f}_i(t, j_1\Delta, j_2\Delta).$$

C. Complexity

Table I shows the space and time complexity of the two algorithms both for the general and the battery case. Regarding run time, both algorithms are quadratic in the number of states and linear in time and the two reward bounds. In the battery case, the step size Δ enters as Δ^{-3} for the discretisation algorithm and only as Δ^{-2} for the Markovian approximation. However, for the Markovian approximation, the step size is also coded into the generator matrix of the new CTMC by multiplying the reward rates with $\frac{1}{\Delta}$ (see the definition of $\mathbf{Q}^*$). The transient solution of the CTMC has a time complexity linear in the *uniformisation constant*. For small Δ, this uniformisation constant gets linear in $\frac{1}{\Delta}$ and we thus obtain the same time complexity as for the discretisation algorithm.

VI. Results

We apply the algorithms described in the previous section to the battery models developed in Section IV. All numerical parameters are taken from there.

Figure 6 reflects the battery lifetime in the simple case where the physical aspects leading to degradation and recovery of battery capacity are not considered. The battery is empty if the accumulated reward of the first (and only) type reaches the battery capacity $C = 800$. On the x-axis we have different time instants t, on the y-axis is the probability that the battery is already empty at time t, that is, Pr $\{Y_1(t) > C\}$. For comparison reasons we first computed the curve for the simple model with the uniformisation-based algorithm by Sericola [19], these values are exact with respect to an error of 10^{-3}. The other two curves, one for the simple and one for the burst model, are computed with the Markovian approximation. As one can see for the simple model, for small t, the values are underestimated while for bigger t they are overestimated. Unfortunately we are not able to provide an exact curve for the burst model, as the run time of that algorithm is prohibitive.

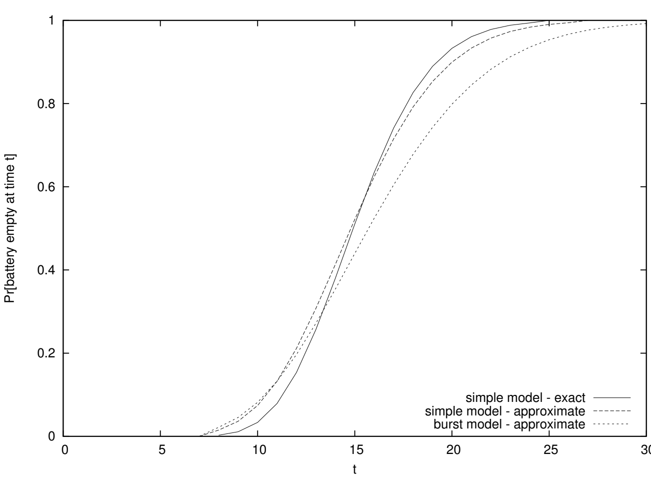

Fig. 6. Battery lifetime without degradation and recovery

For the computation of the curves, we took the steady-state distributions as initial distributions for the two models. Remember that the probability to be in sending mode is the same for both models such that they are comparable. However, the burst model spends more time in the sleep state (37% vs 25%) and so we would expect the curve of the burst model to be always below the curve of the simple model. This is not the case, the curves are crossing, when computed with the Markovian approximation. This seems to be an effect of the approximate nature of the algorithm.

As could be expected, the probability increases with t, reaching values close to one (meaning that the battery is almost surely empty) at about $t = 25$ for the simple model and $t = 30$ for the burst model.

Figure 7 shows the lifetime distribution taking the degradation effect alone, as well as degradation plus recovery into account. For both the simple and the burst model, the lifetime of the battery is decreases if degradation comes into play. This means that the curves have a steeper increase and reach 1 for smaller t than in the simple case depicted in Figure 6. The second pair of curves in Figure 7 shows the lifetime distribution with degradation *and* recovery. As expected, these curves are between the ones in Figure 6 and the ones without recovery.

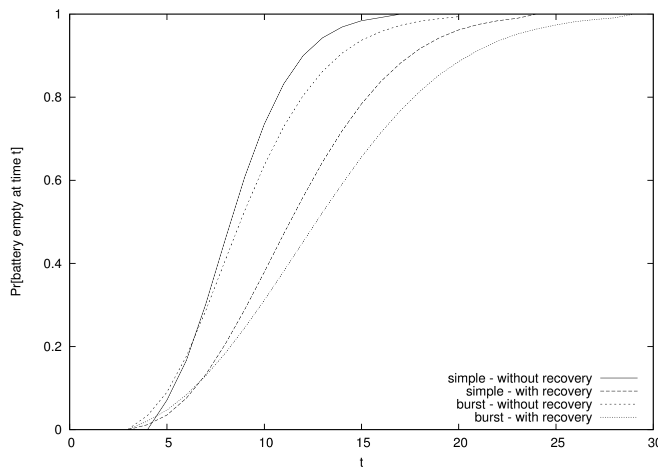

Fig. 7. Battery lifetime without and with recovery

With Figure 8 we illustrate the advantage the burst model has over the simple model if recovery is taken into account: it shows the difference between the curves without and with recovery, normalised to the value without recovery. For increasing t, the value of the burst model is substantially greater than the value for the simple model. This can be attributed to the fact that the burst model spends more time in the sleep state, thereby allowing for longer periods of recovery.

The numbers shown so far have been computed using the Markovian approximation with $\Delta = 16$. It has proven to be a powerful algorithm for the computation of this type of measures. The runtimes never exceeded one minute. Unfortunately we were not able to apply the discretisation algorithm. It requires $\Delta < \frac{1}{6}$ and for the burst model even $\Delta < \frac{1}{188}$, because Δ must be smaller than any diagonal entry in the generator matrix. In the time complexity Δ appears as $\frac{1}{\Delta^3}$, this leads to prohibitive run times. Additionally, the discretisation algorithm is only suitable for integer rewards. However, the degradation rate (3) is in general not an integer value. Scaling rewards and capacity accordingly would result in even longer run times.

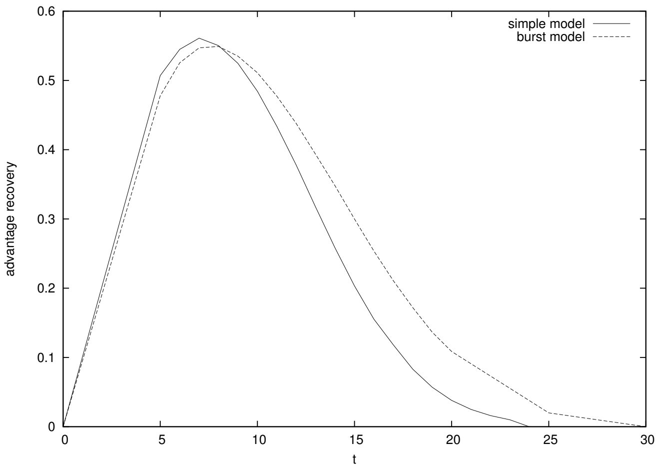

Fig. 8. Benefit of recovery

VII. Conclusion

The aim of this paper has been twofold. First of all we have discussed the increasing importance of the incorporation of battery aspects into system models. In particular, we stressed the need for incorporating the nonlinear aspects of the battery models. We have shown that all these aspects can be taken into account in a time- and reward-inhomogeneous Markov reward model setting. Efficient algorithms to compute the performability distribution in such models were not available. The second goal of this paper was the presentation of two approaches towards such algorithms, namely a Markovian approximation and a discretisation algorithm. Even though the two algorithms have the same space and time complexity, the case study has shown that they are not equally applicable: the discretisation approach suffered from the fact that it requires quite a small step size which in turn led to prohibitive run times. In contrast, the Markovian approximation performed very well. Using this algorithm we could show how the nonlinear aspects of batteries affect the system lifetime for two different workload models. Further work will include the exploration of realistic parameters for the battery model and the evaluation of real world power-aware devices.

References

[1] *IEEE Computer*, vol. 38, no. 11. IEEE Press, 2005.

[2] N. van Dijk, "Uniformisation for nonhomogeneous markov chains," *Operations Research Letters*, vol. 12, 1992.

[3] A. Rindos, S. Woolet, I. Viniotis, and K. Trivedi, "Exact methods for the transient analysis of nonhomogeneous continuous time markov chains," in *2nd International Workshop on the Numerical Solution of Markov Chains*, 1995, pp. 121–133.

[4] A. van Moorsel and K. Wolter, "Numerical solution of non-homogeneous markov processes through uniformisation," in *Proceedings of the 12th European Simulation Multiconference*, 1998, pp. 710–717.

[5] M. Telek, A. Horváth, and G. Horváth, "Analysis of inhomogeneous markov reward models," *Linear Algebra and its Applications*, vol. 386, pp. 383–405, 2004.

[6] B. Haverkort and J. Katoen, "The performability distribution for nonhomogeneous markov-reward models," in *Proceedings 7th Performability Workshop (PMCCS'05)*, 2005, pp. 38–42.

[7] D. Rakhmatov and S. Vrudhula, "An analytical high-level battery model for use in energy management of portable electronic systems," in *Proceedings of the International Conference on Computer Aided Design (ICCAD'01)*, 2001, pp. 488–493.

[8] Overview of lithium ion batteries. [Online]. Available: `http://www.panasonic.com/industrial/battery/oem/images/pdf/Panasonic_LiIon_Overview.pdf`

[9] The spice page. [Online]. Available: `http://bwrc.eecs.berkeley.edu/Classes/IcBook/SPICE/`

[10] K. Lahiri, A. Raghunathan, S. Dey, and D. Panigrahi, "Battery-driven system design: a new frontier in low power design," in *7th Asia and South Pacific Design Automation Conference (ASP-DA'02)*, 2002, pp. 261–267.

[11] C. Chiasserini and R. Rao, "Pulsed battery discharge in communication devices," in *Proceedings of the 5th International Conference on Mobile Computing and Networking*, 1999, pp. 88 – 95.

[12] D. Panigrahi, C. Chiasserini, S. Dey, R. Rao, A. Raghunathan, and K. Lahiri, "Battery life estimation of mobile embedded systems," in *Proceedings of the 14th International Conference on VLSI Design*, 2001, pp. 57 – 63.

[13] J. F. Meyer, "On evaluating the performability of degradable computing systems," *IEEE Transactions on Computers*, vol. 29, no. 8, pp. 720–731, 1980.

[14] ——, "Performability: a retrospective and some pointers to the future," *Performance Evaluation*, vol. 14, no. 3, pp. 139–156, 1992.

[15] B. R. Haverkort, H. Hermanns, J.-P. Katoen, and C. Baier, "Model checking CSRL-specified performability properties," in *Proceedings of the 5th International Workshop on Performability Modeling of Computer and Communications Systems (PMCCS'01)*, 2001, pp. 105–109.

[16] B. R. Haverkort, L. Cloth, H. Hermanns, J.-P. Katoen, and C. Baier, "Model checking performability properties," in *Proceedings of the International Conference on Dependable Systems and Networks (DSN'02)*. IEEE Press, 2002, pp. 102–112.

[17] L. Cloth, "Model Checking Algorithms for Markov Reward Models," Ph.D. dissertation, University of Twente, 2006.

[18] H. Tijms and R. Veldman, "A fast algorithm for the transient reward distribution in continuous-time Markov chains," *Operations Research Letters*, vol. 26, pp. 155–158, 2000.

[19] B. Sericola, "Occupation times in Markov processes," *Communications in Statistics — Stochastic Models*, vol. 16, no. 5, pp. 479–510, 2000.