Fluid Queue Models of Battery Life

Gareth L. Jones, Peter G. Harrison, Uli Harder, Tony Field
Department of Computing, Imperial College London
Huxley Building, 180 Queen's Gate, London SW7 2AZ, UK
E-mail:{gljones,uh,ajf,pgh}@doc.ic.ac.uk

***Abstract*—We investigate how a power-save mode affects the battery life of a device subject to stochastically determined charging and discharging periods. We use a multi-regime fluid queue, imposing a threshold at some value. When the power level falls below the threshold, (for example, 20% of charge remaining) a power-save mode is entered and the rate of discharge decreased. An expression for the Laplace transform of the battery life's probability density function is found and inverted numerically in particular instances. We show the life of battery can be significantly improved by the introduction of the power-saving threshold.**

I. Introduction

Rechargeable batteries are commonplace in modern electronic devices, such as mobile phones, digital cameras, remote sensors and communication satellites. Whilst the physical processes of the battery are well understood a key challenge is to be able to model accurately the remaining charge of a battery. Knowledge of the remaining charge can be usefully used by the control processes of battery-operated devices, for example to determine the optimal sample rate of a sensor between recharging periods or whether to perform non-essential functions whilst a communication satellite is in the earth's shadow.

An electrical current is essentially formed by moving a very large number of electrons along a conductor and it is quite natural to think of the flow of charge as an incompressible fluid. Batteries can then be thought of as 'reservoirs' for this fluid which is 'drained' into appliances, the sinks of the system. The charging process adds fluid to reservoir. In this paper we use a fluid queueing model to investigate how the battery life can be improved by introducing a threshold for the remaining battery charge at which the consumer throttles their rate of consumption. This might be used, for instance, to reduce the light output of a display, the sample rate of a sensor, or the speed of a vehicle.

We use a *multi-regime fluid queue* where the fluid arrival process is governed by a separate Continuous Time Markov Chain and fluid arrival rate vector whose various parameters are allowed to depend on the current reservoir level. Specifically, in this paper we partition a finite reservoir, or 'buffer', $[0, B]$ into two parts, $(0, V)$ and (V, B), and allow the rate parameters of the arrival process to depend on which partition the buffer level is in. We show that the introduction of a threshold for the charge level, below which the device draining the battery reduces its demands, can lead to a significant improvement in battery life.

A. Fluid queueing models

A fluid queueing model is a pair of stochastic processes $\{(X(t), Z(t))|t \geq 0\}$ together with a vector $\mathbf{r}$. Typically, $Z(t)$ is a continuous time Markov chain (CTMC) with generator matrix Q and state space $\mathcal{S} = \{0, 1, 2, \ldots, n-1\}$. We usually call $Z(t)$ the *environment process* or *background process*. When $Z(t)$ is in state i, the evolution of the process satisfies

$$\frac{\mathrm{d}X(t)}{\mathrm{d}t} = r_i$$

when the process is not at a boundary value. $\mathbf{r}$ is a vector with dimension the size of the state space of $Z(t)$. Karandikar [1] provide a good survey of steady state results for these models.

The *busy period* or *wet period* [2] is the time, starting from the instant that fluid first enters the buffer, until the buffer is again empty. This quantity is a special case of a *hitting time*. The hitting time is the length of time, starting at a given fluid level, which we must wait until we hit another specified level. The busy period is the hitting time of 0 starting at 0.

The distribution of the busy period was first published for a discrete model by Borel (1942) and in fluid queues by Kendall [3] and is now regarded as "one of the salient features of queueing theory" [4]. A recent paper by one of the authors has calculated this result in a setting where the background process is a continuous time semi-Markov chain (CTSMC) [5] allowing very general distributions for holding times of filling states. In this paper we restrict ourselves to CTMCs and thus exponentially distributed holding times.

The busy period result is of interest for use in tandem network approximations [6], when modelling the waiting time of non-priority traffic in a priority traffic system [7] and the time until ruin of an insurance firm [8]. In this paper we calculate a hitting time and use it to model the battery life.

B. Multi-regime fluid queues

In multi-regime fluid queues the rate vector $\mathbf{r}$ and generator matrix of the Markov process $Z(t)$ are allowed to depend to some extent on x, the current level of the buffer process [9]. The buffer level space is partitioned and the process is

allowed to depend on which partition the buffer process is in at time t.

Kankaya et al. [9] define a model with K thresholds $0 = V_0 < V_1 < V_2 < V_3 < \cdots < V_{K-1} < V_K = B < \infty$ and when $V_{k-1} < X(t) < V_k$ say that the process $X(t)$ is in regime k at time t. When $X(t)$ is in regime k the process behaves like a fluid queue with background process generator matrix Q_k and rate vector $\mathbf{r}_k$. The model introduced by these authors also allows for delays to be defined on the boundaries. In this work we do not use the full model, instead restricting ourselves to the situation where for a given i only $\mathbf{r}_k$ depends on k. We further impose that for every value of k, either $r_{k,i} > 0$, or for all k either $r_{k,i} > 0$.

In [9] conditions are given to rule out anomalous situations and guarantee the model is well formed. For example, we do not permit the construction of a regime where at the lower boundary we are forced upwards, but at the upper boundary we are forced down as this would lead to the process making "infinitesimal oscillations."

Further, the equilibrium distribution is derived for this model and a numerically stable algorithm proposed. Examples in the paper [9] show situations where this approach is a significant improvement on the widely used spectral approach [10].

We use multi-regime fluid queues to model power reserves in an electronic device subject to random charging periods. For the example introduced in this paper we consider only a single threshold value so drop the subscript and refer to the lower boundary 0, the threshold value V and upper boundary B. During a charging period, power arrives at charge-rate λ and during non-charging periods no power arrives. The device regularly uses power at rate μ_h but when reserves fall below a threshold value V, the device enters a power-save mode, using power at rate μ_ℓ, so we label the two regimes h and ℓ for high and low use.

Entering a power-save mode (with lower performance) but greater longevity may be desireable in many situations. For example, sending data wirelessly consumes a lot of power. Sending less frequent data for a longer period might be more preferable to running the battery low and then having a period where no data is recorded.

C. Related Battery Models

Most other work concentrates on improving battery life due to an interesting feature of discharging batteries which allows the battery to regain some voltage during a period of rest if the battery is not completely discharged. During discharge the voltage of a battery does not remain constant and drops due to the polarization effect. A diffusion process counteracts this during an idle time. To predict battery life (the remaining charge) accurately one needs to account for this effect correctly. This can for instance be done by chemical or electrical models, however they tend to have a large number of battery specific parameters (for references to this approach see for instance section 2 of [11]). These models tend to be very accurate.

Chiasserini and Rao in [12], [13] model the recovery behaviour with a discrete time Markov chain. In the simplest model in [12] the battery life is modelled by $N+1$ steps representing the remaining charge and with constant transition probabilities. Battery recovery is incorporated in this model by constant probabilities to step backwards in the chain. The later model makes this more realistic by having non-constant probabilities that depend on the remaining charge level for the recovery process.

In [11] Jongerden and Haverkort discuss more models like the kinetic battery model (KiBaM) [14] and the diffusion model [15] (see [11] for more references to these two models). The kinetic model uses a tandem of reservoirs of fluid to represent the charge of the battery. The battery is drained from the second reservoir and the recovery effect is modeled by the first reservoir that can only add fluid to the second reservoir at rate dependent on the difference of the height in fluid in the two reservoirs. When the second reservoir is empty the battery is fully discharged. Jongerden and Haverkort show that the kinetic mode is an approximation of the more general diffusion model [15]. They also extend their work to encompass battery scheduling which allows to prolong the battery life by having several cells which are drained in rotation to optimise the battery life by using the recovery effect in [16].

In our model we have a simple charge and drain process and can think of the recovery effect as included in the charging gained during charging periods. Ideally this would be modeled more explicitly to make the model more realistic.

II. Our Battery Model

We model the charge level in a battery in a device subject to random charging and discharging periods, a situation that has not been discussed in the literature previously as far as the authors are aware.

In our model a single fluid reservoir represents the remaining charge of such a rechargeable battery. This battery has a finite capacity of B. The device is in an environment with two states which causes the battery to experience periods of charging followed by periods of discharge. The fluctuations in this process could for instance be caused by varying cloud cover if a photovoltaic device is used to charge the battery. The battery is discharged either at rate μ_h or at rate μ_ℓ, depending on whether the level of charge in the battery is above or below a threshold value V. (To achieve this we assume that the device using the battery can run in two different modes, which use energy at a higher or lower rate respectively.)

The background process $Z(t)$ is a continuous time Markov chain with state space $\mathcal{S}$ and generator matrix Q. We assume that Q is not dependent on the charge level in the battery, but it would be straightforward to implement such

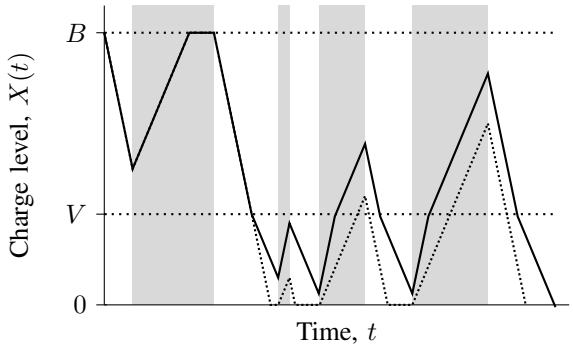

Figure 1. This figure illustrates the evolution of the charge level of the battery in time. B is the maximum charge of the battery, V is the threshold triggering the lower discharge rate. The battery is charging during the grey periods with a rate of $\lambda = 3$ and discharges above V with $\mu_h = 2$. When the threshold V is crossed the discharged rate is lowered to $\mu_\ell = 1$. The dotted line shows what would happen to the charge level if the rate was not changed below the threshold.

dependence and make these parameters partition-dependent [9]. As earlier the rate vector $\mathbf{r}_k$ contains elements $r_{k,i}$ describing the net fluid rate when the the background process $Z(t)$ is in state i and the fluid in regime $k = h$ or ℓ.

A sample trace for the charge level over time in a two-state model is shown in Figure 1. The periods with grey highlight indicate charging periods, while the periods without highlight are non-charging periods. The dotted line shows the behaviour of the battery level over time without the power-saving regime, and the solid line shows the effect of the power-saving regime. In this model there are four net inflow rates as shown in the table below.

$r_{k,i}$	$i = 0$ non-charging period	$i = 1$ charging period
$V < x \leq B$ so $k = h$	$-\mu_h$	$\lambda - \mu_h$
$0 \leq x \leq V$ so $k = \ell$	$-\mu_\ell$	$\lambda - \mu_\ell$

The *battery life* is the first passage time, starting with charge level B, that the process $X(t)$ hits 0.

III. Analysis & Results

We derive and solve differential equations in the style of [17], which the first passage time distribution satisfies. The battery life is a special case of a first passage time, starting at B and ending at 0. We use the same upper and lower boundary conditions defined therein and note that at the threshold value V, the Laplace transform (of the density function with respect to t) must be continuous in both x and w. We later calculate moments of battery life directly from it and numerically invert it for some sample model parameter values.

The hitting time required is

$$T = \inf\{t > 0 \mid X(t) = 0\}$$

and we define the probability distribution function of T, conditioned on the initial state of the background process $Z(t)$ and fluid level,

$$F_i(t, x) = \mathbb{P}(T \leq t \mid Z(0) = i, X(0) = x).$$

Then the backward equation gives[1]

$$F_i(t,x) = (1 + q_{ii}\delta)F_i(t - \delta, x + r_{k,i}\delta) + \sum_{j \in \mathcal{S},\, j \neq i} q_{ij}\delta F_j(t - \delta, x + r_{k,i}\delta) + o(\delta)$$

and it is elementary to obtain the following differential equation for $x, t > 0$ by rearranging algebraically and taking the limit as $\delta \to 0$ (see [17], for example):

$$\frac{\partial F_i(t,x)}{\partial t} - r_{k,i}\frac{\partial F_i(t,x)}{\partial x} = \sum_{j \in \mathcal{S}} q_{ij}F_j(t,x).$$

In general, $\mathcal{S}$ is the state space of the background process $Z(t)$, and in our model $\mathcal{S} = \{0, 1\}$. The boundary conditions are $F_i(t, 0) = 1$ if $t \geq 0$ and $r_{\ell,i}(0) \leq 0$, $\quad F_i(0, x) = 0$ if $x > 0$ $\quad$ and $F_i(0, 0) = 0$ if $r_{\ell,i}(0) > 0$.

Let $H_i(w, x) = \int_0^\infty e^{-wt}F_i(t, x)\mathrm{d}t$ be the Laplace transform of $F_i(t, x)$ with respect to t and $\mathbf{H}(w, x)$ a vector of such transforms indexed by the elements of $\mathcal{S}$, the state space of the background process. Then taking Laplace transforms, we obtain

$$-D_k\frac{\partial \mathbf{H}(w,x)}{\partial x} + w\mathbf{H}(w,x) = Q\mathbf{H}(w,x)$$

where $D_k = \mathrm{diag}(\mathbf{r}_k)$, with boundary condition

$$H_i(w, 0) = 1/w \text{ if } r_{\ell,i} \leq 0$$

by direct calculation. We consider this equation in two intervals, referred to as "low" and "high": $(0, V)$ and (V, B). For each interval, we solve the requisite differential equation over the domain $(0, \infty)$ and then restrict the solution to the interval under consideration. Thus, for $0 < x < \infty$, we solve the pair

$$\begin{aligned} -D_\ell\frac{\partial \mathbf{H}^\ell(w,x)}{\partial x} + w\mathbf{H}^\ell(w,x) &= Q\mathbf{H}^\ell(w,x) \\ -D_h\frac{\partial \mathbf{H}^h(w,x)}{\partial x} + w\mathbf{H}^h(w,x) &= Q\mathbf{H}^h(w,x). \end{aligned}$$

Taking the Laplace transform with respect to x (of the function $H_i^h(w, x + V)$ in the second equation), we obtain

$$\begin{aligned} (wI - D_\ell\theta)\mathbf{H}^{\ell *}(w,\theta) + D_\ell\mathbf{H}^\ell(w,0) &= Q\mathbf{H}^{\ell *}(w,\theta) \\ (wI - D_h\theta)\mathbf{H}^{h *}(w,\theta) + D_h\mathbf{H}^h(w,V) &= Q\mathbf{H}^{h *}(w,\theta) \end{aligned}$$

where the Laplace transforms of the vectors $\mathbf{H}^\ell(w, x)$ and $\mathbf{H}^h(w, x)$ have components

$$\mathbf{H}_i^{\ell *}(w,\theta) = \int_0^\infty e^{-\theta x}H_i^\ell(w,x)\mathrm{d}x$$

[1]The version published in the MASCOTS proceedings erroneously states $1 + q_{ii}$ rather than $1 + q_{ii}\delta$, which has been corrected here.

and

$$\begin{aligned}\mathbf{H}_i^{h*}(w,\theta) &= \int_0^\infty e^{-\theta x} H_i^h(w, x+V)\mathrm{d}x \\ &= e^{\theta V}\int_V^\infty e^{-\theta y} H_i^h(w,y)\mathrm{d}y\end{aligned}$$

respectively. Thus the solution required is given by

$$\begin{aligned}\mathbf{H}^{\ell*}(w,\theta) &= M_\ell^{-1} D_\ell \mathbf{H}^\ell(w,0) \\ \mathbf{H}^{h*}(w,\theta) &= M_h^{-1} D_h \mathbf{H}^h(w,V) \qquad (1)\end{aligned}$$

where $M_k = Q - wI + \theta D_k$ for $k = \ell, h$. We solve each of these equations as follows.

First, $M_k^{-1} = C_k(\theta)/\Delta_k(\theta)$ where $\Delta_k(\theta)$ is the determinant of M_k and $C_k(\theta)$ is the transpose of its cofactor matrix. For simplicity, assuming that the roots (for θ) of the characteristic equation $\Delta_k(\theta) = 0$ are distinct[2], we have

$$\Delta_k(\theta) = \prod_{i\in S} r_{k,i}(\theta - e_{k,i}(w))$$

where the $e_{k,i}(w)$ are the aforesaid roots. Notice that $\{e_i\}$ are the eigenvalues of the matrix $-(Q-wI)D_k^{-1}$ (because $|Q - wI + \theta D_k| = 0$ implies $|-(Q-wI)D_k^{-1} - \theta I| = 0$).

Expanding the terms $C_k(\theta)/\Delta_k(\theta)$ in partial fractions (in each partition), we find that

$$\mathbf{H}^{\ell*}(w,\theta) = \sum_{i\in S} \frac{A_{\ell,i}(w) D_\ell \mathbf{H}^\ell(w,0)}{\theta - e_{\ell,i}(w)}$$

and

$$\mathbf{H}^{h*}(w,\theta) = \sum_{i\in S} \frac{A_{h,i}(w) D_h \mathbf{H}^h(w,V)}{\theta - e_{h,i}(w)}$$

where the elements of the matrices $A_{\ell,i}(w), A_{h,i}(w)$ are the ith coefficients in the partial fraction expansions of the corresponding elements of the transposed cofactor matrices C_ℓ, C_h divided by the determinants $\Delta_\ell(\theta)$, $\Delta_h(\theta)$ (i.e. of the inverse matrices M_ℓ^{-1}, M_h^{-1}), in the lower and upper partitions respectively.

We can now invert the Laplace transform with respect to x by inspection to get:

$$\mathbf{H}^\ell(w,x) = \sum_{i\in S} A_{\ell,i}(w) D_\ell \mathbf{H}^\ell(w,0) e^{e_{\ell,i}(w)x}$$

and

$$\mathbf{H}^h(w,x+V) = \sum_{i\in S} A_{h,i}(w) D_h \mathbf{H}^h(w,V) e^{e_{h,i}(w)x}.$$

[2]It is relatively straightforward to extend the method to multiple roots, see [5]. However, in practice, it is rare to find two numerically almost equal roots, especially taking into account that the characteristic equation's parameters are functions of w.

A. Boundary conditions

When the number of states in the background process $|S| = n$, we require $2n$ unknown quantities in the solution of equations (1). For a given background state i, we assume that the sign of $r_{k,i}$ is the same for all k, so drop the subscript k in this section. Let $n_e(x)$ be the number of "emptying states" i for which $r_i < 0$; similarly, the number of filling states is $n_f = n - n_e$, the states j for which $r_j \geq 0$. Our assumption means that a given background process state is either always an emptying or filling state, we exclude the mixed case. We therefore find n continuity conditions at the threshold $x = V$,

$$\mathbf{H}^h(w,V) = \mathbf{H}^\ell(w,V). \qquad (2)$$

As noted above, there are also n_e equations in the boundary condition at level $x = 0$,

$$H_i^\ell(w,0) = 1/w \text{ if } r_i \leq 0. \qquad (3)$$

We therefore require a further n_f equations, which come from the upper boundary $x = B$.

To find these, let the random variable T_i denote the hitting time of fluid level 0, starting from fluid level B and state i at time 0. Then T_i is the battery life. For background process states i where $r_i \geq 0$,

$$\begin{aligned}H_i^h(w,B) &= \int_0^\infty e^{-wT_i} F_i(t,B)\mathrm{d}t \\ &= E[e^{-wT_i}] \\ &= E[E[e^{-w(Y_i+T_j)} \mid J = j]]\end{aligned}$$

where Y_i is the (residual) holding time in state i and J is the state entered by the CTMC after i. Hence,

$$H_i^h(w,B) = \sum_{j\in S} p_{ij} E[e^{-wY_i}] E[e^{-wT_j}]$$

since the random variables Y_i and T_j are independent, where $p_{ij} = q_{ij}/q_i$, for $i \neq j, q_i = -q_{ii}$, is the transition probability from state i to state j, $p_{ii} = 0$ $(1 \leq i,j \leq n)$. Thus, we may write

$$\begin{aligned}H_i^h(w,B) &= \sum_{j\in S} p_{ij}\frac{q_i}{w+q_i} H_j^h(w,B) \\ &= \sum_{j\in S\,:\,r_j\geq 0} p_{ij}\frac{q_i}{w+q_i} H_j^h(w,B) \\ &\quad + \sum_{j\in S\,:\,r_j<0} p_{ij}\frac{q_i}{w+q_i} H_j^h(w,B).\end{aligned}$$

Defining the sub-vectors

$$\begin{aligned}\mathbf{L}^f(w) &= (H_j^h(w,B) \mid r_j \geq 0) \\ \mathbf{L}^e(w) &= (H_j^h(w,B) \mid r_j < 0)\end{aligned}$$

the $n_f \times n_f$ matrix

$$P^f(w) = (p_{ij}q_i/(w+q_i) \mid 1 \leq i,j \leq n; r_i, r_j \geq 0)$$

and the $n_f \times n_e$ matrix

$$P^e(w) = (p_{ij}q_i/(w+q_i) \mid 1 \le i,j \le n; r_i \ge 0, r_j < 0)$$

we obtain

$$\mathbf{L}^f(w) = \Big(I^f - P^f(w)\Big)^{-1} P^e(w)\mathbf{L}^e(w)$$

where I^f is the $n_f \times n_f$ identity matrix. This is a set of n_f further constraints, completing the specification of the solution. These equations were implemented in Mathematica and used to compute the analytical results presented in the following sections.

B. Cumulative distribution function (CDF)

To investigate the CDF of the battery life, $F_0(B,t)$, we can either invert the Laplace transform $H_0^h(w,B)$, obtained in the previous section, or else simulate the battery system; doing both provides mutual validation of the analytic and simulation models. Our simulation was developed as a Matlab program.

We return to the earlier two-state example with background process $\mathcal{S} = \{0,1\}$ where state 0 is a discharging state and state 1 a charging state. We choose a charging rate of $\lambda = 3$, discharging rate $\mu_h = 2$ in the upper partition (above the threshold value V) and $\mu_\ell = 1$ in the lower partition. The maximum charge of the battery is $B = 1$. We set the background process' generator matrix

$$Q = \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}.$$

For the examples in this paper we calculate the battery life starting with a full battery in a discharging state 0. It would be straightforward to find the same results for other starting situations.

The CDF shown in Figure 2 is plotted for $V = 0$ (dashed), $V = 0.1$, $V = 0.3$ and $V = 0.5$. The value $V = 0$ (dashed) corresponds to the situation where the power-save mode is not used, while the $V = 0.1$, $V = 0.3$ and $V = 0.5$ correspond to 10%, 30% and 50% threshold values. All four CDFs exhibit a jump at a positive t value. This is because there is a minimum time in which a full battery can fully discharge. If the discharge is not interrupted by any charging then it must discharge an amount $B-V$ of charge at rate μ_h and then the remaining V charge at rate μ_ℓ. As we decrease μ_ℓ this minimum time increases. In Figure 2 we see a strict improvement in battery life performance for any increase in the threshold value V.

With this parameterization, when there is no threshold we can read off the CDF in Figure 2 that 95% of battery lives are finished by 3.5 time units and that with the threshold at $V = 0.5$ this is increased to 6 time units.

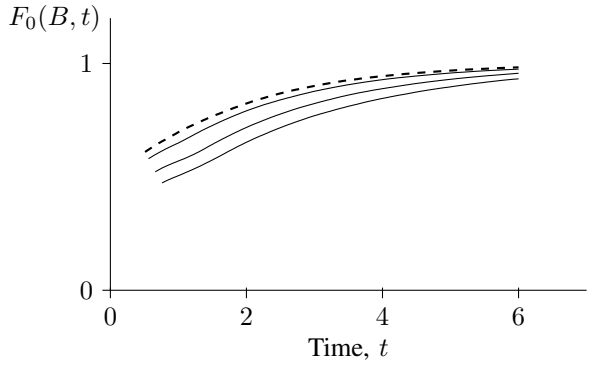

Figure 2. CDF of battery life, $F_0(B,t)$ for different threshold values. The dashed line represents the case of no threshold, lines beneath are thresholds at $V = 0.1$, $V = 0.3$ and $V = 0.5$. Note the jump at the start. This is the probability that the initial non-charging period is long enough for the battery life to end with no charging periods having occurred. The parameters here are set to $\lambda = 3$, $\mu_h = 2$, $\mu_\ell = 1$, $B = 1$ and $\alpha = \beta = 1$.

C. Mean battery life

With the same parameters chosen in the previous section, we investigate the mean battery life as a function of the threshold V for different low discharge rates μ_ℓ. Using symbolic differentiation of the Laplace transform evaluated at $w = 0$, this is plotted in Figure 3; recall that μ_ℓ is the power use during the power-save mode. We set $\mu_h = 2$ so that, for example, $\mu_\ell = 0.5$ corresponds to a power-save state that uses just 25% of the energy used in the high power state. We believe such a parameter is reasonable in the context of reduced data collection and reporting of a wireless sensor. Figure 3 shows that, as the efficiency of the power-save mode increases (a decreasing value of μ_ℓ), the gain from using the threshold increases significantly – more than linearly.

The data on average battery life for different parameters can also help in battery selection. Presumably the user will know something about the environment a sensor will be placed in (thus determining the generator Q which controls the background process). The parameters μ_h and μ_ℓ are constrained by the design of the system and λ is determined by the design of both the system and the environment in which it operates. This leaves V and possibly B for selection.

Suppose the sensor required an average battery life of at least 4.4 time units. Then from Figure 4 we can determine that the battery's capacity would need to be at least 0.6. To give the required battery life, a battery of capacity exactly 0.6 would always have to run at the low power mode, whereas using a battery of capacity 0.8 could deliver the required average life with a threshold as low as 63% of capacity. Figure 4 highlights different combinations of battery capacity B and threshold value V that would satisfy this requirement.

Further moments can be calculated by using the standard

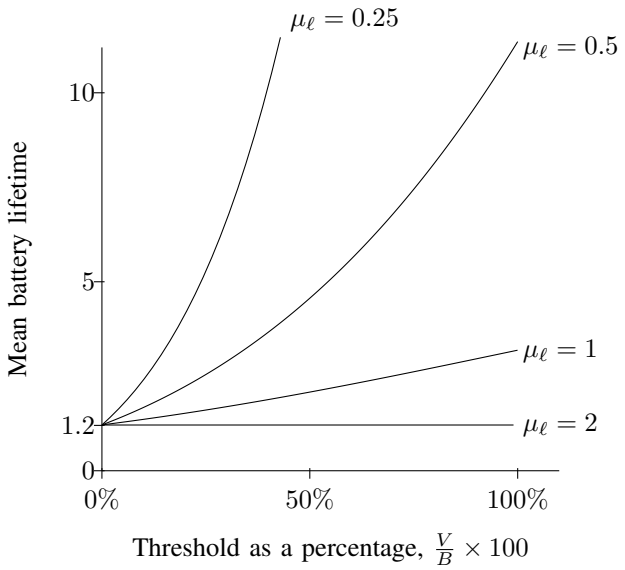

Figure 3. Mean battery life plotted against threshold value V for different values of μ_ℓ, the rate of power use in power-save mode. When $V = 0$ no threshold is used and the device always runs in high power mode. Lines correspond to different efficiencies in power-save mode, with smaller μ_ℓ being more efficient. The line labelled $\mu_\ell = 2$ is flat because in this example $\mu_h = 2$ threshold value makes no difference. The parameters here are set to $\lambda = 3$, $\mu_h = 2$, μ_ℓ as labelled, $B = 1$ and $\alpha = \beta = 1$.

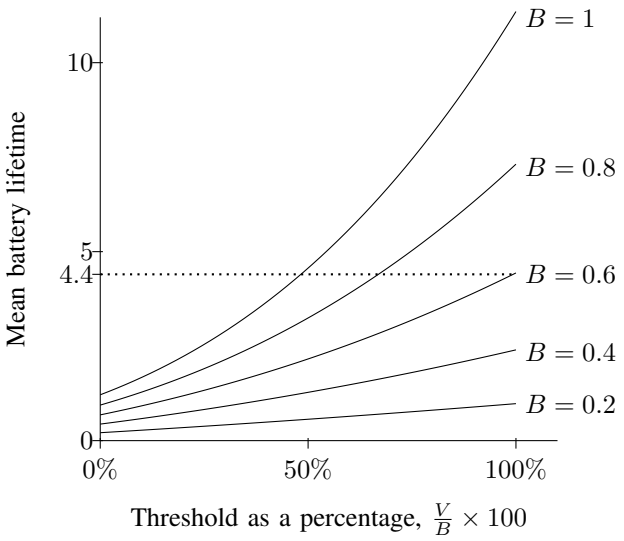

Figure 4. Mean battery life plotted against threshold value V for different battery capacities, B. We plot $B = 0.2, 0.4, 0.6, 0.8, 1$ (from bottom to top). The dashed line shows how larger batteries with a low threshold can have the same mean battery life as smaller batteries with a higher threshold value. The parameters here are set to $\lambda = 3$, $\mu_h = 2$, $\mu_\ell = 1$, B as labelled and $\alpha = \beta = 1$.

relationship that[3]

$$E(T_i^k) = (-1)^k \left. \frac{\mathrm{d}^k w H_i^h(w, B)}{\mathrm{d}w^k} \right|_{w=0}$$

where $H_i^h(w, B)$ is as computed in section III.

D. Simulation

In this section we present simulation results to test the robustness of the analytical model, by investigating scenarios where the model's Markovian assumptions are not satisfied. We consider the case where the charging periods have a heavy tail and are log-normally distributed, which is currently outside the scope of our model, being a non-exponential random variable.

We choose parameters for the log-normal distribution so that the charging period distributions have the same mean value, 1, as the exponential model. Figure 6 shows the CDFs computed for the battery life by simulation. The higher of the two CDFs corresponds to the case where charging periods have (an appropriately scaled) log-normal distribution, while the lower of the two curves corresponds to case where charging periods are exponentially distributed with parameter 1.

Note that both these CDFs start at the same point $(0.95, 0.39)$ as this point corresponds to the situation where the battery life ended without any charging periods. The non-charging period is exponentially distributed with parameter 1 in both cases, so these CDFs meet here.

Figure 5 shows how these different charging period distributions effect the mean battery life. Though both charging distributions have mean 1, the mean battery life is significantly lower when the charging periods have a heavy tail. (The non-charging periods have the same distribution in both cases.) The results show value in extending our model to non-Markovian situations.

Simulation results *not* shown here for the CDF and mean of battery lives serve to (mutually) validate our theoretical work.

IV. Conclusions and Future Work

We have shown that imposing a power-save mode, when the charge in a rechargeable battery falls below a threshold level V, can significantly increase battery life. The more efficient the power-save mode, the greater these gains are, and they grow in a relationship that is faster than linear. We also showed how the value of the threshold level can be set so as to optimize some power requirement of a battery's usage.

The mathematical modelling technique used enables graphs, such as the ones we have displayed, to be plotted rapidly – literally in a few seconds. This is in contrast with simulation, which we used for mutual validation with the analytical model and to examine its robustness in situations where the underlying assumptions did not all hold; specifically non-exponential charging times. The underlying

[3]The version published in the MASCOTS proceedings erroneously states $\frac{\mathrm{d}^k H_i^h(w,B)}{\mathrm{d}w^k}$ rather than $\frac{\mathrm{d}^k w H_i^h(w,B)}{\mathrm{d}w^k}$, which has been corrected here.

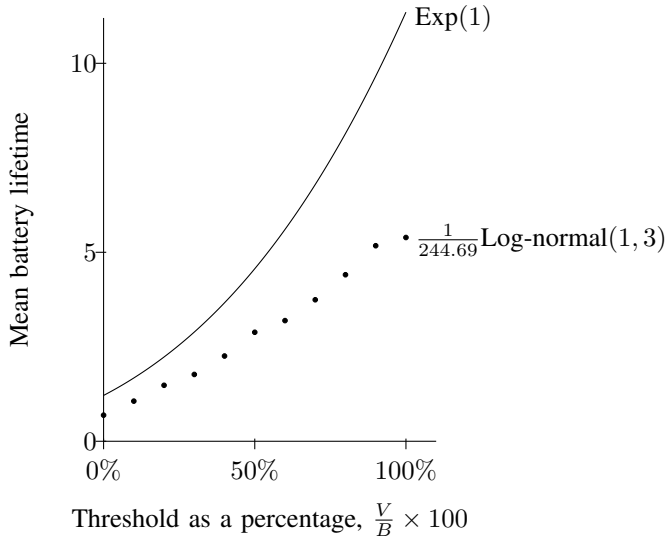

Figure 5. Mean battery life against threshold value V for different charging period distributions. Both on period distributions have mean 1. In the upper curve the distribution is exponential with parameter 1, in the lower points a log-normal distribution. The parameters here are set to $\lambda = 3$, $\mu_h = 2$, $\mu_\ell = 0.5$, $B = 1$, $\beta = 1$ with charging period distribution either Exp(1) or $\frac{1}{244.69}$Log-normal$(1,3)$.

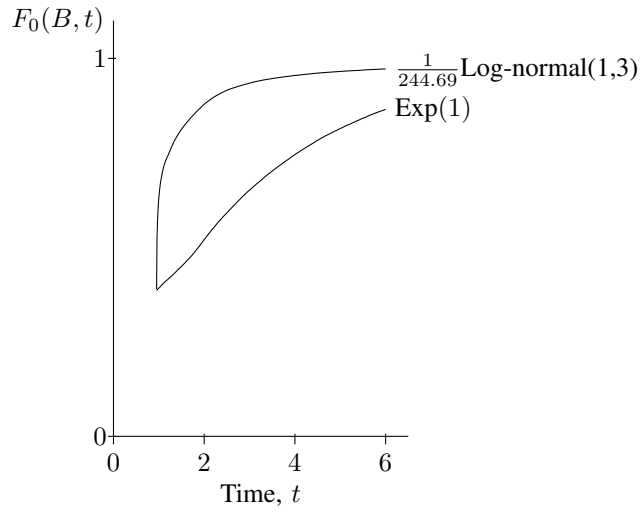

Figure 6. Simulated CDF showing the effect of changing the charging period distribution. All simulations run with $B = 1$, $\lambda = 3$, $\mu_h = 2$, $\mu_\ell = 1$, $\beta = 1$ and $V = 0.5$. The charging period for the lower curve is exponentially distributed with parameter 1 as earlier in this paper. The charging period for the upper curve is a scaled log-normal distribution such that both on period distributions have mean 1.

theory is relatively simple, based on a pair of linear partial differential equations with constant coefficients, one for each partition. The interesting part is the handling of the boundary conditions, where one term is known at the lower (no charge) boundary, continuity holds at the boundary between the partitions (giving two equations) and a further condition holds at the upper (fully charged) boundary to represent the fact that any charging has no effect there. This leads to four linear equations for the four unknown parameters (actually functions of the Laplace-parameter), which are readily solved by standard software (Mathematica in our case), giving highly complex expressions. The resultant Laplace transforms of the density functions can also be differentiated multiple times by mathematical software and so yield the mean and higher moments of the battery life.

It is interesting to compare the work with that of [5], where the busy period of a similar fluid queue is investigated, even when fluid "filling times" (corresponding to charging times here) may have any distribution. With some effort, the method used – using "pseudo-busy periods" which are identically distributed at any fluid level – could be adapted. (This is an important extension as exponential distributions are rarely accurate models of reality.) The finite capacity reservoir/battery means that the pseudo-busy periods are no longer identical and so have to be parameterized with their starting (and finishing) fluid level. We believe that, with care, this can be done and that, in fact, we would then be able to calculate analytically the results we simulated in the robustness tests. However, if the discharge times were to be non-exponential, there is no known analytical approach, these being "emptying" times, so simulation would be the only mode of solution.

One way in which the battery model could be extended is to include the case where the charge rate is smaller than the discharge rate, $\lambda < \mu$, and to apply the model in more complicated situations using a larger state space for $Z(t)$, the background process. In the former case, the battery would always be discharging at some rate, so the model would describe the dynamics of the charge in a standard, non-rechargeable battery. Using more partitions would also allow us to model phenomena like recovery more accurately, as for instance in [11].

Using the entire model presented in [9] requires us to allow the generator matrix Q to depend on the partition and for delays to occur on transition across a boundary between partitions. Service level agreements also need to be incorporated into our model. We have shown how we can find parameters that satisfy certain requirements on mean battery life, but ideally we should be able to cope with more complicated requirements such as "more than 950 devices out of a batch of 1,000 deployed fail within 1,000 time units in less than 1% of batches." Using a power-save mode imposes a cost on the system. Either data collection is less frequent or processing is slower. In future work, we intend

to investigate this problem by finding the optimal threshold V in terms of a utility function modelling this cost.

There are no doubt more applications for our model. Rather than viewing the model as a power-saving regime when reserves get low, if we view the process as modelling the buffer contents of a router/job processing node in a computer network, we can view the discharge rate μ_ℓ as the regular rate of processing and μ_h as a more expensive (but faster rate) – used when the backlog at the node gets too long (and possibly near the capacity), for example. If we extend the model to account for the cost of this extra resource, we can again phrase an optimization problem for the waiting time of jobs arriving at the queue, subject to some restriction on the costs of running the server.

Finally, we intend to incorporate the kinetic battery model (KiBaM) into our model [14]. This could utilize results on tandem fluid queues, which have been researched before, for example in [18], and KiBaM is essentially a tandem pair of fluid queues because the bound charge can only be released into the active charge, but not to the consumer directly. Similarly, we hope to be able to extend the model to one with several cells to investigate how our threshold improves battery performance in that situation.

REFERENCES

[1] R. L. Karandikar and V. G. Kulkarni, "Second-order fluid flow models: reflected Brownian motion in a random environment," *Operations Research*, vol. 43, no. 1, pp. 77–88, 1995.

[2] E. Y. Lee and K. K. J. Kinateder, "The expected wet period of finite dam with exponential inputs," *Stochastic Processes and their Applications*, vol. 90, no. 1, pp. 175–180, Nov. 2000.

[3] D. G. Kendall, "Some Problems in the Theory of Dams," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 19, no. 2, pp. 207–233, 1957.

[4] A. M. K. Tarabia, "A new formula for the busy period of a non-empty multiserver queueing system," *Applied Mathematics and Computation*, vol. 143, no. 2-3, pp. 401–408, Nov. 2003.

[5] A. J. Field and P. G. Harrison, "Busy periods in fluid queues with multiple emptying input states," *Journal of Applied Probability*, vol. 47, pp. 474–497, 2010.

[6] S. Aalto, "Characterization of the Output Rate Process for a Markovian Storage Model," *Journal of Applied Probability*, vol. 35, no. 1, pp. 184–199, 1998.

[7] D. P. Kroese and W. R. W. Scheinhardt, "Joint Distributions for Interacting Fluid Queues," *Queueing Systems*, vol. 37, no. 1, pp. 99–139, 2001.

[8] K. Ito, Ed., *Encyclopedic Dictionary of Mathematics*, 2nd ed. The MIT Press, 1993.

[9] H. E. Kankaya and N. Akar, "Solving Multi-Regime Feedback Fluid Queues," *Stochastic Models*, vol. 24, no. 3, pp. 425–450, Jul. 2008.

[10] V. G. Kulkarni, "Fluid models for single buffer systems," *Frontiers in Queueing: Models and Applications in Science and Engineering*, pp. 321–338, 1997.

[11] M. R. Jongerden and B. R. Haverkort, "Which battery model to use?" *IET Software*, vol. 3, no. 6, pp. 445–457, 2009.

[12] C. F. Chiasserini and R. R. Rao, "A model for battery pulsed discharge with recovery effect," in *IEEE Wireless Communications and Networking Conference*, 1999, pp. 636–639.

[13] ——, "Energy Efficient Battery Management," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 7, pp. 1235–1245, 2001.

[14] J. F. Manwell and J. G. McGowan, "Lead acid battery storage model for hybrid energy systems," *Solar Energy*, vol. 50, no. 5, pp. 399–405, 1993.

[15] D. Rakhmatov, S. Vrudhula, and D. A. Wallach, "Battery lifetime prediction for energy-aware computing," in *International Symposium on Low Power Electronics and Design*. ACM, 2002, pp. 154–159.

[16] M. Jongerden, A. Mereacre, H. Bohnenkamp, B. Haverkort, and J. P. Katoen, "Computing Optimal Schedules of Battery Usage in Embedded Systems," *IEEE Transactions on Industrial Informatics*, no. 99, pp. 1–11, 2010.

[17] N. Barbot, B. Sericola, and M. Telek, "Distribution of Busy Period in Stochastic Fluid Models," *Stochastic Models*, vol. 17, no. 4, pp. 407–427, 2001.

[18] V. Gupta and P. G. Harrison, "Fluid level in a reservoir with an on-off source," *SIGMETRICS Performance Evaluation Review*, vol. 36, no. 2, pp. 128–130, 2008.