AMUEM 2007 – International Workshop
on Advanced Methods for Uncertainty Estimation in Measurement
Sardagna, Trento, Italy, 16-18 July 2007

Prediction of the Lifetime of a Battery-Operated Video Node Using Offline and Online Measurements

G. Zhang, R. Manduchi
Department of Computer Engineering, University of California, Santa Cruz
{gefan,manduchi}@soe.ucsc.edu

Abstract *– This paper presents a model for determining the first and second order statistics of the lifetime of a battery-operated node. This model assumes knowledge of the statistics of duration and consumption of the tasks that can be performed by the system, as well as of the statistics of the task occurrences and of the remaining charge in the battery. The latter may be monitored by an onboard sensor. Knowledge of the remaining lifetime is very important in order to devise an online control strategy assuring that specific requirements are met.*
Keywords *– Wireless camera networks, lifetime prediction.*

I. INTRODUCTION

Wireless networks of cameras have received a good deal of attention recently for applications such as wide area surveillance and environment monitoring [1]. The video nodes in this type of networks are battery-operated, which demands careful strategies of energy management. Indeed, video nodes typically have relatively high power consumption, due to tasks such as image acquisition, image processing/compression, and data transmission. The control policy of a video node thus has a critical role in the average energy drawn from the battery and ultimately in the lifetime of the node itself.

Typical control policies are based on reactive duty cycles. For example, the node may take and process images at regular periods of time. If an image contains interesting or suspicious elements (e.g. an intruder in a controlled space), then the image, a portion thereof, or simply an alarm message may be transmitted to the base station. A video node may also receive alert messages from another node, in which case it may temporarily modify the parameters of its own duty cycle. For example, in order to track a moving intruder, nodes detecting the intruder may send alert messages to the nodes that are most likely to see the intruder next [4], which in turn may start taking images at a higher rate.

A similar policy is adopted by our testbed wireless camera network, "Meerkats". Each Meerkats node is based on the Crossbow's Stargate platform, which has an XScale PXA255 CPU (400 MHz) with 32MB flash memory and 64MB SDRAM. PCMCIA and Compact Flash connectors are available on the main board. The Stargate also has a daughter board with Ethernet, USB and serial connectors. We equipped each Stargate with an Orinoco Gold 802.11b PCMCIA wireless card and a Logitech QuickCam Pro 4000 webcam connected via USB. The QuickCam can capture video with resolution of up to 640x480 pixels. The operating system is Stargate version 7.3 which is an embedded Linux system (kernel 2.4.19).

This extended abstract describes a method to predict the remaining lifetime of a Meerkats node operating under a reactive duty cycle, assuming that the remaining charge in the battery is known. Remaining lifetime estimation is critical for network management. For example, if at some point the estimated lifetime is less than what is required by the application, then some duty cycle parameters (e.g., the rate of image acquisition) may be changed in order to reduce the average energy consumption. It should be clear that external factors (e.g., the rate at which moving objects are detected in the monitored area) play an important role in the energy consumed by a reactive system. Although these factors cannot typically be well modeled a priori, their statistics can be measured online by the system, and adaptive strategies can be put in place.

II. MODELS AND MEASUREMENTS

Our approach to modeling the energy consumed by a video node is based on the definition of a set of M tasks $\mathbf{S}$. At each point in time, the node is assumed to be performing only one such task. The *elementary tasks* for Meerkats are described in [4]; they include processes such as acquiring and processing an image, transmitting a certain amount of data, remaining in 'sleep' mode for a certain period of time, as well as transition tasks such as activating or deactivating the webcam and putting the CPU in 'sleep' mode. A duty cycle is defined as a sequence of elementary tasks, with branching points controlled by external events (the presence of a target in an image or an alert message from another node). In the context of this paper, a 'task' is any maximal sequence of elementary tasks in a duty cycle that does not include a branching point. Note that, whereas an elementary task as defined in [4] typically does not repeat itself twice in a row, 'tasks' as defined here may indeed be repeated consecutively.

Each task S is characterized by two random variables (r.v.)[1]: the energy $e(S)$ consumed by the task, and the duration $t(S)$ of the task. Note that this approach is rather

[1] We denote r.v. by lower case letters and constant values by upper case letters. We denote probability density functions (pdf) with the lower case letter p, and mass distributions with the upper case letter P.

different from the one taken by [2]. The latter models the system behavior as a Markov chain of states with fixed duration and consumption. The problem with the formalism of [2]. is that the duration and consumption of a task (defined as a maximal sequence of identical states) are forced to behave as geometric random variables. This is a critical limitation of this model, which makes it unsuitable to represent the distributions of energy and duration observed in practice.

The goal of the offline measurement phase [4,5,6] is precisely to characterize the statistical properties of $e(S)$ and $t(S)$. The sequence of tasks (states) forms a random process. The state evolution process could be modeled by Markov chain, which implies that the energy consumed after a certain period of time is as a function of a semi-Markov chain, defined by the transition probabilities of the tasks and by the distributions of $e(S)$ and $t(S)$ [3]. However, we use a simpler approach here, and model the task sequence as a Bernoulli scheme. In other words, we assume that the next task is independent of the current and previous tasks. This simplification is justified by two main considerations. First, there are no absorbing states in the sequence. Second, given that the sequence of states before the battery is depleted is typically very long, we are only interested in asymptotical behaviors, which are well described by marginal state distributions. Sec. III shows how the mean and variance of the remaining lifetime t_0, given that the remaining charge in the battery is E_0, can be computed using the distribution of state occurrences, $P_s(S)$, and the mean and variance of $e(S)$ ($\mu_{e|S}, \sigma^2_{e|S}$) and of $t(s)$ ($\mu_{t|S}, \sigma^2_{t|S}$).

The purpose of the online measurement phase described in Sec. II B is to measure the statistics of task occurrences as well as to continuously track E_0 (and thus the expected value of t_0) under varying external conditions.

A. Offline Measurements

The task taxonomy for Meerkats, along with procedures to compute the mean and standard deviation of each task's duration and consumption, are described elsewhere [5,6]. We briefly summarize the measurement methodology here. The HP34401A digital multimeter (DMM) was used to measure the current flow to the system while the different benchmarks were executed, with the system powered at 5.6 Volts by a HP E3631A power supply. The duration and energy consumption for a specific instance of a task were measured based on the time profile of the current measured against a steady state 'idle' profile. Note that the standard deviation of duration and consumption are not negligible for several tasks, which reflects into a possibly large variance of the remaining lifetime estimate.

B. Online Measurements

The online measurement of the marginal probabilities of task occurrences can be performed directly by the system while operating. For the measurement of the remaining charge in the battery we rely on a DS2438 chip, which is available on the main board of the Stargate. Two kernel modules ('onewire' and 'batmon') provide access to the battery monitor chip and retrieve information about the battery's current state. The DS2438 measures the voltage across the Stargate's power source and the current flowing out of the battery and into the Stargate. Current measurement is achieved by measuring the voltage at the ends of a 270 mΩ shunt resistor. This provides a granularity of 0.9 mA with a ±1.8 mA error. For more details, the reader is referred to [6].

III. PREDICTION OF REMAINING LIFETIME

Once the statistics of duration and consumption for each task have been measured, one can attempt to estimate the remaining lifetime based on the available charge E_0 in the battery. We outline our method in the following. A number of simplifying assumptions are used in our computation, three of which are listed below, while the remaining one is deferred till Sec. III A:

1. The energy consumption and duration for a given task are mutually independent normal random variables.
2. The task sequence is a Bernoulli scheme.
3. The energy consumption and duration of the current task only depend on the task type.

We already commented on the second assumption. The third assumption seems reasonable if the tasks are defined carefully. The first assumption, however, is certainly questionable: If a task at some point requires more (or less) time than average, it seems likely that it will use more (or less) energy than average as well. We plan to investigate the correlation between energy and duration of tasks in the future, and to improve our model accordingly. Modelling these two marginal distributions as Gaussians seems

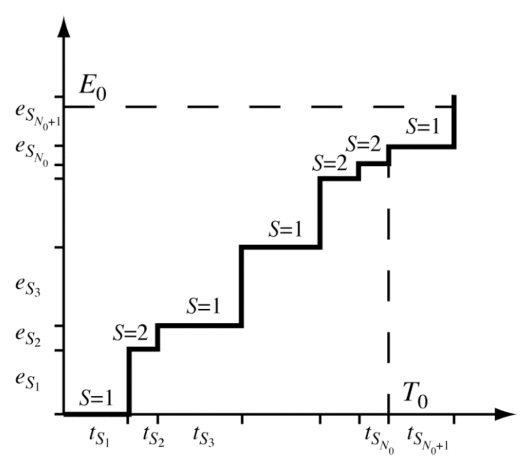

reasonable as a first approximation, as long as their standard deviations are much smaller than their means (since both energy and durations are positive variables).

Our derivation of the statistical properties of the remaining lifetime t_0 given E_0 is broken down into three steps. First we derive the statistical properties (probability density function (pdf) $p_{e|n}(E|N)$, mean $\mu_{e|N}$ and variance $\sigma^2_{e|N}$) of the energy consumed after N tasks in the sequence. Then we compute the distribution $P_{n_0}(N_0)$ of the remaining length n_0 of the task sequence given that the remaining charge is E_0, along with its mean μ_{n_0} and variance $\sigma^2_{n_0}$. Finally, we derive the properties of the lifetime t_0 ($p_{t_0}(T_0)$, μ_{t_0}, $\sigma^2_{t_0}$).

A. Computation of $p_{e|n}(E|N)$, $\mu_{e|N}$, $\sigma^2_{e|N}$

Using the theorem of total probability, and remembering that the task occurrences are independent ($P_{s_1\ldots s_N}(S_1\ldots S_N)=\prod_{i=1}^{N}P_s(S_i)$) we obtain:

$$p_{e|n}(E|N)=\sum_{S_1\in\mathbf{S}}P_s(S_1)\ldots\sum_{S_N\in\mathbf{S}}P_s(S_N)p_{e|s_1\ldots s_N}(E|S_1\ldots S_N) \quad (1)$$

where S_i represents the task at the *i*-th step in the sequence, and $p_{e|s_1\ldots s_N}(E|S_1\ldots S_N)=g(E;\mu_{e|S_1\ldots S_N},\sigma^2_{e|S_1\ldots S_N})$, where $g(x;\mu,\sigma^2)$ is the normal pdf. Under our previous assumptions, $\mu_{e|S_1\ldots S_N}=\sum_{i=1}^{N}\mu_{e|S_i}$ and $\sigma^2_{e|S_1\ldots S_N}=\sum_{i=1}^{N}\sigma^2_{e|S_i}$.

The mean of the energy consumed after N tasks is

$$\begin{aligned}\mu_{e|N}&=\sum_{S_1\in\mathbf{S}}P_s(S_1)\ldots\sum_{S_N\in\mathbf{S}}P_s(S_N)p_{e|s_1\ldots s_N}(E|S_1\ldots S_N)\\&=\sum_{S_1\in\mathbf{S}}P_s(S_1)\ldots\sum_{S_N\in\mathbf{S}}P_s(S_N)\sum_{i=1}^{N}\mu_{e|S_i}=N\sum_{S\in\mathbf{S}}P_s(S)\mu_{e|S}\end{aligned} \quad (2)$$

which is N times the average task consumption $\mu_{e|1}=\sum_{S\in\mathbf{S}}P_s(S)\mu_{e|S}$

The variance of the energy consumed after N tasks can be computed as follows:

$$\sigma^2_{e|N}=\mu_{e^2|N}-\mu^2_{e|N}$$

$$\mu_{e^2|N}=\sum_{S_1\in\mathbf{S}}P_s(S_1)\ldots\sum_{S_N\in\mathbf{S}}P_s(S_N)\left(\sum_{i=1}^{N}\sigma^2_{e|S_i}+\left(\sum_{i=1}^{N}\mu_{e|S_i}\right)^2\right)=$$

$$=N\sum_{S\in\mathbf{S}}P_s(S)\sigma^2_{e|S}+N\sum_{S\in\mathbf{S}}P_s(S)\mu^2_{e|S}+N(N-1)\left(\sum_{S\in\mathbf{S}}P_s(S)\mu_{e|S}\right)^2$$

$$\sigma^2_{e|N}=N\left(\sum_{S\in\mathbf{S}}P_s(S)\sigma^2_{e|S}+\sum_{S\in\mathbf{S}}P_s(S)\mu^2_{e|S}-\left(\sum_{S\in\mathbf{S}}P_s(S)\mu_{e|S}\right)^2\right) \quad (3)$$

Note that only the first component in the sum is a function of the uncertainty about the energy consumption in each task. The sum of the remaining two components is equal to the variance of the mean energy consumed by the individual tasks in the sequence. Note that, in general, the first component can be neglected since it is much smaller than the remaining two.

The pdf of the energy consumed after N tasks is, according to (1), a mixture of M^N Gaussians. In order to simplify further analysis, we now state our fourth assumption:

4. The energy consumed after N tasks can be modelled by a normal r.v. with mean $\mu_{e|N}$ and variance $\sigma^2_{e|N}$.

Numerical evidence shows that this approximation is acceptable for large N.

The mean $\mu_{t|N}$ and variance $\sigma^2_{t|N}$ of the duration of a sequence with N tasks are obtained in a completely analogous way. Here too we will approximate this quantity by a normal r.v.

B. Computation of $P_{n_0}(N_0)$, μ_{n_0}, $\sigma^2_{n_0}$

n_0 is the r.v. representing the number of tasks in the sequence till the battery is depleted. This occurs for $n_0=N_0$ when $\sum_{i=1}^{N_0}E(S_i)\le E_0$ but $\sum_{i=1}^{N_0+1}E(S_i)>E_0$, where E_0 is the remaining battery charge (note that energy only takes non-negative values). Hence,

$$\begin{aligned}P_{n_0}(N_0)&=P\left(\left(\sum_{i=1}^{N_0}E(S_i)\le E_0\right)\wedge\left(\sum_{i=1}^{N_0+1}E(S_i)>E_0\right)\right)\\&=P\left(\sum_{i=1}^{N_0}E(S_i)\le E_0\right)+P\left(\sum_{i=1}^{N_0+1}E(S_i)>E_0\right)\\&-P\left(\left(\sum_{i=1}^{N_0}E(S_i)\le E_0\right)\vee\left(\sum_{i=1}^{N_0+1}E(S_i)>E_0\right)\right)\\&=P_{e|n_0}(e\le E_0|N_0)+\left(1-P_{e|n_0}(e\le E_0|N_0+1)\right)-1\\&=P_{e|n_0}(e\le E_0|N_0)-P_{e|n_0}(e\le E_0|N_0+1)\end{aligned} \quad (4)$$

which , under Assumption 4. above, becomes

$$P_{n_0}(N_0)=G(E_0;\mu_{e|N_0},\sigma^2_{e|N_0})-G(E_0;\mu_{e|N_0+1},\sigma^2_{e|N_0+1}) \quad (5)$$

where $G(x;\mu,\sigma^2)$ is the normal cumulative distribution function (cdf).

The expectation of N_0 is easy to compute:

$$\begin{aligned}\mu_{n_0}&=\sum_{N_0=1}^{\infty}N_0\left(P_{e|n_0}(e\le E_0|N_0)-P_{e|n_0}(e\le E_0|N_0+1)\right)\\&=\sum_{N_0=1}^{\infty}P_{e|n_0}(e\le E_0|N_0)=\sum_{N_0=1}^{\infty}G(E_0;\mu_{e|N_0},\sigma^2_{e|N_0})\end{aligned} \quad (6)$$

A simple approximate relationship between μ_{n_0} and E_0 can be found by neglecting the term $\sigma^2_{e|N_0}$ in (6). This is equivalent to approximating $G(E_0;\mu_{e|N_0},\sigma^2_{e|N_0})$ with a function that is equal to 1 for $E_0\le\mu_{e|N_0}$, and 0 elsewhere. This leads to the following approximation:

$$\mu_{n_0}\approx\frac{E_0}{\mu_{e|1}} \quad (7)$$

which conforms to the intuition that the expected number of remaining tasks is approximately proportional to the remaining battery charge, E_0, where the proportionality coefficient is the inverse of the average task consumption, $\mu_{e|1}$.

The variance of N_0 can be computed using the formula $\sigma_{n_0}^2 = \mu_{n_0^2} - \mu_{n_0}^2$ with

$$\begin{aligned}\mu_{n_0^2} &= \sum_{N_0=1}^{\infty} N_0^2 \left(P_{e|n_0}\left(e \le E_0 \mid N_0\right) - P_{e|n_0}\left(e \le E_0 \mid N_0+1\right)\right) \\ &= \sum_{N_0=1}^{\infty} \left(N_0^2 - \left(N_0 - 1\right)^2 \right) P_{e|n_0}\left(e \le E_0 \mid N_0\right) \\ &= \sum_{N_0=1}^{\infty} \left(2N_0 - 1\right) G\left(E_0; \mu_{e|N_0}, \sigma_{e|N_0}^2\right)\end{aligned} \tag{8}$$

We noticed in our experiments that $\sigma_{n_0}^2$ is approximately proportional to μ_{n_0} for large enough E_0. Further investigation will be needed to provide a formal justification to this observation.

C. Computation of $p_{t_0}(T_0)$, μ_{t_0}, $\sigma_{t_0}^2$

We use the theorem of total probability and average over N_0:

$$p_{t_0}\left(T_0\right) = \sum_{N_0=1}^{\infty} p_{t_0|n_0}\left(T_0 \mid N_0\right) P_{n_0}\left(N_0\right) \tag{9}$$

Now,

$$\begin{aligned}p_{t_0|n_0}\left(T_0 \mid N_0\right) &= \sum_{S_1 \in \mathbf{S}} \cdots \sum_{S_{N_0+1} \in \mathbf{S}} p_{t_0|s_1\ldots s_{N_0+1}}\left(T_0 \mid S_1, \ldots, S_{N_0+1}\right) \\ &\cdot P_{s_1\ldots s_{N_0+1}|n_0}\left(S_1, \ldots, S_{N_0+1} \mid N_0\right) \\ &= \sum_{S_1 \in \mathbf{S}} \cdots \sum_{S_{N_0} \in \mathbf{S}} p_{t_0|s_1\ldots s_{N_0}}\left(T_0 \mid S_1, \ldots, S_{N_0}\right) \\ &\cdot \sum_{S_{N_0+1} \in \mathbf{S}} P_{s_1\ldots s_{N_0+1}|n_0}\left(S_1, \ldots, S_{N_0+1} \mid N_0\right)\end{aligned} \tag{10}$$

where we removed unnecessary dependencies, and noted that the sequence of the first N_0 states completely determines the statistics of the duration T_0 given that n_0 is fixed. Here, $p_{t_0|s_1\ldots s_{N_0}}\left(T_0 \mid S_1, \ldots, S_{N_0}\right) = g\left(T_0; \mu_{t|S_1\ldots S_{N_0}}, \sigma_{t|S_1\ldots S_{N_0}}^2\right)$. Using Bayes' rule, and remembering that the states form a Bernoulli scheme,

$$\begin{aligned}&P_{s_1\ldots s_{N_0+1}|n_0}\left(S_1, \ldots, S_{N_0+1} \mid N_0\right) \\ &= \frac{P_{n_0|s_1\ldots s_{N_0+1}}\left(N_0 \mid S_1, \ldots, S_{N_0+1}\right) P_{s_1\ldots s_{N_0}}\left(S_1, \ldots, S_{N_0}\right) P_{s_{N_0+1}}\left(S_{N_0+1}\right)}{P_{n_0}\left(N_0\right)}\end{aligned} \tag{11}$$

Using results from the previous section, we maintain that

$$\begin{aligned}&\sum_{S_{N_0+1} \in \mathbf{S}} P_{n_0|s_1\ldots s_{N_0+1}}\left(N_0 \mid S_1, \ldots, S_{N_0+1}\right) P_{s_{N_0+1}}\left(S_{N_0+1}\right) \\ &= \sum_{S_{N_0+1} \in \mathbf{S}} \left(G\left(E_0; \mu_{e|S_1\ldots S_{N_0}}, \sigma_{e|S_1\ldots S_{N_0}}^2\right) - G\left(E_0; \mu_{e|S_1\ldots S_{N_0+1}}, \sigma_{e|S_1\ldots S_{N_0+1}}^2\right)\right) \\ &\cdot P_{s_{N_0+1}}\left(S_{N_0+1}\right) \approx \sum_{S_{N_0+1} \in \mathbf{S}} g\left(E_0; \mu_{e|S_1\ldots S_{N_0}}, \sigma_{e|S_1\ldots S_{N_0}}^2\right) \mu_{e|S_{N_0+1}} P_{s_{N_0+1}}\left(S_{N_0+1}\right) \\ &= g\left(E_0; \mu_{e|S_1\ldots S_{N_0}}, \sigma_{e|S_1\ldots S_{N_0}}^2\right) \cdot \mu_{e|1}\end{aligned} \tag{12}$$

and therefore

$$\begin{aligned}p_{t_0|n_0}\left(T_0 \mid N_0\right) &\approx \mu_{e|1} \sum_{S_1 \in \mathbf{S}} \cdots \sum_{S_{N_0} \in \mathbf{S}} P_{s_1\ldots s_{N_0}}\left(S_1, \ldots, S_{N_0}\right) \\ &\cdot g\left(E_0; \mu_{e|S_1\ldots S_{N_0}}, \sigma_{e|S_1\ldots S_{N_0}}^2\right) g\left(T_0; \mu_{t|S_1\ldots S_{N_0}}, \sigma_{t|S_1\ldots S_{N_0}}^2\right)\end{aligned} \tag{13}$$

Thus, $p_{t_0, e_0|n_0}\left(T_0, E_0 \mid N_0\right)$ can be written as $\mu_{e|1}$ times a mixture of M^{N_0} 2-D separable Gaussians. Based on Assumption 4 (extended to the 2-D case), we will approximate this mixture of Gaussians with one Gaussian, having the same mean and covariance as the mixture:

$$p_{t_0|n_0}\left(T_0 \mid N_0\right) \approx \mu_{e|1} g\left(E_0; \mu_{e|N_0}, \sigma_{e|N_0}^2\right) g\left(T_0; \mu_{t|N_0}, \sigma_{t|N_0}^2\right) \tag{14}$$

Finally, we obtain:

$$\begin{aligned}p_{t_0|n_0}\left(T_0 \mid N_0\right) &\approx \frac{\mu_{e|1} g\left(E_0; \mu_{e|N_0}, \sigma_{e|N_0}^2\right) g\left(T_0; \mu_{t|N_0}, \sigma_{t|N_0}^2\right)}{\mu_{e|1} g\left(E_0; \mu_{e|N_0}, \sigma_{e|N_0}^2\right)} \\ &= g\left(T_0; \mu_{t|N_0}, \sigma_{t|N_0}^2\right)\end{aligned} \tag{15}$$

and

$$p_{t_0}\left(T_0\right) \approx \sum_{N_0=1}^{\infty} g\left(T_0; \mu_{t|N_0}, \sigma_{t|N_0}^2\right) P_{n_0}\left(N_0\right) \tag{16}$$

The mean of this pdf is:

$$\mu_{t_0} = \sum_{N_0=1}^{\infty} \mu_{t|N_0} P_{n_0}\left(N_0\right) = \sum_{N_0=1}^{\infty} N \mu_{t|1} P_{n_0}\left(N_0\right) = \mu_{n_0} \mu_{t|1} \tag{17}$$

which is the average number of tasks before depletion, multiplied by the average duration of a task.

It is not difficult to prove that its variance is equal to

$$\begin{aligned}\sigma_{t_0}^2 &= \mu_{n_0} \left(\sum_{S \in \mathbf{S}} P_s(S) \sigma_{t|S}^2 + \sum_{S \in \mathbf{S}} P_s(S) \mu_{t|S}^2 - \left(\sum_{S \in \mathbf{S}} P_s(S) \mu_{t|S} \right)^2 \right) \\ &+ \sigma_{n_0}^2 \left(\sum_{S \in \mathbf{S}} P_s(S) \mu_{t|S} \right)^2 = \sigma_{t|\mu_{n_0}}^2 + \sigma_{n_0}^2 \mu_{t|1}^2\end{aligned} \tag{18}$$

The first term in this sum corresponds to the variance of the time elapsed after μ_{n_0} tasks (compare with (3)). The second term accounts for the variance of n_0. If $\sigma_{n_0}^2$ can be considered approximately proportional to μ_{n_0} (see previous section), then $\sigma_{t_0}^2$ can be considered approximately proportional to μ_{n_0} as well.

D. Simulations

We consider here two very simple duty cycles for the Meerkats testbed. The first duty cycle simply repeats the sequence 'Acquire/compress image', 'Transmit image', 'Wait in idle for 5 seconds'. These three elementary tasks are thus summarized into a single task ($\mathbf{S} = \{A\}$) with the following measured characteristics:

$$\mu_{e|A} = 2.94 \text{ C}, \ \sigma_{e|A}^2 = 1 \cdot 10^{-3} \text{ C}^2$$

$$\mu_{t|A} = 6.30 \text{ s}, \ \sigma_{t|A}^2 = 4 \cdot 10^{-3} \text{ s}^2$$

where 'C' represents 'Coulombs' and 's' are 'seconds'. The sample mean and variance were computed from 100 measurements, after removing the 5% largest and 5% smallest ones in order to reduce the influence of outliers.

The second duty cycle has a branching condition. When an image is taken, if a target of interest has been detected, the image is transmitted and then the system remains idle for a short time (so that another snapshot can hopefully be taken of the target before it leaves the camera's field of view). Otherwise, the system does not transmit the image and remains idle for a longer time. This duty cycle can be formalized as a random draw from two possible tasks ($\mathbf{S}=\{B,C\}$), where B is the same sequence as A but with an idle time of only 1 second, and C is the sequence 'Acquire/compress image', 'Wait in idle for 5 seconds'. Note that, for simplicity's sake, we neglect the cost of the image analysis algorithm to detect the target. We assume that $P(B)=0.2$ and $P(C)=0.8$. The measured characteristics of the tasks B and C are:

$$\mu_{e|B}=1.1\,5\text{C},\ \sigma^2_{e|B}=0.3\cdot10^{-3}\ \text{C}^2$$

$$\mu_{t|B}=2.30\ \text{s},\ \sigma^2_{t|B}=1\cdot10^{-3}\ \text{s}^2$$

$$\mu_{e|C}=2.88\ \text{C},\ \sigma^2_{e|C}=0.5\cdot10^{-3}\ \text{C}^2$$

$$\mu_{t|C}=6.19\ \text{s},\ \sigma^2_{t|C}=1.7\cdot10^{-3}\ \text{s}^2$$

Fig. 2 shows the values of μ_{n_0} and μ_{t_0} for the cases $\mathbf{S}=\{A\}$ and $\mathbf{S}=\{B,C\}$, along with $\pm3\sigma$ error bars. Note that the variance of n_0 and of t_0 are very small for the first case, but sensibly larger in the second case. Since $\sigma^2_{n_0}$ turns out to be approximately proportional to μ_{n_0}, according to the observation at the end of the previous section, $\sigma^2_{t_0}$ is approximately proportional to μ_{t_0}. For $\mathbf{S}=\{A\}$, $\sigma^2_{t_0}\approx1.3\cdot10^{-3}\mu_{t_0}$, while for $\mathbf{S}=\{B,C\}$, $\sigma^2_{t_0}\approx0.85\mu_{t_0}$.

IV. CONCLUSIONS

This paper has presented a model for determining the first and second order statistics of the lifetime of a battery-operated node. This model assumes knowledge of the statistics of duration and consumption of the tasks that can be performed by the system, as well as of the statistics of the task occurrences and of the remaining charge in the battery. The latter may be monitored by an onboard sensor, as in the case of the Meerkats testbed. Knowledge of the remaining lifetime is very important in order to devise an online control strategy assuring that specific requirements are met.

Although rather versatile, this model is based on a number of rather simplistic assumptions, some of which are briefly discussed below.

- The probability distribution of task occurrence is typically not known *a priori* but rather should be inferred by a suitable online mechanism.
- The assumption of normal distribution of duration and consumption for a given task is often not tenable. In many cases, we observed "heavy tail" distributions with a non-negligible amount of outliers.
- The lifetime of a node is not determined solely by the about of charge that it draws. As noted in [3], more complex models of the battery (that take into account, for example, the discharge rate) need to be used for a more realistic modeling.

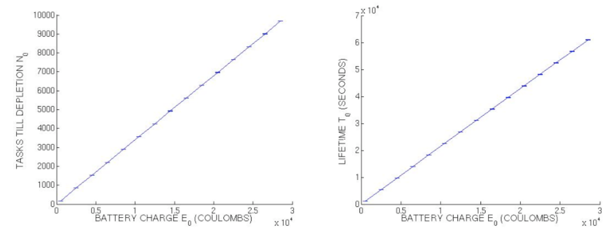

$\mathbf{S}=\{A\}$

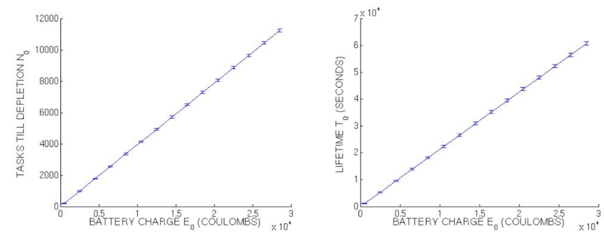

$\mathbf{S}=\{B, C\}$

V. REFERENCES

[1] Workshop on Distributed Smart Cameras (DSC'06), Boulder, CO, October 2006.

[2] L. Benini, A. Bogliolo, G. Paleologo, G. De Micheli, "Policy Optimization for Dynamic Power Management", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 18(6):813-33, June 1999.

[3] A. Bogliolo, L. Benini, E. Lattanzi, G. De Micheli, "Specification and Analysis of Power-Managed Systems", Proceedings of the IEEE, 92(8):1308-46, August 2004.

[4] C. B. Margi, X. Lu, G. Zhang, G. Stanek, R. Manduchi, K. Obraczka, "Meerkats: A Power-Aware, Self-Managing Wireless Camera Network for Wide Area Monitoring", DSC'06, Boulder, CO, October 2006.

[5] C. Margi, R. Manduchi, K. Obraczka, "Energy Consumption Tradeoffs in Visual Sensor Networks", SBRC 2006, Sao Paulo, Brazil, May 2006.

[6] C. Margi, V. Petkov, K. Obraczka, R. Manduchi, "Characterizing Energy Consumption in a Visual Sensor Network Testbed", TridentCom'06, Barcelona, Spain, March 2006.

ACKNOWLEDGMENTS

This work was supported by NASA under contract NNA04CK89A.