Harnessing Battery Recovery Effect in Wireless Sensor Networks: Experiments and Analysis

Chi-Kin Chau, *Member*, *IEEE*, Fei Qin, Samir Sayed, Muhammad Husni Wahab,
Yang Yang, *Senior Member*, *IEEE*

***Abstract*—Many applications of wireless sensor networks rely on batteries. But most batteries are not simple energy reservoirs, and can exhibit *battery recovery effect*. That is, the deliverable energy in a battery can be self-replenished, if left idling for sufficient time. As a viable approach for energy optimisation, we made several contributions towards harnessing battery recovery effect in sensor networks. 1) We empirically examine the gain of battery runtime of sensor devices due to battery recovery effect, and affirm its significant benefit in sensor networks. We also observe a saturation threshold, beyond which more idle time will contribute only little to battery recovery. 2) Based on our experiments, we propose a Markov chain model to capture battery recovery considering saturation threshold and random sensing activities, by which we can study the effectiveness of duty cycling and buffering. 3) We devise a simple distributed duty cycle scheme to take advantage of battery recovery using pseudo-random sequences, and analyse its trade-off between the induced latency of data delivery and duty cycle rates.**

***Index Terms*—Wireless Sensor Networks, Energy Optimisation, Battery Recovery Effect, Duty Cycle.**

I. Introduction

Wireless sensor networks are created by networks of small devices, integrated with tiny embedded processors, radio transceivers, and MEMS micro-sensors. Many applications of sensor networks require batteries as energy source for the sensors. However, small form factors of devices often prohibit the uses of larger capacity batteries. Also, ad-hoc deployment of sensor networks and the inconvenience of battery recollection usually constrain frequent replacements of on-board batteries. Hence, the design of energy-efficient protocols has become a crucial topic in sensor networking.

There are many studies on energy optimization that regard batteries as ideal energy reservoirs, where energy is drained at constant discharging voltage, and can be halted and resumed anytime at the same voltage. However, most commercial batteries are governed by complex intrinsic chemical reactions to produce energy. Such chemical reactions are known by chemical engineers to be dependent on a variety of environmental factors and operational parameters (e.g., discharge duration, current and history) [1], [2].

Particularly, there is a subtle phenomenon called *battery recovery effect*, which refers to the process that the active chemical substances in a battery can be self-replenished if left idling for a sufficient period of time. In this paper, we are motivated to exploit battery recovery effect to optimise the energy efficiency of wireless sensor networks.

There are several immediate questions. 1) Is battery recovery effect sufficiently significant to extend battery runtime? 2) If so, is there a simple but effective approach to take advantage of battery recovery effect in sensor networks? 3) Such an approach may inevitably affect the performance of sensor networks. Then how will the induced performance (e.g., latency of data delivery) be affected?

To address these questions, in this paper we first empirically examine the gain of battery runtime due to battery recovery effect, through test-bed experiments on commercial sensors. Since radio transceivers consume significant portion of energy (even in listening mode), as compared to processing and sensing activities, we especially measure the battery runtime in the presence of duty cycled radio operations under both deterministic and randomised schedules. We found that there is up to a gain of 30%-45% to the normalised battery runtime[1] between duty cycled and continuous radio operations.

We then empirically study the characteristics of battery recovery effect, with respect to different duty cycle schedules. We observe that there exists a *saturation threshold*, beyond which more consecutive idle time will contribute only little to battery recovery. The ramification to sensor networking is that if we carefully adjust the sleep time period of battery before reaching the saturation threshold, we can maximise battery recovery without exacerbating the latency of data delivery.

Next, we model the process of battery recovery considering saturation threshold and random sensing activities. We study one-hop sensor networks by a Markov chain model, and investigate the effectiveness of energy optimisation by duty cycling and buffering. We obtain several analytical insights for the battery runtime, corroborated by simulation.

We then extend our study to multi-hop sensor networks, where each sensor can act as a relay to forward the sensing data for other sensors to the sink. In such setting, there

Manuscript received 1 May 2009; revised 16 February 2010. C.-K. Chau was supported by a Croucher Foundation fellowship. Y. Yang was partially supported by CONET under FP7 with contract number FP7-2007-2-224053 and by the National Natural Science Foundation of China (NSFC) under the grant 60902041.

C.-K. Chau is with Computer Laboratory, University of Cambridge, UK (e-mail: chi-kin.chau@cl.cam.ac.uk).

F. Qin and S. Sayed are with Electronic & Electrical Engineering Department, University College London, UK (e-mail: {f.qin,s.sayed}@ee.ucl.ac.uk).

M. H. Wahab is with the Computer Science Department, University College London (e-mail: h.wahab@cs.ucl.ac.uk).

Y. Yang is with Shanghai Research Center for Wireless Communications (WiCO), SIMIT, Chinese Academy of Sciences, China (e-mail: yang.yang@shrcwc.org).

Digital Object Identifier 10.1109/JSAC.2010.100926.

[1]Normalised battery runtime is the measured battery runtime multiplied by the duty cycle rate.

requires a coordination scheme among the duty cycled sensors, such that each sensor can discover the appropriate timeslot to transmit data without wasting energy to probe the availability of their relays. We adapt a distributed randomised coordination scheme from [3], and extend it to take advantage of battery recovery. In our scheme, each sensor infers the random duty cycle schedule of the relay based on pseudo-random sequence. The random duty cycle schedule is set to let a sensor to sleep for a period of time within the saturation threshold, when it has been active for a certain a period of time. Through simulation, we show that our battery-aware duty cycle scheme can significantly extend the battery runtime.

Finally, since there is an inevitable trade-off between the latency of data delivery and the energy efficiency in duty cycled sensor networks, to shed light on the performance under constrained energy budgets we analyse and bound the latency of our battery-aware duty cycle scheme in multi-hop sensor networks (and also specifically in lattice networks). The analysis is based on the results in our previous work [4].

In summary, our contributions are threefold:

1) In Sec. III, we provide the experimental evidence on the significant gain of battery runtime by battery recovery effect, and report that the gain is characterised by a saturation threshold.
2) In Sec. IV, we model battery recovery and study energy optimisation by duty cycling and buffering, in the presence of saturation threshold and random sensing activities.
3) We present a distributed duty cycle scheme that takes advantage of battery recovery using pseudo-random sequence in Sec. V, and study the induced latency of data delivery in Sec. VI.

We also present the background and related work in Sec. II, and discuss some issues of our work in Sec. VII. Because of space constraint, the proofs of theorems are deferred to [5].

II. Background and Related Work

A. Battery Models

There are many studies on the performance of batteries in chemical engineering [1]. In networking, [6] carried out an empirical study to measure the performance of battery-powered sensors, but did not examine the saturation threshold. Although battery consumption has been modelled extensively in networking, many extant models over-simplified the realistic battery characteristics (e.g., allowing unlimited battery recovery). There are two main models that consider realistic battery characteristics.

First, the kinetic battery models [7], [8] attempt to model the detailed chemical reactions and diffusion process between the electrode and electrolyte in a battery through a set of partial differential equations. These models aim to fully capture the non-linear dynamics in a battery. However, these models are less tractable, and different form factors of batteries can significantly affect the accuracy of the models.

Second, there are stochastic battery models [9]–[11] that capture the battery dynamics using randomised Markovian models. But most of them did not consider the effect of idle time period. While [12] considers idle time period in embedded systems, it did not address the performance in sensor networks. We remark that while all these stochastic battery models attempt to imitate the kinetic battery model with less complexity, the uses of randomised battery recovery is different from the deterministic kinetic battery models. Moreover, these models are also less tractable, with few analytical insights provided.

In this paper, we present a more analysable Markov chain model that simplifies the stochastic battery models [9], [10], [12]. Particularly, our model uses deterministic battery recovery, yet is able to capture realistic battery characteristics, such as limited battery recovery and the effect of idle time period. More importantly, several analytical insights for realistic battery behaviour can be derived from our model.

B. Energy Management and Optimisation

There are many studies on energy optimization in sensor networks, including topology management and network layer optimisation. Particularly relevant to our work are those based on MAC layer, which aim to reduce redundant radio operations in MAC protocols: 1) idle listening (keeping radio on even when no reception), 2) overhearing (reception of a message not intended for the receiver), and 3) protocol overhead (redundant headers or signalling messages). Examples include S-MAC, SEEDEX, O-MAC, RI-MAC [13]–[16]. Listening and reception can consume significant energy in wireless sensors. This paper reduces idle listening and overhearing by duty cycle based on pseudo-random sequence [3] and battery recovery effect. There are other MAC protocols that consider battery characteristics. BAMAC and Bel-MAC [17], [18] relied on exchanging dynamic battery state information to optimise the use of batteries among sensors. But such information cannot be obtained without relying on online measurement of battery. However, our idea of exploiting saturation threshold does not rely on online measurement. Also, [19] considered battery recovery in scheduling and routing, but did not incorporate the information of saturation threshold, unlike our distributed coordination scheme among the duty cycled sensors.

III. Experimental Results

In this section, we present the experimental results from our sensor network test-bed to show the significance of battery recovery effect. These experiments were carried out on two types of commercial sensors from Crossbow: TelosB and Imote2. Both are popular models for wireless sensor networking. TelosB is consisted of MSP430 as MCU and CC2420 as radio transceiver. Imote2 is consisted of PXA271 as CPU and CC2420 as radio transceiver. TelosB allows more energy-saving settings with low energy consumption, whereas Imote2 is equipped with more computation ability with high energy consumption. Here, we mainly study TelosB.

In the experiments, we used an analogue-digital conversion (ADC) interface card and software LabVIEW to measure and record the discharge profiles of communicating sensors. Each sensor was powered by standard AAA NiMH 600 mAh batteries (TelosB has two and Imote2 has three). When the supply voltage of the battery is lower than a certain threshold (called stop voltage), the device can no longer operate, which

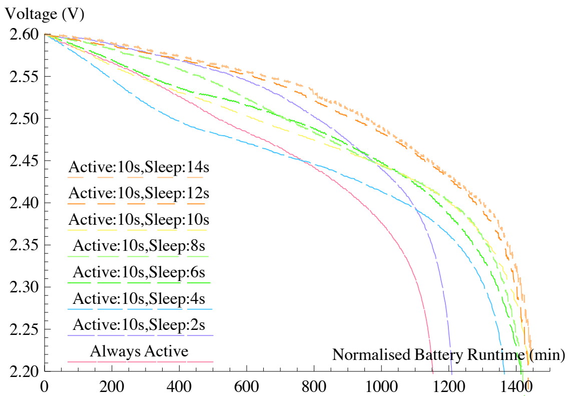

Fig. 1. The discharge profiles of TelosB in deterministic duty cycle schedules of radio transceiver w.r.t. different sleep time periods.

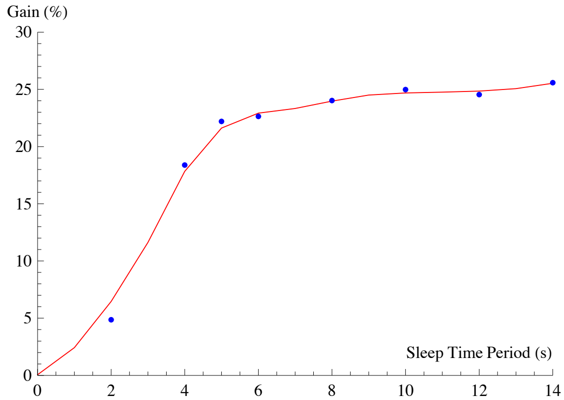

Fig. 2. The gain of normalised battery runtime of Fig. 1, as compared to the continuous (always active) radio operation.

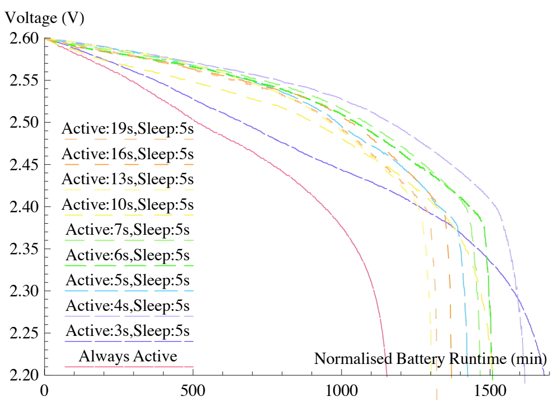

Fig. 3. The discharge profiles of TelosB in deterministic duty cycle schedules of radio transceiver w.r.t. different active time periods.

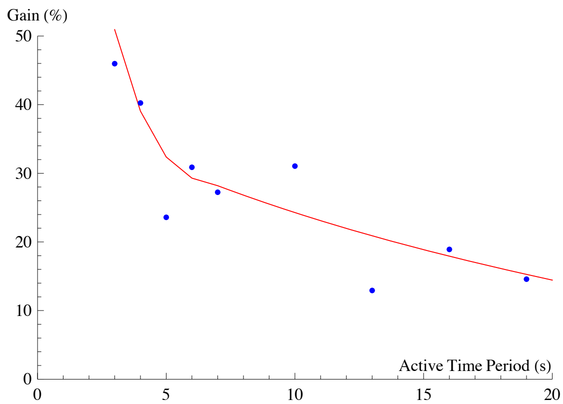

Fig. 4. The gain of normalised battery runtime of Fig. 3, as compared to the continuous (always active) radio operation.

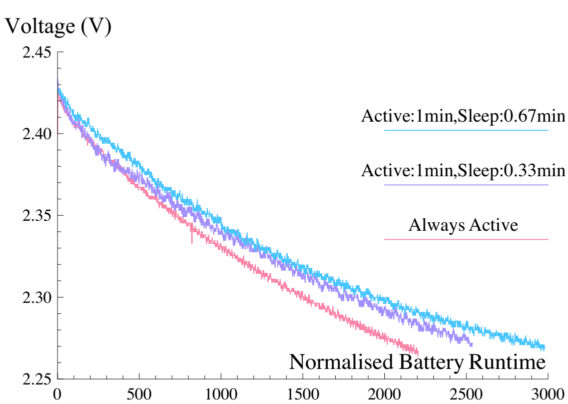

Fig. 5. The discharge profiles of Imote2 in deterministic duty cycle schedules of radio transceiver w.r.t. different sleep time periods.

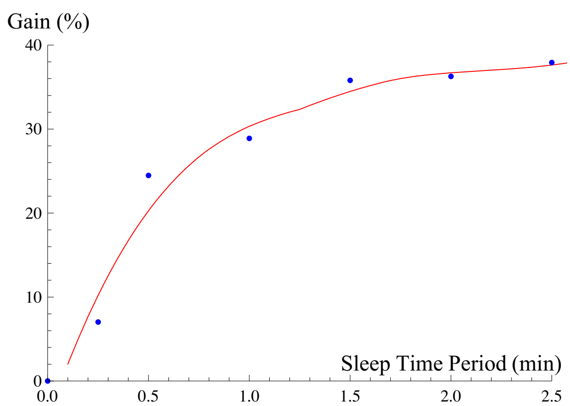

Fig. 6. The gain of normalised battery runtime of Fig. 5, as compared to the continuous (always active) radio operation.

is considered as completely discharged. We set different duty cycle rate on the sensors by putting the radio transceiver in active and sleep modes periodically, and measure the induced battery runtime. The duty cycle rate is defined as the fraction of active time periods, and the *normalised battery runtime* is the measured battery runtime multiplied by the duty cycle rate.

First, we tested on deterministic duty cycle schedules. Figs. 1 and 3 show the discharge profiles of the transmitter in normalised battery runtime for TelosB[2] and Fig. 5 for Imote2. Figs. 1, 3, 5 show the respective gain of normalised battery runtime, compared to the continuous radio operation.

The key observations from the experiments are:

1) There are clear signs of battery recovery effect. With the same active time period, longer sleep time period

[2]Note that because TelosB has a much longer battery runtime, we used a lower sampling rate to facilitate data logging, which gives an apparently smooth profile curve. However, the discharge profiles should be as ragged as the ones of Imote2 in Fig. 5.

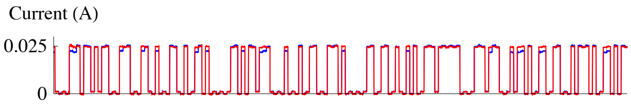

Fig. 7. We use the same pseudo-random sequence to generate the random duty cycle schedules for both transmitter and receiver. This is a snapshot of the current profile of a pair of transmitter-receiver. The higher current indicates the periods of higher energy consumption in active slots.

induces longer normalised battery runtime, and hence larger deliverable energy of a battery.

2) The effect of sleep time period is non-linear. It appears that sleep time period more than a certain threshold will contribute much less to battery recovery, which we call a *saturation threshold*.
3) The effect of active time period is also non-linear. Very small active time period appears to have large gain of normalised battery runtime (up to 45% in TelosB for active/sleep time as 3 sec/5 sec).
4) Even the sensor is in sleep mode on radio transceiver, there is still energy consumption due to the timer and other processing activities. TelosB consumes 6.1μA in sleep mode, whereas Imote2 consumes 0.38mA. We observe that battery recovery can take place under low battery consumption, and the impact of background consumption is not substantial to battery recovery.

Next, we study the gain of randomised duty cycle schedules. We used the same pseudo-random sequence to generate the random duty cycle schedules for both transmitter and receiver (see Fig. 7 for a snapshot). Fig. 8 shows two discharge profiles of transmitter and receiver both with the same duty cycle rate as 0.5 (one sets each random slot as 0.5 sec, another as 5 sec). We observe that the gains of normalised battery runtime between the 5-sec and 0.5-sec cases are 27% and 25% for transmitter and receiver respectively. Note that the normalised battery runtime of 5-sec random slots is comparable to the deterministic duty cycle schedule (active and sleep time period as 5 sec) in Fig. 4, while the one of 0.5-sec random slots is comparable to the always active schedule. The implications are 1) the gain of battery runtime is determined by how the active and sleep time periods last, rather by the duty cycle rate alone, and 2) the randomised duty cycle schedules can have a comparable gain as the deterministic duty cycle schedules.

Note that the actual battery runtime is much longer than the normalised battery runtime (which is multiplied by the duty cycle rate). Hence, duty cycle schedules not only prolong the network lifetime by spending time in sleep mode, but also increase the deliverable energy by allowing battery recovery.

Although our measurement of the gains of battery recovery may differ in other environmental settings (e.g., different temperature), the insights revealed by our experiments will still be useful to the modelling and optimisation of battery recovery in sensor networks. In general, under a pulsed (duty cycle) discharge profile, a battery is able to recover charge during idle time periods, which effectively increases the deliverable energy of the battery. But the effectiveness of battery recovery is critically determined by the active and sleep time periods. The presence of saturation threshold appears universal in different duty cycle schedules. We envision that one can

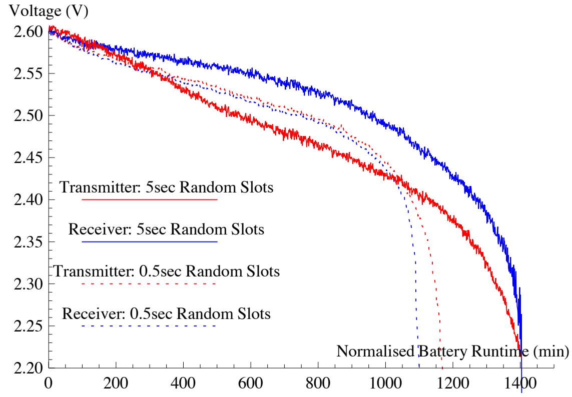

Fig. 8. Two discharge profiles a pair of transmitter-receiver of TelosB in pseudo-random duty cycle schedules of radio transceiver (duty cycle rate as 0.5): one sets each random slot as 0.5 sec, while another sets each random slot as 5 sec.

evaluate the saturation threshold (in some intervals) for certain environments in a-priori manner through experiments. Next, we design network protocols that exploit battery recovery by taking the estimated saturation threshold as a parameter.

IV. Model of Battery Recovery

Our experiments corroborated the presence of battery recovery effect and saturation threshold. It is useful to employ a simple model to capture these essential characteristics qualitatively, which enables further analysis and larger scale simulation on various battery consumption patterns. Hence, we are motivated to present a simplified Markov chain model, where time is discretised as slots, to capture the battery consumption by transmitting random sensing data. In this model, we assume that the sensing data arriving to a sensor in the previous slots, will be transmitted in the current slot.

The state of a battery is characterised by a tuple $\langle n, c, \mathtt{t}\rangle$, where $n, c, \mathtt{t}$ are non-negative integers. c is the *theoretical capacity* determined by the amount of chemicals in the electrode and electrolyte, n is the *nominal capacity* determined by the amount of available active chemicals for chemical reactions in the battery, and $\mathtt{t}$ is the number of idle slots since last discharging. The use of $\langle n, c\rangle$ follows a popular battery model described in [9], [10]. But we introduce $\mathtt{t}$ in this paper, as related to the saturation threshold.

In the discharging process, both n and c are decreasing. The amount of available active chemicals constrains the energy of a battery can deliver, despite the presence of unused chemicals in the battery. Thus, we require $n \le c$. But when the battery stops discharging, there is a recovery process, as a diffusion process between electrode and electrolyte to replenish available active chemicals, which effectively increases n (which however cannot increase theoretical capacity c). There is a saturation threshold $\mathtt{t}_{\rm sat}$ for $\mathtt{t}$, such that more consecutive idle slots $\mathtt{t} > \mathtt{t}_{\rm sat}$ will not contribute more recovery. Here, we normalise the unit of n and c, such that at each idle slot, n can be recovered by only one unit.

In this section, we study this battery model in a one-hop sensor network, where the multi-hop study will deferred to the next section. Each sensor is a leaf node, whose task to transmit its detected sensing data to the sink. We assume the sink is always active, but the sensor may follow duty cycle schedule on its radio operations. And the sensor is a simple transmitter, requiring no feedback from the sink. For generality, we assume that sensing data arrival process follows a Poisson distribution. The consumption of a battery is proportional to the amount of sensing data for transmission to the sink in a slot.

Next, we study two scenarios: 1) the always-active case where the data is transmitted in the following slot, and 2) the duty cycling case where the data is buffered for a certain period before transmission.

A. Always-active Case without Duty Cycling

We model the scenario without duty cycle as a Markov chain $\mathcal{M}_{\rm bat}$ with the set of states as:

$$\Big\{\langle n, c, \mathtt{t}\rangle : n, c, \mathtt{t} \text{ are non-negative integers, and } n \leq c\Big\}$$

$\langle 0, c, \mathtt{t}\rangle$ is an absorption state that will not transit to other states, corresponding to a complete discharged battery. Due to Poisson arrival of data, the transition probabilities from states of $n \geq 1$ are obtained as:

a) *Discharging*:

$$\langle n, c, \mathtt{t}\rangle \xrightarrow{\mathtt{a}} \langle n\text{-}k, c\text{-}k, 0\rangle$$

for $k \geq 1$ and $n\text{-}k \geq 1$, and the transition probability is $p^{\lambda}_{\rm po}(k) \triangleq \frac{\lambda^k e^{-\lambda}}{k!}$, where λ is the average Poisson arrival rate of data in a slot.

b) *Completely Discharged*:

$$\langle n, c, \mathtt{t}\rangle \xrightarrow{\mathtt{b}} \langle 0, c\text{-}n, 0\rangle$$

and the transition probability is $\sum\limits_{k=n}^{\infty} p^{\lambda}_{\rm po}(k) = \sum\limits_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!}$. That is, when consumption exceeds nominal capacity n, the battery will be completely discharged.

c) *Idling with Recovering*:

$$\langle n, c, \mathtt{t}\rangle \xrightarrow{\mathtt{c}} \langle n\text{+}1, c, \mathtt{t}\text{+}1\rangle$$

for $c \geq n\text{+}1$ and $\mathtt{t} < \mathtt{t}_{\rm sat}$, and the transition probability is $p^{\lambda}_{\rm po}(0) = e^{-\lambda}$. That is, there is recovery when there is no battery consumption, and the consecutive idle duration is lesser than the saturation threshold $\mathtt{t}_{\rm sat}$.

d) *Idling without Recovering*:

$$\langle n, c, \mathtt{t}\rangle \xrightarrow{\mathtt{d}} \langle n, c, \mathtt{t}\text{+}1\rangle$$

for $c < n\text{+}1$ or $\mathtt{t} \geq \mathtt{t}_{\rm sat}$, and the transition probability is $p^{\lambda}_{\rm po}(0) = e^{-\lambda}$. That is, there is no recovery when the nominal capacity reaches the theoretical capacity c, or the consecutive idle duration reaches the saturation threshold $\mathtt{t}_{\rm sat}$.

e) Otherwise, the transition probability for all other pairs of states is 0.

In Figs. 9-10, we illustrate the state transitions in a graphical manner, with $\mathtt{t}_{\rm sat} = 1$.

We remark that there are several simplifications in this model. First, we assume the saturation threshold $\mathtt{t}_{\rm sat}$ is independent of n and c. Second, the recovery process is the same for any state below the saturation threshold. While these assumptions cannot reflect the complete non-linear dynamics in real batteries, we intend to capture the essence of generic battery behaviour. With the minimum number of parameters (only $\mathtt{t}_{\rm sat}$), our model does not rely on experiments to determine further parameters as in a more complete model that otherwise critically depend on the types of batteries and other environmental factors[3]. Moreover, our model enables us to obtain simple analytical insights for harnessing battery recovery.

Let $\mathtt{W}^{\mathtt{t}}_m(n)$ be the *expected battery runtime* starting at state $\langle n, n\text{+}m, \mathtt{t}\rangle$ in Markov chain $\mathcal{M}_{\rm bat}$ (i.e., the expected number of slots a battery can last until reaching an absorption state of completely discharged).

By the definition of state transitions in $\mathcal{M}_{\rm bat}$ and the linearity of expected values, it follows that

$$\mathtt{W}^{0}_{0}(n) = 1 + \sum_{k=0}^{n\text{-}1} p^{\lambda}_{\rm po}(k) \mathtt{W}^{0}_{0}(n\text{-}k), \quad \mathtt{W}^{\mathtt{t}}_{0}(n) = \mathtt{W}^{0}_{0}(n) \qquad (1)$$

That is, if the consumption is of k units and $k \leq n$, then the expected battery runtime will be $\mathtt{W}^{0}_{0}(n\text{-}k) + 1$. Otherwise, it will be completely discharged and the expected battery runtime is 1. For $m \geq 1$, Eqn. (2) follows. Eqns. (1)-(2) give a recursive relationship of $\mathtt{W}^{\mathtt{t}}_m(n)$. But in general, $\mathtt{W}^{\mathtt{t}}_m(n)$ appears to have no simple closed-form solution. However, several general analytical results can be derived as follows.

Theorem 1: The expected battery runtime is bounded by:

$$\mathtt{W}^{\mathtt{t}}_m(n) \leq \mathtt{W}^{0}_m(n) \leq \frac{n\text{+}m\text{-}1}{\lambda} + \frac{1}{1 - e^{-\lambda}}$$

This upper bound is intuitive, because the battery runtime should not exceed $\Theta(n\text{+}m) = \Theta(c)$. Indeed, the upper bound is tight as shown as follows.

Theorem 2: When λ is small or n is large, for some constants α,

$$\mathtt{W}^{0}_m(n) \approx \frac{n\text{+}m}{\lambda} + \alpha$$

We verified Theorems 1-2 for some specific n and m in Fig. 13, where $m = 40$. We observe that the upper bound is indeed tight when λ is small or n is large. Theorem 1 can be regarded as a capacity of battery runtime. We are interested in the gap between the actual battery runtime $\mathtt{W}^{0}_m(n)$ and the capacity of battery runtime. Particularly, we can characterise $\mathtt{W}^{0}_m(1)$ when m is large.

Theorem 3: Considering $\mathtt{t}_{\rm sat} = 1$,

$$\mathtt{W}^{0}_m(1) \leq \lim_{m\to\infty} \mathtt{W}^{0}_m(1) = \frac{e^{2\lambda}}{-e^{\lambda} + e^{2\lambda} - \lambda}$$

Theorem 3 shows that $\mathtt{W}^{0}_m(1) = \mathrm{O}(1)$ for a fixed λ. However, Theorem 1 shows that the upper bound can be $\mathrm{O}(m\text{+}n)$. Therefore, there is a gap with an order as $\mathrm{O}(m\text{+}n)$. This suggests that there is a large potential for exploiting the undelivered energy of a battery. Hence, we are motivated to make use of duty cycling and buffering to improve energy efficiency by harnessing battery recovery.

[3]Other models (e.g., [9], [10], [12]) attempt to incorporate more parameters, such as complex randomised state transitions to capture non-linear battery recovery behaviour. However, these models appear inconvenient for analysis, and little analytical insights have been obtained.

$$\mathtt{W}^{\mathtt{t}}_{m}(n) = \begin{cases} 1 + p^{\lambda}_{\rm po}(0)\mathtt{W}^{\mathtt{t+1}}_{m\text{-}1}(n{+}1) + \sum\limits_{k=1}^{n\text{-}1} p^{\lambda}_{\rm po}(k)\mathtt{W}^{0}_{m}(n\text{-}k) & \text{if } \mathtt{t} < \mathtt{t}_{\rm sat} \\ 1 + p^{\lambda}_{\rm po}(0)\mathtt{W}^{\mathtt{t}_{\rm sat}}_{m}(n) + \sum\limits_{k=1}^{n\text{-}1} p^{\lambda}_{\rm po}(k)\mathtt{W}^{0}_{m}(n\text{-}k) & \text{if } \mathtt{t} = \mathtt{t}_{\rm sat} \end{cases} \quad (2)$$

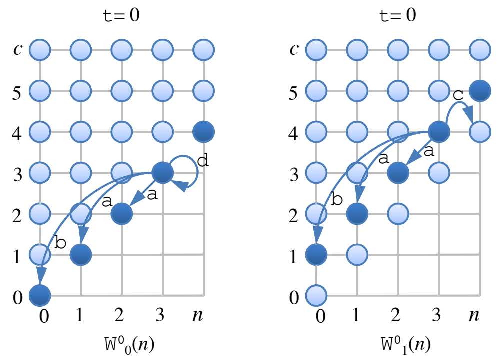

Fig. 9. State transitions for Markov chain $\mathcal{M}_{\rm bat}$ at $\mathtt{t} = 0$.

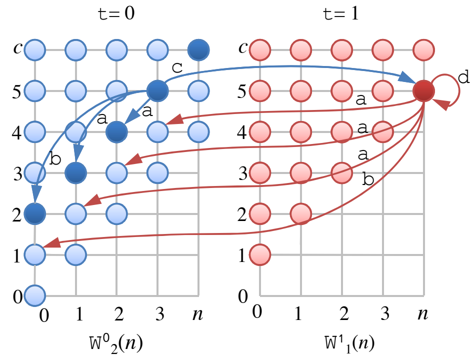

Fig. 10. State transitions for $\mathtt{t} = 1$, when $\mathtt{t}_{\rm sat} = 1$

B. Duty Cycling Case with Buffering

In this section, we analyse the effectiveness of duty cycling and buffering on the battery model. We consider a simple strategy that the sensor will sleep for $\mathtt{b}_{\rm buf}$ slots, every time after the battery is consumed for transmissions. During the sleep time periods, the sensing unit is still on, and the data is collected for buffering within the $\mathtt{b}_{\rm buf}$ slots. To take the maximal advantage of battery recovery, we set $\mathtt{b}_{\rm buf} \le \mathtt{t}_{\rm sat}$.

To capture the above buffered operations, we define Markov chain $\mathcal{M}_{\rm buf}$, where each state is a tuple $\langle n, c, \mathtt{b}\rangle$, and $\mathtt{b} \ge 1$ means that the battery has been in buffered state for $\mathtt{b}$ slots. If $\mathtt{b} + 1 > \mathtt{b}_{\rm buf}$, then the buffer will not hold any data and proceed to immediate transmissions. Thus, the set of states is:

$$\Big\{\langle n, c, \mathtt{b}\rangle : n, c, \mathtt{b} \text{ are non-negative integers, and } n \le c\Big\}$$

We next analyse the expected battery runtime of $\mathcal{M}_{\rm buf}$ in a similar manner as $\mathcal{M}_{\rm bat}$. However, the straightforward definitions of state transitions of $\mathcal{M}_{\rm buf}$ appear rather intractable (even numerically). For the ease of analysis, in the following we define the state transitions of $\mathcal{M}_{\rm buf}$ in a simplified manner, such that the battery consumption is recorded during the buffering slots before actual transmissions, and hence, the battery runtime of $\mathcal{M}_{\rm buf}$ is smaller than the actual battery runtime. This still serves as a lower bound for the actual battery runtime, and does not affect the insights that we obtain.

Now the transition probabilities of $\mathcal{M}_{\rm buf}$ are defined as:

`a')` *Discharging*:

$$\langle n, c, 0\rangle \xrightarrow{\mathtt{a}'} \langle n\text{-}k, c\text{-}k, 1\rangle$$

for $k \ge 1$ and $n\text{-}k \ge 1$, and the transition probability is $p^{\lambda}_{\rm po}(k) = \frac{\lambda^k e^{-\lambda}}{k!}$.

`b')` *Completely Discharged*:

$$\langle n, c, \mathtt{b}\rangle \xrightarrow{\mathtt{b}'} \langle 0, c\text{-}n, 0\rangle$$

and the transition probability is $\sum\limits_{k=n}^{\infty} p^{\lambda}_{\rm po}(k) = \sum\limits_{k=n}^{\infty} \frac{\lambda^k e^{-\lambda}}{k!}$.

`c')` *Idling*:

$$\langle n, c, 0\rangle \xrightarrow{\mathtt{c}'} \begin{cases} \langle n{+}1, c, 0\rangle & \text{if } c \ge n{+}1 \text{ (recovery)} \\ \langle n, c, 0\rangle & \text{if } c < n{+}1 \text{ (no recovery)} \end{cases}$$

and the transition probability is $p^{\lambda}_{\rm po}(0) = e^{-\lambda}$.

`d')` *Buffering*:

$$\langle n, c, \mathtt{b}\rangle \xrightarrow{\mathtt{d}'} \begin{cases} \langle n\text{-}k{+}1, c\text{-}k, \mathtt{b}^{+}\rangle & \text{if } c \ge n{+}1 \text{ (recovery)} \\ \langle n\text{-}k, c\text{-}k, \mathtt{b}^{+}\rangle & \text{if } c < n{+}1 \text{ (no recovery)} \end{cases}$$

for $1 \le \mathtt{b} \le \mathtt{b}_{\rm buf}$ and $n\text{-}k \ge 1$, and the transition probability is $p^{\lambda}_{\rm po}(k) = \frac{\lambda^k e^{-\lambda}}{k!}$. We define $\mathtt{b}^{+} \triangleq \mathtt{b}{+}1 (\bmod\ \mathtt{b}_{\rm buf}{+}1)$ (i.e., the number of buffering slots increases by 1, until reaching $\mathtt{b}_{\rm buf}$, and then proceed to transmissions by setting $\mathtt{b} = 0$).

`e')` Otherwise, the transition probability for all other pairs of states is 0.

Note that transition $\mathtt{d}'$) is defined to capture battery consumption in the buffering slots before actual transmissions. In Figs. 11-12, we illustrate the state transitions in a graphical manner, with $\mathtt{b}_{\rm buf} = 1$.

Similarly, let $\mathtt{B}^{\mathtt{b}}_{m}(n)$ be the expected battery runtime starting at state $\langle n, n{+}m, \mathtt{b}\rangle$ in $\mathcal{M}_{\rm buf}$. We obtain:

$$\mathtt{B}^{\mathtt{b}}_{0}(n) = \mathtt{W}^{0}_{0}(n) \quad (3)$$

For $m \ge 1$, see Eqn. (4).

Theorem 4: If $\mathtt{b}_{\rm buf} = \mathtt{t}_{\rm sat}$, buffering is always better: $\mathtt{B}^{\mathtt{b}}_{m}(n) \ge \mathtt{W}^{\mathtt{t}}_{m}(n)$ for any $0 \le \mathtt{b} \le \mathtt{b}_{\rm buf}$ and $0 \le \mathtt{t} \le \mathtt{t}_{\rm sat}$.

Theorem 4 affirms the usefulness of duty cycling and buffering. For a more quantitative measure, in Fig. 13, we numerically solve the solutions of Eqns. (3)-(4) for specific m and n. We observe that there is a significant improvement of the battery runtime. We examine the gain of using duty cycling and buffering in Fig. 14, where we identify the specific n's in Fig. 13 that maximise the difference between $\mathtt{W}^{0}_{40}(n)$ and $\mathtt{B}^{0}_{40}(n)$ for different values of λ, $\mathtt{b}_{\rm buf}$, $\mathtt{t}_{\rm sat}$. We observe that the maximum gain is up to 200%. The effectiveness decreases as $\mathtt{b}_{\rm buf}$ and $\mathtt{t}_{\rm sat}$ increase, because the larger saturation threshold implies more the likely battery recovery is, and duty cycling and buffering will not improve too much.

$$\mathtt{B}^{\mathtt{b}}_{m}(n) = \begin{cases} 1 + \sum_{k=0}^{n-1} p^{\lambda}_{\rm po}(k)\mathtt{B}^{\mathtt{b}^{+}}_{m\text{-}1}(n\text{-}k\text{+}1) & \text{if } \mathtt{b}_{\rm buf} \geq \mathtt{b} > 0 \\ 1 + p^{\lambda}_{\rm po}(0)\mathtt{B}^{0}_{m\text{-}1}(n\text{+}1) + \sum_{k=1}^{n-1} p^{\lambda}_{\rm po}(k)\mathtt{B}^{1}_{m}(n\text{-}k) & \text{if } \mathtt{b} = 0 \end{cases} \tag{4}$$

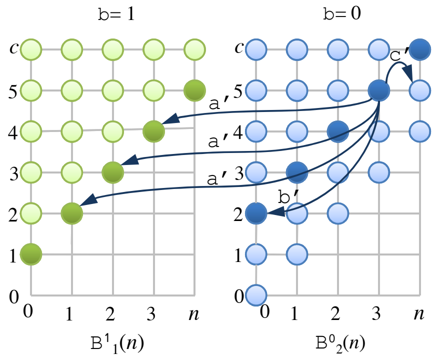

Fig. 11. State transitions for Markov chain $\mathcal{M}_{\rm buf}$ at $\mathtt{b} = 0$.

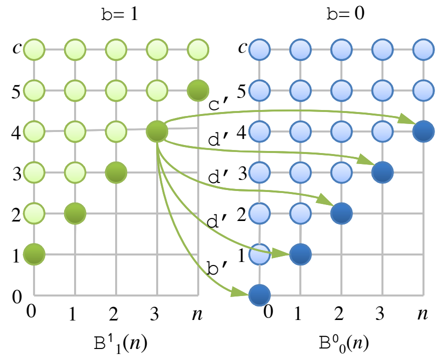

Fig. 12. State transitions for $\mathtt{b} = 1$, when $\mathtt{b}_{\rm buf} = 1$.

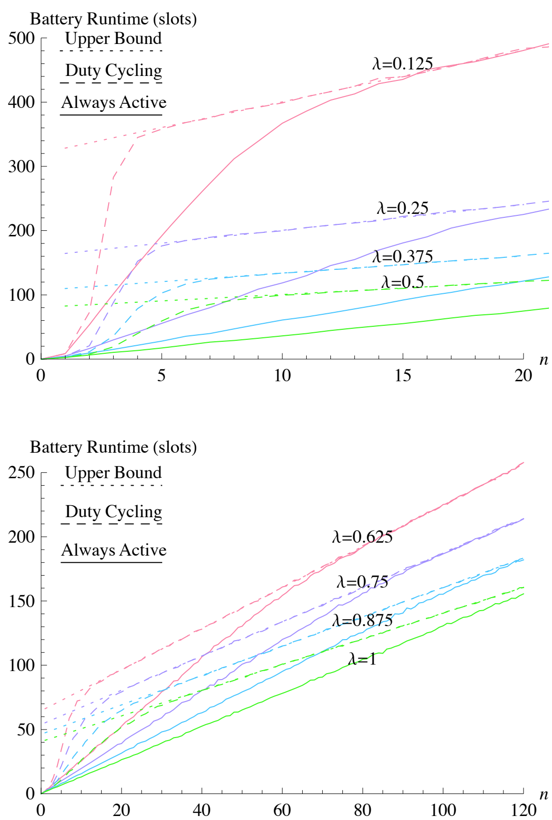

Fig. 13. When $\mathtt{b}_{\rm buf} = \mathtt{t}_{\rm sat} = 1$, we plot of the upper bound $\frac{n+m-1}{\lambda} + \frac{1}{1-e^{-\lambda}}$, along with the battery runtime of always-active case $\mathtt{W}^{0}_{40}(n)$, and the duty cycling case $\mathtt{B}^{0}_{40}(n)$.

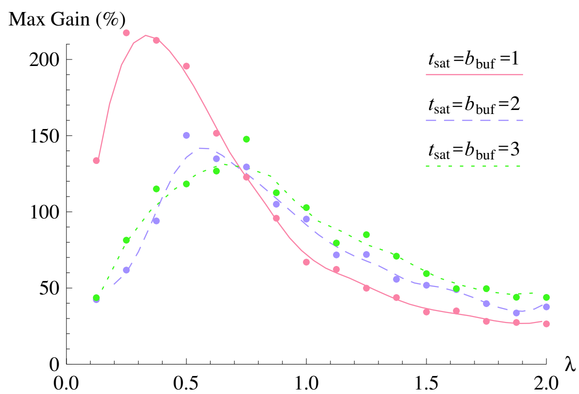

Fig. 14. The maximum gain of the duty cycling case is shown against the always-active case over all n. We obtain smooth curves using moving average.

V. Multi-hop Sensor Networks

In multi-hop sensor networks, a sensor not only transmits, but also listens and receives from its neighbours. Since listening still consumes considerable energy, we are motivated to use duty cycle to regulate radio, such that in sleep mode the radio is off, while in active mode the radio is on for all operations. As in Sec. IV, this can reduce unnecessary idle listening and harness battery recovery.

However, this also creates a problem of requiring a co-ordination scheme among the duty cycled sensors, such that each sensor can discover if its neighbour is active. A simple scheme is to employ a global coordinator that assigns a periodic/deterministic duty cycle schedule to each sensor in a-priori manner such that all sensors are provided the knowledge of duty cycle schedules of their neighbours. However, this scheme suffers from serious scalability issue, and cannot cope with the rapid changing network topology due to ad hoc sensor deployments. One may also consider contention-based approaches, such that all sensors are active time periodically at the same time and contend for transmission opportunities (e.g., S-MAC). However, this incurs significant collisions and overhearing, consuming considerable energy.

Another approach is to turn radio off and on at each slot independently and randomly. Then, packets can only be forwarded when the transmitting and receiving sensors are both active. This requires no global coordinator, and is self-configuring, with reduced collisions and overhearing. But random duty cycle also suffers from control problem — a sensor needs to probe the availability of its neighbours.

Here, we adapt a distributed randomised approach from [3], by which each sensor infers the random duty cycle schedule of the neighbours based on pseudo-random sequence. If the transmitter knows the seed and cycle position of the pseudo-random sequence generator used by the receiver to generate the random duty cycle schedule, it can deterministically predict the active schedules of the receiver without probing. Pseudo-random sequence has been used in other MAC protocols (e.g., SEEDEX, O-MAC [14], [15]) for reducing MAC overhead and overhearing. Next, we modify pseudo-random duty cycle to exploit the saturation threshold of battery recovery.

A. Pseudo-Random Duty Cycle Scheme

As in Sec. IV, we assume that when a sensor switches to sleep mode, only the radio transceiver is off, keeping the processing and sensing units on for sensing activities. Hence, it will not undermine the sensing functionality of the sensor network. All sensors will buffer all their outgoing data until the slot when their respective receivers are active. There is a unique sink in the network to collect all the data.

The pseudo-random duty cycle scheme is as follows:

1) At bootstrapping, the sensors exchange the seed, cycle position and duty cycle rate (as the probability threshold of a slot if it is active or asleep) of their pseudo-random sequence generators with neighbours. At each slot, a sensor determines its state (being active or asleep), and the states of all its neighbours in the next slot.
2) To forward packets, a sensor will wait until there is an active neighbour[4] on its shortest path to the sink, and then transmit the packets with corresponding receiver ID in the header. Uniform random tie-breaking is used if there are more than one active neighbour. The receiver of the corresponding ID will carry out the forwarding of the packets until reaching the sink.

The advantages of using pseudo-random sequence are that 1) the active schedules are essentially uncorrelated and random among the sensors, reducing MAC collisions and overhearing; 2) a sensor can dynamically adjust its duty cycle rate by changing the probability threshold of pseudo-random sequence generator, and 3) the active schedules can be compactly distributed in the networks, represented by the tuple of the seed, cycle position and duty cycle rate. Like other synchronised MAC protocols (e.g., S-MAC), pseudo-random duty cycle requires synchronisation, which will be addressed in Sec. VII. And we assume the arrival packet rate is not high, and the collision of simultaneous transmitters is negligible[5].

B. Battery Recovery Awareness

To extend the pseudo-random duty cycle scheme to take advantage of battery recovery, we propose a simple scheme by *forced sleep*. Suppose that a sensor has been active for more than $\mathsf{w}_{\max}$ consecutive slots from the current slot, then it must go to sleep for the next b_{buf} slots for some $\mathsf{b}_{\text{buf}} \leq \mathsf{t}_{\text{sat}}$. This allows sufficient battery recovery to improve the deliverable energy of a battery. After b_{buf} slots, the sensor decides its active schedule according to its pseudo-random sequence generator. We suppose that $\mathsf{w}_{\max}$ and b_{buf} are known to all sensors. Each sensor can still predict its neighbours' battery-aware duty cycle scheme accordingly. A typical setting will be $\mathsf{w}_{\max} = 1$ and $\mathsf{b}_{\text{buf}} = \mathsf{t}_{\text{sat}}$.

C. Evaluations

We define the *network lifetime* as the expected time that there is a relaying sensor running out of its battery. Particularly, we evaluate the network lifetime in 2D lattice network, where we index each sensor as (i, j) by integers i, j. There is a link between nodes (i, j) and (i', j'), if ($|i\text{-}i'| = 1$ and $j = j'$) or ($|j\text{-}j'| = 1$ and $i = i'$). Without loss of generality, we denote $(0, 0)$ as the sink.

We suppose that each sensor employs greedy forwarding (i.e., forwarding packets to the neighbour that is on the shortest paths to the sink and is active in the next earliest slot) and random tie-breaking when there are more one eligible neighbour. We assume that the duty cycle rates of all sensors are the same. In lattice network, the routing algorithm proceeds as follows:

1) For $i \geq 1$, sensor $(i, 0)$ will forward to $(i\text{-}1, 0)$, as there is only one shortest path.
2) For $i, j \geq 1$, (i, j) will randomly forward to $(i\text{-}1, j)$ or $(i, j\text{-}1)$ with equal probability.

We employ the same discrete battery model as Sec IV for each sensor. We assume that transmission consumes the same energy as reception, and the sensing events of all sensors follow the an independent Poisson distribution $p^{\lambda}_{\text{po}}(k)$. We defer the study of correlated sensing events to the future work.

In multi-hop sensor network, the Markov chain is difficult to analyse. Hence, we rely on simulation to study the network lifetime. By simulation, we compare the network lifetime of battery-aware, pseudo-random duty cycle schemes and simple always-active case. To demonstrate the effectiveness, we select some typical parameters, such as the duty cycle rate as 0.5, $m = 40$ and $n = 150$. Fig. 15 show that the gain of pseudo-random duty cycle scheme is up to 50%, while the gain of battery-aware duty cycle scheme is up to 100%. And the network lifetime of battery-aware duty cycle scheme is generally longer than the one of pseudo-random duty cycle scheme.

VI. Latency Analysis

Increasing the sleep time period of a sensor to maximise battery recovery will inevitably increase the latency of delivering a packet to the sink. In this section, we provide analytical results for the latency of data delivery in sensor networks with both battery-aware and pseudo-random duty cycle.

Suppose node i is waiting to forward data, which has a set of neighbours $\mathcal{N}_i$ and degree as $d_i = |\mathcal{N}_i|$. Each of these neighbours is performing pseudo-random duty cycle with probability ρ_{dc}, such that in one time slot, each node is active with i.i.d. probability ρ_{dc}, and is asleep with probability $1\text{-}\rho_{\text{dc}}$.

Let $L^{(i)}$ be the random number of slots for i before a neighbour becomes active. Therefore,

[4]Note that when even the sensor's own pseudo-random duty cycle schedule prescribes the slot is to sleep, it will still be active to transmit data to the active neighbour.

[5]Even the packet rate is moderate, independent random duty cycle on sensors already cut down the number of potential collisions. Also, one can use a standard MAC protocol within a slot to solve the problem of collision.

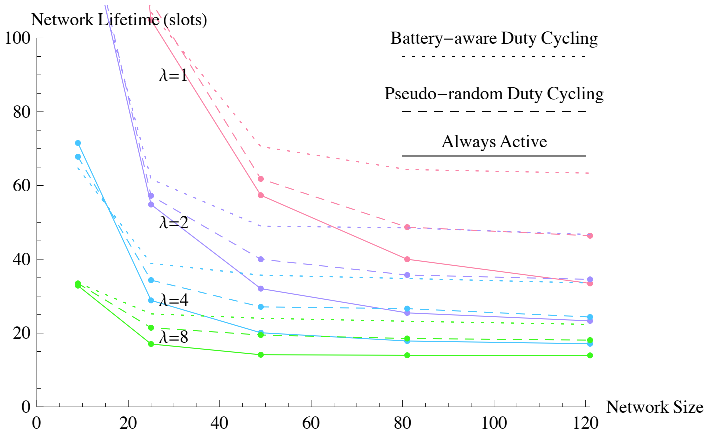

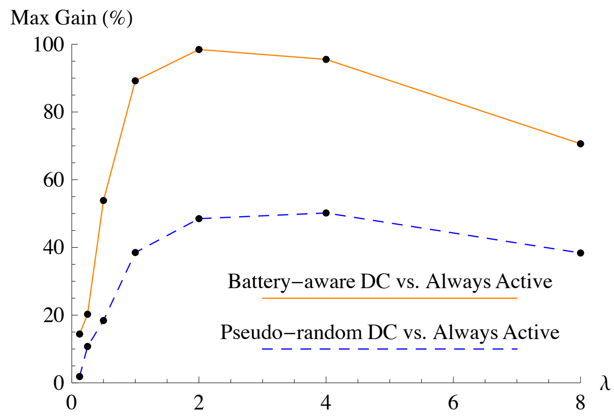

Fig. 15. In a lattice with $\mathsf{b}_{\mathrm{buf}} = \mathsf{t}_{\mathrm{sat}} = 3$, we obtained through simulation the network lifetime of battery-aware duty cycle scheme, pseudo-random duty cycle scheme and simple always-active case. We set the duty cycle rate as 0.5, $m = 40$, $n = 150$, $\mathsf{w}_{\max} = 1$.

$L^{(i)} = \min\{L_1, L_2, \ldots, L_{d_i}\}$, where L_j is the random waiting time for neighbour $j \in \mathcal{N}_i$ to become active.

Theorem 5: $\mathbb{E}[L^{(i)}] = \dfrac{1}{1-(1-\rho_{\mathrm{dc}})^{d_i}}$

Note that the expected per-hop latency decreases quickly with increasing node degree d_i. For very low duty cycle rate (i.e., very small value ρ_{dc}), we obtain an approximation as:

$$\mathbb{E}[L^{(i)}] \approx \frac{1}{1-(1-\rho_{\mathrm{dc}} d_i)} = \frac{1}{\rho_{\mathrm{dc}} d_i}$$

Theorem 6: For 2D lattice, let $\ell(i,j)$ be the end-to-end latency from (i,j) to $(0,0)$.

(1) $\mathbb{E}[\ell(i,0)] = 1 + \dfrac{i\text{-}1}{\rho_{\mathrm{dc}}}$ and $\mathbb{E}[\ell(0,j)] = 1 + \dfrac{j\text{-}1}{\rho_{\mathrm{dc}}}$

(2) For $i,j \geq 1$, $\mathbb{E}[\ell(i,j)] \leq 1 + \dfrac{i\text{+}j\text{-}1}{\rho_{\mathrm{dc}}}$

In battery-aware duty cycle scheme, a sensor that has been active for more than $\mathsf{w}_{\max}$ consecutive slots from the current slot will go to sleep for the next $\mathsf{b}_{\mathrm{buf}}$ slots.

Theorem 7: In battery-aware pseudo-random duty cycle with forced sleep,

$$\mathbb{E}[L_j] \leq \mathsf{b}_{\mathrm{buf}} + \frac{1}{\rho_{\mathrm{dc}}} \quad \text{and} \quad \mathbb{E}[L^{(i)}] \leq \mathsf{b}_{\mathrm{buf}} + \frac{1}{1-(1-\rho_{\mathrm{dc}})^{d_i}}$$

For 2D lattice, let $\ell(i,j)$ be the expected end-to-end latency from (i,j) to $(0,0)$.

(1) $\mathbb{E}[\ell(i,0)] \leq 1 + (i\text{-}1)\big(\mathsf{b}_{\mathrm{buf}} + \frac{1}{\rho_{\mathrm{dc}}}\big)$ and $\mathbb{E}[\ell(0,j)] \leq 1 + (j\text{-}1)\big(\mathsf{b}_{\mathrm{buf}} + \frac{1}{\rho_{\mathrm{dc}}}\big)$

(2) For $i,j \geq 1$, $\mathbb{E}[\ell(i,j)] \leq 1 + (i\text{+}j\text{-}1)\big(\mathsf{b}_{\mathrm{buf}} + \frac{1}{\rho_{\mathrm{dc}}}\big)$

We next present a good approximation of $\mathbb{E}[L^{(i)}]$. When $\mathsf{w}_{\max} = 1$, we can approximate $\mathbb{E}[L_j] \approx \mathsf{b}_{\mathrm{buf}} + \frac{1}{\rho_{\mathrm{dc}}}$, Next, we consider a Markov chain of the state of neighbour j. Let p_j^{act} be the probability that at any slot j is active. Since it is a 2-state Markov chain, by standard Markov chain theory on the recurrent times, $p_j^{\mathrm{act}} = \frac{1}{\mathbb{E}[L_j]} \approx \frac{1}{\mathsf{b}_{\mathrm{buf}} + \frac{1}{\rho_{\mathrm{dc}}}}$. Since the pseudo-random duty cycle of all neighbours are independent, we can obtain the following approximation in a similar as Theorem 5:

$$\mathbb{E}[L^{(i)}] \approx \frac{1}{1-\left(1-\frac{1}{\mathsf{b}_{\mathrm{buf}}+\frac{1}{\rho_{\mathrm{dc}}}}\right)^{d_i}}$$

In Fig. 16, we observe that the approximation is indeed accurate as compared to simulation.

The above theorems enable sensor network designers a useful tool to balance and optimise the trade-off between improving battery runtime and the incurred latency of data delivery. Here we discuss a useful application of our theorem. Suppose that we design a sensor network with a latency constraint. We let the maximum tolerable latency as $\mathbb{E}[L^{(i)}]$ and $\mathbb{E}[\ell(i,j)]$, and obtain the minimum $\mathsf{b}_{\mathrm{buf}}$ from Theorem 7. Such $\mathsf{b}_{\mathrm{buf}}$ can be used to infer the battery runtime from experimental data.

VII. Discussion

A. Significance of Battery Recovery

Energy efficiency can be improved from both the supply side and the demand side. Improvements from the supply side include new battery materials and technologies. However, the development of new practical batteries in recent years has been slow. Although fuel cells may bring a large leap in energy supply, we argue that they are not available in physical format that are appropriate for economical sensor networks that are required to deploy in large scale and outdoors in long unattended periods. Another conceptual alternative can be energy harvesting (e.g., solar cells). But the availability of energy source will significantly constrain the applications of sensor networks. Whatever energy source is used, efficient energy management from the demand side is always desirable. In this work, we have recently examined by experiments the benefit of harnessing battery recovery effect, and found a significant extension of battery runtime (up to 45% normalised gain) by taking advantage of the intrinsic battery characteristics.

B. Applications

In this paper, we found that appropriate (deterministic and randomised) duty cycle schedules can increase the deliverable energy of a battery without jeopardising latency of data delivery. This is particularly useful to sensor network applications with energy and latency constraints, such as security and emergence surveillance. In these applications, duty cycle schedules may be used to prolong the network

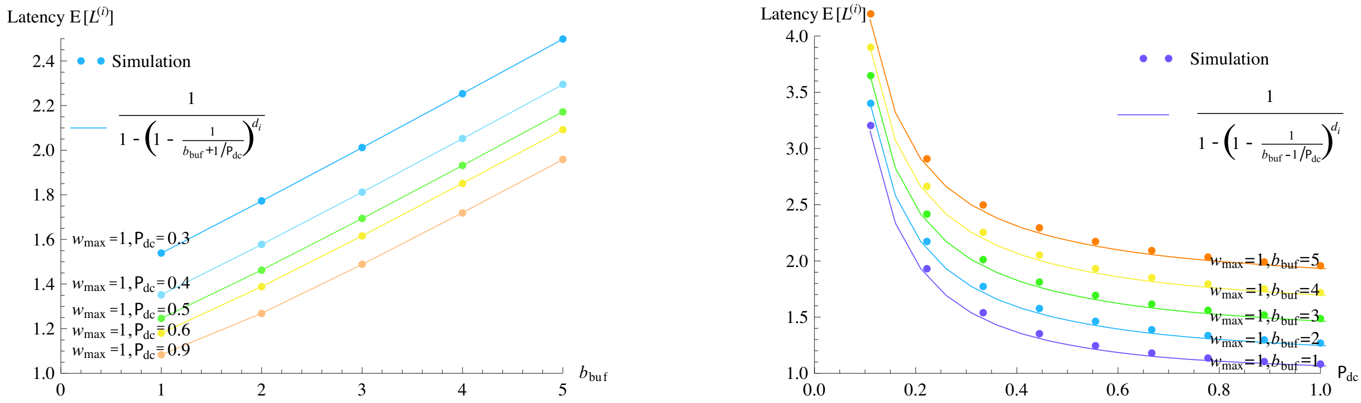

Fig. 16. The figure shows the per-hop latency $\mathbb{E}[L^{(i)}]$ with $d_i = 4$ and $\mathsf{w}_{\mathsf{max}} = 1$. The coloured dots are obtained from simulation, while the solid lines are by equation $\frac{1}{1-\left(1-\frac{1}{\mathsf{b}_{\mathsf{buf}}+\frac{1}{\rho_{\mathrm{dc}}}}\right)^{d_i}}$.

lifetime. However, as shown by our experiments, duty cycle rate is not the only factor to determine the effectiveness of battery recovery effect. Even with the same duty cycle rate, some sleep/active time periods can have a significant gain by the battery recovery effect. If we carefully set the sleep time periods within the saturation threshold, we can maximise battery recovery without exacerbating the latency of data delivery. Other applications that can tolerate latency (e.g., monitoring daily temperature) usually have very long sleep time period, where battery recovery is fully utilised.

C. Synchronisation Protocols

Like other synchronised energy optimising protocols (e.g., S-MAC [13]), we assume that the sensors synchronise their clock with neighbours using a popular low-overhead time synchronisation protocol (e.g., FTSP [20]) in sporadic intervals. The need for time synchronisation is due to the possible clock drift on the sensors that can cause errors in the duty cycle schedules. Small and low-end sensors may suffer from inconsistent clocks because of several factors (temperature, hardware jitter, instability of the clock crystal). However, we argue that active/sleep time periods are much longer than the clock drift (within 40 μs per second as measured in [20]). It does not require frequent time synchronisation. Moreover, relatively accurate time synchronisation is only required for per-hop basis. Our duty cycle does not require network-wide time synchronisation that comes with a high complexity.

VIII. Conclusion

This paper examines the gain of battery recovery effect and provides analytical results to shed light on harnessing the battery recovery in sensor networks. In particular, we analyse battery recovery in the presence of saturation threshold and random sensing activities, based on our experiments. We derive upper bounds of battery runtime and study the benefit of duty cycling and buffering. We then propose a more energy-efficient duty cycle scheme that is aware of battery recovery, by extending the pseudo-random duty cycle scheme. We provide analytical results that predict the latency of data delivery in sensor networks when considering battery recovery optimisation. In future work, we aim to study a broader scope of optimising battery recovery effect in conjunction with a variety of qualities of service observed in sensor networks, such as coverage, connectivity, reliability. We will also consider the impact of interference in wireless communications on battery recovery effect, and correlated sensing activities.

Acknowledgment

The authors thank Yunsheng Wang and Xin Yin for helps in the experiments, and the reviewers for helpful comments.

References

[1] D. Linden, *Handbook of Batteries and Fuel Cells*. McGraw-Hill, 1994.
[2] I. Buchmann, *Batteries in a Portable World: A Handbook on Rechargeable Batteries for Non-Engineers*. Cadex Electronics Inc, 2001.
[3] J. Redi, S. Kolek, K. Manning, C. Partridge, R. Rosales-Hain, R. Ramanathan, and I. Castineyra, "Javelen: An ultra-low energy ad hoc wireless network," *Ad Hoc Networks Journal*, vol. 5, no. 8, 2008.
[4] P. Basu and C.-K. Chau, "Opportunistic forwarding in wireless networks with duty cycling," in *Proc. ACM Workshop on Challenged Networks (CHANTS)*, September 2008.
[5] C.-K. Chau, F. Qin, S. Sayed, M. H. Wahab, and Y. Yang, "Harnessing battery recovery effect in wireless sensor networks: Experiments and analysis," University of Cambridge, Tech. Rep., 2009.
[6] C. Park, K. Lahiri, and A. Raghunathan, "Battery discharge characteristics of wireless sensor nodes: An experimental analysis," in *Proc. IEEE Secon*, 2005, pp. 430–440.
[7] M. Doyle and J. S. Newman, "Analysis of capacity-rate data for lithium batteries using simplified models of the discharge process," *J. Applied Electrochem*, vol. 27, pp. 848–856, July 1997.
[8] T. F. Fuller, M. Doyle, and J. S. Newman, "Modeling of galvanostatic charge and discharge of the lithium polymer insertion cell," *J. Applied Electrochem*, vol. 140, pp. 1526–1533, 1993.
[9] C. F. Chiasserini and R. R. Rao, "Improving battery performance by using traffic shaping techniques," *IEEE J. Sel. Areas in Commun.*, vol. 19, pp. 1385–1394, July 2001.
[10] ——, "Energy efficient battery management," *IEEE J. Sel. Areas in Commun.*, vol. 19, pp. 1235–1245, July 2001.
[11] S. Sarkar and M. Adamou, "A framework for optimal battery management for wireless nodes," *IEEE J. Sel. Areas in Commun.*, vol. 21, no. 2, 2002.
[12] V. Rao, G. Singhal, A. Kumar, and N. Navet, "Battery model for embedded systems," in *Proc. Intl. Conf. on VLSI Design*, 2002.
[13] W. Ye, J. Heidemann, and D. Estrin, "Medium access control with coordinated adaptive sleeping for wireless sensor networks," *IEEE/ACM Trans. Netw.*, vol. 12, pp. 493–506, June 2004.
[14] R. Rozovsky and P. R. Kumar, "SEEDEX: A MAC protocol for ad hoc networks," in *Proc. ACM MobiHOC*, 2001.
[15] H. Cao, K. Parker, and A. Arora, "O-MAC: A receiver centric power management protocol," in *Proc. IEEE ICNP*, 2006.

[16] Y. Sun, O. Gurewitz, and D. Johnson, “RI-MAC: a receiver-initiated asynchronous duty cycle MAC protocol for dynamic traffic loads in wireless sensor networks,” in *Proc. ACM SenSys*, 2008.
[17] S. Jayashree, B. Manoj, and C. S. R. Murthy, “On using battery state for medium access control in ad hoc wireless networks,” in *Proc. ACM MOBICOM*, 2004.
[18] M. Dhanaraj, S. Jayashree, and C. S. R. Murthy, “A novel battery aware mac protocol for minimizing energy $\times$ latency in wireless sensor networks,” in *Proc. HiPC*, 2005.
[19] C. F. Chiasserini, P. Nuggehalli, V. Srinivasan, and R. R. Rao, “Energy-efficient communication protocols,” in *Proc. DAC*, 2002.
[20] M. Maróti, B. Kusy, G. Simon, and A. Lédeczi, “The flooding time synchronization protocol,” in *Proc. ACM SenSys*, 2004.

Chi-Kin Chau is currently a Research Associate at University of Cambridge as a primary researcher for the International Technology Alliance in Network and Information Science (ITA) program that is jointly funded by the U.S. Army Research Laboratory and the U.K. Ministry of Defense. Previously, he was a Croucher Foundation Research Fellow at the EE Dept., University College London, awarded a two-year fellowship by the Croucher Foundation Hong Kong. He received the Ph.D. from the Computer Laboratory, University of Cambridge, and the B. Eng. (First-class Honors) from the Dept. of Information Engineering, the Chinese University of Hong Kong. His research interests concern diverse areas of networking, communications, distributed systems, stochastic networks and randomized algorithms.

Fei Qin received the B.Eng in Information Engineering from HUST, 2004 and M.Eng. in Electronic Engineering from BIT, 2006. He is currently funded by DHPA scholarship to work as PhD Candidate with Dr. John Mitchell and Dr. Yang Yang in EE Dept., University College London. Prior to that, he served as Sr. Application Engineer at Crossbow Technology, Beijing Rep. Office.

Samir Sayed received his B.S and M.S. degrees in Electronics and Communications Engineering from Helwan University, Egypt, in 1997 and 2003, respectively. Currently, he is a Ph.D. candidate at the Department of EE Dept., University College London. His present research interests include wireless networking, medium access control, quality of service, and cooperative communications.

Muhammad Husni Wahab received the BSc in Electronics and Computer Engineering from Carnegie Mellon University and MSc from Glasgow University in 2008. He worked on machine learning and networking protocols. He is now a PhD student in the Computer Science department, University College London, working on machine learning for vision systems under the Vision and Imaging Science group.

Yang Yang is currently the Assistant Director and Vice CTO at Shanghai Research Center for Wireless Communications (WiCO), SIMIT, Chinese Academy of Sciences. Prior to that, he was a Senior Lecturer with EE Dept., University College London, United Kingdom. His general research interests include wireless ad hoc and sensor networks, wireless mesh networks, and 3G/4G mobile systems.