Energy-Efficient Communication Protocols

Carla F. Chiasserini
Dip. di Elettronica Politecnico
di Torino
C.so Duca degli Abruzzi 24
10129 Torino, Italy
chiasserini@polito.it

Pavan Nuggehalli
Dept. of ECE
UCSD
9500 Gilman Dr.
La Jolla, CA 92093
pavan@cwc.ucsd.edu

Vikram Srinivasan
Dept. of ECE
UCSD
9500 Gilman Dr.
La Jolla, CA 92093
vikram@cwc.ucsd.edu

ABSTRACT

Wireless networking has experienced a great deal of popularity, and significant advances have been made in wireless technology. However, energy efficiency of radio communication systems is still a critical issue due to the limited battery capacity of portable devices. In this paper, we deal with the charge recovery effect that takes place in electrochemical cells and show how we can take advantage of this mechanism to increase the energy delivered by a battery. Then, we present energy-aware traffic shaping techniques, as well as scheduling and routing algorithms, which exploit the battery recovery effect.

Categories and Subject Descriptors

C.2.2 [**Computer-Systems Organization**]: Network Protocols

General Terms

Algorithms

Keywords

Wireless networks, energy efficiency, battery charge recovery

1. INTRODUCTION

Wireless devices have become an integral part of modern society. These devices are employed in cellular, ad hoc and sensor networks to name a few. A common thread restraining the true potential of these applications is the limitation in energy supply which are predominantly provided by batteries. For example, in cellular and ad hoc networks, the addition of wireless interfaces to portable devices allows for connectivity to the Internet. In order to provide ubiquitous and continuous access to Internet-based services, a significant energy supply is required from batteries. The same

requirement exists in sensor networks, which are often deployed in inaccessible locations. In this case, the nodes batteries must have a long lifetime since it may be impossible to replace them.

The major contributions to energy consumption in wireless devices are: (i) the power consumed by the digital part of the device circuitry; (ii) the power consumed by the electronics of the transceiver in transmitting and receiving mode; (iii) the output transmission power.

For devices that operate in the ISM bands, typical values of the power amplifier output range over the interval [1 mW,1000 mW], while the transceiver power consumption is of the order of tens to hundreds of mW [1]. The total power consumption depends on the amount of transmitted, received, and processed traffic. It follows that the greater the energy capacity of the battery, the longer the run time of a wireless device. On the other hand, battery features such as long lifetime, light weight and small size are highly desirable in portable wireless devices. An obvious solution to the problem would be to increase the energy delivered by a battery by improving battery design. However, in the last few years, advances in battery technology have not kept pace with the wireless revolution. An alternative strategy consists in developing energy-efficient protocols and architectures which enhance the devices lifetime.

Various MAC protocols [2, 3] and schemes for power management control during transmissions [5, 6] have been proposed in the literature. Dynamic power management policies have been proposed in [7, 8]. The work in [9] focuses on data link layer mechanisms and show that, if energy is a constraint, then retransmissions should be minimized. In [10, 11], energy aware routing protocols have been proposed. Traditional routing protocols typically implement the shortest path algorithm; thus, it is likely that a few nodes will lie on many routes resulting in their quick death and loss of network connectivity. On the contrary, the solutions presented in [10, 11] distribute the load more evenly among the available routes leading to network longevity.

In this paper, we present a different approach to energy saving. Our goal is to understand the battery behavior and leverage this understanding to develop new energy-efficient protocols. In Section 2, we show that battery life is dependent on its usage pattern. Batteries recover charge during idle periods. This phenomenon can be exploited by the upper layers to design battery aware scheduling and traffic shaping algorithms. Moreover in general, all levels of the protocol stack may contribute to energy saving by intelli-

gently using some knowledge about the behavior of the traffic, the radio channel and the network. A scheduling scheme and a routing protocol exploiting the battery recovery effect are described in Section 3.

2. THE RECOVERY MECHANISM IN BATTERIES

Our research focuses on the understanding of the intrinsic behavior of batteries and on the use of this understanding to develop a stochastic battery model and new energy-efficient protocols.

A battery consists of one or more electrochemical *cells*, organized in an array. Each cell consists of an anode, a cathode and the electrolyte that separates the two electrodes. The electrical current obtained from a cell results from electrochemical reactions occurring at the electrode-electrolyte interface [12, 13, 14]. At zero current, the concentration of the active species in the cell is uniform. As the discharge current increases, the active species are consumed at the electrode-electrolyte interface by electrochemical reactions, and replaced by new active materials that move from the electrolyte solution to the electrode through *diffusion*. As the intensity of the current is increased, the deviation of the concentration from the average becomes more significant and the state of charge as well as the cell voltage decrease. This phenomenon is called *rate capacity effect*. Beyond a threshold value, called the *limiting* current, the diffusion phenomenon is unable to compensate for the depletion of active materials and the cell voltage drops below the usable value even though the *theoretical* available energy stored in the cell may not have been exhausted. The amount of electrical energy that can be obtained from a battery cell is fundamentally limited by the type and mass of the electrodes and electrolytes. In practice, the delivered specific energy (ratio of the total delivered energy to the weight of the cell) depends on the intensity of the discharge current, the power level drained from the cell and whether the discharge is constant or pulsed.

The typical way to drain power from cells is at a constant level of current (constant current discharge). In this case, the relationship between specific energy and specific power (ratio of the delivered power to the weight of the cell) of different battery technologies, is displayed in the so-called Ragone plot, as shown in Fig. 1. The fact that in the plot the curves "lean" to the left shows that a high specific energy can be obtained only if the discharge is at low power levels. However, even if one settles for a low rate of discharge, the energy delivered by a battery under constant discharge is typically only 10-30% of the theoretical value. Although improvements in battery technology are being made, they tend to lag behind the demand.

Some of the adverse consequences of constant current discharge can be overcome when the discharge is pulsed. If a cell is allowed to relax long enough after delivering a pulse, the diffusion process, mentioned above, compensates for the depletion of the active materials that takes place during the current drain. The concentrations gradient of the electrolyte species decreases, and charge recovery takes place at the electrode. Thus, by using a pulsed discharge instead of a constant current discharge, it turns out that: (i) for a fixed power level, the delivered specific energy is greater; (ii) for a constant delivered specific energy, a higher specific power
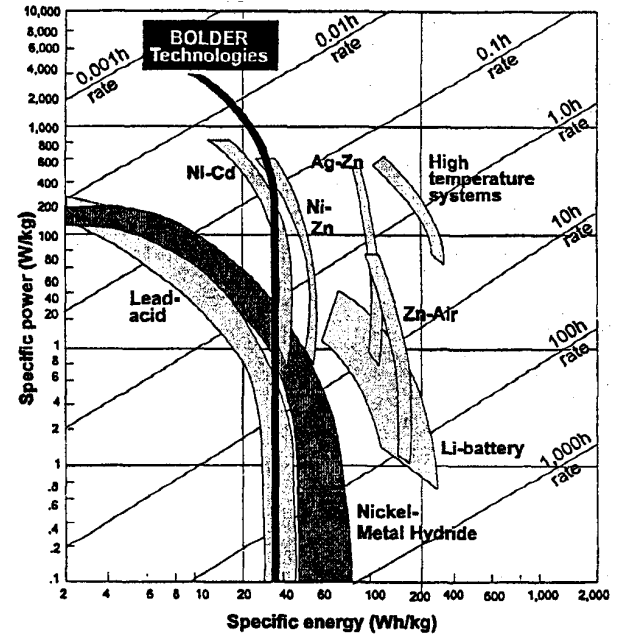

Figure 1: Ragone plots for various battery chemistries.

can be drained from the cell. Several findings [12, 15, 16, 17] quantify these advantages. For instance, in batteries characterized by a relatively low conductivity, e.g., lithium-polymer cells [18], pulsed discharge can increase the delivered specific power. Lithium-polymer cells have lightweight and flat formats and fit well in extra-thin cellular phones, but ion diffusion through the conducting polymers is slow and this limits the rate at which current can be withdrawn from a cell. However, in [19] a lithium-polymer cell was discharged at up to 50 mA/cm^2, about fourteen times the typical constant discharge rate, using pulses of 10 ms duration followed by a 50 ms rest period.

These facts suggest that the burstiness of the battery discharge process may be a more significant determinant of the delivered energy than the initial charge stored in the cell. That being so, we are led to explore the possibility of discharge shaping to maximize the battery energy efficiency. In order to investigate this possibility systematically, we need models for representing the battery behavior.

2.1 Modeling Battery Recovery

We start by considering an electrochemical battery model [20] that represents the underlying electrochemical phenomena in great detail. The electrochemical model has been modified to let the discharge of the cell be driven by a stochastic process, and can be numerically solved by using a program developed by Newman *et al.* [21]. Results are obtained assuming that the time scale is divided into time slots with duration equal to 1 ms and that the discharge process is composed of pulses occurring at stochastic time instants according to a Bernoulli distribution. Pulses have a constant duration equal to 1 ms (i.e., equal to one time slot). We take the current discharge at each pulse as a varying parameter of the system. The specific delivered energy is derived as a function of the number of discharge pulses per time slot. By fixing the current discharge, we obtain that as the number of

current pulses per time slot decreases, the specific delivered energy dramatically increases since the chance to recover for the cell increases. This proves that a significant improvement in performance of real batteries is possible when a stochastic pulsed discharge is used instead of a constant current discharge. By varying the current discharge, results show that higher current pulses may degrade performance, even if the number of higher current pulses is really small. This phenomenon takes place whenever the current drained from a cell exceeds the specified limiting value and, therefore, the concentration gradient of active materials becomes significant. Thus, applications that involve the use of different levels of power (e.g., routing in ad hoc networks, power control in CDMA systems, etc.) should carefully administer the available battery energy. The set of results that we could derive through the electrochemical model however is limited because the computation time for the model solution is exceedingly large. Since our objective is to define a battery model for studying the energy efficiency of communication protocols, we need a more tractable model [22, 23].

We therefore develop a stochastic model, which is much simpler and represents the battery behavior mathematically in terms of parameters that can be related to physical characteristics of the electrochemical cell. The stochastic model captures the recovery effect that is observed when a pulsed discharge is applied, and neglects all the other phenomena represented in the electrochemical model. The model characterizes a fully charged cell by a certain theoretical energy and by a nominal energy, the latter representing the energy that could be extracted using a constant discharge profile. Notice that the nominal energy is much less than the theoretical energy. The recovery effect is modeled as a decreasing exponential function of the state of charge and delivered energy. To more accurately model real cell behavior, the exponential decay coefficient is assumed to take different values as a function of the delivered energy. In each time slot, some energy units are lost if any request of current supply arrives and has to be satisfied, otherwise the battery may recover or remain in the same state. The resulting cell behavior is a transient stochastic process that starts from the state of full charge and terminates when the cell is completely discharged, or the theoretical available energy is exhausted. The stochastic model has been matched to the electrochemical model for the particular case of a lithium-ion battery, and a comparison between the curves obtained from the stochastic and the electrochemical model has showed a very good agreement [22].

2.2 Application: An Example

By using the stochastic battery model, we are able to explore ways in which the energy efficiency of mobile wireless communications can be enhanced through the use of improved energy-efficient battery management techniques. We consider an array of L cells connected in parallel, that can be selectively scheduled. Fig. 2 illustrates the case for L=2. A delivered energy equal to N can be drained from each cell by means of a constant current discharge. However, by using a proper pulsed discharge technique, the region of delivered energy can be increased. The goal is to find methods to extend the region of delivered energy up to the maximum limit, i.e., the theoretical available energy T. Cells can be scheduled based on the discharge demand and their own status, so that the demand can be distributed in a manner that
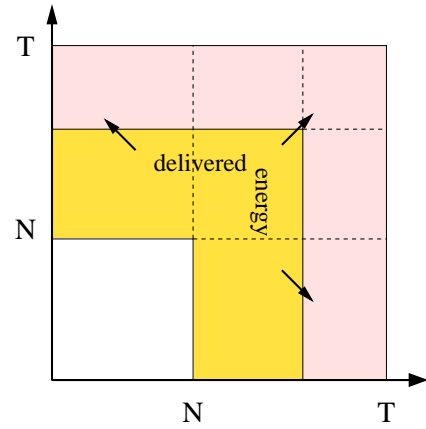

Figure 2: Margin of improvement for the energy delivered by an array of L=2 cells.

provides adequate rest periods for individual cells. We apply load balancing algorithms to battery management and propose both a *delay-free* and *delayed delivery* approach.

The former implements a scheduling scheme among the cells and provides the power supply as soon as required. A technique based on such an approach is the round robin scheduling, i.e., the discharge request arrivals are directed to the cells by switching from a cell to the next one. By using the stochastic battery model, we can evaluate the improvement in delivered battery energy that is achieved.

The latter approach uses scheduling algorithms in conjunction with discharge shaping, so that an additional delay is introduced in the power supply. In this case, the battery is always able to deliver its maximum available energy even for high discharge demand. This approach involves a coordination among the cells of the array and drains current from the cells according to their state of charge. The goal is to monitor the cells status and make them recover as much as they need to obtain the maximum available energy from the discharge process. We define a lower threshold for the cell state of charge such that, whenever the state of charge drops to this value, the cell is considered *not active* and current cannot be drawn off until recovery occurs. The event of the cell state of charge dropping to a certain threshold can be easily revealed by the control apparatus present in smart battery packages; the acquisition of the information is therefore considered instantaneous. A scheme, that can be implemented to schedule a cell among the set of the active ones, consists in selecting the cell that currently has the best state of charge. Whenever the state of charge of a cell drops to the value of threshold, the cell is removed from the set of active cells and allowed to recover. If none of the cells is active, the discharge requests that arrive are buffered and served as soon as a sufficient amount of charge is recovered. By applying this scheme, a battery is able to deliver the maximum available energy at the cost of a fairly small additional delay and complexity.

3. EXPLOITING THE RECOVERY EFFECT IN COMMUNICATION PROTOCOLS

Typically, energy efficiency can be obtained at the expense of other performance measures, such as delay and

throughput. We provide two examples in which we show how communication protocols can exploit battery behavior to increase the delivered energy. The first example concerns scheduling and can be incorporated in the physical layer. The second example deals with routing, allowing for energy efficiency at the network layer.

3.1 Scheduling

We address the problem of designing an energy aware transmission schedule for a wireless node subject to some delay constraints [24]. We do so by combining two unconnected concepts, namely the battery recovery mechanism and channel coding [4]. In [4], it was shown that channel coding can be used to conserve energy by transmitting at reduced power levels over longer durations. The channel coding idea favors extending transmission durations for as long as possible. On the other hand, battery dynamics require the transmitter to be idle to allow for recovery. In other words, the device saves energy by transmitting in bursts and recovering in the idle periods. Delay constraints prohibit the use of arbitrarily large transmission or idle times. Clearly the two energy conserving ideas above are at variance with each other. Thus, a strategy which is based entirely on either one or the other idea is not optimal. We pose the following question: can these two concepts be amicably merged to devise a delay constrained energy-efficient transmission schedule?

In [4], they observe that, in a variety of channel coding situations, the energy required to transmit a bit is a strictly convex and decreasing function of its transmission time. Using this insight, they propose an optimal off-line lazy scheduling policy. The policy is off-line in the sense that the arrival times of all the packets in a bounded interval $(0, T)$ are known a priori. The aim of the policy is to schedule all departures before the deadline T. They contend that, for energy consumption to be minimized, a transmitter should attempt to transmit a packet for as long as possible subject to the deadline constraint, T. Motivated by the off-line analysis, they propose an online algorithm which they show to be near optimal. They then consider the problem of scheduling packets over an infinite time horizon with an average queuing delay constraint. They favorably compare a certain heuristically derived scheduling policy with a constant service policy that results in the same average delay. However they leave open the problem of devising an optimal scheduling policy for an arbitrary average delay constraint.

We first tackle the deadline constraint problem by adapting the optimal off line scheduling policy in [4] to a discrete-time scenario. We then assume a simple battery recovery model and derive the optimal transmission strategy. Our results show that energy savings of more than 50% are feasible by accounting for energy recovery.

Next, we consider the average delay case. In this model, packets arrive at a transmitter's queue according to a Poisson process. Packets are served on a first come, first serve basis. We use constrained dynamic programming (CDP) [25], to devise an optimal transmission strategy which uses only the channel coding concept and ignores battery dynamics. The resulting transmission policy is called the non-battery aware policy. We then consider a stochastic model of battery behavior incorporating battery recovery. This allows the transmitter to conserve energy by scheduling idle periods after a transmission. However, incorporating a realistic model of energy recovery along with the channel cod-

ing concept in the CDP framework is computationally prohibitive. We resolve this by making simplifying assumptions about battery recovery. This leads to policies which may be sub-optimal. We call these the battery aware policies. For a given average delay constraint, we evaluate the battery aware policy and the non-battery aware policy against a constant service policy using simulation in Fig. 3. Note that with tight delay constraints, the non-battery aware policy can save as much as 40% over a constant service policy. With the battery aware policy, this figure can improve up to 90% for low arrival rates. As the delay constraint is relaxed, the constant service policy is seen to do almost as well as the non-battery aware policy. However, the gains from the battery aware policy become more substantial for low arrival rates.
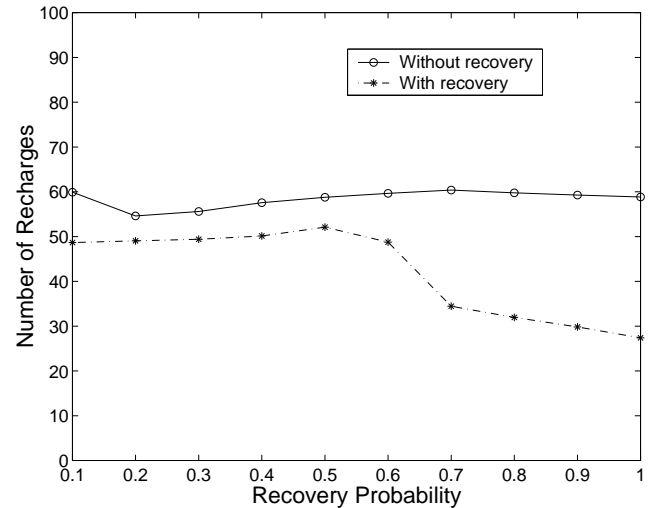

Figure 3: Arrival Rate = 0.1, Average Delay Constraint = 1.5. Each curve shows the number of recharges required per 100 constant service policy recharges. As arrival rate increases, the battery aware policy outperforms the non-battery aware policy for high values of β.

3.2 Routing

We present a routing protocol that aims at maximizing the network lifetime, i.e., the time period from the instant when the network starts functioning to the instant when the first network node runs out of energy [11]. We consider the battery capacity of the nodes as a common resource of the system and we improve the network lifetime by efficiently exploiting the available battery capacity.

We develop a routing protocol, named BEE (Battery Energy-Efficient) protocol [26], that aims to balance the battery consumption among all network nodes. The scheme selects routes whose links have a low energy cost; e.g., a route with 3 hops may be preferred to a route with 2 hops if the maximum level of power per link is lower. Also, the BEE protocol favors the nodes with low battery status by routing the multihop traffic through nodes with high battery capacity and allowing the others to benefit of the recovery effect.

We define K as the set of the nodes in the network, S as the set of traffic sources and G as the set of destination nodes. Nodes belonging to S can direct the traffic to any

of the destination nodes using the other nodes as relays. We consider that time is discretized into unitary intervals corresponding to the packet transmission time. Nodes are assigned an initial amount of energy, denoted by B_i; the instantaneous battery status is denoted by b_i with $i \in K$. The transmission range of each node is limited to a certain value, indicated by ρ, so that a radio link can be established between any pair of nodes (i, j) only if the distance d_{ij} is less than or equal to ρ. We indicate as R_i the set of nodes whose distance from i is less than ρ. The energy spent for each data transmission from node i to node j is [11]: $e_{ij} = (d_{ij}/\rho)^4$ if $j \in R_i$, and infinite otherwise. The energy consumed by node j ($j \in R_i$) to receive a data unit is considered to be ten times lower than the transmission energy e_{ij}.

We discretize the transmission energy to any node in R_i ($i \in K$) into few levels ranging from e_{min} to e_{max}. To represent the benefit of the recovery phenomenon, the battery status of a node is increased whenever the node's transmission rate is sufficiently low. We model the energy accrue at node i as a function of the transmission rate λ_i and of the mean energy necessary to node i to transmit a packet ($\overline{e_i}$). We indicate the energy increase function as $\Gamma(\lambda_i, \overline{e_i})$ with $\overline{e_i} = \frac{1}{R_i} \sum_{j \in R_i} e_{ij}$. $\Gamma(\cdot, \cdot)$ is obtained through the stochastic battery model described in Section 2.1. The stochastic model was matched to the electrochemical model [21] for the case of impulses energy equal to $\overline{e_i}$ drained from the battery at rate λ_i.

Similarly, whenever a node needs to use a transmission energy level $e_{ij} > e_{min}$, the node's battery status is decreased by $e_{ij} + \Phi(e_{ij} - e_{min})$, where $\Phi(\cdot)$ takes into account the rate capacity effect and is derived from the electrochemical model results. Thus, at each time unit the evolution of the battery status is,

$$b_i = \begin{cases} b_i + \Gamma(\lambda_i, \overline{e_i}) & \text{if node } i \text{ is idle} \\ b_i - [e_{ij} + \Phi(e_{ij} - e_{min})] & \text{else} \end{cases} \quad (1)$$

The routing algorithm assigns to each route a cost function taking into account both energy transmission and battery behavior. Given the generic source s and the generic destination g, the cost function associated to the k-th route, r_{sg}^k, is

$$F_k = \sum_{l_{ij} \in r_{sg}^k} [\Psi(\lambda_i)e_{ij} + p_{ij}] - \min_{i \in r_{sg}^k} b_i \quad (2)$$

where: (i) l_{ij} is the link between nodes i and j belonging to route r_{sg}^k. (ii) $\Psi(\lambda_i)$ is a weighting function that emphasizes the energy loss of the source node and depends on the source transmission rate; we have: $\Psi(\lambda_i) = A \cdot \lambda_i$ for $i = s$, with A being a proper constant value greater than 1, and $\Psi(\lambda_i) = 1$ otherwise. (iii) p_{ij} is the energy penalty due to the rate capacity effect experienced at node i when a power level higher than the mean power value is required to transmit a packet over link l_{ij}; we have $p_{ij} = max\{0, e_{ij} - \overline{e_i}\}$. (iv) $\min_{i \in r_{sg}^k} b_i$ is the minimum value of battery status among the nodes belonging to route r_{sg}^k.

Whenever source s has a block of data to transmit toward destination g, s evaluates the cost function for all the possible loop-free routes and selects the route r_{sg}^m such that: $F_m = \min_{r_{sg}^k} F_k$. We notice that every time the cost function has to be computed the information about the nodes battery status have to be updated.

If the set of routes over which we have to minimize the cost

function is very large, the complexity of this scheme may become unacceptable for networks whose topology varies often. A simpler and faster algorithm is obtained as follows. We consider that the node in charge of selecting a route (i.e., either the source or the destination) chooses c routes independently and uniformly at random from all the available routes. The cost function is then evaluated over the c routes, and the selected path is the one among the c routes which minimizes F.
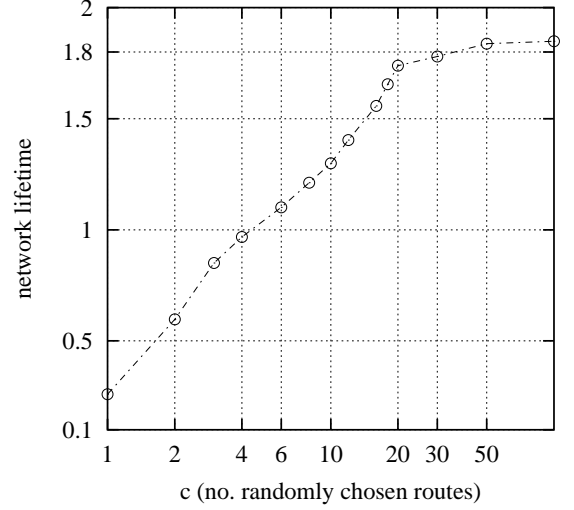

Figure 4: Network lifetime derived through the BEE scheme as a function of c, normalized to the value obtained through the MTE algorithm.

Fig. 4 shows the network lifetime obtained through the BEE scheme. We assume that the network consists of 15 nodes randomly scattered all over a circular area. The maximum transmission range is ρ=2.0 and the length of the routes connecting any pair source-destination can be equal to four hops at most. The initial battery status is set equal to 1 for all the nodes in the network. There are five source nodes with traffic rate $\lambda_0 = \lambda_1 = \lambda_2 = 0.4$ and $\lambda_3 = \lambda_4 = 0.3$, and two fixed destination nodes. A source node updates the cost function and selects a new route to a destination every time it finishes to transmit a block of five packets. Results show the network lifetime obtained through the BEE scheme as c varies, normalized to the value of lifetime derived through the MTE algorithm [27]. For $c > 4$, the BEE algorithm outperforms the MTE technique; the maximum value of network lifetime is obtained when the cost function is evaluated over all the possible routes. In this case, the factor of improvement with respect to the MTE scheme is equal to 1.8. These results suggest that, depending on the level of acceptable complexity, an optimal value of c can be selected.

4. CONCLUSIONS

We addressed the problem of energy efficiency in wireless networks. We proposed strategies that exploit battery dynamics to increase the energy delivered by a battery. We found that the longer the idle times of a battery-powered device, the greater the charge recovery effect and the delivered energy. In order to introduce idle times in the devices'

activity, energy-aware communication algorithms can be developed at each level of the protocol stack. In this paper, we focussed on scheduling and routing algorithms, and showed the improvement in performance that can be achieved by implementing such schemes.

5. ADDITIONAL AUTHORS

Additional authors: Ramesh R. Rao (Dept. of ECE, UCSD, email: rao@cwc.ucsd.edu).

6. REFERENCES

[1] V. Rodoplu and T.H. Meng, "Minimum energy mobile wireless networks," *IEEE J. Selected Areas in Commun.*, vol. 17, no. 8, pp. 1333–1344, August 1999.

[2] J.-C. Chen, K.M. Sivalingam, P. Agrawal, and R. Acharya, "Scheduling multimedia services in a low-power MAC for wireless and mobile ATM networks," *IEEE Transactions on Multimedia,* vol. 1, no. 2, pp. 187–201, June 1999.

[3] A. K. Salkintzis and Christodoulos Chamzas, "An In-Band Power-Saving Protocol for Mobile Data Networks," *IEEE Trans. on Commun.*, vol. 46, pp. 1194–1205, Sept. 1998.

[4] B. Prabhakar, E. Uysal, and A. El Gamal, "Energy-efficient transmission over a wireless link via lazy packet scheduling," *Proc. of Infocom 2001,* April 2001.

[5] N. Bambos and J.M. Rulnick, "Mobile power management for maximum battery life in wireless communication networks," *Proc. of IEEE Infocom'96,* pp. 443–450, March 1996.

[6] B. Narendran, J. Sienicki, S. Yajnik, and P. Agrawal, "Evaluation of an adaptive power and error control algorithm for wireless systems," *Proc. of IEEE ICC'97,* pp. 349–355, 1997.

[7] T. Simunic, L. Benini, P. Glynn, and G. De Micheli, "Dynamic power management for portable systems," *Proc. of IEEE/ACM MobiCom'2000,* Boston, MA, USA, August 2000.

[8] L. Benini, A. Bogiolo, and G. De Micheli, "A survey of design techniques for system-level dynamic power management," *IEEE Transactions on VLSI Systems,* vol. 8, no. 3, pp. 299–316, June 2000.

[9] M. Zorzi and R.R. Rao, "Error control and energy consumption in communications for nomadic computing," *IEEE Transactions on Computers,* vol. 46, no. 3, pp. 279–289, March 1997.

[10] S. Singh, M. Woo, C. S. Raghavendra, "Power-aware Routing in Mobile and Ad Hoc Networks," *Proc. of 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking,* Dallas, TX, pp. 181–190, October 1998.

[11] J.-H. Chang and L. Tassiulas, "Energy conserving routing in wireless ad-hoc networks," *Proc. of Infocom 2000,* Tel Aviv, Israel, March 2000.

[12] H.D. Linden, *Handbook of batteries,* 2nd ed., McGraw-Hill, New York, 1995.

[13] B. Nelson, R. Rinehart, and S. Varley, "Ultrafast pulse discharge and recharge capabilities of thin-metal film battery technology," *11th IEEE International Pulsed Power Conference,* Baltimore, pp. 636–641, June 1997.

[14] 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, no. 7, pp. 846–856, July 1997.

[15] T.F. Fuller, M. Doyle, and J.S. Newman, "Relaxation phenomena in lithium-ion-insertion cells," *J. Electrochem. Soc.,* vol. 141, no. 4, pp. 982–990, April 1994.

[16] R.M. LaFollette, "Design and performance of high specific power, pulsed discharge, bipolar lead acid batteries," *10th Annual Battery Conference on Applications and Advances,* Long Beach, pp. 43–47, January 1995.

[17] R.M. LaFollette and D. Bennion, "Design fundamentals of high power density, pulsed discharge, lead-acid batteries. II Modeling," *J. Electrochem. Soc.,* vol. 137, no. 12, pp. 3701–3707, December 1990.

[18] J.W. Halley and B. Nielsen, "Simulation studies of polymer electrolytes for battery applications," *MRS Symposium Proceedings,* vol. 496, pp. 101–107, December 1997.

[19] H.S. Choe and K.M. Abraham, "Synthesis and characterization of $LiNiO_2$ as a cathode material for pulse power batteries," *MRS Symposium Proceedings,* vol. 496, pp. 303–308, December 1997.

[20] M. Doyle, T.F. Fuller, and J.S. Newman, "Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell," *J. Electrochem. Soc.,* vol. 140, pp. 1526–1533, 1993.

[21] J.S. Newman, *FORTRAN programs for simulation of electrochemical systems,* http://www.cchem.berkeley.edu/~jsngrp/ (Accessed July 1999).

[22] C.F. Chiasserini and R.R. Rao, "Energy efficient battery management," *IEEE J. Selected Areas in Commun.*, vol. 19, no. 7, pp. 1235–1245, July 2001.

[23] C.F. Chiasserini and R.R. Rao, "Improving battery performance by using traffic shaping techniques," *IEEE J. Selected Areas in Commun.*, vol. 19, no. 7, pp. 1385–1394, July 2001.

[24] P. Nuggehalli, V. Srinivasan, and R.R. Rao, "Delay Constrained Energy Efficient Transmission Strategies for Wireless Devices," *Proc. of Infocom 2002,* New York City, NY, June 2002.

[25] L. I. Sennott, "Computing Average Optimal Constrained Policies in Stochastic Dynamic Programming," *Probability in the Engineering and Informational Sciences,* vol. 15, pp. 103–133, 2001.

[26] C.F. Chiasserini and R.R. Rao, "Routing protocols to maximize battery efficiency," *Proc. of IEEE MILCOM 2000,* Los Angeles, USA, October 2000.

[27] T. Shepard, "Decentralized channel management in scalable multihop spread spectrum packet radio networks," *Tech. Rep. MIT/LCS/TR-670,* Massachusetts Institute of Technology Laboratory for Computer Science, July 1995.