Energy Minimization for Real-Time Data Gathering in Wireless Sensor Networks

Yang Yu, *Member, IEEE,* Viktor K. Prasanna, *Fellow, IEEE,* and Bhaskar Krishnamachari, *Member, IEEE*

***Abstract*— This paper studies the challenging problem of energy minimization for data gathering over a multiple-sources single-sink communication substrate in wireless sensor networks by exploring the energy-latency tradeoffs using rate adaptation techniques. We consider a real-time scenario for mission-critical applications, where the data gathering must be performed within a specified latency constraint. We first propose an off-line numerical optimization algorithm with performance analysis for a special case with a complete binary data gathering tree. Then, by discretizing the transmission time, we present a simple, distributed on-line protocol that relies only on the local information available at each sensor node. Extensive simulations were conducted for both long and short-range communication scenarios using two different source placement models. We used the baseline of transmitting all packets at the highest speed and shutting down the radios afterwards. Our simulation results show that compared with this baseline, up to 90% energy savings can be achieved by our techniques (both off-line and on-line), under different settings of several key system parameters.**

***Index Terms*— Energy efficiency, rate adaptation, real-time data gathering, wireless sensor networks.**

I. INTRODUCTION

ENERGY-EFFICIENCY is a key concern in wireless sensor networks (WSNs) [2]. One useful mechanism for energy-efficient communication is to explore the energy-latency tradeoffs by adjusting the transmission time [3]. An important observation is that in many channel coding schemes, the transmission energy can be significantly reduced by lowering transmission power and increasing the duration of transmission [3]. Rate adaptation techniques (e.g., modulation scaling [4]) have been proposed for implementing such tradeoffs.

In this paper, we exploit the energy-latency tradeoffs in the context of data gathering in WSNs. Typical communication patterns in data gathering involve multiple data sources and one data sink, forming a reverse-multicast structure, called the *data gathering tree* [5]. Data aggregation along such a tree [5] is particularly useful in eliminating data redundancy and reducing the communication load. We consider a real-time mission-critical scenario where the raw data gathered from the source nodes must be aggregated and transmitted to the sink within a specified latency constraint. Our objective is to minimize the overall energy cost of the sensor nodes in the data gathering tree subject to the latency constraint. Although our problem is formulated as a convex programming problem which is solvable in polynomial time by using general optimization tools, we propose more a time-efficient algorithm in this paper by exploiting special properties of the problem. Such properties include the convexity of the energy function of wireless communication and the tree structure of the underlying communication substrate.

It is important to evaluate the usefulness of the latency-energy tradeoffs by examining the sources of latency and energy costs for data gathering in WSNs. We assume a low-duty cycle WSN with sleep scheduling so that nodes are completely shut down in idle state. In such a system, besides the time cost for packet transmission, the latency for data gathering can be further decomposed into queuing delay, channel access delay, re-transmission delay, and the delay for waking up sleeping nodes. We argue that under several reasonable assumptions, the packet transmission delay is significant and worth trading for energy. In the following paragraph, we explain our argument in detail.

First, due to the application-specific design of WSNs, most traffic throughout the network is due to transporting the gathered data to the base station. It is also anticipated that many applications for WSNs require transmission of tens to hundreds of bytes per second [6]. In such a light-traffic scenario, queuing delay is not as a major concern as it is in traditional wireless ad hoc networks. Second, we assume the availability of a collision-free medium access control (MAC) protocol (e.g., using multi-packet reception (MPR) techniques [7], [8]), so that channel access delay due to collision detection and avoidance is negligible. However, our techniques are not directly applicable to TDMA-based protocols, due to the extra waiting time for transmission slots, or to contention-based MAC protocols, due to the latency caused by packet collisions. Third, the number of expected re-transmissions is actually a function of the Bit Error Rate at the receiving node, which in turn determines the energy-latency tradeoffs for packet transmission (see Section III-C for details). Therefore, it is convenient to explicitly incorporate the tradeoffs between expected number of re-transmissions and energy into our work. Fourth, we assume that a full duty cycle, ultra-low power wakeup radio [9] is available for each sensor node. Thus, sleeping sensor nodes can be

Manuscript received November 11, 2004; revised July 31, 2005; accepted February 13, 2006. The associate editor coordinating the review of this paper and approving it for publication was W. Liao. This work was supported in part by the NFS grants 0330445, 0325875, and 0347621. A preliminary version of this paper appears in IEEE InfoCom 2004 [1]. This version contains additional materials, including modified data aggregation model, performance analysis for a special case, improved on-line protocol, and more comprehensive simulation results.

Y. Yu is with the Application Research Center, Motorola Labs, Schaumburg, IL 60008 USA (e-mail: yang@motorola.com). The presented work was performed when he was at the University of Southern California.

V. K. Prasana and B. Krishnamachari are with the Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089-2562 USA (e-mail: {prasanna, bkrishna}@usc.edu).

Digital Object Identifier 10.1109/TWC.2006.xxxxx

woken up for packet transmission with almost no delay and energy penalties. Also, the typical startup time for sensor nodes is around 100 μSec [6], while the time for transmitting a packet of 200 bytes using 1 Mbps is 200 μSec. Based on the above observations, the time for packet transmission in light-traffic WSN applications constitutes a significant portion of the overall delay.

Since we assume that sensor nodes are completely shut down in idle state, the main source of energy cost is due to packet transmission for data gathering. It is therefore crucial to explore the energy-latency tradeoffs of packet transmission in such a context.

Technical Approach Overview: We first present an off-line numerical optimization algorithm, where the structure of the data gathering tree and the energy characteristics of all sensor nodes are known *a priori*. We also analyze the performance of our algorithm for a special case over a complete binary data gathering tree.

We then approximate the transmission time using a set of discrete values and describe a simple, localized on-line protocol. The key idea is to iteratively identify the sensor node with the highest energy gradient (to be defined later) in the tree and reduce its energy cost when allowed by the latency constraint. In this protocol, each sensor node only needs to perform simple operation based on its local information and the piggybacked information from data messages.

Finally, we evaluate the performance of our techniques through extensive simulations. The simulations were conducted for both long and short-range communications. We used the baseline where all sensor nodes transmit the packets at the highest speed (8 bits/symbol) and shut down the radio afterwards. Our simulation results from the scenarios we studied show that compared with this baseline, up to 90% energy savings could be achieved by our off-line algorithm and the on-line protocol. We also investigate the impact of several key network and radio parameters.

Paper Organization: We briefly discuss the related work in Section II. We describe our underlying network model and the energy-latency tradeoffs for wireless communication in Section III. The packet transmission problem is then defined in Section IV. The off-line algorithm and on-line protocol for the problem are presented in Sections V and VI, respectively. Simulation results are shown in Section VII. Finally, concluding remarks are made in Section VIII.

II. Related Work

The construction of the data gathering tree has been studied under various circumstances. For example, several localized tree topology generation mechanisms are compared by Zhou *et. al.* using metrics including node degree, robustness, and latency [10]. When the joint entropy of multiple information sources is modeled as a concave function of the number of sources, a randomized logarithmic approximation algorithm is developed by Goel *et. al.* [11]. By considering a simplified compression model, where the entropy conditioning at nodes only depends on the availability of side information, Cristescu *et. al.* proposed a 2-approximation performance for minimizing the overall cost of data gathering [12]. A nice analysis of the impact of spatial correlation on several practical schemes for tree construction [13] indicates that a simple cluster-based routing scheme performs well regardless the correlation among sources. All these works provide the underlying communication substrate above which our algorithm and protocol can be applied for further energy reduction.

From wireless communication perspective, rate adaptation has been widely studied to optimize spectral efficiency (e.g., network throughput) subject to the channel conditions in cellular networks or local-area wireless networks (e.g., [14], [15]). Several recent works [3], [4], [16], [17], which are closely related to our paper, have studied the application of rate adaptation for energy conservation. For a single-hop link, the problem of minimizing the energy cost of transmitting a set of packets subject to a specified latency constraint was studied by Prabhakar *et. al.* [3]. An extension of the problem [16] investigates the packet transmission from one single transmitter to multiple receivers. The concept of modulation scaling was first proposed by Schurgers *et. al.* [4]. For a single-hop link, policies for adjusting the modulation level are developed for cases where no real-time requirements are imposed [4]. Modulation scaling is also used for balancing the energy cost of a linear path [17].

The real-time latency constraint considered in this paper requires the use of global time-synchronization [18]. Our real-time scenario is similar to the epoch-based data gathering scheme [19]. However, prior work has not considered the possibility of using packet-scheduling techniques that trade latency for energy in such a scenario.

To the best of our knowledge, our paper is the first work that addresses packet scheduling in a general tree structure while considering a real-time latency constraint. The challenges of our problem are multi-fold. First, the energy functions can vary for different links. Thus, general optimization techniques instead of explicit solutions are required. Second, the latency constraint for data gathering in real applications is typically given by considering the data gathering tree as a whole. It is difficult to directly apply the techniques in [3], [16], as they require a latency constraint over each link. Finally, compared with most of the previous work, we use a more general and accurate energy model for wireless packet transmission in WSNs. Specifically, the transmission energy does not monotonically decrease as the transmission time increases – the transmission energy may increase when the transmission time exceeds some threshold value [2]. We refer to this general model that we will use as the *non-monotonic energy model*.

III. Models and Assumptions

A. Data Gathering Tree

Let $T = (V, E)$ denote the data gathering tree, where V denotes the set of n sensor nodes, $\{V_i : i = 1, \ldots, n\}$, and E denotes the set of directed communication links between the sensor nodes. Let M denote the number of leaf nodes in the tree. Without loss of generality, we assume that the sensor nodes are indexed in the topological order with $V_1, \ldots, V_M$ denoting the M leaf nodes and V_n denoting the sink node. Every link in E is represented as a (source, destination) pair (i, j), implying that V_j is the parent of V_i.

Let T_i denote the subtree rooted at any node, V_i, with $T_n = T$. A path in T is defined as a series of alternate nodes and edges from any leaf node, $V_i, i \in \{1, \ldots, M\}$, to V_n, denoted as p_i. We use the notation $V_j \in p_i$ to signify that node V_j is an intermediate node of path p_i.

Raw data is generated by a set of source nodes from V. Data aggregation is performed by all non-sink and non-leaf nodes (referred to as *internal nodes*). We assume that aggregation at an internal node is performed only after all input information is available at the node – either received from its children, or generated by local sensing for a source node. The aggregated data is then transmitted to the parent node. Let s_i denote the size of the packet transmitted by V_i to its parent. We discuss the computation of data aggregation to determine s_i in the next section.

For ease of analysis, it is assumed that raw data is available at all source nodes at time 0. Let Γ denote the latency constraint, within which data from all source nodes needs to be aggregated and transmitted to the sink node.

We assume a simplified communication model with a collision-free medium access control (MAC) layer that ensures no collision or interference at a node. This MAC layer can be realized by multi-packet reception (MPR) techniques through frequency, code, or spatial diversity [7], [8]. We also assume that sensor nodes are completely shut down in idle state and can be woken up for packet transmission using ultra-low power wakeup radios with almost no delay and energy penalties. A promising technique for such a wakeup radio has been discussed by Zhao *et al.* [9]. Also, sensing and computation cost for data aggregation are considered to be negligible.

B. Data Aggregation Paradigm

While there may be transients during the creation phase of a data gathering tree, we assume that this tree, once formed, lasts for a reasonable period of time and provides the routing substrate over which aggregation can take place during data gathering. Various techniques have been previously proposed for computing aggregates, or joint information entropy, from multiple source nodes. In our study, we adopt the model proposed by Pattem *et. al.* [13] where the joint entropy (or total compressed information) from multiple information sources is modeled as a function of the inter-source distance d, and a pre-specified correlation parameter c, that characterizes the extent of spatial correlation between data. Specifically, let H_1 denote the data size generated from any single source. The compressed information of two sources is calculated as [13]:

$$H_2 = H_1 + \frac{d}{d+c} H_1 \ . \tag{1}$$

We assume that the correlation parameter c is the same for any set of sources. Pattem's study shows that the above spatial correlation model closely matches the daily rainfall precipitation for the pacific northwest region over a period of 46 years [13]. Moreover, based on (1), a recursive calculation of the total compressed information of multiple sources is developed [13]. Due to space limitation, we omit the details here.

Although we use the expression in (1) as a typical aggregation function, please note that our technique is not limited to this function alone. The only requirement is that we can derive the value of s_i's based on the functions. Thus, different functions can be used to specify the aggregation at different sensor nodes.

C. Energy-Latency Tradeoffs in Wireless Communication

We model the transmission energy using the example of modulation scaling based on the Quadrature Amplitude Modulation (QAM) scheme [4]. Note that the techniques presented in this paper are extendible to other modulation schemes as well as other techniques that provide energy-latency tradeoffs, such as code scaling [14]. Consider a packet of s bits to be transmitted between two sensor nodes. Assuming that the symbol rate, R, is fixed, the transmission time, τ, can be calculated as [4] $\tau = \frac{s}{b \cdot R}$, where b is the modulation level of the sender in terms of the constellation size (number of bits per symbol). The corresponding transmission energy can be modeled as the sum of output energy and electronics energy. To illustrate the key energy-latency tradeoffs, we abstract the energy cost as a function of τ [4], denoted as $w(\tau)$:

$$w(\tau) = [C \cdot (2^{\frac{s}{\tau \cdot R}} - 1) + F] \cdot \tau \cdot R \ , \tag{2}$$

where C is determined by the quality of transmission (in terms of Bit Error Rate) and the noise power, and F is a device-dependent parameter that determines the power consumption of the electronic circuitry of the sender. Further, the output power, P_o, and the electronic power, P_e, can be calculated as [4] $P_o = C \cdot R \cdot (2^b - 1)$ and $P_e = F \cdot R$.

Note that different assumptions about the radio characteristics, including power consumption and data rate, may significantly affect the analysis of various energy-saving mechanisms. In this work, we consider the radio modules described in [2], [6]. Typically, for *short-range* communication with $R = 1$ Mbaud, we have $P_e \approx 10$ mW and $P_o \approx 1$ mW at 4-QAM (i.e., 2 Mbps). Although the above numbers are better than currently available radios for commercial sensor nodes (e.g., Berkeley motes typically support data rate up to 100 Kbps with higher power assumption), radio devices with the above specifications are anticipated in the near future.

From the equations for P_o and P_e, we derive $C \approx 3 \times 10^{-10}$ and $F = 10^{-8}$. We consider a d^2 path loss model for signal propagation, where d is the communication radius. Assuming that it takes 10 pJ/bit/m^2 by the amplifier to transmit one bit at an acceptable quality [20], we infer that $d = \sqrt{50} \approx 7$ m (from $\frac{1 \text{ mW}}{2 \times 10^6 \text{ bit/sec}} = 10$ pJ/bit/m$^2 \times d^2$). In our study, we also consider one more case of communication in WSNs with longer radius. Specifically, we set the communication range to 30 m, implying $P_o = 10$ pJ/bit/m$^2 \times 30^2$ m$^2 \times 2 \times 10^6$ bit/sec $= 18$ mW at 4-QAM, and consequently $C = 6 \times 10^{-9}$. We refer to this communication scenario as *long-range* communication. Note that these numbers for communication radii are for illustrative purpose only – to show the different weights of C against F with respect to variations in communication radius. They may vary according to different radio devices and operating environments.

Fig. 1 plots the energy functions with $b \in [2, 8]$ for the long and short range communication scenarios. In practice, b is

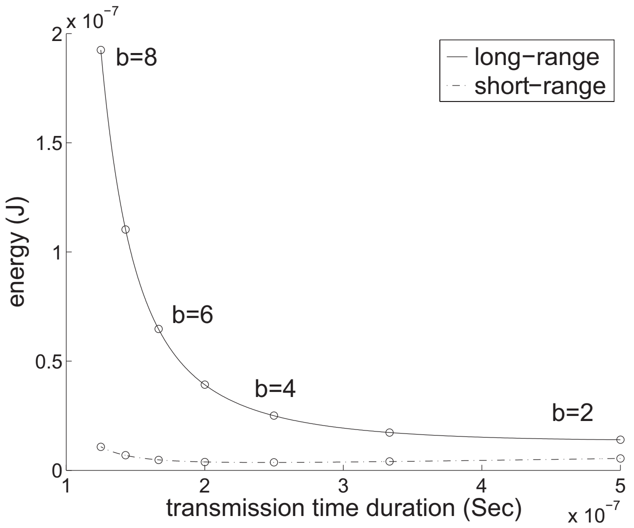

Fig. 1. Energy-latency tradeoffs for transmitting one bit.

typically set to positive even integers, as indicated by circles in the figure. We can observe a 10-fold energy reduction for long-range communication by varying b from 8 to 2, and a 3-fold energy reduction for short-range communication by varying b from 8 to 4. Intuitively, it is more beneficial to explore the energy-latency tradeoffs for long-range communication. However, we demonstrate in Section VII that more than 50% energy savings can still be achieved by our techniques for short-range communication.

For simplicity, we do not consider amplifier inefficiency in our work, which typically increases the actual power consumption of the amplifier by a factor of 5 [21]. When amplifier inefficiency is considered, the output power will dominate more over the electronics power. Thus, we can expect that our techniques will be more useful for reducing the total energy cost for packet transmission.

IV. Problem Definition

A schedule of packet transmission is defined as a vector $\vec{\tau} = \{\tau_i : i = 1, \ldots, n-1\}$, where τ_i is the time duration for packet transmission over link (i, j). Since a sensor node can transmit its packet only after receiving all input packets from its children, the start time of each transmission is implicitly determined by $\vec{\tau}$. The transmission latency of a path, p_i, is denoted as L_i and calculated as $L_i = \sum_{j:V_j \in p_i} \tau_j$. A schedule is feasible if for any $p_i \in T$, we have $L_i \leq \Gamma$.

Let $w_i(\tau_i)$ denote the energy function of V_i, with m_i denoting the value of $\tau_i \in (0, \Gamma]$ when $w_i(\cdot)$ is minimized. Note that $w_i(\cdot)$ may vary for different nodes due to variations in packet size and transmission radius (in other words, such information is implicitly embedded into $w_i(\cdot)$). We now formally state the *packet transmission problem* (PTP) as follows:

Given:
a. a data gathering tree T consisting of n sensor nodes,
b. energy functions for each link $(i, j) \in E$, $w_i(\tau_i)$, and
c. the latency constraint, Γ;
find *a schedule of packet transmission, $\vec{\tau} = \{\tau_i : i = 1, \ldots, n-1\}$, so as to minimize*

$$f(\vec{\tau}) = \sum_{i=1}^{n-1} w_i(\tau_i) \tag{3}$$

subject to

$$\forall p_i \text{ in } T, L_i = \sum_{j:V_j \in p_i} \tau_j \leq \Gamma \, . \tag{4}$$

We will present both an off-line algorithm and an on-line distributed protocol for PTP. In the off-line case, we assume that the structure of the data gathering tree and the energy functions for all sensor nodes are known a priori; whereas in the on-line case, each sensor node only has local knowledge about its own radio status and can communication with its parent and children.

V. A Numerical Optimization Algorithm for Off-Line PTP

In this section, we first describe an extension of the MoveRight algorithm [16] to get optimal solutions for the off-line version of PTP. We then analyze the performance of the proposed technique for a special case over a complete binary data gathering tree.

A. A Numerical Optimization Algorithm

Since we must have $\tau_i \leq m_i$ in an optimal solution to PTP, the latency of a path does not necessarily equal Γ. Moreover, let V_i denote an internal node. We show the following necessary and sufficient condition for the optimality of the off-line PTP problem.

Lemma 1: A schedule, $\vec{\tau}^*$, is optimal for off-line PTP iff

1) for any node V_i with $\tau_i^* < m_i$, the length of at least one path that contains V_i equals Γ; and
2) for any internal node, V_i, we have

$$\dot{w}_i(\tau_i^*) = \sum_{(j,i) \in E} \dot{w}_j(\tau_j^*) \, , \tag{5}$$

where $\dot{w}_i(\cdot)$ is the first derivative of $w_i(\cdot)$.

We now extend the MoveRight algorithm from [16] to solve off-line PTP in a general-structured tree with non-monotonic energy functions. The pseudo code for the extended MoveRight algorithm (EMR-Algo) is shown in Fig. 2, which is found at the top of the next page. In the figure, τ_i^k denotes the value of τ_i in the k-th iteration. Initially, we set the starting time for all packet transmission to zero – the transmission time for all the links to the sink (i, n) is set to $\min\{\Gamma, m_i\}$, while the transmission time for the rest of the links is set to 0 (Steps 2 and 3). The main idea is to iteratively increase (move right) the starting times of packet transmissions, so that each move locally optimizes the overall energy function. Finally, this iterative local optimization leads to a globally optimal solution.

The best($\cdot$) function returns the transmission time for node V_i and its children so that Lemma 1 holds for the subtree formed by V_i and its parent and children, with respect to the invariant that $\tau_j^k \leq m_j$ for any node V_j in the subtree. When the best($\cdot$) function is called upon the subtree around V_i, the

Begin

1. Set $k \leftarrow 0$ // initialize iteration counter
2. **For** $(i,n) \in E$, set $\tau_i^k \leftarrow \min\{\Gamma, m_i\}$ // initialize transmission time for links to the sink
3. **For** $(i,j) \in E$ such that $j \neq n$, set $\tau_i^k \leftarrow 0$ // initialize transmission time for other links
4. Set $flag \leftarrow 0$ // flag to keep track of convergence in the iterations
5. **While** $flag = 0$
6. $\quad k \leftarrow k+1$ // increment the iteration counter by 1
7. $\quad$ **For** each V_i with i from $n-1$ downto M+1 // perform local optimization for each internal node
8. $\qquad (\{\tau_j^k\}_{(j,i)\in E}, \tau_i^k) \leftarrow$ best$(\{\tau_j^{k-1}\}, \tau_i^{k-1})$ // move right the start time of transmission from V_i
9. $\quad$ **For** $(i,n) \in E$
10. $\qquad$ Set $\tau_i^k \leftarrow \min\{m_i, \Gamma - (\max_{V_i \in p_j}\{L_j\} - \tau_i^k)\}$ //increase the transmission time for links to the sink
11. $\quad$ if $\vec{\tau}^k = \vec{\tau}^{k-1}$, $flag \leftarrow 1$ // check convergence

End

Fig. 2. Pseudo code for EMR-Algo.

transmission for all the links not within the subtree remain fixed, i.e., the starting time of transmissions from the children of V_i and the ending time of the transmission from V_i are fixed. We can prove that the starting time of the transmission from V_i will never be decreased by calling the best$(\cdot)$ function (Theorem 1 in [22]). Thus, in the best$(\cdot)$ function, the locally optimal starting time of the transmission from V_i is obtained by a binary search between the original starting time and the ending time of the transmission. Step 10 is important as it moves right the complete time of transmissions on links to the sink. This movement stops when the latency constraint is reached.

The correctness of EMR-Algo can be proved by exploring the convexity property of the energy functions. Due to space limitation, we omit the proofs of the correctness of EMR-Algo and Lemma 1. Complete proofs can be found in [22].

B. Performance Analysis for a Special Case

We consider a special case where the data gathering tree is a complete binary tree with all the leaf nodes being the source nodes. Let d denote the depth of the tree, with 2^i nodes at depth $i \leq d$. Let s denote the size of data packet from all sources. We assume a perfect data aggregation at all nodes such that the packets on all edges are of size s (i.e., $c = \infty$ in (1)).

For simplicity, we assume that the communication environment and radio device for all nodes are identical in terms of parameters R and C. We assume a long-range communication scenario where F is negligible. Due to the structure of a complete binary tree, it can be inferred that nodes at the same depth shall have the same transmission duration in the optimal solution. Let τ_i denote the transmission time of nodes at depth i, where $1 \leq i \leq d$.

Based on (2), we can state the PTP problem in this special case, denoted as PTP-SP to find a schedule $\{\tau_i : i = 1, \ldots, d\}$ to:

$$\min \quad \sum_{i=1}^{d} 2^i [C(2^{\frac{s}{\tau_i R}} - 1) \cdot \tau_i \cdot R] \tag{6}$$

$$\text{subject to} \quad \sum_{i=1}^{d} \tau_i \leq \Gamma \, . \tag{7}$$

When $\Gamma \to \infty$, a lower bound on the cost of PTP-SP equals $(2^{d+1} - 2)sC \ln 2$. Thus, we set Γ in a more interesting range: $[\frac{ds}{8R}, \frac{ds}{2R}]$. The boundaries of the range are obtained by setting modulation level of all sensor nodes to 8 and 2, respectively. We also consider the baseline where all sensor nodes transmit with a modulation level of 8 and shut down afterwards, i.e., $\tau_i = \frac{s}{8R}$ for all $i \leq d$. The cost of such a baseline is $(2^{d+1} - 2)\frac{255sC}{8}$, which is an upper bound on the cost of PTP-SP.

Based on Lemma 1, it can be shown that the optimal schedule to PTP-SP shall satisfy

$$2^i \cdot C \cdot R(2^{\frac{s}{\tau_i R}} - 1 - 2^{\frac{s}{\tau_i R}} \cdot \ln 2 \cdot \frac{s}{\tau_i R}) = \lambda \tag{8}$$

where λ is Lagrange multiplier determined by the constraint $\sum_{i=1}^{d} \tau_i = \Gamma$. Let $b_i = \frac{s}{\tau_i R}$ be the modulation level for nodes at depth i. Since $b_i \geq 2$, we can approximate (8) as $2^{b_i} = \frac{\lambda}{2^i CR}$. This approximation is similar to the inverse-log scheduling in [23], which is shown in [23] to be a close approximation for (8) when b_i is large. This approximation leads to $\tau_i = \frac{s}{R \ln \frac{\lambda}{2^i CR}}$, where λ is determined by $\sum_{i=1}^{d} \tau_i = \Gamma$.

To solve λ, let $x = \ln \frac{\lambda}{CR}$ and $\kappa = \frac{\Gamma R}{s}$. The constraint $\sum_{i=1}^{d} \tau_i = \Gamma$ can be written as

$$\sum_{i=1}^{d} \frac{1}{x - i} = \kappa \, , \tag{9}$$

which is essentially a polynomial equation in x of degree d. Since $\tau_i > 0$ for all $i \leq d$, we consider the only root that satisfies $x > d$. Let $Harm(i) = 1 + \frac{1}{2} + \ldots + \frac{1}{i}$ denote the Harmonic function. We have $\sum_{i=1}^{d} \frac{1}{x-i} \approx Harm(x-1) - Harm(x-d-1)$. We use the approximation $Harm(i) = \ln(i + \frac{1}{2}) + \gamma$ for large i, where γ is the Euler-Mascheroni constant. Thus, we derive

$$x \approx \frac{(2d+1)e^{\kappa} - 1}{2(e^{\kappa} - 1)} \, , \tag{10}$$

where e is the natural number.

Thus, when d is large, we have an approximated optimal schedule:

$$\tau_i = \frac{s}{R(x - i)} \tag{11}$$

for $i = 1, \ldots, d$, with an energy cost that equals

$$\begin{aligned} C_{appr} &= \sum_{i=1}^{d} 2^i C(2^{\frac{s}{\tau_i R}} - 1)\tau_i R \\ &= sC \sum_{i=1}^{d} \frac{2^i(2^{x-i} - 1)}{x - i} \\ &\approx sC2^x \sum_{i=1}^{d} \frac{1}{x-i} \qquad \text{(when } d \text{ is large)} \\ &= sC2^x \kappa \qquad \text{(from (9))}\,, \qquad (12) \end{aligned}$$

where $\kappa = \frac{\Gamma R}{s}$ and x is given by (10). This gives an improvement over the baseline by a factor around $\frac{2^{d+6-x}}{\kappa}$.

In Fig. 3(a), with $C = 6 \times 10^{-9}$, $R = 10^6$, $s = 200$, and $d = 6$, we plot C_{appr} and the cost obtained by EMR-Algo as a function of Γ as well as the lower and upper bounds on the cost of PTP-SP. We observe that when Γ is small, C_{appr} and the cost of EMR-Algo are very close. When Γ is large, although there is noticeable difference between C_{appr} and the cost of EMR-Algo, the ratio of their improvement over the upper bound is actually very close. We also notice that when Γ is set to minimal, there is still improvement by C_{appr} and EMR-Algo over the upper bound. This is because for C_{appr} and EMR-Algo, we have relaxed the constraint that b_i must be in $[2, 8]$. We will show in Section VII-C that for our on-line protocol which considers such a constraint, the resulting energy savings is still comparable to that of EMR-Algo.

We also plot the energy conservation of E_{appr}, which is defined as the percentage of energy savings by E_{appr} over the upper bound. We observe an energy conservation from 30% to 90% for E_{appr}. Although these numbers are based on the above special case, they confirm our simulation results for general trees in Section VII very well. Thus, our analysis on this special case gives meaningful insight on the energy conservation that can be achieved by our technique in general scenarios.

In Fig. 3(b), we also plot τ_i's with respect to variations in Γ given by (11). It can be observed that when Γ increases, transmission time of nodes with larger depth increases faster than that of nodes with smaller depth does. This is because the number of nodes increases exponentially with depth. Thus, more transmission time is desired for nodes with large depth to sustain Lemma 1.

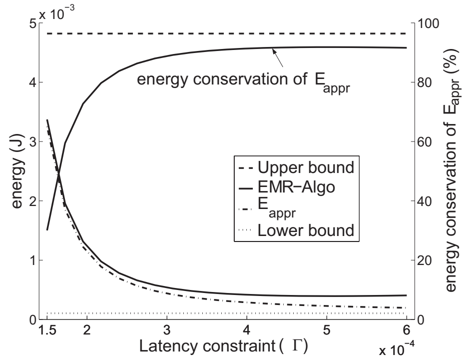

(a) Comparison between E_{appr} and the cost of EMR-Algo.

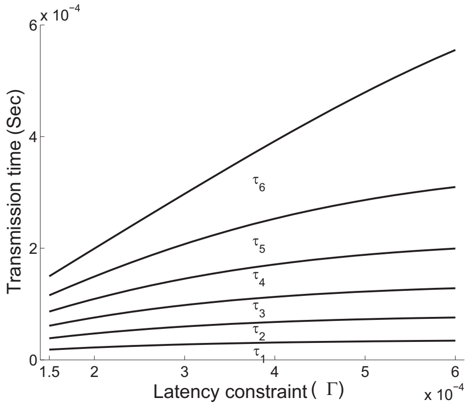

(b) Transmission time of sensor nodes.

Fig. 3. Performance analysis for a special case over a complete binary data gathering tree.

VI. Distributed On-Line Protocol

We first discretize the transmission time of each sensor node using a list of increasing transmission time that can be chosen to transmit its packet. We define the *energy gradient* of a sensor node as the amount of energy that can be saved by increasing its transmission time to the next level. The key idea of the protocol is to iteratively identify the sensor nodes with the largest positive energy gradient and increase their transmission time if the latency constraint allows.

To facilitate the on-line scheduling, we further make the following assumptions:

1) Some local unique neighbor identification mechanisms are available at each sensor node for identifying the parent and children.
2) Every node V_i can derive the time cost for data gathering within the subtree rooted V_i.
3) Every sensor node can measure its current energy gradient.
4) Interference among sensor nodes is handled by MPR techniques.

The local identifier in assumption 1 is commonly implemented in protocols such as Directed Diffusion [24]. Assumption 2 can be fulfilled by attaching a time stamp to each packet from the leaf nodes (we shall be assuming that time synchronization schemes, such as [18], are available). In assumption 3, the energy gradient of a sensor node can be determined using the system parameters provided by the hardware vendors and the operating configuration of the system, such as the bit rate. Assumption 4 can be satisfied by using for example, CDMA technique.

Moreover, we define the *latency laxity* of a node as the maximal amount of time that can be used to increase the transmission time of the node without violating the latency constraint. Let x_i denote the latency laxity of V_i. The latency laxity of each node is dynamically maintained during the protocol to verify if the transmission time of the node can be safely increased.

In the following, we first describe the local data structure maintained at each sensor node. A distributed adaptation policy for minimizing the energy cost is then presented.

Local Data Structure: Each sensor node, V_i, maintains a simple local data structure (r, τ_i, τ_d). The flag r equals one if V_i is the node with the highest *positive* energy gradient in subtree T_i, and zero otherwise. Field τ_i is the time cost for transmitting the packet from V_i to its parent, while τ_d records the time cost of the longest path, *excluding* τ_i, in T_i.

The local data structure is maintained as follows. Every leaf node piggybacks its energy gradient to the outgoing packet. Once a sensor node, V_i, receives packets from all its children, the node compares the energy gradients piggybacked to each packet and the energy gradient of its own. The value of r at V_i is then set accordingly. If V_i is not the sink, the largest energy gradient from the above comparison is piggybacked to the packet sent to the parent of V_i. The above procedure continues until all the sensor nodes have the correct value of r. Fields τ_i and τ_d can be easily maintained based on the above assumptions.

Adaptation Policy: The sink node periodically disseminates a feedback packet to its children that contains the value of its local τ_d and the difference between Γ and τ_d, denoted as δ. Basically, δ is the latency laxity of nodes on the longest path of the data gathering tree.

Once a sensor node V_i receives the feedback packet from its parent, it performs the following adaptation. To distinguish from the field τ_d in V_i's local data, let τ_d' denote the field τ_d in the feedback packet. First, the latency laxity of V_i can be calculated as $x_i = \delta + \tau_d' - (\tau_i + \tau_d)$. This is because $\tau_i + \tau_d$ is the time cost of T_i; τ_d' is the time cost of the longest path in the subtree rooted at V_i's parent (excluding the transmission time of V_i's parent); and δ is actually the latency laxity of nodes on this longest path. Then, V_i takes one of the following actions.

1) If $\delta < 0$, the transmission time for packet from V_i is decreased by a factor of β, where β is a user-specified parameter. The feedback packet is then forwarded to all of V_i's children.
2) If $r = 1$ and x_i is large enough to accommodate the increase of V_i's transmission time to the next level, increase V_i's transmission time to the next level. The local data structure at V_i is updated accordingly; and the feedback packet is suppressed.
3) Otherwise, the feedback packet is updated by setting $\delta = x_i$ and $\tau_d' = \tau_d$. The updated packet is then forwarded to all children of V_i.

The rationale behind the above adaptation policy is that when the latency constraint is violated, all the sensor nodes send out packets with an increased rate (action 1). If V_i is the node with the largest positive energy gradient in T_i and the latency laxity allows, the transmission time of V_i is increased (action 2). Otherwise, the latency laxity of V_i is recorded in the feedback packet and the sensor nodes in T_i are recursively examined (action 3).

During each dissemination of the feedback packet, the proposed on-line protocol increases the transmission time for at most one sensor node per path. Such an increment is guaranteed not to violate the latency constraint. Therefore, the on-line protocol converges after the latency constraint is reached by all paths, or $\tau_i = m_i$, for each $V_i \in V$. We assume that each sensor node has q discretized transmission times. Before the protocol converges, a feedback packet would increase the transmission time for at least one sensor node when it traverses the data gathering tree. Thus, the protocol converges after the dissemination of at most nq feedback packets, where n is the number of sensor nodes in the tree. Moreover, in our protocol, the information needed for maintaining local data structures is piggybacked on existing data packets. The feedback packet from the sink contains only two fields that need no more than 4 bytes each. Thus, the overhead of the protocol is relatively low compared with the cost for data packets.

The above protocol does not require the discretized transmission time to be evenly distributed. For example, the set of transmission times in modulation scaling is generated by $\frac{s}{b \cdot R}$ for $b = 2, 4, \ldots$. The distance between adjacent transmission times actually decreases with b.

VII. Simulation Results

A. Simulation Setup

To conduct the simulations, a simulator was developed using the PARSEC [25] software, which is a discrete-event simulation language. We used the specific time and energy models described in Section III-C. For the off-line EMR-Algo, we assumed that the modulation level of all sensor nodes was continuously adjustable in $(0, \infty)$. This gave an upper bound for the performance of the modulation scaling used in our on-line simulations, where the modulation levels were even integers in $[2, 8]$, indicating a corresponding data rate in $\{2, 4, 6, 8\}$ Mbps. For both off-line and on-line cases, the baseline is to transmit all packets at 8 Mbps and shutdown the radios afterwards. This policy was used, for example, in the PAMAS [26] and DMAC protocols [27]. The performance metric was defined as the percentage of energy savings achieved by using our techniques, compared with the baseline.

A sensor network was generated by randomly scattering 200 sensors in a unit square. The sink node was put at the left-bottom corner of the square. The neighbors that a sensor node could directly communicate was determined by a connectivity parameter, $\rho \in (0, 1]$. Specifically, two sensor nodes could communicate with each other only if the distance between them was within ρ. Note that ρ is a purely relative measurement in the unit square. Consider the example when $\rho = 0.1$. In the case of short-range communication, ρ is translated to 7 m and the square scales to 70 m $\times$ 70 m. In the case of long-range communication, ρ is translated to 30 m and the square scales to 300 m $\times$ 300 m. The size of raw data from all source nodes was set to 200 bits.

We used two models for generating the location of the source nodes, namely the random sources (RS) model and the event radius (ER) model. In the RS model, N (the number of sources) out of 200 sensor nodes were randomly selected to be the sources, whereas in the ER model, all sources were located within a distance S (essentially the sensing range) of a randomly chosen "event" location. For both models, the Greedy Incremental Tree (GIT) algorithm [5] was used for constructing the data gathering tree. In Fig. 4, we illustrate two example data gathering trees generated based on the RS and ER models, respectively.

The energy function used in the simulation was in the form of (2). Unless otherwise stated, we set $R = 10^6$ and $F = 10^{-8}$ for all the sensor nodes, while the value of C of a sensor

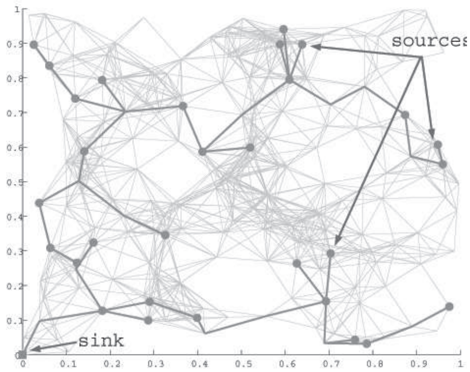

(a) Random sources model (number of sources $N = 30$).

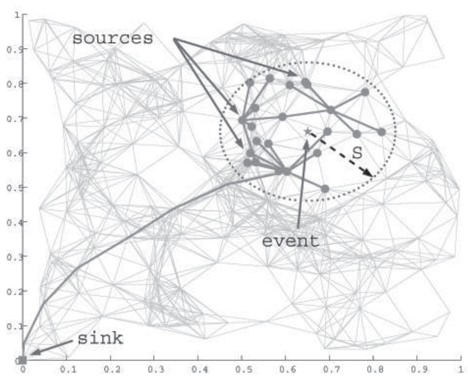

(b) Event radius model (sensing range $S = 0.2$).

Fig. 4. Two example data gathering trees generated by the random sources and event radius models, respectively (connectivity parameter $\rho = 0.15$).

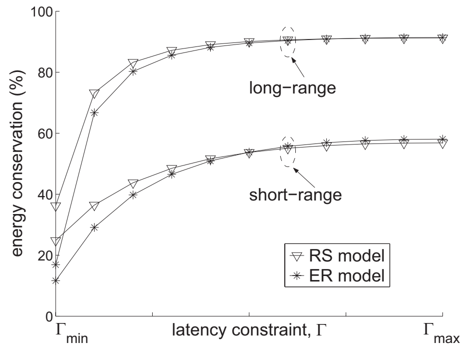

(a) Performance overview.

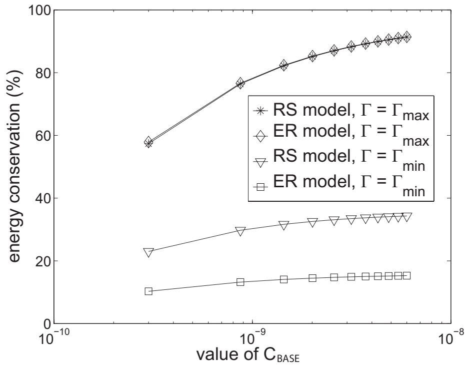

(b) Impact of C_{base}

Fig. 5. Performance of EMR-Algo (correlation parameter $c = 0.5$, connectivity parameter $\rho = 0.15$, number of sources $N = 30$, sensing range $S = 0.15$).

node was determined by the distance from the node to its parent in the tree. Specifically, we assumed a d^2 power loss model, where d was the distance between a node and its parent. Then, for node V_i, we have $C_i = C_{base} \cdot (\frac{d}{\rho})^2$. Based on our analysis in Section III-C, C_{base} was set to 6×10^{-9} for the long-range communication and 3×10^{-10} for the short-range communication.

During our simulation, the latency constraints Γ was determined as follows. We define the shortest time cost, Γ_{min} of a gathering tree as the transmission latency of the longest path in the tree when all sensor nodes transmit at 8 Mbps. On the other hand, the longest time cost, Γ_{max} of the gathering tree is defined as the transmission latency of the longest path in the tree when every sensor node V_i sends its packet using time $\min\{m_i, \frac{s_i}{2}\}$. In the above definition, the term m_i comes from the fact that it is not energy beneficial for V_i to transmit its packet using time beyond m_i, and the term $\frac{s_i}{2}$ is based on modulation level 2. Γ was then adjusted between Γ_{min} and Γ_{max}.

While the simulations in [1] focus on the performance of off-line algorithms, we believe that the validation of the performance of the on-line protocol is of more practical significance, which is the emphasis in our simulation. The presented data is averaged over more than 150 problem instances and has a 95% confidence interval with a 10% (or better) precision.

B. Performance of EMR-Algo

Performance Overview: In Fig. 5(a), we illustrate the performance of EMR-Algo in the RS and ER models when varying Γ from Γ_{min} to Γ_{max}. The first thing to notice is that when Γ approached Γ_{max}, EMR-Algo achieved more than 90% energy saving for the long-range communication, and around 50% for short-range. Even when $\Gamma = \Gamma_{min}$, EMR-Algo could still save more than 30% of the energy for the long-range communication and 20% for short-range. The numbers for long-range communication conform the analysis in Section V-B quite well.

The reason for successful energy saving even when $\Gamma = \Gamma_{min}$ is two-fold. First, modulation level in EMR-Algo is allowed to be varied between $(0, \infty)$ instead of $[2, 8]$. Second, Γ equals the transmission time of the longest path in the data gathering tree. Thus, energy can still be reduced for nodes not on the longest path. On one hand, when there exists only one path in the tree, few energy can be saved when $\Gamma = \Gamma_{min}$. On the other hand, when the tree forms a star-like structure, all links, except the longest ones, can be optimized for energy savings when $\Gamma = \Gamma_{min}$. This also explains the performance degradation of EMR-Algo in the ER model, compared with the performance in the RS model — the gathering tree for the ER model forms a small cluster connected to the sink by a linear array of sensor nodes, while the tree for the RS model is more close to a star-like structure (Fig. 4).

Impact of the radio parameters: In Fig. 5(b), we show the impact of radio parameter C_{base}. In the figure, the x-axis represents the value of C_{base} from 3×10^{-10} to 6×10^{-9} in logarithmic scale. As expected, the energy conservation increased with C_{base}.

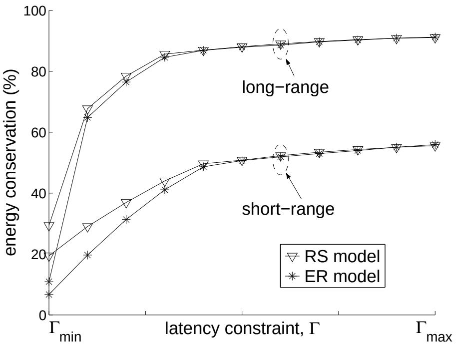

Fig. 6. Performance of the on-line protocol (correlation parameter $c = 0.5$, connectivity parameter $\rho = 0.15$, number of sources $N = 30$, sensing range $S = 0.15$).

To evaluate the impact of symbol rate R, we varied R from 10 KBaud to 1 MBaud. Considering the modulation level in $[2, 8]$, the above range of R reflectd a bit rate of 20 to 80 Kbps when $R = 10$ KBaud and 2 to 8 Mbps when $R = 1$ MBaud. Our simulation results show a constant amount of performance improvement by EMR-Algo throughout the variation of R. This is understandable, since from (2), the performance ratio of EMR-Algo to the baseline is determined by b, but not R.

We also investigated the impact of start-up energy of radios, which was estimated as 1 μJ [2]. Our simulation results show that while the impact of the start-up energy to the long-range communication was almost negligible, a decrease of 6-15% energy conservation was observed for the short-range communication.

C. Performance of the On-Line Protocol

We focus on the results for the RS model; similar analysis can be mode for the ER model.

Energy Conservation: We show the energy savings achieved by the on-line protocol in Fig. 6. When the latency constraint approaches Γ_{min}, there was slight performance degradation compared with EMR-Algo (Fig. 5(a)). Specifically, for the RS model, we observed around 4% less energy conservation for long-range communication and 3% for short-range communication. This was quite reasonable considering the fact that only 4 options were available to set the transmission time for each sensor in the on-line protocol, instead of the continuous adjustment of the transmission time in EMR-Algo.

Impact of Network Parameters: Fig. 7(a) shows the energy conservation achieved by our online protocol with respect to variations in c and Γ. It was observed that for a fixed Γ, the resulting energy gain slightly increased when c increased. This was because smaller value of c caused larger size data packets after aggregation. Thus, the energy cost of links close to the sink node dominated the overall energy cost of the tree. It was however difficult to reduce the energy cost of these links since they had a high likelihood of lying on the longest path of the tree.

Fig. 7(b) plots the performance of our protocol with respect to variations in N and Γ. It can be seen that when Γ was close to Γ_{min}, the energy gain of the protocol increased as the number of sources increased. This was because a larger number of sources offerd more opportunities for the optimization of links on paths other than the longest one. Fig. 7(c) demonstrates the performance of our protocol with respect to variations in ρ and Γ. It can be observed that the energy saving increased when ρ increased. This was understandable since a large ρ reduces the height of the data gathering tree (the extreme case is a star-like tree formed by setting $\rho = 1$).

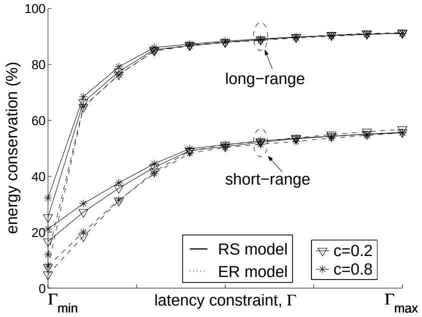

(a) Impact of the correlation parameter c ($\rho = 0.15$, $N = 30$, $S = 0.15$).

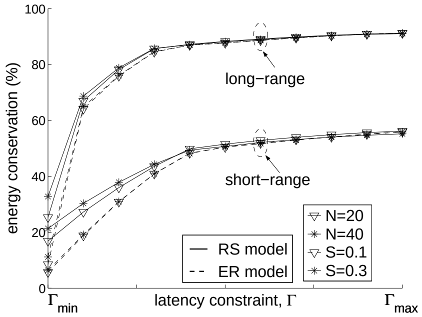

(b) Impact of the number of sources N or sensing range S ($c = 0.5$, $\rho = 0.15$,).

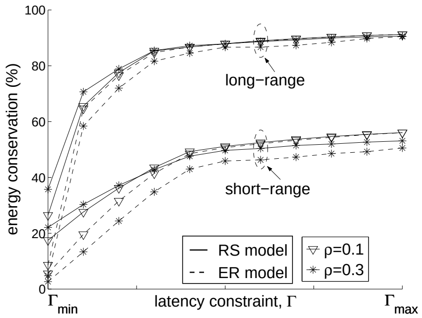

(c) Impact of the connectivity parameter ρ ($c = 0.5$, $N = 30$, $S = 0.15$).

Fig. 7. Impact of network parameters (c: correlation parameter, ρ: connectivity parameter, N: number of sources, S: sensing range).

We have also examined the running time behavior of our on-line protocol through specific scenarios. Despite of the nq feedback packets required for the worst case analysis in Section VI, our results indicate that the actual convergence time for the protocol was very short for the simulated scenarios. Due to space limitations, we omit the results here (please refer to [22] for details).

VIII. CONCLUDING REMARKS

In this paper, we have studied the problem of scheduling packet transmissions over a data gathering tree in wireless sensor networks by exploring the energy-latency tradeoffs. We have provided an off-line numerical optimization algorithm with performance analysis for a special case over a complete binary data gathering tree. We have also proposed a distributed on-line protocol that relies on local information only. Our simulation results show that up to 90% energy savings could be achieved by the off-line algorithm and the on-line protocol. We have investigated the performance of our algorithms with different settings of several key system parameters.

REFERENCES

[1] Y. Yu, B. Krishnamachari, and V. K. Prasanna, "Energy-latency tradeoffs for data gathering in wireless sensor networks," in *Proc. IEEE InfoCom*, Mar. 2004, vol. 1.

[2] V. Raghunathan, C. Schurgers, S. Park, and M. B. Srivastava, "Energy-aware wireless microsensor networks," *IEEE Signal Processing Mag.*, vol. 19, no. 2, pp. 40–50, Mar. 2002.

[3] B. Prabhakar, E. Uysal-Biyikoglu, and A. E. Gamal, "Energy-efficient transmission over a wireless link via lazy packet scheduling," in *Proc. IEEE InfoCom*, Apr. 2000, vol. 1, pp. 386-394.

[4] C. Schurgers, O. Aberhorne, and M. B. Srivastava, "Modulation scaling for energy-aware communication systems," in *IEEE International Symposium on Low Power Electronics and Design*, Aug. 2001, pp. 96–99.

[5] B. Krishnamachari, D. Estrin, and S. Wicker, "The impact of data aggregation in wireless sensor networks," in *Proc. International Workshop on Distributed Event-Based Systems*, July 2002, pp. 568-575.

[6] A. Wang, S.-H. Cho, C. G. Sodini, and A. P. Chandrakasan, "Energy-efficient modulation and MAC for asymmetric microsensor systems," in *IEEE International Symposium on Low Power Electronics and Design*, Aug. 2001, pp. 106-111.

[7] K. Sohrabi and G. J. Pottie, "Performance of a novel self-organization protocol for wireless ad hoc sensor networks," in *Proc. IEEE Vehicular Technology Conference*, Sep. 1999, vol. 2, pp. 1222-1226.

[8] L. Tong, Q. Zhao, and G. Mergen, "Multipacket reception in random access wireless networks: From singal processing to optimal medium access control," *IEEE Commun. Mag.*, vol. 39, no. 11, pp. 108–112, Nov. 2001.

[9] L. C. Zhong, R. Shan, C. Guo, and J. Rabaey, "An ultra-low power and distributed access protocol for broadband wireless sensor networks," in *Proc. IEEE Broadband Wireless Summit*, May 2001.

[10] C. Z. Zhou and B. Krishnamachari, "Localized topology generation mechanisms for self-configuring sensor networks," in *Proc. IEEE GlobeCom*, Dec. 2003, vol. 3, pp. 1269-1273.

[11] A. Goel and D. Estrin, "Simultaneous optimization for concave costs: Single sink aggregation or single source buy-at-bulk," in *Proc. ACM-SIAM Symposium on Discrete Algorithms*, Jan. 2003, pp. 499-505.

[12] R. Cristescu, B. Beferull-Lozano, and M. Vetterli, "On network correlated data gathering," in *Proc. IEEE InfoCom*, Mar. 2004, vol. 4, pp. 2571-2582.

[13] S. Pattem, B. Krishnamachari, and R. Govindan, "The impact of spatial correlation on routing with compression in wireless sensor networks," in *Proc. ACM/IEEE International Symposium on Information Processing in Sensor Networks*, Apr. 2004, pp. 28-35.

[14] E. Armanious, D. D. Falconer, and H. Yanikomeroglu, "Adaptive modulation, adaptive coding, and power control for fixed cellular broadband wireless systems," in *Proc. IEEE Wireless Communications and Networking Conference*, Mar. 2003, vol. 1, pp. 238-242.

[15] O. B. Akan and I. F. Akyildiz, "ARC: The analytical rate control scheme for real-time traffic in wireless networks," *IEEE/ACM Trans. Networking*, vol. 12, no. 4, pp. 634-644, June 2004.

[16] A. E. Gamal, C. Nair, B. Prabhakar, E. Uysal-Biyikoglu, and S. Zahedi, "Energy-efficient scheduling of packet transmissions over wireless networks," in *Proc. IEEE InfoCom*, June 2002, vol. 3, pp. 1773-1782.

[17] Y. Yu and V. K. Prasanna, "Energy-balanced multi-hop packet transmission in wireless sensor networks," in *Proc. IEEE GlobeCom*, Dec. 2003, vol. 1, pp. 480-486.

[18] J. Elson, L. Girod, and D. Estrin, "Fine-grained network time synchronization using reference broadcasts," in *Proc. Symposium on Operating Systems Design and Implementation*, Dec. 2002, pp. 147-163.

[19] S. R. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong, "TAG: A Tiny AGgregation service for ad-hoc sensor networks," in *Proc. Symposium on Operating Systems Design and Implementation*, Dec. 2002, pp. 131-146.

[20] W. Heinzelman, A. P. Chandrakasan, and H. Balakrishnan, "An application specific protocol architecture for wireless microsensor networks," *IEEE Trans. Wireless Commun.*, vol. 1, no. 4, pp. 660–670, Oct. 2002.

[21] R. Min, "Energy and Quality Scalable Wireless Communication." Ph.D. diss., MIT, 2003.

[22] Y. Yu, B. Krishnamachari, and V. K. Prasanna, "Exploring energy-latency tradeoffs for data gathering in wireless sensor networks," Department of Electrical Engineering, University of Southern California, Tech. Rep. CENG-2003-05, 2003. [Online]. Available: `http://halcyon.usc.edu/~yangyu/data/TR_CENG200305.pdf`

[23] Y. Yao and G. B. Giannakis, "Energy-efficient scheduling for wireless sensor networks," *IEEE Trans. Comput.*, vol. 53, no. 8, pp. 1333–1342, Aug. 2005.

[24] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Proc. ACM/IEEE International Conference on Mobile Computing and Networking*, Aug. 2000, pp. 56-67.

[25] "PARSEC Project." [Online.] Available: `http://pcl.cs.ucla.edu/projects/parsec`

[26] C. S. Raghavendra and S. Singh, "PAMAS – Power aware multi-access protocol with signaling for ad hoc networks," *ACM SIGCOMM Computer Commun. Review*, vol. 28, no. 3, pp. 5–26, July 1998.

[27] G. Lu, B. Krishnamachari, and C. Raghavendra, "An adaptive energy-efficient and low-latency mac for data gathering in sensor networks," in *Proc. International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks*, Apr. 2004.

Yang Yu is a Senior Research Engineer in the Application Research Center of Motorola Labs. He received his B.S. and M.S. degrees in Computer Science from shanghai JiaoTong university, China in 1996 and 1999 respectively. He received his Ph.D. in Electrical Engineering from Uninversity of Southern California (USC) in 2005. Dr. Yu's research focus includes algorithm design and analysis for energy-efficient information processing and routing in wireless sensor networks. He is a member of IEEE and ACM.

Viktor K. Prasanna (V. K. Prasanna Kumar) received his BS in Electronics Engineering from the Bangalore University and his MS from the School of Automation, Indian Institute of Science. He obtained his Ph.D. in Computer Science from the Pennsylvania State University in 1983. Currently, he is a Professor in the Department of Electrical Engineering as well as in the Department of Computer Science at the University of Southern California (USC). He is also an associate member of the Center for Applied Mathematical Sciences (CAMS) at USC. He served as the Division Director for the Computer Engineering Division during 1994-98. His research interests include parallel computation, computer architecture, VLSI computations, and high performance computing for signal and image processing, and vision.

Dr. Prasanna has published extensively and consulted for industries in the above areas. He has served on the organizing committees of several international meetings in VLSI computations, parallel computation, and high performance computing. He is the Steering Co-chair of the International Parallel and Distributed Processing Symposium (IPDPS) [merged IEEE International Parallel Processing Symposium (IPPS) and the Symposium on Parallel and Distributed Processing (SPDP)] and is the Steering Chair of the International Conference on High Performance Computing (HiPC). He is in the steering committee of the IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS). He serves on the editorial board of the *Journal of Parallel and Distributed Computing*. He has served on the editorial boards of IEEE TRANSACTIONS ON COMPUTERS, IEEE TRANSACTIONS ON VLSI SYSTEMS, and the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS. He is the Editor-in-Chief of the IEEE TRANSACTIONS ON COMPUTERS. He was the founding Chair of the IEEE Computer Society Technical Committee on Parallel Processing. He is a Fellow of the IEEE.

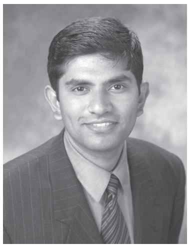

Bhaskar Krishnamachari is an Assistant Professor in the Department of Electrical Engineering-Systems at the University of Southern California (USC), with a joint appointment in Computer Science. He obtained his B.E. in Electrical Engineering from The Cooper Union for the Advancement of Science and Art with a four-year scholarship in 1998, and his M.S. and Ph.D. in Electrical Engineering from Cornell University in 1999 and 2002 respectively with a full graduate fellowship. Dr. Krishnamachari's research focus is on modeling, design and analysis of algorithms for information routing in wireless sensor networks. He has received a 2004 NSF CAREER award, the 2005 USC Viterbi School of Engineering outstanding junior faculty award, and is the first holder of the Philip and Cayley MacDonald Early Career Chair at the USC Viterbi School of Engineering. He serves on the editorial boards of *Elsevier Ad Hoc Networks Journal*, the *ACM Mobile Computing and Communications Review*, and the *EURASIP Journal on Wireless Communcations and Networking*. He is the author of a book titled *Networking Wireless Sensors*, published by Cambridge University Press.