

Optimizing Transmission and Shutdown for Energy-Efficient Real-time Packet Scheduling in Clustered Ad Hoc Networks

Sofie Pollin,^{1,2} Bruno Bougard,^{1,2} Rahul Mangharam,^{1,3} Francky Catthoor,^{1,2} Ingrid Moerman,^{1,4} Ragunathan Rajkumar,³ and Liesbet Van der Perre¹

¹Wireless Research, IMEC, 3001 Leuven, Belgium
Emails: pollins@imec.be, bougardb@imec.be, catthoor@imec.be, vdperre@imec.be

²ESAT/INSYS, Katholieke Universiteit Leuven, 3001 Leuven, Belgium

³Real-Time & Multimedia Systems Laboratory, Carnegie Mellon University, Pittsburgh, PA 15213, USA
Emails: rahulm@ece.cmu.edu, raj@ece.cmu.edu

⁴INTEC, Universiteit Gent, 9000 Gent, Belgium
Email: ingrid.moerman@intec.ugent.be

Received 30 June 2004; Revised 22 March 2005

Energy efficiency is imperative to enable the deployment of ad hoc networks. Conventional power management focuses independently on the physical or MAC layer and approaches differ depending on the abstraction level. At the physical layer, the fundamental tradeoff between transmission rate and energy is exploited, which leads to transmit as slow as possible. At MAC level, power reduction techniques aim to transmit as fast as possible to maximize the radios power-off interval. The two approaches seem conflicting and it is not obvious which one is the most appropriate. We propose a transmission strategy that optimally mixes both techniques in a multiuser context. We present a cross-layer solution considering the transceiver power characteristics, the varying system load, and the dynamic channel constraints. Based on this, we derive a low-complexity online scheduling algorithm. Results considering an M -ary quadrature amplitude modulation radio show that for a range of scenarios a large power reduction is achieved, compared to the case where only scaling or shutdown is considered.

Keywords and phrases: clustered ad hoc networks, energy efficiency, lazy scheduling, shutdown, schedule-based MAC.

1. INTRODUCTION

Ad hoc wireless networks consist of a group of autonomous mobile nodes configuring themselves to form a network that is adapted to the environment and the current needs. A broad range of applications is possible, going from low-rate sensor monitoring applications [1] to high-rate multimedia applications [2]. Both monitoring and multimedia applications are delay sensitive and an appropriate QoS architecture is needed to take care of this in dynamic environments.

On the other hand, ad hoc networks are severely constrained in terms of energy. Wireless communication allows untethered operation, which implies the need for battery-

powered devices. Due to the slow advances in battery technology compared to the growth in system power requirements [3], the use of ad hoc networks is limited by short battery lifetimes. It has already been shown in several design cases [4, 5] that the most critical energy consumers in a wireless node are the radio electronics. Reducing the radio power dissipation is hence crucial to enable the deployment of ad hoc networks with satisfactory lifetime.

Currently, energy-efficient radio communication is tackled differently depending on the level of abstraction. At the physical layer, one tends to exploit the fundamental tradeoff that exists between transmission rate and energy [6, 7]. The information theory has shown that the capacity of the wireless channel increases monotonically with the signal-to-noise ratio [8]. Hence, downscaling the transmission rate—that is, reducing the required channel capacity—allows decreasing the signal-to-noise ratio and therefore the signal power. This

This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

leads to the “lazy scheduling” approach [7], which consists of transmitting with the lowest power over the longest feasible duration.

From a network point of view, the “lazy scheduling” results in a selfish behavior of the individual nodes. A schedule, energy-optimal for one user—that is, which maximizes its timeshare of the wireless channel—might be heavily sub-optimal for the network, since other nodes contending for the channel will have to delay their transmission or speed it up if they have to meet a deadline. Moreover, “lazy scheduling” only optimizes the transmit power. More specifically, it minimizes only the contribution of the electronics whose power consumption is a function of the transmit power. Yet, in low- and middle-range radios, as mostly considered in ad hoc networks, an important part of the power dissipation—that is, the contribution of the frequency synthesizer, the up-conversion mixers, and the filters—is not proportional to the transmit power [9]. This motivates the approaches based on radio shutdown that tend to minimize the duty cycle of the radio circuitry, and therefore transmit as fast as possible. As a result, they give other nodes the maximum timeshare of the channel, showing inherently altruistic behavior. Approaches exist that jointly consider the medium access and routing [10, 11, 12] but neglect the physical layer aspects.

At first sight, the “lazy scheduling” and the shutdown approaches seem conflicting. In this paper, we show that they actually correspond to two extreme cases and that the optimal transmission strategy in a multiuser scenario consists of a cross-layer combination of both approaches. Our contribution in this paper is a solution to determine a transmission strategy with a small and bounded deviation from the global optimum, to be applied to ad hoc wireless networks where individual nodes cooperate. As practical radio implementations only allow a discrete set of transmission schemes, the discrete nature of the problem is taken into account in the system model and solution. We assume the channel is only divided in time, hence no spatial reuse or interference is considered. The core of the scheduling algorithm consists of computing per user a set of transmit opportunities that represent optimally the tradeoff between the transmission time and energy consumption. Then, these are combined across users to determine the schedule with the minimal network energy consumption. The proposed algorithm is adaptive: depending on the traffic constraints and on the current channel states of the users, more transmission scaling or shutdown is considered. This is illustrated using discrete-event simulations under varying traffic loads and node mobility.

Obtaining cooperation in a distributed and multiuser context is not trivial. Approaches based on gaming theory exist to achieve energy efficiency and fairness between rational users [13]. However, the control overhead can be significant to achieve those equilibriums. Scalability and energy-efficiency concerns suggest a hierarchical organization of ad hoc networks. In those cluster-based approaches, a *cluster leader* (CL) is present to be in charge of the clusters maintenance and communication, and is able to enforce solidarity between the users when needed. The CL can be periodically elected not to overload one single node [14]. Therefore, for

the remainder of this paper, we focus on clustered ad hoc networks. The CL is always on to collect the requirements of the other nodes, and to distribute the optimal schedule. We assume that each node in a cluster can overhear the other nodes, hence 1-hop communication is applied within each cluster. Only one cluster is considered in this work. A possible extension would be to employ a scheme similar to [15], and also exploit diversity across clusters.

The remainder of the paper is organized as follows. In Section 2, a detailed overview of work related to the contributions and specific focus of this work is given. Section 3 elaborates on the energy and performance radio model and on the data link control protocol. Taking into account all practical overheads, we present in Section 4 the tradeoff between rate scaling and shutdown. An algorithm is proposed in Section 5 to determine a close-to-optimal time allocation across all users and give results for a multiuser scenario. Finally, conclusions are drawn in Section 6.

2. RELATED WORK

The battery constraints of wireless ad hoc networks have already triggered a lot of research ranging from low-power circuits for analog front end [16], power-aware digital circuitry and embedded software [17] to energy-efficient protocols for medium access control [11, 18]. These works propose solutions that may differ significantly depending on the considered level of abstraction.

At the physical layer, one tries to exploit the fundamental tradeoff that exists between the transmission rate and signal-to-noise ratio [8]. This leads to the so-called “lazy scheduling” approach of Uysal-Biyikoglu et al. [7]. The approach has been extended in [6] to encounter first the discrete nature of the radio settings and second the nonproportionality of the radio circuitry consumption with the transmitted power. Discrete rate scaling is achieved by adapting the constellation size of the modulation, leading to dynamic modulation scaling (DMS), or by changing the code rate (dynamic code scaling, DCS).

From a network point of view, the “lazy scheduling” concept translates in trading off bandwidth (in terms of transmission time) to power. To that extent, it is not trivial to generalize it to the multiuser context. Uysal-Biyikoglu et al. have proposed a generalized version of their algorithm (right-flow) for a broadcast channel and to the multiaccess channel assuming a centralized medium access control protocol [19]. In [20], a practical multiuser lazy scheduling scheme called L-CSMA/CA is proposed. This scheme relies on a CSMA/CA distributed medium access control and considers a finite discrete set of possible transmission rates. For applications with periodic traffic and stringent instantaneous delay requirements, real-time energy-aware packet scheduling is proposed in [21]. In this work, a share of the channel is allocated to each flow depending on its deadline and worst-case data requirements. Depending on its current data requirements, each node makes optimal use of its timeshare, and scales down the transmission rate if possible. Although significant energy gains are achieved, this does not necessarily result in

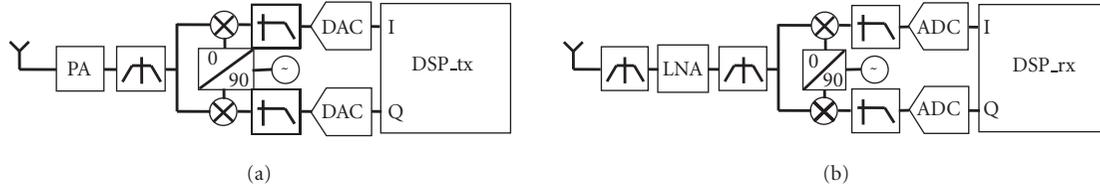


FIGURE 1: (a) The tx and (b) the rx path considered.

the most energy-efficient schedule from network point of view, as it is not exploiting multiuser channel or traffic diversity.

To reduce the part of the energy consumption that is fixed and not related to the transmitted power, the sole option is to minimize the radio duty cycle, shutting down the circuitry as much as possible (sleep mode). However, a node cannot receive data when turned off, hence effective use of the sleep mode requires a significant degree of coordination between nodes. To take care of this coordination at the medium access level, both contention- and schedule-based solutions have been proposed. PAMAS [18] is one of the earliest contention-based energy-efficient protocols that avoids overhearing among neighboring nodes by using out-of-band paging to coordinate the shutdown. TRAMA is a time-slotted, schedule-based MAC that allows nodes to switch to a low power mode when they are not transmitting or receiving [22]. It uses a distributed election scheme based on information about the traffic at each node to determine which node can transmit at a particular timeslot.

To our knowledge, the joint optimization of the a priori contradictory “lazy scheduling” and shutdown approaches has not been studied yet in the dynamic multiaccess context. Although, in [6], a general framework is provided to derive the operating regions when a transceiver should sleep or use transmission scaling, a solution to optimize both in a scenario with multiuser channel or traffic diversity is not proposed. In [9, 23], a transmission strategy, combining transmission rate scaling and sleep duration optimization is studied with and without coding. An offline optimization algorithm is proposed but the scope is limited to a single-user link or a multiuser link with a fixed timeshare for each user. As a result, no solidarity exists between the users in achieving global energy gains in a dynamic environment. In [24], it is shown that the fixed circuit power consumption has a large impact when optimizing the energy consumption across both physical and MAC layers in IEEE 802.11 DCF wireless LANs. However, no shutdown is taken into account in the optimization.

3. SYSTEM MODEL

Prior to analyzing the problem stated above, appropriate energy and performance models have to be defined. We carry out the analysis for modulation scaling. We assume M -ary quadrature amplitude modulation (MQAM), as it is a common case for benchmarking [6, 9]. By varying the modulation order M , the transmission rate can be scaled down.

Other physical layers can be used too, without impact on our algorithm as shown in previous work [25, 26]. The proposed algorithm is general and flexibly adapts to the run time load and physical layer details. In this section, we detail the energy consumption and performance models of the MQAM physical layer. More specifically, we derive the relation that gives the data rate (R), the packet error probability (P_e), and the transmit and receive energies per packet (E_{pt} and E_{pr}) as functions of the transmit power (P_{tx}), the discrete scaling parameter (M) and the transmitter characteristics.

3.1. MQAM radio model

Energy model

Assume that a node can be in one of four modes: (1) a transmit mode, when the transmit part of the radio, including the power amplifier that drives the antenna is on; (2) a receive mode, when the complete receive path of the transceiver is fueled; (3) an idle mode when the receiver is listening to the channel; and (4) a sleep mode, when the complete radio, including the frequency synthesizer is switched off. Let's denote P_{on_tx} , P_{on_rx} , P_{idle} , and P_{sl} , the power consumption in each mode, respectively. The sleep mode power P_{sl} is typically very small when CMOS technology is used [27], so that we neglect it in our model: $P_{sl} \approx 0$. Also, the receiver energy consumption being dominated by the analog part, we can assume that $P_{idle} \approx P_{on_rx}$. Considering the transmit mode, P_{on_tx} corresponds to the DC power of the circuitry (Figure 1), that is, the digital signal processing to produce the baseband signal (P_{dsp_tx}), the digital-to-analog converter (P_{DAC}), the frequency synthesizer to generate the carrier (P_{syn}), the mixers (P_{mix}), and image rejection filters (P_{filt_tx}) to operate the frequency upconversion, and finally the power amplifier (P_{PA}) that drives the current to the antenna. We consider a direct-conversion architecture, so that only one frequency synthesizer and two mixers are required. Hence, P_{on_tx} is given by the following sum:

$$P_{on_tx} = P_{dsp_tx} + 2P_{DAC} + P_{syn} + 2P_{mix} + P_{filt_tx} + P_{PA}. \quad (1)$$

The five first terms of the sum do not vary with the transmit power and the rate scaling parameter. For simplicity, we will refer to this power as P_{elec_tx} . The last term, P_{PA} however depends on the transmit power P_{tx} . We can assume that P_{PA} is, at first order, proportional to the transmit power. We define η as the PA power efficiency:

$$P_{PA} = \frac{P_{tx}}{\eta}. \quad (2)$$

TABLE 1: Parameter values used in our experiment.

Energy model	Performance model	MAC model
P_{tx} (dBm) [0 to 36] (step 0.5)	$A_1 = -40$ dB	$L = 1000$ B
$M[1, 2, 4, 6]$	$K = -4$	$T_{IFS}=10 \mu s$
$W = 1$ MHz	$d = [10-50]$ m	$L_{ACK} = L_{POLL} = 36$ B
$P_{elec_tx} = P_{elec_rx} = 100$ mW	$kT = -174$ dBm/Hz	$L_{header} = L_{NULL} = 20$ B
$T_{wake_up} = 100 \mu s$	$N_f = 10$ dB	$L_{control} = 1$ B
$\eta = 0.3$	$\eta_{IL} = -5$ dB	PER = $10e-3$

From (1) and (2), considering the definition of P_{elec_tx} , we can express P_{on_tx} as

$$P_{on_tx} = P_{elec_tx} + \frac{P_{tx}}{\eta}. \quad (3)$$

Similarly, the receiver DC power can be expressed as a function of the powers of the low-noise amplifier (P_{LNA}), the frequency synthesizer, the downconversion mixers (P_{mix}), the image rejection filters (P_{filt_rx}), the analog-to-digital converter (P_{ADC}), and the digital signal processing (P_{dsp_rx}):

$$P_{on_rx} = P_{LNA} + P_{syn} + 2P_{mix} + 2P_{filt_rx} + 2P_{ADC} + P_{dsp_rx}. \quad (4)$$

We summarize the notation by introducing

$$P_{on_rx} = P_{elec_rx}. \quad (5)$$

From the knowledge of the expression of P_{on_tx} , P_{on_rx} and neglecting P_{sl} , we can compute the energy needed to transmit and receive a packet of L bits:

$$\begin{aligned} E_{tx}(M, P_{tx}) &= P_{on_tx} T_{on}, \\ E_{rx}(M, P_{tx}) &= P_{on_rx} T_{on}. \end{aligned} \quad (6)$$

T_{on} is the time the transmitter or the receiver has to be switched on to, respectively, send or receive the packet. It depends on the modulation scaling parameter M and the packet size L . Assuming a constant bandwidth W (Hz), the symbol rate (or baud rate) for an MQAM modulation is limited to $R_s = W$ (baud). For a constellation size of M , $b = \log_2 M$ bits are transmitted per symbol. Hence, T_{on} is given by

$$T_{on}(M) = \frac{L}{W \log_2 M}. \quad (7)$$

Finally, from (3), (5), (6), and (7), we obtain the expression of E_{tx} and E_{rx} (parameters are listed in Table 1):

$$\begin{aligned} E_{tx}(M, P_{tx}) &= \left(P_{elec_tx} + \frac{P_{tx}}{\eta} \right) \times \frac{L}{W \log_2 M}, \\ E_{rx}(M, P_{tx}) &= P_{elec_rx} \times \frac{L}{W \log_2 M}. \end{aligned} \quad (8)$$

Performance model

Next to the energy model, it is mandatory to derive a performance model that relates the transmit power P_{tx} and the scaling parameter M to the packet error probability. Indeed, to achieve reliable transmission, a corrupted packet has to be retransmitted, which obviously affects the radio energy consumption.

First, the signal-to-noise ratio per symbol (E_s/N_0) at the receiver has to be related to the transmitted power. This requires taking assumptions on the channel. We assume a narrowband flat fading channel is encountered. Also, considering a slowly varying network topology, we can assume that the channel attenuation (due to the path loss and the fading) is constant during a scheduling cycle. The received power is typically expressed as a function of the distance d by (10), where A_1 is the path loss for a distance of 1 m, K is the path loss exponent, α is the random short time fading gain, and η_{IL} represents the implementation loss. E_s/N_0 is given by (10), where k is the Boltzmann constant, T the temperature, and N_f the receiver noise figure:

$$P_r = \alpha A_1 d^K \eta_{IL} P_{tx}, \quad (9)$$

$$\frac{E_s}{N_0} = \frac{P_r}{P_n} = \frac{\alpha A_1 d^K \eta_{IL} P_{tx}}{kT N_f}. \quad (10)$$

With MQAM signaling, assuming an Additive White Gaussian Noise (AWGN) channel, the symbol error probability is bounded by [28]

$$P_M(M, P_{tx}) \leq 2 \cdot \text{erfc} \left(\sqrt{\frac{3}{2(M-1)} \times \frac{E_s}{N_0}} \right). \quad (11)$$

On an AWGN channel, without coding, the symbols errors are noncorrelated, so the packet error probability per transmission can be directly derived from the symbol error probability:

$$P_e(M, P_{tx}) = 1 - (1 - P_M(M, P_{tx}))^{L/b}. \quad (12)$$

Power ratio

The energy saving potential of transmission scaling compared to shutdown depends largely on the relative impact of

the fixed circuit energy consumption to the scalable transmitter power consumption. Given (9) and (10), this ratio (C) can be written as

$$C(d) = \frac{P_{\text{elec_tx}} \times \eta \times \alpha A_1 d^K \eta_{\text{IL}}}{E_s/N_0 \times WkTN_f} = C_{\text{im}} \times d^K. \quad (13)$$

For a given transceiver, it depends on the distance d and on the target performance through the signal-to-noise ratio per symbol (E_s/N_0). Let's fix E_s/N_0 to the value needed to achieve a target packet error rate (PER) of $10e-3$ with $M = 6$.¹ Then, we see that C depends on a transceiver-dependent constant C_{im} and the distance only.

Depending on the value of C , the fixed or the variable part of the power consumption will be dominant. Consider an ad hoc networking scenario where the mobile users are moving around. Clusters are formed dynamically by the hierarchical routing protocol, and the cluster ranges and node density can vary drastically depending on the current node distribution. As such, the underlying scheduling scheme should track at run time the instantaneous C (depending on a node-specific C_{im} and varying distance) of each node, in order to determine the most energy-efficient schedule. Also, the mobility of the different users can be uncorrelated, leading to multiuser diversity that should be exploited to achieve the best possible energy savings.

We carry out the analysis for different ratios to cover different cluster topologies. Using discrete-event simulations, we show results for scenarios where the nodes move around, or have fixed positions. In the next subsection, we show how the node information exchange is implemented and what is the resulting protocol overhead. Next, we show how the optimal schedule can efficiently be determined at run time.

3.2. Data link control protocol

Next to the performance and energy consumption behavior of the radio, the medium access protocol has to be characterized. We consider a centrally controlled protocol as depicted in Figure 2. Periodically, a cluster leader (CL) is elected to be responsible for the cluster scheduling. This CL communicates with the other mobile users (MUs) every scheduling period. To minimize the cost of waking up the radio, all communications of a single MU should be grouped together in the scheduling period. Also, the total time needed for each communication should be known in advance, such that all other MUs can be put asleep during that time. Hence, before each communication round, the schedule has to be determined that allocates to each MU a transmit opportunity TXOP (when to start transmitting and for how long). This optimal timeslot, however, varies with the current data requirements, distance and C_{im} of each MU.

Indeed, the distance and traffic requirements vary and cannot be predicted. To cope with unpredictable traffic

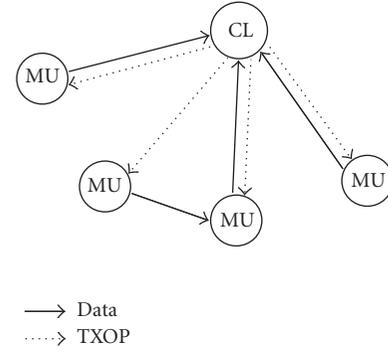


FIGURE 2: Centrally controlled LAN topology illustrating uplink and peer-to-peer communication.

arrivals, it is possible to introduce a look-ahead buffer, during which traffic to be scheduled in the future is captured. This is also proposed in [7, 20]. However, the solution proposed in [20] requires a communication step after each look-ahead period to communicate the data requirements of each user and determine the schedule, prior to the actual data exchanges. It is obvious that, when considering shutdown too, this approach is not optimal as it requires users to wake up more often than needed for the data exchanges alone. It would however be much more practical, for a clustered topology where all traffic is received or overheard by the CL taking the scheduling decision, to piggyback the control information on the periodic data exchanges.

The piggybacking mechanism that enables optimal scaling and shutdown is illustrated in Figure 3. The CL collects the data requirements X_i , which denotes the number of L -sized packets to send, for each MU_i during the period $[D, 2D]$. The scheduling decision is taken at time $2D$. Next, during $[2D, 3D]$, the CL will piggyback the resulting schedule on the data and acknowledgements transmitted during that scheduling period. Finally, during $[3D, 4D]$, each node can send the data it buffered during the initial period $[\varepsilon, D+\varepsilon]$. We note that ε is different and varying for each node, depending on the TXOP allocation for that node. It can be seen that the packet delay is bounded to $[4D-\varepsilon]$ with this scheme.

It should be clear that this delay look-ahead buffer solves the problem of the unpredictable traffic arrivals, without introducing significant communication and wake up costs. Considering the distance MU-CL, introducing this look-ahead delay will result in constraints on the maximum speed of the users. Consider a maximum delay of $4D = 100$ milliseconds, an MU at a speed of 5 km/h will have traveled 0.14 m during that period, which we will show to be negligible.

We want to determine the total energy and time needed to send a packet with a given packet error rate (PER). The protocol overhead introduced by this piggybacking mechanism in addition to the protocol overhead of a centralized and reliable MAC protocol as depicted in Figure 4 is very small. Using the MAC scheme discussed above, for uplink

¹As such, depending on the actual M used for the transmission, the actual power ratio will not be smaller than C .

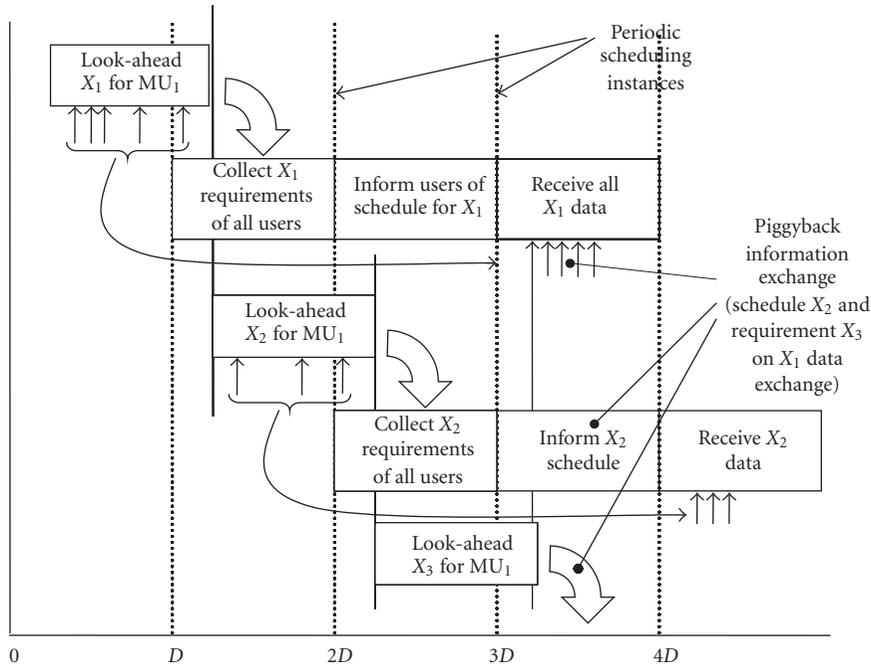


FIGURE 3: The three phases of the delay look-ahead mechanism to obtain optimized transmission rate scaling and shutdown for multiple users: (1) collect data requirements of all users, (2) inform users of schedule, and (3) receive data. All control information is piggybacked on the periodic data transfer to minimize control communication overhead.

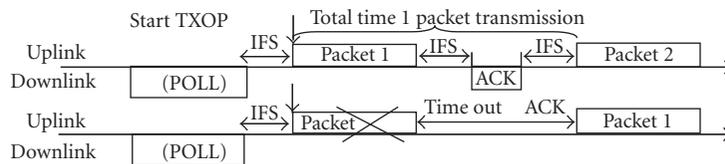


FIGURE 4: Timing of successful and failed uplink packet transmission under a MAC polling scheme.

communication, we can suppress the *POLL* message in most cases. Only in the case no data or ACK between CL and MU are scheduled in a given scheduling period, an additional *POLL* (L_{POLL}) or *NULL* packet with size (L_{NULL}) is needed. In the most efficient case, to implement the control information exchange, it is only needed to foresee an additional 8 bits (L_{control}) for this case study. This is sufficient to communicate a maximum distance of 50 m between CL and MU (see later) and a maximum buffer size of 31 packets. For the exact protocol overheads, we refer to Table 1. This overhead is sent using the same configuration as the data. If there is no data to send (e.g., *NULL* packet), the basic settings $M = 1$ and max P_{tx} are used. Next, using the buffer scheme of Figure 3, the communication is scheduled so that each node is only awake, that is, only consumes energy, when communicating. The wake up energy cost is paid once each scheduling period, and is hence not considered in the per-packet analysis. This leads to the following expressions for the energy for a successful or failed uplink packet transmission, taking into account the overhead of header (L_{header}), messages and

interframe spaces (T_{IFS}) (Table 1, Figure 4):

$$\begin{aligned}
 E_{\text{good_towardsCL}}(M, P_{\text{tx}}) &= E_{\text{tx}}(M, P_{\text{tx}}) \times \frac{L + L_{\text{Header}}}{L} \\
 &+ \left(\left(2 \times T_{\text{ifs}} + T_{\text{on}}(M) \times \frac{L_{\text{ACK}}}{L} \right) P_{\text{on_rx}} \right), \\
 &= E_{\text{bad_CL}}(M, P_{\text{tx}}), \\
 T_{\text{good_CL}}(M) &= T_{\text{on}}(M) \times \frac{L + L_{\text{Header}} + L_{\text{ACK}}}{L} + (2 \times T_{\text{ifs}}) \\
 &= T_{\text{bad_CL}}(M).
 \end{aligned} \tag{14}$$

For peer-to-peer communication, the energy consumed by the receiving node is of interest too. The overhead of the *POLL* or control message to inform the peers of the schedule is not included in the per packet values, and should be added once per scheduling period. This leads to the following

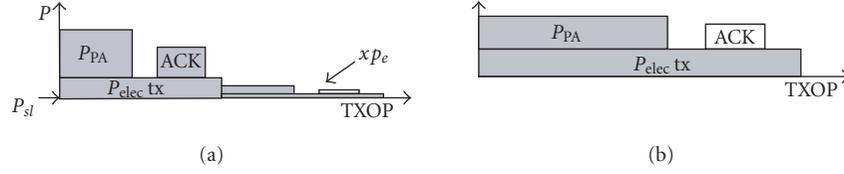


FIGURE 5: Expected Energy consumption and TXOP as a function of variable and fixed energy consumption and the number of retransmissions. (a) A single retransmission is foreseen, and the energy cost is scaled with the probability that this retransmission should happen (as the node could shut down otherwise). (b) No retransmissions are foreseen, as the target PER can be guaranteed by a sufficiently large output power P_{tx} .

expressions for 1 packet, with an increased fixed energy consumption compared to the scenario where data is forwarded to the CL:

$$\begin{aligned}
 E_{\text{bad_peer}}(M, P_{tx}) &= E_{\text{bad_CL}}(M, P_{tx}) + T_{\text{bad_peer}}(M) \times P_{\text{on_rx}}, \\
 E_{\text{good_peer}}(M, P_{tx}) &= E_{\text{bad_peer}}(M, P_{tx}) + \frac{L_{\text{ACK}}}{L} E_{\text{tx}}(M, P_{tx}), \\
 T_{\text{good_peer}}(M) &= T_{\text{bad_peer}}(M) = T_{\text{good_CL}}(M).
 \end{aligned} \tag{15}$$

The expressions for transmission from CL to MU are straightforward. In the remainder of this section, we omit the scenario indices.

When targeting a certain degree of reliability, that is, PER, potential packet retransmissions must be considered in the timeslot. This will allow to determine the total timeslot and expected energy for transmitting a packet with given PER under the given scenario constraints (e.g., distance). The resulting PER when sending a packet with error rate P_e and maximum m retransmissions is

$$P(m, M, P_{tx}) = P_e(M, P_{tx})^{m+1}. \tag{16}$$

Knowing the target degree of reliability by the deadline, the transmit opportunity (TXOP) to be allocated to an MU to send a unit of data L is determined for the worst-case number of retransmissions m needed (17). This might result in channel idle time considering the possibility that a retransmission is not needed. However, we want to determine in advance a schedule that guarantees for each packet the target PER. As a result, the potential allocation of unneeded transmission time to an MU cannot be avoided. Indeed, if probabilistic events would cause the schedule to vary, it would be impossible to determine an optimal schedule in advance and put the nodes to sleep² the time they are not allocated

transmit time (Figure 5):

$$\text{TXOP}(m, M, P_{tx}) = T_{\text{good}}(M, P_{tx}) + m \times T_{\text{bad}}(M, P_{tx}). \tag{17}$$

Considering that the MU is only awake to transmit or retransmit a packet, and sleeps immediately after successful transmission of all queued packets, we can calculate the expected energy consumption for one packet. We consider the expected values, as the number of retransmissions that will be needed is an average variable. Equation (18) scales the energy due to retransmissions with the probability they should happen, that is, the probability that the previous $(j - 1)$ th transmission failed (Figure 5):

$$\begin{aligned}
 E(m, M, P_{tx}) &= (1 - P(m, M, P_{tx})) \times E_{\text{good}}(M, P_{tx}) \\
 &\quad + E_{\text{bad}}(M, P_{tx}) \times (m + 1) \times P(m, M, P_{tx}) \\
 &\quad + E_{\text{bad}}(M, P_{tx}) \times (1 - P_e(M, P_{tx})) \\
 &\quad \times \sum_{j=1}^m P(j - 1, M, P_{tx}) j.
 \end{aligned} \tag{18}$$

4. SYSTEM ENERGY VERSUS TRANSMIT OPPORTUNITY TRADEOFF

In the previous section, expressions are given for the expected energy $E(m, M, P_{tx})$ and timeslot $\text{TXOP}(m, M, P_{tx})$ to communicate a unit of data L , and the resulting error rate $P(m, M, P_{tx})$. They can be determined for each configuration of the output power P_{tx} and scaling parameter M , and each number of retransmissions m , for a given C_{im} and d . In this section, we want to obtain the set of useful points, to be considered by the run-time scheduling algorithm, for each given C_{im} and d .

When determining the expected Energy and TXOP for each configuration (m, M, P_{tx}) , a cloud of discrete points in the Energy-TXOP plane is obtained (Figure 6). However, the only useful points are those that represent the optimal trade-off between Energy and TXOP for a given target error rate P , that is, the points that are closest to the origin (lowest energy and timeslot). Indeed, for each timeshare of the channel allocated to a user, we are interested in the configuration point that achieves the lowest possible energy within this

²It is possible to share retransmission time for packets of the same cluster head. This additional optimization is not considered in this paper.

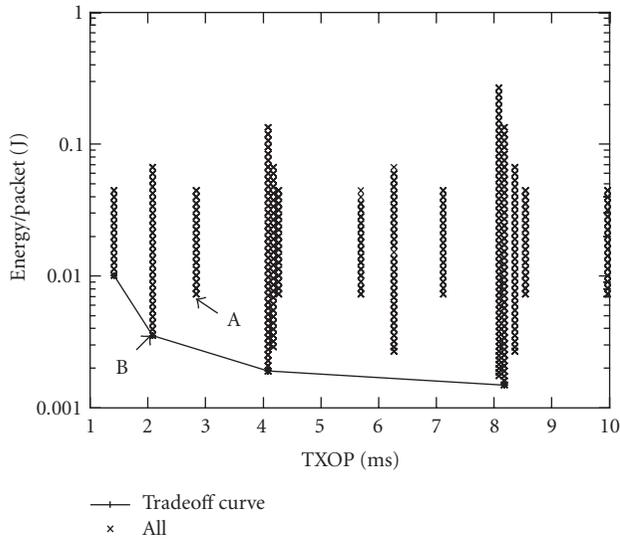


FIGURE 6: Optimal energy versus TXOP to send a unit L of data for different transceiver ratios for distance = 35 m, compared to all points in the energy-TXOP plane that are obtained by varying the different scaling parameters (P_{tx} and M) or the number of retransmissions m , which satisfy the target PER constraint.

timeshare. Consider *configuration A* on Figure 6. This configuration should never be allocated, as for each timeshare it fits in, there exists another configuration that also fits the timeshare and achieves a lower average energy consumption (*configuration B* in this case).

We approximate this complete set of useful points with the piecewise linear interpolation of the convex minorant of the point cloud. The considered tradeoff is then that part of the minorant that is monotonically decreasing (Figure 6). This pruned piecewise linear interpolation of the convex minorant will be called the Energy-TXOP tradeoff curve in the remainder of this paper. Only the discrete points can be allocated in practical transceivers. In fact, this discrete set of optimal configuration points can be determined at the design time (or during a calibration step) of the transceiver. Although the models used in this paper enable an analytical computation of the optimal curves, real system implementations incur lots of complex interactions between both analog and digital components, making the exact tradeoff analytically intractable. As will be shown later, this tradeoff curve captures all information needed to determine efficiently and dynamically the optimal schedule across nodes.

The optimal points should be determined for a range of power ratios, as the value that is of interest depends on the run time operating conditions due to topology variations. Targeting a practical implementation of the algorithm, we only consider a discrete set of calibration curves. Considering a fixed C_{im} per node, a discrete set of distances should be determined to do the calibration. Determining the optimal discrete set of distances for which the calibration step should be performed clearly involves a tradeoff. The larger the set of

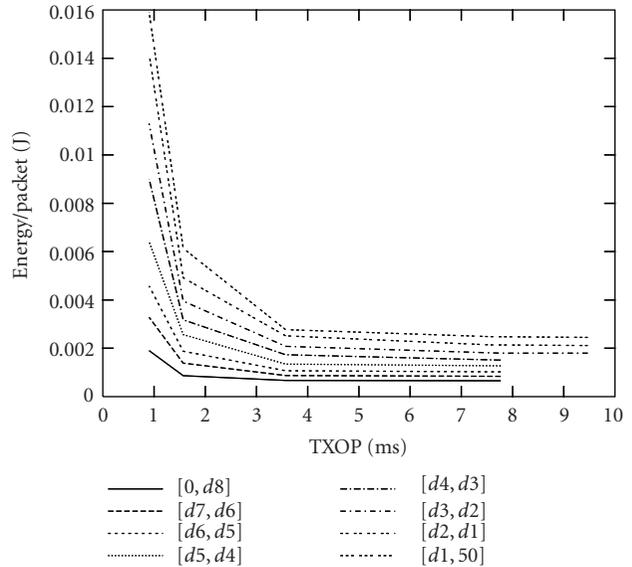


FIGURE 7: Optimal energy versus TXOP for different distances determined according to (19). Based on these curves, we will derive the scheduling algorithm.

curves, the more calibration time will be needed, and more memory to store the databases. Moreover, the overhead to communicate the current distance will increase with finer granularity. On the other hand, a more accurate adaptation to the actual distance will result in more precise adaptation of the output power to the current distance (for the target PER and delay constraint). Also, as the optimal combination of shutdown and scaling depends on the power ratio C , it is also affected by this discretization.

Considering a maximum MU-CL distance of, for example, 50 m, we want to determine the set of discrete distances $\{d_i\}$ that guarantee a bounded suboptimal power consumption at each moment in time. For each actual distance, we use the precomputed curve for a distance that is “just larger” than the actual distance. Allocating a transmit power for a larger distance than the actual one will result in an excessive power allocation, which we want to bound by x . Following this strategy, we determine the optimal set of distances $\{d_i\}$ as:

$$d_0 = 50 \text{ m},$$

$$(d_{i+1})^{-\kappa} = \frac{(1 - xC(d_i))}{(1 + x)} \times (d_i)^{-\kappa}, \quad (19)$$

where x is a positive value smaller than 1 denoting the power loss that can be tolerated between two discrete optimal curves. Enough curves are determined when $xC(d_i) > 1$, that is, the fixed part of the power consumption is dominant so it is not needed to consider smaller distances. In Figure 7, the curves for a maximum distance of 50 m and $x = 0.15$ are

plotted. Only 8 different calibration curves are needed, resulting in only 3 bits required to communicate the distance.

It can be seen that, for smaller d , the Energy-TXOP tradeoff curve spans a much smaller range in energy—that is, downscaling is not beneficial. Indeed, it has been shown that the gains that can be achieved by scaling down the transmission power are small [9]. On the other hand, when the transmit power dominates, a large gain in energy can be achieved when scaling down.

Using this information, we target a TXOP allocation that adapts optimally to the varying distance and data requirements typically encountered in wireless ad hoc networks. Each node is only awake to serve its own data requirements, wasting no energy in overhearing traffic of the other nodes. In the next section, it is shown how the optimal cluster transmission strategy is determined.

5. NETWORK OPTIMAL TRANSMISSION ALLOCATION

Based on the Energy-TXOP tradeoff for each MU, we want to determine the set of transmit opportunities that minimizes the total network energy consumption for the current aggregate data requirement X , which denotes the number of L -sized packets to be transmitted during the next scheduling period D . In the first subsection, we derive an algorithm to compute, based on per packet tradeoff curves of the different MUs, a solution that deviates by a small and bounded offset from the global optimal solution. Second, results are illustrated for a range of scenarios implemented in a discrete-event simulator.

5.1. Cluster TXOP allocation

To determine the optimal transmission strategy for the cluster, we build the aggregate Energy-TXOP tradeoff curve for the whole cluster, based on the aggregate traffic load X and the Energy-TXOP tradeoff curve for each MU. To emphasize the difference between the cluster and per-node tradeoff we call the former $\text{Energy}_{\text{cluster}}\text{-TXOP}_{\text{cluster}}$ and the latter $\text{Energy}_i\text{-TXOP}_i$ tradeoff curve, for a network consisting of N mobile users MU_i , $1 \leq i \leq N$. Each MU_i has data requirement X_i , the aggregate requirement is $X = \sum_{i=1}^N X_i$.

Each MU_i considers, depending on its current distance, its tradeoff curve representing a set of j points, $(E_{i,j}, \text{TXOP}_{i,j})$, $0 \leq j \leq Q$. Each curve is a set of maximal Q (minimal 0) segments with a negative slope:

$$\begin{aligned} s_{i,j} &= |\Delta E_{i,j} / \Delta \text{TXOP}_{i,j}|, \\ \Delta E_{i,j} &= E_{i,j} - E_{i,j-1}, \\ \Delta \text{TXOP}_{i,j} &= \text{TXOP}_{i,j} - \text{TXOP}_{i,j-1}. \end{aligned} \quad (20)$$

Within a tradeoff curve, the segments are ordered according to increasing TXOP or decreasing Energy. Because of the convexity of the curve, the segments are as such ordered according to decreasing negative slope, that is, the energy that can be gained when increasing the allocated timeslot with a time unit decreases. For each curve, the starting point of the

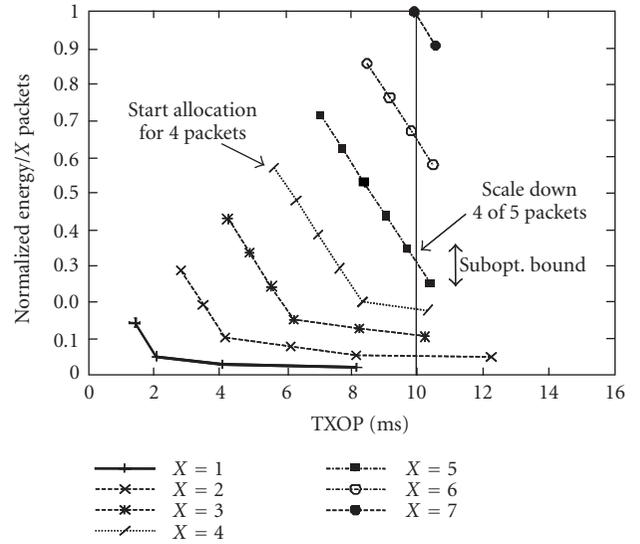


FIGURE 8: Aggregate Energy-TXOP for identical cluster heads, data requirement X from 1 to 7 and scheduling period $D = 10$ milliseconds. Starting from the curve for one packet for a single MU network (lowest curve), the aggregate curves are plotted to send up to 7 packets for that MU within the scheduling period D or equivalently to send 1 packet for 7 MUs with the same per-packet curve (same C_{im} and distance).

first segment $\text{TXOP}_{i,0}$ corresponds to the smallest timeslot allocation with the largest energy consumption.

Based on the $\text{Energy}_i\text{-TXOP}_i$ tradeoff curves and data requirements X_i , we determine the cluster $\text{Energy}_{\text{cluster}}\text{-TXOP}_{\text{cluster}}$ tradeoff consisting of a set of points k , using the following greedy algorithm (See Figure 8 for $X_i = 1$ to 7 and a single MU_i). First the start allocation for the network is determined. This allocation gives to each MU the minimal time needed to satisfy its requirements,³ at maximal energy consumption. In next rounds of the algorithm, energy will be saved by repeatedly allocating more time to some users.

(1) Allocate each MU_i its minimal required $\text{TXOP}_{i,j}$, that is, $\text{TXOP}_{i,0}$. Multiply this timeslot with the total load for this MU_i , to obtain the total timeslot needed for that node in the cluster: $\text{TXOP}_{\text{cluster},i,0} = X_i \times \text{TXOP}_{i,0}$, where $k = 0$ refers to the current (first) point added. This corresponds to an average energy consumption of $E_{\text{cluster},i,0} = X_i \times E_{i,0}$ for that node. Knowing the requirements for each node i , we can construct the first point $k = 0$ of the cluster $\text{Energy}_{\text{cluster}}\text{-TXOP}_{\text{cluster}}$ tradeoff: $(E_{\text{cluster},k}, \text{TXOP}_{\text{cluster},k})$:

$$\begin{aligned} E_{\text{cluster},0} &= \sum_{i=1}^N E_{\text{cluster},i,0}, \\ \text{TXOP}_{\text{cluster},0} &= \sum_{i=1}^N \text{TXOP}_{\text{cluster},i,0}. \end{aligned} \quad (21)$$

³We assume it is always possible to construct this first point. Hence, no overload is taken into account.

The first point is the sum of the per-node minimal resource requirements, resulting in the maximum energy consumption for the cluster. After determining the first point of the curve, we will construct the whole cluster curve allowing for optimal decrease of the energy consumption. We will add points k to the $\text{Energy}_{\text{cluster}}\text{-TXOP}_{\text{cluster}}$ curve, using the segments $s_{i,j}$ of the per MU_i individual curves. MU_i with no segment $s_{i,j}$ are not longer considered, as their only TXOP ($= \text{TXOP}_{i,0}$) has already been allocated. As the curve for each MU_i consists of different segments depending on their current distance and C_{im} , the loop j across the segments will be different for each MU_i . Hence, from now, we denote $j(i)$. After this initialization, we set $j(i) = 1$ for each node i ; $k' = 0$ for the cluster, that is, k' denotes the last added point to the aggregate optimal curve.

(2) Search across the set of current segments $s_{i,j(i)}$ those with the largest negative slope S . As such, we are sure that the best possible energy saving is obtained across the cluster. For each MU_i with current slope $s_{i,j(i)} = S$ and for each of its packets X_i ,⁴ a new point is added to the aggregate tradeoff curve, resulting in segments $s_{\text{cluster},k} = |\Delta E_{\text{cluster},k} / \Delta \text{TXOP}_{\text{cluster},k}|$, where each increment can be understood as increasing the time allocated to one packet of one MU_i , hence $\Delta \text{TXOP}_{\text{cluster},k} = \Delta \text{TXOP}_{i,j(i)}$. This results in a network energy decrease $\Delta E_{\text{cluster},k} = \Delta E_{i,j(i)}$. The result of this step is a set of network allocation vectors with lower aggregate expected energy but a larger time allocation:

$$\begin{aligned} (E_{\text{cluster},k}, \text{TXOP}_{\text{cluster},k}), \quad \forall k \mid k' < k \leq \left(k' + \sum_{i \mid s_{i,j(i)}=S} X_i \right), \\ E_{\text{cluster},k} &= E_{\text{cluster},k-1} - \Delta E_{\text{cluster},k}, \\ \text{TXOP}_{\text{cluster},k} &= \text{TXOP}_{\text{cluster},k-1} + \Delta \text{TXOP}_{\text{cluster},k}, \end{aligned} \quad (22)$$

where k' denotes the number of points after the previous step. The sum of the number of packets across the selected MU_i 's corresponds to the number of points added in this step. After adding all points, the current set of segments is updated. This means that for each MU_i that was treated in this step, the next segment of its tradeoff curve (if it exists) is considered: $j(i) \leftarrow (j(i) + 1)$, for all $i \mid (s_{i,j(i)} = S)$. Also the aggregate curve counter is updated: $k' = k$.

(3) Repeat step 2 until all segments $s_{i,j(i)}$ for all MU_i are treated. A network tradeoff curve with maximum QXX points is constructed, Q denoting the maximum number of segments per $\text{Energy}_i\text{-TXOP}_i$ curve for each MU_i .

Knowing the cluster $\text{Energy}_{\text{cluster}}\text{-TXOP}_{\text{cluster}}$ curve, the network allocation vector corresponds to the point with the largest aggregate $\text{TXOP}_{\text{cluster},k}$ that is smaller than the scheduling period D , as illustrated in Figure 8 for $D = 10$ milliseconds. It is clear that for larger data requirements, less downscaling is possible. The figure represents a set of

⁴The exact order to add extra time for each packet of different mobile users should be random to achieve fairness.

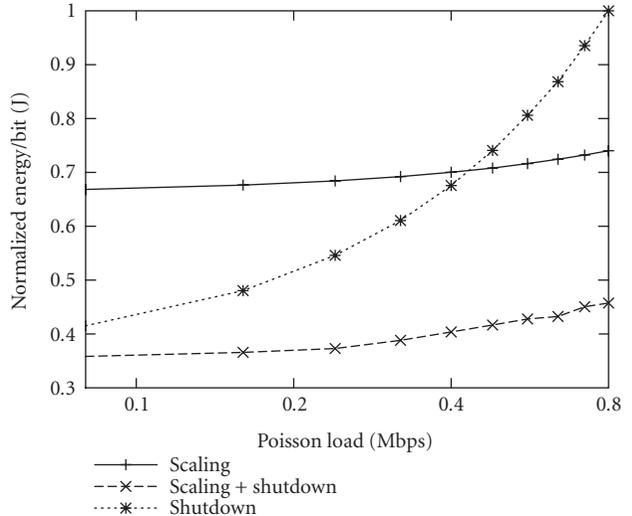


FIGURE 9: Normalized energy per bit for a topology of 5 nodes, $D = 100$ milliseconds, distance 33 m, for a range of poisson loads.

aggregate $\text{Energy}_{\text{cluster}}\text{-TXOP}_{\text{cluster}}$ curves for a single MU_i with data requirement X_i ranging from 1 to 7 packets per period. The complexity to construct the aggregate curve is $O(NQ \log(N))$.

It can be shown that solving this kind of discrete optimization problems with a greedy approach (e.g., according to steepest decreasing slope) based on the convex piecewise-linear interpolation of the tradeoff results in a solution that is bounded suboptimal. This can be understood intuitively, as shown in Figure 8. As the solution relies on the convex piecewise-linear interpolation of the tradeoff, each discrete point of the aggregate curve corresponds to an optimal allocation, but only for a scheduling period D that is exactly equal to $\text{TXOP}_{\text{cluster},k}$ of the selected point k . However, most often, a point has to be taken with a value that is slightly smaller than D . The greedy search based on pruned convex tradeoff curves however does not guarantee that there does not exist a solution with $\text{TXOP}_{\text{cluster},\text{optimal}}$ that is larger than $\text{TXOP}_{\text{cluster},k}$ but smaller than D (and has a smaller energy consumption $E_{\text{cluster},\text{optimal}}$). However, due to convexity, this point has to be above the piecewise linear tradeoff curve. Consequently, it can be seen that the worst case difference between $E_{\text{cluster},\text{optimal}}$ and $E_{\text{cluster},k}$ is bounded by the ΔE_{max} across all segments of the curve, which is relatively small and depends on the granularity of the system parameters considered.

5.2. Results

To illustrate the strengths of the proposed scheme over a range of load scenarios and node topologies, we have implemented it in the discrete-event simulator *ns-2* [29]. The implementation reflects the full energy and performance behavior of the MQAM radio as presented in Section 3.1. Next, the delay look-ahead scheduling protocol presented in Section 3.2 has been implemented on top of a centrally

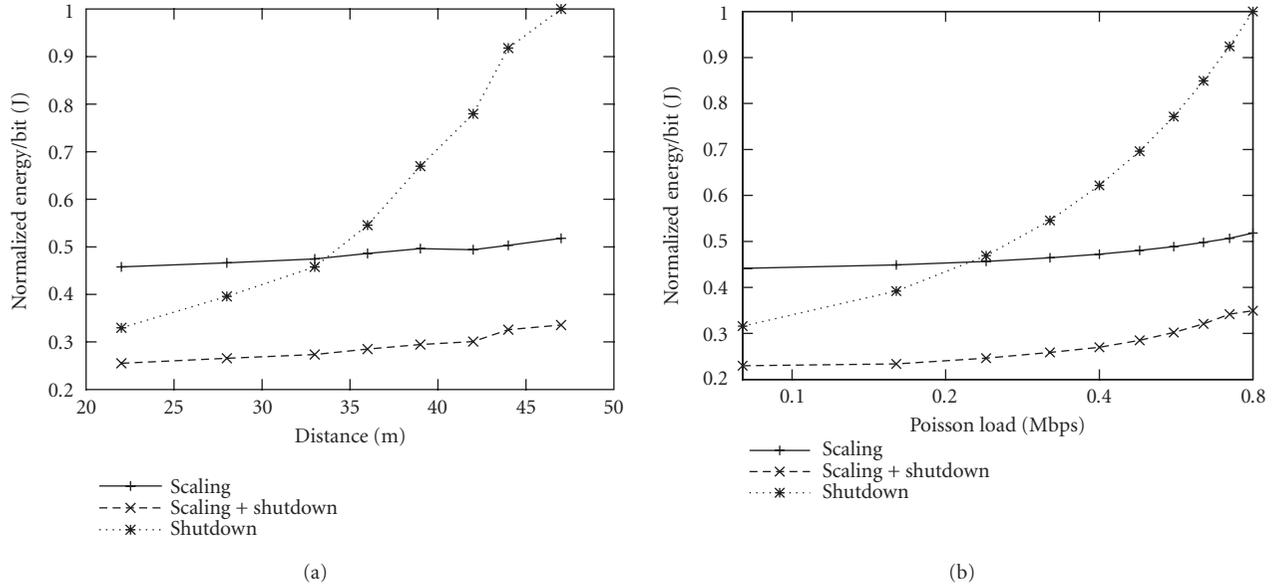


FIGURE 10: Normalized energy per bit for a topology of 5 nodes, $D = 100$ milliseconds (a) with Poisson load of 0.4 Mbps for a range of distance, (b) moving around randomly for a range of Poisson loads.

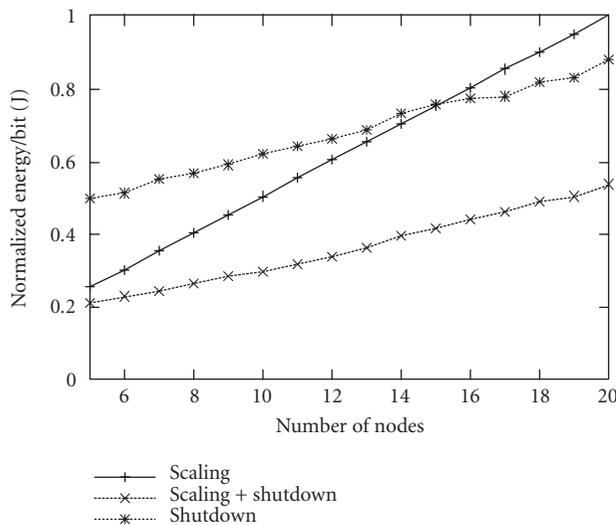


FIGURE 11: Normalized energy per bit for a topology with a range of nodes, with aggregate CBR load of 1.6 Mbps, distance of 33 m, $D = 100$ milliseconds.

controlled reliable MAC scheme. The exact overhead considered for the MAC protocol is given in Table 1. When there is no data available, a NULL packet is sent. The proposed scheme is compared with energy management techniques that use scaling or shutdown only. In the shutdown only protocol, we do adapt the output power to the given distance (but do not scale down the transmission rate).

Simulations have been carried out for a range of mobile users, with identical C_{im} , but with possible different and varying CL-MU distances. The scheduling database has been

generated according to the parameters listed in Table 1 and using (1)–(19). This results in a database for the distances [22, 29, 33, 37, 40, 42, 45, 47] m. Using the broad range of scenarios possible with this discrete-event simulation tool, we mainly want to show that the proposed algorithm indeed optimally adapts to the instantaneous scenario constraints, exploiting more scaling or shutdown depending on the scenario, to achieve maximum energy savings.

First, we show that depending on the current traffic load, shutdown or scaling achieves larger energy savings. The proposed algorithm, however, adapts and achieves for each load instance the best possible gains. Figure 9 shows the energy consumptions of the proposed scheme, compared to shutdown or scaling only, for a Poisson load up to 0.8 Mbps, and a distance of 33 m. It can be seen that when the load is small, more shutdown should be used. However, when the load increases, the use of transmission scaling becomes more and more useful. The proposed scheme however adapts and achieves at each moment a smaller energy consumption.

Next, we consider the effect of mobility on the energy consumption. As mentioned before in Section 3.1, a larger distance corresponds to a more dominant transmission power. To that extent, the gains of shutdown compared to scaling also vary with distance, as illustrated in Figure 10a for a CBR load of 0.4 Mbps over 5 users at varying (fixed) distance. In Figure 10b, the energy is plotted over a range of Poisson loads, for 5 users with mobility 2 km/h, walking around in a square of 50 m by 50 m, with the CL in the origin. The mobility pattern has been generated using the *setdest* tool for *ns-2*. It can be seen that, when introducing mobility and hence larger distances than the 33 m of Figure 9, the overall gains of scaling are larger, resulting in the crossing of the “scaling” and “shutdown” curves for a lower load.

The proposed scheme however adapts and exploits the possibilities to save energy for each distance and load optimally.

Finally, we investigate the effect of increasing the number of users (Figure 11). It can be seen, for an aggregate CBR load of 1.6 Mbps (or 37.5%) that the energy consumed when using the “scaling” energy management technique increases linearly with the number of nodes (for the same aggregate network load). This is because the idle and receiver energy will scale linearly with the number of nodes, irrespective of the aggregate load. When adding the possibility to shutdown, the energy increase with increasing number of nodes is much slower. In this case, each node is asleep when the others transmit. The energy increase is hence only due to increase wake up cost, and the increased probability to send a NULL packet when the queue is empty (as the per-node load decreases). It should be noted that it depends on the network density to decide whether the “shutdown” or “scaling” solution is the most energy efficient. The proposed adaptive solution, however, takes advantage of both techniques in each situation.

6. CONCLUSIONS

In this paper, we propose a transmission strategy that combines close-to-optimally “lazy scheduling” and shutdown, two energy management techniques that seem contradictory. The former exploits the fundamental tradeoff between the time and energy needed to send a unit of data, and hence maximizes the transmission duration to minimize the transmit energy consumption. The latter minimizes the fixed circuit energy consumption, hence decreasing the transceiver on time as much as possible. We show that the optimal transmission strategy in a multiuser scenario is a combination of both approaches. Moreover, the optimal combination differs depending on the instantaneous scenario traffic and channel constraints.

First, we derive a solution to determine a transmission strategy with a worst-case deviation from the optimal strategy that is bounded. As practical radio implementations only allow a discrete set of transmission schemes, this discrete nature of the problem is taken into account in the system model and solution. The proposed algorithm is adaptive: depending on the traffic constraints and on the relative impact of the transmission power to the circuit energy consumption, more transmission scaling or shutdown is considered. We show that the algorithm indeed results in significant energy savings for a range of traffic loads and transceiver characteristics, using discrete-event simulation. It adaptively combines and trades off the gains that can be achieved when scaling or shutting down only, and hence significantly outperforms those energy management techniques in each scenario. Moreover, it optimally exploits multiuser diversity by scaling down the rate of those users where the instantaneous gains are the largest.

ACKNOWLEDGMENT

The work presented in this paper is partly based on results published in EWSN '05.

REFERENCES

- [1] D. Estrin, R. Govindan, J. S. Heidemann, and S. Kumar, “Next century challenges: Scalable coordination in sensor networks,” in *Proc. 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99)*, pp. 263–270, Seattle, Wash, USA, August 1999.
- [2] H.-L. Chao and W. Liao, “Credit-based slot allocation for multimedia mobile ad hoc networks,” *IEEE J. Select. Areas Commun.*, vol. 21, no. 10, pp. 1642–1651, 2003.
- [3] T. E. Starner, “Powerful change part 1: batteries and possible alternatives for the mobile market,” *IEEE Pervasive Computing*, vol. 2, no. 4, pp. 86–88, 2003.
- [4] A. Chandrakasan, R. Amirtharajah, S. Cho, et al., “Design considerations for distributed microsensor systems,” in *Proc. IEEE Custom Integrated Circuits Conference (CICC '99)*, pp. 279–286, San Diego, Calif, USA, May 1999.
- [5] J. M. Rabaey, M. J. Ammer, J. L. da Silva Jr., D. Patel, and S. Roundy, “PicoRadio supports ad hoc ultra-low power wireless networking,” *IEEE Computer*, vol. 33, no. 7, pp. 42–48, 2000.
- [6] C. Schurgers, *Energy-aware wireless communications*, Ph.D. thesis, University of California, Los Angeles, Calif, USA, 2002.
- [7] E. Uysal-Biyikoglu, B. Prabhakar, and A. El Gamal, “Energy-efficient packet transmission over a wireless link,” *IEEE/ACM Trans. Networking*, vol. 10, no. 4, pp. 487–499, 2002.
- [8] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, no. 1, 3, pp. 379–423, 623–656, 1948.
- [9] S. Cui, A. J. Goldsmith, and A. Bahai, “Modulation optimization under energy constraints,” in *Proc. IEEE International Conference on Communications (ICC '03)*, Anchorage, Alaska, USA, May 2003.
- [10] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, “Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks,” in *Proc. 7th Annual International Conference on Mobile Computing and Networking (MobiCom '01)*, pp. 85–96, Rome, Italy, July 2001.
- [11] W. Ye, J. S. Heidemann, and D. Estrin, “An energy-efficient MAC protocol for wireless sensor networks,” in *Proc. IEEE 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '02)*, vol. 3, pp. 1567–1576, New York, NY, USA, June 2002.
- [12] Y. Xu, J. S. Heidemann, and D. Estrin, “Geography-informed energy conservation for ad hoc routing,” in *Proc. 7th Annual International Conference on Mobile Computing and Networking (MobiCom '01)*, pp. 70–84, Rome, Italy, July 2001.
- [13] V. Srinivasan, P. Nuggehalli, C. F. Chiasserini, and R. R. Rao, “Cooperation in wireless ad hoc networks,” in *Proc. IEEE 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '03)*, vol. 2, pp. 808–817, San Francisco, Calif, USA, March–April 2003.
- [14] Z. Cai, M. Lu, and X. Wang, “Channel access-based self-organized clustering in ad hoc networks,” *IEEE Transactions on Mobile Computing*, vol. 2, no. 2, pp. 102–113, 2003.
- [15] Y. Yu, B. Krishnamachari, and V. K. Prasanna, “Energy-latency tradeoffs for data gathering in wireless sensor networks,” in *Proc. 23rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '04)*, vol. 1, pp. 244–255, Hong Kong, China, March 2004.
- [16] A.-S. Porret, T. Melly, C. C. Enz, and E. A. Vittoz, “A low-power low-voltage transceiver architecture suitable for wireless distributed sensors network,” in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS '00)*, vol. 1, pp. 56–59, Geneva, Switzerland, May 2000.
- [17] R. Min and A. Chandrakasan, “A framework for energy-scalable communication in high-density wireless networks,”

in *Proc. International Symposium on Low Power Electronics and Design (ISLPED '02)*, pp. 36–41, Monterey, Calif, USA, August 2002.

- [18] S. Singh and C. Raghavendra, "PAMAS: Power aware multi-access protocol with signaling for ad hoc networks," *ACM Computer Communication Review*, vol. 28, no. 3, pp. 5–26, 1998.
- [19] A. El Gamal, C. Nair, B. Prabhakar, E. Uysal-Biyikoglu, and S. Zahedi, "Energy-efficient scheduling of packet transmissions over wireless networks," in *Proc. IEEE 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '02)*, vol. 3, pp. 1773–1782, New York, NY, USA, June 2002.
- [20] R. R. Kompella and A. C. Snoeren, "Practical lazy scheduling in sensor networks," in *Proc. 1st International Conference on Embedded Networked Sensor Systems (ACM SenSys '03)*, pp. 280–291, Los Angeles, Calif, USA, November 2003.
- [21] C. Schurgers, V. Raghunathan, and M. B. Srivastava, "Modulation scaling for real-time energy aware packet scheduling," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM '01)*, vol. 6, pp. 3653–3657, San Antonio, Tex, USA, November 2001.
- [22] V. Rajendran, K. Obraczka, and J. J. Garcia-Luna-Aceves, "Energy-efficient collision-free medium access control for wireless sensor networks," in *Proc. 1st International Conference on Embedded Networked Sensor Systems (ACM SenSys '03)*, pp. 181–192, Los Angeles, Calif, USA, November 2003.
- [23] S. Cui, A. J. Goldsmith, and A. Bahai, "Energy-constrained modulation optimization for coded systems," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM '03)*, vol. 1, pp. 372–376, San Francisco, Calif, USA, December 2003.
- [24] J. Zhao, Z. Guo, and W. Zhu, "Power efficiency in IEEE 802.11a WLAN with cross-layer adaptation," in *Proc. IEEE International Conference on Communications (ICC '03)*, vol. 3, pp. 2030–2034, Anchorage, Alaska, USA, May 2003.
- [25] S. Pollin, B. Bougard, R. Mangharam, et al., "Optimizing transmission and shutdown for energy-efficient packet scheduling in sensor networks," in *Proc. European Workshop on Wireless Sensor Networks (EWSN '05)*, Istanbul, Turkey, January–February 2005.
- [26] R. Mangharam, S. Pollin, B. Bougard, et al., "Optimal fixed and scalable energy management for wireless networks," in *Proc. IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '05)*, Miami, Fla, USA, March 2005.
- [27] B. Razavi, *Design of Analog CMOS Integrated Circuits*, McGraw-Hill, New York, NY, USA, 2001.
- [28] J. Proakis, *Digital Communications*, McGraw-Hill, New York, NY, USA, 1995.
- [29] "ns-2 Network Simulator," <http://www.isi.edu/nsnam/ns>.

Sofie Pollin received the M.S. degree in electrical engineering from the Katholieke Universiteit Leuven, Belgium, in 2002. In October 2002, she joined the Wireless Research group, the Interuniversity Microelectronics Center (IMEC), and started her Ph.D. thesis at the Electrical Engineering Department, the Katholieke Universiteit Leuven. Her current research focuses on the cross-layer design and implementation of adaptive and low-power wireless networking systems. In the summer of 2004, she was a Visiting Scholar at National Semiconductor, Santa Clara, Calif, and in the summer of 2005 at UC Berkley.



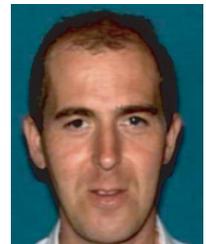
Bruno Bougard received the M.S. degree in electrical engineering from the Polytechnic Institute of Mons, Belgium, in 2000. He joined the Interuniversity Microelectronics Center (IMEC), Leuven, Belgium, in June 2000, as a Research Engineer in the Wireless Research Group. His current research focuses on design methodologies for low-power wireless communication systems. He previously contributed as a system architect to the design, the optimization, and the characterization of low-power, high-data-rate turbo decoder architecture. Since 2002, he has been a Research Assistant of the Fund for Scientific Research (Belgium) and a Ph.D. candidate at the Electrical Engineering Department, the Katholieke Universiteit Leuven, Belgium, still carrying out his research at IMEC.



Rahul Mangharam is a Ph.D. student in the Department of Electrical and Computer Engineering, Carnegie Mellon University, USA. His interests are in scheduling algorithms for wireless and embedded systems. He was a Visiting Scholar in the Wireless Systems Group at IMEC, Belgium, in 2003. In 2002, he was a member of technical staff in the Ultra-Wide Band Wireless Group at Intel Labs. He has worked on ASIC chip design at Marconi Communications (1999) and Gigabit Ethernet at Apple Computer Inc. (2000).



Francky Catthoor is a Fellow at IMEC, Heverlee, Belgium. He is also an IEEE Fellow. He received the Engineering degree and a Ph.D. degree in electrical engineering from the Katholieke Universiteit Leuven, Belgium, in 1982 and 1987, respectively. Between 1987 and 1999, he has headed research domains in the area of architectural and system-level synthesis methodologies, within the DESICS (formerly VSDM) Division at IMEC. His main current research activities belong to the field of architecture design methods and system-level exploration for power and memory footprint within real-time constraints, oriented towards data storage management, global data transfer optimization, and concurrency exploitation. Platforms that contain both customizable/configurable architectures and (parallel) programmable instruction-set processors are targeted. Also deep-submicron technology issues are included.



Ingrid Moerman was born in Gent, Belgium, in 1965. She received the Eng. degree in electrotechnical engineering and the Ph.D. degree from the Ghent University, Gent, Belgium, in 1987 and 1992, respectively. Since 1987, she has been with the Interuniversity Microelectronics Center (IMEC), the Department of Information Technology (INTEC), the Ghent University, where she conducted research in the field of optoelectronics. In 1997, she became a permanent member of the research staff at IMEC. Since 2000, she has been a part-time Professor at the Ghent University. Since 2001, she has switched



her research domain to broadband communication networks. She is currently involved in the research and education on broadband mobile and wireless communication networks and on multimedia over IP. Her main research interests related to mobile and wireless communication networks are adaptive QoS routing in wireless ad hoc networks, personal networks, body area networks, wireless access to vehicles (high bandwidth & driving speed), protocol boosting on wireless links, design of fixed access/metro part, traffic engineering and QoS support in the wireless access network. She is an author or coauthor of more than 300 publications in the field of optoelectronics and communication networks.

Ragunathan Rajkumar is a Professor in the Departments of Electrical and Computer Engineering and of Computer Science, Carnegie Mellon University. He obtained his B.E. (honors) degree from the University of Madras in 1984, and his M.S. and Ph.D. degrees from Carnegie Mellon University in 1986 and 1989, respectively. His research interests include all aspects of embedded real-time systems as well as QoS support in operating systems and networking. He was also the primary founder of TimeSys Corporation, a vendor of embedded Linux and Java products. He has chaired several international conferences and has authored a book and more than 90 publications in conferences and journals.



Liesbet Van der Perre received the M.S. degree and the Ph.D. degree in electrical engineering from the Katholieke Universiteit Leuven, Belgium, in 1992 and 1997, respectively. Her work in the past focused on system design and digital modems for high-speed wireless communications. She was a System Architect in IMEC's OFDM ASICs development and a Project Leader for the Turbo codec. Currently, she is the Scientific Director of wireless research in IMEC.

