

Latency-Sensitive Power Control for Wireless Ad-hoc Networks

Mohamed R. Fouad, Sonia Fahmy, Gopal Pandurangan
Department of Computer Science, Purdue University
250 N. University St., West Lafayette, IN 47907–2066, USA
E-mail: {mrf, fahmy, gopal}@cs.purdue.edu *

ABSTRACT

We investigate the impact of power control on latency in wireless ad-hoc networks. If transmission power is increased, interference increases, thus reducing network capacity. A node sending/relaying delay-sensitive real-time application traffic can, however, use a higher power level to reduce latency, if it considers information about load and channel contention at its neighboring nodes. Based on this observation, we formulate a new distributed power control protocol, *Load-Aware Power Control (LAPC)*, that heuristically considers low end-to-end latency when selecting power levels. We study the performance of LAPC via simulations, varying the network density, node dispersion patterns, and traffic load. Our simulation results demonstrate that LAPC achieves an average *end-to-end* latency improvement of 54% over the case when nodes are transmitting at the highest power possible, and an average end-to-end latency improvement of 33% over the case when nodes are transmitting using the lowest power possible, for uniformly dispersed nodes in a lightly loaded network.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; D.4.4 [Communications management]: Network communication; D.4.8 [Performance]: Simulation

General Terms

Algorithms, Design, Performance

Keywords

ad-hoc networks, wireless networks, power control, network connectivity, latency, medium access control

1. INTRODUCTION

Power control is the process whereby each node selects an appropriate transmission power level, out of a discrete set of levels, each

*– This research has been sponsored in part by NSF grant ANI-0238294 (CAREER).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MSWiM'05 Workshop on QoS and Security in Wireless and Mobile Networks, October 10–13, 2005, Montreal, Quebec, Canada.
Copyright 2005 ACM 1-59593-188-0/05/0010 ...\$5.00.

time it transmits a packet. This process can be a function of several factors, such as the residual energy, neighbor proximity, and the type of application. In general, the higher the transmission power, the wider the area covered by that transmission, and, accordingly, the higher the number of nodes directly accessible, as depicted in Fig. 1. Cisco Systems Aironet [1] and Crossbow MICA2 motes [2] are two examples of systems that allow power level selection.




Figure 1: The range increases with transmission power increase

Power control has been extensively studied in the literature [3, 4, 5, 6, 7, 8, 9]. Most of these studies, however, aim at prolonging the network lifetime by reducing battery power consumption at each wireless node. Little work has considered end-to-end latency, which includes propagation and transmission delays, as well as the delay encountered by the packets at the intermediate nodes on their route from the source to the destination, as a result of the node queuing policy and/or interference and channel conditions.

Clearly, latency minimization and energy efficiency are conflicting goals. Reducing the end-to-end latency can be of a great significance to real-time applications or to sensor applications during certain *critical situations*. For example, consider a sensor/actuator network of accelerometers and dampers installed in a building to detect and react to unusual motion. Accelerometer data indicating an earthquake or a strong wind-storm requires an immediate response (e.g., on the order of tens of milliseconds) by the dampers that react to this data. In this scenario, latency considerations, at least temporarily, override energy efficiency. Similarly, in military applications such as tracking enemy vehicles, it is imperative that latency be maintained within certain bounds. Communication among soldiers in the battlefield can also have severe consequences if it gets delayed. Finally, multimedia transport over wireless networks must support applications such as multimedia streaming, video-conferencing, and surveillance, with strict latency constraints.

In this paper, we investigate latency versus energy efficiency tradeoffs encountered when making power control decisions. We define a simple distributed protocol that allows increasing the power

levels in order to reduce latency when required. We note, however, that increasing the power level reduces latency only under light loads or with non-contention-based MAC protocols. Since latency is the time for a packet to travel from the source to the destination (which may be several hops away), our objective is to reduce *end-to-end* latency. Latency reduction is achieved since increasing the power level will enable the packet to reach the destination in a fewer number of hops under low load/contention conditions. Otherwise, the increased interference and reduced capacity can increase the queuing delays and increase the end-to-end latency, even if the number of hops is reduced. We therefore augment our distributed protocol with a heuristic that considers neighbor load information when making power control decisions. Note also that our proposed protocol can be implemented on top of contention-based or non-contention-based MAC protocols. Carrier sensing is part of the underlying MAC protocol. What our protocol primarily does is that each node estimates the power level that it will be using for transmission by exploiting its knowledge of the load history at neighboring nodes. The node is then free to communicate using whichever MAC protocol it is configured to use at the elected power level.

The remainder of this paper is organized as follows. Section 2 gives our network model and performance measures. Section 3 examines related work and the impact of power control on network connectivity, MAC layer protocols, interference, network throughput, routing, and latency. Section 4 states the problem that we address in this paper and describes our approach. Section 5 discusses our simulation model and results. Finally, Section 6 gives brief concluding remarks and future work plans.

2. NETWORK MODEL

We model an ad-hoc wireless network as a graph $G(V, E)$ where V is the set of all nodes (vertices) and E is the set of all links (edges) at certain power levels. We make the following assumptions: (i) The nodes $v \in V$ are quasi-stationary. This is typical for several types of ad-hoc networks, including sensor networks; (ii) All nodes have similar capabilities and equal significance; and (iii) Each node has a fixed number of transmission power levels. Note that in our model, no assumptions are made about any of the following: (i) homogeneity of node dispersion in the field; (ii) location-awareness, i.e. being equipped with GPS-capable antennae; (iii) network density or diameter; and (iv) distribution of energy consumption among nodes.

We study the following performance measures: (i) **Network lifetime**: In our simulations, we measure the network lifetime as the time until the *first* node depletes its energy, and (ii) **End-to-end latency**: The end-to-end latency is the average delay encountered by all packets on their routes from source $S_i \in V$ to destination $D_i \in V$. This delay can be attributed to: (1) *queuing delays*, which denote the time spent by the packet in the output queues of the nodes, (2) *transmission delays*, which denote the amount of time a network of bandwidth W bps needs to transmit a packet of length L bits. This is equal to $\frac{L}{W}$, and (3) *propagation delays*, which denote the time needed by the packet to propagate in the wireless medium for all the hops between the source and the destination. This delay depends on the distance d between the source and destination for every hop, and the speed S at which the packet travels, which should be no faster than the speed of light (approximately 3×10^8 m/s).

3. RELATED WORK

In this section, we summarize the impact of power control on

connectivity, medium access control protocols, interference, throughput, routing protocols, and latency.

3.1 Power Control and Connectivity

Reducing energy consumption requires that nodes lower their transmission power to a minimum, without partitioning the network. Gupta et al. [10] give the critical power a node in an ad-hoc network needs to transmit at, in order to ensure that the network is connected with probability one, as the number of nodes in the network grows to ∞ . In [3], the problem of adjusting the transmission power to control network connectivity is addressed. The problem is formulated as a constrained optimization problem with the connectivity as its constraint and the power used as its objective function. In [4], the effects of using different power levels on the average energy consumption and network throughput are investigated. ‘‘Clusters’’ where a node adapts its transmission power so as to establish connectivity are defined. In [11], it is proposed that each node should make local decisions, which collectively guarantee global connectivity. Based on directional information, a node grows its transmission power until it finds a neighbor node in each direction. Hu [12] gives an algorithm to choose logical links in a wireless network based on Delaunay triangulations.

3.2 Power Control and Medium Access Control

Power control affects the medium access control (MAC) layer protocols, especially contention-based MAC protocols such as the IEEE 802.11b MAC.

As the transmission range is reduced, nodes contending for the channel can send at a lower rate, and minimize the MAC layer contention [6]. Monks et al. [13] propose a power-controlled multiple access wireless MAC protocol (PCMA) which generalizes the existing collision avoidance protocols. In [14], a transmit power control (TPC) mechanism is proposed to address the tradeoff between the MAC TPC and the physical layer (PHY) transmission rate. When a node is transmitting, and based upon the data transmission status, a look-up table (constructed offline) is consulted and an optimal rate-power is obtained that would maximize the energy efficiency. When the load is light and given the bursty nature of traffic, a node may be better off using only PHY rate adaptation without TPC, however.

In [15], a power control protocol is proposed wherein RTS/CTS packets are transmitted at the highest power level and DATA/ACK are transmitted at a lower power level to save energy. Collisions resulting from link asymmetry are resolved by having the source nodes periodically transmit DATA packets at the highest power level so that nodes in the neighborhood are informed that the medium is busy.

3.3 Power Control and Interference

A wireless signal intended for a specific receiver may cause interference at other receivers, thus reducing the signal to noise ratio at these receivers. This reduces the receiver information processing capacity. Many iterative power control algorithms have been developed to achieve the minimal interference at non-involved nodes. Bambos et al. [8] develop a power control scheme which provides protection for links that are currently operational, that is, their signal-to-interference ratios (SIRs) are maintained above a certain threshold at all times. However, these kinds of protocols are deterministic in the sense that they require prior knowledge or perfect estimates of quantities such as the SIR. Motivated by the fact that these quantities are difficult to estimate, Ulukus and Yates [16] present a new power control algorithm that makes use of available

measurements, and then converges stochastically to the optimum power.

More recently, Gobriel et al. [17] study the tradeoff between the low transmission power and the high probability of collision per message arising from increasing the number of hops on the path from source to destination. They come to the conclusion that sending the data packet to the nearest neighbor is not always optimal. They do not, however, account for the required latency when selecting the transmission power level.

3.4 Power Control and Network Throughput

Network throughput can be defined as the number of bytes (or number of packets) delivered per second. Intuitively, the higher the transmission range, the higher the interference, and the more likely the packets will get dropped and retransmitted. In [7], the network throughput is shown to be, in some sense, inversely proportional to the transmission range of the nodes, and hence nodes should transmit at the lowest power possible. Selecting the optimal transmission range to maximize throughput is investigated in [18], but network connectivity is not considered in these studies.

3.5 Power Control and Routing

Power control significantly impacts multi-hop routing. Fig. 2 illustrates the routes when nodes are using high transmission powers versus those when nodes are using low transmission powers. Several recent studies have considered increasing node and network lifetime by using power-aware metrics for routing. Singh et al. [9] present five routing metrics based on power consumption at nodes. A dynamic power routing scheme is proposed in [19]. This scheme incorporates physical layer and link layer statistics to conserve power. Heinzelman et al. [20] propose a clustering-based routing protocol called LEACH (Low-Energy Adaptive Clustering Hierarchy) that utilizes randomized rotation of cluster heads to evenly distribute energy consumption among the nodes in the network.

In [21], two centralized algorithms are given for dynamically optimizing the network end-to-end delay in response to changing traffic conditions. The algorithms alter the network topology or add a new node. These algorithms, however, rely on the existence of a centralized node capable of collecting traffic information and deforming the entire network topology accordingly. The algorithms are also shown not to be highly effective in densely connected or lightly loaded networks. Tickoo et al. [22] argue in favor of variable-range transmission routing protocols. They provide an asymptotic measure of the average variable-range transmission and traffic capacity in wireless ad hoc networks. The analytic bound is used to prove that variable-range transmission based routing not only consumes less transmission power, but also increases the capacity of the network.

3.6 Power Control and Latency

The tradeoff between energy and delay has increasingly become the focus of many recent studies. Yu et al. [23] present algorithms that minimize the overall energy consumption of the data aggregation tree of a multiple source single-sink sensor network, subject to overall latency constraints. A load-adaptive power control mechanism is proposed in [24], where each node decides the transmission power level that guarantees a certain predefined signal to noise ratio (SNR) at the intended receiver. Their work, however, assumes the existence of a centralized scheduler capable of granting transmission requests. Requests are granted simultaneously when their intended destinations are sufficiently apart from the source of any other simultaneous transmission.

Since the transmission range affects the interference among nodes, which in turn affects the network throughput, recent papers have considered the throughput-delay tradeoffs. El Gamal et al. [25] derive asymptotic bounds for the optimal throughput-delay tradeoff in both static and mobile ad hoc networks. Three different schemes are presented that achieve such an optimal throughput-delay tradeoff. An upper bound on the optimal throughput-delay tradeoff over all scheduling policies is established in [26]. A new scheduling scheme is developed that achieves the upper bound, scaled down by a logarithmic factor.

4. LOAD-AWARE POWER CONTROL (LAPC)

In this section, we consider the latency versus energy efficiency tradeoffs encountered when making power control decisions. Table 1 compares low power transmission (when nodes transmit at the lowest possible power that guarantees connectivity), and high power transmission (when nodes transmit at the highest possible power). From the table, it is clear that sending at high power can reduce latency, but interference can cause retransmissions and delays.

Nodes in an ad-hoc network can use the same power level, or use different power levels. *Common transmission power* denotes the case when all nodes transmit at the same power level. This power level is often set to the lowest power level that guarantees network connectivity. In clustered networks where nodes are not homogeneously dispersed, common power is inefficient, in the sense that nodes inside the same cluster need a much lower power level to communicate with each other. Therefore, nodes need to transmit at variable power levels [5]. Table 2 compares common and variable transmission power schemes.

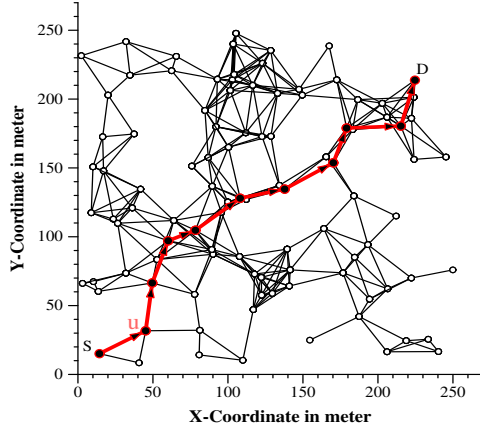
As previously mentioned, most prior work on power control optimizes energy consumption. In real-time applications or when an application is operating in delay-sensitive modes, however, end-to-end latency is of paramount importance. Responsiveness to sensed data in a sensor network, especially during critical or dangerous situations such as earthquakes or military battles, temporarily overrides energy concerns. Motivated by these scenarios, we consider the problem from a different perspective: can we reduce end-to-end latency by intelligent power control?

The answer to this question is a “conditional yes” if the node can obtain information about load at neighboring nodes. With the aid of this local information, a node can adjust its own power level to reduce latency. We refer to this heuristic approach as *Load-Aware Power Control (LAPC)*. The required load information can, for example, be piggy-backed on the periodic routing updates that each node receives from its neighbors to construct routing tables at different power levels (Fig. 2). Note that routing updates do not consume significant bandwidth, as discussed in [6]. Assuming each node broadcasts a routing update packet of 1000 bytes every 5 seconds¹ at each power level², the overhead per node per power level is $\frac{1000}{5} \times 8 = 1.6$ kbps. Assuming each node has 7 power levels, the total bandwidth consumed by the routing update packets per node is $7 \times 1.6 = 11.2$ kbps. Finally, assuming each node has, on the average, 6 contending neighbors, the total bandwidth consumed is $6 \times 11.2 = 67.2$ kbps. This value is approximately 0.61% of the theoretical bandwidth of today’s 802.11b wireless cards: 11 Mbps³.


¹This rate suffices to change approximately 10 routing table entries with 100 bytes each, which is reasonable for accommodating mobility and topology changes.

²Sending only at the highest power level will not give the receiver enough information to build the routing tables at each power level.

³In practice, the 802.11b bandwidth is much lower after accounting




(a) A graph when the range is 48 m



(b) A graph when the range is 88 m

Figure 2: Graphs for different ranges showing the smaller number of hops with larger ranges



(a) Accessible neighbors at different power levels

Transmission Power	Busy Nodes	Free Nodes	% Busy
P_1	1	2	33.33%
P_2	3	5	37.5%
P_3	7	9	43.75%

(b) Percentage of “busy” neighbors at different power levels

Figure 3: Increasing β allows higher transmission power

The load at neighbors of a node $v \in V$ gives a basic indication of whether interference with contention-based MAC protocols is likely, if v transmits at a certain power level. The neighbors of v at a certain power level P_i are its main contenders for bandwidth in the wireless medium. If none of them has data to transmit or receive, it is unlikely that the transmission of v will interfere with other transmissions. It is therefore reasonably safe for v to transmit at this power level, with the goal of reducing end-to-end latency.

We generalize this simple heuristic in a distributed algorithm, which we refer to as LAPC. In LAPC, a node can transmit at a power level P_i if the ratio of its “busy” neighbors to the total num-

ber for the MAC and PHY overhead. The routing updates will comprise 1.68% of bandwidth for 4 Mbps bandwidth, which is still quite small.

ber of neighbors accessible at this power level is at most a pre-specified ratio β . We refer to this ratio β as the *safety ratio*. By “busy,” we mean that the node is attempting to transmit or receive data at that time. When transmitting a packet, a node continues increasing its transmission power level as long as this level new satisfies the LAPC condition. As the ratio β increases, the fraction of “busy” neighbors at the next higher power level becomes more likely to satisfy the LAPC condition. Fig. 3 depicts an example. According to the percentage of “busy” neighbors given in this scenario, setting β to a value less than 33.33% implies that none of the three transmission power levels satisfies the LAPC condition and, hence, the sender refrains from sending data. If β exceeds 33.33%, the sender may start sending at P_1 . Increasing β beyond this value will not change the transmission power level until it reaches 37.5% when P_2 becomes a valid transmission power. Likewise, P_3 will become valid for transmission as β reaches 43.75%. Thereafter, the sender will keep on transmitting at P_3 as β approaches 100%. Hence, increasing β gives a node more *freedom* to select a higher transmission power level. Fig. 4 gives three possible scenarios for LAPC at different values of β . The figure shows how a node determines at which power level it will transmit, when the percentages of “busy” neighbors at the first three power levels are as shown in Fig. 3.

Observe that we set β as a percentage of neighbors, rather than as a fixed number. Fixing β to a certain number of “busy” neighbors, say m , would impose an additional burden on nodes that have a smaller number of neighbors than m . The number of “busy” neighbors of these nodes would always meet the LAPC criterion and, therefore, these nodes would always transmit at the highest power level. This causes them to deplete their energy quite rapidly. In contrast, nodes that have a larger number of neighbors than m are less vulnerable, since they typically have many “busy” neighbors. Therefore, such nodes may not increase their transmission power level. Nodes in dense networks have to always contend with a significant number of nodes for the medium. Therefore, we select the safety ratio β as a percentage of the neighbors at a certain power level.

We define the empirically optimal safety ratio β^* as the empirical value of β that gives the minimum end-to-end latency for a particular scenario. The value of β^* can be determined by monitoring the end-to-end latency encountered by the packets as the safety

Table 1: Effect of Transmission Power

	<i>Low power</i>	<i>High power</i>
<i>Network lifetime</i>	Reduces energy consumption and consequently increases the network lifetime.	Higher energy consumption leads to shorter network life.
<i>Interference</i>	Lower interference due to lower density.	Higher interference.
<i>Network traffic carrying capacity (throughput)</i>	Higher capacity due to the lower interference.	Higher interference reduces the capacity.
<i>Packets loss rate</i>	Lower loss rate.	Higher loss rate due to higher interference.
<i>Network sparsity</i>	Sparse network.	Dense network.
<i>Network connectivity</i>	The likelihood that the network is connected shrinks as the power is lowered.	High power increases network connectivity by adding more links.
<i>Routing</i>	Low overhead in constructing routing tables due to the reduced network density. Low power levels are commensurate with power optimal routing.	Higher overhead in constructing routing tables due to the availability of more routing choices.
<i>Network diameter</i>	A smaller number of links leads to larger diameter.	Smaller network diameter due to increased node degrees.
<i>Number of hops to destination</i>	Larger number of hops to the destination.	Fewer hops.
<i>End-to-end latency</i>	Lower at high network loads.	Lower at low network loads. Interference with high load imposes longer delays.

Table 2: Common Transmission Power Versus Variable Transmission Power

	<i>Common power</i>	<i>Variable power</i>
<i>Nodes dispersion</i>	Best when nodes are homogeneously dispersed.	Best when nodes are clustered.
<i>MAC layer</i>	Having nodes transmit at the same power ensures that if a node u is in the range of a node v , then v is in the range of u . In other words, both nodes can hear each other.	Problematic in the sense that if u hears v , this does not necessarily imply that v can hear u .
<i>Network traffic carrying capacity (throughput)</i>	Per node throughput of $O(\frac{1}{\sqrt{n \log n}})$ [27].	Per node throughput of $O(\frac{1}{\sqrt{n}})$.
<i>Energy consumption</i>	Some nodes may consume more energy than necessary to deliver data to the next hop.	Optimal in the sense that a node can consume only enough energy to carry its data to the next hop.
<i>Routing</i>	Common power ensures the bidirectionality of links, which is implicitly assumed in most distributed routing algorithms (e.g., Bellman-Ford).	The fact that links are not bidirectional imposes restrictions on routing algorithms.

ratio β increases from 0 to 1.

5. PERFORMANCE EVALUATION

Using the general-purpose simulation package *CSIM*, we have constructed a simulation model for wireless ad hoc networks⁴. We implement the proposed transmission power control protocol LAPC, as well as the lowest transmission power control protocol. For the sake of comparison, results are also compared with the case when nodes are transmitting at the highest power level, which represents a worst case scenario. Our simulation model is used to study the impact of deploying power control techniques on the network lifetime and end-to-end latency. We simulate an arbitrary numbers of nodes uniformly or non-uniformly dispersed in a rectangular area.

We use a simple model for energy dissipation in our simulations. All nodes start with the same amount of energy. We consider the energy dissipated by electronic devices in transmitters and receivers, $E_{electronic}$. We also consider the energy dissipated by

⁴Our simulation code is available upon request.

amplifiers in transmitters that allow the transmitted signal to travel the required distance, $E_{amplifier}$.

We assume a free space channel model, where the power loss is proportional to r^2 (r is the transmission range). Thus, when handling a packet of length l bits, the energy dissipated by transmitters and receivers is:

$$\begin{aligned} E_{Xmit} &= E_{amplifier} + E_{electronic} = l \times (\eta \times r^2 + \zeta) \\ E_{Recv} &= E_{electronic} = l \times \zeta \end{aligned}$$

where η (*Joule/bit/m²*) is the energy consumed by the power amplifier to transmit 1 bit for a distance of 1 meter and ζ (*Joule/bit*) is the energy consumed by the electronics to handle 1 bit. A network is considered to be “alive” as long as none of its nodes has completely depleted its energy.

5.1 Parameters and Modeling Assumptions

Throughout each simulation run, variable length packets are generated in the network according to a Poisson distribution with pa-

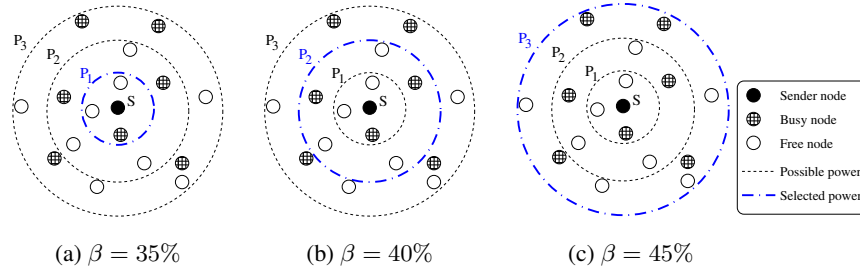


Figure 4: The LAPC protocol at different values of β

parameter λ (the arrival rate at *each node* is $\frac{\lambda}{n}$). Simulations are conducted for different values of λ and of β ($1 \leq \lambda \leq 10,000$ and $0 \leq \beta \leq 1$). Nodes are assumed to support 802.11b compliant wireless cards that have bandwidth capabilities up to 11 Mbps. The end-to-end latency is computed as the sum of transmission, propagation, and queuing delays. The transmission delay is computed according to the bandwidth of the wireless links and the packet length. The propagation delay is computed according to the distance between the source and destination. The queuing policy at each node is FIFO. Note that contention-based MAC protocols represent a worst case interference scenario. We expect the performance of LAPC to significantly improve with MAC protocols such as TDMA or CDMA that limit interference. A pseudo-random number generator is used to generate random values used for node locations, packet lengths, packet inter-arrival times, packet source and destination, and the “BackOffTime” used in the MAC protocol. Three different seeds are used for each set of parameters, and their results are averaged. Node mobility and changes in topology are assumed to occur at a relatively large time scale, and therefore are not taken into account (routing tables are static). Nodes run a simplified model of the 802.11b MAC protocol with typical parameters. Table 3 gives a summary of the simulation parameters.

Table 3: Summary of Simulation Parameters

Parameter	Value
Network area size	$250 \times 250 \text{ m}^2$
Number of nodes n	50, 100, 150, 200, and 5000
Transmission ranges	50, 60, \dots , 110
Packet length	Uniform [0,2000] bytes
<i>TIMESLOT</i>	50 μsec
<i>TIMEOUT</i>	0.1 sec
Contention window size CW	[32,128]
<i>DIFS</i>	128 μsec
Initial node energy	1 Joule
η	25 nJoule/bit/ m^2
ζ	40 nJoule/bit/ m^2

5.2 Estimating Load at Neighboring Nodes

Since periodic routing updates occur at relatively large time scales compared to packet transmission rate, a transmitting node might not have recent information about the load at its neighboring nodes. To mitigate this problem, we slowly decay load information. Every node keeps a history of n entries about each of its neighbors in the past n routing updates. Each entry x_{ij} equals 1 if the neighbor v_i is “busy” at time j ; otherwise $x_{ij} = 0$. When a node has a packet

to transmit, and since LAPC depends primarily on the load at the neighbors at that time, it uses this history to estimate if a neighbor is “busy.” A neighbor v_i is estimated to be “busy” with probability p , where p is the weighted average of $x_{ij} = \sum_{j=1}^n w_j x_{ij}$. The percentage of “busy” nodes is the number of neighbors that are estimated to be “busy” divided by the total number of neighbors. The weights w_j are assigned to favor the most recent information since $w_j = \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots, \frac{1}{2^{n-1}}, \frac{1}{2^{n-1}}$ from the most recent to the least. To conserve storage, a node should only maintain the latest status along with a value capturing the history for each neighbor. The weighted average of these two values gives an estimate of the probability p .

In summary, estimating the load at the neighbors is achieved by using the history of each neighbor at various time instants to forecast its status at the next instant. As a general rule from estimation theory, estimates can be misleading. However, Fig. 5 depicts the percentage of times that such load estimates were correct in our simulations. By correct, we mean that the neighbor was “busy” when the estimate indicated so. This percentage of times is referred to as the *hit ratio*. According to the figure, we were able to obtain accurate estimates up to 85% accuracy. This method works well because our arrivals are Poisson. More sophisticated methods for estimating load will be the subject of our future work.




Figure 5: Accuracy of estimated loads at neighbors

5.3 Simulation Results

5.3.1 Uniformly dispersed nodes

Fig. 6 depicts the measured end-to-end latency and network lifetime at different loads when nodes are uniformly dispersed. Figs. 6(a) and (b) show that higher power schemes perform best in terms of end-to-end latency at low loads. As the network load starts to increase, however, packets encounter higher queuing delays at intermediate nodes due to the increased contention with CSMA/CA. The resulting interference increases the likelihood of packet drop

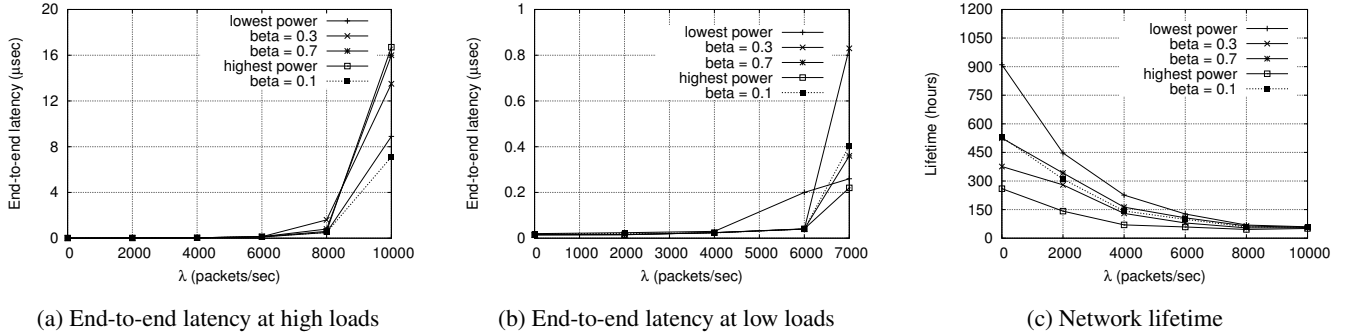


Figure 6: Simulation results at different system loads when nodes are uniformly dispersed

and retransmission. Consequently, the lowest power schemes generally outperform the highest power schemes at high network loads. Observe that $\beta = 0.1$ performs best in terms of end-to-end latency.

Fig. 6(c) shows that, as expected, the lowest power schemes outperform the highest power schemes in terms of power consumption and accordingly network lifetime. The load-aware power scheme lies between the highest power schemes and the lowest power schemes. This is intuitive since the higher the transmission power level a node selects, the sooner it will deplete its energy. The difference between all the schemes diminishes with increased load, however.

5.3.2 Non-uniformly dispersed nodes and dense networks

We also studied the cases when n nodes are non-uniformly dispersed (i.e., *clustered*) and when $n = 5000$ nodes are uniformly dispersed (dense networks). We do not give the results here due to space constraints. Our main findings are that packets are always subject to higher latencies in the non-uniform case due to the higher interference the packets encounter as a result of using higher transmission power to maintain connectivity. With our experimental setup, dense networks outperform both clustered and sparse networks in terms of the latency the packets encounter. This is attributed to the fact that we use the same system overall load λ , so the per-node load in dense networks is lower than that of sparse networks.

5.3.3 Empirically optimal safety ratio β^*

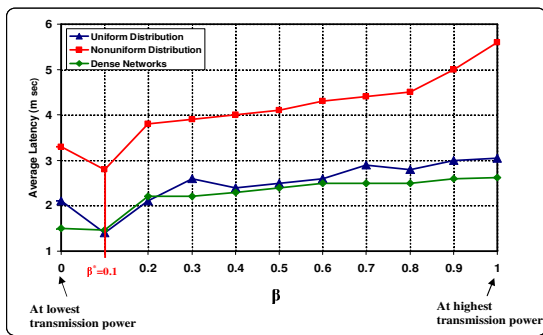


Figure 7: Average latency for different β (safety ratio) values

Fig. 7 depicts the averaged latency results for the 3 cases studied (uniform, non-uniform, dense), for each safety ratio β . These results are for traffic loads ranging between 1 and 10,000, in 1000

packets/sec increments. The figure illustrates that there is an empirical optimal value $\beta^* \simeq 10\%$ for minimizing latency under this range of loads and for this simulation scenario. This means that a node selects the highest transmission power level at which no more than 10% of its neighbors at this power level are “busy.” For $\beta < \beta^*$, nodes tend to send at low power levels and packets encounter a larger number of hops on their route to the destination. For $\beta > \beta^*$, nodes tend to transmit at high power levels and accordingly packets encounter a higher level of interference. Note that the value of β^* will typically decrease as the overall network load increases. Under light loads, it is advantageous to use a large β , as latency can be significantly reduced. This does not apply to heavily loaded networks with contention-based MAC protocols, where β should be small.

Table 4: LAPC Latency Improvement for Loads Up to 10,000 packets/sec

Improvement of LAPC over	Uniform dispersion	Non-uniform dispersion	Dense networks
Lowest power schemes	33%	18%	3%
Highest power schemes	54%	50%	44%

In addition, the figure shows that, when nodes are uniformly dispersed there is a relative improvement in terms of the average latency of approximately 54% when nodes are transmitting using a safety ratio of $\beta^* = 10\%$ over the case when nodes are transmitting using the safety ratio of $\beta = 99\%$. Table 4 summarizes the relative latency improvement of LAPC over low and high transmission power schemes for different network types.

6. CONCLUSIONS AND FUTURE WORK

We have studied the latency versus energy efficiency tradeoffs in wireless ad-hoc networks. Based upon our observations, we have proposed a heuristic for local power control, LAPC, for delay-sensitive applications or delay-sensitive modes of applications, e.g., sensor applications. LAPC aims at reducing the latency encountered by packets in the network by adjusting the power level up or down based upon local (neighborhood) load information. Simulation results show that the best performance in terms of end-to-end latency for the range of loads studied is achieved when nodes transmit with the highest power level at which no more than 10% of their neighboring nodes are contending for the medium.

We plan to extend our work by: (1) conducting more extensive analysis, and testbed measurements. This includes validating our conjecture that LAPC performs well with non-contention-based MACs, and evaluating energy consumption and latency in this case. This also includes analytically deriving bounds on the latency under various loads; (2) investigating possible 801.11 MAC layer protocol improvements to overcome the variable power control problems discussed in table 2. These problems essentially arise from the fact that nodes are using different transmission powers, and if one node hears the other, that does not necessarily imply the reverse; (3) taking node mobility and topology changes into consideration; (4) investigating metrics that can be locally propagated for nodes to estimate loads at their neighbors; and (5) evaluating when power level should be based upon the proximity of the next hop destination. In this case, schemes such as COMPOW [6] are not efficient, but schemes such as CLUSTERPOW [5] are rather complex. We will explore how LAPC can take destination proximity into account when making power control decisions.

7. REFERENCES

- [1] "Cisco Systems," <http://www.cisco.com/>, 2005.
- [2] "CrossBow," <http://www.xbow.com/>, 2005.
- [3] R. Ramanathan and R. Rosales-Hain, "Topology control of multihop wireless networks using transmit power adjustment," in *Proceedings of IEEE INFOCOM*, 2000, pp. 404–413.
- [4] T. A. ElBatt, S. V. Krishnamurthy, D. Connors, and S. Dao, "Power management for throughput enhancement in wireless ad-hoc networks," in *IEEE International Conference on Communications*, 2000, pp. 1506–1513.
- [5] V. Kawadia and P. R. Kumar, "Power control and clustering in ad hoc networks," in *Proceedings of IEEE INFOCOM*, 2003.
- [6] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol," *European Wireless Conference*, 2002.
- [7] V. Kawadia, S. Narayanaswamy, R. Rosovsky, R. S. Sreenivas, and P. R. Kumar, "Protocols for media access control and power control in wireless networks," in *Proceedings of 40th IEEE Conference on Decision and Control*, 2001, pp. 1935–1940.
- [8] N. Bambos, S. Chen, and G. Pottie, "Radio link admission algorithms for wireless networks with power control and active link quality protection," in *Proceedings of IEEE INFOCOM*, 1995, pp. 97–104.
- [9] S. Singh, M. Woo, and C. S. Raghavendra, "Power-aware routing in mobile ad hoc networks," in *Proceedings of ACM MOBICOM*, 1998, pp. 181–190.
- [10] Piyush Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications*, 1998.
- [11] R. Wattenhofer, L. Li, P. Bahl, and Y-M. Wang, "Distributed topology control for power efficient operation in multihop wireless ad hoc networks," in *Proceedings of IEEE INFOCOM*, 2001, pp. 1388–1397.
- [12] L. Hu, "Topology control for multihop packet radio networks," in *IEEE Transactions on Communications*, 1993, pp. 1482–1493.
- [13] J. P. Monks, V. Bhargavan, and W.-M. Hwu, "A power controlled multihop access protocol for wireless packet networks," in *Proceedings of IEEE INFOCOM*, 2001, pp. 219–228.
- [14] D. Qiao, S. Choi, A. Jain, and K. G. Shin, "Miser: An optimal low-energy transmission strategy for ieee 802.11a/h," in *Proceeding of the 9th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, September 2003, pp. 161–175.
- [15] Eun-Sun Jung and Nitin Vaidya, "A power control mac protocol for ad hoc networks," in *ACM MOBICOM*, 2002.
- [16] S. Ulukus and R. Yates, "Stochastic power control for cellular radio systems," in *IEEE Transaction on Communications*, 1998, pp. 784–798.
- [17] S. Gobriel, R. Melhem, and D. Mosse, "A unified interference/collision analysis for power-aware adhoc networks," in *Proceedings of IEEE INFOCOM*, March 2004.
- [18] H. Takagi and L. Kleinrock, "Optimal transmission ranges for randomly distributed packet radio terminals," in *IEEE Transactions on Communications*, 1984.
- [19] M. W. Subbarao, "Dynamic power-conscious routing for MANETs: An initial approach," in *IEEE Vehicular Technology Conference*, 1999, pp. 1232–1237.
- [20] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan, "Energy-efficient communication protocol for wireless microsensor networks," in *Proceedings of the IEEE Hawaii International Conference on System Sciences*, 2000, pp. 4–7.
- [21] A. Basu, B. Boshes, S. Mukherjee, and S. Ramanathan, "Network deformation: Reducing end-to-end delay in multihop wireless networks," in *Proceedings of the ACM 10th Annual International Conference on Mobile Computing and Networking (MobiCom 2004)*, 2004.
- [22] J. Gomez and A. Campbell, "A case for variable-range transmission power control in wireless multihop networks," in *Proceedings of IEEE INFOCOM*, March 2004.
- [23] Y. Yu, B. Krishnamachari, and V. Prasanna, "Energy-latency tradeoffs for data gathering in wireless sensor networks," in *Proceedings of IEEE INFOCOM*, March 2004.
- [24] I. Rubin, R. Khalaf, A. Moshfegh, and A. Behzad, "Delay-throughput performance of load-adaptive power-controlled multihop wireless networks with scheduled transmissions," in *Proceedings of the 58th IEEE Semiannual Vehicular Technology Conference (VTC)*, October 2003.
- [25] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah, "Throughput-delay trade-off in wireless networks," in *Proceedings on IEEE Infocom 2004*, March 2004.
- [26] X. Lin and N. B. Shroff, "On the fundamental relationship between achievable capacity and delay in mobile wireless networks," in *National Science Foundation Workshop on Pervasive Computing and Networking*, April 2004.
- [27] Piyush Gupta and P. R. Kumar, "The capacity of wireless networks," in *IEEE Transactions on Information Theory*, 2000, vol. IT-46, pp. 388–404.