Wireless Link Scheduling under a Graded SINR Interference Model*

Paolo Santi
IIT-CNR
Pisa, Italy
paolo.santi@iit.cnr.it

Ritesh Maheshwari
Dept. of Computer Science
Stony Brook University
ritesh@cs.sunysb.edu

Giovanni Resta
IIT-CNR
Pisa, Italy
g.resta@iit.cnr.it

Samir Das
Dept. of Computer Science
Stony Brook University
samir@cs.sunysb.edu

Douglas M. Blough
School of ECE
Georgia Tech
doug.blough@ece.gatech.edu

ABSTRACT

In this paper, we revisit the wireless link scheduling problem under a graded version of the SINR interference model. Unlike the traditional thresholded version of the SINR model, the graded SINR model allows use of "imperfect links", where communication is still possible, although with degraded performance (in terms of data rate or PRR). Throughput benefits when graded SINR model is used instead of thresholded SINR model to schedule transmissions have recently been shown in an experimental testbed. Here, we formally define the wireless link scheduling problem under the graded SINR model, where we impose an additional constraint on the minimum quality of the usable links, (expressed as an SNR threshold β_Q). Then, we present an approximation algorithm for this problem, which is shown to be within a constant factor from optimal. We also present a more practical greedy algorithm, whose performance bounds are not known, but which is shown through simulation to have much better average performance than the approximation algorithm. Furthermore, we investigate, through both simulation and implementation on an experimental testbed, the tradeoff between the minimum link quality threshold β_Q and the resulting network throughput.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Networks

General Terms

Algorithms, Experimentation, Performance, Theory

*The work of D. Blough and S. Das was supported by NSF Grant CNS–0721596.

FOWANC'09, May 18, 2009, New Orleans, Louisiana, USA.

1. INTRODUCTION

Ever since Gupta and Kumar's classical result [10] showing that the capacity of multi-hop wireless networks does not scale linearly with the number of nodes, researchers have studied a multitude of ways to increase throughput in such networks. Many of the considered approaches focus on increasing the concurrency of communications in the wireless medium, by separating communications either in frequency or in space. Concurrent communications can be separated in frequency by using multiple channels, with or without multiple radios. A variety of techniques exist for improving spatial separation, or spatial reuse, within the wireless channel. These techniques apply different methods for reducing or eliminating interference. For example, directional antennas focus the signal in a certain direction, thereby preventing a transmission from interfering with other communications outside of the focused area. Transmission power control can reduce the overall "inteference footprint" of a communication. More recently, MIMO technology has been considered, both as a means to improve throughput on individual links and for its ability to suppress interference, thereby permitting increased spatial reuse.

Another research focus has been TDMA approaches for multi-hop wireless networks [14]. TDMA has been considered for its ability both to improve throughput and to provide fairness to flows of differing lengths, which has been shown to be a significant problem in CSMA/CA-based wireless multi-hop networks. TDMA has been adopted for use in the IEEE 802.16 standard for WiMax [1]. With use of TDMA, comes the opportunity to develop scheduling algorithms that can carefully separate communications in space, thereby maximizing concurrency and presumably throughput, as well.

When considering the scheduling of transmissions in a multi-hop wireless network, it is necessary to model interference. Over time, the research in wireless scheduling has considered more accurate interference models. Early work used a simple k-hop interference model, while later work employed more accuracte distance-based models. In the last few years, several papers [3, 4, 9] have considered transmission scheduling under more accurate physical interference models, which are based on signal to interference and noise ratio (SINR) at the receiver. However, all of this recent

work has considered that packets are successfully received only when SINR exceeds a given threshold, and assumes that packet reception rate (PRR) is zero below this threshold. In reality, PRR falls off gradually with decreasing SINR. This phenomenon has been well documented in the literature, and the SINR region corresponding to imperfect (but considerably greater than 0) reception rates is known as *transitional region*, or *gray region* [15, 16, 24]. In this paper, we formulate a graded physical interference model, which accounts for this more accurate relationship between PRR and SINR, and we investigate the question of whether scheduling algorithms can effectively use links that operate below the SINR threshold in order to increase spatial reuse and thereby improve throughput.

To the best of our knowledge, non-thresholded SINR-based interference models have been seldom used in the wireless scheduling literature, with a few notable exceptions [7, 8]. However, the emphasis in [7, 8] is on jointly optimizing routing, scheduling, and transmit power in order to minimize the total average transmit power, given some constraints on the minimum data rate achieved on each link. Furthermore, the approach of [7] is based on convex programming, which has exponential time complexity, while that of [8] is based on solving a complex fixed point equation.

In addition to developing a graded SINR model of packet reception, we consider the design of scheduling algorithms that take advantage of this more accurate model. We present a scheduling algorithm, $GradedSINR$, and we prove that this algorithm is within a constant factor of optimal in terms of the length of the schedules it produces. We then turn to evaluation of the graded SINR model and associated scheduling algorithms in practical settings. We present both simulation-based results and an experimental evaluation carried out using a Mote-based testbed. Simulation results demonstrate that *throughput increases of up to 50%* are possible relative to thresholded SINR models. The throughput increase achieved is dependent upon node density, with more improvement seen for sparse networks, where additional opportunities for spatial reuse can have more impact. However, even in dense networks, throughput improvements of almost 20% were achieved. Results from the Mote-based testbed confirm that use of lower-quality links can improve throughput. In fact, even greater throughput increases, up to 70% improvement over using only 100% quality links, were seen in the testbed. This proof-of-concept implementation also demonstrates the practicality of our approach. Thus, we believe that graded SINR-based scheduling algorithms hold great promise for dramatically improving performance of TDMA-based multi-hop wireless networks.

2. THE GRADED SINR MODEL

The graded SINR model is motivated by the observation that, in a practical scenario, the packet reception rate (PRR) vs. SINR is not sharply thresholded, but rather presents a smooth transition between close to 0 and close to 1 reception rate. The region in which packet reception is not perfect is known as the *transitional region*, or *gray region* in the literature [15, 16, 24], and typically spans 5 to 10 dBs.

In graded SINR models, originally proposed in [15, 16], the PRR achievable on a certain link is a function of the SINR value experienced at the receiver. The PRR vs. SINR curve has the following properties in these models: *i*) the PRR is 0 when the SINR is below a certain value, which we denote β_0 in the following; *ii*) the PRR is 1 when the SINR is above a certain value, which we denote β_1 in the following; *iii*) the PRR is an *increasing* function of the SINR in the transitional region. We adopt this model in this paper, with the further requirement (needed for technical purposes) that $\lim_{x\to\beta_0^+} f(x) = 0$ and $\lim_{x\to\beta_1^-} f(x) = 1$, where $f()$ is the function representing the PRR vs. SINR curve in the $[\beta_0, \beta_1]$ interval. We also assume in the following that $\beta_0 > 0$, which is perfectly reasonable in a realistic scenario (note that SNR and SINR are expressed as linear ratios, not in dB). An example of such a function when f is a linear function is shown in Figure 1. For clarity of presentation, in the following we extend function $f()$ as follows: $f(x) = 0$ if $x \leq \beta_0$, and $f(x) = 1$ if $x \geq \beta_1$.

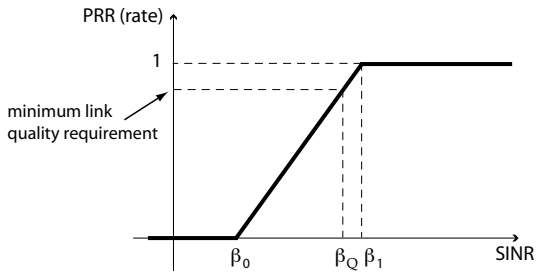

Figure 1: The graded SINR model. The *rate* is intended as normalized w.r.t. maximal possible rate W_{max}.

In the following, we denote a communication link as $l_i = (s_i, r_i)$, where s_i is the sender and r_i is the receiver node. According to our model, the PRR experienced on link l_i, *in the absence of interference,* is given by $f(SNR_i)$, where SNR_i is the signal-to-noise ratio at node r_i. Formally, $SNR_i = \frac{P_i}{N}$, where P_i is the received power at node r_i of the signal transmitted by node s_i, and N is the background noise power.

In presence of multiple concurrent transmissions on links $l_1, \ldots, l_k$, the PRR on link $l_i = (s_i, r_i)$ is given by $f(SINR_i)$, where $SINR_i$ is the signal-to-noise-and-interference ratio measured at r_i when all the s_js are transmitting. Formally,

$$SINR_i = \frac{P_i}{N + \sum_{j \neq i} P_j} \ ,$$

where P_j denotes the received power at node r_i of the signal transmitted by node s_j, for each $j \neq i$.

It is worth observing the similarities between the graded SINR model and the generalized physical interference model (see, e.g., [13]), according to which the *data rate* W_i observed on link l_i is given by Shannon's channel capacity formula, i.e.,

$$W_i = B \log_2 (1 + SINR_i) \ , \qquad (1)$$

where B is the channel bandwidth[1]. The graded SINR interference model introduced above can be interpreted in terms of data rate as follows. Assume the channel has a maximal nominal data rate W_{max}. We can interpret the PRR vs. SINR curve as a data rate vs. SINR curve. The idea is that, when the SINR value is below the minimum threshold β_1 required for successful transmission of a packet at rate

[1]Note that the SINR in the formula for W_i is expressed as a linear ratio, not dBs.

W_{max}, PHY layer parameters such as coding (e.g., increasing bit redundancy in packet transmission) and/or symbol sending rate are modified, so that packets can be successfully received at r_i. Hence, we can view the situation as if packets are always correctly received when transmitted on link l_i, but with the achieved data rate depending on the experienced SINR value at r_i. Given this interpretation, the main difference between the graded SINR model and the generalized physical interference model is that actual data rate on link l_i is 0 under the graded SINR model when $SINR_i \leq \beta_0$, while it is always greater than 0 with a positive received signal power under the generalized physical interference model.

Unless otherwise stated, in the following we will use the data rate interpretation of the graded SINR model, since it eases the derivation of clean approximation bounds for the considered scheduling problem. In order to keep the values of $f()$ in the $[0, 1]$ interval, we will interpret function $f()$ as giving, for a certain SINR value, the resulting data rate normalized with respect to the maximal nominal data rate W_{max}. Hence, the actual data rate on link l_i with SINR value $SINR_i$ will be $f(SINR_i) \cdot W_{max}$. To simplify notation and when clear from the context, in the following we will sometimes overload the $f(SINR_i)$ notation to denote the data rate on link l_i. Since the data rate interpretation of the graded SINR model assumes accurate control of PHY layer parameters is possible, which is not always the case in practical scenarios, in the experimental setup reported in Section 5 we have used the original, PRR-based interpretation of the graded SINR model.

3. SCHEDULING ALGORITHM

3.1 Problem Formulation

The problem we consider is often referred to as the *wireless link scheduling problem* in the literature, although we are introducing here an extension to deal with the case of imperfect link transmissions (or, equivalently, flexible data rates). We are given a set of links $L = \{l_1, \ldots, l_n\}$ to schedule, with $l_i = (s_i, r_i)$. Note that the s_is and r_is are not necessarily distinct, i.e., a single node can be involved in multiple communication as either transmitter, or receiver, or both.

A link l_i is assigned a weight d_i, which represents the current traffic demand on link l_i. To simplify presentation, in the following we assume unit link demands (i.e., $d_i = 1$ for $i = 1, \ldots, n$). However, the presented results can easily be extended to the case of arbitrary integer demands by replacing each link l_i with d_i copies of the same link.

Links experience different SNRs at the receiver end; i.e., when only transmitter s_i is transmitting, node r_i will experience a certain SNR value SNR_i, which is in general different for different receivers. However, in the following we assume the $SNR_i \geq \beta_Q$ for each i, for some constant value $\beta_Q \geq \beta_0$. Given our model, this is equivalent to assuming that data rate (or, equivalently, PRR) on each link is at least k, for some constant $k > 0$ (see Figure 1). The SNR lower bound β_Q on the links to schedule is introduced to reflect the fact that, in practical situations, only relatively high quality links (i.e., with acceptable data rate, or with a PRR considerably above 0) are used to transmit packets.

The problem we consider is to schedule links in set L in such a way that: *i*) the demands on each link are satisfied, and *ii*) the length of the schedule is minimum. Note that, with respect to classical scheduling problems with non-graded interference models (e.g., with the physical interference model [3, 9]), *we do not impose any feasibility constraint on the schedule*. This is because under the graded SINR model every transmission set is feasible. What changes is the data rate (equivalently, PRR) experienced on each link, which is dictated by the graded SINR model, and can actually be 0 on some of the links. Feasibility of the schedule is, in a sense, captured by condition *i*), which states that demands on each link must be satisfied. This implies that, if we define a unit of time as the time needed to send a unit of demand (packet) from transmitter to the intended receiver at rate W_{max}[2], the total time allocated for transmission on link l_i in the schedule must be sufficient to send the packet to destination at the achieved data rate. To be specific, if link l_i is scheduled for time intervals $t_1, \ldots, t_h$, and the experienced SINR values at r_i during these intervals are $SINR_{i,1}, \ldots, SINR_{i,h}$, we must have $\sum_{j=1,h} f(SINR_{i,j}) \cdot W_{max} \cdot t_j \geq S$, where S is the packet size in bits.

The computational complexity of the link scheduling problem has been studied under different interference models, and has been proven to be NP-hard in many cases, e.g. the physical interference model [9] and most hop-based interference models [22]. While similarities of the graded SINR model with the physical interference model suggest that the problem might remain NP-hard also in the graded SINR model, a formal proof of this fact is beyond the scope of this paper and, to the best of our knowledge, the problem remains open.

In the following, we present an algorithm for this problem and prove a worst-case bound on its performance with respect to performance of an optimal scheduling algorithm. In order to prove the approximation bound, we adopt the classical model for radio signal propagation in wireless networks, which is referred to as the *log-distance path loss model*. In this model, the radio signal strength (power) at a distance d from the transmitter is given by P/d^α, where P is the transmission power and $\alpha > 2$ is the path loss coefficient [20] (the actual value of the constant α depends on the environment – e.g., indoor or outdoor). Our results should be easily extensible to more general radio propagation models that account for irregular radio coverage area, such as the cost-based model proposed in [21], which approximates log-normal shadowing propagation. In the following, we assume all nodes use the same transmit power, an arbitrary constant P.

3.2 Algorithm GradedSINR

Algorithm $GradedSINR$, which is reported in Figure 2, is based on the simple idea of grouping links with similar SNR values in the same class, and scheduling them in consecutive slots. Link classes are defined as follows: link class C_k, with $k = 1, \ldots, \bar{k}$, contain links l_js such that

$$(1+\epsilon)^{k-1}\beta_Q \leq SNR_j < (1+\epsilon)^k \beta_Q \ , \qquad (2)$$

where ϵ is an arbitrary constant greater than or equal to $1/7$ and $\bar{k} = \lfloor \log_{1+\epsilon}(P/\beta_Q N) \rfloor + 1$. Note that: *i*) all links belong

[2]Note that this time in practice depends also on the transmitter/receiver separation. Hence, time unit can be interpreted as the *maximum* over set L of the time needed to send a packet from transmitter to receiver.

Algorithm *GradedSINR*:

Input: A set L of n links with unit demand
Output: A schedule $S_1, \ldots, S_{\bar{k}}$ under graded SINR model

1. $t = 1$
2. Let $C = \{C_1, \ldots, C_{\lfloor \log_{1+\epsilon}(P/\beta N)\rfloor = \bar{k}}\}$ be link classes defined as in (2)
3. **for each** $C_k \neq \emptyset$, with $1 \leq k \leq \bar{k}$
4. Partition network deployment region into squares of width $\mu_k \cdot D_{k+1}$
5. 4-color the squares such that no two adjacent squares have the same color
6. **for** $j = 1, \ldots, 4$
7. Select color j
8. **repeat**
9. For each square A of color j, choose a link $l_i \in C_k$ with receiver in A; $L_j^k = L_j^k \cup \{l_i\}$
10. $t = t + 1$; $S_t = L_j^k$
11. set duration of slot S_t to $1/f((1+\epsilon)^{k-2}\beta_Q)$
12. **until** all links of C_k in selected squares are scheduled
13. **return** $S_1, \ldots, S_{\bar{k}}$

Figure 2: The *GradedSINR* Algorithm.

to one of the C_ks, since the minimum SNR value of links is β_Q, and the maximum SNR value is P/N;[3] *ii*) the number $\bar{k}$ of link classes is a *constant*, i.e., it does not depend on the number n of links to schedule.

Note that, under our working assumption of log-distance radio propagation with path loss exponent $\alpha > 2$, links in the k-th SNR class have length

$$D_{k+1} = \left(\frac{P}{(1+\epsilon)^k \beta_Q N}\right)^{\frac{1}{\alpha}} < L_k \leq \left(\frac{P}{(1+\epsilon)^{k-1}\beta_Q N}\right)^{\frac{1}{\alpha}} = D_k \,.$$

When considering links in class $1 \leq k \leq \bar{k}$, the deployment region is divided into square cells of side $\mu_k D_{k+1}$, where constant μ_k is defined as follows:

$$\mu = 2\left(\frac{64(1+\epsilon)^{k-1}\beta_Q(\alpha-1)}{\alpha-2}\right)^{\frac{1}{\alpha}} .$$

Cells in the same class are then 4-colored in such a way that no two adjacent cells have the same color. Then, at Steps 6–12 links are greedily scheduled in successive slots, with the property that only links with the same color whose receivers are in different cells are assigned to the same slot.

At Step 11, the duration of slots whose links are in class k is set to $1/f((1+\epsilon)^{k-2}\beta_Q)$, which, as shown in the following, is sufficient to send a unit of demand along the scheduled links. In fact, cell dimensioning is such that, under the hypothesis fulfilled by *GradedSINR* that no two links with receivers in the same cell of color j are scheduled concurrently, the minimum SINR value at each scheduled receiver is at least $(1+\epsilon)^{k-2}\beta_Q$.

We now formally prove that the schedule computed by *GradedSINR* satisfies the traffic demands of all links in L.

THEOREM 1. *Assume that* $\frac{1}{7} \leq \epsilon \leq 63$ *and* $\beta_Q \geq 1$. *Then, the schedule computed by Algorithm GradedSINR satisfies the traffic demands of all links in* L.

[3]By fundamental laws of physics, the received signal power can be at most as large as the transmitted power.

PROOF. Let us consider a slot containing links in class C_k, for some $1 \leq k \leq \bar{k}$. We now upper bound the interference experienced by a receiver r in a certain cell $\mathcal{C}$ in the partitioning obtained for class C_k. Once we focus on a receiver r_i in specific cell $\mathcal{C}$, the cells containing receivers of the interfering links can be arranged in circumcentric square frames around $\mathcal{C}$. The inner frame contains $3^2 - 1^2 = 8$ cells, the second frame contains $5^2 - 3^2 = 16$ cells, and in general the h-th frame will contain $(2h+1)^2 - (2h-1)^2 = 8 \cdot h$ cells. The generic receiver contained in the h-th frame will be at least $(2h-1)\mu_k D_{k+1}$ apart from r_i. Considering that in class k all links have a length smaller that D_k, the minimum distance between r_i and a *sender* relative to frame h is $(2h-1)\mu_k D_{k+1} - D_k = (2h-1)\mu_k D_k/(1+\epsilon)^{1/\alpha} - D_k = D_k((2h-1)\mu_k(1+\epsilon)^{-1/\alpha} - 1)$. Hence, the total interference I_r experienced by r_i can be upper bounded by

$$\begin{aligned}
I_r &< \sum_{h=1}^{\infty} \frac{8h \cdot P}{D_k^\alpha \cdot ((2h-1)\mu_k(1+\epsilon)^{-1/\alpha} - 1)^\alpha} && (3)\\
&\leq \frac{8P}{D_k^\alpha} \sum_{h=1}^{\infty} \frac{h}{(\frac{1}{2}(2h-1)\mu_k(1+\epsilon)^{-1/\alpha})^\alpha} && (4)\\
&= \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu_k^\alpha D_k^\alpha} \sum_{h=1}^{\infty} \frac{h}{(2h-1)^\alpha} && (5)\\
&\leq \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu_k^\alpha D_k^\alpha} \sum_{h=1}^{\infty} \frac{h}{(2h-h)^\alpha} && (6)\\
&= \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu_k^\alpha D_k^\alpha} \sum_{h=1}^{\infty} \frac{1}{h^{\alpha-1}} && (7)\\
&\leq \frac{8(1+\epsilon)P}{(1/2)^\alpha \mu_k^\alpha D_k^\alpha} \cdot \frac{\alpha-1}{\alpha-2} && (8)
\end{aligned}$$

where (4) follows because $x - 1 > x/2$ for $x > 2$ and indeed $(2h-1)\mu_k(1+\epsilon)^{-1/\alpha}$ is greater than 2 under the theorem assumptions, and (8) follows from a known bound on Riemann's zeta function.

The SINR for the receiver r_i can thus be bounded by

$$\begin{aligned}
SINR_i &\geq \frac{\frac{P}{D_k^\alpha}}{I_r + N} \geq \frac{\frac{P}{D_k^\alpha}}{\frac{8(1+\epsilon)P}{(1/2)^\alpha \mu_k^\alpha D_k^\alpha} \cdot \frac{\alpha-1}{\alpha-2} + N} = \\
&= \frac{\frac{P}{D_k^\alpha}}{\frac{P}{8(1+\epsilon)^{k-2}\beta_Q D_k^\alpha} + N} = \frac{(1+\epsilon)^{k-1}\beta_Q N}{\frac{(1+\epsilon)^{k-1}\beta_Q N}{8(1+\epsilon)^{k-2}\beta_Q} + N} = \\
&= \frac{(1+\epsilon)^{k-1}\beta_Q}{\frac{(1+\epsilon)}{8} + 1} = \frac{8 \cdot (1+\epsilon)}{(1+\epsilon)+8} \cdot (1+\epsilon)^{k-2}\beta_Q \geq \\
&\geq (1+\epsilon)^{k-2}\beta_Q \,, \qquad (9)
\end{aligned}$$

where (9) follows since $\epsilon \geq \frac{1}{7}$.

Since link l_i in class C_k, for some $1 \leq k \leq \bar{k}$, is scheduled in a slot of duration $1/f((1+\epsilon)^{k-2}\beta_Q)$, and the (normalized w.r.t. W_{max}) data rate on link l_i is at least $f((1+\epsilon)^{k-2}\beta_Q)$ (recall that $f()$ is an increasing function of SINR), we have that at least one unit of demand can be transmitted on link l_i in the scheduled slot, and the theorem follows. □

DEFINITION 1. *Given a set* L *of links to schedule, the* SNR density *for link class* C_k, *with* $1 \leq k \leq \bar{k}$, *is the maximal number of receivers in a cell of class* C_k, *and is denoted* Δ_k.

DEFINITION 2. *Given a set L of links to schedule, the* normalized SNR density *for L, denoted $\Psi(L)$, is defined as*

$$\Psi(L) = max_{1 \le k \le \bar{k}} \left\{ \frac{\Delta_k}{f((1+\epsilon)^{k-2}\beta_Q)} \right\} .$$

We now prove an upper bound on the length of the schedule computed by Algorithm $GradedSINR$.

THEOREM 2. *The schedule computed by Algorithm GradedSINR has $O(\Psi(L))$ length.*

PROOF. Links in class C_k, for $1 \le k \le \bar{k}$, whose receivers are in a cell of color, say, j, are scheduled in parallel if they are in different cells; hence, the number of slots needed to accommodate all links in class C_k is the number of colors (four) times the number of receivers in the maximally occupied cell, i.e., Δ_k. Since slot duration for links in class k is $1/f((1+\epsilon)^{k-2}\beta_Q)$, total schedule length is upper bounded by $\sum_{k=1}^{\bar{k}} 4 \cdot \frac{\Delta_k}{f((1+\epsilon)^{k-2}\beta_Q)} \le 4 \cdot \bar{k} \cdot \Psi(L) \in O(\Psi(L))$ since $\bar{k}$ is a constant. □

We are now ready to prove the approximation bound for Algorithm $GradedSINR$.

THEOREM 3. *Algorithm GradedSINR computes a schedule whose length is within a factor $O(1)$ from optimal.*

PROOF. Let us consider a link class $C_{\tilde{k}}$ for which the normalized SNR density $\Psi(L)$ is achieved, and let $L_{\tilde{k}} = l_1, \ldots, l_{\Delta_{\tilde{k}}}$ be links in class $C_{\tilde{k}}$ whose receivers are in a maximally occupied cell. Call this cell the *critical cell*. We lower bound the time needed to schedule links in $L_{\tilde{k}}$ only. Clearly, since the optimal schedule must accommodate a possibly larger set of links, the computed lower bound applies also to the optimal schedule for link set L.
We start by proving an upper bound on the number of feasible transmissions with receivers belonging to the critical cell, under the assumption that the feasible rate on the links is at least $f(\beta)$, for some $0 < \beta_0 < \beta < (1+\epsilon)^{\tilde{k}}\beta_Q$. Note that β must be greater than β_0 in order to have a non-zero data rate on the link, and that the maximum data rate of links in class $C_{\tilde{k}}$ is $< (1+\epsilon)^{\tilde{k}}\beta_Q$. In particular, we prove that no more than

$$q_{\tilde{k},\beta} = ((1+\epsilon)^{1/\alpha} + \sqrt{2}\mu_{\tilde{k}})^{\alpha} \cdot \frac{(1+\epsilon)^{\tilde{k}}\beta_Q - \beta}{\beta(1+\epsilon)^{\tilde{k}}\beta_Q}$$

such transmissions can occur in parallel. The value of $q_{\tilde{k},\beta}$ is obtained by solving the following inequality

$$\frac{\frac{P}{(D_{\tilde{k}+1})^{\alpha}}}{N + x \cdot \frac{P}{(\sqrt{2}\mu_{\tilde{k}} D_{\tilde{k}+1} + D_{\tilde{k}})^{\alpha}}} = \frac{(1+\epsilon)^{\tilde{k}}\beta_Q N}{N + x \cdot \frac{(1+\epsilon)^{\tilde{k}}\beta_Q N}{(\sqrt{2}\mu_{\tilde{k}} + (1+\epsilon)^{1/\alpha})^{\alpha}}} < \beta \tag{10}$$

which, after straightforward algebraic manipulation, leads to

$$x < ((1+\epsilon)^{1/\alpha} + \sqrt{2}\mu_{\tilde{k}})^{\alpha} \cdot \frac{(1+\epsilon)^{\tilde{k}}\beta_Q - \beta}{\beta(1+\epsilon)^{\tilde{k}}\beta_Q}$$

from which the above value of $q_{\tilde{k},\beta}$ is obtained. Inequality (10) comes from assuming the largest possible received power at the numerator, and the minimum possible contribution to interference from links whose receiver end is in the critical cell.

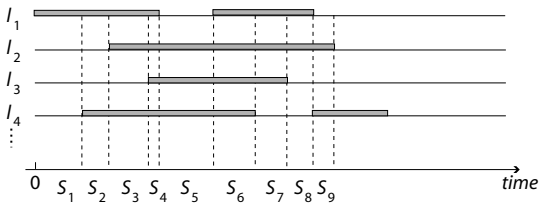

Figure 3: Example of possible link schedule under the graded SINR model. The data rate on, e.g., link l_1 is different in slot $S_1, S_2, S_3, S_4, S_6, S_7, S_8$ in which it is activated.

Let us consider the schedule computed by the optimal algorithm, and let $x > 0$ be the minimum data rate of a link in the optimal solution. Define $\bar{\beta}$ as the SINR value corresponding to data rate x according to function $f()$, i.e., $f(\bar{\beta}) = x$. Given the previous result, we have that at most $q_{\tilde{k},\bar{\beta}}$ links, each with rate $\ge x$, can be scheduled in parallel. The data rate on each of these links is at most $f((1+\epsilon)^{\tilde{k}}\beta_Q)$, since all the links in the critical cell belongs to class $C_{\tilde{k}}$. Since values $q_{\tilde{k},\beta}$ are a *decreasing* function of β, we have that the maximum demand that can be satisfied in a unit of time in the optimal schedule is $q_{\tilde{k},\bar{\beta}} \cdot f((1+\epsilon)^{\tilde{k}}\beta_Q)$. Since the total demand of links in the critical cell is $\Delta_{\tilde{k}}$, we have that the length of the optimal schedule is at least

$$\frac{\Delta_{\tilde{k}}}{q_{\tilde{k},\bar{\beta}} \cdot f((1+\epsilon)^{\tilde{k}}\beta_Q)} .$$

We now have that the ratio between the schedule length of the optimal solution and that of the schedule computed by $GradedSINR$ is

$$O\left(\frac{\Psi(L) \cdot q_{\tilde{k},\bar{\beta}} \cdot f((1+\epsilon)^{\tilde{k}}\beta_Q)}{\Delta_{\tilde{k}}} \right) =$$

$$= O\left(\frac{\Delta_{\tilde{k}}}{\Delta_{\tilde{k}}} \cdot \frac{q_{\tilde{k},\bar{\beta}} \cdot f((1+\epsilon)^{\tilde{k}}\beta_Q)}{f((1+\epsilon)^{k-2}\beta_Q)} \right) = O(1)$$

since function $f(x)$ has values in the interval $(0, 1]$ when $x > \beta_0$, and $(1+\epsilon)^{k-2}\beta_Q > \beta_0$. This concludes the proof of the theorem. □

Note the importance of the result stated in Theorem 3: under the graded SINR model, different transmission sets can be active at different times, possibly using flexible slots of very different time duration (see Figure 3). Hence, finding the optimal schedule in such a large set of possible solutions appears to be a very difficult task (although not yet formally proved to be NP-hard). Theorem 3 states that by imposing a strict structure on the schedule (all links of the same SNR class are scheduled in contiguous slots of fixed duration), we can still obtain a solution which is close to optimal (in asymptotic sense). This is especially important since, while general schedules allowed under the graded SINR model as the one depicted in Figure 3 can be difficult to realize in a practical setting (due to, e.g., required PHY layer parameter tuning while a packet is in the air), the well structured schedule computed by $GradedSINR$ can be implemented more easily in a practical setting.

4. SIMULATION-BASED EVALUATION

In this section we extensively evaluate the performance of scheduling algorithms based on the graded SINR model through simulation. The main goals of the evaluation are: (1) to identify the throughput maximizing configuration of the link quality threshold β_Q under different node density and topology/radio propagation scenarios, and (2) to quantify the potential throughput advantages of using the graded SINR model compared to a strict threshold-based SINR model. In view of (1), the interaction between scheduling and routing has to be considered: in fact, as the link quality threshold is varied, different sets of links are made available to the routing protocol and possibly used to route messages to the destinations. Hence, what specific routing protocol is used is an important choice that eventually determines the traffic load experienced on the available links.

In general, maximum throughput can be obtained only by jointly optimizing routing and scheduling, possibly exploiting multi-path routes (see, e.g., [2, 6]). However, joint routing and scheduling optimization under the graded SINR model is an open problem that is beyond the scope of this paper. Here, we are concerned with optimizing the scheduling step after a certain routing algorithm has been executed, and link demands generated. Hence, in our simulations, we will consider a simple (yet significant) routing algorithm coupled with a traffic generation method tailored to a wireless mesh network scenario, and use these two components to generate the link traffic demands given as input to the various scheduling algorithms considered.

4.1 Simulation setup

The simulation setup is tailored to a wireless mesh network scenario. A set of n nodes is deployed in a square region of side L. Two deployment methods are considered: grid-like, and uniform random. After node deployment, the $n \times n$ link matrix M is generated, where entry $m_{i,j}$ of the matrix represents the channel gain between transmitter node i and receiver node j. Channel gains are computed based on node positions, and on the radio propagation model. Radio signal propagation obeys log-normal shadowing, with path loss exponent α, for some $\alpha > 2$, and variance σ.[4] After the channel matrix is generated, a fixed fraction of the nodes (0.1) is selected as gateway nodes, according to a uniform random distribution. For each non-gateway node, a traffic demand is generated by randomly and uniformly choosing an integer in the interval $[1, 5]$. The generated traffic is directed to gateways, according to the following routing scheme. First, an *available link matrix* AM is obtained from M by retaining entries $m_{i,j}$ such that the SNR value at the receiver node j is at least β_Q, where β_Q is the desired link quality threshold. The other entries in matrix AM are set to 0, in order to prevent the routing algorithm from using the corresponding links. Using matrix AM, the routing algorithm then builds shortest path trees rooted at the gateway nodes to set up the routing paths. In case of ties, the gateway to which a specific node sends its traffic is selected uniformly at random. The link demands, which constitute the input to the scheduling algorithms, are then computed based on the node traffic demands and the chosen shortest path trees.

The metric used to build the shortest path trees is hop-count. Although very simple, this metric is used by most of the current routing algorithms for wireless multi-hop networks (e.g., DSR [12] and AODV [19]). Furthermore, when coupled with a link quality criterion, using minimal hop routes tends to reduce the total demand on the links, while only marginally sacrificing link throughput (if the link quality threshold is relatively high). For this reason, we believe shortest path routing based on hop-count is a reasonable heuristic to achieve a relatively high network throughput.

When using the graded SINR model, function $f()$ dictating the SINR (in dB) vs. link data rate relationship is defined as follows: $f(x) = 0$ if $x \leq \beta_0 = 10dB$, $f(x) = 1$ if $x \geq \beta_1 = 25dB$, and $f()$ linearly varies between these two values for $\beta_0 \leq x \leq \beta_1$. This setting is coherent with the SINR vs. PRR measurements for WLAN environments reported in [16], as well as with Shannon's capacity formula for intermediate SINR values[5]. We recall that the data rates returned by function $f()$ are normalized with respect to the maximum nominal bit rate of the link, set to $55Mbs$ in our experiments.

4.2 Simulated scheduling algorithms

In addition to Algorithm $GradedSINR$, we have also implemented Algorithm *I-GradedSINR*, which is an optimized version of *GradedSINR*, as well as a greedy algorithm called *GreedyGraded*. We do not give details of *GradedSINR* due to length limitations. *GreedyGraded* is inspired by the algorithm used in [16] to evaluate throughput in the WLAN experimental testbed. More specifically, *GreedyGraded* orders links randomly, and considers them sequentially. When a specific link l has to be scheduled, the currently formed slots are scanned, and, for each of them, the duration of the slot if link l were to be added is computed. Similarly to *I-GradedSINR*, the duration of a slot is set to the minimum value needed to transmit a packet along all active links and, hence, is determined by the SINR value of the weakest active link. Note then that the duration of the slot if l were to be added is in general longer than that of the original slot, since adding l to the slot would degrade SINR values (and, consequently, data rates) at the receiver nodes. Let $S(l)$ be the currently formed slot such that adding l to the slot increases slot duration of the minimal amount of time $T(l)$. The value of $T(l)$ is compared with $1/f(SNR(l))$, i.e., the duration of a slot in which only link l is active. If $T(l) < 1/f(SNR(l))$, then link l is added to slot $S(l)$, otherwise a new slot is formed at the end of the schedule with only link l active. This process is repeated until all links have been scheduled.

In order to understand the relative benefits of the graded SINR model vs. the commonly used, thresholded version of the model, we have also implemented the *GreedyPhysical* algorithm of [3], which is a simple greedy algorithm that schedules a link in the first available slot(s), subject to the condition that the resulting transmission set is feasible under the thresholded SINR model. We recall that relative benefits of graded vs. thresholded SINR model have been recently quantified in about 30% throughput improvements in an experimental testbed [15], although these improvements refer to a different but related scheduling problem (single slot scheduling, also referred to as one-shot scheduling).

[4]We have repeated the simulations with log-distance path loss propagation, obtaining similar results.

[5]We recall that a logarithmic SINR vs. data rate relationship in the linear scale as in equation (1) is equivalent to a linear relationship in dB scale.

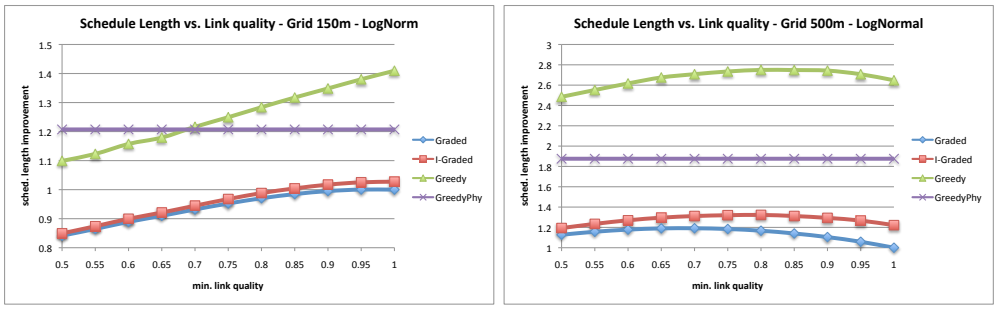

Figure 4: Schedule length improvement for increasing link quality threshold in the dense (left) and sparse (right) grid-like deployment scenario.

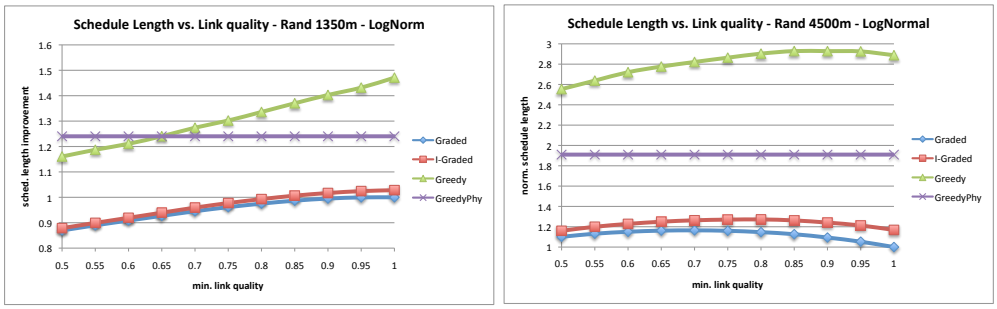

Figure 5: Schedule length improvement for increasing link quality threshold in the dense (left) and sparse (right) random deployment scenario.

4.3 Simulation results

In a first set of simulations, we distributed $n = 100$ nodes in a grid-like fashion. Two grid steps are considered, to mimic relatively dense and relatively sparse network deployments. Considering that PHY layer parameters are set as follows: path loss exponent $\alpha = 3$, transmit power $100mW$ ($20dBm$), and noise power $-90dBm$, we have a resulting nominal transmission range (in absence of shadowing and interference) of about $680m$ to obtain the maximum data rate of $55Mbs$ (which, we recall, requires a SINR $\geq 25dB$). Hence, we set the grid step in the dense deployment to $150m$, and to $500m$ in the sparse deployment. In both cases, internode separation is randomly perturbed by up to 10% to avoid artificial discretization effects. The shadowing parameter σ is set to $4dB$.

In both deployments, 10 nodes are randomly chosen as gateways, and node traffic, routing, and link demands generation is performed as described in the previous sections. The schedule lengths computed by $GradedSINR$, $I\text{-}GradedSINR$, and $GreedyGraded$ for a given link quality threshold β_Q are returned as the simulation result[6]. The simulator returns also the schedule computed by $GreedyPhysical$, which is based on the thresholded SINR model, and hence invariant to changes in the link quality threshold β_Q. We have generated 1000 different deployments for both the dense and the sparse scenarios, and considered link quality thresholds corresponding to link data rates ranging from 50% to 100% of the maximum nominal rate. Simulation results are shown in Figure 4.

We have also considered a random node deployment scenario, in which nodes are distributed uniformly at random in a certain square region. Similarly to the case of grid-like deployment, we have set the side of the deployment area to relatively small ($1350m$) and relatively large ($4500m$) values, to mimic relatively dense and relatively sparse deployments. In case of sparse deployments, we check that each non-gateway node has a path composed of only 100% quality links to at least one gateway node, so that demands can be fully satisfied under the thresholded SINR model. Any deployments not meeting this criterion are discarded. The results of this second set of simulations, also averaged over 1000 experiments, are reported in Figure 5.

The plots reported in Figures 4 and 5 report the average throughput length improvement of the various scheduling algorithms, which is normalized with respect to the schedule length of the sequential schedule when only 100% quality links are used. As seen from the figures, the trends for the grid-like and random scenarios are similar. In all cases, $GreedyGraded$ was by far the best scheduling algorithm, achieving as high as a near three-fold throughput improvement with respect to the sequential schedule[7]. The other scheduling algorithms for the graded model, for which, we recall, we have provable performance guarantees with respect to optimal, achieve only marginal throughput improvements (below 1.4), with $I\text{-}GradedSINR$ consistently performing slightly better than $GradedSINR$. Note that, due to the large size of the cell partitioning used in $GradedSINR$, the schedule computed by this algorithm always coincided with the sequential schedule; i.e., due to the very conservative

[6]Parameter ϵ in algorithms $GradedSINR$ and $I\text{-}GradedSINR$ is set to $1/2$.

[7]In the rest of this section, throughput improvements are always considered to be with respect to the sequential schedule using only 100% quality links.

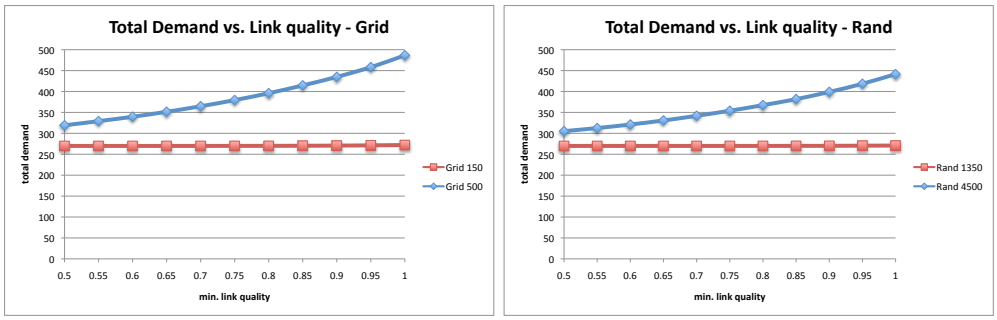

Figure 6: Total link traffic demand in the grid-like (left) and random (right) deployment scenario.

choice of the cell size driven by worst-case considerations, $GradedSINR$ was unable to achieve any spatial reuse. This lack of spatial reuse, coupled with possible usage of lower quality links (hence, longer slots) when the lower bound on link quality is below 100%, and with the fact that the total demand does not depend on the link quality threshold in dense deployments (see Figure 6), explains the relative throughput *degradation* with respect to the sequential schedule with 100% quality links experienced by $GradedSINR$ in dense deployments when using weak links is allowed.

When comparing the dense and sparse scenarios, we observe higher throughput improvements in the sparse scenario: close to three-fold improvements are achieved in the sparse deployments, compared to no more than 1.5-fold improvements in the dense setting. This difference is due to the additional opportunities for spatial reuse in a larger, i.e. sparser, network deployment.

Concerning the impact of link quality threshold on schedule length, we observe a clear effect of node density for the more aggressive scheduling algorithm, namely $GreedyGraded$: when the network is dense, the schedule length tends to increase as the lower bound on link quality decreases, implying that allowing use of relatively weak links is detrimental for network throughput. To the contrary, in sparse network deployments, using relatively weak links can improve throughput: about 10% (5%) further improvements are observed when the link quality threshold is reduced from 100% to 80% (85%) in the grid-like (random) case. In both cases, further reducing the link quality threshold has negative effects on throughput.

The radically different behavior in case of dense or sparse networks can be explained by the data reported in Figure 6, which shows the total traffic demand as a function of the link quality threshold. In case of dense deployments, the total demand does not depend on the link quality threshold, indicating that, even for the most stringent link quality requirement, relatively short paths to the gateways are available. The throughput degradation that is observed in case of lower link quality thresholds is due to the fact that the routing algorithm is oblivious to link quality when building the shortest path tree; hence, if relatively weak links are included in the tree, the average slot duration is increased (lower link rates) which, coupled with the unchanged total demand, results in an overall throughput degradation. On the other hand, in sparse network deployments total traffic demand considerably increases as the link quality threshold increases, indicating the short paths to the gateways can be found only if relatively weak links are used. Although usage of weak links tends to increase average slot duration, the lower total demand compensates this increase with a reduction in the total number of slots, resulting in an overall throughput increase. However, if very weak links ($\leq 75\%$) are used, the reduction in total traffic demand is no longer sufficient to compensate for the increased average slot duration, resulting in an overall throughout degradation.

Finally, we comment on the relative throughput benefits of using the graded vs. thresholded SINR interference model: with similar greedy approaches to schedule links, we observe a throughput improvement of $GreedyGraded$ over $GreedyPhysical$ of about 18% for dense deployments, and about 50% for sparse deployments. This is true, even though we are using a routing algorithm that is oblivious to link quality (except in a relatively crude way, through use of the link quality threshold). Hence, we expect even larger throughput improvements can be attained when using a link-quality-aware routing algorithm. Study of this aspect is left for future work. Nevertheless, throughput improvements of up to 50%, even with the simple routing algorithm used herein, show that very substantial benefits can be achieved through use of the graded SINR model.

5. EXPERIMENTAL EVALUATION

The main purpose of the experimental evaluation is to study how the choice of link quality threshold affects throughput in a real network. We use TelosB motes [18] that are equipped with CC2420 radio [5]. The radio is compliant with the IEEE 802.15.4 [11] PHY layer standard in the 2.4 GHz ISM band and operates at a fixed nominal bit rate of 250 Kbits/s. We have implemented a simple TDMA protocol in TinyOS-2.0 [23] in which motes transmit at designated time instants without performing carrier sensing or backoff as in the default MAC implementation in TinyOS.

The data rate is fixed due to the choice of hardware. For simplicity, we also fix the transmission power to -32.5 dBm uniformly on all nodes. Hence, in this section we use the PRR interpretation of the graded SINR interference model. Furthermore, we focus our attention on the simpler and more practical (as well as best performing on the average) greedy approach for transmission scheduling.

The setup of the experimental testbed is similar to the one used in simulations. More specifically, we deploy $n = 20$ nodes, placed randomly on a 10 foot by 3 foot tabletop in an office environment. Through extensive measurements we derive the input parameters to the "routing/scheduling" module, which are the following: 1) the $n \times n$ channel gain matrix CG, reporting the channel gain between each possible node

pair; 2) the $n \times 1$ noise vector NV, reporting the noise level at each node; and 3) the PRR vs. SINR function $f()$. The measurement methodology used to collect 1)–3) is similar to the one used in [15]. The PRR vs SINR function we obtained is similar to the one presented in [15]. The function has a graded region from -3 to +5 dB. Beyond a SINR of 5 dB, links always have PRR close to 100%.

The input parameters are fed to a centralized node (a PC) which runs the "routing/scheduling" module as follows. Similarly to the simulation-based evaluation, two of the nodes are randomly selected as gateways. Non-gateway nodes are assigned an integer demand chosen at random in the interval $[1, 5]$. Then, given the link quality threshold SNR_Q and matrix CG, the set of available links is determined, and shortest path trees routed at the gateway nodes are built. Given node demands and the set of routes, link demands are computed, which are fed to the scheduling algorithm.

Given the PRR interpretation of the graded SINR interference model, a re-design of the link scheduling algorithm is needed. In particular, variable slot duration is no longer needed, since link data rate is fixed and the same for all links. However, a packet scheduled for transmission along link l in slot S under the PRR interpretation is received only with probability $p_{l,S}$, with $0 \leq p_{l,S} = PRR_{l,S} \leq 1$, where $PRR_{l,S}$ is the PRR on link l in slot S. Packet transmissions in a specific slot can then be interpreted as Bernoulli trials with a certain, fixed success probability[8]. If the schedule is repeated N times, by the LLN we have that the expected number of successful transmissions along link l in slot S converges to $N \cdot p_{l,S}$ as N grows larger. Hence, the expected long-term effective data rate on link l in slot S is $p_{l,S}$. Based on this observation, the greedy scheduling algorithm described below considers that an amount of demand equal to $p_{l,S}$ is satisfied when link l is scheduled for transmission in slot S.

The scheduling algorithm is as follows. The approach is again greedy: links are initially ordered, and are processed sequentially. The algorithm keeps extracting elements from the list of links to be scheduled, till the demand on all links is satisfied. The main difference with *GreedyGraded* is that a single link might be considered repeatedly when building the schedule (see below).

When link l is considered, the algorithm sequentially scans all currently built slots. For each slot S, the algorithm first checks whether adding l to the slot would keep it "feasible" (this is a soft notion of feasibility, described below); if the slot remain "feasible", the algorithm computes a "fitness" measure, namely the difference between the increase in expected throughput due to adding the new link, and the throughput decreases on the already scheduled links. The throughput of a slot S is the sum of all $p_{l,S}$ values on the scheduled links. If the "fitness" of the slot is positive (i.e., we have a throughput increase by adding l to the slot), then the slot is a candidate slot for link l. After scanning all currently available slots, the algorithm adds l to the slot S with best positive fitness $fit(S)$. If $fit(S) < 0$, a new slot is created at the end of the schedule, and link l only is put in the new slot.

Once link l has been included in a slot, link demands are updated as follows.
*Case*1. Link l is added to an existing slot S: the demand of l is decreased of $p_{l,S}$; furthermore, the demands of all links in $S \setminus \{L\}$ is *increased* of $(p_{l,S\setminus\{L\}} - p_{l,S})$. This is to possibly account for PRR degradation of links in $S\setminus\{L\}$ due to adding l to the slot. Note that if the demand on some of these links were 0 (link already successfully scheduled), a new instance of the link with the remaining demand has to be included again in the list of links to be scheduled.
*Case*2. Link l is added to a new slot S': the demand of l is decreased of $p_{l,S'} = PRR(l)$, since only link l is scheduled in S'.

The soft notion of "feasibility" used in the algorithm is an optimization aimed at ensuring that the demand on a link is decreased of a significant amount when scheduled in a slot. In particular, we define a set of transmissions $l_1, \ldots, l_k$ to be *feasible* if $p_{l_i,\{l_1,\ldots,l_k\}} \geq PRR_q$ for each i, where $PRR_q < PRR_Q$ is a PRR quality threshold (e.g., 0.5). Note that this threshold is different (and lower) than the quality threshold used to define which links are "good" and usable by a routing algorithm. In fact, the latter threshold refers to the link quality based on the SNR, while the formed on the link quality based on the (lower) SINR value when all scheduled links are simultaneously transmitting.

[8]This holds true only under the assumption that the radio environment is relatively stable.

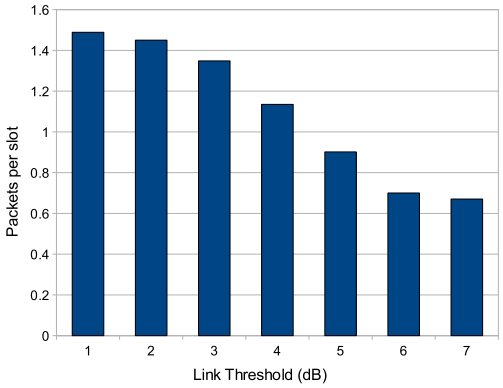

Figure 7: Normalized aggregate throughput at the gateway nodes as a function of the link quality threshold.

5.1 Experimental Results

Different schedules are obtained by choosing different link quality thresholds. Once the schedule is computed, it is fetched to the testbed nodes, which repeatedly execute the schedule and transmit packets. Each schedule is repeated 100 times. The outcome of an experiment is the aggregate throughput measured at the two gateway nodes. Note that, sometimes links can be over-scheduled. This means that the sum of PRRs of a link scheduled in different slots might exceed the weight on that link. Thus, as a result, the number of packets successfully received at the gateways might exceed the number of packets *scheduled* to be received. We do not consider these extraneous packets in our calculation of the throughput.

We present the results of our testbed experiments in Figure 7. The X-axis enumerates the various schedules generated with different link quality thresholds. The link quality thresholds are varied from SNR values of 1 dB to upto 7 dB. On Y-axis we plot the throughput in terms of packets

successfully received at the gateways normalized with the schedule length, or as packets per slot. As can be seen, using a lower link quality threshold – even in the transition region – results in improving the throughput. Infact, 70% better throughput is obtained by using weak links (a link quality threshold of 1 dB) compared to very strong links (7 dB). We conjecture that this is because, by letting the routing protocol utilize weak links, a packet ends up taking fewer number of hops to the gateways – thus making the schedule more compact. Lowering link threshold further does not give any performance benefit in our testbed giving same results as for the threshold of 1 dB. These results show that the *GreedyGraded* algorithm works quite well in a real *mesh* network scenario, where packets are routed towards gateways, giving high end-to-end throughput even with relatively weak links.

6. CONCLUSIONS AND FUTURE WORK

We believe this paper delivers several contributions, and opens numerous avenues for further research. From the methodological point of view, the paper encompasses all stages of the "from ideas to testbed implementation" process: 1) starting from the formalization of a new interference model and related problem definition; 2) continuing with presentation of algorithms with proven approximation bounds for the problem considered; then 3) evaluating performance through simulation, as well as presenting a more practical variation of the scheduling algorithm; and finally 4) implementing the practical version of the scheduling algorithm in an experimental testbed, and evaluating its performance in a practical setting.

Several questions are left open by this paper, which can be considered only as a starting point towards a better understanding of the possible benefits of allowing use of "imperfect" links on the resulting network throughput. In particular, the problem of routing and scheduling for throughput optimization under the graded SINR model should be considered. Furthermore, a better understanding of the impact of node density on routing/scheduling performance is needed. From the experimental viewpoint, an assessment of whether the throughput measured at the gateway nodes is not only increased, but also proportional to actual node demands is needed. Such an assessment would make our proposed scheduling approach a promising candidate as a building block for providing strong QoS guarantees in a wireless multi-hop network. Finally, implementing the proposed scheduling techniques with a high data rate technology (e.g., WiFi) is another challenge to be undertaken.

7. REFERENCES

[1] `http://ieee802.org/16/`

[2] M. Alicherry, R. Bathia, L. Li, "Joint Channel Assignment and Routing for Throughput Optimization in Multi-Radio Wireless Mesh Networks", *Proc. ACM Mobicom*, pp. 58–72, 2005.

[3] G. Brar, D. Blough, and P. Santi, "Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks," *Proc. ACM MobiCom*, pp. 2–13, 2006.

[4] G. Brar, D. Blough, and P. Santi, "The SCREAM Approach for Efficient Distributed Scheduling with Physical Interference in Wireless Mesh Networks", *Proc. IEEE ICDCS*, 2008.

[5] *CC2420 Radio Datasheet*, Texas Instruments, Oct. 2005.

[6] D. Chafekar, V.S. Anil Kumar, M.V. Marathe, S. Parthasarathy, A. Srinivasan, "Approximation Algorithms for Computing Capacity of Wireless Networks with SINR Constraints", *Proc. IEEE Infocom*, pp. 1166–1174 , 2008.

[7] R.L. Cruz, A.V. Santhanam, "Optimal Routing, Link Scheduling and Power Control in Multi-Hop Wireless Networks", *Proc. IEEE Infocom*, pp. 702–711, 2003.

[8] R.L. Cruz, A.V. Santhanam, "Hierarchical Link Scheduling and Power Control in Multihop Wireless Networks", *Proc. Allerton Conference on Communications, Control and Computing*, 2002.

[9] O. Goussevskaia, Y.V. Oswald, R. Wattenhofer, "Complexity in Geometric SINR", *Proc. ACM MobiHoc*, pp. 100–109, 2007.

[10] P. Gupta and P.R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. on Info. Theory,* Vol. 46, No. 2, pp. 388–404, 2000.

[11] IEEE Computer Society LAN/MAN Standards Committee, "802.15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (LR-WPANS)," 2003.

[12] D.B. Johnson, D.A. Maltz, "Dynamic Source Routing in Ad Hod Wireless Networks", *Mobile Computing*, n. 353, pp. 153–181, 1996.

[13] A. Keshavarz-Haddad, R. Riedi, "On the Broadcast Capacity of Multihop Wireless Networks: Interplay of Power, Density and Interference," *Proc. IEEE Secon*, pp. 314-323, 2007.

[14] R. Nelson and L. Kleinrock, "Spatial-TDMA: A Collison-free Multihop Channel Access Protocol," *IEEE Transactions on Communication,* Vol. 33, pp. 934–944, Sept. 1985.

[15] R. Maheshwari, S. Jain, S.R. Das, "A Measurement Study of Interference Modeling and Scheduling in Low-Power Wireless Networks", *Proc. 6th ACM Conference on Embedded Networked Sensor Systems (ACM SenSys 2008)*, Raleigh, NC, Nov 2008.

[16] R. Maheshwari, J. Cao, S.R. Das, "Physical Interference Modeling for Transmission Scheduling on Commodity WiFi Hardware", *Proceedings of Infocom 2009.*

[17] T. Moscibroda and R. Wattenhofer, "The Complexity of Connectivity in Wireless Networks," *Proceedings of Infocom 2006.*

[18] Moteiv Corporation, *tmote sky: Ultra low power IEEE 802.15.4 compliant wireless sensor module*, November 2006.

[19] C. Perkins, E. Belding-Royer, S. Das, "Ad Hoc On-Demand Distance Vector (AODV) Routing", *IETF draft*, `http://www.ietf.org/rfc/rfc3561.txt`, 2003.

[20] T.S. Rappaport, *Wireless Communications*, Prentice Hall, Upper Saddle River, NJ, 2002.

[21] C. Scheideler, A. Richa, P. Santi, "An O(log n) Dominating Set Protocol for Wireless Ad Hoc Networks under the Physical Interference Model", *Proc. ACM MobiHoc*, pp. 91–100, 2008.

[22] G. Sharma, R. Mazumdar, N. Shroff, "On the Complexity of Scheduling in Wireless Networks", *Proc. ACM Mobicom*, pp. 227–238, 2006.

[23] "TinyOS community forum," http://www.tinyos.net.

[24] M. Zuniga, B. Krishnamachari, "Analyzing the Transitional Region in Low Power Wireless Links", *Proc. IEEE Secon*, pp. 517–526, 2004.