

Computationally Efficient Scheduling with the Physical Interference Model for Throughput Improvement in Wireless Mesh Networks *

Gurashish Brar
School of Electrical and
Computer Engineering,
Georgia Institute of
Technology

gsbrar@ece.gatech.edu

Douglas M. Blough
School of Electrical and
Computer Engineering,
Georgia Institute of
Technology

dblough@ece.gatech.edu

Paolo Santi
Istituto di Informatica e
Telematica del CNR,
Pisa, Italy

paolo.santi@iit.cnr.it

ABSTRACT

Wireless mesh networks are expected to be widely used to provide Internet access in the near future. In order to fulfill the expectations, these networks should provide high throughput simultaneously to many users. Recent research has indicated that, due to its conservative CSMA/CA channel access scheme and RTS/CTS mechanism, 802.11 is not suitable to achieve this goal.

In this paper, we investigate throughput improvements achievable by replacing CSMA/CA with an STDMA scheme where transmissions are scheduled according to the physical interference model. To this end, we present a computationally efficient heuristic for computing a feasible schedule under the physical interference model and we prove, under uniform random node distribution, an approximation factor for the length of this schedule relative to the shortest schedule possible with physical interference. This represents the first known polynomial-time algorithm for this problem with a proven approximation factor.

We also evaluate the throughput and execution time of this algorithm on representative wireless mesh network scenarios through packet-level simulations. The results show that throughput with STDMA and physical-interference-based scheduling can be up to three times higher than 802.11 for the parameter values simulated. The results also show that our scheduling algorithm can schedule networks with 2000 nodes in about 2.5 minutes.

Categories and Subject Descriptors

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

*This research was supported in part by the National Science Foundation under Grants ENG-0225417 and INTL-0405157. The work of Paolo Santi was supported in part by CNR under the Short-Term Mobility program.

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.

MobiCom'06, September 23–26, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-286-0/06/0009 ...\$5.00.

General Terms

Algorithms, Performance, Theory

Keywords

Wireless mesh networks, STDMA scheduling, physical interference model, network capacity

1. INTRODUCTION

In wireless mesh networks, wireless backbone nodes (also called *wireless routers* in the following) must convey a large amount of traffic generated by wireless clients to a few nodes that act as gateways to the Internet. For these networks, the main design concern is increasing the traffic carrying capacity of the wireless backbone as much as possible, while energy conservation (one of the main design goals in ad hoc/sensor networks) is not an issue – wireless routers typically do not have energy constraints. In order to meet the high traffic demand, wireless backbone nodes can be equipped with multiple radios and/or operate on multiple channels.

In order to understand whether forthcoming wireless mesh networks will actually be able to meet traffic demands, several recent papers have investigated the potential benefits on network capacity of using multiple radio/channels. This type of investigation requires addressing at least three problems – routing, channel allocation, and transmission scheduling – which have been studied either separately or jointly. For instance, [32] is concerned with routing, [21, 23] with channel allocation, [3, 12, 18] with scheduling, while [1, 13, 15, 17, 29] jointly consider the three problems.

The main factor that limits capacity in mesh networks (and in wireless networks in general) is *interference*, which is a consequence of using a shared communication medium. Hence, an accurate modeling of interference is fundamental in order to derive theoretical and/or simulation-based results of some practical relevance. In the literature, two main interference models have been proposed [8]: the *protocol* and the *physical* interference models.

In the former model, a communication between nodes u and v is successful if no other node within a certain *interference range* from v (the receiver) is simultaneously transmitting. Due to its simplicity, and to the fact that this model can be used to mimic the behavior of CSMA/CA networks

such as IEEE 802.11,¹ the protocol interference model has been mostly used in the literature.

In the physical interference model, a communication between nodes u and v is successful if the SINR (Signal to Interference and Noise Ratio) at v (the receiver) is above a certain threshold, whose value depends on the desired channel characteristics (e.g., data rate). This model is less restrictive than the protocol interference model: it may occur that a message from node u to node v is correctly received even if there is a simultaneous transmitting node w close to v (for instance, because node u is using a much larger transmit power than node w). As a result, higher network capacity can in general be achieved by applying the physical interference model. Note that the physical interference model is representative of a scenario that *does not* use CSMA techniques; instead, transmissions should be carefully scheduled so that only sender/receiver pairs that do not conflict with each other transmit simultaneously. In other words, the physical interference model is suited for use with TDMA-like channel access schemes.

Recent research indicates that CSMA/CA is not suitable to meet the high traffic demand of wireless mesh networks. The reason for this is that CSMA/CA is a very conservative mechanism: due to the combination of carrier sensing and collision avoidance techniques, many network nodes are silenced when a certain communication takes place. The ‘silenced area’ grows quadratically (in an idealized setting) with the transmission range, hence CSMA/CA becomes more and more conservative (with a negative impact on capacity) as the nodes’ transmission ranges increase. Since wireless routers typically have relatively long ranges (ranges of several hundreds meters or a few kilometers can be expected [4, 22]), CSMA/CA is deemed not adequate to meet the high traffic demand of wireless mesh networks. This is the reason why existing implementations of 802.11-based mesh networks disable the collision avoidance mechanism (i.e., the RTS/CTS message exchange) [4], or completely new TDMA-like MAC protocols are proposed for mesh networks [22, 31].

The above discussion motivates use of the physical interference model in investigations of the capacity of wireless mesh networks. A major difficulty lies in the complexity of handling physical interference. Let us consider the problem of scheduling, which is fundamental to understand whether actual capacity improvements can be achieved by the physical interference model. Assume that traffic demands on each network link are known (for instance, by executing an interference-aware routing algorithm – see, e.g., [1]), and consider the problem of optimal conflict-free scheduling of these transmissions. This problem is known to be NP-hard even with the simple protocol model. However, in case of protocol interference, the authors of [1] introduce a polynomial time algorithm that computes a conflict-free schedule that is at most a constant factor away from the optimal schedule. Unfortunately, whether there exists a polynomial-time algorithm that is within a provable approximation factor of optimal for scheduling under the physical interference model is an open problem.

In this paper, we solve this open problem, presenting a polynomial time heuristic called Algorithm GreedyPhysical that computes a conflict-free schedule under the physical

interference model. We formally prove an approximation factor for the schedule computed by our algorithm relative to the optimal schedule in case of unplanned (uniform random) network deployments.

The approach presented in this paper makes the relative advantage of the physical interference model over the protocol model tangible. We present extensive simulation results that: 1) compare schedule length under the two models, and 2) evaluate throughput under a TDMA approach with our scheduling algorithm and compare it to 802.11 throughput using a packet-level simulator with accurate models for 802.11 and physical interference. The results show significant performance benefits are obtainable with scheduling approaches based on physical interference.

2. RELATED WORK

Our scheduling approach can be applied within the spatial reuse TDMA (STDMA) access scheme defined by Nelson and Kleinrock [19]. Most of the scheduling mechanisms proposed for STDMA use the protocol interference model. Only a few previous works have considered physical interference in this context [5, 6, 7, 12, 18, 28].

In [12], Jain, et al., formulate the problem of scheduling under physical interference as an LP problem. Unfortunately, this formulation requires computing the set of all possible sets of non-conflicting transmissions, which requires exponential time and is, therefore, computationally intractable for more than a trivial number of nodes. Due to this requirement, the authors were not able to provide any simulation-based evaluation of scheduling with physical interference, not even for very small networks. The work of [6] also provides an exponential-time LP formulation. The approach of [28] is similar to that of [12] in that all maximal sets of non-conflicting transmissions are considered, and this approach also requires exponential time.

In [5], Gronkvist and Hansson describe the use of physical interference in STDMA and compare it to an approach that considers interference up to a certain distance from a node and ignores it beyond that point. The scheduling algorithm in this work is based on the algorithm presented in [7]. Neither [5] nor [7] provides an evaluation of the algorithm’s time complexity nor compares its performance to the optimal.

In [18], Moscibroda and Wattenhofer consider the problem of scheduling with physical interference under the assumption that traffic demands are the same on every network link. Although the analysis and algorithms presented in [18] are interesting, they are not representative of the typical wireless mesh network scenario, in which traffic demands on the links are in general very different – typically, links closer to a gateway node experience higher traffic demands than links on the border of the network.

In summary, our approach provides the first computationally efficient scheduling algorithm for the physical interference model with non-uniform link demands for which an approximation factor relative to the optimal schedule has been proved. We also use this algorithm to evaluate, through simulations, performance of physical-interference-based scheduling relative to traditional protocol-model-based scheduling for a variety of scenarios that are representative of different wireless mesh network deployments. These results demonstrate that significant throughput improvements can be obtained with physical-interference-based scheduling in wireless mesh networks.

¹This can be done by using the concept of interference range both at the sender and at the receiver of a communication (see [1]).

3. NETWORK AND INTERFERENCE MODELS

Consider a mesh network composed of n wireless routers. The links among the wireless routers are represented by a physical communication graph $G_p = (V, E)$, where V is the set of routers, and edge $(u, v) \in E$ if and only if a link between nodes u and v exists *in absence of interference* from other network nodes. We do not assume any specific radio propagation model, nor that all nodes use the same transmission power. Hence, unidirectional links can be present in the physical topology of the network.

We assume that link-layer reliability is required for both CSMA/CA and STDMA. Thus, an ACK packet is generated by each receiver for every data packet it receives. This implies that, although unidirectional links can be present, they are not used by the protocols. Thus, in the logical topology of the network, represented by a logical communication graph G , every edge represents a bi-directional link. This assumption is made so as to provide a fair throughput comparison between the STDMA approach and 802.11, which mandates link-layer reliability. We note that removing this assumption and making use of unidirectional links within STDMA could potentially further improve throughput relative to 802.11, albeit at the expense of reliability.

We assume that each edge $e = (u, v)$ in G is labeled with a weight w_e , which represents the traffic demand on the link for a single STDMA scheduling period. Note that, since edge e is undirected, w_e represents the aggregated traffic in both directions. We are not concerned with how weights w_e are generated: our approach can be applied for any weight values, i.e., for arbitrary traffic patterns. In practice, the traffic demand on each link depends on the distribution and traffic pattern of wireless clients, and on the routing algorithm used to convey data from/to the clients to/from the gateway nodes. We assume that the demands are known a priori because, in this paper, we are primarily concerned with investigating the potential for throughput improvements with STDMA scheduling, and not with implementation of practical protocols. However, if demands do not fluctuate rapidly, it should be possible to measure the demands in an actual system and periodically adjust the input values to the scheduling algorithm accordingly.

In this paper, we consider both the protocol and the physical interference models. To be specific, we use the following variation of the protocol model (introduced in [1]) to accurately mimic the behavior of CSMA/CA networks such as 802.11. Let r_T and r_I denote the *transmission range* and the *interference range* (or *carrier sensing range*), respectively. The transmission range r_T represents the maximum distance up to which a packet can be received, while the interference range r_I represents the maximum distance up to which a node sensing the channel detects an ongoing transmission². The interference range is larger than the transmission range, i.e., $r_I = q \cdot r_T$, with $q > 1$ ($q = 2$ is a typical value for 802.11 networks). Note that the actual values of the transmission and interference ranges depend on the transmission power used by the nodes. Since in our model we allow nodes to use different transmission powers, we denote the transmission and interference ranges of node x as $r_T(x)$ and $r_I(x)$,

²Implicit in the definition of transmission and interference ranges is the assumption of a deterministic signal propagation model, such as the log-distance path model.

respectively. Due to carrier sensing and RTS/CTS/ACK exchange, a transmission along link $e = (u, v)$ (in either direction) blocks all simultaneous transmissions within the interference ranges of both nodes u and v . To be specific, denoting with $D(x, r_I(x))$ the disk of radius $r_I(x)$ centered at node x , all transmissions such that one of (or both) the endpoints of the link lies within $IR_e = D(u, r_I(u)) \cup D(v, r_I(v))$ are blocked by the transmission along link (u, v) . Here, IR_e stands for Interference Region associated with link e .

In the physical interference model [8], successful reception of a packet sent by node u to node v depends on the SINR at v . To be specific, denoting by $P_v(x)$ the received power at v of the signal transmitted by node x , a packet along link (u, v) is correctly received if and only if:

$$\frac{P_v(u)}{N + \sum_{w \in V'} P_v(w)} \geq \beta,$$

where N is the background noise, V' is the subset of nodes in V that are transmitting simultaneously, and β is a constant that depends on the desired data rate, the modulation scheme, etc.

To be coherent with our assumption of link-layer reliability, we extend the physical interference model as follows. We assume that a packet sent by node u is correctly received by node v if and only if the packet is successfully received by v , and the ACK sent by node v is correctly received by node u . Furthermore, for a transmission from node x to node y that is concurrent with the packet on (u, v) , we account for the interference both from node x 's data packet *and* from node y 's ACK. Since only one of x and y transmits at a time but their data and ACK packets could both overlap with either the data packet or the ACK along (u, v) , we choose the maximum of the interferences from x and y when calculating the total interferences at u and v . Note that which of the two (x or y) contributes the maximum interference could be different at u and v . In light of this discussion, a packet sent along link (u, v) (in either direction) is correctly received if and only if:

$$\frac{P_v(u)}{N + \sum_{(x,y) \in E'} \max(P_v(x), P_v(y))} \geq \beta,$$

and

$$\frac{P_u(v)}{N + \sum_{(x,y) \in E'} \max(P_u(x), P_u(y))} \geq \beta,$$

where set E' contains all links that have transmissions concurrent with the one on (u, v) .

Given an interference model, we can build a *conflict graph*, which can be used to determine whether a certain set of transmissions can occur simultaneously. In the protocol interference model, the conflict graph G_{Prot} is constructed by having a node for each link in G , and by adding the undirected edge (e_1, e_2) if and only if links e_1 and e_2 conflict with each other (according to the rule described above). It is easy to see that a set of transmissions along links e_1, \dots, e_k is *feasible* if and only if the corresponding nodes in G_{Prot} form an independent set³. This leads to the definition of feasible scheduling under the protocol interference model:

DEFINITION 1. *Let G be the communication graph with traffic demands w_e on each link, and let G_{Prot} be the conflict*

³Given a graph $G = (V, E)$, a set $V' \subset V$ is an *independent set* of G if and only if no edge in G has both endpoints in V' .

graph under the protocol interference model. A schedule S composed of T time slots t_1, \dots, t_T is feasible for G if and only if the following conditions are satisfied:

- the set of transmissions scheduled at each time slot t_i is an independent set of $G_{P_{rot}}$, and
- each link e is scheduled for at least w_e time slots.

Note that in the above definition we have made the assumption that demands w_e on the links are expressed as a multiple of the time slot used in the scheduling. This can be easily accomplished in practice, once the link data rates and the length of the time slot are known.

Let us now consider the physical interference model. A major factor contributing to the complexity of dealing with physical interference is the fact that the notion of ‘conflict between transmissions’ cannot be modeled as a binary relation between edges, as it was the case with the protocol interference model. For instance, it might be the case that transmission along links e_1, e_2 can successfully occur simultaneously if these are the only scheduled transmissions, while one of them (or both) might be compromised if other nodes are allowed to transmit simultaneously. In other words, whether a certain set of scheduled transmissions is feasible can be determined only by considering the SINR ratios at all the endpoints of the scheduled links, and not by analyzing edge-to-edge mutual interference as was the case with the protocol interference model.

An edge-based notion of conflict graph for the physical interference model has been introduced in [12], where a node is introduced for each link in the network, and a weighted edge is inserted between nodes e_1, e_2 (which represent links in the communication graph) that can interfere with each other, where the weight of the link represents the fraction of the maximum permissible noise+interference level at the node at the receiver end of link e_2 that is contributed by activity on link e_1 . This edge-based notion of conflict graph is quite complicated, and can be considerably simplified by observing that, under the assumption that no per-packet transmit power control is used, the notion of interference in the physical model is associated with *nodes*, and not with *edges*. In other words, the interference level generated at a receiver v by a transmitter u , which is transmitting to a third node, say w , is independent of w , because u uses the same transmission power for all communications (of course, due to our link-layer reliability assumption, the interference generated by w when sending its ACK packet must also be included). This leads to the following definition of a node-based conflict graph in the physical interference model. The conflict graph G_{Phy} has the same node set V as the communication graph, and edge set $E' \supseteq E$, where directed edge $e_i = (u, v) \in E'$ represents the fact that the signal transmitted by node u potentially contributes to the interference level measured at node v . Edge (u, v) is assigned a weight w_{uv}^i , which corresponds to the received power at node v of the signal transmitted by node u , i.e., $w_{uv}^i = P_v(u)$. Note that edges in the conflict graph G_{Phy} are directed, and that $w_{uv}^i \neq w_{vu}^i$ in general.

Concerning the relative complexity of G_{Phy} as introduced here and the edge-based notion of conflict graph introduced in [12], we observe that in the physical interference model every transmitter in the network contributes (possibly with a very small, but greater than zero, amount) to the interference level measured at any receiver. Hence, *every pair*

of links in the network potentially conflict with each other, and both G_{Phy} and the interference graph defined in [12] are complete graphs. The complete graph on m nodes has $m(m-1) = \Theta(m^2)$ edges. The fundamental difference between G_{Phy} and the conflict graph of [12] is that in our case G_{Phy} has a node for every *node* in the network, while the conflict graph of [12] has a node for every *link* in the network. Hence, G_{Phy} has $\Theta(n^2)$ edges, as compared to the $O(n^4)$ of the conflict graph of [12], resulting in a significant reduction in computational complexity just to construct the conflict graph, and even greater complexity improvements for scheduling based on the conflict graph.

Given the communication graph $G = (V, E)$ and the conflict graph $G_{Phy} = (V, E')$, we can determine whether a certain set of transmissions $E'' = \{e_1, \dots, e_k\} \subseteq E$ is feasible as follows. Let us denote with $V(E'') \subseteq V$ the set of all nodes $u \in V$ such that u is the endpoint of at least one edge in E'' . First, we observe that if a node in $V(E'')$ is an endpoint of more than one edge in E'' , then E'' is not feasible, since a node cannot transmit and receive on different links simultaneously. Hence, a necessary condition for transmission set E'' to be feasible is that E'' is a *matching*⁴ on the communication graph G .

PROPOSITION 1. *Given a communication graph $G = (V, E)$, a set $E'' = \{e_1, \dots, e_k\} \subseteq E$ of transmissions is feasible only if E'' is a matching of G .*

In order to determine whether a matching E'' on G is feasible, we have to check for every node $u \in V(E'')$ whether the SINR at u is above the desired threshold β . In order for E'' to be feasible, the condition on the SINR must be fulfilled at every node in $V(E'')$. When computing the SINR for node u which is incident into edge $e_i = (u, v) \in E''$, only the nodes in $V(E'') - \{v\}$ contribute to the interference level at u . For any such node x which is an endpoint of link $e' = (x, y)$, x 's contribution to the interference measured at node u equals w_{xu}^i if $P_u(x) \geq P_u(y)$, and it is 0 otherwise (recall the use of the max operator in our definition of the physical interference model). Summarizing, we have the following characterization of a feasible set of transmissions under the physical interference model.

DEFINITION 2. *Given a communication graph $G = (V, E)$ and a conflict graph $G_{Phy} = (V, E')$, a set $E'' = \{e_1, \dots, e_k\} \subseteq E$ of transmissions is feasible under the physical interference model if and only if:*

- E'' is a matching of G , and
- for every $u \in V(E'')$, with $e_i = (u, v) \in E''$,

$$\frac{P_u(v)}{N + \sum_{(x,y) \in E'' - \{(u,v)\}} \max(w_{xu}^i, w_{yu}^i)} \geq \beta .$$

We can now define a feasible schedule under the physical interference model.

DEFINITION 3. *Let G be the communication graph with traffic demands w_e on each link, and let G_{Phy} be the conflict graph under the physical interference model. A schedule S composed of T time slots t_1, \dots, t_T is feasible for G if and only if the following conditions are satisfied:*

⁴A matching on a graph $G = (V, E)$ is a set of edges $E' \subseteq E$ such that no two edges in E' share a common endpoint.




Figure 1: Example Communication Graph




Figure 2: Conflict graphs in the protocol (left) and physical (right) interference models. In case of G_{Phys} , only the outgoing edges from node u_1 are shown. The weight of edge (u_1, u_i) represents the received signal strength at node u_i when only node u_1 is transmitting and is expressed in dBm.

- the set of transmissions scheduled at each time slot t_i is feasible under the physical interference model, and
- each link e is scheduled for at least w_e time slots.

We now present an example of a communication graph, and the corresponding conflict graphs in the protocol and physical interference models. Consider the node deployment of Figure 1, and assume all nodes have transmit range $r_T = 200m$, and interference range $r_I = 400m$. The conflict graph with protocol interference is shown on the left side of Figure 2: since all nodes interfere with all other nodes, the conflict graph G_{Prot} is a complete graph. This implies that no two transmissions can occur concurrently in the network for this example. Let us now consider the physical interference model. Assume that radio signal propagation obeys the log-distance path model with path loss exponent $\alpha = 3.5$, and that nodes have transmit power equal to 50mW. Part of the physical conflict graph G_{Phys} resulting from this setting is shown on the right side of Figure 2. Since every node potentially interferes with every other node in the network, G_{Phys} is also a complete graph. However, transmissions along links $e_1 = (u_1, u_2)$ and $e_7 = (u_5, u_6)$ can occur simultaneously (in either direction) in the physical interference model. In fact, with the above setting and assuming a typical value of the background noise $N = -90dBm$, the SINR at node u_1 (and, because of symmetry, also at the other three nodes involved in the transmissions) is about 16dB, which is an adequate value for supporting 11 Mbps data rate on both links [10, 20].

Algorithm GREEDYPHYSICAL:

Input: a weighted communication graph G
and the conflict graph G_{Phys}
Output: a feasible schedule S_{Phys} of length T_{Phys}
under physical interference model

1. set available slots to \emptyset ;
 2. order the links in E according to the interference number. Let e_1, \dots, e_m be the resulting ordering;
 3. for $i = 1, \dots, m$
 - 3a. schedule link e_i in the first w_{e_i} available slots t_k such that the resulting set of scheduled transmission is feasible with the physical interference model
 - 3b. if currently available slots are not sufficient to schedule w_{e_i} slots for e_i , add new slots at the end of the schedule S_{Phys} and schedule link e_i alone in these slots
 4. return schedule S_{Phys} , and schedule length T_{Phys} ;
-

Figure 3: Algorithm GreedyPhysical

4. SCHEDULING ALGORITHM

In this section, we present a computationally efficient heuristic referred to as Algorithm GreedyPhysical for building a feasible schedule under the physical interference model, and we formally prove an approximation bound for our algorithm in the case of nodes distributed uniformly at random in a square of unit area.

Algorithm GreedyPhysical is reported in Figure 3. GreedyPhysical is a simple heuristic, which is similar to the one used in [1] to prove the approximation bound for the scheduling problem with protocol interference. First, the links to be scheduled are ordered according to the *interference number*, which is defined as follows. The interference number of edge $e = (u, v)$ is the number of edges e_i in the communication graph such that e_i does not have u or v as an endpoint and $E = (e, e_i)$ is infeasible. All of these e_i are definitely prevented from communicating at the same time as e . The idea behind the notion of the interference number is to measure, to some extent, the amount of interference generated in the network by a communication going on along a certain link. Links are ordered in decreasing order, i.e. links which generate relatively higher interference (which are more difficult to schedule) are scheduled first. Note that other criteria, such as traffic demand, could have been used to order links. The approximation bound proved in the following is valid independently of the actual ordering used for the links.

After the links have been ordered, they are scheduled in a greedy way, starting from the ‘largest’ link. A selected link e_i is scheduled in the first w_{e_i} available slots such that no conflict with previously scheduled transmissions occurs according to the physical interference model. If there are not enough slots in the current schedule, empty slots are added at the end of the schedule, and link e_i is scheduled in these slots. It is immediate to see that Algorithm GreedyPhysical computes a feasible schedule S_{Phys} of some length T_{Phys} according to the physical interference model.

We now prove that the schedule computed by GreedyPhysical is at most a factor $O(n^{1-\frac{2}{\psi(\alpha)+\epsilon}} (\log n)^{\frac{2}{\psi(\alpha)+\epsilon}})$ away

from the optimal schedule, where $\epsilon > 0$ is an arbitrarily small positive constant and $\psi(\alpha)$ is a constant which depends on α . This result holds under the assumption that network nodes are distributed uniformly at random in a square of unit area⁵, and that radio signal propagation obeys the log-distance path model with path loss exponent α . Similarly to [8], we assume $\alpha > 2$.

Random node deployment is representative of a scenario in which positioning of the wireless routers forming the mesh is not carefully planned, as is typically the case in community networking. In this scenario, users in a community (neighborhood, or even an entire city area) decide to share their wireless resources (access points) so that to create a Wireless Community Network (WCN). Several WCN are being created worldwide (see [30] for a list), the most well known examples being the Seattle and Houston Wireless initiatives [11, 27]. In WCNs, network deployment is typically an unplanned and uncoordinated process, since users autonomously decide to add or remove wireless resources from the network. As a consequence, the topology of the resulting mesh network is typically unstructured. Note that recent research has observed that unplanned mesh networks also have the potential to fulfill user requirements (in terms of perceived bandwidth) [4].

The following details the uniform random assumption and corresponding critical transmission range requirement needed to prove our approximation result:

- A1 a number n of nodes is distributed uniformly at random in a square area of unit side. All the nodes have the same transmission range r . In order to ensure connectivity with high probability⁶ of the generated mesh network, we set $r = r(n)$ to the critical transmission range for connectivity, i.e. (see [26], Corollary 4.1.2, pg. 42)

$$r(n) = \sqrt{\frac{\ln n + f(n)}{\pi n}}, \quad (1)$$

where $f(n)$ is an arbitrary function such that $f(n) \rightarrow \infty$ as $n \rightarrow \infty$.

Before presenting the proof, we give a high level sketch intended to provide an intuitive understanding. The main difficulty in trying to prove an approximation bound for the scheduling problem with the physical interference model lies in the fact that any given scheduled link e potentially conflicts with every other link in the network. This is because, even if the interference induced by the sender/receiver side of e might be very low on a far away link e' , there might be many interferers in the network, and the small interference induced on e' by e can be sufficient to drive the SINR at one of the endpoints of link e' below the acceptable threshold β . Hence, for any scheduled link e , we cannot separate the deployment area into a ‘close-in’ region (where the interference generated by e matters) and a ‘far away’ region (where the interference generated by e is negligible). This separation between ‘close-in’ and ‘far away’ region is a fundamental building block in the proof of the approximation bound for the scheduling problem with the protocol interference model reported in [1]. To circumvent this problem, we

⁵With some minor technicalities, the proof can be easily adapted to the case where the deployment area is a disk of unit area.

⁶In this paper, w.h.p. means with probability that converges to 1 as $n \rightarrow \infty$.

note that while the above observation holds true in a general mesh network scenario where the position and transmit power of nodes can be arbitrary, it might be the case that under certain circumstances we are indeed able to separate the deployment area into a ‘close-in’ and a ‘far away’ region even in the physical interference model. The idea is to use a *close-in radius* r_{in} centered at each receiver of a communication, and to show that, under Assumption A1, the aggregated interference level generated by nodes outside r_{in} can be neglected. Then we can use an argument similar to the one used in [1] to prove the approximation bound: for any scheduled link e , at most all links within the close-in region centered at e are silenced by the transmission along e , while all links outside the close-in region are not influenced by the transmission along e . Since GreedyPhysical, when considering link e , schedules this link as soon as possible, and scheduling e in a certain slot possibly prevents only links within the close-in region of e to be scheduled concurrently, an approximation bound on GreedyPhysical’s performance can be obtained by counting the maximum number of links that are silenced by a transmission along any link, which is upper bounded by the number of nodes within its close-in region.

We start with some lemmas needed in the proof.

LEMMA 1. Assume the deployment area $A = [0, 1]^2$ is divided into $C = \frac{1}{r^2}$ square cells of side $r = r(n)$, and that we are under the conditions of assumption A1. Let us set $f(n) = \bar{f}(n) = (\pi - 1) \ln n$ in (1). Then, the number of nodes in every one of the C cells is $\Theta(\log n)$ w.h.p.

PROOF. Let us first prove that the number of nodes in every cell is $O(\log n)$ w.h.p. We use the following result from occupancy theory (see [14], Thm. 2, pg. 96):

PROPOSITION 2. Assume n nodes are distributed uniformly at random into C cells, and let $\eta = \frac{n}{C}$. If $\frac{\eta}{\ln C} \rightarrow x$ as $n, C \rightarrow \infty$ and $k = k(\eta, C)$ is chosen so that $Cp_k \rightarrow \lambda$, where x and λ are positive constants and $p_k = \frac{\eta^k}{k!} e^{-\eta}$, then

$$\text{Prob}\{Max \leq k + g\} \rightarrow \exp\left\{-\frac{\lambda\gamma^{g+1}}{1-\gamma}\right\},$$

where Max denotes the number of nodes in the maximally occupied cell and γ is the root of the equation

$$\gamma + x(\ln \gamma - \gamma + 1) = 0$$

in the interval $0 < \gamma < 1$.

Setting C and $r = r(n)$ as in the statement of Lemma 1, we have that $\eta = \ln n$ and $\frac{\eta}{\ln C} \rightarrow x = 1$ as $n, C \rightarrow \infty$. By setting $k = \ln n$, we have $Cp_k \rightarrow \lambda = e$, and we can apply the result above to prove that

$$\text{Prob}\{Max \leq \ln n + g\} \rightarrow \exp\left\{-\frac{e\gamma^{g+1}}{1-\gamma}\right\}. \quad (2)$$

Since $x = 1$ implies that $\gamma = \frac{1}{e}$, we can rewrite (2) as follows

$$\text{Prob}\{Max \leq \ln n + g\} \rightarrow \exp\left\{-\frac{1}{e^{g-1}(e-1)}\right\}. \quad (3)$$

It is easy to see that the term on the right hand side of (3) converges to 1 if and only if $g = g(n)$ is an arbitrary function of n such that $g(n) \rightarrow \infty$ as $n \rightarrow \infty$. The proof of the first part of the theorem follows by selecting as $g(n)$ an arbitrary

function in $o(\log n)$ such that $g(n) \rightarrow \infty$ as $n \rightarrow \infty$ (e.g., $g(n) = \ln \ln n$).

The proof that the number of nodes in every cell is $\Omega(\log n)$ w.h.p. is similar, using Thm. 5, page 111, of [14]. \square

Before presenting the next lemma, we recall that in this section we are assuming $\alpha > 2$.

LEMMA 2. Assume that Assumption A1 holds, and let $f(n)$ in (1) be defined as in Lemma 1. Consider a family of functions $g_i(n)$ defined as follows:

$$g_i(n) = \left(\frac{n}{\ln n} \right)^{\frac{1}{2} - \frac{1}{i}}.$$

If $i > \psi(\alpha)$, where $\psi(\alpha)$ is a constant that depends on α defined as follows:

$$\psi(\alpha) = \frac{1}{2} + \frac{\sqrt{9\alpha^2 - 20\alpha + 4}}{2(\alpha - 2)},$$

then the total interference level generated by transmitters at distance $d(n) \cdot r(n)$ from a certain receiver, with $g_i(n) \ll d(n) = O(g_{i+1}(n))$, converges to 0 as n goes to infinity, w.h.p.

PROOF. Let us consider an arbitrary intended receiver u , and let v be the intended transmitter for u . Since the communication range is r , the distance between u and v is at most $r = r(n) = \sqrt{\frac{\ln n}{n}}$. Let us consider the contribution to the interference level measured at u coming from nodes at distance $d(n) \cdot r(n)$ from u , where $d(n)$ is defined as in the statement of the lemma. Each of these nodes w_j contributes an interference level of $P/(d(n)r(n))^\alpha$ at u , where P is the transmit power of the node. Since we are considering the SINR ratio, and P is the same for all the nodes in the network under assumption A1, we can rewrite the interference level generated at u by any of the w_j 's as $1/d(n)^\alpha$. Let us now evaluate how many nodes are within distance $d(n) \cdot r(n)$ from u . By using the same cell subdivision as in the proof of Lemma 1, we can prove that a circle of radius $d(n) \cdot r(n)$ centered at u intersects $\Theta(d^2(n))$ cells of side $r(n)$ (this is easy to see by circumscribing a square to a circle of radius $d(n) \cdot r(n)$). Hence, by Lemma 1 the number of nodes within distance $d(n) \cdot r(n)$ from u is $\Theta(d^2(n) \log n)$ w.h.p. In order to prove the Lemma, we have then to show that

$$\lim_{n \rightarrow \infty} \frac{d^2(n)}{d^\alpha(n)} \ln n = 0$$

We observe that $d^2(n) = O(g_{i+1}^2)$ and $d^\alpha(n) \gg g_i^\alpha(n)$, so in order to prove the lemma it is sufficient to show that

$$\lim_{n \rightarrow \infty} \frac{g_{i+1}^2(n)}{g_i^\alpha(n)} \ln n = 0$$

The above can be rewritten as follows:

$$\lim_{n \rightarrow \infty} \frac{\ln n}{\left(\frac{n}{\ln n} \right)^{\frac{\alpha}{2} - \frac{\alpha}{i} - 1 + \frac{2}{i+1}}}$$

If $i > \psi(\alpha)$, then $\frac{\alpha}{2} - \frac{\alpha}{i} - 1 + \frac{2}{i+1} > 0$ and the above limit is 0. This completes the proof of the lemma. \square

For the next lemma, we need an additional definition.

DEFINITION 4. Consider an arbitrary set of links $E'' = \{e_1, \dots, e_k\}$ in the network, and a link $e_i = (u, v) \in E''$.

The set of transmitters for the endpoint u of link e , denoted by $T_u(E'')$, is defined as $\{x \in V | ((x, y) \in E'' - \{(u, v)\}) \text{ and } (P_u(x) \geq P_u(y))\}$.

LEMMA 3. Assume that Assumption A1 holds, and that $f(n)$ is set as in Lemma 1. Let us consider an approximate version of the physical interference model, in which the contribution to the interference level generated by all transmitters beyond distance $g_{i^*}(n) \cdot r(n)$ from a certain receiver is ignored, where $g_{i^*}(n)$ is defined as in the statement of Lemma 2 and $i^* = \psi(\alpha) + \epsilon$, for some arbitrarily small constant $\epsilon > 0$. Then, any set of transmissions which is feasible under this approximate physical interference model is also feasible under the complete physical interference model w.h.p.

PROOF. Let us consider the SINR at a receiver u , and consider a hierarchy \mathcal{G} of functions $g_{i^*}(n), g_{i^*+1}(n), g_{i^*+2}, \dots$, where $g_i(n)$ is defined as in Lemma 2. To ease notation, denote with $D(u, x)$ the ball of radius x centered at u , with $\Gamma(u, 0)$ the ball of radius $g_{i^*}(n) \cdot r(n)$ centered at u , and with $\Gamma(u, j+1)$ the annulus $D(u, g_{i^*+j+1}(n)) - D(u, g_{i^*+j}(n))$. Now, let E' be an arbitrary set of links in the network. By hierarchy \mathcal{G} , and observing that: 1) the transmitter/receiver distance is at most $r(n)$, and 2) the distance between a node in $\Gamma(u, j+1)$ and u is at least $g_{i^*+j}(n)$, we can decompose the contribution to the interference level measured at u according to \mathcal{G} as follows:

$$\text{SINR}(u) \geq \frac{P}{r(n)^\alpha} / \left[N + \sum_{w \in \Gamma(u, 0) \cap T_u(E')} \frac{P}{d_{uw}^\alpha} + \sum_{w \in \Gamma(u, 1) \cap T_u(E')} \frac{P}{(g_{i^*}(n)r(n))^\alpha} + \dots \right],$$

where d_{uw} is the distance between w and u . We can rewrite the right hand term of the above inequality as follows:

$$\left[\frac{Nr^\alpha(n)}{P} + \sum_{w \in \Gamma(u, 0) \cap T_u(E')} \left(\frac{r(n)}{d_{uw}} \right)^\alpha + \sum_{w \in \Gamma(u, 1) \cap T_u(E')} \frac{1}{g_{i^*}^\alpha(n)} + \dots \right]^{-1}. \quad (4)$$

By Lemma 2, we know that each of the terms

$$\sum_{w \in \Gamma(u, j) \cap T_u(E')} \frac{1}{g_{i^*+j}^\alpha(n)}$$

converges to 0 as $n \rightarrow \infty$. However, there are an infinite number of terms in the hierarchy \mathcal{G} , so this is not sufficient to prove that the sum of the terms converges to 0 as well, which would imply that the contribution to the interference of the nodes beyond distance $g_{i^*}(n) \cdot r(n)$ can be ignored.

However, let us consider the function $\bar{g}(n) = \left(\frac{n}{\ln n} \right)^{\frac{1}{2}}$. It follows directly that a ball of radius $\bar{g}(n) \cdot r(n)$ corresponds to the entire deployment area (in asymptotic terms), i.e. a ball with this radius includes all n nodes. Let us now consider the function $g_{i_\eta}(n) = \left(\frac{n}{\ln n} \right)^{\frac{1}{2} - \eta}$, for some $\eta > 0$. Consider the contribution to the interference level of all nodes beyond distance $g_{i_\eta}(n) \cdot r(n)$ from u . Then,

$$\lim_{n \rightarrow \infty} \frac{n}{g_{i_\eta}(n)^\alpha} = 0 \quad (5)$$

whenever $\eta < \frac{1}{2} - \frac{1}{\alpha}$. Now, denote by i' the minimum value in the sequence $i^*, i^* + 1, i^* + 2, \dots$ such that $\frac{1}{i'} < \frac{1}{2} - \frac{1}{\alpha}$ (note that this value always exists and is finite). By (5), we can truncate the hierarchy at i' . Thus, in the denominator of (4), the interference contribution of all nodes beyond distance $g_{i^*}(n) \cdot r(n)$ from u can be expressed as a *constant* sum of terms, each of which individually converges to 0 as $n \rightarrow \infty$. We have thus proved that the total contribution to the interference level measured at u due to transmitters located beyond distance $g_{i^*}(n) \cdot r(n)$ from u can be neglected w.h.p. It follows that, w.h.p., any set of feasible transmissions under the approximate physical interference model in which only the interference generated by all the transmitters within distance $g_{i^*}(n) \cdot r(n)$ from a receiver is also feasible under the complete physical interference model in which all transmitters are considered. \square

THEOREM 1 (APPROXIMATION BOUND). *Let G be a communication graph with given link demands and let G_{Phys} be its corresponding conflict graph. Let T_{opt} be the minimum possible value of T such that a schedule of length T is feasible for G under the physical interference model, and let T_{Phys} be the length of the schedule computed by GreedyPhysical. Under assumption A1, $\frac{T_{Phys}}{T_{opt}} \in O(n^{1 - \frac{2}{\psi(\alpha) + \epsilon}} (\log n)^{\frac{2}{\psi(\alpha) + \epsilon}})$, for any arbitrarily small constant $\epsilon > 0$, w.h.p.*

PROOF. By Lemma 3, for any scheduled link $e = (u, v)$ we can divide the deployment area into a ‘close-in’ region (the union B of the balls of radius $g_{i^*}(n) \cdot r(n)$ centered at u and v , respectively), in which the transmission along link e generates non-negligible interference level, and a ‘far away’ region (the rest of the deployment area, i.e. $A - B$) which is not affected by the transmission along e . Let us now consider a certain set E' of links which are assigned the same slot at the end of GreedyPhysical’s execution. It is immediate to see that E' is a feasible set of transmissions. Let $\{e_1, \dots, e_k\}$ be the links in E' , and consider the union R of their ‘close-in’ regions. We now upper bound the maximum number M of links that can be scheduled concurrently in R under the physical interference model. Since a node cannot transmit and receive in the same slot, it follows that an upper bound to M is given by the maximum number of nodes within R , which in turn is upper bounded by k times the maximum number of nodes within the ‘close-in’ region of e , for any possible scheduled link e . Observing that the ‘close-in’ region of any link in E' is at most as large as the union of two balls of radius $g_{i^*}(n) \cdot r(n)$, and that each of these balls contains $\Theta(g_{i^*}(n)^2 \log n) = O(n^{1 - \frac{2}{\psi(\alpha) + \epsilon}} (\log n)^{\frac{2}{\psi(\alpha) + \epsilon}})$ nodes w.h.p. by Lemma 1, we can conclude that $M \in O(k \cdot n^{1 - \frac{2}{\psi(\alpha) + \epsilon}} (\log n)^{\frac{2}{\psi(\alpha) + \epsilon}})$. Observe that GreedyPhysical schedules k transmissions in the slot, while at most $O(k \cdot n^{1 - \frac{2}{\psi(\alpha) + \epsilon}} (\log n)^{\frac{2}{\psi(\alpha) + \epsilon}})$ can be scheduled by any algorithm obeying the physical interference model, and that the above argument can be applied to every set of links scheduled by GreedyPhysical in the same slot. Furthermore, by the nature of GreedyPhysical, either R covers the entire region or any links outside of R have zero demand (otherwise GreedyPhysical would have scheduled at least one link outside of R in the slot). We can therefore conclude that $\frac{T_{Phys}}{T_{opt}} \in O(n^{1 - \frac{2}{\psi(\alpha) + \epsilon}} (\log n)^{\frac{2}{\psi(\alpha) + \epsilon}})$. \square

The approximation bounds obtained for some values of α are reported in Table 1. Note that we can get arbitrarily

close to these bounds (to an arbitrarily small constant $\epsilon > 0$).

α	bound
2.1	$n^{0.714} (\ln n)^{0.286}$
2.5	$n^{0.460} (\ln n)^{0.540}$
3	$n^{0.333} (\ln n)^{0.666}$
4	$n^{0.219} (\ln n)^{0.780}$
5	$n^{0.164} (\ln n)^{0.836}$
6	$n^{0.131} (\ln n)^{0.868}$

Table 1: Approximation bounds for various α .

THEOREM 2 (TIME COMPLEXITY). *Let $G = (V, E)$ be a communication graph with traffic demands w_e on each link; let $n = |V|$, $m = |E|$, and let $TD = \sum_{e \in E} w_e$ be the total traffic demand in the network. Then, Algorithm GreedyPhysical executed on G has $O(m \cdot TD \cdot n)$ time complexity.*

PROOF. Ordering the links at Step 2 requires $O(m \log m)$ time. For each considered link e_i , finding the first w_{e_i} available slots requires at most $O(TD \cdot n)$ time, since the maximum number of slots in the schedule is TD (which corresponds to scheduling each link sequentially), and verifying the feasibility of a set of transmissions requires $O(n)$ time – in fact, since any node can be either a transmitter or a receiver in a feasible slot, $O(n)$ links are present in a feasible set of transmissions. Steps 3a and 3b are repeated for each of the m links, hence the time complexity of step 3 is $O(m \cdot TD \cdot n)$. It follows that Algorithm GreedyPhysical has $O(m(\log m + TD \cdot n))$ time complexity, which is equivalent to $O(m \cdot TD \cdot n)$. \square

Considering that in the worst case $m \in O(n^2)$, the time complexity of GreedyPhysical can be rewritten as $O(n^3 \cdot TD)$. Concerning TD , we observe the following. Let l_{min} and l_{max} denote the minimum and the maximum of the load offered by nodes to the network. Under the assumption that $\frac{l_{max}}{l_{min}}$ is upper bounded by a constant, we have that the aggregated demand on each link of the network is $O(n)$. Since there are at most $O(n^2)$ links in the network, we have that $TD \in O(n^3)$, and GreedyPhysical’s time complexity is $O(n^6)$ for arbitrary networks. However, if tree-based routing is used to convey traffic to the gateway(s), there are $O(n)$ active links in the network, and TD becomes $O(n^2)$, implying an $O(n^4)$ time complexity. If balanced trees are used, there are $O(\log(n))$ levels in each tree, and the aggregated traffic at each level is $O(n)$. Hence, $TD \in O(n \log n)$, and the time complexity of GreedyPhysical is $O(n^3 \log n)$.

5. SIMULATION RESULTS

5.1 Simulation Setup

We present two sets of simulation results. The first set compares the schedule lengths generated by Algorithm GreedyPhysical and the algorithm of Alicherry, et al. [1], for the protocol interference model. The second set of results compares the throughput of Algorithm GreedyPhysical against a real protocol-interference-based approach, namely 802.11. The throughput simulations were done using the GTNeTS simulator [24], which is a packet-level simulator with a complete 802.11 MAC model. We did not compare GreedyPhysical against other physical-interference-based scheduling algorithms because, as detailed in related work, the existing

algorithms are not computationally feasible even for a small number of nodes. We also do not report throughput results of STDMA with protocol-interference-based scheduling, because we found that the schedules generated were almost always infeasible under a realistic interference model and the resulting throughputs were thus extremely low.

To simulate Algorithm GreedyPhysical, we modified GT-NeTS to include calculation of physical interference over all nodes (not just those within the carrier sensing range, as is done by most network simulators). For both sets of simulations, the approaches were evaluated under the log-normal shadowing model for signal propagation with path loss exponents of two and three. We also consider both planned (grid) and unplanned (uniform random distribution) deployments.

The following approach was used to generate the link demands for all simulations. Each node is assigned a demand that is uniformly distributed from 1 to 10 units. For each node, a shortest path to a nearest gateway node is found. The collection of shortest paths forms a set of shortest-path trees, one for each gateway. Demand is then aggregated on the links of these shortest-path trees, i.e. the demand on a link of the tree is the sum of the demands of all nodes in the sub-tree below that link. Note that we have not tried to jointly optimize schedule length and routing, nor to use an interference-aware routing mechanism. Both of these approaches are the subject of future research and could improve performance under physical interference even more compared to the results presented in this paper.

In order to evaluate the approaches for different node densities, we used the standard approach of fixing the number of nodes and scaling the size of the region in which the nodes are distributed. We used 100 nodes in the schedule length and throughput simulations. We also performed a set of simulations to evaluate the execution time of Algorithm GreedyPhysical in which the number of nodes ranged up to 2000.

We use the following notation for simulation parameters: n denotes the number of nodes, GW denotes the number of gateways, and α represents the path loss exponent.

5.2 Evaluation of Schedule Lengths

95% confidence intervals were computed for all data points reported in this section. The number of simulations was chosen in order to reduce the confidence interval to a meaningful value. We first study the performance improvement of scheduling with physical interference versus node density for a fixed number of gateways, and we then investigate how the results change as the number of gateways is varied.

Figure 4 shows the percentage improvement, relative to the protocol model, in the schedule length generated by Algorithm GreedyPhysical for a path loss of three and several different scenarios. The scenarios shown in the figure include a planned grid deployment, an unplanned deployment with node locations generated uniformly at random and independently, and the same unplanned deployment but with heterogeneous transmit powers. Figure 4 shows that performance improvements of up to 45% are achieved compared to the optimistic protocol-model-based schedule. The performance benefits of the physical model become smaller as the density increases, approaching a steady-state value of about 7% improvement. For very high densities, all nodes are within range of all gateways and the schedule generated under the protocol model is completely serialized. Such high

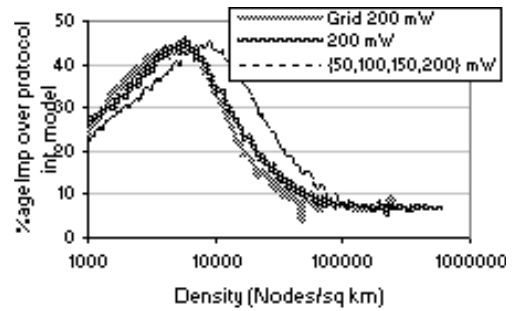


Figure 4: Schedule Length Improvement in Physical Interference Model: Log Normal Shadowing, $\alpha = 3$, $GW = 10$, $n = 100$

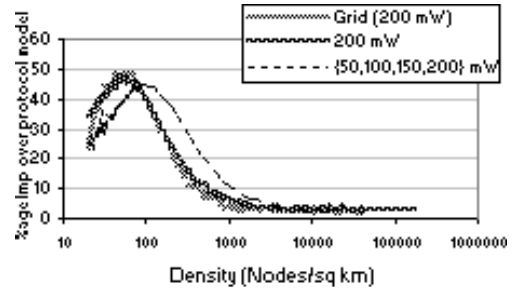


Figure 5: Schedule Length Improvement in Physical Interference Model: Log Normal Shadowing, $\alpha = 2$, $GW = 10$, $n = 100$

densities generate strong interference in the physical model also, but we are still able to achieve some concurrency in transmissions even in this situation. In terms of the different scenarios simulated, the relative performance compared to protocol-based scheduling for the grid deployment and the random deployment with common transmission power are similar. The curve for heterogeneous transmission powers has almost identical shape to the other two but is shifted right (to higher densities). Due to some nodes having lower transmission powers, a level of interference that is equivalent to the other two cases at a given density is generated at a higher density for the heterogeneous case.

Figure 5 shows the schedule length improvement with a path loss exponent of two. The results are similar to those for a path loss exponent of three. Peak improvements for the homogeneous settings are actually slightly higher, around 50%. However, for very high densities, the improvements are lower, only around 3%. Since a path loss of two is the worst case for physical interference (it generates the largest amount of interference for a given transmission), these results are a strong indicator of the potential performance benefits of scheduling with physical interference, particularly when node densities are not extremely high.

Figure 6 shows the percentage improvement for physical-interference-based scheduling versus the number of gateway nodes, for two different density values and a path loss of three. The figure shows that adding more gateways, up to a point, improves performance relative to the protocol model, particularly for higher densities. For the lower density value, the performance improvement stays fairly constant above 10 gateways, finally dropping off slightly beyond 70 gateways.

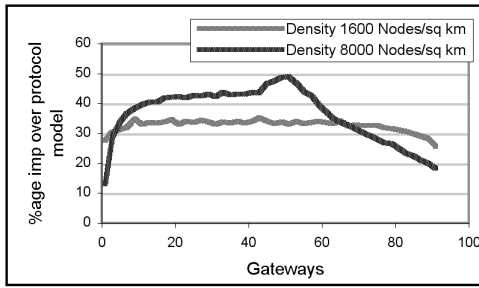


Figure 6: Schedule Length Improvement vs. Number of Gateways in Physical Interference Model: Log Normal Shadowing, $\alpha = 3$, $n = 100$

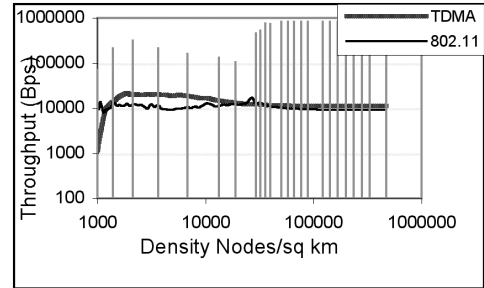


Figure 8: Throughput vs. Density for 802.11 and STDMA with GreedyPhysical Scheduling: Grid, Log Normal Shadowing, $\alpha = 3$, $GW = 10$, $n = 100$, data rate = 1MBps

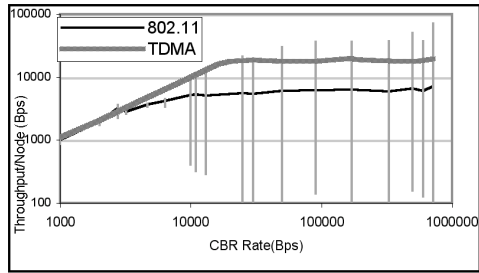


Figure 7: Throughput vs. Data Rate for 802.11 and STDMA with GreedyPhysical Scheduling: Grid, Log Normal Shadowing, $\alpha = 3$, $GW = 10$, $n = 100$, density = 8,000 nodes/sq km

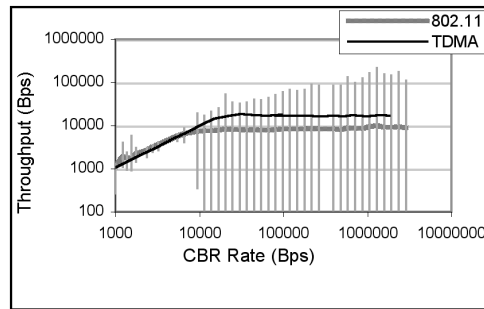


Figure 9: Throughput vs. Data Rate for 802.11 and STDMA with GreedyPhysical Scheduling: Uniform Random, Log Normal Shadowing, $\alpha = 3$, $GW = 10$, $n = 100$, density = 8,000 nodes/sq km

At the higher density value, the improvement with 10 gateways is 38%, which increases to a maximum of close to 50% with 50 gateways. Once each node is within one hop of a gateway, adding further gateway nodes does not improve performance relative to the protocol model. This point occurs earlier for higher densities, because each transmission reaches a higher number of nodes than a transmission that occurs at a lower density.

5.3 Evaluation of Throughput

The simulations in this section were done with GTNetS and use a path loss exponent of three, which is more representative of real environments than a path loss of two.

Figure 7 shows the average throughput per node of 802.11 and a STDMA approach, which is based on the schedule produced by Algorithm GreedyPhysical, versus the data rate of each node for a density of 8,000 nodes/km². Note that the figure uses a log-log scale. As can be seen in the figure, the throughputs of the approaches are identical at lower data rates, as is to be expected. However, with 802.11, the network reaches saturation significantly earlier and at a much lower throughput than with STDMA with GreedyPhysical scheduling. The throughput of STDMA is approximately three times that of 802.11 for high data rates.

Figure 8 shows the average throughputs of 802.11 and STDMA versus density, again on a log-log scale. For densities in the range of 1,400 to 10,000 nodes/km², STDMA with GreedyPhysical scheduling has average throughputs from 1.5 to 2 times better than are achieved with 802.11. At very high densities, STDMA's average throughput is about

25% higher than that of 802.11. Figure 8 also illustrates an important difference between 802.11 and STDMA for wireless mesh networks, related to fairness. The vertical bars in the figure show the range of node throughputs for 802.11, indicating that the minimum throughputs are either zero or very close to zero, while some nodes get many times the mean value. With 802.11, nodes that are far from a gateway are starved, while ones that are one hop from a gateway get very high throughput. This behavior is contrasted with that of the STDMA approach where every node achieves the same throughput.

Figures 9 and 10 show the throughput vs. data rate and throughput vs. density curves for the uniform random node deployment. The results for uniform random deployments are similar to those produced with the grid deployment. The throughput vs. data rate figure matches fairly closely Figure 7 for grid deployments. In this case, the performance of STDMA with physical scheduling is about twice that of 802.11 at high data rates, which is somewhat lower than with grid deployments. The throughput vs. density curves for uniform random deployments are very close to those for grid deployments. Again, the STDMA throughput is as high as twice that of 802.11 and at steady-state densities, it is 25% higher. The fairness problem with 802.11 is clearly illustrated in this scenario, as well.

These GTNetS results are somewhat optimistic for STDMA relative to a practical setting in that they ignore overhead for




Figure 10: Throughput vs. Density for 802.11 and STDMA with GreedyPhysical Scheduling: Uniform Random, Log Normal Shadowing, $\alpha = 3$, $GW = 10$, $n = 100$, data rate = 1MBps

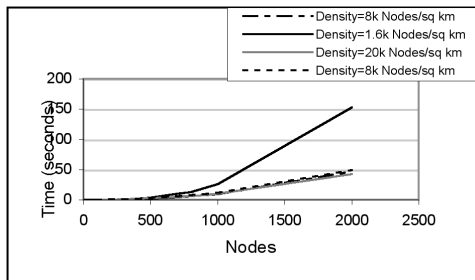


Figure 11: Execution Time of Algorithm GreedyPhysical vs. Number of Nodes: Uniform Random, Log Normal Shadowing, $\alpha = 3$, $GW = n/10$

achieving synchronization (required in case GPS is not available on all nodes) and the scheduler has global information about the topology and link demands. Nevertheless, the results suggest that significant throughput improvements relative to CSMA/CA (up to about three times) might be possible in certain situations if effective and low-overhead distributed scheduling protocols can be developed.

5.4 Execution Time of GreedyPhysical

One of the primary advantages of our scheduling approach is that it is the first physical-interference-based scheduler with a proven polynomial-time complexity. Figure 11 shows the execution time of Algorithm GreedyPhysical vs. the number of nodes for four different node densities. Even for 2000 nodes, the execution time is only about 2.5 minutes.

The execution time in Figure 11 seems to exhibit a polynomial growth rate as analyzed in Section 4. That analysis also pointed out that execution time is sensitive to the link demands. Figure 12 shows the total demand corresponding to the execution time values reported in Figure 11. At the lowest density, paths to a gateway require multiple hops for many nodes and this results in aggregated demand on the links near the gateway. This increases total demand substantially in this case and drives up the execution time accordingly. For the three higher densities, most nodes were within one hop of a gateway and, as a result, total demand remained lower, as did execution time.

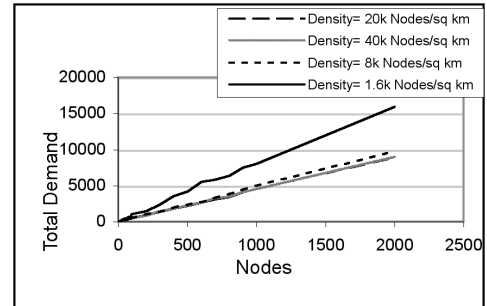


Figure 12: Total Demands for Figure 11

6. CONCLUSION

We have presented a polynomial-time algorithm for scheduling under the physical interference model, proved an approximation bound on its performance relative to optimal, and evaluated it extensively through simulation under representative wireless mesh network scenarios. There are several interesting avenues of future work opened by this research. The first is the possibility to integrate our efficient scheduling algorithm with routing to provide even better performance. One could consider these as a joint optimization problem on the length of the schedule, or simply use an interference-aware routing algorithm instead of the current minimum-hop routing algorithm employed in our work. Another interesting possibility is to develop a distributed scheduler based on our algorithm. While distributed scheduling algorithms that try to account for interference exist, none of these are based on a centralized algorithm with proven performance. A fundamental problem facing such an approach is how to gather interference measurements from nodes for which communication links do not exist. We have shown in this work that such nodes can not be safely ignored in constructing a schedule, and yet no work that we know of has addressed this problem.

7. REFERENCES

- [1] 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.
- [2] Azulstar Website: <http://www.azulstar.com/coverage-map.html>.
- [3] N. Ben Salem, J.P. Hubaux, "A Fair Scheduling for Wireless Mesh Networks", *Proc. IEEE Workshop on Wireless Mesh Networks (WiMesh)*, 2005.
- [4] J. Bicket, D. Aguayo, S. Biswas, R. Morris, "Architecture and Evaluation of an Unplanned 802.11b Mesh Networks", *Proc. ACM Mobicom*, pp. 31–42, 2005.
- [5] J. Gronkvist and A. Hansson, "Comparison Between Graph-Based and Interference-Based STDMA Scheduling," *Proc. MobiHoc 2001*, pp. 255–258.
- [6] J. Gronkvist, J. Nilsson, and D. Yuan, "Throughput of Optimal Spatial Reuse TDMA for Wireless Ad-Hoc Networks," *Proc. IEEE Vehicular Technology Conference*, pp. 2156–2160, 2004.
- [7] J. Gronkvist, "Traffic Controlled Spatial Reuse TDMA in Multi-hop Radio Networks," *Proc. Int'l.*

- Symp. on Personal, Indoor, and Mobile Radio Communications*, pp. 1203–1207, 1998.
- [8] P. Gupta and P.R. Kumar, “The Capacity of Wireless Networks,” *IEEE Trans. Info. Theory*, Vol. 46, No. 2, pp. 388–404, 2000.
- [9] P. Gupta, P.R. Kumar, “Critical Power for Asymptotic Connectivity in Wireless Networks”, *Stochastic Analysis, Control, Optimization and Applications*, Birkhauser, Boston, pp. 547–566, 1998.
- [10] I. Haratcherev, J. Taal, K. Langendoen, R. Lagendijk, H. Sips, “Automatic IEEE 802.11 rate control for streaming applications”, *Wireless Communications and Mobile Computing*, Vol. 5, pp. 421–437, 2005.
- [11] Houston Wireless Website: <http://www.houstonwireless.org/>.
- [12] K. Jain, J. Padhye, V. Padmanabhan, L. Qiu, “Impact of Interference on Multi-Hop Wireless Network Performance”, *Proc. ACM Mobicom*, pp. 66–80, 2003.
- [13] M. Kodialam, T. Nandagopal, “Characterizing the Capacity Region in Multi-Radio Multi-Channel Wireless Mesh Networks”, *Proc. ACM Mobicom*, pp. 73–87, 2005.
- [14] V.F. Kolchin, B.A. Sevast’yanov, V.P. Chistyakov, *Random Allocations*, V.H. Winston and Sons, Washington D.C., 1978.
- [15] P. Kyasanur, N. Vaidya, “Capacity of Multi-Channel Wireless Networks: Impact of Number of Channels and Interfaces”, *Proc. ACM Mobicom*, pp. 43–57, 2005.
- [16] P. Kyasanur, X. Yang, N. Vaidya, “Mesh Networking Protocols to Exploit Physical Layer Capabilities”, *Proc. IEEE Workshop on Wireless Mesh Networks (WiMesh)*, 2005.
- [17] P. Kyasanur, N. Vaidya, “Routing and Interface Assignment in Multi-Channel Multi-Interface Wireless Networks”, *Proc. IEEE WCNC*, 2005.
- [18] T. Moscibroda, R. Wattenhofer, “The Complexity of Connectivity in Wireless Networks”, *Proc. IEEE Infocom 2006*, to appear.
- [19] R. Nelson and L. Kleinrock, “Spatial-TDMA: A Collision-free Multihop Channel Access Protocol,” *IEEE Trans. on Communication*, Vol. 33, pp. 934–944, Sept. 1985.
- [20] D. Qiao, S. Choi, “Goodput enhancement of IEEE 802.11a wireless LAN via link adaptation”, *Proc. IEEE ICC*, pp. 1995–2000, 2001.
- [21] B. Raman, “Channel Allocation in 802.11-based Mesh Networks,” *Proc. IEEE Infocom 2006*, to appear.
- [22] B. Raman, K. Chebrolu, “Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks” *Proc. ACM Mobicom*, pp. 156–169, 2005.
- [23] A. Raniwala, T. Chiueh, “Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Networks”, *Proc. IEEE Infocom*, pp. 2223–2234, 2005.
- [24] G. Riley, “The Georgia Tech Network Simulator,” *ACM SIGCOMM MoMeTools Workshop*, 2003.
- [25] P. Santi, D.M. Blough, “The Critical Transmitting Range for Connectivity in Sparse Wireless Ad Hoc Networks”, *IEEE Transactions on Mobile Computing*, Vol. 2, n. 1, pp. 1–15, January-March 2003.
- [26] P. Santi, *Topology Control in Wireless Ad Hoc and Sensor Networks*, John Wiley & Sons, Chichester, UK, 2005.
- [27] Seattle Wireless Website: <http://www.seattlewireless.net/>.
- [28] O. Somarriba, *Multihop Packet Radio Systems in Rough Terrain*, Tech.lic. Thesis, Radio Communication Systems, Royal Institute of Technology, Stockholm, Sweden, Oct. 1995.
- [29] J. Tang, G. Xue, W. Zhang, “Interference-Aware Topology Control and QoS Routing in Multi-Channel Wireless Mesh Networks”, *Proc. ACM MobiHoc*, pp. 68–77, 2005.
- [30] Wireless Communities Website: <http://wiki.personaltelco.net/index.cgi/WirelessCommunities>.
- [31] Z. Wu and D. Raychaudhuri, “D-LSMA: Distributed Link Scheduling Multiple Access Protocol for QoS in Ad-hoc Networks,” *Proc. Globecom 2004*, pp. 1670–1675.
- [32] Y. Yuan, H. Yang, S. Wong, S. Lu, W. Arbaugh, “ROMER: Resilient Opportunistic Mesh Routing for Wireless Mesh Networks”, *Proc. IEEE Workshop on Wireless Mesh Networks (WiMesh)*, 2005.