Engineering wireless mesh networks

Girard, Andre.; Rosenberg, Catherine.; Luo, Jun.

2008

Rosenberg, C., Luo, J., & Girard, A. (2008). Engineering wireless mesh networks. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (19th:2008:Cannes)

<https://hdl.handle.net/10356/94404>

<https://doi.org/10.1109/PIMRC.2008.4699918>

IEEE International Symposium on Personal, Indoor and Mobile Radio Communications © 2008 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE. This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder. <http://www.ieee.org/portal/site>.

Downloaded on 26 Apr 2023 20:36:19 SGT

Engineering Wireless Mesh Networks

(Invited Paper)

Catherine Rosenberg Jun Luo
Department of Electrical and Computer Engineering
University of Waterloo
Waterloo, Ontario, Canada N2L 3G1
Emails: {cath, j7luo}@engmail.uwaterloo.ca

André Girard
INRS-EMT and GERAD
Montréal, Québec, Canada H5A 1K6
Email: andre.girard@gerad.ca

Abstract—Wireless mesh networks are considered as a potentially attractive alternative to provide broadband access to users. They have been studied extensively by the research community since they raised a lot of new issues due to their unique characteristics. Here, we focus on scenarios where these networks are installed and managed to provide broadband access to a set of fixed nodes. While a lot of research has been done on this type of networks, there are very few insightful engineering results that can help network operators deploy and manage such networks. It is the objective of this paper to present some major engineering insights on such networks. We limit our scope to networks that are single rate and in which all nodes use the same transmit power. In particular, we quantify the advantage of multi-hop over single-hop. We illustrate the importance of multi-path routing over single path routing, and of optimal routing versus min-hop routing. We revisit the notion of spatial reuse. Finally we present results showing the importance of selecting an appropriate interference model.

Index Terms—Wireless mesh network, throughput, interference model, routing, scheduling.

I. INTRODUCTION

Wireless mesh networks (WMN) are considered as a potentially attractive alternative to provide broadband access to users. They have been studied extensively by the research community since they raised a lot of new issues due to their unique characteristics. We focus on scenarios where such networks are installed and managed to provide broadband access, through a gateway connected to the Internet, to a set of fixed mesh routers called nodes in the following, each serving many users (see Fig. 1 for a typical scenario). These networks are

are expected to carry high throughput. Hence, they need to be “configured” rather carefully to deliver reasonable performances. In particular, such networks will most probably use conflict-free scheduling as opposed to random access protocols since it is well-known that the latter yield rather poor overall throughput. Configuration could mean many different things in different contexts. We assume that the positions of the nodes and the gateway are known in advance, that the traffic flows are either from the nodes to the gateway or/and from the gateway from the nodes (not from a node to another node), and hence configuration refers mostly to admission control, routing (since the network is multi-hop), power and rate control, and scheduling. Because the traffic is aggregated, the network is “managed”, and the nodes are fixed, which means that the channel gains can be assumed quasi time-invariant, (see [4]), the configuration is static and will change only when new nodes are installed or nodes fail.

From a modeling standpoint, many interesting problems can be formulated to find the “optimal” configuration(s) of a wireless mesh network. Clearly the notion of optimality depends on the problem being defined and the most general problem is a joint routing, scheduling, power control, rate control and even (nodes and gateway) positioning problem. The reader is referred to [7], [1], [10] for papers on that topic. In some of these studies, simplistic assumptions on the interference model are taken. In others, sub-problems are considered (e.g., the routing is fixed or there is no power control). However, most of the studies usually stop short of showing the engineering insights that such models can provide to WMN operators. In [8], we have formulated and solved a joint routing, power and rate control and scheduling problem based on max-min throughput. In [9], we have developed very powerful computational tools to solve this joint problem. These tools allow us to present here optimal configuration results for networks of up to 50 nodes, which we believe is a first. It is the objective of this paper to present the main engineering insights that were obtained as a results of our studies reported in [8] and [9]. We limit our scope to networks that are single rate and in which all nodes use the same transmit power. We focus on the following four issues:

1. Can we quantify the advantage of multi-hop over single-hop?

Fig. 1. A WMN with uplink and downlink flows.

usually multi-hop, power constrained, interference-limited and

2. How more efficient is multi-path routing versus single path routing and how good is min-hop routing when compared with optimal routing?
3. What is the relationship between spatial reuse and performance?
4. What is the impact of different interference models on performance and configurations?

We organize the remainder of this paper as follows. In Section II, we present our modeling framework. Detailed studies to give insights on the above four questions are presented in Sections III to VII. We finally conclude our paper in Section VIII.

II. MODELING FRAMEWORK

We model a WMN as a set $\mathcal{N}$ of *nodes* and a set $\mathcal{L}$ of *directional links*, with $|\mathcal{N}| = N$ and $|\mathcal{L}| = L$. Each node $i \in \mathcal{N}$ has a location (x_i, y_i) . We denote by $\mathcal{L}_i$ the set of links incident (inbound or outbound) to a node i . A link $l \in \mathcal{L}$ is identified by its source-destination pair. Let $\mathcal{F} : |\mathcal{F}| = F$ denote the set of *flows*. A flow $f \in \mathcal{F}$ is identified by its source-destination pair (f_s, f_d) and has a rate λ_f . In the following subsections, we will describe the network operations, the physical layer model, the interference models, and finally the joint routing and scheduling problem that we are solving to configure the network. The descriptions are based on the assumptions made earlier (single power, single rate, etc.). The models proposed and analyzed in [8], [9] are much more general.

A. Network Operation

To bridge the gap between the optimization problem that we will present in Section II-D, which is based on a fluid model, and the “realistic” WMN that we want to configure, we present below a brief description of the network operation as we see it and highlight how our model departs from this reality. We assume that time is slotted, exactly one packet can be transmitted per time slot and that there is no packet loss. The time slots are organized in frames.

Routing: We consider multi-path routing in this paper except when otherwise stated. We denote by $\mathcal{R}_f$ the set of all routes that can be used by f and by $\mathcal{R}_f^l$ the set of routes of f going through link l . The fraction of flow f routed on $r \in \mathcal{R}_f$ is ϕ_f^r , hence $\sum_{r \in \mathcal{R}_f} \phi_f^r = 1$. Let $\phi = [\phi_f^r]_{r \in \mathcal{R}_f, f \in \mathcal{F}}$.

Scheduling: Denote by $\mathcal{S}$ the power set of $\mathcal{L}$. A link schedule is an $|\mathcal{S}|$ -dimensional vector $\alpha = [\alpha_s]_{s \in \mathcal{S}}$ such that $\alpha_s > 0$ if $s \in \mathcal{S}$ is scheduled, otherwise $\alpha_s = 0$. We interpret α_s as the fraction of time slots in the frame allocated to a link set s . We will define later the notion of independent sets and $\alpha_s > 0$ only if s is one. We require that in a frame $\sum_{s \in \mathcal{I}} \alpha_s \leq 1$. Our problem formulation defined later does not explicitly assign time slots within a frame to independent sets but computes the values of α_s . Although there are many ways to assign time slots within each frame, a specific assignment does not affect network throughput although it does have an impact on other QoS requirements such as delay and jitter. We do not address the scheduling problem to this level of details

in this paper; interested readers are referred to [2] and the references therein.

Admission control: The joint routing and scheduling problem that we will formulate in Section II-D yields the optimal rate vector λ^* used for admission control, i.e., a flow f is rate-controlled to λ_f^* .

B. Physical Layer Model

Each link $l \in \mathcal{L}$ is identified by $o(l)$ its transmitter and $d(l)$ its receiver. P is the transmit power used by all nodes (including the gateway) and c denotes the data rate of the links in bits per second since we have assumed that there is only one modulation/coding scheme.

We assume that in a given time slot, a packet transmission on a link l is successful if the *signal to interference plus noise ratio* (SINR) is greater than a threshold β , namely:

$$\gamma_l = \frac{G_{ll}P}{N_0 + \sum_{l' \in \mathcal{A}_l} G_{l'l}P} \geq \beta \quad \forall l \in \mathcal{L} \quad (1)$$

Here G_{ll} denotes the channel gain on l , $G_{l'l}$ the channel gain from $o(l')$ to $d(l)$, $\mathcal{A}_l$ is the set of links $l' \neq l$ that are active in the time slot under consideration, and N_0 is the average thermal noise power in the operating frequency band. The channel gain of a link l of length d is assumed to be given by $K_l(d/d_0)^{-\eta}$, where K_l is taken to be 1 here without loss of generality, d_0 is the close-in reference distance and η is the path loss exponent.

C. Interference Models

We now present several interference models that yield different link conflict structures. The essence of each conflict structure is the concept of *independent set* (ISet)¹: a set of links that can operate at the same time without interfering with each other. We denote by $\mathcal{I}$ the set of all ISets and by $\mathcal{I}_l$ the set of ISets that contain link l .

1) *Node Exclusion (NX) Model:* Two links $l = (i, j)$ and $l' = (i', j')$ do interfere with each other if $i = i' \vee i = j' \vee j = i' \vee j = j'$. Therefore, a set $s \subseteq \mathcal{L}$ is an ISet only if:

$$i \neq i' \wedge i \neq j' \wedge j \neq i' \wedge j \neq j' \quad \forall l, l' \in s. \quad (2)$$

Remark: This NX model is **common** to all other interference models, as the radio constraints it represents, namely that a node cannot transmit and receive at the same time, transmit to multiple nodes at the same time or receive from different nodes at the same time, are our assumption.

2) *Interference Range Model:* This model has been proposed by [7]. Denote by d_{ij} the distance between nodes i and j . For two links $l = (i, j)$ and $l' = (i', j')$, let $d(l', l) = d_{i'j}$. Given the transmit power P and the SINR threshold β , we denote by r the maximum transmission range $r = d_0(\beta N_0/P)^{-1/\eta}$. We also define a parameter $\sigma \geq 1$ to represent the fact that the interference range is usually

¹This terminology could be slightly misleading since it is not equivalent to the independent set concept in graph theory. However, we use it in order to be consistent with the literature.

larger than the transmission range. Therefore, the links l and l' belong to the same ISet s iff the NX condition (2) is met and both conditions $d(l', l) > \sigma r$ and $d(l, l') > \sigma r$ hold.

Note that under this interference model, a set of links is an ISet if the two conditions above are met for all pairs of links in the set.

3) *Capture Threshold Model*: This is the model used in the *ns2* network simulator [3]. In this model, taking into account our assumption that all the transmit powers are the same, link l' interferes with link l if

$$\frac{G_{ll'}}{G_{ll}} < \beta \quad (3)$$

Therefore, the links l and l' belong to the same ISet s iff the NX condition (2) is met and both conditions $G_{ll'}/G_{ll} \geq \beta$ and $G_{ll''}/G_{ll'} \geq \beta$ hold.

Note that under this interference model, a set of links is an ISet if the two conditions above are met for all pairs of links in the set.

4) *Additive Interference Model*: Whereas the two previous interference models are only concerned with pairwise (binary) link conflict, this model [5] captures a more realistic situation: the interference to a certain link is the cumulative interference from multiple links that are active at the same time. Hence, under this interference model, a set $s \subseteq \mathcal{L}$ is an ISet iff:

$$\frac{G_{ul}P}{N_0 + \sum_{l' \in s: l' \neq l} G_{ul'}P} \geq \beta \quad \forall l \in s. \quad (4)$$

and the NX condition (2) is met.

D. Joint Routing and Scheduling Problem (JRSP)

We are now ready to formulate a class of joint routing and scheduling problems. Note that since we have assumed that all nodes use the same transmit power and the same modulation/coding scheme, this class of problems does not include power control (see [8] for a more general version of this problem). While the network operation was described above in terms of time slots and packets, the Joint Routing and Scheduling problem (JRSP) is based on a fluid model. The inputs of this class of problems are:

- $\mathcal{N}$ a set of N nodes along with its location vector $[(x_1, y_1), \dots, (x_N, y_N)]$, where $(x_i, y_i) \in \mathbb{R}^2$ for $i \in \mathcal{N}$.
- $\mathcal{L}$ a set of L links given a transmit power P and a modulation/coding scheme characterized by a normalized rate and a SINR threshold β .
- $\mathcal{F}$ a set of F flows characterized by its vector $[(f_s^1, f_d^1), \dots, (f_s^F, f_d^F)]$, where $(f_s^i, f_d^i) \in \mathcal{N}^2$ and $f_s^i \neq f_d^i$ for $f^i \in \mathcal{F}$, and its rate vector $\lambda = [\lambda_1, \dots, \lambda_F]$.
- $\mathbf{U}(\lambda)$ a utility (vector) function that measures the satisfaction (of the flows) on the rate vector λ .

Given a certain interference model (taken from those defined in Section II-C), we are interested in **maximizing** the utility $\mathbf{U}(\lambda)$ by optimizing over all possible routing ϕ and scheduling α . The solution of the problem includes the following information:

- λ^* the flow rate allocation vector $(\lambda_1^*, \dots, \lambda_F^*)$.
- ϕ^* the routing control vector along with the corresponding routing paths for all flows.
- α^* the scheduling vector along with the corresponding independent sets $\mathcal{I}^* = \{s | s \in \mathcal{I}, \alpha_s > 0\}$.

The detailed mathematical formulation of the *joint routing and scheduling problem* (JRSP) is omitted; interested readers are referred to our previous work [8] in which we focus on maximizing the minimum flow rate. In the following, rather than focusing on how to solve JRSP, we investigate in detail the engineering insights provided by the solutions of JRSP for different network topologies that are described in the next section. These solutions are obtained thanks to our recently developed computation tools [9].

III. NETWORK SCENARIOS AND PRELIMINARY RESULTS

Among the many network topologies that we have studied, we choose the four randomly generated ones shown in Fig. 2 as representatives. For each network, we call the flow pattern

Fig. 2. Four randomly generated WMN topologies: (a)-(b) 30-node and (c)-(d) 50-node

in which every node has a flow to the gateway (represented by a “square” node in the picture) as *converging* and the pattern in which the gateway has a flow to each node as *diverging*. We assume that $d_0 = 0.1\text{m}$ and $\eta = 3$. As mentioned earlier, we consider one particular utility $\mathbf{U}(\lambda) = \min(\lambda_1, \dots, \lambda_F)$, and we call λ^* the *max-min throughput* of the WMN. Also, we fix the rate $c = 1$, take $\beta = 6.4\text{dB}$, and we investigate the optimal throughput as a function of P .

The main engineering insights for this type of networks that were reported in [8] can be summarized as follows:

1. The max-min throughput of such networks is a non-decreasing function of P . Hence, even if WMNs are interference-limited, it pays to use high transmit power.

2. The largest achievable max-min throughput (recall that it is a per-flow throughput) is $1/N$ if the flow pattern is diverging or converging and $1/2N$ if the flow pattern is both diverging and converging. This max-min throughput can be achieved in a single-hop setting if the transmit power is greater or equal to $P_{SH} = \beta N_0 (D/d_0)^n$ where D is the largest distance between the gateway and a node. Note that the throughput is limited by the fact that the gateway cannot receive (or transmit) more than one packet at a time.
3. For a given P , the max-min throughput can usually be obtained for more than one optimal configuration. Usually these configurations are so complex that no simple rule can be deduced from them. For example, it is not possible to say in general whether it is better to use power for range or for making the links more robust against interference, whether multi-path routing is crucial, and whether spatial reuse is an important indicator of good performance.

In the next sections, we provide more insights by addressing the four questions posed in the introduction. Except when stated otherwise, all the results are obtained with the additive interference model.

IV. THE MULTI-HOP ADVANTAGE

Given a WMN of N nodes and one gateway, multi-hop communication obviously allows connectivity at much lower transmit powers than single-hop communication. In this section, we will discuss and quantify another advantage of multi-hop over single-hop under the additive interference model: multi-hop communication enables us to obtain the maximum achievable throughput as defined above at much lower transmit power than single-hop communication. Hence, using multi-hop, we gain both by providing connectivity at low power, something which we cannot do with single-hop, and by offering the maximum achievable throughput at much lower power which comes at the cost of a more complex network operation. Single-hop communication is much simpler since it does require routing and the scheduling is a simple round-robin while using multi-hop involves the need for routing and more complex scheduling. However the results below indicate that this complexity is worth it.

Let P_{SH} be the transmit power that allows every node to have a single-hop connection with the gateway, and let $\bar{P}$ be the minimum transmit power for which the maximum achievable throughput can be obtained via multi-hopping. Both P_{SH} and $\bar{P}$ can be found at the intersection of the two vertical lines shown in Fig. 5 with the x -axis. We characterize this “multi-hop advantage” by $P_{SH}/\bar{P}$. Table I shows the above quantities for different networks. We see that multi-hop networking achieves the maximum achievable max-min throughput with a transmit power often 4 (or more) times lower than the power needed for single-hop communication. This is made possible by allowing spatial reuse, i.e., the activation of more than one link at a time. Using single-hop communication at P_{SH} , the bottleneck is the gateway as only one link can be

Network	P_{SH} (dBm)	$\bar{P}$ (dBm)	$P_{SH}/\bar{P}$ (dB)
Rand30a (converging)	-22.50	-28.25	5.75
Rand30a (diverging)	-22.50	-29.00	6.50
Rand30b (converging)	-23.25	-28.75	5.50
Rand30b (diverging)	-23.25	-28.75	5.50
Rand50a (converging)	-18.25	-26.25	8.00
Rand50a (diverging)	-18.25	-25.50	7.25
Rand50b (converging)	-19.00	-25.75	6.75
Rand50b (diverging)	-19.00	-27.25	8.25

TABLE I
MULTI-HOP ADVANTAGE: P_{SH} , $\bar{P}$, AND $P_{SH}/\bar{P}$.

active at a point in time. Surprisingly, as we will show in Section VI, this result is obtained with relatively low spatial reuse.

Note that Table I has been obtained assuming that $\beta = 6.4$ dB. Table II shows, for a particular network, that the multi-hop advantage as defined above is rather sensitive to the value of β . This is not surprising since the larger the β the lower the potential for spatial reuse.

β (dBm)	P_{SH} (dBm)	$\bar{P}$ (dBm)	$\frac{P_{SH}}{\bar{P}}$ (dB)
6.4	-22.50	-28.25	5.75
9.4	-19.50	-24.25	4.75
11.2	-17.75	-18.75	1.00
16.4	-12.50	-12.50	0.00
18.2	-10.75	-10.75	0.00

TABLE II
MULTI-HOP ADVANTAGE AS A FUNCTION OF β FOR RAND30A.

V. THE MULTI-PATH ADVANTAGE AND WHAT ABOUT MIN-HOP ROUTING?

In this section, we are trying to address the following three questions:

1. How much do we gain in throughput by allowing each flow to be routed on as many routes as necessary?
2. Is min-hop routing a good routing scheme for scheduled mesh network?
3. Is the impact of cross-layer design on performance important?

To answer the first question, we formulated a single-path version of our JRSP. The problem becomes an integer program that is much more difficult to solve and requires a different set of computational tools. Figure 3 shows the max-min throughput obtained for a 30-node randomly generated network for both the single path and the original multi-path problem. Clearly, multi-path does not bring a significant increase in throughput since the single-path max-min throughput is never more than 2% below the multi-path value. This is true for all the scenarios that we have studied with converging and diverging flow patterns.

Answering the two other questions is much more difficult. We try to quantify the importance of cross-layer design through the choice of the routing. Many networks use min-hop routing because it is simple to compute and implement. This

Fig. 3. Optimal multi-path and single path routings vs. min-hop routings: Rand30a with converging traffic.

can be done either without any consideration for lower layers, for example by using a simple Dijkstra's algorithm or by using some information about the lower layers to find a "good" min-hop solution. We want to compare these two options with the multi-layer design that we have presented here.

Note that there might exist many min-hop paths between a source-destination pair. We computed the max-min throughput for a min-hop routing obtained by using Dijkstra's algorithm, i.e., we formulated a pure scheduling problem by fixing the routing in JRSP. This min-hop is represented by the lowest curve in Fig. 3. We have also formulated a "cross-layer" min-hop problem in which we try to compute for each P the best possible min-hop path for each flow. This problem is computationally hard to solve and this is why the corresponding curve in Fig. 3 is limited to a small portion of the power range. From this, we can see that simple min-hop routing can be very inefficient while the "best" min-hop routing, i.e., the one selected by taking cross-layer into consideration, is much better but still somewhat far from optimal.

VI. REVISITING SPATIAL REUSE

It is a common belief that the advantage of multi-hopping stems from spatial reuse and that the more spatial reuse the better. In this section, we revisit this notion in some details. Spatial reuse is the ability for a network to schedule multiple links at once without creating harmful interference. In that sense, it is related to the size of the independent sets (ISets). For a given P , we can compute the feasible links, i.e., those that have a length $d \geq d_0(\beta N_0/P)^{-1/\eta}$, and from them the ISets using (4) and (2). As discussed in [8], the number of ISets is non-decreasing with P . However, we found that for a large number of randomly generated networks, the maximum size of an ISet is constant for all P 's that yield connectivity. This is not true for grid networks for which the maximum size of an ISet is very dependent on P . For example, for the 30-node randomly generated networks illustrated in Figure 2, the

maximum size of an ISet is 8 while it is 12 for the 50-node networks. It was our conjecture that an optimal configuration would rely heavily on those large independent sets which would mean that scheduling would be a very complex and opaque process. Surprisingly, this was not true, at least for the diverging and converging traffic patterns that we have studied.

For each network scenario shown in Fig. 4, we first compute the optimal throughput curves (a function of the transmit power) without any restrictions on the size of the independent sets. Then we compute the throughput obtained by restricting

Fig. 4. Optimal max-min throughput under constraints on the maximum size of ISets.

the size of the ISets that can be used to be less or equal to 1 (respectively, 2, 3, and 4). Comparing the case $\text{ISet} \leq 2$ (a maximum of 2 links can be scheduled at the same time) with $\text{ISet} = 1$ (corresponding to no spatial reuse), it is obvious that there is a big advantage to allow some level of spatial reuse. However, the gain obtained by allowing more spatial reuse, e.g., $\text{ISet} \leq 3$ is not high as compared to using $\text{ISet} \leq 2$. In fact, in all the four cases, the max-min throughput obtained by limiting the size of the ISets to 2 is never more than 10% below the optimal value. Moreover, $\text{ISet} \leq 4$ yields a throughput that is barely distinguishable from the optimal one. This is rather surprising since it seems to indicate that even moderate spatial reuse is enough to reach excellent throughput. We believe that the reason for that is that our traffic patterns are very much gateway-centric and hence as discussed in [8], the performance of the network can only be improved by trying to "always" schedule one link to or from the gateway. This result is very important since computing the throughput by limiting the ISet size to 2 or 3 makes the computation and possibly the network operation much simpler.

Finally, note that as mentioned in Section IV, the multi-hop advantage is obtained with ISets of size 3 and in some cases, even with ISets of size 2.

VII. IMPACT OF THE INTERFERENCE MODEL

In this section, we investigate the impact of using different interference models on the optimal solution of JRSP. As the additive interference model captures a more realistic interference relation, we use it as the benchmark and compare the interference range and capture threshold models against it. We have presented similar results in [6] for a network with nodes regularly deployed on a grid. Here we consider a scenario where 30 nodes are arbitrarily deployed and we report our results in Fig. 5.

Fig. 5. Impact of the interference model on the max-min throughput: Rand30a with converging traffic.

We first observe that compared with the additive interference model, using the capture threshold model tends to overestimate the optimal throughput and using the interference model tends to underestimate the optimal throughput. The reason for the overestimation is that the over-simplified binary interference relation of the capture model leads to not only more independent sets but also to independent sets of larger size than for the additive interference model. It is important to note that even when both the additive interference and capture threshold models yield the same optimal throughput, the optimal configurations are quite different: most of the independent sets used by the capture threshold model are infeasible under the additive interference model. The underestimation obtained by using the interference range model is surprising since one would expect that dealing with pairwise interference would increase the number of independent sets and hence yield a higher throughput. However, since the interference range model does not take into account the respective strength of the links that it compares, the end-result turns out to be an underestimation of the throughput.

VIII. CONCLUSIONS

We have been able to give interesting insights on the engineering of scheduled wireless mesh networks. In particular, we were able to quantify the advantage of multi-hop over single-hop. We quantify on medium size WMNs, how multi-hop networking achieves the maximum achievable max-min throughput with a transmit power often four times lower than the power needed for single-hop communication as long as the SINR threshold is not too high. This is made possible by allowing spatial reuse, i.e., the activation of more than one link at a time. We also showed that while multi-path routing is not providing much higher max-min throughput than single-path, cross-layer design, i.e., joint routing and scheduling is much better than using any min-hop routing combined to an optimal scheduling. The study on spatial reuse showed us that even moderate spatial reuse is enough to reach excellent throughput. This has a potential operational impact that we will investigate further. Finally, we illustrated the importance of selecting an appropriate interference model.

ACKNOWLEDGMENTS

This work has been supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) and by Research in Motion.

REFERENCES

- [1] M. Alicherry, R. Bhatia, and L. Li. Joint Channel Assignment and Routing for Throughput Optimization in Multi-Radio Wireless Mesh Networks. In *Proc of the 11th ACM MobiCom*, 2005.
- [2] P. Dujkic and S. Valaee. Link Scheduling for Minimum Delay in Spatial Re-use TDMA. In *Proc. of the 26th IEEE INFOCOM*, 2007.
- [3] K. Fall and K. Varadhan, editors. *The ns Manual*. The VINT Project, UC Berkeley, LBL, USC/ISI, and Xerox PARC, 2008. Available from http://nslam.isi.edu/nslam/index.php/User_Information.
- [4] R. Gallager. A Minimum Delay Routing Algorithm using Distributed Computation. *IEEE Trans. on Communications*, 25(1):73–85, 1977.
- [5] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Trans. on Information Theory*, 46(2):388–404, 2000.
- [6] A. Iyer, C. Rosenberg, and A. Karnik. What is the Right Model for Wireless Channel Interference. In *Proc. of the 3rd ICST QShine (journal version submitted to IEEE Trans. on Wireless Communications)*, 2006.
- [7] K. Jain, J. Padhye, V.N. Padmanabhan, and L. Qiu. Impact of Interference on Multi-hop Wireless Network Performance. In *Proc. of the 9th ACM MobiCom*, 2003.
- [8] A. Karnik, A. Iyer, and C. Rosenberg. Throughput-Optimal Configuration of Wireless Networks. *IEEE/ACM Trans. on Networking (to appear)*, 2007.
- [9] J. Luo, C. Rosenberg, and A. Girard. Joint Scheduling, Routing, Power Control and Rate Adaptation in Fixed Wireless Mesh Networks: Algorithms and Engineering. Technical report, University of Waterloo, 2008. Available from <http://www.eng.uwaterloo.ca/~j7luo/documents/JPAIgo.pdf>.
- [10] J. Tang, G. Xue, and W. Zhang. Cross-Layer Optimization for End-to-End Rate Allocation in Multi-Radio Wireless Mesh Networks. *Springer Wireless Networks (to appear)*, 2007.