

A Convex Optimization Approach to Decentralized Rate Control in Wireless Networks with Partial Interference

Lei Wang, David Ripplinger, Anurag Rai, Sean Warnick, and Daniel Zappala

Abstract—Because wireless networks use shared communication channels, contention and interference can significantly degrade throughput and fairness. Optimal rate control algorithms can be designed for wireless networks by modeling the contention between routers using capacity constraints and solving a convex optimization problem. In this work we develop a more accurate network model that directly incorporates partial interference as a receiving constraint, so that it is modeled separately from contention. We show that using this new model leads to a convex optimization problem when formulated using link rates, but it is non-convex when formulated using flow rates. We then use numerical results to illustrate situations when modeling partial interference separately yields significantly higher effective utilities in practical network topologies.

I. INTRODUCTION

Wireless mesh networking is a promising technology that can be used to provide inexpensive communications infrastructure for both developing countries and for urban environments. A typical wireless mesh network consists of stationary routers connected to each other through wireless links. Communication paths, called flows, in a typical mesh network span multiple wireless hops before reaching the wired network or destination.

It is well known that using a shared communications channel in multi-hop wireless networks can significantly degrade throughput and fairness. Because wireless transmissions are broadcast in nature, competing transmissions cause contention and interference, which often results in poor throughput and unfairness between competing flows [1], [2].

Some research seeks to improve the capacity of a mesh network by introducing multiple radios operating on different frequencies, or by manipulating the transmission power of the radios in order to reduce contention [3], [4]. We view this work as complementary to ours. We limit the scope of our problem to designing a transmission rate controller for either communication flows or links, where algorithms manipulating routes, radio frequencies, antennae direction, or transmission power are already employed, and such algorithms are quasi-static with respect to the rate controller.

One approach to solving this problem is to form a model representing the wireless network that imposes constraints on the rates, and then design the controller to compute a distributed algorithm that solves an optimization problem. In wireless networks, these constraints are primarily imposed

by the channel capacity among competing nodes [5]. By formulating the problem so that it is convex, well established techniques can, in most cases, be applied to design a distributed algorithm that computes the rates [6].

The objective function in the optimization problem must yield high throughput while also providing fairness between flows. It is well known that these two goals are conflicting for wireless networks — to provide fairness, some flows may need to sacrifice throughput. A systematic way of achieving a balance between the two is by constructing the objective function as a sum of rate utility functions, which are concave and strictly increasing. The type of fairness is determined by the utility function’s degree of concavity. Several notions of fairness and their corresponding utility functions are well established in the literature [7], [8].

Traditional graph-based models impose resource constraints using a contention graph, where communication links between nodes are represented by vertices and contention is represented by edges. Two links are said to contend with each other if they cannot be active at the same time without causing collisions or interference. Maximal cliques are the largest sets of links where each link in a set contends with all others in that set. Within each maximal clique, only one link can transmit at a time [9], [10], [11].

While this approach does lead to convex problem formulations, it does not accurately represent the interference relationship among links. Contention occurs when sending nodes can sense one another and thus take turns, whereas interference occurs when two *non-contenting* senders transmit simultaneously, causing a packet to be corrupted at a receiving node. A recent measurement study reveals that interference is typical and *partial* [12]. This means that transmissions from an interfering node may corrupt only a fraction of the packets received at a remote node. Modeling interference as contention results in a misleading optimization problem, where capacity may be wasted and actual receiving rates may be far from fair.

Several recent papers have considered partial interference when modeling wireless mesh networks [13], [14], but they are designed for offline prediction of the expected performance of a network, given a vector of input traffic demands. We are concerned with online rate control algorithms, where rates must be adjusted dynamically as flows enter and leave the network.

In this paper, we seek to answer a fundamental question in modeling wireless mesh networks as applied to rate control: is it important to model partial interference separately from contention? We first present the framework for traditional

Supported by AFRL FA8750-09-2-0219 and NSF CNS-0917240.

The authors are with the Computer Science Department, Brigham Young University, Provo, UT 84602, USA. lei@cs.byu.edu, dalerip@gmail.com, anurag1985@gmail.com, sean@cs.byu.edu, zappala@cs.byu.edu




Fig. 1. A network topology graph.

binary interference models, and explain their limitations. Next, we present our new partial interference model, and formulate the optimization problem. We show that this leads to a convex problem when using a link-based formulation and restricting the objective to proportional fairness, but leads to a non-convex problem when using a flow-based formulation. We then use standard techniques to derive a distributed algorithm for the convex problem. Finally, we use numerical results in common network topologies to show that the binary interference models are overly conservative, resulting in sub-optimal rates and significant degradation in network utility. Based on our results, we illustrate the conditions under which the effects of partial interference cannot be ignored and should be modeled accurately.

II. BINARY INTERFERENCE MODEL

Existing graph-based models conservatively characterize interference as a binary effect [10], [11]. Under this model, an interfering node is assumed to corrupt all of the packets received at a remote node, while non-interfering nodes have no effect. Binary interference is represented in a contention graph by simply treating it as contention, that is, if one link interferes with another, neither may send at the same time.

A. Resource Constraints

To illustrate how the binary interference model derives resource constraints, it is useful to consider an example. Fig. 1 shows a sample network topology, denoting active transmission links, transmission ranges, carrier sense ranges, and interference ranges. Nodes within transmission range of each other, such as A and B, can exchange messages. Nodes within carrier sense range, such as D and F, can detect that the channel is busy when one of them is transmitting, but do not necessarily exchange messages. Interference occurs when a node such as D is within interference range of a receiving node such as C. Nodes that are not within carrier sense range of each other can still cause interference, meaning that they cannot simply exchange messages or sense each other to coordinate their use of the channel. In practice, the carrier sense range is typically larger than the transmission range, and the interference range is larger than both. A link is a directional line of communication between two adjacent nodes. A flow is a sequence of links used for communication between users at two distinct nodes.

A contention graph transforms this representation of a network into a new graph that represents the contention and interference constraints [9], [10]. Fig. 2 shows a contention graph for Fig. 1. Vertices in the contention graph correspond




Fig. 2. A contention graph for the sample network topology, based on the binary interference model.

to wireless links, and an edge between two links indicates that the links cannot be active at the same time, due to contention or interference.

Once a contention graph is created, resource constraints can be determined using maximal cliques. Any pair of links in the same maximal clique is prohibited from transmitting concurrently in order to avoid collisions. Thus, for each clique there is a resource constraint, in terms of air time, represented by the clique capacity. Fig. 2 shows three maximal cliques for our sample network, with corresponding resource constraints, where s_l is the sending rate of link l , and c_j is the capacity of clique j . Clique capacities and rates are typically normalized to a value between zero and one. Clique capacities are usually between 0.8 and 0.9, depending on the efficiency of the MAC protocol.

B. Objective Function

When using a binary interference model, the modeler has a choice of objective functions. One could choose to optimize as a function of link utilities, assuming that all links have an infinite backlog of packets to send. A typical objective function in this case is

$$\max f(s) = \sum_{l \in L} U(s_l),$$

where s is the vector of sending rates for links, L is the set of links in the network, and U is the utility function. In this case, the modeler is choosing to maximize the network utility, regardless of how flows map to links. Alternatively, one could choose to optimize as a function of flow utilities, assuming that all flows have an infinite backlog of packets to send. Note that this assumption constrains hops along a flow to only send as much data as they receive. In this case, the objective function becomes

$$\max f(s) = \sum_{t \in T} U(s_t),$$

where s is the vector of source rates for flows, and T is the set of flows in the network. This optimizes based on the rate given to each individual flow.

Flow-based utilities may be more favorable than link-based utilities because they more accurately model realizable rate allocations. If links do not actually have an infinite backlog of data to send, but are limited by how much data earlier links in a flow send, link-based utilities would yield rate allocations for some links that are not realizable. This




Fig. 3. A contention graph for the sample network topology, based on the partial interference model.

can easily lead to situations where capacity is wasted that could have been assigned to other links.

Flow-based utilities may also be preferred because they provide flow-level fairness. Link-based rate control is oblivious to the number of flows using the link. For example, if two links are in the same clique, they may each be assigned a rate of 0.5. However, one link may have 10 flows and another link only 1 flow. In this case, the 10 flows would get a significantly unfair allocation of resources. Despite these disadvantages, allocating based on links may be preferable in some situations, such as when opportunistic routing causes several paths to be used simultaneously [15].

The binary interference model leads to convex problems when modeling fairness based on either links or flows, as long as the utility function is strictly concave [7]. With binary interference, constraints are always linear functions. Thus, standard techniques in convex programming can be used to derive distributed algorithms [6]. However, this model suffers from the problem that finding the set of maximal cliques for a graph is NP-hard. Distributed rate control algorithms often use approximations, for example assigning all links within two hops of each other to the same clique [16]. This is a very conservative heuristic that may overly restrict rates.

III. PARTIAL INTERFERENCE MODEL

To explore the impact of partial interference on optimal rate control for wireless networks, we have developed a partial interference model. In this case, an interfering node may corrupt a fraction of the packets received at a remote node. Partial interference is not represented in the contention graph, but is instead represented in a directional interference map, and is incorporated as an additional constraint or as part of the objective function.

A. Resource Constraints

To model partial interference accurately, we separate contention constraints from interference constraints. Contention is represented as an undirected edge between two vertices (links), and interference is modeled as a directional edge from the interfering link to the receiving link that is affected by the interference. The modified contention graph corresponding to Fig. 1 is shown in Fig. 3. Maximal cliques are then determined as before. Clique constraints, as shown in the figure for Clique 1 and Clique 2, are the same as in the binary interference model, but the constraint between links 2 and 3 is modeled separately.

Interference relationships impose constraints on the receiving rates, as illustrated in the figure, where r_l is the effective receiving rate of link l , s_l is the sending rate of link l , d_l is the inherent loss of the link (e.g. due to obstacles or noise), and the term $(1 - a_{il}s_i)$ is the loss of link l due to interference from an interfering node i . This constraint is taken from a recent measurement study of wireless mesh networks showing that partial interference from one link can be modeled as a linear function [12]. The interference factor a_{il} represents the degree of partial interference inflicted by the interferer. It is in the range $0 \leq a_{il} \leq 1$, and is directional, meaning that a_{il} may be significantly different from a_{li} . Interfering factors can be experimentally measured between any pair of links in a network by methods suggested in [12], [13], constructing an interference map for the network. A study shows that interfering transmissions are independent of each other, and the joint impact of interferers to a receiving node is merely the product of their isolated impacts [12].

The partial interference model is less conservative than other graph-based approaches because more links are assumed to transmit concurrently. For example, consider links 2 and 3 in Fig. 3, and suppose link 3 corrupts 40% of packets received at link 2. If both links transmit at the clique capacity 1, then the sum of their effective receiving rates becomes $1 + (1 - 0.4) = 1.6$. In the binary interference model, these links can not transmit at the same time because they are in the same clique, resulting in a total effective receiving rate of 1. The partial interference model thus allows for significantly higher utilization of the network.

It is important to recognize that even complete interference cannot properly be modeled as contention. That is, a link i will not become a contender to a remote link l even if the interference factor $a_{il} = 1$. Consider again the relationship between links 2 and 3 in Fig. 3. Suppose $a_{32} = 1$. If interference is modeled as contention, then both links will transmit at a rate of 0.5. However, the total effective receiving rate will be $r_2 + r_3 = (0.5)(0.5) + 0.5 = 0.75$. With the partial interference model, it is easy to see that link 2 should send at the full rate regardless of link 3's rate. If link 3 continues to send at a rate of 0.5, then the total effective receiving rate will be $(1)(0.5) + 0.5 = 1$. Thus the partial interference model will result in higher utility.

B. Objective Function

When modeling partial interference, it is more accurate to optimize over receiving rates, because the sending and the receiving rates may be significantly different. Based on a recent study [12], we can model the receiving rate of a link by multiplying the individual interference factors:

$$r_l = d_l s_l \prod_{i \in I(l)} (1 - a_{il} s_i), \quad (1)$$

where $I(l)$ is the set of all links interfering with l . The objective function becomes

$$f(r) = \sum_{l \in L} U(r_l),$$

where r is the vector of link receiving rates and L is the set of all links. The objective function could also be defined in terms of the flows instead of the links.

The multiplicative term that arises in this model may make the optimization problem non-convex, sacrificing well-established techniques in solving convex problems. We will show in the following section that a link-based formulation that incorporates partial interference is still a convex problem if proportional fairness is used. We will also demonstrate that flow-based formulations lose their convexity when considering partial interference.

We are thus faced with a tradeoff between convexity and accuracy when modeling partial interference. A flow-based formulation achieves realizable rate allocations, with a sense of fairness that is more closely correlated to user experience in a network, but at a price of losing convexity in the optimization problem. A link-based formulation, on the other hand, can be convex, but the derived rate allocations may not be realizable for a set of flows, and may give users a sense of unfairness, especially when several flows use the same link or when some flows have more hops than others.

IV. OPTIMIZATION PROBLEM AND SOLUTION

In this section, we first present a link-based problem formulation for optimal rate control in a wireless network, and methods to cast it as a convex problem. We also demonstrate that a flow-based formulation leads to a non-convex problem. At the end of this section, we show a distributed rate control algorithm derived from the link-based convex problem.

A. Link-Based Problem Formulation

Consider the problem of finding an optimal rate allocation that maximizes the sum of link utilities in a wireless mesh network, which consists of a set L of stationary links. We use the following assumptions in our formulation.

- Contention between links is binary (either fully contending or not at all) and symmetric.
- Links have infinite backlog of packets to send.
- The impacts of interference from links are independent and linear with respect to each interferer's sending rate, as described in (1).

For the ease of reading, we provide a list of notations and their descriptions for this problem in Table I.

Given a contention graph with maximal cliques C and an interference map A , the optimization problem maximizes the sum of link utilities, which are functions of link *receiving* rates, in a wireless mesh network:

$$\mathbf{P} : \max_s f(r) = \sum_{l \in L} U(r_l) \quad (2)$$

subject to:

$$s_l \geq 0, \forall l \in L, \quad (3)$$

$$r_l = d_l s_l \prod_{i \in I(l)} (1 - a_{il} s_i), \forall l \in L, \quad (4)$$

$$\sum_{l \in L(j)} s_l \leq c_j, \forall j \in C. \quad (5)$$

L	The set of links in the network.
s_l	The sending rate of link l .
r_l	The receiving rate of link l .
d_l	The delivery ratio of link l .
A	The interference map, or set of interference factors a_{il} in the network, where a_{il} represents link i interfering with link l .
$I(l)$	The set of links that interfere with link l .
$F(l)$	The set of links that link l interferes with.
C	The set of maximal cliques in the contention graph.
$L(j)$	The set of links in maximal clique j .
c_j	The effective capacity of maximal clique j .
$C(l)$	The set of maximal cliques that contain link l .
$U(\cdot)$	The utility function for each link or flow.

TABLE I

NOTATION USED IN THE LINK-BASED FORMULATION

We assume the utility function U of a link is continuously differentiable, strictly concave, monotonically increasing, and approaches negative infinity as the argument approaches zero from the right.

Problem \mathbf{P} is non-convex because of the multiplicative term in constraint 4. However, the problem can be reformulated by substituting (4) into the objective function. Depending on the utility function, the problem may or may not be convex. If we seek to maximize network utility while maintaining proportional fairness, then we let $U(\cdot) = \ln(\cdot)$, and the problem is convex. The objective function becomes

$$f(s) = \sum_{l \in L} \left(\ln s_l + \ln d_l + \sum_{i \in I(l)} \ln (1 - a_{il} s_i) \right). \quad (6)$$

Note that the terms can be reordered and that maximizing (6) gives the same optimal rates whether or not the delivery ratios d_l are considered, so that the objective function may be reformulated as

$$f'(s) = \sum_{l \in L} \left(\ln s_l + \sum_{i \in F(l)} \ln (1 - a_{li} s_l) \right). \quad (7)$$

Thus, we can reformulate problem \mathbf{P} as a convex problem \mathbf{P}'

$$\mathbf{P}' : \max f'(s) \quad (8)$$

subject to:

$$s_l \geq 0, \forall l \in L, \quad (9)$$

$$\sum_{l \in L(j)} s_l \leq c_j, \forall j \in C. \quad (10)$$

B. Non-convexity in the Flow-Based Formulation

The flow-based formulation differs from the link-based formulation in that it maximizes end receiving rates of multi-hop transport-layer flows, where only the source link of each flow has an infinite backlog. For the ease of reading, we provide a list of notations and their descriptions for this problem in Table II.

The optimization problem for this formulation is

$$\mathbf{Q} : \max_s f(r) = \sum_{t \in T} U(r_{end}^t) \quad (11)$$

T	The set of transport-layer flows in the network.
$T(l)$	The set of flows traversing link l .
r_{end}^t	The end receiving rate of flow t .
s_k^t	The sending rate of flow t at hop k .
r_k^t	The receiving rate of flow t at hop k .
$k(l, t)$	Hop k of flow t , corresponding with link l .
$h(t)$	The length of flow t in hops.

TABLE II
NOTATION USED IN THE FLOW-BASED FORMULATION

subject to:

$$s_k^t \geq 0, \quad \forall t \in T, \quad k = 1, \dots, h(t), \quad (12)$$

$$s_k^t = r_{k-1}^t, \quad \forall t \in T, \quad k = 2, \dots, h(t), \quad (13)$$

$$\sum_{l \in L(j)} \sum_{t \in T(l)} s_{k(t,l)}^t \leq c_j, \quad \forall j \in C, \quad (14)$$

where each r_k^t is a function of sending rates, according to (4).

Constraint (13) makes problem **Q** non-convex. This constraint arises because in this formulation there is no longer an assumption that all links have an infinite backlog of packets, as interference at an earlier hop causes the next hop to have less data available to send. In order to achieve feasible rate allocations for flows, we require that a flow is transmitted at the same rate as it is received at each hop, otherwise congestion or starvation would occur along its path.

Because this problem is non-convex, a distributed solution may not be as straightforward or efficient. We note, however, that problem **Q** could be approximated by relaxing the condition in (13) and including all hops in the objective function, with carefully assigned weights on each utility, thus casting the problem in a form similar to **P'**. This could be a topic of further research, but we do not pursue it here.

C. Distributed Algorithm

We derive a distributed algorithm to solve problem **P'**, based on the methods presented in [6]. This problem meets Slater's condition [17], giving us strong duality. We seek to solve the problem in a distributed fashion by finding the solution to the dual using Lagrangian relaxation. The Lagrangian of problem **P'** is

$$\begin{aligned} L(s, \lambda) &= f'(s) + \sum_{j \in C} \lambda_j \left(c_j - \sum_{l \in L(j)} s_l \right) \\ &= f'(s) - \sum_{j \in C} \sum_{l \in L(j)} \lambda_j s_l + \sum_{j \in C} c_j \lambda_j \\ &= f'(s) - \sum_{l \in L} s_l \sum_{j \in C(l)} \lambda_j + \sum_{j \in C} c_j \lambda_j \\ &= \sum_{l \in L} g(s_l, \lambda) + \sum_{j \in C} c_j \lambda_j, \end{aligned}$$

where λ_j are Lagrange multipliers for constraints (5) and

$$g(s_l, \lambda) = \ln s_l + \sum_{i \in F(l)} \ln(1 - a_{li} s_l) - s_l \sum_{j \in C(l)} \lambda_j. \quad (15)$$

Note that $g(s_l, \lambda)$ is concave in s_l and approaches $-\infty$ to the left and right, so that for a given λ there is always a unique maximizer

$$\bar{s}_l(\lambda) = \arg \max_{s_l} g(s_l, \lambda). \quad (16)$$

This can easily be found by taking the derivative of g with respect to s_l and setting it equal to zero:

$$\frac{1}{\bar{s}_l} - \sum_{i \in F(l)} \frac{a_{li}}{(1 - a_{li} \bar{s}_l)} = \sum_{j \in C(l)} \lambda_j. \quad (17)$$

We can use efficient algorithms, such as Newton's method, to solve for the optimal rates according to (17). Link l does not need any information other than the λ s for each maximal clique to which l belongs in order to evaluate (17).

The dual function to problem **P'** is given by

$$\begin{aligned} Z(\lambda) &= \max_s L(s, \lambda) \\ &= \max_s \sum_{l \in L} g(s_l, \lambda) + \sum_{j \in C} c_j \lambda_j \\ &= \sum_{l \in L} \max_{s_l} g(s_l, \lambda) + \sum_{j \in C} c_j \lambda_j \\ &= \sum_{l \in L} g(\bar{s}_l(\lambda), \lambda) + \sum_{j \in C} c_j \lambda_j, \end{aligned}$$

and the dual problem is

$$\mathbf{D} : \min Z(\lambda) \quad (18)$$

subject to:

$$\lambda \geq 0. \quad (19)$$

We use the gradient projection method to iteratively obtain the optimal λ for the problem. From Danskin's theorem [18], we know that

$$\begin{aligned} \frac{\partial Z}{\partial \lambda_j} &= \frac{\partial}{\partial \lambda_j} \left[f'(s) + \sum_{i \in C} \lambda_i \left(c_i - \sum_{l \in L(i)} s_l \right) \right]_{s=\bar{s}} \\ &= c_j - \sum_{l \in L(j)} \bar{s}_l. \end{aligned} \quad (20)$$

Using a step size γ in the negative direction of the gradient gives the algorithm

$$\lambda_j(k+1) = \max \left(0, \lambda_j(k) - \gamma \left(c_j - \sum_{l \in L(j)} \bar{s}_l(k) \right) \right), \quad (21)$$

where

$$\bar{s}_l(k) = \bar{s}_l(\lambda(k)).$$

Each link in a maximal clique can share with each other their current values of \bar{s}_l , which is only local information, and then compute the next λ_j .

The convergence of the algorithm is well established in the literature, even when it is asynchronous. Once λ converges to the optimal solution, λ^* , of the dual problem, the optimal solution, s^* , to the primal problem is given by

$$s^* = \bar{s}(\lambda^*).$$

V. NUMERICAL RESULTS

We seek to determine in what situations the partial interference model outperforms binary interference models, and by how much. We use MATLAB to numerically compute solutions to the rate optimization problem for several different wireless networks. We use network topologies that represent basic situations — these can be thought of as building blocks out of which larger topologies can be formed.

We introduce three binary interference models that we compare with the partial interference (PI) model. The interference-as-contention (IC) model replaces any interference mappings with contention, no matter how small the interference factor a . The interference-ignored (II) model simply ignores any interference mappings and models only contention. The adaptive contention (AC) model follows the IC model or the II model, depending on which model has higher performance. Thus the AC model gives binary contention the benefit of the doubt — it ignores interference when this provides good performance and models it as contention otherwise.

A. Performance metric

To compare these different models, we define a performance metric that is based on the objective function of the PI model, using receiving rates. We justify this by recognizing that receiving rates are what ultimately matters for users of the network. Data that is sent but is lost due to interference is not considered useful. Thus the comparison should be made between the performance observed with PI-derived receiving rates r^* and the receiving rates r' actually obtained by the other model from its sending rates s' , according to the PI constraint on receiving rates.

For ease of interpretation, we consider the ratio R of performances P , that is,

$$R = P(r^*)/P(r'). \quad (22)$$

Thus, the comparison will simply read that the PI model outperforms the other model R times.

The PI model uses proportional fairness, so that its objective function is

$$f(r) = \sum_{l \in L} \ln r_l. \quad (23)$$

However, scores obtained from $f(r)$ range from $-\infty$ to zero, making it non-intuitive to ascertain how significant a better score might be in comparison to a worse score. We introduce the performance function

$$P(r) = e^{f(r)/|L|}, \quad (24)$$

where $|L|$ is the number of links in the network. Note that if $f(r^*) > f(r')$, then clearly $P(r^*) > P(r')$, maintaining the ordering of feasible rate vectors r , based on the objective function scores. Furthermore, note that P turns out to be the geometric mean of receiving rates, which ranges between zero and one, and is normalized with respect to the size of the network. Therefore, we study the ratio R of performances, as denoted in (24), between the PI model and other models for various network topologies.

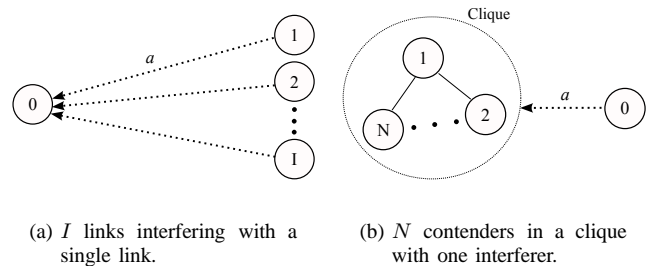


Fig. 4. Topologies used for numerical results.


B. Results

We consider three generic network topologies and plot R for each topology and for each contention model being compared with the PI model. Each topology is represented in the figures as a combined contention graph and interference map, according to the PI model. Clique capacities in each topology are all $c = 0.85$.


In all cases, the IC model never does as well as the PI model because modeling interference as contention is too conservative. For low values of interference, it is better to let links send at faster rates and suffer some packet loss. At high values of interference, it is better to have the interfered link send at a faster rate than the interferer, to provide better performance and fairness. However, modeling interference as contention is often better than ignoring it when interference is high. Thus in most cases, the combined AC model follows the II model for low values of interference and follows the IC model for high values of interference.

Fig. 4(a) shows the first topology, where I links interfere with a single link with a common interference factor a , but do not interfere with each other. Fig. 5(a), 5(b), and 5(c) plot R for this topology for the PI model against the IC, II, and AC models, respectively. The dotted curves show where R begins to be greater than one. Interestingly, the PI model and the II model perform exactly the same for values of a below 0.59. This is because, for low values of a , the cost of interference is offset by the gain of the interferer sending at full capacity. Thus, both the PI model and the II model calculate sending rates at full capacity for each link. For larger values of a and I , the PI model outperforms the binary interference models more than 1.5 times.


Fig. 4(b) shows the second topology, where a single link has interference factor a on N links that contend in a single clique. Fig. 6(a), 6(b), and 6(c) plot R for this topology for the PI model against the IC, II, and AC models, respectively. The dotted curves show where R begins to be greater than one. The PI model starts performing better than the II model at much lower values of a when N is large. This is due to the fact that the contending links already have small rates as a consequence of sharing the medium. Utilities are lowered much more by interference when sending rates are small. Thus, even for low values of a , the PI model does not calculate sending rates at full capacity. However, for higher values of a and N , the PI model outperforms the IC model



(a) Ratio R of performance between the PI model and the IC model.




(b) Ratio R of performance between the PI model and the II model. The dotted curve marks where R begins to be greater than one.




(c) Ratio R of performance between the PI model and the AC model. The dotted curve marks where R begins to be greater than one.


Fig. 5. Numerical results for I interferers on one link.



(a) Ratio R of performance between the PI model and the IC model.



(b) Ratio R of performance between the PI model and the II model. The dotted curve marks where R begins to be greater than one.



(c) Ratio R of performance between the PI model and the AC model. The dotted curve marks where R begins to be greater than one.

Fig. 6. Numerical results for N contenders with one interferer.

only about 1.1 times.

To demonstrate the worth of the PI model, we consider a topology combining features of the first two, where I links have a fixed interference factor $a = 0.4$ on N links that contend in a single clique. Fig. 7(a), 7(b), and 7(c) plot R for this topology for the PI model against the IC, II, and AC models, respectively. Experimental results show that it is typical for interference factors to range anywhere between zero and one in a real network, with usually at least one interferer on a link having a factor of at least $a = 0.8$, so choosing $a = 0.4$ in this topology is a somewhat conservative comparison [12]. The combined effect of several interferers and several contenders causes the PI model to perform significantly better than the binary interference models.

VI. LIMITATIONS


We model partial interference according to a well-known measurement study [12], and this model might generate inaccurate interference degradation in some scenarios. For example, if links 1 and 2 share the same sending node, and both have interference factor a on link 0, the actual receiving

rate of link 0 is $r_0 = s_0(1 - a(s_1 + s_2))$, since links 1 and 2 cannot send at the same time. Our model predicts $r_0 = s_0(1 - as_1)(1 - as_2)$, so some accuracy is sacrificed in that regard.


Our model is also limited by our restriction to objective functions that provide proportional fairness. We make this restriction in order to guarantee convexity in the optimization problem.

VII. CONCLUSIONS


We have presented a partial interference model for optimal rate control in multi-hop wireless networks. This new model directly incorporates partial interference as a receiving constraint, so that it is modeled separately from contention. This leads to a convex rate control problem when formulated with link rates, but a non-convex problem when formulated with flow rates. Using standard techniques, we derive a distributed rate control algorithm for the link formulation. Numerical results on typical network topologies show that it is never correct to model partial interference as contention. In some cases, when partial interference is low, it can be



(a) Ratio R of performance between the PI model and the IC model.



(b) Ratio R of performance between the PI model and the II model.



(c) Ratio R of performance between the PI model and the AC model.

Fig. 7. Numerical results for I interferers and N contenders, with $a = 0.4$.

safely ignored. However, as the number of contenders and the number of interferers grows, it becomes increasingly important to model partial interference accurately, even for small levels of interference.

The partial interference model can be used as the basis for a number of interesting wireless network problems. Joint routing and rate control, link-layer scheduling, channel assignment for multiple radio networks, and transmission power control problems can all benefit from an accurate model of relationships among links. In addition, the rate control problem we have presented here needs additional work. We plan to design an efficient implementation and to verify numeric performance gains on a wireless mesh testbed.

REFERENCES

- [1] J. Shi, O. Gurewitz, V. Mancuso, J. Camp, and E. Knightly, "Measurement and modeling of the origins of starvation in congestion controlled mesh networks," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, April 2008, pp. 1633–1641.
- [2] S. Rangwala, A. Jindal, K.-Y. Jang, K. Psounis, and R. Govindan, "Understanding congestion control in multi-hop wireless mesh networks," in *MobiCom '08: Proceedings of the 14th ACM international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2008, pp. 291–302.
- [3] P. Gupta and P. Kumar, "The capacity of wireless networks," *Information Theory, IEEE Transactions on*, vol. 46, no. 2, pp. 388–404, mar 2000.
- [4] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," in *Broadband Networks, 2004. BroadNets 2004. Proceedings. First International Conference on*, Oct. 2004, pp. 344–354.
- [5] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255–312, 2007.
- [6] S. Low and D. Lapsley, "Optimization flow control. i. basic algorithm and convergence," *Networking, IEEE/ACM Transactions on*, vol. 7, no. 6, pp. 861–874, Dec. 1999.
- [7] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *Networking, IEEE/ACM Transactions on*, vol. 8, no. 5, pp. 556–567, Oct 2000.
- [8] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An axiomatic theory of fairness in network resource allocation," in *IEEE INFOCOM*, 2010.
- [9] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan, "Achieving mac layer fairness in wireless packet networks," in *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2000, pp. 87–98.
- [10] L. Chen, S. Low, and J. Doyle, "Joint congestion control and media access control design for ad hoc wireless networks," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 3, March 2005, pp. 2212–2222 vol. 3.
- [11] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," *Wirel. Netw.*, vol. 11, no. 4, pp. 471–487, 2005.
- [12] D. Niculescu, "Interference map for 802.11 networks," in *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2007, pp. 339–350.
- [13] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan, "A general model of wireless interference," in *MobiCom '07: Proceedings of the 13th annual ACM international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2007, pp. 171–182.
- [14] Y. Li, L. Qiu, Y. Zhang, R. Mahajan, and E. Rozner, "Predictable performance optimization for wireless networks," in *SIGCOMM '08: Proceedings of the ACM SIGCOMM 2008 conference on Data communication*. New York, NY, USA: ACM, 2008, pp. 413–426.
- [15] S. Chachulski, M. Jennings, S. Katti, and D. Katabi, "Trading structure for randomness in wireless opportunistic routing," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 169–180, 2007.
- [16] B. Bensaou and Z. Fang, "A Fair MAC Protocol for IEEE 802.11-Based Ad Hoc Networks: Design and Implementation," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 8, pp. 2934–2941, August 2007.
- [17] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [18] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation*. Prentice-Hall, 1989.