Throughput-optimal Configuration of Wireless Networks

Aditya Karnik, Aravind Iyer and Catherine Rosenberg

***Abstract*— In this paper, we address the following two questions: (i) given a set of nodes with arbitrary locations, and a set of data flows, what is the max-min achievable throughput? and (ii) how should the network be configured to achieve the optimum? We consider these questions from a networking standpoint, and pose them as an optimization problem to determine jointly optimal flow routes, link activation schedules and physical layer parameters. We obtain some interesting insights into the interplay of the achievable throughput, routing and transmit power and modulation schemes. Determining the achievable throughput is computationally hard in general, however, using a smart technique we obtain numerical results for different scenarios of interest. We believe that our optimization based framework can also be used as a tool, for planning and capacity provisioning of wireless networks.**

I. Introduction

Characterizing the "capacity" of a wireless network has turned out to be a difficult problem owing to the intricacies of communication over a wireless medium. Beginning with [1], a popular approach has been to characterize the asymptotic scaling of capacity in the number of nodes (e.g., [2]). By asking for bounds only in an order sense, it has been possible to derive the trend of capacity scaling, even in the information theoretic sense [3]. However, although the knowledge of a capacity scaling law is quite valuable, it lends no insights into actual numbers for network capacity based on current technologies, or into the impact of macroscopic parameters such as transmit power budget and the availability of modulation schemes, on the network capacity. There has also been some work on optimally operating a wireless network; for example, in [4] a routing and power control policy to stabilize the network is devised and in [6] a power transmission strategy to minimize the power cost subject to average rate constraint is found. Owing to their focus on the algorithmic perspective, these works also do not shed any light into the network capacity, and an optimal configuration of a given wireless network to achieve it. This is the problem we seek to address in this paper.

Specifically, we seek a "capacity" (i.e., maximum throughput) result as well as a particular network "configuration" achieving it. The interest in such a configuration is not only theoretical but practical as well- we believe that it may be used as a design or guideline while deploying infrastructure-based wireless network, for example, broadband mesh networks. Therefore, we consider this problem from a networking standpoint rather than an information theoretic one, and establish explicit, rather than asymptotic, bounds for specific instances of node locations and flows, under currently implementable physical layer technologies. Thus the first question we address is the following.

Q1. Given a set of nodes with arbitrary locations, and a set of data flows specified as source-destination pairs, what is the *maximum achievable throughput*, under certain constraints on the radio parameters (in particular, on transmit power)?

In this work, our notion of the maximum throughput is the *max-min flow rate*, *i.e.*, the maximum of the minimum end-to-end flow throughputs that can be achieved in the network. This is an appropriate notion of capacity from a networking perspective since it may represent the aggregate bandwidth demands of subscriber stations in an IEEE 802.16-like access network, or the sampling rate at which sensors produce information about their environment, in a sensor network. We assume point-to-point links which can operate in accordance with their perceived signal-to-interference-and-noise ratio (SINR). We also assume that transmissions are co-ordinated (possibly, though not necessarily, by a central controller) so as to be conflict-free through activation schedules. This is not only to establish "capacity" but also because broadband wireless standards such as IEEE 802.16 provide for such mechanisms. Thus, our bounds reflect the intrinsic limits posed by the underlying conflict structure of parallel transmissions, as dictated by the SINR calculations. The second question we answer is the following.

Q2. How should the network be configured to achieve this maximum? Specifically, by *configuration*, we mean the complete choice of the set of links (*i.e.*, topology), the routes, link schedules, and transmit powers and modulation schemes for each link.

We assume that the wireless nodes have an arbitrary set of transmit power levels, and an arbitrary set of modulation and coding strategies. All wireless links are considered to be operated at some required bit-error-rate (BER). A transmission on a wireless link is said to succeed, if the SINR perceived by the link remains greater than an SINR threshold which depends on the transmit power, modulation and coding, and the requisite BER. We show that this requirement results into a *conflict set* for each wireless link. For a given link, the corresponding conflict set is a collection of subsets of links, such that at least one link from each subset cannot be activated *simultaneously* with the given link (in the sense of satisfying the BER requirement).

To resolve Q1 and Q2 we take a constructive approach, i.e., we explicitly construct a network that has the maximum throughput. Since the problem *is to construct a throughput-optimal network*, the idea is to pose it as a problem of optimal resource allocation and routing on a "dummy network" specified by the complete graph on given wireless nodes.

We discuss two static optimization formulations for this problem. They are complementary in the sense that one throws light on the routing perspective while the other on the scheduling perspective. It can be shown that determining the maximum throughput is a computationally hard problem in general. However, using a smart enumerative technique, for specific instances of the problem computations can be greatly reduced. We provide results for different scenarios of interest such as a sensor network on a grid, and a mesh of nodes deployed in hexagonal cells.

A static optimization setting similar to ours has been discussed in [5], [6]. In [6], the authors focus on a low SINR regime with the link capacities being linear functions of the perceived SINR. In our work, we consider modulation and coding schemes as determining the link capacities, rather than approximations of the Shannon capacity function. In [5], the authors use a conflict graph formulation to model the schedulability relationships between the wireless links. Our results are based on a more general conflict set formulation, and while they certainly support those in [5], they further provide important insights into the impact of macroscopic parameters such as transmit power budget and availability of modulation schemes, on the throughput performance. Further, the upper bounds on throughput derived in [5] are rather loose and consider only the cliques formed by singleton conflict set members.

The paper is organized as follows. The problem formulation and the main results are discussed in Section II, followed by a few case-studies in Section III. Finally, we conclude in Section IV.

II. Analytical Formulations

We are interested in answering Q1 and Q2 posed in Section I. Seen together they are equivalent to constructing a throughput-optimal network given the arbitrary sets of nodes, flows among them, and available radio parameters. The idea is to pose this problem as one of optimal resource allocation and routing on a "dummy network" specified by the complete graph on the given wireless nodes. An optimal solution of this problem completely characterizes a throughput-optimal configuration of these nodes; the set of links with positive transmission power allocated to them represents the network topology whereas the flow routes, link schedule and radio parameters at each link specify an optimal network configuration. Such a resource allocation problem is, however, intrinsically complicated owing to the interdependence of routing, scheduling and radio parameters.

In this paper, we address a special case of this general problem under the following assumptions. The channel conditions are time-invariant and known (this is not a restrictive assumption in view of relatively static fixed wireless channels [8]). The choice of transmission power and modulation-coding schemes is to be made from given finite sets. Finally, transmissions are co-ordinated through the link activation schedules. Under these assumptions, the problem of resource allocation and routing can be cast simply as an optimal routing problem. The idea is to replace a link say, between nodes i and j, by multiple "artificial links" each one corresponding to one combination of transmit power levels and modulation-coding schemes available at i. Thus, optimal selection of power and modulation is translated into optimal selection of "links".

Some basic notation is in order. N denotes the number of given static wireless nodes; their set is denoted by $\mathcal{N}$. $\mathcal{L}$ denotes the set of all possible links among these nodes; $\mathcal{L}$ includes all artificial links as discussed above. Cardinality of $\mathcal{L}$ is denoted by L. Links are assumed to be directed; $l \in \mathcal{L}$ is also represented as (l_o, l_d), where l_o and l_d denote the originating and the destination nodes resp. $\mathcal{L}_i^O$ (resp. $\mathcal{L}_i^I$) denotes the set of links outgoing from node i (resp. incoming to i). Let P_l denote the transmission power on link l and z_l the corresponding modulation-coding scheme.

Link Conflict Structure: In this non-information-theoretic setup, communication errors cannot be completely eliminated; hence the "success" is in the sense of achieving a specified bit (or packet) error rate. The SINR on link l, γ_l, at a given time instant is given by $\frac{G_{ll}P_l}{N_0+\sum_{l'} G_{l'l}P_{l'}}$ where the summation in the denominator is over the links transmitting simultaneously with l. G_{ll} (resp. $G_{l'l}$) denotes the channel gain on link l (resp. from link l' to l where it is understood that it refers to the gain from l'_o to l_d). As assumed earlier, $G_{\{.\}}$s are known. A bit error rate (BER) specification translates into an SINR threshold, $\beta(z_l)$ for each z_l on link l; for narrow-band systems $\beta(z_l) > 1$. This essentially means that a transmission on link l is considered to be successful if γ_l is at least $\beta(z_l)$ for the duration of transmission. Let c_l denote the bit rate on link l. If z_l is feasible, i.e., if it achieves the BER specification in the absence of co-channel interference on link l, then c_l equals the bit rate provided by z_l. Otherwise, c_l is 0.

The condition for a successful transmission, $\gamma_l \geq \beta(z_l)$, on link l defines what we call its "conflict set", $\mathcal{D}_l$. Each $D \in \mathcal{D}_l$ is a subset of $\mathcal{L}$ with the interpretation that if *all* the links in D are transmitting simultaneously with link l then the transmission on link l fails. This can be seen as follows. Let $\mathbf{v}$ be an L-dimensional $\{0, 1\}$ vector and let $\mathcal{V}_l = \left\{\mathbf{v} : \frac{G_{ll}P_l}{N_0+\sum_{l'} G_{l'l}P_{l'}v_{l'}} < \beta(z_l)\right\}$. Then each $D \in \mathcal{D}_l$ corresponds to a $\mathbf{v} \in \mathcal{V}_l$ that cannot be represented as (modulo 1) addition of any other vectors in $\mathcal{V}_l$; $l' \in D$ if $v_{l'} = 1$; thus D's are minimal subsets of interfering links. Therefore, to guarantee a successful transmission on link l, at least one link from each $D \in \mathcal{D}_l$ must be silent. Thus, under a realistic physical layer model, conflicts among links may be more complicated than those representable by k-hop neighbourhoods for each link or by the two-circle model. Moreover, the conflict relation may not even be "binary" to be represented by a conflict graph. Interestingly, however, the conflict sets can be seen as specifying *multiple conflict graphs*; in each of these graphs for each link l one interferer is selected from each $D \in \mathcal{D}_l$ and an edge established (the requirement that "at least one link from D must be silent" is, thus, satisfied).

Link Scheduling: Let $\mathcal{S}$ denote the power set of $\mathcal{L}$;

cardinality of $\mathcal{S}$ is denoted by S. Sets in $\mathcal{S}$ will be arbitrarily ordered and indexed $1, 2, \ldots, S$; a generic set will be denoted by A. A link activation schedule is an S-dimensional vector $\alpha := (\alpha_1, \alpha_2, \ldots, \alpha_S)$ such that $\alpha_i \geq 0,\ i = 1, \ldots, S$ and $\sum_{i=1}^{S} \alpha_i = 1$; α can be interpreted either as a probability vector denoting probability of activating the set of links in A_i for the duration of a slot or as a time allocation map denoting the fraction of time links in A_i transmit simultaneously in a frame. $A \in \mathcal{S}$ will be called an "independent set" for the conflict structure imposed by $\mathcal{D}_l$, if for every $l \in A$ at least one link from each $D \in \mathcal{D}_l$ is not in A[1]. Denote by $\mathcal{I}$ the set of independent sets in $\mathcal{S}$ and by $\mathcal{A}$ the set of non-conflicting schedules, i.e., $\{\alpha | \alpha_i \geq 0, i \in \mathcal{I}, \alpha_i = 0, i \notin \mathcal{I}, \sum_{i=1}^{S} \alpha_i = 1\}$. It is known from [7] that link l has capacity C_l if and only if $C_l = c_l \sum_{i \in \mathcal{I}_l} \alpha_i$ for some $\alpha \in \mathcal{A}$. $\mathcal{I}_l$ denotes the set of independent sets which contain l. To make dependence on α explicit we denote the capacity of link l by $c_l(\alpha)$.

Flows and Routing: Data transfer requirements are specified in terms of M flows; the set of flows is denoted by $\mathcal{M}$. Each flow $f \in \mathcal{M}$ is identified with a source-destination pair $(f_s, f_d), f_s, f_d \in \mathcal{N}$. r generically denotes a route and specifies a sequence of links from f_s to f_d. $\mathcal{R}_f$ is the set of possible routes f can be routed on. $\mathcal{R}_f^l$ denotes the set of routes of f going over link l. λ_f denotes the flow rate of f. $\phi_f^r \geq 0$ denotes the fraction of traffic of flow f routed on $r \in \mathcal{R}_f$. Clearly $\sum_{r \in R_f} \phi_f^r = 1$.

Remark: It must be noted that the conflict structure and hence non-conflicting schedules are implicitly dependent on the underlying placement of the nodes. □

The following is the formal max-min throughput optimization problem.

$$\begin{array}{ll} \max \lambda & (1) \\ \sum_{f \in F} \lambda_f \left(\sum_{r \in \mathcal{R}_f^l} \phi_f^r \right) \leq c_l(\alpha) & l = 1, \ldots, L \\ \sum_{r \in \mathcal{R}_f} \phi_f^r = 1, \phi_f^r \geq 0 & f = 1, \ldots, M \\ \lambda \leq \lambda_f & f = 1, \ldots, M \\ \lambda \geq 0, \alpha \in \mathcal{A} & \end{array}$$

The following is straightforward to show; detailed proofs have been provided in the extended version [9]. An optimal solution exists for (1). Denote it by λ^*. λ^* can be obtained as a solution to the parameterized version of (1); parameterization is with respect to the routing variables $\phi := (\phi_f^r)_{f \in F, r \in R_f}$. Let $\Phi = \{\phi | \sum_{r \in \mathcal{R}_f} \phi_f^r = 1, \phi_f^r \geq 0\, f = 1, \ldots, M\}$.

To simplify the exposition, first let us assume that the conflict structure is specified by one conflict graph G. Let vertex l in graph G be assigned a weight equal to $\frac{\sum_{f \in F} \sum_{r \in \mathcal{R}_f^l} \phi_f^r}{c_l}$ for given $\phi \in \Phi$. Let $\mathcal{C}$ be the set of cliques in G with $w_c(\phi)$ denoting the weight of clique $c \in \mathcal{C}$; by weight of a clique we mean the sum of the weights of vertices in that clique. Thus, $w_c(\phi) = \sum_{l \in c} \frac{\sum_{f \in F} \sum_{r \in \mathcal{R}_f^l} \phi_f^r}{c_l}$.

[1]Recalling the fact that $\mathcal{D}_l$'s realize multiple conflict graphs, an independent set A is a graph-theoretic independent set in one of those graphs.

Proposition 2.1: For some $\kappa \in (0, 1]$ which depends on the conflict graph G,

$$\frac{\kappa}{\min_\phi \max_c w_c(\phi)} \leq \lambda^* \leq \frac{1}{\min_\phi \max_c w_c(\phi)} \qquad (2)$$

Proof: See Section V. □

Remark: If G is perfect, $\kappa = 1$, hence, $\lambda^*(\phi) = \frac{1}{\max_c w_c(\phi)}$. □

Let $\hat{\phi} \in \arg\min_{\phi \in \Phi} \max_{c \in C} w_c(\phi)$. We shall refer to $\hat{\phi}$ as "min-max routing". Then Proposition 2.1 can equivalently be stated as follows.

Proposition 2.2: For some $\kappa \in (0, 1]$ which depends on the conflict graph G,

$$\kappa \lambda^* \leq \lambda^*(\hat{\phi}) \leq \lambda^*$$

Proof: See Section V. □

In general, the conflict structure specified in terms of the conflict sets, is not necessarily representable by a single conflict graph. So we take the following approach to arrive at Proposition 2.1 in a general setting. The idea is to "embed" multiple conflict graphs specified by the conflict sets in a larger conflict graph. This is done by considering multiple copies of each link, each copy basically realizing one combination of activation constraints given by its conflict set. Recall that multiple combinations of these constraints give rise to multiple conflict graphs. For example, for link m let $\mathcal{D}_m = \{\{l_1\}, \{l_2, l_4\}, \{l_3, l_5\}\}$. Then the conflict graph is constructed by replacing link m by a clique of size 4, with copies of link m as vertices, say, m^1, m^2, m^3 and m^4, with edges to $\{l_1, l_2, l_3\}$, $\{l_1, l_4, l_3\}$, $\{l_1, l_2, l_5\}$, $\{l_1, l_4, l_5\}$ respectively. In general, if $\mathcal{D}_m = \{D_1, D_2, \ldots, D_k\}$, then link m would be replaced by a clique of size $|D_1| \times |D_2| \times \ldots \times |D_k|$. Since the links $l_1, \ldots, l_5$ themselves may have similar copies, by an edge between say m^1 and l_5, we mean edges from m^1 to all the "virtual" copies of l_5. Again using the idea of "artificial" links, in the "extended" network link m is now replaced by its virtual copies $m^1, \ldots, m^4$. By appropriately redefining the routing variables ϕ, the optimization problem over this extended network and the corresponding conflict graph has the same form as (1), and thus, Proposition 2.1 also holds in a general setting, but on an extended conflict graph.

Now, in the case of perfect (extended) conflict graphs, the optimal routing is "min-max" whereas in general, due to Proposition 2.2, the min-max routing, $\hat{\phi}$, is guaranteed to result in a throughput which is within a constant factor of the optimal (in the sense of (1)). The routes in $\hat{\phi}$ can be shown to *minimize the total cost of air-time in the maximum weighted cliques* they pass through. The air-time of link l is the time to send one unit of data, i.e., $\frac{1}{c_l}$ and the air-time of a clique in G is simply the sum of air-times of its constituent links. When $c_l = 1$ (normalized) for all links, the routes in $\hat{\phi}$ employ the shortest path through the maximum weighted cliques. This is a generalization of an obvious result that if G is complete, i.e., if only one link, among links of equal data rates, can transmit at a time, the minimum hop routing is optimal. We believe that min-max routing can provide interesting insights into the structure of optimal routing and possibly into design of routing algorithm. We do not explore these issues in this paper.

(1) can be formulated in an alternative way, emphasizing the scheduling aspect. Let x_l^f denote the flow variable associated with flow $f \in \mathcal{M}$ on link l. Then throughput optimization can be cast as the following linear program.

$$\max \lambda \tag{3}$$

$$\sum_{l \in \mathcal{L}_i^o} x_l^f - \sum_{l \in \mathcal{L}_i^I} x_l^f = \begin{cases} 0 & i \notin \{f_s, f_d\} \\ \lambda_f & i = f_s \\ -\lambda_f & i = f_d \end{cases}$$
$$i = 1, \ldots, N$$
$$\sum_{f \in F} x_l^f \leq c_l \sum_{k \in I_l} \alpha_k \quad l = 1, \ldots, L$$
$$\sum_{k \in \mathcal{I}} \alpha_k = 1$$
$$0 \leq \lambda \leq \lambda_f \qquad f = 1, \ldots, M$$

The equivalence of (3) and (1) is direct. Therefore, we denote an optimal solution of (3) by λ^* as well. An interesting characterization of λ^* can be obtained from the dual of (3) as follows.

$$\lambda^* = \min_{\nu \in V} \max_{1 \leq i \leq \mathcal{I}} \sum_{l \in i} \nu_l c_l \tag{4}$$

where the set of dual variables $V = \{(\mu, \nu, u) \in R^{NM} \times R_+^L \times R_+^M | \sum_{f \in \mathcal{F}} u_f = 1, u_f \leq \mu_{f_s,f} - \mu_{f_d,f},\ f \in \mathcal{F};$ and $\mu_{l_o,f} - \mu_{l_d,f} \leq \nu_l,\ f \in \mathcal{F}, l \in \mathcal{L}\}$.

It can be shown that for $i, j \in \mathcal{I}$ with $\alpha_i^* > 0$ and $\alpha_j^* > 0$, $\sum_{l \in i} \nu_l^* c_l = \sum_{l \in j} \nu_l^* c_l$. Thus, if ν_l is interpreted as the cost of using link l, then the cost of using the total data-rate of every actively used independent set is equalized.

Remark: Note the interesting complementary characterization of λ^* in terms of flow routes & cliques in G (Equation 2), and link schedules & independent sets in G (Equation 4). These two results also imply that computation of λ^* is a hard problem. We do not formally prove this fact in this paper. □

III. Numerical Results

In this section, we present some numerical results of interest by solving the problem (3) numerically. The most computationally expensive part of our calculations, is determining the set of independent sets $\mathcal{I}$ used in (3). Although this problem is computationally hard in general, we use a smart enumerative technique to compile the set of non-conflicting link activations, for several cases of our interest. For computing the numerical results, we make the following additional assumptions:

A1. The channel gains are modeled by isotropic path loss. To be precise, denoting the channel gain from a point x to a point y by G_{xy}, we assume that $G_{\{.\}}$ is given by $G_{xy} = (\frac{|x-y|}{d_0})^{-\eta}$. Here, η is the path loss exponent, and is typically between 2 and 4, and d_0 is the far-field crossover distance.[2]

[2]Note that we do not take into account the location-dependent shadowing component of the channel gain. We feel this is reasonable for two reasons. Firstly, the shadowing component is relatively static and not time-varying [8]. Secondly, our interest is more in observing overall trends, and distilling structural properties, rather than predicting exact values.

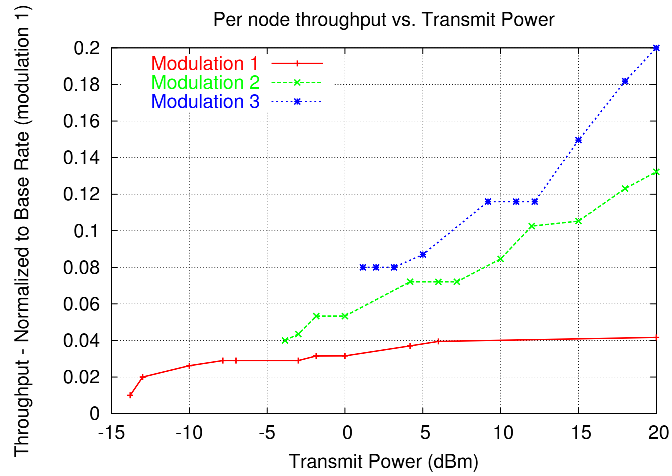

Fig. 1. Variation of λ^* with Transmit Power (in dBm)

A2. All nodes are physically separated by at least a distance d_{min}. This assumption is necessary since $d^{-\eta}$ grows unbounded as d approaches zero, thereby yielding arbitrarily high channel gains.

A3. The network is confined to a bounded area in space, say a square of size $L \times L$.

As mentioned earlier, we operate under a given fixed BER specification. Let us also denote the lowest SINR threshold for the given BER (which usually corresponds to the lowest rate modulation scheme) as β_{min}.

Proposition 3.1: Under the assumptions A1-A3, the maximum size of an independent set (or the maximum number of links that can be scheduled simultaneously) is bounded above by a constant B which depends only on η, d_{min}, L and β_{min}.

Proof: See Section V. □

Now, exploiting the fact that with a minimum node separation, there is a bound on the number of links that can be simultaneously activated in a given region, we enumerate only those subsets of links that are of a size smaller than the bound, and check whether those subsets are "independent sets". Once the set of non-conflicting schedules $\mathcal{A}$ has been so characterized, solving the problem is just a matter of solving the linear program in (3).

Let us now summarize the physical layer parameters used to derive the results in this section. The channel attenuation is modeled by an isotropic path-loss with an exponent of 4. The far-field crossover distance is taken as 0.1 m. Three modulation schemes are considered with normalized data-rates of 1, 4 and 8, and SINR thresholds of 10 dB, 20 dB and 25 dB, respectively. All the nodes operate at the same transmit power, and use the same modulation and coding scheme, unless otherwise stated. In what follows, the term "transmission range" (for a given power and modulation scheme), is used to refer to the maximum transmitter-receiver separation under which successful packet decoding remains possible for that power and modulation, in the absence of any co-channel interference.

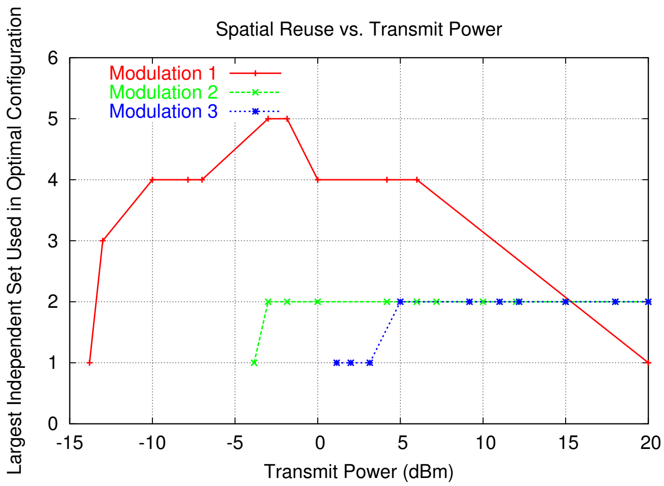

Fig. 2. Variation of Spatial Reuse with Transmit Power (in dBm)

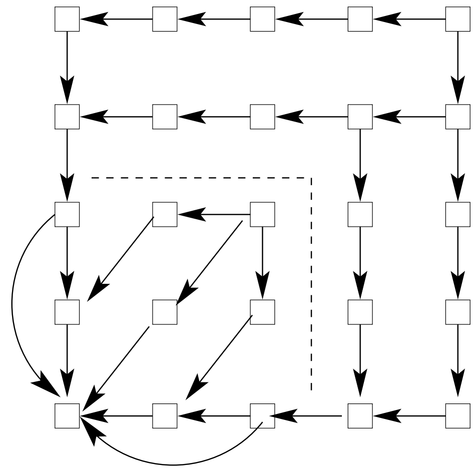

Fig. 3. Illustration of Optimal Routing: Modulation 1, Transmit Power -1.85 dBm.

Our first study investigates the achievable throughput of a 24 node sensor network deployed as a 5x5 grid (see Figure 3 and 4). All the nodes generate the same amount of traffic intended for the sink/gateway node at the bottom left corner (which generates no traffic). The separation between adjacent nodes along the grid side (not diagonal) is 8 m. The optimal max-min throughput is plotted as a function of the transmit power in Figure 1, for different modulation and coding schemes. The leftmost point on each of the three sets of curves, indicates the minimum transmit power at which the network is connected. As can be seen from the figure, lower rate modulation schemes provide connectivity at low transmit powers, but cannot obtain any significant gains in throughput at higher transmit powers. Indeed, for any value of the transmit power, the highest rate modulation scheme under which the network is connected, should be used. For modulation 1, the network operates as a single-hop network at 20 dBm, with each sensor directly communicating with the sink node. This is also the point of maximum throughput. Observe that about 50% or 70% of the maximum throughput is achieved at much lower powers (-13 dBm and -7 dBm, respectively).

To better understand the reason behind the initial steep increase and gradual flattening out of the throughput curves, let us look at Figure 2 which shows the size of the largest independent set used in the optimal configuration (which is a measure of spatial reuse) as a function of the transmit power. Focusing on the curve for modulation 1, we can observe that when the network is just barely connected, all the links formed are quite weak, and have little immunity to interference. Hence, they can only be scheduled one at a time (as indicated in Figure 2). As the transmit power increases, the spatial reuse in the network steadily improves as the links become more and more immune to interference. Beyond a certain transmit power, as longer links start getting created, it becomes more favorable for the data to be routed over longer links which means fewer hops, at the cost of spatial reuse. Here the corresponding increase in throughput is not as dramatic. Also, observe that using higher rate modulation schemes, means less immunity to interference, and reduced spatial reuse, although that is more than made up by the data-rate increase.

Now let us look at Figure 3 and Figure 4 which depict the optimal routing under different choices of transmit power and modulation and coding scheme. The range of a sensor node in the configuration in Figure 3 is 16 m which is twice the grid side, and in Figure 4 it is 12.7 m which is more than the grid diagonal. Note that in both the cases, the optimal routes are not minimum hop, for all the nodes. Within the region indicated by the dotted line in both the figures, the nodes use minimum hop paths, although in Figure 3 some sensors split their data along multiple paths not all of which are minimum hop. Also, observe that the nodes along the diagonal, beyond the dotted line, route their data along the periphery of the network. This is an illustration of what may be termed *interference-avoiding* routing.

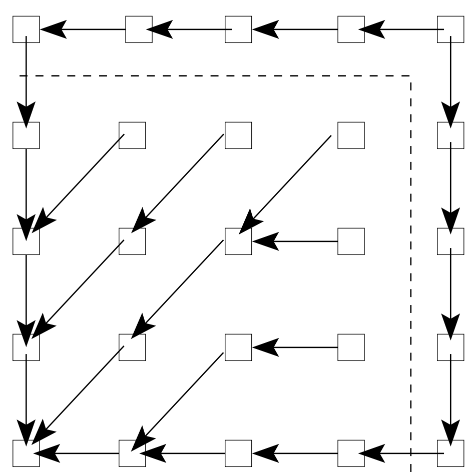

Fig. 4. Illustration of Optimal Routing: Modulation 2, Transmit Power 4.185 dBm.

The next example (see Figure 5) provides a clearer illustration of interference-avoiding routing. The topology used in

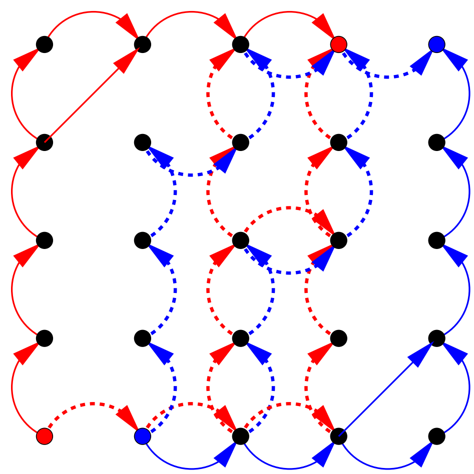

Fig. 5. Illustration of Optimal Routing: Modulation 1, Transmit Power 3 dBm.

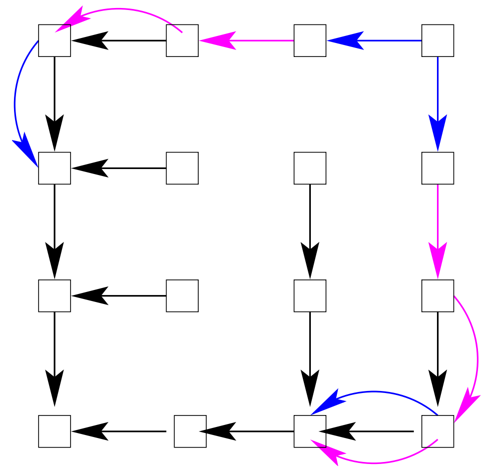

Fig. 6. 4×4 grid. 2 power levels and 2 modulation schemes at each node. Modulation 1 rate is 4 times that of modulation 0. $\lambda^* = 0.112$.

this example is identical. However, there are only two data-flows (indicated as red-straight edges and blue-curved edges). The red flow originates at the bottom left corner node, and is destined to the red node adjacent to the top right corner. Likewise, for the blue flow. The optimal routes which achieve the max-min throughput for the two data-flows, are indicated by the red and blue links in Figure 5. The dotted links carry less than 15% of the total traffic, and more than 80% of the total traffic is carried along the periphery over the solid links. In this case, the range of each node is 11.3 m which is equal to the grid diagonal. However, the routing uses only two diagonal links, and is far from minimum hop. Although some data is routed along common paths and links, the bulk of the data is routed so that the flows "avoid" each other.

We now consider an example in which 2 power levels and 2 modulation-coding schemes are available at each node. Figure 6 shows an optimal configuration for a network of 15 nodes placed on a grid and sending data to a sink/gateway at the bottom left corner. Power levels are -3 dBm and 2 dBm. Modulation 0 has rate 1 and SINR threshold 10 dBm whereas modulation 1 threshold is 20 dBm and rate 4. The other parameters are identical to the example in Figure 3. The transmission range for -3 dBm power is 14.95 m with modulation 0 and 8.36 m with modulation 1. For 2 dBm power these numbers are 19.85 m and 11.41 m resp. We find that the maximum throughput λ^* equals 0.112. Observe from Figure 6 that the optimal routing is an interesting balance of links of different physical layer parameters, and is not simply a minimum hop routing. Links operating at -3 dBm and modulation 1 are not utilized at all, the reason being that though they fetch high data rate, they are highly susceptible to interference. A link employing 2 dBm power and modulation 1 (shown as black-solid edges) provides the highest data rate but also has relatively bigger conflict set (its range, 11.41 m, is not very large in comparison to the grid side, 8 m). Hence such links are used to provide "fat pipes" when the flows start aggregating near the sink. Interestingly, flows from the 4 nodes in the center are also routed on short high data rate links. We believe that this is in accordance with the "shortest air-time path routing" (see Section II)- significantly higher data rate on these links means low air-time, thus, making several short high rate hops efficient in comparison to one long low rate hop.

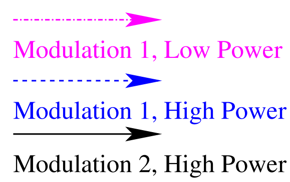

Fig. 7. Legend for Figure 6

Finally, we consider the problem of constructing a throughput-optimal network of 36 subscriber stations deployed in hexagonal cells (of side 8 m) and connected to the Internet through a base-station at the center. Each node has one up-link flow, i.e., to the base-station, and one down-link flow, i.e., from the base-station. Since the Internet traffic is asymmetric in general, the up-link flow rate is chosen to be 30% of the down-link flow. Figure 8 shows an optimal network when each node can use one power level (-13 dBm) and one modulation scheme of rate 1. These parameters yield a transmission range of 8.41 m. The red-solid (resp. blue-dashed) segments represent the down-link (resp. up-link) flow. The optimal up-link flow rate λ_u^* is 0.00433 and the optimal down-link flow rate λ_d^* is 0.0144. While these numbers as such may not be of practical interest, this example shows two things: (i) the routes obtained are quite unlike the tree-based structures proposed for routing over IEEE 802.16-like networks, indicating the importance of taking interference into account; and (ii) our optimization/computational framework can be utilized as a tool in planning and designing such IEEE 802.16 based access networks.

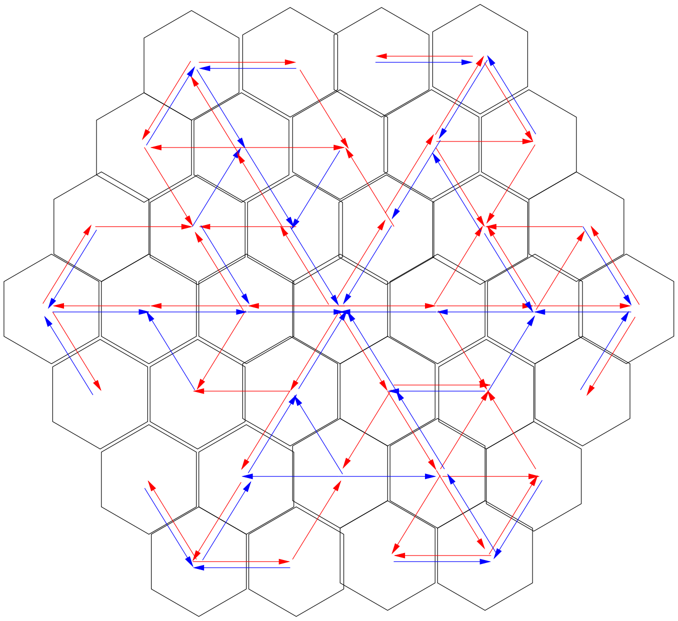

Fig. 8. Optimal network of base-stations deployed in hexagonal cells. Transmit power is -13 dBm. Up-link rate $\lambda_u^* = 0.00433$ and down-link rate $\lambda_d^* = 0.0144$.

IV. Conclusions

Our work addresses the following two questions concerning the optimal throughput of a given wireless networks: (i) what is the max-min throughput for an arbitrary set of nodes and data flows? and (ii) what is the optimal network configuration to achieve the same? We answer these questions via an optimization framework, using a conflict set formulation motivated by signal decoding in the presence of noise, to model the wireless channel interference. By means of several numerical case-studies, we show how this framework can be used to provide insights into dimensioning transmit power budgets and modulation strategies, and also as a direct tool, for capacity planning and configuration of wireless networks.

V. Proofs

Proof of Proposition 2.1: $\lambda_f^*(\phi)_{f\in\mathcal{F}}$ are optimal, hence feasible flow rate in the parameterized problem. This implies that $\sum_{l\in c} \frac{\sum_{f\in\mathcal{F}} \lambda_f^*(\phi)\left(\sum_{r\in\mathcal{R}_f^l} \phi_f^r\right)}{c_l} \leq 1$ for each $c \in \mathcal{C}$. Recall that this is a necessary condition for schedulability in terms of clique feasibility. Since $\lambda^*(\phi) \leq \lambda_f^*(\phi)$ for each $f \in \mathcal{F}$, it follows that

$$\lambda^*(\phi) \leq \frac{1}{\max_c w_c(\phi)}$$

On the other hand,

$$\sum_{l\in c} \frac{\sum_{f\in\mathcal{F}} \lambda_f(\phi)\left(\sum_{r\in\mathcal{R}_f^l} \phi_f^r\right)}{c_l} \leq \kappa \qquad (5)$$

implies there exist $(\lambda, \bar{\lambda}, \alpha)$ realizing flow rates $(\lambda_f)_{f\in\mathcal{F}}$. This is, therefore, a sufficient condition for feasible flow rates under given routing variables. Clearly over all flow rates satisfying (5) the optimal solution of the parameterized problem is $\hat{\lambda}(\phi) := \frac{\kappa}{\max_c w_c(\phi)}$. Since $\lambda^*(\phi) \geq \hat{\lambda}(\phi)$, it follows that

$$\frac{\kappa}{\max_c w_c(\phi)} \leq \lambda^*(\phi) \leq \frac{1}{\max_c w_c(\phi)}$$

Proposition now follows by taking the maximum over all routing variables $\phi = (\phi_f^r)_{f\in F, r\in R_f}$. This is due to the fact that λ^* can be obtained as a solution to the parametrized version of (1), where the parametrization is with respect to the routing variables ϕ. $\square$

Proof of Proposition 2.2: Right inequality is straightforward. For the left inequality note that

$$\begin{aligned} \lambda^*(\hat{\phi}) &\geq \frac{\kappa}{\max_c w_c(\hat{\phi})} & (6)\\ &= \frac{\kappa}{\min_\phi \max_c w_c(\phi)} & (7)\\ &\geq \kappa\lambda^* & (8)\end{aligned}$$

(6) follows from the clique sufficiency condition (see Proof of Proposition 2.1), (7) from the definition of $\hat{\phi}$ as the "min-max routing" and (8) from the left inequality in Proposition 2.1. $\square$

Proof of Proposition 3.1: We show this by using a packing argument to place the transmitters of links, as closely as possible, such that the links still form an independent set. Consider two links, say l and m. Without loss of generality, let the transmitter of link m, m_o, use higher transmit power. Let us denote the transmitter-receiver separation of link l by x, and the distance of m_o from l_d by y. Now, under the assumptions A1-A3, the minimum value of y for a given x, so that link m does not interfere with link l is bounded below by $x\beta_{min}^{1/\eta}$. This can be derived as follows. Use P to denote the transmit power of link m. Then the signal strength of link l is upper bounded by $P(x/d_0)^{-\eta}$, and the interference perceived by link l is lower bounded by $P(y/d_0)^{-\eta}$. This gives an upper bound on the SINR of link l, and therefore a lower bound of y so that l does not encounter a packet decoding failure. Now, by the triangle inequality,

$$\begin{aligned} |m_o - l_o| &\geq |m_o - l_d| - |l_d - l_o| & (9)\\ &= y - x \geq x(\beta_{min}^{1/\eta} - 1) & (10)\\ &\geq d_{min}(\beta_{min}^{1/\eta} - 1) & (11)\end{aligned}$$

where the last inequality follows from the fact that nodes are separated by a minimum distance d_{min}. Thus, for the links l and m to be independent, the distance between their transmitters has to exceed $d_{min}(\beta_{min}^{1/\eta}-1)$. This is equivalent to embedding each transmitter at the center of an exclusion disc of radius $\frac{1}{2}d_{min}(\beta_{min}^{1/\eta} - 1)$, and requiring discs to be non-intersecting, in order for the corresponding links to be independent. Since the network is confined to a region of area L^2, the maximum size of an independent set is upper bounded by,

$$B = \frac{4L^2}{\pi(d_{min}(\beta_{min}^{1/\eta} - 1))^2}$$

□

ACKNOWLEDGMENTS

This work was supported in part by a research grant from NSF (contract no. 0087266), by an NSERC Discovery Grant (contract no. RGPIN42023-37200), and by a grant from Bell University Laboratories (BUL).

REFERENCES

[1] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, Vol. IT-46, no. 2, pp. 388-404, March 2000.

[2] V. Mhatre and C. Rosenberg. The Capacity of Random Ad Hoc Networks under a Realistic Link Layer Model. *submitted to IEEE Transactions on Information Theory, 2005.*

[3] L.-L. Xie and P. R. Kumar. A Network Information Theory for Wireless Communication: Scaling Laws and Optimal Operation. *IEEE Transactions on Information Theory*, Vol. 50, no. 5, pp. 748-767, May 2004.

[4] M. J. Neely, E. Modiano, and C. E. Rohrs. Dynamic power allocation and routing for time-varying wireless networks. *IEEE Journal of Selected Areas in Communications*, 23(1):89–103, January 2005. Special Issue on Wireless Ad Hoc Networks.

[5] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. In *Proceedings of the ACM/IEEE MobiCom 2003*, 2003.

[6] R. L. Cruz and A. V. Santhanam. Optimal routing, link scheduling and power control in multi-hop wireless networks. In *Proceedings of the IEEE INFOCOM 2003*, April 2003.

[7] A. Bar-Noy, A. Mayer, B. Schieber, and M. Sudan. Guaranteeing Fair Service to Persistent Dependent Tasks. *SIAM Journal of Computing, vol. 27, No. 4, pp. 1168-89, August 1998.*

[8] V. Erceg, L. Greenstein, Y. Tjandra, S. Parkoff, A. Gupta, B. Kulic, A. Julius and R. Bianchi. An Empirically-based Path Loss Model for Wireless Channels in Suburban Environments. *IEEE J. Sel. Areas Commun., vol. 17, pp 1205-1211, July 1999.*

[9] A. Karnik, A. Iyer and C. Rosenberg. Throughput-optimal Configuration of Wireless Networks. *submitted to IEEE Trans. on Networking*.