Joint Scheduling, Power Control and Routing in Symmetric, One-dimensional, Multi-hop Wireless Networks

Bozidar Radunovic, Jean Le Boudec

Joint Scheduling, Power Control and Routing in Symmetric, One-dimensional, Multi-hop Wireless Networks

Božidar Radunović and Jean-Yves Le Boudec
bozidar.radunovic@epfl.ch, jean-yves.leboudec@epfl.ch
EPFL, CH-1015 Lausanne, Switzerland *

Abstract

We are interested in finding a jointly optimal scheduling, routing and power control that achieves max-min fair rate allocation in a multi-hop wireless network. This is a highly complex non-convex optimization problem and it has been previously solved only for small networks. We restrict ourselves to symmetric networks with ring and line topologies, and we numerically solve the problem for a large number of nodes. We model point-to-point links as single user Gaussian channels where nodes cannot send and receive at the same time. This type of channel approximates the performance of CDMA networks and performs better than the equivalent 802.11 network.

We show that for smaller transmission powers it is optimal to relay over other nodes whereas for high powers it is optimal to send data directly to a destination. We also show when this transition occurs. We analyze the optimal schedule and find that if a node is active, it should send at the maximum power. When a transmission power is small, the optimal schedule is that every second node is sending, and as the power grows, the distance between active nodes grows. Furthermore, in large networks the distance between nodes sending at the same time is never larger than 4.5 times the size of links used (number of nodes spanned by one transmission link), and it converges to that value for large transmission powers.

1 Introduction

It has recently become evident that a traditional layering network approach, separating routing, scheduling, flow and power control, is not efficient for ad-hoc wireless networks [1]. This is primarily due to the interaction of links through interference, which implies that a change in power allocation or schedules on one link can induce changes in capacities of all links in the surrounding area and changes in the performance of flows that do not pass over the modified link.

There are several papers that study variations of the joint optimization problem. In [12] the joint scheduling and power control problem is considered in networks with QoS constraints, where a minimum signal to noise ratio is defined for every link. Given these set of constraints, they find an optimal scheduling and power allocation that satisfies constraints and minimizes dissipated power. A similar model with a minimum SIR constraint is analyzed in [5]. In [7], the authors find optimal scheduling and power control that maximizes total throughput of a network, given the power constraints for each user. They solve the problem for small networks and they demonstrate a distributed algorithm that finds an approximate solution for large networks based on hierarchy. Joint routing and scheduling is considered in [3], where a total throughput of a network is maximized when given a set of links that cannot be active at the same time.

In [10] the authors define a very general model of a wireless network that covers both routing, scheduling and power control. They take the total network capacity as a performance measure. However, the complexity of the model is such that even with the linear objective it can handle less than 10 nodes.

Per-link utility fairness in multi-hop wireless networks is discussed in [13] and is achieved by different scheduling. A similar approach for max-min fairness can be found in [4, 14].

In this paper we focus on best-effort networks, and we want to find scheduling, power allocation and routing that achieves the max-min fair rate allocation. This is a highly complex non-convex optimization problem for a general network topology. A similar framework in [10] has been solved for up to 10 nodes. It is thus difficult to draw a general conclusion about network design from such small networks.

In order to obtain results for larger networks, we focus on one-dimensional network topologies, where all nodes are aligned on a straight line. These topologies represent a large class of existing networks, from car

*The work presented in this paper was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

networks on highways to networks on coasts or mountain valleys. If a network has a finite size we replace a line with a ring in order to avoid border effects.

We also assume that a network is symmetric, that is, all nodes are equally spaced and each node has data to send to its d hop away neighbor on the right. This simplification allows us to immediately identify the optimal schedule, as will be shown later. The symmetry argument is not always realistic, but we think it gives us an important insight into design of homogeneous networks. We model a link between two nodes as a point to point single user Gaussian channel with deterministic fading and a unique power limit for all nodes. This model closely approximates CDMA networks and performs better than the equivalent 802.11 network.

Finally, we allow the two most common routing policies. One is that each node sends data directly to its destination (DIR), hence there is no relaying. The other is that each node forwards data to the destination by relaying it over all intermediate nodes, thus using minimum energy path (MER). There are clearly other relay routing policies that might perform better than minimum energy routing. However, minimum energy routing is frequently used, simple to implement, and performs comparably to direct routing.

The first question we answer in this paper is, in a given network, should it relay packets or not. In other words, which is better for optimal scheduling and power allocation, minimum energy routing or direct routing? We also characterize the optimal scheduling and the optimal power allocation for a given routing. From these characterizations we can derive heuristics that can be applied on homogeneous one-dimensional networks.

This paper is organized as follows: In section 2 we present a general model and its simplification in the case of a symmetric one dimensional network. In section 3 we analyze the proposed model. In section 4 we present numerical results and in section 5 we conclude and give directions for future work.

2 Assumptions and Modeling

2.1 MAC Model

Let us consider a wireless network on a plane represented by a set of nodes $\{1, \cdots, n\}$, and their positions. Let us denote by $d(i, j)$ distance between nodes i and j. Let $\vec{p}$ represent a vector of powers allocated to each of the n nodes, where power p_i of each node is limited by a maximum power P_i. We use a deterministic fading model where fading between two nodes depends only on their distance. A signal emitted at power P from node i is received by node j at power $P' = h_{ij} * P$ where $h_{ij} = d(i, j)^{-\alpha}$, and α is a constant fading factor.

We further assume noise from others sources as well as background noise to be Gaussian, hence the total interference at a receiver is also Gaussian. Given a power allocation $\vec{p}$, we have that the SIR ratio at a link (i, j) is

$$\gamma_{ij}(\vec{p}) = \frac{h_{ij} p_i}{N + \sum_{k \neq i} h_{kj} p_k},$$

where N is a background noise.

Furthermore, data is separately encoded on each hop, and we assume that a receiver cannot decode third party communications, hence it treats it as noise. Then each hop can be modeled as a single user Gaussian channel. The Shannon capacity bound for this type of channel can be achieved using contemporary CDMA technology. It is also higher than the rate achieved by 802.11 since 802.11 has zero throughput if SNR is lower than the threshold and has constant throughput if SNR is above the threshold. From the information theoretic point of view, hop by hop coding is clearly not the optimal solution. However, it is used in most of today's wireless networks. Also, the capacities of more advanced relay or broadcast Gaussian channels are known only in some special cases [11], and these cases are not applicable in our model. Finally, we assume that a destination can receive data from or send data to only one source at a time, and it cannot send while receiving.

Given the above assumptions, the maximum achievable rate between nodes i and j given a fixed power allocation $\vec{p}$ is $x_{ij} = R_{ij}(\vec{p}) = 1/2 \log(1 + \gamma_{ij}(\vec{p}))$ bits/s/Hz, and we denote by $\vec{x} = \vec{R}(\vec{p})$ a vector of rates on each point to point link achieved with a fixed power allocation $\vec{p}$. Finally, we observe that for any two feasible power allocations $\vec{p}_1$ and $\vec{p}_2$ and corresponding rate allocations $\vec{x}_1$ and $\vec{x}_2$, we can also achieve a convex combination of corresponding rates $\vec{x} = \alpha\vec{x}_1 + (1 - \alpha)\vec{x}_2$ by time-divisioning. Assuming that this time-divisioning is averaged on a long term, a set of feasible link rate allocations is

$$\mathcal{X} = \text{Hull}(\bigcup_{\vec{p} \in \mathcal{P}} \vec{R}(\vec{p})).$$

2.2 Routing

Communication can be done by letting each source transmit data directly to its destination. We call this routing policy direct routing (**DIR**), and it represents a single-hop network where each source sends data directly to its destination without relaying.

As mentioned in the introduction, the main question is whether one can improve performance by relaying data over intermediate nodes. Since in this setting each node can communicate to every other node, there is an exponential number of possible relaying paths. Additionally, each source can use several paths in parallel, depending on the load on each path.

In this paper we focus on minimum energy routing (**MER**) as a relay routing policy. This is a single-path routing policy where each source uses the path that minimizes the sum of fadings over each hop; this is indeed the path that minimizes energy consumption of nodes for a single source destination pair, ignoring global side-effects. Although other relay routing policies might yield a better performance, we analyze this one in order to simplify the model, because it is the most straightforward to implement, and because it outperforms direct routing in large number of cases. When MER performs better than DIR we say that it is beneficial to relay.

2.3 Topology

The above problem is difficult to solve in the general case (discussion for arbitrary networks with up to 6 nodes are given in [10]). In order to analyze networks with a larger number of nodes, we restrict our attention to ring and line topologies with a high level of symmetry, depicted on Fig. 1.

Figure 1: Analyzed topologies: ring and line.

Ring topology represents an n-sided regular polygon with distance l between nodes. Line topology is a limiting case when the number of nodes in a ring tends to infinity while l remains fixed.

2.4 Other assumptions and notations

The maximum power for all nodes is the same, and equal to P. Each node is a source of data, and its destination is d-hop away on the right, where $d < n/2$. We also assume that all nodes use the same routing policy, either MER or DIR, and we see immediately that with a minimum energy routing, each node will send all data only to its one-hop neighbor on the right.

In the paper, we use the following notation: n - number of nodes in a ring, d - distance between a source and a destination, P - maximum transmitting power, R - radius of a ring, α - fading coefficient, **MER** and **DIR** - minimum energy and direct routing policies.

2.5 Performance objectives

In most of the papers concerning optimal performance of wireless networks, such as [12, 5, 7, 3, 10], the authors maximize total throughput of the network. This approach can lead to gross unfairness, where users with worse channels might not get any throughput. Per-link utility fairness has been considered in [13], and per-link max-min fairness in [4, 14]. However, this still doesn't guarantee per-flow fairness for long flows.

Our performance objective is to find a scheduling, routing and power allocation such that long term average flow rates will have max-min fair rate allocation. This means that not one rate of a single flow can be increased without decreasing an already smaller flow rate. Since the feasible region is a convex set, if follows from [2] that the max-min fair rate allocation exists.

3 Analysis

In this section we present several analytical findings for our model.

3.1 Max-Min Fairness

The first proposition is about a property of the max-min fair rate allocation.

Proposition 1 *In the above defined wireless network, if a rate allocation is max-min fair then each source - destination pair has the same rate.*

Proof: Suppose this is not true, and that for some flows i and j we have $f_i > f_j$. Let B_i and B_j be sets of bottleneck links of flow i and j, respectively. Then there exists a time slot where $p_l > 0$ for each $l \in B_i$. By decreasing each $p_l, l \in B_i$ by $\epsilon > 0$, we will decrease rate f_i to a smaller f_i'. We do the same for other flows that share the same bottlenecks, hence they will all have the same new rate f_i'. This in turn increases rates of links in B_j, hence we can increase f_j as well. If ϵ is sufficiently small, we obtain a new feasible rate allocation where we increased a rate of a flow without decreasing a rate of an already smaller flow, which contradicts with the definition of max-min fairness.

q.e.d.

Since all flow rates are the same, we can write $f = f_i$. Due to symmetry in routing, one can easily verify that all link rates are the same, hence we can also write $x_i = x$. For the direct routing, a link capacity corresponds to the rate of a flow, hence the rate of flow $f = x$. For the minimum energy routing we have d flows sharing the same link, hence $f = x/d$.

3.2 Scheduling

We next analyze scheduling and describe the feasible link rate allocation set $\mathcal{X}$. In the case of MER routing, we have n one-hop links, and in the case of DIR routing we have n d-hop links. According to the Carathéodory theorem, every point $\vec{x}$ in the set of feasible link rate allocations set $\mathcal{X}$ can be expressed as a linear combination of $n+1$ points in $\bigcup_{\vec{p}\in\mathcal{P}} \vec{R}(\vec{p})$,

$$\vec{x} = \sum_{i=1}^{n+1} \alpha_i \vec{R}(\vec{p}_i). \tag{1}$$

Vectors $\vec{p}_i$ represent power allocations within slots, and vector α describes frequencies of these slots.

Let us call $\mathcal{R}_i(\vec{p})$ a rotation of $\vec{p}$ such that for all $j \in 0 \cdots n-1$,

$$(\mathcal{R}_i(\vec{p}))_j = (\mathcal{R}_i(\vec{p}))_{(j+i) \bmod n}.$$

For the ring topology, we have the following proposition:

Proposition 2 *In the above depicted scenario, the optimal schedule consists of n rotationally symmetric power vectors $\vec{p}_i = \mathcal{R}_i(\vec{p})$ that are equally frequent, that is $\alpha_i = 1/n$.*

Proof: Let $\vec{x}$ be the vector of the optimal link rates. From proposition 1 we know that all links have the same rate, that is for all i and j, $(\vec{x})_i = (\vec{x})_j$.

Since $\vec{x}$ is rotationally symmetric we can achieve the same rate by rotating each power allocation by an arbitrary k, hence from (1):

$$\vec{x} = \sum_{i=1}^{n+1} \alpha_i \vec{R}(\mathcal{R}_k(\vec{p}_i)).$$

Now, if $\vec{x}_i = 1/n \sum_{k=1}^{n} \vec{R}(\mathcal{R}_k(\vec{p}_i))$ then all $\vec{x}_i$ are also rotationally symmetric. Since

$$\vec{x} = \sum_{i=1}^{n+1} \alpha_i \vec{x}_i \leq \max_i \vec{x}_i,$$

we conclude that for some $\vec{p}$ we can represent a max-min fair allocation as

$$\vec{x} = \frac{1}{n} \sum_{k=1}^{n} \vec{R}(\mathcal{R}_k(\vec{p})).$$

q.e.d.

Since all link rates are the same, we can also write

$$x(\vec{p}) = \frac{1}{n}\sum_{k=1}^{n} R_{1,(d+1)}(\mathcal{R}_k(\vec{p})), \qquad (2)$$

and we see that vector $\vec{p}$ that maximizes $x(\vec{p})$ is the optimal power allocation that achieves per-flow max-min fair rate allocation.

3.3 Power Allocation

In the previous section we showed that the optimal scheduling strategy is to allocate equal time to each rotation of a single power allocation vector $\vec{p}$. We now want to characterize the optimal allocation $\vec{p}$. We noticed that for an arbitrary feasible power allocation $\vec{q}$, if we fix $q_2, \cdots, q_n$, the rate is maximized if q_1 is either 0 or P. We were not able to formally prove this statement. However, we found empirically that the rate is a quasi-convex function [9] of the power q_1, which implies the $0/P$ property:

Claim 1 *Let us consider $x(q_1)$ as a function only of the first component of vector $\vec{q}$. Then, for arbitrary values of components $q_2, \cdots, q_n$ of $\vec{q}$, function $x(q_1)$ is quasi-convex.*

One way to prove that $x(q_1)$ quasi-convex is to show that if $\frac{\partial x}{\partial q_1} = 0$ then necessarily $\frac{\partial^2 x}{\partial q_1^2} \geq 0$. We numerically tested this claim for rings with n up to 6. We also observed that for large n, when a ring can be approximated with a line, only a few of the closest neighbors significantly contribute to the interference. Specifically, if we take the most dense power allocation where every second node is sending at maximum power, and we take $\alpha = -4$, then nodes further than 3 hops away from a destination contribute to the overall interference 0.47%, and for $\alpha = -2$, it contributes about 15%. Using this approximation we numerically verified this claim for large n as well.

Proposition 3 *If claim 1 is true, then in the optimal power allocation $\vec{p}$ each power p_i is either 0 or P.*

Since claim 1 is true for arbitrary $q_2, \cdots, q_n$, it is also valid for $p_2, \cdots, p_n$, hence we conclude p_1 in the optimal allocation $\vec{p}$ has to be either 0 or P. The same reasoning applies to all coordinates of $\vec{p}$.

q.e.d.

Restricting ourselves to power allocations where all nodes either send at the full power or not send at all, we have the following claim:

Claim 2 *Optimal power allocation $\vec{p}$ consists of several groups of adjacent active nodes. Each group has a size from 1 to d, and the differences in sizes of any two groups is at most 1. The difference in distances between two adjacent groups of active nodes for any two pairs of groups is at most 1. Especially, in case of MER routing, the size of all active groups is one.*

An example of the optimal power allocation is illustrated on Fig. 2. We numerically tested this claim and found it true for networks of up to 20 nodes. According to claim 2 we can then define t as an average size of a group of active nodes, and s as an average distance between the first active nodes of two adjacent groups. Since these parameters vary for at most one, they depict well the optimal power allocation, especially in large networks, and thus in the line case. Note that for MER routing t will always be 1, and for DIR we have $t \leq d$.

4 Results

In this section we present numerical solutions to the joint optimization problems for various values of parameters. These parameters are a number of nodes n, a distance to the destination d, a relative maximum power $Pl^{-\alpha}/N$ and a fading coefficient α. Again, the two main questions that we seek to answer are when it is beneficial to relay (i.e., when is MER routing better than DIR), and what is the optimal power allocation, for various values of system parameters.

In order to obtain numerical solutions we used different techniques. We first solved the problem over all possible power allocations, using eq. (2) as a rate equation. Since $R_{1,(d+1)}(x) = \log(1 + x)$, it is easy to show that this optimization problem is a d.c. programming problem [8]. Although we solved this problem with a branch and bound approach, the solution was too complex to be applied for networks larger than a few nodes. These solutions however verified claim 1. We next used claim 1 and we solved the problem by searching over 2^n possible power allocations for networks with up to 20 nodes. Finally, we used claim 2 to solve the problem for a line and an arbitrary large ring (since a very large ring can be approximated with a line). In the following sections, we present numerical results and conclusions. Where it is not explicitly stated, we assume $\alpha = -4$.

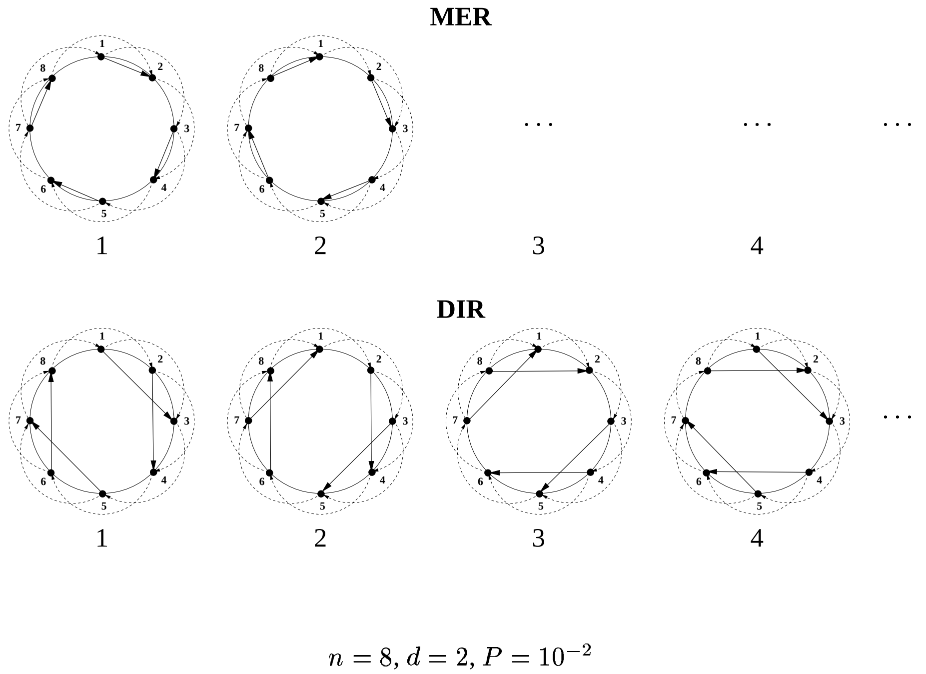

Figure 2: Illustration of claim 2. On the top, the optimal scheduling for the two routing policies is given, for $n = 8$, $d = 2$, $P = 10^{-2}$. Dashed arrows depict flows, and solid arrows depict active links. Numbers below are time slot numbers. Each link is either inactive, or active at the full power. There are 8 allocations in total, but the first 4 are exactly the same as the last 4, for both routing policies. These 8 allocation repeat in a row, each taking the equal time slot, and each one is a rotation of the previous one. In MER routing, distance s between active nodes is 2. In DIR routing, there are 2 groups of $t = 2$ active nodes and distance between them is $s = 4$.

$n = 8, d = 2, P = 10^{-2}$

MER	DIR
10101010	11001100

$n = 18, d = 2, P = 10^{-2}$

MER	DIR
101010101010101010	110011001100011000

Figure 3: Short representations of the optimal policies policies for $n = 8$ and $n = 18$: one rotation of the optimal power allocation is given, each link denoted with 0 if inactive and 1 if active. The left case $n = 8$ is the one depicted on Fig. 2. In MER routing, every second node is active. In DIR routing, for $n = 8$ there are 2 groups of 2 active nodes, and distance between them is 4. For $n = 18$, there are 4 groups of 2 active nodes and distances between them are 4 and 5.

4.1 The Ring Case

We first consider an $n = 18$ node ring. We search over 2^n possible power allocations to find the optimal one for both MER and DIR routing. The results are depicted on Fig. 4. On the left we see the rate per flow, and on the right we see the rate per distance per flow, as defined in [6]. For small enough powers it is better to use minimum energy routing, hence to relay, and for large power it is optimal to use direct routing.

Fig. 5 shows at what power limit the transition from MER to DIR occurs, for different network sizes and flow lengths. The larger the network is and the shorter the flows are, the more spacial reuse and more incentive to relay there is. We also see that transition power is log-linear with respect to the size of the network n and is super log-linear with respect to the flow length d.

Another interesting question to consider is the optimal power allocation, given the system parameters. An example is given on figure 6, for an 18 node ring with flow length of 3. We see that for very small powers, as many nodes as possible are active at the same time, because the interference is small. In MER this means every second node, hence $s = 2$. A group cannot be larger than one, since link size is 1, hence $t = 1$. In DIR, it means first d nodes are active and the next d nodes are receiving, hence inactive, and so on, leading to $t = d$ and $s = 2d$. When power grows, the distance between active groups s grows, and group size t shrinks. Eventually, for very large powers, we will have only one active node at a time for both routing policies. We then have a log-linear increase of rate with power in both policies, since now power allocation, scheduling and topology are fixed and a change in power limit directly increases rate. It is also interesting to notice that DIR is never better than MER

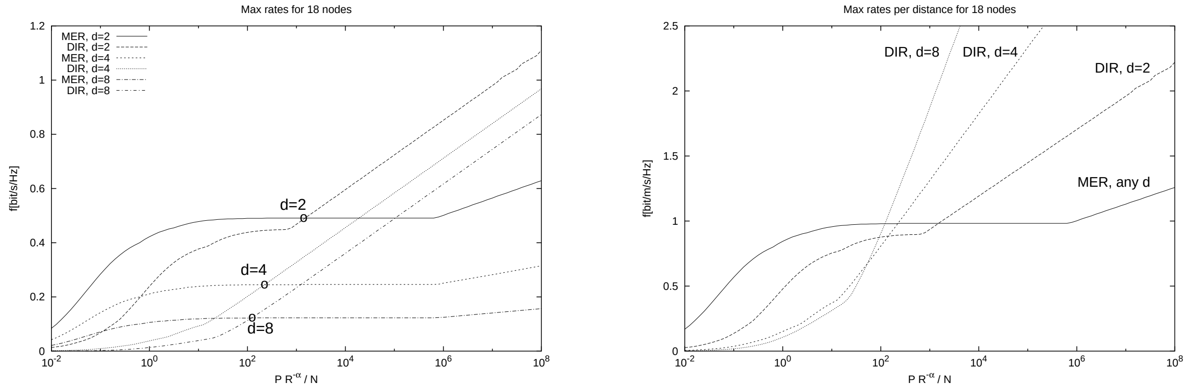

Figure 4: On the left, maximum rates for DIR and MER routing and different values of d are depicted, for a ring of 18 nodes. We see that for small transmitting powers MER routing is better then DIR, and for large powers DIR becomes better. The transition power depends on d. On the right, maximum rate per distance is depicted. We see that rate per distance is the same for MER, regardless of d. This is not the case for DIR.

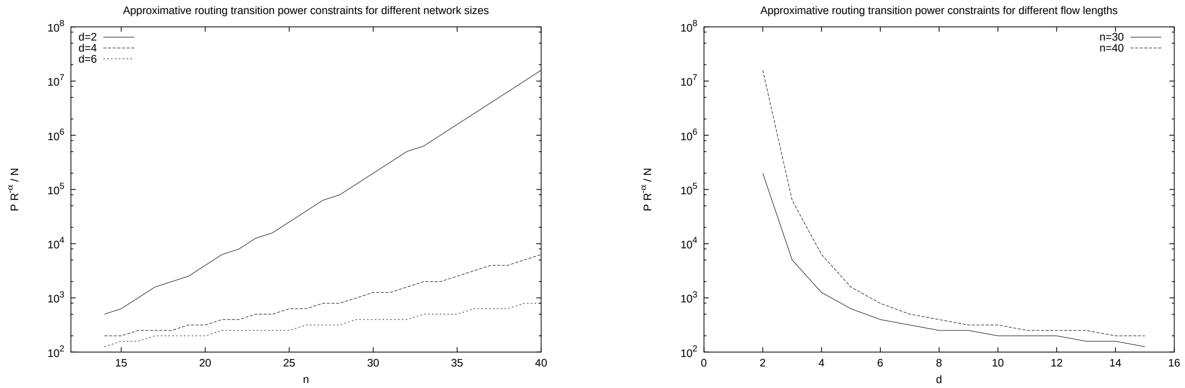

Figure 5: The two figures depicts at which power constraint a transition from minimal-energy routing to direct routing occurs, as a function of n and d.

P	MER	s	DIR	(s,t)	Better
10^{-2}	**101010101010101010**	2	111000111000111000	(6,3)	MER
10^{-1}	**100100100100100100**	3	110000110000110000	(6,2)	MER
1	**100100100010001000**	3.6	100000000100000000	(9,1)	MER
10	**100010001000010000**	4.5	100000000100000000	(9,1)	MER
10^5	100010001000010000	4.5	**100000000000000000**	(18,1)	DIR
10^6	100000000000000000	18	**100000000000000000**	(18,1)	DIR

Figure 6: Optimal scheduling for different power limits, for both MER and DIR (n=18, d=3). The first column gives the power limit. The second gives the optimal schedule for MER routing, where 1 means a node is on and 0 means a node is off, and s is the average distance. The third column gives the optimal schedule for DIR and the average distance between adjacent active groups s and average size of an active group t(t is always 1 for MER). The last column tells which routing performs better. The optimal joint scheduling and routing policy is shown in boldface.

when $t > 1$ since the interfering node is closer to the destination than the sender.

The number of active nodes can be seen on Fig. 7. This number is the same for all slots, and since all active nodes use the same power, it is directly proportional to the dissipated energy. We see that for smaller powers, both routing policies dissipate approximately the same energy. When the power gets larger, DIR uses only 1 active node while MER uses 4, hence MER dissipates 4 times more energy at a lower rate (see Fig. 4). For even larger powers, energy dissipation is the same since only one node is active at a time.

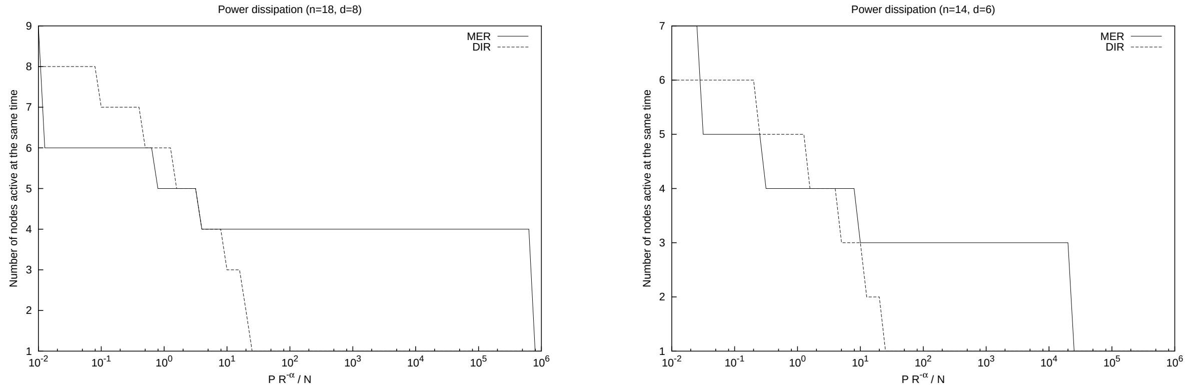

Figure 7: Number of nodes active in one time slot. Since the emitting powers or all active nodes are the same, and the number of active nodes is the same in all slot, this number is directly proportional to dissipated energy.

4.2 The Line Case

Next, we consider an infinite line with nodes equally spaced on a distance l. Since this is the limiting case of rings when n tends to infinity and distance l between nodes remains constant, these results also apply to rings with a large n. On the left of Fig. 8 we see maximum rates for the two routing policies. We first note that DIR routing is never better than MER (it is actually only slightly better, and the reason for this lies in the optimal power allocation, as is discussed later). This is in accordance with Fig. 5, where we see that a transition power constraint grows exponentially with a number of nodes. From Fig. 8 we see that for smaller rates, MER performs better, and when power constraint is sufficiently high, DIR and MER performs the same. The same is observed for different values of α. The transition when DIR becomes as good as MER is depicted on the right of Fig. 8. We see that the transition function is log-log linear with respect to d, where the slope depends on α (the higher α is, the more spacial reuse is possible, and transition occurs later). An example of the optimal schedule for the line case with $d = 3$ is given on the Fig. 9.

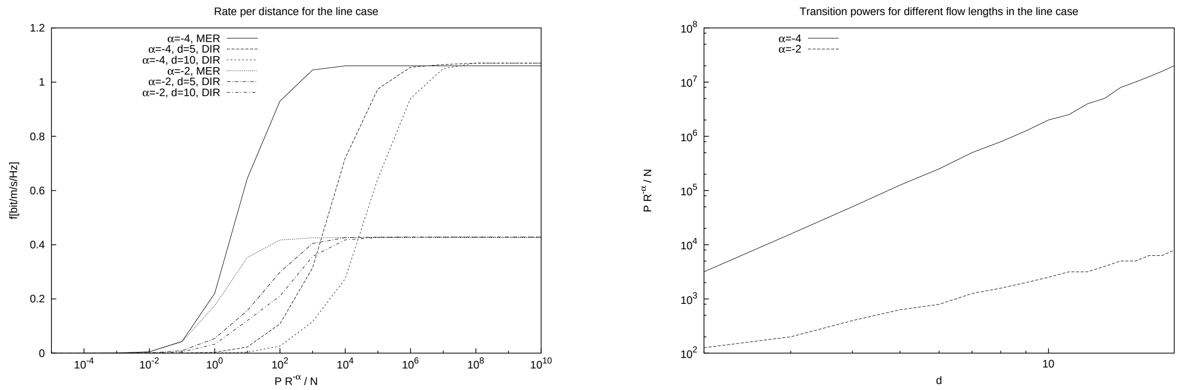

Figure 8: On the left, rates per distance achieved in the line case, for both routing policies, different flow lengths and power constraints. On the right, transition power for different flow length, on a log-log scale.

The fact that DIR is never better than MER in the line case is somewhat different from the ring case. In the ring case (Fig. 4), each node was given a fixed time slot of a positive length. When transmitting power was high, optimal power allocation contained only one active node, and by increasing the transmission power it was able to increase arbitrary its rate. This was true for both MER and DIR, whereas the DIR rate in that case grew faster since relaying was expensive. In the line case, this scenario is no longer possible, since we always have more than one node active at the same time and the increase in the transmission powers will eventually be canceled by interference.

Similar phenomenon can be observed on a transition diagram for the line (Fig. 8 on the right) and for the ring (Fig. 5 on the right). For the line case, as the flow length increases, a transition power increases, whereas for the ring this dependency is inverse. This is again because for finite rings and DIR routing, the higher d is, the smaller the power for which optimal power allocation has only one active node is, and the higher the derivative of the rate over transmitting power is (as can be seen on Fig. 4). As it is not possible to have only one active node in an infinite network, the transition diagram changes and relaying becomes cheaper (a transition power grows with d).

P	MER	s	DIR	(s,t)	Better
10^{-2}	**...1010...**	2	...111000111000...	(6,3)	MER
1	**...100100...**	3	...11100001110000...	(7,3)	MER
10	**...100100...**	3	...11000001100000...	(7,2)	MER
100	**...10001000...**	4	...1000000010000000...	(8,1)	MER
10^3	**...10001000...**	4	...100000000100000000...	(9,1)	MER
10^4	**...1000010000...**	5	...1000000000010000000000...	(11,1)	MER
10^5	**...1000010000...**	5	**...10000000000001000000000000...**	(13,1)	both
10^6	**...1000010000...**	5	**...10000000000001000000000000...**	(13,1)	both

Figure 9: Optimal scheduling in the line case, for different power limits, for both MER and DIR routing (d=3), in the same format as on the Fig. 6. Dots before and after an array of ones and zeros means that the schedule consists of the same pattern infinitely repeating. While MER routing is better in the first six cases, both routing performs the same in the last two cases.

Since in the line case MER is always optimal, we conclude that rate per distance is constant, regardless of the flow length. This in turn means that network transport capacity is roughly $\Theta(n)$ compared to the transport capacity $\Theta(\sqrt{n})$ of an arbitrary network on a unit disk, as in [6]. This is due to the fact that in our networks traffic is distributed locally ($d \ll n$) whereas in [6] sources and destinations are distributed uniformly throughout the network.

Finally, we analyze the optimal power allocation for the line case. On Fig. 10 we see how it evolves as n grows, for MER. For small powers we have $s = 2$. We see that as n grows, a power for which we have only one active nodes grows. For powers slightly lower than this transition one, we see that s is approximately 4.5. In the line case, we also have $s = 2$ for small powers, and it converges to around 4.5 as the power goes to infinity, as can be seen on the left of Fig. 11. Since we allow only integer average distances in our approximate model (that is, all adjacent groups of active nodes are equally spaced), this limiting s is approximated with 5. We also see from Fig. 11 that the same phenomenon occurs for $\alpha = -2$, hence it is not sensitive to the fading factor.

In the case of DIR routing, a similar behavior can be observed. For very small powers, the distance between groups of active nodes is $s = 2d$, $t = d$ and with this allocation each node either sends or receives data. As the power grows, s increases and d decreases. As the power tends to infinity, t becomes 1 and s becomes $4.5 * d$. Therefore, for very large powers we can approximate a line (l, d) with a line $(ld, 1)$. Here we also see why there is a difference between rates on the left of Fig. 8. Specifically, optimal s is 4.5 and in the MER case this is approximated with 5; it is thus obvious that in the DIR case, the error of the integer approximation of the optimal $s = 4.5 * d$ will be smaller than in the MER case, hence the rate will be higher. Again, the conclusions are the same for $\alpha = -2$.

In the line case, as in the ring case, we can analyze the energy dissipated by a node by counting a fraction of active nodes in the same slot. On the right of Fig. 11 we see the fraction of nodes active in one slot. Since all active nodes sends with the same power and since this fraction is constant in all slots, it is directly proportional to the energy a node dissipates. Optimal power allocation for MER does not depend on flow length d, hence neither does energy dissipated. The total rate however decreases as d increases due to relaying. On the other hand, for high powers DIR can achieve the same rates as MER, but using d times less energy. For low powers DIR achieves lower rates than MER with approximately the same energy dissipation.

5 Conclusions and Future Work

We have given a general model for a joint scheduling, power allocation and routing optimization problem. We solved the problem for symmetric one-dimensional networks, for both direct and minimum energy routing policies. We found that for small power constraints it is better to relay, and for large power constraints it is better to send data directly to destinations. We also describe for which power constraints this transition occurs. In the case of an infinitely large network, minimum energy routing is always as good as direct routing. However, in contrast with their names, direct routing dissipates less energy than minimum energy routing for large power, while achieving the same rate.

We characterized optimal scheduling and power allocation for both routing policies, and we found that due to the symmetry, an ideal schedule consists of n equal time slots, and in each time slot we use a rotation of a single power allocation. We further show that in this single optimal power allocation, each node either sends at the full power or does not send at all. Nodes that send at the same time are grouped in equally sized sets of adjacent nodes. Distances between two adjacent groups is the same for all pairs of adjacent groups. The groups' sizes and

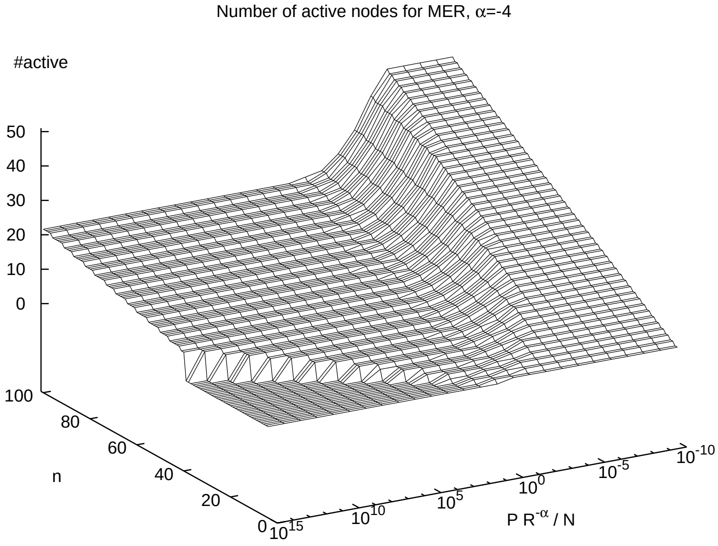

Figure 10: Number of nodes transmitting at the same time (inverse of s) for MER. For large n it converges to between 4 and 5 nodes. Since distance between active nodes grows with power, DIR will have the same scheduling behavior for large powers. Note that the flow length does not change scheduling policy for minimal energy routing, but only total rate.

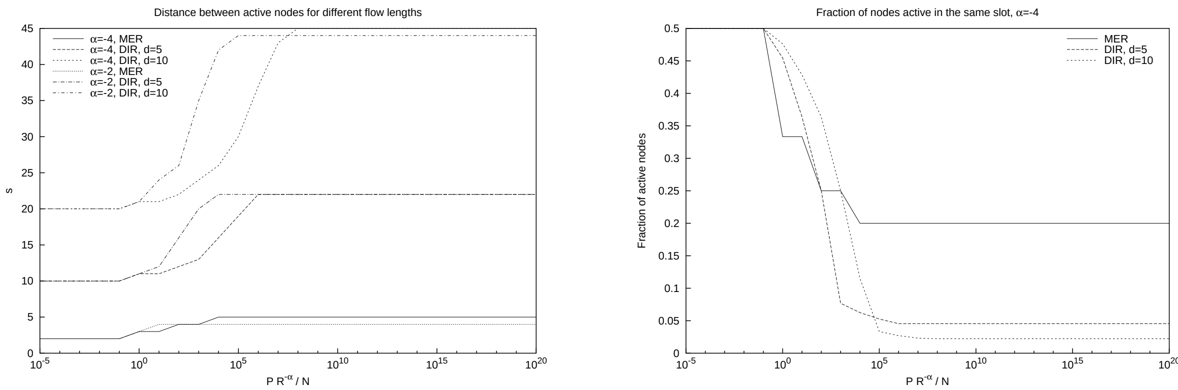

Figure 11: On the left is the average distance s between adjacent groups of active nodes. For small powers, s is 2 for MER and $2d$ for DIR. It converges to approximately 4.5 for MER and $4.5d$ for DIR. On the right is the fraction of nodes that are active in the, same slot, which is directly proportional to dissipated energy.

spacing depends on the power constraint, and for large networks and large power constraints group sizes converge to 1 and the distance between groups converges to 4.5 times the size of a used link (1 for MER, d for DIR).

In the future, we intend to study the sensitivity of presented results to the model assumptions, in particular the symmetry in topology and traffic matrix, and to see if and how much can scheduling, power allocation and routing strategies suggested in this paper improve performance of an arbitrary network. We would also like to extend this analysis to two dimensional networks.

References

[1] A. Goldsmith, S.B.Wicker. "Design challenges for energy-constrained ad hoc wireless networks". *IEEE Wireless Communications*, 9(4):8–27, August 2002.

[2] B. Radunović, J.Y. Le Boudec. "A Unified Framework for Max-Min and Min-Max Fairness with Applications". In *Proceedings Allerton'02*, 2002.

[3] L. Tassiulas, A. Ephremides. ”Jointly optimal routing and scheduling in packet radio networks”. *IEEE Transactions on IT*, 38(1):165–168, January 1992.

[4] L. Tassiulas, S. Sarkar. ”Max-min Fair Scheduling in Wireless Networks”. In *Proceedings INFOCOM'02*, 2002.

[5] N. Bambos, S. Chen, G. Pottie. ”Radio Link Admission Algorithms for Wireless Networks with Power Control and Active Link Quality Protection”. In *Proceedings INFOCOM'95*, 1995.

[6] P. Gupta and P.R. Kumar. ”The Capacity of Wireless Networks”. *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000.

[7] R. Cruz, A.V. Santhanam. ”Optimal Link Scheduling and Power Control in CDMA Multihop Wireless Networks”. In *Globecom'02*, 2002.

[8] R. Horst, P. Pardalos, N. Thoai. *Introduction to Global Optimization*. Kluwer Academic Publishers, 2000.

[9] S. Boyd, L. Vandenberghe. *Convex Optimization*. http://www.stanford.edu/class/ee364/, 2003.

[10] S. Toumpis and A.J. Goldsmith. ”Capacity regions for wireless ad hoc networks”. *IEEE Transactions on Wireless Communications*, to appear.

[11] T. Cover and J.A. Thomas. *Elements of Information Theory*. John Whiley & Sons, 1991.

[12] T. Elbatt, A. Ephremides. ”Joint Scheduling and Power Control for Wireless Ad-hoc Networks”. In *Proceedings INFOCOM'02*, 2002.

[13] T. Nandagopal, et al. ”Achieving MAC Layer Fairness in Wireless Packet Networks”. In *Proceedings MobiCom'00*, Boston, Massachusetts, August 2000.

[14] X. Huang, B. Bensaou. ”On Max-min Fairness and Scheduling in Wireless Ad-Hoc Networks: Analytical Framework and Implementation”. In *Proceedings MobiHoc'01*, Long Beach, California, October 2001.

Biography

Božidar Radunović graduated in 1999 from School of Electrical Engineering, University of Belgrade. He is currently pursuing his PhD at EPFL. His interests are in the architecture and design of wireless multi-hop networks. His email address is bozidar.radunovic@epfl.ch.

Jean-Yves Le Boudec graduated from Ecole Normale Superieure de Saint-Cloud, Paris, received his doctorate in 1984 from the University of Rennes, France. In 1987 he joined Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory where he was manager of the Customer Premises Network Department. In 1994 he became professor at EPFL, where he is now full professor. His interests are in the architecture and performance of communication systems. His e-mail address is jean-yves.leboudec@epfl.ch.