

Rate Performance Objectives of 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 consider the question of what performance metric to maximize when designing ad-hoc wireless network protocols such as routing or MAC. We focus on maximizing rates under battery lifetime and power constraints. Commonly used metrics are total capacity (in the case of cellular networks) and transport capacity (in the case of ad-hoc networks). However, it is known in traditional wired networking that maximizing total capacity conflicts with fairness, and this is why fairness oriented rate allocations, such as max-min fairness, are often used. We review this issue for wireless ad-hoc networks. Indeed, the mathematical model for wireless networks has a specificity that makes some of the findings different. It has been reported in the literature on Ultra Wide Band that gross unfairness occurs when maximizing total capacity or transport capacity, and we confirm by a theoretical analysis that this is a fundamental shortcoming of such metrics in wireless ad-hoc networks, as it is for wired networks. The story is different for max-min fairness. Although it is perfectly viable for a wired network, it is much less so in our setting. We show that, in the limit of long battery lifetime, the max-min allocation of rates always leads to strictly equal rates, regardless of the MAC layer, network topology, choice of routes and power constraints. This is due to the “solidarity” property of the set of feasible rates. This results in all flows receiving the rate of the worst flow, and leads to severe inefficiency. We show numerically that the problem persists when battery lifetime constraints are finite. This generalizes the observation reported in the literature that, in heterogeneous settings, 802.11 allocates the worst rate to all stations, and shows that this is inherent to any protocol that implements max-min fairness. Proportional fairness is an alternative to max-min fairness that approximates rate allocation performed by TCP in the Internet. We show by numerical simulations that proportional fairness of rates or transport rates is robust and achieves a good trade-off between efficiency and fairness, unlike total rate or maximum fairness. We thus recommend that metrics for the rate performance of mobile ad-hoc networking protocols be based on proportional fairness.

Index Terms—System design, Mathematical programming/optimization, wireless, max-min, utility fairness, best-effort

I. INTRODUCTION

A. Rate-based Performance Metrics with Power Constraints

We consider the question of what metric to use when evaluating the performance of ad-hoc wireless network protocols such as routing or MAC. We focus on maximizing rates under battery lifetime and power constraints. Typical application examples are networks of wireless laptops and PDAs; this is also the framework used by many papers analyzing various models of physical layers (purely information theoretic approach [12], [5], CDMA [18], UWB [6]). In contrast, some sensor networks put more emphasis on minimizing energy under minimum rate constraints. We study here the former and leave the latter to a companion paper.

For cellular wireless networks, a frequently chosen performance metric is total capacity, i.e. the sum of the rates of all flows. An extension that maximizes a weighted sum of the rates is applied in CDMA/HDR [16]. In multi-hop wireless networks, the same metrics are used, but also transport capacity, a variant popularized by Gupta and Kumar in [17]. This is in fact a weighted sum of rate, where weights are the distances between the source and the destination of each flow.

B. The Tension Between Efficiency and Fairness

The tradition in wired networking has also focused on performance metrics that incorporate some form of fairness. Indeed, it is known that considering only total capacity yields gross unfairness if implemented in a wired network [14]. Therefore, different performance metrics that account for fairness have been developed. A typical example is max-min fairness [4], which is used in many existing networking protocols, including the ABR mode of ATM [1]. This is an egalitarian approach by which the rate of a flow can be increased only when it is not possible to increase the rate of an already smaller flow. Max-min fairness is often viewed as an extreme fairness; this justifies using a *fairness index*, which measures the departure from max-min fairness (it

is a slight variant of the fairness index defined by Jain in [19]; see Section III-D).

Max-min fairness is also used, often implicitly, in many existing wireless multi-hop network protocols (e.g. [13], [26]). In fact, as we show, 802.11 essentially implements max-min fairness. However, in wireless networks, there is still no tradition of evaluating a system in light of both total rate and fairness. It turns out that the issue is significantly different than in wired networking, due to the peculiarities of the mathematical models for wireless networks. In particular, we find that the allocations that implement max-min fairness have fundamental efficiency problems. This is due to the “solidarity” property of the set of feasible rates (Section IV-A).

Another way to reduce the tension between efficiency and fairness is to use a weighted sum of the rates as a design objective. The most well known example of this type of criteria in wireless networks is transport capacity [17] where each flow is assigned a weight equal to the distance between the source and the destination of the flow. We show in Fig. 4 in Section VI that this approach does not reconcile the tension.

C. Utility Fairness

Utility fairness is often used as an alternative, less egalitarian approach to max-min fairness. It corresponds to the “utility” metric $\sum_j U(x_j)$ where x_j is the rate of flow j and $U()$ is a concave function (called the utility function); $U()$ is interpreted as user satisfaction by Kelly et al. [7]. Maximizing the utility metric is known, in wired networking, to be fairer than maximizing the total capacity, but less egalitarian than a max-min fair allocation. The Internet congestion control performed by TCP approximates some form of utility fairness. A special case, widely used in economy, is *proportional fairness*, which has $U(x) = \ln(x)$ [7]. Variants of utility fairness are used in existing wireless multi-hop network protocols as well (e.g. [15]).

Note that the utility approach can easily be extended to account for power and energy not in the form of constraints as we do here, but through a cost function subtracted from the utility metric. This is explored for example by Baldi et al [3]. We leave such metrics out of the scope of this paper, as we focus on rate-based metrics with power constraints.

D. Reported Facts In The Context of Wireless Networks

The tension between efficiency and fairness was reported by Tse and Hanly in [5] for a cellular network. A strategy that maximizes the total capacity is such that a node with the best channel conditions in a given slot

should send data. Nodes that are farther away will less frequently satisfy this constraint, but will still have a positive throughput, due to the random part of fading. However, if a node is very far away from the base station, its average rate is going to be very small and essentially it will not be able to communicate. In [5], a remedy is found by assigning weights to node rates, such that a level of fairness is assured. The implicit assumption in this type of network is that an area with mobile nodes is well covered with base stations, so there is no big variation in distances from mobiles to closest base stations.

However, variations in the distances between sources and destinations in the case of ad-hoc networks are typically much higher since a node does not talk to the closest base-station but to an arbitrary destination in a network. This makes it difficult to remedy fairness with weights, and longer flows risk low or zero throughput. Indeed, it has been observed in the context of Ultra Wide Band by Cuomo et al [6] that the unfairness of total capacity persists in wireless networks and some long distant flows obtain zero throughput.

A performance anomaly was reported by Berger-Sabatel et al in [8] for 802.11. There, several nodes talk to a base station. One of them is far away and codes for 1 Mb/s while others are near and code for 11 Mb/s. Still, on average, all nodes achieve the same throughput of approximately 1 Mb/s. We show later in this paper that this anomaly is in fact not an abnormal behavior, rather a fundamental property of max-min fairness for wireless networks, regardless of any underlying physical, MAC or routing protocol.

E. Modeling of Ad-hoc Wireless Networks

We are interested in a model of a wireless network in order to analyze efficiency and fairness of different design criteria for various network technologies. General models of wireless networks that incorporate various physical layers, and MAC and routing protocols are discussed in [18], [23], [11]. Maximizing a weighted sum of the rates as a criteria, is considered in [18], [23] for a very general model of a network. Proportional fairness and maximizing the minimal rate in a network (a weaker version of max-min fairness) are analyzed in [11]. However, the latter considers only a subset of possible routing and MAC protocols, those that can be transformed to convex problems.

We define a model of an ad-hoc wireless network that allows the most general assumptions on a physical layer (including variable rate 802.11, UWB or CDMA), MAC and routing protocols. And for a given network

topology and traffic demand, we characterize a set of feasible end-to-end rate and transport rate allocations. Next, we find the optimal allocations on the two sets with respect to the three design criteria considered. In some numerical examples, where it is not possible to find an exact solution of the optimization problem due to its non-convexity, we consider an approximation that is close to the optimal solution and that allows us to accurately characterize the efficiency and fairness of the optimum.

F. Our Findings

We prove that under a general model of an ad-hoc wireless network, in the limit of long battery lifetime, max-min fairness leads to equal rates of all flows, regardless of network topology, routing or power constraints. This means that all rates are equal to the rate of the worst flow, making the network very inefficient. The same happens when considering transport rates. We show numerically that the problem persists with battery lifetime constraints. This conclusion is in sharp contrast with the findings from the framework of wired networks, where max-min fairness is widely used. Also, this generalizes the result in [8]; it shows that their finding is not a unique property of 802.11 and that any protocol that strives for max-min fairness will have the same problem.

We also prove that, for sufficiently high powers, a protocol that maximizes total capacity always starves flows with bad channel conditions for sufficiently high powers, that is, only the most efficient flow gets a positive rate and all other flows have a zero rate. We verify numerically, on a large number of random networks, that this unfairness occurs not only at the limit, but also with realistic transmission power constraints. This generalizes results in [6], showing that this unfairness property is not a problem of UWB but rather of the design criteria. We also show that the use of transport capacity, although fairer than total capacity, does not completely compensate unfairness, and can also assign zero rates to the worst flows.

We further show that for very small battery lifetimes, the max-min fair, proportionally fair and rate maximizing allocations are equivalent. In this limiting setting, fairness is not an issue and any of these metrics can be used in a design. However, we find that this, in general, does not hold for realistic power constraints.

Finally, we show that proportional fairness is a robust trade-off between fairness and efficiency, insensitive to different transmission power and long-term average power constraints, and network topologies. Thus an ideal candidate metric when designing or evaluating a

performance of an ad-hoc wireless network is the sum of the logarithms of the achieved rates over all source destination pairs. This also suggests that 802.11 should be redesigned with proportional fairness as a design objective, in order to avoid inefficiencies observed in [8].

G. Organization of This Paper

The outline of the paper is as follows. The next section describes system assumption. In Section III we give mathematical formulation of the model of a network. In Sections IV, V and VI we present findings related to max-min fairness, maximizing total capacity and proportional fairness objectives, respectively. In Section VII we discuss the influence of long-term average power constraints. In the last section we give conclusions and directions for further work. Proofs are moved to the appendix for the sake of clarity. They can also be found in [21].

II. SYSTEM ASSUMPTIONS

We analyze an arbitrary ad-hoc wireless network that consists of a set of nodes, and each two nodes that directly exchange information are called a link. For each pair of nodes we define a signal attenuation, that is a level of signal received at the receiver, assuming the sender is sending with unit power. This attenuation is usually a decreasing function of a link size due to power spreading in all directions, but here we assume it is an arbitrary number defined for each pair of nodes. We assume the network is located on a finite surface and that all attenuations are strictly positive, hence every node can be heard by any other node in the network and there is no clustering.

There is also a random component of the signal attenuation, due to changes of characteristics of paths the signal takes. This component is called random fading. It usually has a smaller order of magnitude than the constant attenuation and we do not model it here.

We next give properties of the physical model of communications on links.

A. Physical Model Properties

All physical links are point-to-point, this means each link has a single source and a single destination. There are more advanced models such as broadcast channel and relay channel [24] that attain higher performances, but they are not used in most of the contemporary networks, and their performance is in general not known and is still an open research issue. Broadcast is used by 802.11 by MAC layer control packets, but this is an aspect we do not model here.

A node can either send to one next hop or receive from one at a time. There are more complex transmitter or receiver designs that can overcome these limitations. An example is a multi-user receiver that could receive several signals at the time. This would change the performance of links having a common destination, but would not change the interactions over a network. However, these more complex techniques are not used in contemporary multi-hop wireless networks (like 802.11, UWB, bluetooth or CDMA) due to high transceiver complexity, and we do not analyze them here.

We model rate as a strictly increasing function $r(\text{SNR})$ of the signal-to-noise ratio at the receiver, which is a ratio of received power by the total interference perceived by the receiver including the ambient noise and the communications of other links that occur at the same time. This model corresponds to a large class of physical layer models, for example:

- Shannon capacity of a Gaussian channel [24]: $r(\text{SNR}) = 1/2 \log_2(1 + \text{SNR})$.
- Ultra-wide band [6]: $r(\text{SNR}) = K \times \text{SNR}$.
- CDMA HDR [16]: $r(\text{SNR})$ is a stair function of SNR.
- variable rate 802.11 [9]: $r(\text{SNR})$ is a stair function of SNR.

We note that in the last two models, rate is not a strictly increasing function of SNR but in most applications can be approximated as such. On the contrary, in the basic model of 802.11 (e.g. [17]), the rate is assumed to be constant hence this model does not fit this framework.

B. MAC Protocol

We further assume a slotted protocol. In each slot a node can either send data, receive or stay idle, according to the rules defined above. Each slot has a power allocation vector associated with it, which denotes what power is used for transmitting by the source of each link. If a link is not active in a given slot, its transmitting power is 0. A schedule consists of an arbitrary number of slots of arbitrary lengths.

We assume an ideal MAC protocol that calculates the optimal transmission power of each link in each slot in a centralized manner and according to a predefined metric. This is equivalent to a network where nodes dispose of an ideal control plane with zero delay and infinite throughput to negotiate schedule and power allocation. A more realistic MAC protocol would introduce some errors and delays, but a good approximation should be close to the ideal case. Also, by considering an ideal protocol, we focus our analysis on properties of performance metrics, and not artifacts of leaks in protocol design. Our

assumption corresponds to neglecting the overhead (in rate and power) of the actual MAC protocol.

C. Routing Protocol and Traffic Flows

We assume an arbitrary routing protocol. Flows between sources and destinations are mapped to paths, according to some rules specific to the routing protocol. At one end of the spectrum, nodes do not relay and only one-hop direct paths are possible. At the other end, nodes are willing to relay data for others and multi-hop paths are possible. There can be several parallel paths. All these cases correspond to different constraint sets in our model, as explained in Section III-A. Sources can send to several destinations (multicast) or to one (unicast).

D. Power Control

There are three types of power constraints in a wireless network: peak constraint, short-term average constraint and long-term average constraint. Here we describe them in detail:

Peak power constraint: Given a noise level on a receiver, a sender can decide which codebook it will use to send data over the link during one time slot. Different symbols in the codebook will have different powers. The maximum power of a symbol in a codebook is then called *peak power*. It depends on the choice of the physical interface and its hardware implementation and we cannot control it. It limits the choice of possible codebooks, and it puts restrictions on the available rate. For example, the rate of an UWB link, given the average SNR on the receiver, depends on the shape of the pulse, thus on the peak power level of the pulse [25]. In our model, the peak power constraint is integrated in a rate function, given as an input.

Transmission power : We assume a slotted system. In each slot a node chooses a codebook and its average power, and sends data using this codebook within the duration of the slot. We call transmission power, or simply transmission power the average power of a symbol in the codebook. This is a short-term average power within a slot, since a codebook is fixed during one slot. We assume that this transmission power is upper-bounded by P^{MAX} . This power limit is implied by technical characteristics of a sender and by regulations, and is not necessarily the same for all nodes. For example, this is the only power constraint that can be set by users on 802.11 equipment.

Long-term average power : While transmitting a burst of data (made of a large number of bits), a node uses several slots, and possibly several different codebooks. Each of these codebooks has its transmission

power. We call the long-term average power the average of transmission powers during a burst, and we assume it is limited by \bar{P}^{MAX} . Long-term average power is related to the battery lifetime in the following way:

$$T_{\text{lifetime}} \approx \frac{E_{\text{battery}}}{\bar{P}^{MAX} \times u}$$

where T_{lifetime} is the battery lifetime, E_{battery} is the battery energy, \bar{P}^{MAX} is the long-term average power constraint and u is the fraction of time a node has data to send (or activity factor, measured in Erlangs). The approximation corresponds to ignoring overhead spent managing the sleep / wakeup phases, etc. \bar{P}^{MAX} is thus set by a node to control its lifetime; it can vary from a node to a node.

We incorporate explicitly in our model the transmission power and the long-term average power constraints. The peak power is incorporated implicitly through the choice of the rate function.

III. MATHEMATICAL FORMULATION OF THE FEASIBLE SETS AND OF THE METRICS

A. Feasible Set of Rates

We model the wireless network as a set of I flows, L links and N time-slots. Every flow can use one or several paths (multicast or unicast). There are P paths ($P \leq 2^L$).

- $\mathbf{f} \in \mathbb{R}^I$ is the vector of average rates achieved by flows,
- $\bar{\mathbf{x}} \in \mathbb{R}^L$ is the vector of average rates that are achievable on links
- $\mathbf{y} \in \mathbb{R}^P$ is the vector of average rates used on paths
- F (flow matrix) is such that $F_{i,p} = 1$ if path p belongs to flow i , and 0 elsewhere. We have $\mathbf{f} = F\mathbf{y}$.
- R (routing matrix) is such that $R_{p,l} = 1$ if path p uses link l . We have $\bar{\mathbf{x}} \geq R\mathbf{y}$. The matrix R is defined by the routing algorithm.
- $h_{l_1 l_2}$ is the attenuation of a signal from the source of link l_1 to the destination of link l_2 .

Next, we assume that a schedule consists of time slots $n = 1 \dots N$ of duration α_n . We normalize these lengths such that $\sum_{n=1}^N \alpha_n = 1$. Let us call \mathbf{p}^n the vector of transmission powers assigned to links in slot n , and let \mathbf{SNR}^n be the vector of signal-to-noise ratios at receivers of links, induced by \mathbf{p}^n . If we denote by $r(\mathbf{SNR})$ the rate function, then the rate achievable on links in slot n on link l is $\mathbf{x}_l^n = r(\mathbf{SNR}_l^n)$. The vector of average rates on links is thus $\bar{\mathbf{x}} = \sum_{n=1}^N \alpha_n \mathbf{x}^n$. Since \mathbf{x}^n has dimension L (where L is a number of links), by virtue of Carathéodory theorem, it is enough to consider $N = L + 1$ time slots

of arbitrary lengths α in order to achieve any point in the convex closure of points \mathbf{x}^n .

We are interested in the set \mathcal{F} of feasible average flow rates. It is the set of $\mathbf{f} \in \mathbb{R}^I$ such that there exist a schedule α , a set of power allocations \mathbf{p}^n , corresponding set of rate allocations \mathbf{x}^n and average rates and powers $\bar{\mathbf{x}}$ and $\bar{\mathbf{p}}$, such that the following set of equalities and inequalities are satisfied:

$$\begin{aligned} \mathbf{f} &= F\mathbf{y} & \mathbf{SNR}_l^n &= \frac{\mathbf{p}_l^n h_{ll}}{N + \sum_{k \neq l} \mathbf{p}_k^n h_{kl}} \\ \bar{\mathbf{x}} &\geq R\mathbf{y} & \mathbf{x}_l^n &= r(\mathbf{SNR}_l^n) \\ \mathbf{p}_l^n &\leq P_l^{MAX} & \bar{\mathbf{x}} &= \sum_{n=1}^{L+1} \alpha_n \mathbf{x}^n \\ \bar{\mathbf{p}}_l &\leq \bar{P}_l^{MAX} & \bar{\mathbf{p}} &= \sum_{n=1}^{L+1} \alpha_n \mathbf{p}^n \\ & & 1 &= \sum_{n=1}^{L+1} \alpha_n \end{aligned} \quad (1)$$

B. Feasible Set of Transport Rates

In [17], the transport rate of a flow is defined as the rate of a flow multiplied by the distance covered by the flow between the source and destinations (call this $\text{len}(i)$ for flow i). Therefore, the set of feasible transport rates \mathcal{T} is defined as

$$\mathcal{T} = \{\mathbf{t} \in \mathbb{R}^I \mid (\exists \mathbf{f} \in \mathcal{F}) \mathbf{t}_i = \mathbf{f}_i \text{len}(i)\} \quad (2)$$

C. Design Criteria

Given network technology, for each topology and traffic demand there is a given set of feasible rates \mathcal{F} and a set of transport rates \mathcal{T} . We consider optimizing the system according to one of the following criteria:

1) Rate Criteria:

- **capacity:** maximize $\sum_{i=1}^I \mathbf{f}_i$ over all $\mathbf{f} \in \mathcal{F}$
- **max-min fairness:** find the max-min fair rate vector \mathbf{f}^* in \mathcal{F}
- **proportional fairness:** maximize $\sum_{i=1}^I \ln(\mathbf{f}_i)$ over all $\mathbf{f} \in \mathcal{F}$

The first and the third criteria are defined by concave maximization problems over \mathcal{F} that is convex and compact; thus they always have a solution. Max-min fairness is defined as follows [4]. We say a point \mathbf{x}^* is max-min fair on some set \mathcal{X} iff for all $\mathbf{x} \in \mathcal{X}$ and all index i $\mathbf{x}_i > \mathbf{x}_i^*$ implies that there exists an index j such that $\mathbf{x}_j < \mathbf{x}_j^* \leq \mathbf{x}_i^*$ i.e. increasing some component \mathbf{x}_i^* must be at the expense of decreasing some already smaller component \mathbf{x}_j^* . The max-min fair allocation does not always exist, but if it exists it is unique. It always exists if \mathcal{X} is convex and compact [20], which is the case here.

The max-min fair allocation does not have $\mathbf{x}_i^* = \mathbf{x}_j^*$ in general for $i \neq j$, even on convex sets (see [4] for some examples).

In general, the rate vectors that satisfy each of the three criteria are significantly different, as illustrated by the examples in the rest of the paper.

2) *Transport Rate Criteria*: Similarly, we define:

- **transport capacity**: maximize $\sum_{i=1}^I \mathbf{f}_i \text{len}(i)$ over all $\mathbf{f} \in \mathcal{F}$
- **transport-max-min fairness**: find the max-min fair transport rate vector \mathbf{t}^* in \mathcal{T}

Transport proportional fairness leads to the same objective as proportional fairness (up to a constant) and need not be considered separately. This is a nice feature of the proportional fairness criterion. In contrast, the rates that maximize transport capacity [resp. are transport max-min fair] differ from the rates that maximize capacity [resp. are max-min fair]. Existence and unicity hold for transport criteria in the same way as for rate criteria.

D. Performance Indices

In the rest of this paper we evaluate the properties of the optimal rates that correspond to each of the criteria above. It is convenient to use indices that quantify efficiency and fairness.

The **efficiency index** of a feasible rate \mathbf{f} in a given feasible set \mathcal{F} is $\frac{\sum_{i=1}^I \mathbf{f}_i \text{len}(i)}{\sum_{i=1}^I \mathbf{f}_i^c \text{len}(i)}$, where \mathbf{f}^c is the rate vector that maximizes capacity in \mathcal{F} . It is always between 0 and 1.

Similarly, the **transport efficiency index** of \mathbf{f} in \mathcal{F} is $\frac{\sum_{i=1}^I \mathbf{f}_i \text{len}(i)}{\sum_{i=1}^I \mathbf{f}_i^t \text{len}(i)}$, where \mathbf{f}^t is the rate vector that maximizes transport capacity in \mathcal{F} .

The max-min fairness index ϕ of a feasible rate \mathbf{f} in \mathcal{F} is defined as $\cos^2 \alpha$, where α is the angular deviation from \mathbf{f} to the max-min fair allocation \mathbf{f}^* in \mathcal{F} . Thus $\phi = \frac{(\mathbf{f}^T \mathbf{f}^*)^2}{(\mathbf{f}^T \mathbf{f})(\mathbf{f}^{*T} \mathbf{f}^*)}$. The max-min fairness index is between 0 and 1; it is equal to 1 if \mathbf{f} is proportional to the max-min fair allocation of rates. The smaller it is, the less fair the allocation is. When the number of flows L is large, the minimum value of the max-min fairness index is close to 0.

Our max-min fairness index coincides with Jain's definition of fairness index [19] in the case where the max-min fair allocation \mathbf{f}^* has all components equal. Otherwise, it differs.

The **max-min fairness index** of \mathbf{f} in \mathcal{F} is thus defined as $\frac{(\sum_i \mathbf{f}_i^* \mathbf{f}_i)^2}{(\sum_i \mathbf{f}_i^{*2})(\sum_i \mathbf{f}_i^2)}$, where \mathbf{f}^* is the max-min fair element of \mathcal{F} .

Similarly, the **transport max-min fairness index** of \mathbf{f} in \mathcal{F} is $\frac{(\sum_i \mathbf{t}_i^* \mathbf{f}_i \text{len}(i))^2}{(\sum_i \mathbf{t}_i^{*2})(\sum_i (\mathbf{f}_i \text{len}(i))^2)}$, where \mathbf{t}^* is the max-min fair element of \mathcal{T} .

E. Performance Metrics

The indices defined above require computing the reference rate vector that is optimal with respect to a design criterion, and depend on the set of rate vectors that is

being considered. In contrast, metrics are defined as a function of the rate alone, independent of any set of rate vectors. For completeness, we now give the metrics that correspond to the design criteria defined above. They may be useful in practical situations where, unlike in this paper, the computation of the reference rates is not feasible. This occurs for example when a protocol is given by its implementation in a simulator and the feasible set is hard to define explicitly.

For a rate vector \mathbf{f} , the **capacity metric** is $\sum_{i=1}^I \mathbf{f}_i$ and the **transport capacity metric** is $\sum_{i=1}^I \mathbf{f}_i \text{len}(i)$. They both measure the efficiency of \mathbf{f} .

A metric that corresponds to max-min fairness is more difficult to define. Many authors use Jain's fairness index defined above, but this is not always appropriate. Indeed, it measures the deviation from an ideal rate vector where all components are equal, and this is not necessarily the fairest vector. A more accurate, but more complex, metric uses leximin ordering [2], [22]. It is not a real number in the usual sense. Instead, the **fairness metric** $f(\mathbf{f})$ of a rate vector \mathbf{f} is the list of all its components in increasing order, and we say that a rate vector \mathbf{f}^1 is fairer than a rate vector \mathbf{f}^2 if $f(\mathbf{f}^1)$ is larger than $f(\mathbf{f}^2)$ in lexicographic order. The max-min fair vector is the fairest, in the sense of this metric. Similarly, the **transport fairness metric** is defined as the order statistic of the vector of transport rates $(\mathbf{f}_i \text{len}(i))_i$.

The **logarithmic utility** is $\sum_{i=1}^I \ln(\mathbf{f}_i)$. Again, while we can analogly define transport logarithmic utility, this is essentially the same metrics as logarithmic utility since they differ only by a constant additive factor, and we do not consider it here. Proportional fairness provides a combined measure of efficiency and fairness. It is maximized by the proportionally fair allocation. In this paper, we show that the design criteria based on proportional fairness is the best in the sense of robustness against efficiency or fairness anomalies. This suggests using logarithmic utilities as a metric of choice for evaluating ad-hoc wireless networks.

IV. MAX-MIN FAIRNESS

In this section we analyze properties of the max-min fair allocation. We show that there exists a class of convex sets with a property that a max-min fair vector on such a set has all components equal. We then show that a set of feasible rates in any wireless network without long-term average power constraints, modeled by (1) admits this property, implying that the rates in max-min fair allocation have to be equal.

A. Solidarity Property and Equality

Let us consider a class of sets in \mathbb{R}^n with a property that for any feasible point we can trade a sufficiently small value of one component for a sufficiently small value of another component. More precisely, we define the solidarity property as follows:

Definition 1: A subset \mathcal{X} of \mathbb{R}^n has the *solidarity property* iff for all $i, j, i \neq j$, for all $\mathbf{x} \in \mathcal{X}$ such that $\mathbf{x}_i > 0$, and for all $\epsilon > 0$ small enough, there exist positive $0 \leq \alpha_i < \epsilon$, $0 < \alpha_j < \epsilon$ such that $\mathbf{y} = \mathbf{x} - \alpha_i \mathbf{e}_i + \alpha_j \mathbf{e}_j$ belongs to \mathcal{X} .

Not all sets have solidarity property. In particular, not all convex set have solidarity property. Simple examples of networks with feasible rate sets with and without solidarity property are given on Fig. 1.

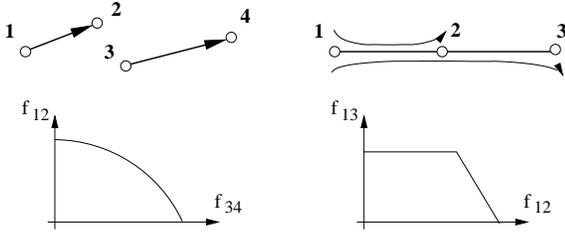


Fig. 1. On the left, there is an example of a wireless network whose set of feasible rates has solidarity property. Rate of flow 12 is constrained by $f_{12} \leq r(P_{12}h_{12}/(N + P_{34}h_{31}))$. Rate of flow 34 is constrained by $f_{34} \leq r(P_{34}h_{34}/(N + P_{12}h_{14}))$. It is always possible to increase the rate of one flow at the expense of the other. On the right there is an example of a wired network whose feasible rate set does not have solidarity property. Flows f_{12} and f_{13} are constrained by $f_{12} + f_{13} \leq x_{12}$ and $f_{13} \leq x_{23}$. When flow 13 hits limit on link 23 it cannot be further increased by decreasing rate of flow 12.

A characteristic of a set with solidarity property is that all components of the max-min fair vector are equal. This is formulated in the following proposition:

Proposition 1: If a set \mathcal{X} has the solidarity property, then the max-min fair allocation \mathbf{x} on \mathcal{X} has all components equal: $\mathbf{x}_i = \mathbf{x}_j$ for all i, j , if the max-min fair allocation on \mathcal{X} exists.

B. Solidarity of The Feasible Rate Set of A Multi-hop Wireless Network

The feasible set of a wired network is given with a set of linear constraints. It is convex, but in general it does not have solidarity property, as can be seen on the right of Fig. 1. In the case of an ad-hoc wireless network, defined under the framework from Section III, we show that the feasible rate set of any such network without long-term average power constraints, has solidarity property.

Proposition 2: Any feasible rate set \mathcal{F} given by a set of equalities and inequalities (1), assuming $\bar{P}_l^{MAX} > P_l^{MAX}$ for all links l , has solidarity property. Also, a

feasible transport rate set given by (2) has the solidarity property.

C. Equality of Max-min Fair Rates

Consider an arbitrary network where long-term average power constraints are larger than transmission power constraints. It is easy to verify that the feasible set given by constraints (1) is convex, hence according to [20] it has the max-min fair allocation. Since this set also has solidarity property, we have the following:

Corollary 1: The max-min fair rate allocation of any network given by constraints (1), with no long-term average power constraints ($\bar{P}^{MAX} \geq P^{MAX}$), has all rates equal. The max-min fair transport rate allocation has all transport rates equal.

Equality of rates implies that all flows, including the most inefficient ones, have an equal rate. This can be very inefficient in a heterogeneous network. For example, if one node is almost disconnected, then it will receive a rate close to zero. According to corollary 1, all other flows will have the same rate.

Another example is given in Fig. 2. On the left, we show an example of a network where 12 nodes are randomly placed on a square 100m x 100m. The source and the destination of each flow are joined with a line. Each flow can use either the direct route or the minimum energy route. In this example, we set all transmission power constraints to be equal to $P^{MAX}/N = 90\text{dB}$, where N is a white background noise. The actual SNR on each receiver depends on the distance between the source of the link and the destination of the link. For example, according to the UWB indoor path loss model [10], if a source sends to a destination which is 10 m away with maximum power and $P^{MAX}/N = 90\text{dB}$, we have SNR at the receiver around 10 dB. This in turn leads to the rate of 3 Mb/s within the framework of [6].

On the right of Fig. 2, we see the optimal rate allocations with respect to the three metrics, for this example. We see that when maximizing total capacity, one flow has a high rate, and the rates of others are zero. In the case of max-min fairness, all rates are the same. Proportional fairness exhibits larger variation in rates than max-min fairness, but it does not starve the least efficient flows. But, it is more efficient than max-min fairness. We also illustrate the corollary 1 on more random examples on Fig. 3 in Section VI.

From corollary 1 we also see that in the case of the max-min fair transport rate allocation, all transport rates are equal. Obviously, the rates themselves are not equal as the flow lengths differ. Still in this case, as can be seen in the numerical examples from Fig. 3 in Section

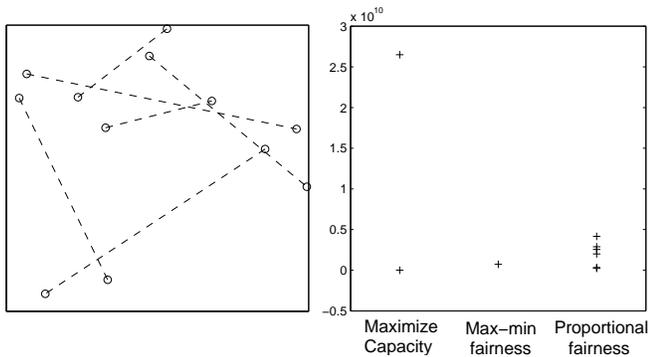


Fig. 2. Left: example of random network topology. 12 nodes are randomly placed on a 100m x 100m grid. Nodes are depicted with circles, and source-destination pairs are joined with lines. Each flow uses the optimal routing (the direct or the minimum-energy route). Right: corresponding rate distribution (total capacity, max-min fairness and proportionally fairness).

VI, the corresponding rate allocation suffers from the same inefficiency problem.

D. Influence of Long-Term Average Power Constraint

Corollary 1 holds when long-term average power constraints are greater or equal to transmission power constraints. When long-term average power constraints are smaller than the transmission power constraints, the max-min fair rates are not equal anymore. However, we see that for high transmission power constraints ($P^{MAX} \geq 40\text{dB}$, see Fig. 3 in Section V-B) and high long-term average power constraints ($\bar{P}^{MAX}/P^{MAX} \geq 0.5$, see Fig. 5, Section VII) the max-min fair rate allocation is still inefficient.

In proposition 5 in Section VII we show that for very small long-term average power constraints, the optimal allocation becomes independent of the choice of the metric.

Overall, these arguments show that max-min fairness is not an appropriate metric even when long-term average power constraints exist.

E. An Application to An 802.11 Network

An example of the above findings can be seen in [8]. Consider an 802.11 network where several nodes send data directly to a single destination (base-station). Assume node 1 is far away and it codes for 1Mb/s, and the others are close enough to codes for 11Mb/s. One would expect that node 1 achieves a smaller rate than other nodes. However, as shown in [8], this is not the case, and all nodes achieve an effective throughput of around 1Mb/s.

According to the analysis done in [8], when a node gets an access to the network, it sends a packet of a fixed

size, thus the occupancy time is inversely proportional to the coding rate. In other words, a node sends the same amount of data during a channel use, regardless of its coding rate. Let us consider a discrete random process X_t representing a user that occupies a channel during the t -th channel use. According to eq. (7) and (8) from [8], X_t is an i.i.d. uniform random process, and all nodes have an equal probability to get network access when the network is idle. This leads us straightforward to the following proposition:

Proposition 3: An 802.11 network in DCF mode where all nodes talk directly to a single destination (hence there is no hidden terminal problem), implements max-min fairness

In other words, the equality of rates observed in [8] is not solely a property of 802.11 physical layer but rather of max-min fairness that is obtained in this specific example. This means that any other protocol that would implement max-min fairness, and would fit in the framework of eq. (1), would have the same inefficiency problem.

F. When Max-min Fairness Does Not Lead To Equality

We note that the assumptions of corollary 1 is not true in general for any convex set, but only for those that have solidarity property. To illustrate this, we give a few counter examples:

- *Wired Networks:* The corollary does not hold for a class of wired networks. For an example, see Fig. 1 on the right.
- *Clustered Networks:* The corollary does not hold for a clustered wireless network. Assume a simple network of two links, link (1, 2) and link (3, 4), and assume it is clustered such that nodes 3 and 4 does not hear node 1 and 2 and vice versa (meaning that $h_{13} = h_{14} = h_{23} = h_{24} = 0$). Then rates f_{12} and f_{34} are not going to be equal.
- *Long-Term Average Power Constraint:* The corollary does not hold if long-term average power constraint is smaller than transmission power constraints, as shown in Section VII.

V. MAXIMIZING TOTAL CAPACITY

A. Asymptotic results

As discussed previously, maximizing total capacity metric is efficient but may lead to high unfairness, especially in the case of large transmission power constraints. In order to demonstrate this, we first look at the asymptotic case and we show that total capacity metric becomes totally unfair as transmission power tends to infinity.

At this point, we need an additional assumption on the rate function $\lim_{\text{SNR} \rightarrow \infty} r(\text{SNR}) = \infty$, that is we can increase the rate of a link arbitrarily high by sufficiently increasing the signal-to-noise ration on this link. We also assume here no long-term average power constraint, hence $\bar{P}^{MAX} \geq P^{MAX}$.

In order to simplify the presentation, we assume that all transmission power constraints are the same, that is for all l , $P^{MAX} = P_l^{MAX}$. This can be generalized for non-uniform power constraints, assuming that when P^{MAX} goes to infinity there exist fixed positive numbers γ_l such that $P_l^{MAX}/P^{MAX} > \gamma_l$.

Proposition 4: Assume that when the signal-to-noise ratio at a receiver SNR_i tends to infinity the rate of link i , $r(\text{SNR}_i)$ also tends to infinity, and assume $\bar{P}^{MAX} \geq P^{MAX}$. In a limiting case when $P^{MAX} \rightarrow \infty$, there will be one or more flows that have the same rate $f = O(r(P^{MAX}))$ and all the others will have a rate that is $o(f)$. The same happens when considering transport rates.

Proposition 4 tells us that if a signal-to-noise ratio is high enough, then only the most efficient flows are going to divide all the capacity of the medium, while all other flows will starve. In what follows we illustrate that the same problem occurs within the realistic signal-to-noise setting.

B. Numerical results

In the above section we have seen that an increase in transmission power constraint will eventually lead to all but some flows having zero rates. It is not clear what the realistic values of the constraints for which this phenomenon occurs are. From [6] we see the phenomenon has been observed in a realistic network, and in this section we investigate in which transmission power region it occurs.

In order to analyze the behavior of total capacity performance metric for a realistic power setting, we numerically evaluated it on random network topologies. We adapted the framework from [18], which assumes a rate is a linear function of the signal-to-noise ratio at a receiver (this also corresponds to an UWB model from [6]). As noted in [18], the optimization problem has exponential complexity so it was not possible to run simulations for more than 12 nodes. We generated 150 random network topologies with 12 nodes uniformly distributed on a square of 100m x 100m. Half of them are sources sending data each to its own destination among the other half. All nodes are assumed to have the same transmission power constraints. We are looking for a routing, scheduling, and power control that maximizes

the total capacity. An example of such a network described above, and the optimal end-to-end rates with respect to different objectives can be seen on Fig. 2.

For each flow we consider a multi-path routing with a set of routes that comprise nodes that are on the shortest path between the source and the destination. This is a suboptimal set of routes since in the case of high congestion in one area of a network, the optimal path may avoid that area even if it is not the shortest one. However, in most cases this heuristic is a good approximation, and it simplifies our calculation. Furthermore, running tests on several random topologies, we concluded that in all cases the optimal route among those is either the minimum energy route (relaying over intermediate nodes that minimizes total dissipated power), or the direct route (send directly to the destination without relaying). Since constraining on these two routes for each flow further reduces the complexity of optimization, we used this heuristics to produce the results.

In Fig. 3, on the top left, we show average fairness indices of the optimal rates with respect to total capacity and proportionally fair metrics, as well as the confidence intervals. On the x-axis, a ratio between maximal transmitting power and noise in dB is given.

From the numerical results depicted in Fig. 3, on the top left, we see that maximizing total capacity leads to an acceptable fairness in the case of small transmission power limits. However, for large transmission power limits we see that maximizing total capacity exhibits high unfairness, which leads to only one flow having non-zero rate, as predicted by proposition 4. These results confirm unfairness observations made in [6], and show they are a consequence of the performance metrics rather than UWB protocol particularities. All these results are for unlimited battery lifetime constraints. However, the unfairness exists for limited battery lifetime; for details, see Section VII.

Next, we used the three metrics to find the optimal solutions on the set of transport rates \mathcal{T} . We then calculated the transport fairness and the transport efficiency indices of the corresponding rate allocations. This can be seen in Fig. 3, middle. We see that for small powers maximizing transport capacity is a bit fairer or equally fair as proportional fairness. For high powers and more realistic rates, it becomes significantly less fair than proportional fairness, as suggested by proposition 4.

We also analyzed the fairness index of the optimal rates in a case of random non-uniform networks. We again considered a square area 100m x 100m, and we divided it into 4 equal sub-squares 50m x 50m each. We placed 12 uniformly distributed nodes in total in upper left and lower right sub-squares. Each node chose

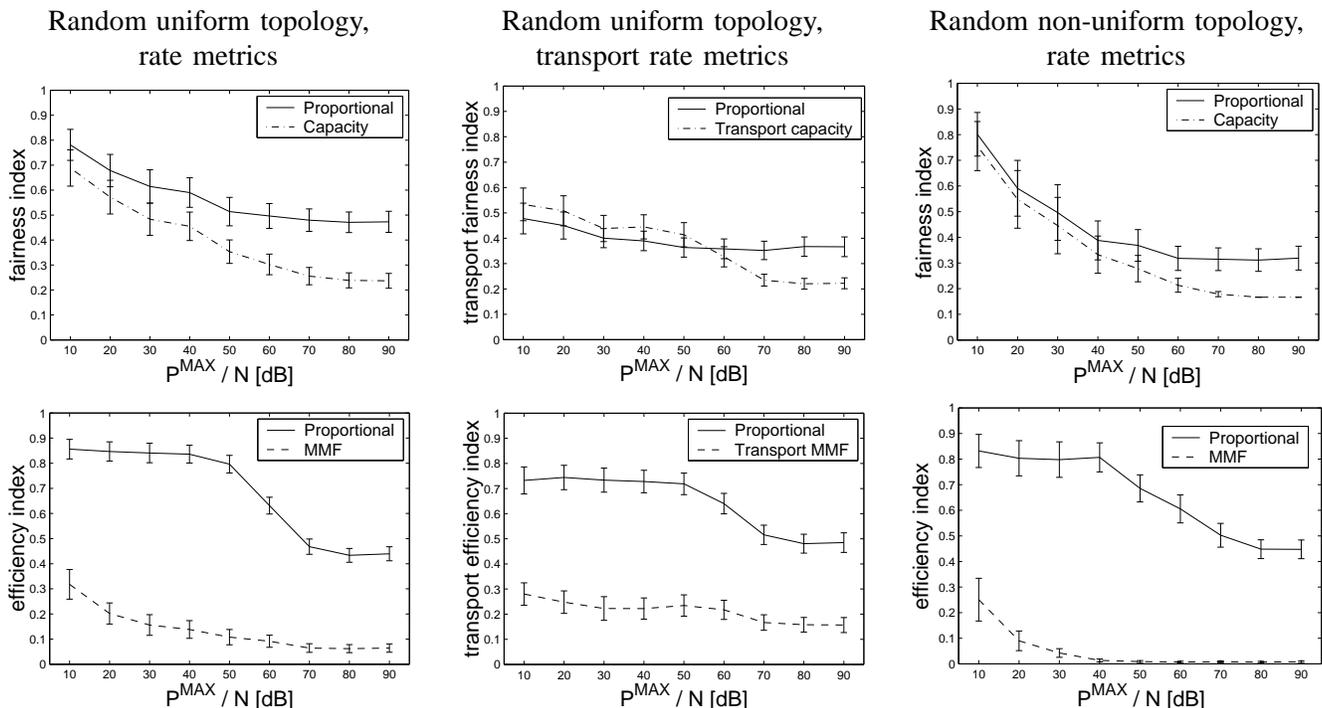


Fig. 3. The fairness and efficiency indices of optimal rates with respect to different performance objectives, versus the ratio between maximal transmitting power and noise: On the top, the fairness index is given; on the bottom is the efficiency one. On the left, we see the indices for the three performance objectives applied to the set of feasible rates. In the middle, we see the indices for the three performance objectives applied to the set of feasible transport rates. In both cases nodes are uniformly spread over the entire 100m x 100m square. Values of P^{MAX}/N on the x axis represent realistic values that can be found on existing UWB or CDMA systems. On the right we see the indices for the three performance objectives applied to the set of feasible rates in the case when nodes are distributed only in the upper-left and lower-right quarters of the square. In all cases we put no constraints on long-term average power. All figures show 95% confidence intervals.

uniformly one destination among all other nodes. We thus had several short and several long flows, and a hot-spot in the center of the big square. The results are depicted in Fig. 3 on the top right.

VI. PROPORTIONAL FAIRNESS

As has been seen in the previous sections (e.g. Fig. 2), both maximizing total capacity and max-min fairness suffer from either inefficiency or unfairness. In this section we analyze in detail proportional fairness and we show that it represents a robust compromise between efficiency and fairness. We numerically evaluated the efficiency and fairness of proportional fairness metric using the same setting as in Section V-B.

It is shown in [18] that an optimal power allocation strategy for maximizing total capacity is either to send with maximal power or not to send at all. It is not clear if the same strategy is optimal for proportional fairness. Optimization over instantaneous powers is a non-convex optimization, hence a difficult problem [18], [23], [11]. We solve this problem for random topologies with a small number of nodes and show that in all cases the strategy from [18] is nearly optimal. Therefore, we use

it as a heuristic when calculating the proportionally fair rate allocation. We also use the same routing heuristic as in the above case.

The fairness index of proportionally fair rate allocation is depicted in Fig. 3 on the top left. It can be seen that it is robust and remains constant for all values of transmission power constraint. On Fig. 3 on the bottom left, we depict the efficiency index, which is a ratio between the total capacity of the optimal allocation under given metric and the maximal total capacity that can be achieved in a given network (when maximizing total capacity metric). We see that the proportionally fair rate allocation remains up to 10 times more efficient than the max-min fair allocation.

We analyzed the efficiency index of the optimal rates in a case of random non-uniform networks, as above, and the results are depicted in Fig. 3 on the bottom right. The fairness index is given on Fig. 3 on the top right. The results are similar to those from the symmetric case, and the same conclusions hold.

In the middle of Fig. 3 we depict the fairness and efficiency properties of the optimal transport rates on set \mathcal{T} . Transport max-min fairness is again much less

efficient than proportional fairness. Maximizing transport capacity is fairer comparing to proportional fairness on set \mathcal{T} than maximizing total capacity on set \mathcal{X} . For small powers, it is even fairer. For high powers and more realistic rates, it becomes almost twice less fair.

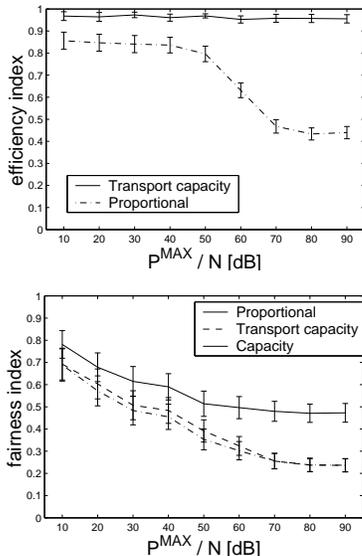


Fig. 4. Efficiency index (top) and fairness index (bottom) of the rates that maximizes total capacity and transport capacity, and the proportionally fair rates. Results are obtained on random uniform network topologies, taking $P^{MAX}/N = 90$ dB.

This finding poses the question if maximizing transport capacity might in some cases be an appropriate metric with respect to the rate efficiency and the rate fairness indices. In other words can maximizing transport capacity reconcile the rate unfairness of total capacity objective? According to proposition 4, maximizing transport capacity also exhibits high unfairness for large transmission power constraints. We give numerical examples for realistic transmission power constraints on random uniform network topologies in Fig. 4. We see that the rate that maximizes transport capacity is only marginally more fair and marginally less efficient than the one that maximizes total capacity. The unfairness becomes the same in both cases for high powers, as suggested by proposition 4. Again, proportional fairness represents a much better compromise between efficiency and fairness than the total capacity based metrics.

VII. INFLUENCE OF LONG-TERM AVERAGE POWER CONSTRAINT

In the previous sections we have seen that when we do not put constraints on battery lifetime, hence we have no long-term average power constraint, in the case of max-min fairness this leads to the equal rates of all flows and in the case of maximizing total capacity to zero rates

of some flows. On the other hand, in the limiting case when \bar{P}^{MAX} is small enough, the optimal allocation is the same regardless of the choice of the performance metric.

Proposition 5: Suppose the rate function $r(\text{SNR})$ is a concave function and $r(0) = 0$. Then for every network of there exist some $\epsilon > 0$ such that for all $\bar{P}^{MAX} \leq \epsilon$ optimal rates for maximizing total capacity, max-min fairness and proportional fairness objectives are the same.

Properties of the three metrics for more realistic long-term average power constraints can be seen on the Fig. 5.

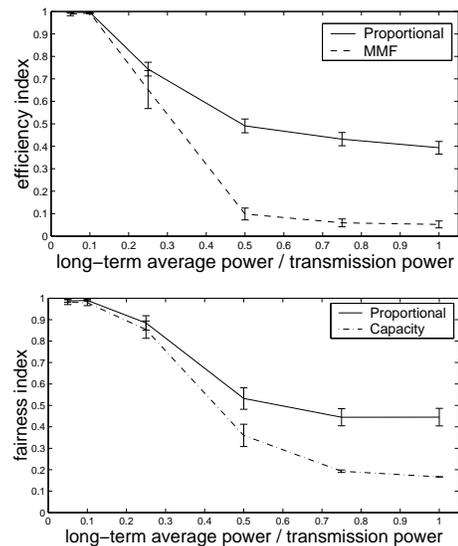


Fig. 5. Efficiency index (top) and fairness index (bottom) of max-min and proportional fairness for finite long-term average power constraint.

As can be seen, the unfairness of total capacity and inefficiency of max-min fairness are visible. When we further decrease long-term average power constraints, maximizing total capacity, max-min and proportional fairness become the same, as suggested by proposition 5.

VIII. CONCLUSION

We analyzed three rate-based performance metrics: total capacity, max-min fairness and proportional fairness, within the framework of ad-hoc wireless networks. We defined a general model of such a network, which incorporates all existing physical layers (CDMA, UWB, variable rate 802.11, etc.), and allows for arbitrary scheduling, routing and power control policy. We then evaluated the three metrics on this model.

We found that max-min fairness yields equal rates to all flows, when users are not implying battery lifetime constraints. In a heterogeneous network, this means

that the rates of all flows are equal to the rate of the smallest flow, which makes a network very inefficient. This confirmed and generalized the findings from [8]. In presence of long-term average power constraints, the max-min fair rate does not necessarily have this property but the inefficiency persists. Also the rate allocation of all flows depends on long-term average power constraints of a single user, which is an undesirable property of a performance objective.

We proved that for large enough power constraints maximizing total capacity gives zero rates to all but the most efficient flows. We showed that this type of unfairness occurs on most of the networks for realistic power constraints. This is a confirmation and a generalization of the findings from [6]. Like in the case of max-min fairness, this phenomenon is somewhat remedied in the case of small long-term average power constraint, but remains. We also showed that in the case of small long-term average power constraints, the optimal rate allocation depends more on these constraints than on the choice of the performance metrics.

Finally, we analyzed the proportionally fair rate allocation on a large number of arbitrary networks with variable transmission power and long-term average power constraints. We found that in all cases it maintains fairness while it achieves relatively high efficiency. We also find it robust with respect to changes in topology and power constraints. These properties make it the optimal performance metric when evaluating or designing a MAC or a routing protocol for an ad-hoc wireless network.

All the metrics analyzed in this paper are rate-based performance metrics. The power constraints were considered explicitly rather than through performance metrics. Still, powers can be incorporated in all three types of metrics analyzed here. A future work would be to analyze what is the ideal power-based and combined performance metric for an ad-hoc wireless network.

REFERENCES

- [1] A. Charny. An algorithm for rate allocation in a packet-switched network with feedback. *M.S. thesis*, MIT, May 1994.
- [2] A. Mas-Colell, M. Whinston, J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [3] Baldi, De Nardis, Di Benedetto. "Modeling and Optimization of UWB Communication Networks Through a Flexible Cost Function". *IEEE Journal on Selected Areas in Communications*, 20(9):1733–44, December 2002.
- [4] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, 1987.
- [5] D. Tse, S. Hanly. "Multi-access Fading Channels - part I - Polymatroid Structure, Optimal Resource Allocation and Throughput Capacities". *IEEE Transactions on Information Theory*, 44(7):2796–2815, November 1998.
- [6] F. Cuomo, et al. "Radio Resource Sharing for Ad Hoc Networking With UWB". *IEEE Journal on Selected Areas in Communications*, 20(9):1722–1732, December 2002.
- [7] F. P. Kelly, A.K. Maulloo and D.K.H. Tan. "Rate control in communication networks: shadow prices, proportional fairness and stability". *Journal of the Operational Research Society*, 49:237–252, 1998.
- [8] G. Berger-Sabbatel, F. Rousseau, M. Heusse, A. Duda. "Performance anomaly of 802.11b". In *INFOCOM'03*, San Francisco, California, April 2003.
- [9] Holland, Vaidya, Bahl. A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks, *mobicom01*. In *MOBIHOC'01*, 2001.
- [10] IEEE P802.15 Working group. "The UWB Indoor Path Loss Model". *Technical report P802.15-02/277r0-SG3a*, June 2002.
- [11] Julian, Chiang, O'Neill, Boyd. "QoS and Fairness Constrained Convex Optimization of Resource Allocation for Wireless Cellular and Ad Hoc Network". In *INFOCOM*, 2002.
- [12] L. Tassiulas, A. Ephremides. "Jointly optimal routing and scheduling in packet radio networks". *IEEE Transactions on Information Theory*, 38(1):165–168, January 1992.
- [13] L. Tassiulas, S. Sarkar. "Max-min Fair Scheduling in Wireless Networks". In *Proceedings INFOCOM'02*, 2002.
- [14] Le Boudec, J.-Y. "A Tutorial on Rate Adaptation, Congestion Control and Fairness in The Internet". *Technical Report*, http://ica1www.epfl.ch/PS_files/LEB3132.pdf.
- [15] Nandagopal, Kim, Gao, Bharghavan. "Achieving MAC Layer Fairness in Wireless Packet Network". In *MOBICOM*, 2000.
- [16] P. Bender, et al. "CDMA/HDR: A Bandwidth-Efficient High-Speed Wireless Data Service for Nomadic Users". *IEEE Communications Magazine*, pages 70–77, July 2000.
- [17] P. Gupta and P.R. Kumar. "The Capacity of Wireless Networks". *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000.
- [18] R. Cruz, A.V. Santhanam. "Optimal Link Scheduling and Power Control in CDMA Multihop Wireless Networks". In *Globecom'02*, 2002.
- [19] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley-Interscience, 1991.
- [20] Radunović, B., Le Boudec, J.-Y. "A Unified Framework for Max-Min and Min-Max Fairness with Applications". In *Proceedings Allerton'02*, 2002.
- [21] Radunović, B., Le Boudec, J.-Y. "Rate Performance Objectives of Multi-hop Wireless Networks". *Technical Report IC/2003/45*, http://lcawww.epfl.ch/Publications/Radunovic/TR03_045.ps, 2003.
- [22] S. Sarkar and L. Tassiulas. "Fair Allocation of Discrete Bandwidth Layers in Multicast Networks". In *INFOCOM'00*, pages 1491–1500, 2000.
- [23] S. Toumpis and A.J. Goldsmith. "Capacity regions for wireless ad hoc networks". *IEEE Transactions on Wireless Communications*, to appear.
- [24] T. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
- [25] Win M.Z., Scholtz R.A. "Ultra-wide bandwidth time-hopping spread-spectrum impulse radio for wireless multiple-access communications". *IEEE Transactions on Communications*, 48(4):679–689, April 2000.
- [26] 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.

APPENDIX

A. Proof of proposition 1

Let us denote by \mathbf{x} a max-min fair allocation on \mathcal{X} and let us assume the contrary, that there exists i and j such that $\mathbf{x}_i - \mathbf{x}_j > 2\epsilon$ for some $\epsilon > 0$. Then, according to the solidarity property, there exists \mathbf{y} such that $\mathbf{x}_i \geq \mathbf{y}_i > \mathbf{x}_i - \epsilon > \mathbf{x}_j + \epsilon > \mathbf{y}_j > \mathbf{x}_j$, and $\mathbf{y}_k = \mathbf{x}_k$ for all $k \neq i, k \neq j$ which contradicts with the definition of max-min fairness. ■

B. Proof of proposition 2

We proceed by contradiction. Consider a feasible rate \mathbf{f} that does not satisfy solidarity property for some coordinates i and j . We also assume \mathbf{f} is Pareto optimal, since if it is not, the contradiction with the definition of solidarity property is straightforward. Let us denote by $(\mathbf{y}, \bar{\mathbf{x}}, \alpha, (\mathbf{p}^i)_{i=1 \dots n+1})$ the values of slack variables, used in the constraint set given by eq. (1), that satisfies these constraints for rates \mathbf{f} .

Since \mathbf{f} is Pareto optimal, there exist a set of rows $K(i)$ in matrix R such that for each $k \in K(i)$, there exists path m that belongs to flow i (i.e. $F_{mi} = 1$), path m passes over link k (i.e. $R_{km} = 1$), and has strict equality $\bar{\mathbf{x}}_k = (R\mathbf{y})_k$ (else, we can increase \mathbf{f}_i at no cost). The same holds for $K(j)$. Let us find link $k \in K(i)$ such that path m passing over k has a positive rate $\mathbf{y}_m > 0$. If there exists path n belonging to flow j (i.e. $F_{nj} = 1$), such that $R_{kn} = 1$ (paths n and m have a common link k), then for any $\epsilon < \mathbf{y}_m$ we can construct \mathbf{f}' such that $\mathbf{f}'_i = \mathbf{f}_i - \epsilon$ by decreasing \mathbf{y}_m , $\mathbf{f}'_j = \mathbf{f}_j + \epsilon$ by increasing \mathbf{y}_n , and $\mathbf{f}'_k = \mathbf{f}_k$ for all $k \neq i, k \neq j$; we have $\mathbf{f}' \in \mathcal{F}$ leading to contradiction. We thus have $K(j) \cap K(i) = \emptyset$.

We pick link $l \in K(j)$, $l \neq k$, a slot s when link k is active, and divide it in two slots, s_1 and s_2 of lengths $\alpha_{s_1} > 0$ and $\alpha_{s_2} = \alpha_s - \alpha_{s_1}$ respectively. In the first slot we keep the same scheduling as in slot s , and in the second slot we turn off link k and increase the power of link l such that $p_l^{s_2} \leq P_l^{MAX}$ and the interferences perceived by other active users is smaller than in the original scheduling of slot s (note that $p_l^{s_2} > 0$). This is always possible since $\bar{P}_l^{MAX} \geq P_l^{MAX}$ for all l . With this new scheduling all links have the same or higher rates, except for link k whose rate has decreased by $\epsilon^k(\alpha_{s_2}) \geq 0$ and link l whose rate has increased by $\epsilon^l(\alpha_{s_2}) \geq 0$, with equalities for $\alpha_{s_2} = 0$. Both ϵ^k and ϵ^l are linear functions of α_{s_2} hence there exists ϵ small enough such that for the new slot length $\alpha_{s_2} < \alpha_s$ we have $\epsilon^k(\alpha_{s_2}) < \epsilon$ and $\epsilon^l(\alpha_{s_2}) < \epsilon$.

We thus have a new average link rate allocation $\bar{\mathbf{x}}'$ such that $\bar{\mathbf{x}}_k - \epsilon < \bar{\mathbf{x}}'_k < \bar{\mathbf{x}}_k$ and $\bar{\mathbf{x}}_l < \bar{\mathbf{x}}'_l < \bar{\mathbf{x}}_l + \epsilon$. Now we can increase \mathbf{f}_j increasing some \mathbf{y}_n passing over link

l by some positive $\alpha_j = \bar{\mathbf{x}}'_l - \bar{\mathbf{x}}_l < \epsilon$, by decreasing \mathbf{f}_i ; decreasing some \mathbf{y}_m passing over link k by some positive $\alpha_i = \bar{\mathbf{x}}'_k - \bar{\mathbf{x}}_k < \epsilon$. This is exactly a solidarity property, hence the contradiction.

The same reasoning holds for a set of transport rate, hence the second part of the statement. ■

C. Proof of proposition 4

We first propose a lemma that characterize the optimal schedule and power allocation when transmission power limit tends to infinity.

Lemma 1: Let \mathbf{p}^s be the optimal power allocation in slot s given transmission power limit P^{MAX} . For all slots s there exists link i such that both are true:

- 1) There exists $\epsilon_i > 0$ such that for all Ω_i there exists $P^{MAX} > \Omega_i$ such that $\mathbf{p}_i^s / P^{MAX} > \epsilon$.
- 2) For all $j \neq i$ and for all $\epsilon_j > 0$ there exists Ω_j such that for all $P^{MAX} > \Omega_j$ we have $\mathbf{p}_j^s / \mathbf{p}_i^s < \epsilon$.

Proof: We begin by showing that first statement is true using contradiction. Suppose that for some slot s and for each link i and all $\epsilon_i > 0$ there exists Ω_i such that for some $P^{MAX} > \Omega_i$ we have $\mathbf{p}_i^s / P^{MAX} < \epsilon$. Let us choose an arbitrary link j and increase its power allocation in slot s to $\mathbf{p}_i^s = P^{MAX}$. We then have the following

$$\frac{\text{SNR}_j^s}{\text{SNR}_i^s} = \frac{P^{MAX} h_{jj} N + \sum_{k \neq i} \mathbf{p}_k^s h_{ki}}{\mathbf{p}_i^s h_{ii} N + \sum_{k \neq j} \mathbf{p}_k^s h_{kj}} \quad (3)$$

$$= \frac{P^{MAX}}{\mathbf{p}_i^s} K > \epsilon_i K. \quad (4)$$

where K is a fixed constant. Therefore, we can make new SNR_j^s arbitrary higher than any signal-to-noise ratio in slot s . Due to the assumption on the rate function, the same way we can make a rate of link j in slot s arbitrary larger than rates of other links in slot s , as well as the sum of rates of all links in slot s . In particular, if link j connects a source and a destination of a flow, by increasing \mathbf{p}_j^s to P^{MAX} we increased the total rate, which contradicts with the initial assumption.

Next, we show the second part of the statement, again by contradiction. We suppose that in some slot s there exists link j such that for some ϵ_j and for all Ω_j there exists P^{MAX} such that $\mathbf{p}_j^s / \mathbf{p}_i^s > \epsilon_j$. Again, we consider a new power allocation where $\mathbf{p}_i^s = P^{MAX}$ and all the other powers are zero. We have the following

$$\begin{aligned} \frac{\text{SNR}_l^s}{\text{SNR}_i^s} &= \frac{P^{MAX} h_{ll} N + \mathbf{p}_j^s h_{ji} + \sum_{k \neq i, j} \mathbf{p}_k^s h_{ki}}{\mathbf{p}_i^s h_{ii} N} \\ &> \frac{P^{MAX} \mathbf{p}_j^s}{\mathbf{p}_i^s} K > P^{MAX} \epsilon_j K. \end{aligned}$$

Here K and ϵ_j are fixed constants and for an arbitrary Ω_j there exists $P^{MAX} > \Omega_j$ that satisfies the above inequality. This in turn means that we can make SNR_l^s arbitrary larger than SNR_i^s . The same applies for SNR_j^s . We can do similarly for a link $k \neq i, k \neq j$ by virtue of (4). As we shown above, if l is a link between a source and a destination of a flow, the new allocation increases total rate which contradicts with the initial assumption. ■

Proof of proposition 4: Consider a link i . From (1) we have the following inequality $\sum_{p \ni i} y_p \leq \sum_s \alpha_s \mathbf{x}_i^s$. By lemma 1 we know that in the optimal power allocation, in each slot there is exactly one link whose power is $O(r(P^{MAX}))$ and all other links have powers $o(r(P^{MAX}))$. Therefore, we can assign all time to the power allocation achieving the highest rate $\sum_{p \ni i} y_p \leq (\sum_s \alpha_s) \max_t \mathbf{x}_i^t$. We might assume equality, since we otherwise can assign all extra time to other power allocations. Also, we can divide the new slot into sub-slots, each serving only one path, hence we can write $y_p = \alpha_{s(i,p)} \mathbf{x}_i^{s(i,p)}$.

Suppose we have an additional time $\Delta\alpha$ to serve path y_p . We need to spread it on all links belonging to path y_p such that each link i gets $\alpha_{s(i,p)} / \sum_{j \in p} \alpha_{s(j,p)}$ fraction of it, and the overall increase in rate of p is $\Delta\alpha \left(\alpha_{s(i,p)} \mathbf{x}_i^s / \sum_{j \in p} \alpha_{s(j,p)} \right)$. Now, since the total capacity is a sum of the rates on all paths, in order to maximize total capacity we will assign time only to links of those paths that have the highest increase factor, and will not serve the other paths letting them have zero rate.

The same happens in the case of transport rates, since increase factor is the same as above, multiplied by a length of the corresponding flow. Consequently, the corresponding rates will also tend to zero. ■

D. Proof of proposition 5

Consider a case when only one link is sending data. It can use any power $p \in [0, P^{MAX}]$, and a fraction of time it is active α has to be $\alpha p \leq \bar{P}^{MAX}$. In order to maximize rate, we have $\alpha = \bar{P}^{MAX} / p$ and the average rate is $\bar{x}(p) = r(p/N) \bar{P}^{MAX} / p$. Since rate is a concave function of SNR, we have for all x $r(p) \leq r'(p) * p$, and the $\bar{x}(p)$ is maximum when $p = P^{MAX}$. Therefore, if a node alone is sending, it should use maximum power for its transmissions.

Suppose there are L links in the network. When $\bar{P}_l^{MAX} \leq P_l^{MAX} / L$ for all l , then a node sending at the maximum power cannot send for more than $1/L$ fraction of time. It is thus optimal, regardless of a performance metric used, that only one node sends at a time (due to zero interference) at the maximum power,

hence the optimal rate allocation is the same for all three performance metrics. ■