Why Max-min Fairness Is Not Suitable For 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 issue of which criteria to use when evaluating the design of a wireless multihop network. It is known, and we illustrate in this paper, that maximizing the total capacity, or transport capacity, leads to gross imbalance and is not suitable. An alternative, which is often used in networking, is to consider the max-min fair allocation of rates, or of transport rates per node. We apply max-min fairness to the class of wireless, multi-hop networks for which the rate of a wireless link is an increasing functions of signal-to-noise ratio. This class includes CDMA and UWB. We show that, for a network in this class, the max-min fair allocation of bit or transport rates always gives the same rate to all flows. We show on one example that such an allocation is highly undesirable when the network is asymmetric. Another form of fairness, utility fairness, does not appear to have the same problem.

Index Terms— wireless, max-min, utility fairness, best-effort

I. INTRODUCTION

In many works concerning the design of a wireless network, the goal is to maximize the total throughput of the network (e.g. [8], [5]). Now it is known, and we illustrate later in this paper, that considering total throughput (or total capacity) as a performance measure in a network with best-effort traffic leads to gross imbalance (by shutting down more expensive users). Another commonly used metric in wireless network is the transport capacity, defined in [7] as the sum of bits $times$ distances over which they are carried per second. We show that this metric suffers from the same imbalance problem as total capacity.

A classical solution to this problem it to evaluate a network design under the assumption that it provides max-min fairness [3]. A rate allocation is said to be max-min fair if one cannot increase a rate of one flow without decreasing an already smaller rate. Max-min fairness is the target criterion used in ATM standards for allocating rates in a best effort mode. Max-min fairness is also widely used in wireless networking [6], [10]. An alternative approach is utility fairness, originating in principles of economy. Each user is assigned a concave and increasing function of its rate, called utility, and the system maximizes the sum of utilities of all users. The most used form of utility fairness is proportional fairness, [4], where the utility is the log function.

In this paper we focus on max-min fairness. We consider the class of wireless network technologies for which the rate

of a wireless link is an increasing functions of signal-to-noise ratio. This class includes CDMA and UWB. Our main finding is that, for a multi-hop network in this class, the max-min fair allocation of bit or transport rates always gives the same rate to all flows, whatever the routing policy. We show on one example that such an allocation is highly undesirable when the network is asymmetric. Utility fairness, does not appear to have the same problem.

II. DEFINITIONS AND NOTATIONS

Consider a network with n flows, identified with source-destination pairs $\{(s_1, d_1)...(s_n, d_n)\}$. There are m possible wireless links and l paths. For a flow i, let f_i be its rate and $t_i = f_i \|s_i - d_i\|$ its "transport rate" ($\|s_i - d_i\|$ is the line of sight distance from source to destination). The vector $\mathbf{f}$ of all rates is feasible if there exists a power allocation, scheduling and allocation of rate to paths that achieve it. We call $\mathcal{F}$ the set of feasible rate vectors and $\mathcal{T}$ the set of all feasible transport rate vectors (defined precisely below). The total rate of the network is $\sum_{i=1}^{n} f_i$ and $\max_{\mathbf{f} \in \mathcal{F}} \sum_{i=1}^{n} f_i$ is the network capacity; the total transport rate of the network is $\sum_{i=1}^{n} t_i$ and $\max_{\mathbf{t} \in \mathcal{T}} \sum_{i=1}^{n} t_i$ is the network transport capacity.

A rate vector $\mathbf{f} \in \mathcal{F}$ is said to be weighted max-min fair [3] with weights w_i if for any other vector $\mathbf{f}' \in \mathcal{F}$ such that for some i, $f_i'/w_{i'} > f_i/w_i$, there exists j such that $f_j'/w_j < f_j/w_j \leq f_i/w_i$. We use an analog definition for a max-min fair transport rate allocation over set $\mathcal{T}$. A weighted max-min fair allocation, if it exists is unique. If $w_i = 1$ for all i we simply say that the rate vector is max-min fair. Finally, a utility fair rate (resp. transport rate) allocation is the one that maximizes $\sum_{i=1}^{n} u(f_i)$ over $\mathcal{F}$ (resp. over $\mathcal{T}$) for a given concave utility function u.

The network model is as follows. At the receiver of link j, $j = 1...m$, the useful power is p_j and $I_j(\mathbf{p})$ is the total power of all interfering signals plus background noise. We assume that the useful rate on link j is $x_j = r(p_j/I_j)$ where $r(x)$ is an arbitrary, but fixed increasing function. This defines the class of wireless networks that allow interference, such as ultra-wide band ($r(x)$ is a linear function), CDMA, or the theoretical case where the link rate if the Shannon capacity of a point to point Gaussian interference channels ($r(x) = \log(1 + x)$).

We assume that the power received on link j is limited to P_j^{max} and let $\mathcal{P} = [0, P_1^{max}] \times \cdots \times [0, P_n^{max}]$. The set $\mathcal{X}$ of achievable link rates $\mathbf{x}$ is the convex closure $\{\mathbf{r}(\mathbf{p}) \mid \mathbf{p} \in \mathcal{P}\}$ and any $\mathbf{x} \in \mathcal{X}$ can be written as $\mathbf{x} = \sum_{i=1}^{n+1} \alpha_i \mathbf{r}(\mathbf{p}^i)$, for some $\mathbf{p}^i \in \mathcal{P}$, and $\sum_{i=1}^{n+1} \alpha_i = 1$. We call α a scheduling and vectors $\mathbf{p}^i$ power allocation policies that achieves $\mathbf{x}$.

We put no restriction on routing, and a flow i may be routed over several parallel paths. More precisely, an arbitrary multi-path routing policy can be expressed with a set of linear constraints $\mathbf{f} = A\mathbf{y}$, $\mathbf{x} \geq B\mathbf{y}$, where $\mathbf{y}$ is the vector of rates on paths, $a_{ik} = 1$ if path k belongs to a flow i and $b_{jk} = 1$ if path k traverses link j. Matrices A and B define the routing policy. If they are such that the set of paths is all possible paths between all sources and destinations, then we have a network model with unconstrained routing. At the other end of the spectrum, we find the more traditional single path, multihop routing model. All of these fit in our framework. Given A and B, the set of feasible rates $\mathcal{F}$ is defined as the set of $\mathbf{f}$ for which there exists $(\mathbf{x}, \alpha, \mathbf{p})$ that satisfy all the constraints enumerated in this section.

III. Equality of Max-min Fair Rates

Proposition 1: For the class of wireless networks defined above, there exist a unique max-min fair allocation of rates [resp. of transport rates], and it gives equal rates [resp. transport rates] to all flows.

The proof is in appendix. If $\mathbf{f}$ is the max-min fair rate allocation, then the corresponding vector of transport rates $\mathbf{t}$ is weighted max-min fair with weights $\|s_i - d_i\|$. There is an analog statement for transport rates (with inverse weights).

The proposition implies that max-min fair allocation equalizes rates to the value of the worst link. We illustrate this numerically in the next section.

IV. Application to an Asymmetric Example

Finding an allocation that satisfies max-min fairness, utility fairness, or maximization of total rate for a general network is a computationally difficult problem [9]. In order to illustrate our proposition, we consider an asymmetric, but regular network that can easily be analyzed.

There are $n + 2$ nodes, as shown on (Fig. 1). Node r_i talks to its right-hand neighbor r_{i+1} (node r_{n-1} talks to r_0), and node d_1 talks to node d_2. All nodes use direct links to talk to their peers. This implicitly defines the matrices A and B. We compute the rates and transport rates that satisfy the following criterion : max-min fairness, utility fairness, and maximization of sum of rates. The computation is based on the following


Fig. 1. The network example used for a numerical application. Nodes $r_0, \cdots, r_{n-1}$ form a ring of equally spaced nodes, node d_2 is located in the center of the ring, and node d_1 is at a distance z from d_2 on the direction orthogonal to the plane the ring.

proposition. Let us $\mathcal{R}_i(\mathbf{p})$ be the permutation of the power

allocation $\mathbf{p}$ that corresponds to the rotation with axis d_1d_2 that maps r_0 to r_1 axis, i.e. $(\mathcal{R}_i(\mathbf{p}))_{r_j} = \mathbf{p}_{r_{(j+i) \bmod n}}$, $(\mathcal{R}_i(\mathbf{p}))_{d_j} = \mathbf{p}_{d_j}$.

Proposition 2: For any of the criteria: max-min fair, utility fair and maximization of sum of rates, there exists a scheduling and power allocation that maximizes the metric, has all rates on the ring equal, and has rotationally symmetric time slots in the form

$$\mathbf{x} = \sum_{i=1}^{3} \alpha_i \sum_{j=1}^{n} \mathbf{R}(\mathcal{R}_j(\mathbf{p}^i)). \qquad (1)$$

The proof is in the appendix. We numerically illustrate these findings for a point-to-point Gaussian channel model of each link, which corresponds to the $r()$ function $x = \log(1 + SNR)$. The attenuation of signal between source i and destination j is $h_{ij} = \|i - j\|^{-4}$, $n = 8$ and $P_{max}/N = 10^2$. Utility fairness is proportional fairness. The results are on shown on Figs. 2 and 3. The figures illustrate that the max-min fair allocation (of rates or transport rates) equalize all sources to the worst case. The rates tend to 0 when the distance z goes to ∞ – clearly an undesirable behaviour. The maximization of total rate (or transport rate) does exactly the opposite: it shuts down either the sources on the ring, or off-ring, depending on the distance z. The off-ring source is shut down when z is large (for maximization of total rate) or when z is either very small or large (for maximization of sum of transport rate). This is equally undesirable (gross unfairness towards links with large attenuation). The proportionally fair allocation does not appear to have this problem. On-ring sources get a rate of the same order of magnitude regardless of z, whereas the off-ring rate x_d still tends to zero when z goes to infinity, thus protecting the performance of the network from links of very bad quality.[1]

V. Conclusion

We have shown that for a very general model of multi-hop wireless networks, the rates of all flows in max-min fair allocation with optimal scheduling and power allocation have to be equal, regardless of node positions, traffic matrix and routing policy. We show that this is unacceptable for a largely asymmetric networks. We also show that a performance metric that maximizes the sum of rates is also not acceptable. In contrast, utility fairness does not appear to have such problems. The same holds for transport rates.

References

[1] B. Radunović, J.Y. Le Boudec. "A Unified Framework for Max-Min and Min-Max Fairness with Applications". In *Proceedings Allerton'02*, 2002.
[2] B. Radunović, J.Y. Le Boudec. "Joint Scheduling, Power Control and Routing in Symmetric, One-dimensional, Multi-hop Wireless Networks". In *WiOpt*, INRIA, Nice, France, March 2003.
[3] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, 1987.
[4] 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.

[1]Note that with maximization of sum of rates, the rate allocation constructed in proposition 2 is not the unique maximizer. But the other allocations have $x_d = 0$ when $\max_{0 \leq i \leq n-1} x_{r_i} > x_d$, and the effect of unfairness is even larger.

Fig. 2. Rates achieved for max-min fairness (left), maximization of sum of rates (center) and utility (proportional) fairness (right). For max-min fairness, the on-ring rates x_r are the same as off-ring rate (this follows from proposition 1). For maximization of sum of rates, either the on-ring rates are $x_r = 0$ and the off-ring rate x_d is > 0, or vice-versa, depending on the the size z of link (d_1, d_2). For proportional fairness, the on-ring rates x_r are always positive, and the off-ring rate x_d goes to zero when $z \to \infty$.


Fig. 3. Same as Fig 2 but for transport rates. The main difference is for maximization of transport rates: the on-ring transport rates x_r are 0 and the off-ring transport rate x_d is > 0 when the size z of link (d_1, d_2) is nor too large nor too small, and vice-versa otherwise.

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

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

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

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

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

[10] 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

Proof of Proposition 1: The feasible set $\mathcal{F}$ is convex, which by [1] shows the existence. Let $(\mathbf{f}, \mathbf{x}, \alpha, (\mathbf{p}^i)_{i=1\cdots n+1})$ be such that max-min fair rate allocation is achieved. It follows from [1] that the set of rates $\mathbf{f}'$ that are feasible when $(\mathbf{x}, \alpha, (\mathbf{p}^i)_{i=1\cdots n+1})$ are fixed can be represented as $\{\mathbf{f} \mid C\mathbf{x} \geq D\mathbf{f}\}$, for appropriate matrices C and D. Then for $\mathbf{f}' = \mathbf{f}$ and for each flow i there exists a set B_i of all bottlenecks inequalities associated to flow i, such that for all $j \in B_i$, $(C\mathbf{x})_j = (D\mathbf{f})_j$, $D_{ji} > 0$, and $f_j \leq f_k$ for all k such that $D_{ki} > 0$. For each inequality $j \in B_i$ we can also identify a set of links $X_b(j, i)$ comprising the jth bottleneck, that is for all $k \in X_b(j, i)$, $(C\mathbf{x})_{jk} > 0$. If we increase the rate of any link in the set $X_b(j, i)$, inequality j is not saturated any more, and flow i looses bottleneck j.

We proceed by contradiction. Thus there are some flows i and j such that $f_i > f_j$. We pick an arbitrary $m \in B_i$ and link $l \in X_b(m, i)$ such that $x_l > 0$, and we also pick for each $k \in B_j$ a link $h_k \in X_b(k, j)$. We pick a slot s when link l 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 l and increase the power of link h_1 such that $p_{h_1}^{s_2} \leq P_{h_1}^{max}$ and the interferences perceived by other active users is smaller than in the original scheduling of

slot s. With this new scheduling all links have the same or higher rates, except for link l whose rate has decreased by $\epsilon_1^l(\alpha_{s_1}) > 0$ and link h_1 whose rate has increased by $\epsilon_1^h(\alpha_{s_1}) > 0$, with equalities for $\alpha_{s_1} = 0$. Therefore, we can obtain an arbitrary small ϵ_1^l and ϵ_1^h by choosing sufficiently small α_{s_1}. We repeat this process for all $h_k, k \in B_j$.

In the new allocation $\mathbf{x}'$ we have $x_l' = x_l - \epsilon^l$ where $\epsilon^l = \sum_{k \in B_j} \epsilon_k^l$, and $x_{h_k}' = x_{h_k} + \epsilon_k^h$. For a sufficiently small ϵ^l and ϵ_k^h we can obtain a new feasible rate allocation $\mathbf{f}'$ such that $f_i'(\epsilon^l) < f_i$, $f_j'(\epsilon_k^h) > f_j$ and $f_k' = f_k$ for all $k \neq i, j$. This contradicts the max-min property of rate allocation $\mathbf{f}$. ■

Proof of Proposition 2: Let us first consider only nodes $r_0, \cdots, r_{n-1}$ and consider (d_1, d_2) as a noise, and let $\mathbf{x} = \sum_{i=1}^{n+1} \alpha_i \mathbf{R}(\mathbf{p}^i)$ represent any achievable rate allocation. We can then construct a new rate allocation by rotating the optimal one $\mathbf{x}' = \sum_{i=1}^{n+1} \alpha_i \sum_{j=1}^{n} \frac{1}{n} \mathbf{R}(\mathcal{R}_j(\mathbf{p}^i))$. This allocation is feasible, and we have $x_i' = \sum_{j=1}^{n} \frac{1}{n} x_j$.

For maximization of sum of rates, define $S(\mathbf{x}) = \sum_{i=1}^{n} x_i$; we have $S(\mathbf{x}') = \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{1}{n} x_j = S(\mathbf{x})$. Similarly, for utility fairness, let $U(\mathbf{x}) = \sum_{i=1}^{n} u(x_i)$ where $u(x)$ is a strictly concave function. Thus $S(\mathbf{x}') = \sum_{i=1}^{n} u(\sum_{j=1}^{n} \frac{1}{n} x_j) \geq \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{1}{n} u(x_j) = S(\mathbf{x})$, with equality when $\mathbf{x}$ and $\mathbf{x}'$ are equal up to a rotation. It thus follows that the optimal scheduling and power allocation on a ring for utility fairness and sum of rates is of the form $\mathbf{x} = \sum_{j=1}^{n} \frac{1}{n} \mathbf{R}(\mathcal{R}_j(\mathbf{p}))$, where we call *ring rate* x_r the equal rates of all links. The same thing is shown in [2] for max-min fairness. Therefore we can restrict ourselves to the ring rates that are functions of only one power allocation.

Considering further link (d_1, d_2) we see that the interference it perceives during rotational scheduling described above is constant. We thus now have only two-dimensional rate space (x_r, x_d), and each rate can be described as in Eq 1. ■