

Spatial fairness in linear wireless multi-access networks

P.M. van de Ven^{1,2}
D. Denteneer³

J.S.H. van Leeuwaarden^{1,2}
A.J.E.M. Janssen^{2,4}

April 8, 2010

Abstract

Multi-access networks may exhibit severe unfairness in throughput. Recent studies show that this unfairness is due to local differences in the neighborhood structure: Nodes with less neighbors receive better access. We study the unfairness in saturated linear networks, and adapt the multi-access CSMA protocol to remove the unfairness completely, by choosing the activation rates of nodes appropriately as a function of the number of neighbors. We then investigate the consequences of this choice of activation rates on the network-average saturated throughput, and we show that these rates perform well in a non-saturated setting.

1 Introduction

Multi-access protocols such as CSMA [11] as used in the IEEE 802.11 standard have gained much popularity for their ability to regulate the access of network nodes to a shared communication channel in a simple and fully distributed fashion. A major drawback of these protocols, however, is that they can exhibit severe unfairness, in the sense that some of the nodes get starved, while others receive good access. One of the main causes of this unfairness is that all nodes are treated the same, irrespective of their position in the network. We propose a way of compensating for the possible spatial disadvantages of nodes by enhancing the multi-access protocols with local information about the immediate neighborhood of nodes.

Unfairness in wireless networks is an active topic of research. Wang and Kar [18] considered three nodes on a line that only block their direct neighbors, and showed that the middle node is starved when the activation rate of all three nodes increases. Such unfairness has been studied for more general networks by Durvy *et al.* [4, 5] and Denteneer *et al.* [3]. We study the same model as in [4, 5, 3], a network with n nodes on a line, in which active nodes block a certain subset of other nodes. Unblocked nodes become active, and active nodes deactivate, after exponential times. Under these assumptions, the n -dimensional process that describes the activity of nodes is a reversible continuous-time Markov process that possesses an elegant product-form expression for its stationary distribution. So this idealized linear network strikes a good balance between tractability and complexity, as it retains the essential features of competition and unfairness.

We assume that an active node blocks the first β nodes on both sides, and we say that nodes that might block each other are *neighbors*. Results from [4, 5, 3] suggest that the unfairness observed in this model is due entirely to boundary effects. To deal with the unfairness, we shall modify the model in one important way: Instead of letting nodes activate at the same rate, we allow for node-specific activation rates. By choosing the activation rate ν_i of node i as a particular function of the number of its neighbors, we can guarantee that all nodes in the network have the

¹Eindhoven University of Technology, Department of Mathematics and Computer Science, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

²Eurandom, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

³Philips Research Europe, 5656 AA Eindhoven, The Netherlands

⁴Eindhoven University of Technology, Department of Electrical Engineering, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

same throughput, completely removing the unfairness. Our main contribution is that we prove that this fair choice of activation rates $\nu_i = \nu_i^*$ takes on the extremely simple form

$$\nu_i^* = \alpha(1 + \alpha)^{\gamma(i) - \gamma(1)} \quad (1)$$

with $\gamma(i)$ the number of neighbors of node i , and α any positive constant. We note that these node-dependent activation rates are still in line with the distributed nature of the multi-access protocol, as ν_i^* only requires the number of neighbors, which a node can obtain locally by sensing its direct environment. By choosing the activation rates according to (1) we essentially ensure long-term fairness. The first result in this direction is due to Kelly [9], who considered a tree network (of which the linear network is a special case) with nearest-neighbor blocking ($\beta = 1$).

If we choose the activation rates as in (1) and let $\alpha \rightarrow \infty$, the throughput approaches $1/(\beta + 1)$, the highest possible per-node throughput. Recently, Jiang and Walrand [8], Rajagopalan and Shah [14] and Liu *et al.* [12] proposed protocols for determining the activation rates that can attain this maximal throughput in general networks. To this end, they introduced adaptive mechanisms that converge over time to the throughput-optimal activation rates. The main difference between our approach and the results in [8, 14, 12] is that we obtain an explicit expression for the throughput-optimal and fair activation rates, rather than relying on an algorithm to find these. In fact, our explicit solution can be used as a benchmark for validating, in this particular case, the various algorithms proposed in [8, 14, 12]. Besides giving insight into the operation of the network, these explicit rates also provide immediate optimal performance of the network, rather than first going through a startup stage during which nodes have to learn the right activation rates. The results discussed in our paper are only valid for linear networks, although the Markov random field approach from [9] allows us to extend our results to various other topologies, see Section 7.

The paper is structured as follows. In Section 2 we introduce the linear network in more detail. In Section 3 we study some of the key features of the unfairness that arises when all nodes have equal activation rates. Section 4 contains the proof of the proposition that the activation rates in (1) yield equal throughputs. In Section 5 we investigate the impact of the rates in (1) on the network-average throughput, and in Section 6 we discuss the extension to non-saturated networks. Section 7 presents some conclusions and further research directions.

2 Model description

We consider n nodes on a line, and assume that all nodes are saturated (i.e. have an infinite supply of packets available for transmission). Nodes activate according to a β -hop blocking model, so a transmitting node blocks the first β nodes on both sides. When node i is blocked, it remains silent until all nodes within distance β are inactive, at which point it tries to activate after an exponentially distributed (backoff) time with mean $1/\nu_i$. Node i activates if it is still unblocked when the backoff timer runs out. If a node finds itself blocked when the backoff timer expires, it waits until all neighboring nodes are inactive once more and then repeats the backoff procedure. Without loss of generality, we assume that transmissions last for an exponentially distributed time with unit mean. Under these assumptions, the n -dimensional process that describes the activity of nodes is a continuous-time Markov process.

Each state of the Markov process is described as

$$\omega = (\omega_1, \dots, \omega_n) \in \{0, 1\}^n,$$

where $\omega_i = 1$ when node i is active. Let $\Omega \subseteq \{0, 1\}^n$ be the set of all feasible states. Call $\omega \in \Omega$ *feasible* if no two 1's in ω are β positions or less apart, i.e., $\omega_i \omega_k = 0$ if $1 \leq |i - k| \leq \beta$.

The Markov process that describes the activity of nodes is then fully specified by the state space Ω and the transition rates

$$r(\omega, \omega') = \begin{cases} \nu_i & \text{if } \omega' = \omega + e_i, \\ 1 & \text{if } \omega' = \omega - e_i, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Here e_i denotes the vector with all zeros except for a 1 at position i .

Alternatively, we can express the set of feasible states as all states that satisfy a certain system of linear equations. Let A be an $(n - \beta) \times n$ matrix where each row contains $\beta + 1$ consecutive 1's, in the following way:

$$A = \begin{pmatrix} 1 & 1 & \dots & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & \dots & 1 & 0 & \dots & 0 \\ & & \ddots & & & \ddots & & \vdots \\ 0 & \dots & 0 & 1 & 1 & \dots & 1 & 0 \\ 0 & 0 & \dots & 0 & 1 & 1 & \dots & 1 \end{pmatrix}. \quad (3)$$

Now we can write the state space as $\Omega = \{\omega \in \{0, 1\}^n \mid A\omega \leq C\}$, where C is the all-1 vector of size n . This characterization has a natural interpretation as a set of capacity constraints, and nodes can activate only when enough capacity is available. We allocate unit capacity to each node, and use the convention that whenever a node is active it uses its own capacity, as well as the capacity of all its neighbors to the left. The i -th row of A thus represents the capacity required when node i is active. The constraints that arise from the last β nodes on the line are redundant, and ignoring these leads to the matrix A in (3).

From the description of the state space as a set of capacity constraints, it is clear that our model belongs to the general class of loss networks, see Kelly [10]. Loss networks are known to be reversible, and possess product-form solutions. For our Markov process, this product-form solution is given by the stationary measure π on Ω for which

$$\pi(\omega) = \begin{cases} Z_n^{-1} \prod_{i=1}^n \nu_i^{\omega_i} & \text{if } \omega \text{ is feasible,} \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

and where Z_n is the normalization constant that makes π a probability measure. This result is well-known in the context of wireless networks, see e.g. [2, 3, 4, 18].

Our main concern is with the long-term behavior of nodes, characterized by their throughputs. A common throughput-degrading phenomenon in wireless networks is collisions, which may occur when multiple nearby nodes transmit simultaneously, causing these transmissions to fail. However, we assume that the blocking range β is large enough to rule out collisions, so any activity of a node contributes to its throughput. Although a model without collisions might seem limited, numerous simulation studies show that choosing the blocking range just large enough to preclude collisions gives very good performance, see e.g. [7, 17, 19, 21]. In fact, the forthcoming paper [16] shows that this choice is throughput-optimal when the activation rates are sufficiently large. We study the throughput vector $\theta = (\theta_1, \dots, \theta_n)$, where θ_i represents the fraction of time node i is active. Denote the total number of feasible states by K , let $\Omega = \{\Omega_1, \dots, \Omega_K\}$, and introduce the $n \times K$ incidence matrix X such that $X_{ik} = 1$ when the i th element in the state Ω_k equals 1. Then

$$\theta = X \cdot \Pi, \quad (5)$$

with $\Pi = (\pi(\Omega_1), \dots, \pi(\Omega_K))$.

By exploiting the structure of the network, we can construct alternative expressions for the throughput in (5). Specifically, we shall make use of the observation that if node i is active, nodes to the left of i behave independently from nodes to the right of i . This leads to the following theorem.

Theorem 1. *Define the sequence $(Z_i)_{i=-\infty}^{\infty}$ such that $Z_i = 1$ for $i \leq 0$, and*

$$Z_i = 1 + \nu_1 + \dots + \nu_i, \quad i = 1, 2, \dots, \beta + 1, \quad (6)$$

$$Z_i = Z_{i-1} + \nu_i Z_{i-\beta-1}, \quad i = \beta + 2, \beta + 3, \dots \quad (7)$$

Let the vector of activation rates $\nu = (\nu_1, \dots, \nu_n)$ be symmetric. Then

$$\theta_i = \nu_i \frac{Z_{i-\beta-1} Z_{n-i-\beta}}{Z_n}, \quad i = 1, \dots, n. \quad (8)$$

Proof. By conditioning on whether node i is active, we can decompose the activity of the network into two parts, separated by this active node (see [2], Equation (15)),

$$\theta_i = \nu_i \frac{Z_{1:i-\beta-1} Z_{i+\beta+1:n}}{Z_{1:n}}, \quad (9)$$

where $Z_{i:j}$ is the normalization constant of a network consisting only of nodes i, \dots, j . For simplicity we denote $Z_i := Z_{1:i}$, and the symmetry of ν implies

$$Z_{i:n} = Z_{1:n-i+1}. \quad (10)$$

Substituting (10) into (9) yields the expression for θ_i in (8). By conditioning on the activity of node i , we immediately get the recursion relation (7) for the Z_i . \square

3 Unfairness

We now venture deeper into the issue of unfairness, and assume all nodes to have equal activation rates $\nu_i = \sigma$. As observed by Durvy *et al.* [4] and Denteneer *et al.* [3], the throughput distribution in this setting is highly unfair, in the sense that some nodes have a systematically higher throughput than others. This unfairness can be explained by the *node-in-the-middle* phenomenon discussed for example in Wang and Kar [18] and Garetto *et al.* [6] for the case $n = 3, \beta = 1$: The middle node is in an unfavorable position as it has to wait for both outer nodes to deactivate, whereas these boundary nodes only compete for transmission with the middle node.

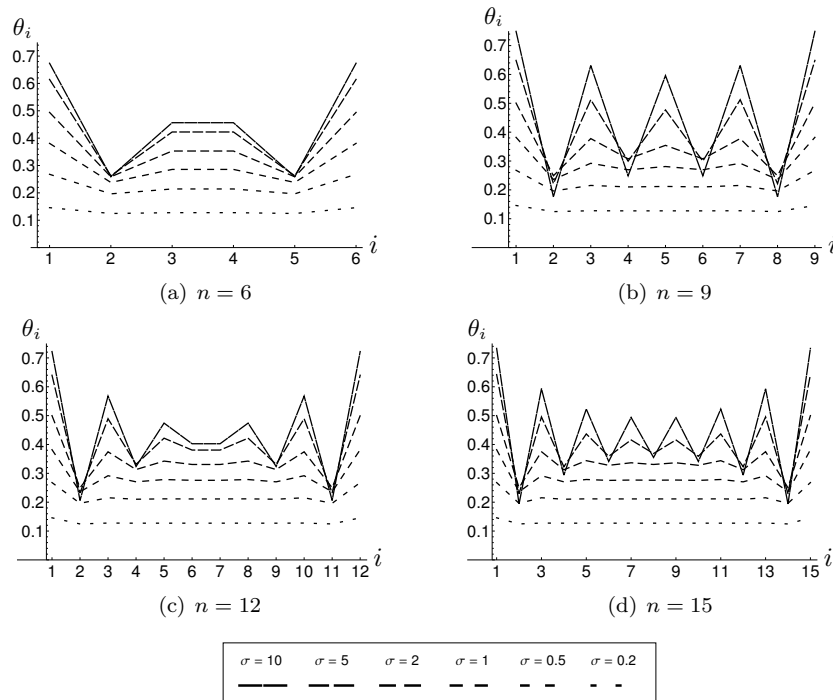


Figure 1: The per-node throughput for $\beta = 1$ and various values of n and σ

In order to compute the throughput from (8), we require information on the Z_i . For general activation rates ν_i these cannot be solved explicitly from the recursion (6)-(7) (an exception to this are the rates in (1), see Section 4). However, when all nodes have equal activation rate σ we can rewrite the Z_i as (see A)

$$Z_i = \sum_{j=0}^{\beta} c_j \lambda_j^i, \quad i = 0, 1, \dots, \quad (11)$$

where λ_j are the roots of

$$\lambda^{\beta+1} - \lambda^\beta - \sigma = 0, \quad (12)$$

and the c_j are given by

$$c_j = \frac{\lambda_j^{\beta+1}}{(\beta+1)\lambda_j - \beta}. \quad (13)$$

Moreover, when the index grows large ($i \rightarrow \infty$), the normalization constant (11) is dominated by the largest root λ_0 , i.e.,

$$Z_i \sim c_0 \lambda_0^i, \quad i \rightarrow \infty. \quad (14)$$

We shall use this asymptotic relation at several places, see (16) and (56).

For ease of presentation, we restrict ourselves to $\beta = 1$ for the remainder of this section. Figures 1(a)-1(d) show the per-node throughput for various values of n and σ . All figures display a similar pattern, with the outer nodes having the highest throughput. Moreover, the throughput is symmetric in all figures, and exhibits some form of oscillatory behavior. These observations are formalized in the following result.

Proposition 1. *For $\nu_i = \sigma > 0$, $i = 1, \dots, n$ and $\beta = 1$, the throughput has the following properties:*

- (i) *Symmetric:* $\theta_i = \theta_{n-i+1}$, $i = 1, 2, \dots, n$.
- (ii) *Alternating and converging:* $(-1)^i(\theta_{i+1} - \theta_i)$ is positive and decreasing for $i = 1, 2, \dots, \lfloor n/2 \rfloor$.

Proposition 1 is proven in B.

In Figure 1 we see that for nearest-neighbor blocking, the largest difference in throughput is between θ_1 and θ_2 . Proposition 1(ii) confirms that this is the most unfair situation, and it persists even in larger networks where the node-in-the-middle problem is mitigated by the activity of the remaining nodes. In fact, for large networks we have the following result.

Proposition 2. *For $\nu_i = \sigma > 0$, $i = 1, \dots, n$ and $\beta = 1$,*

$$\frac{\theta_1}{\theta_2} \sim \frac{1 + \sqrt{1 + 4\sigma}}{2}, \quad n \rightarrow \infty. \quad (15)$$

Proof. We have that $\theta_1 \propto Z_{n-2}$ and $\theta_2 \propto Z_{n-3}$. Using (14) we obtain

$$\frac{\theta_1}{\theta_2} \sim \lambda_0, \quad n \rightarrow \infty. \quad (16)$$

For $\beta = 1$ we can explicitly solve (12) to obtain λ_0 , and the result follows. \square

We note that for $\beta = 1$, alternative descriptions of Z_i exist of the forms

$$Z_i = (-\sigma)^{\frac{1}{2}(i+1)} U_{i+1}(\sqrt{-1/4\sigma}), \quad (17)$$

where $U_n(x)$ is the n th Chebyshev polynomial of the second kind, and

$$Z_i = \sum_{j=0}^{\lfloor \frac{i+1}{2} \rfloor} \binom{i+1-j}{j} \sigma^j. \quad (18)$$

The latter expression can be interpreted as the summation over all possible combinations of nodes that can be active simultaneously.

Figure 1 shows another interesting property of this network. Increasing σ leads in many cases to a higher throughput for each of the nodes. Hence, in such situations, one may want to increase σ further. However, we also observe that there exists a critical value σ^* , such that at least one of the throughputs θ_i decreases as σ increases beyond σ^* . The characterization of this critical value is a possible topic for future research.

Results similar to those presented in this section can be obtained for $\beta \geq 2$. As an example, Figures 2(a)-2(b) show the per-node throughput for $n = 9$ and $\beta = 2, 3$, obtained by iterating the recursion (6)-(7). Both figures exhibit similar oscillatory behavior as observed for $\beta = 1$, although the oscillation period increases with β .

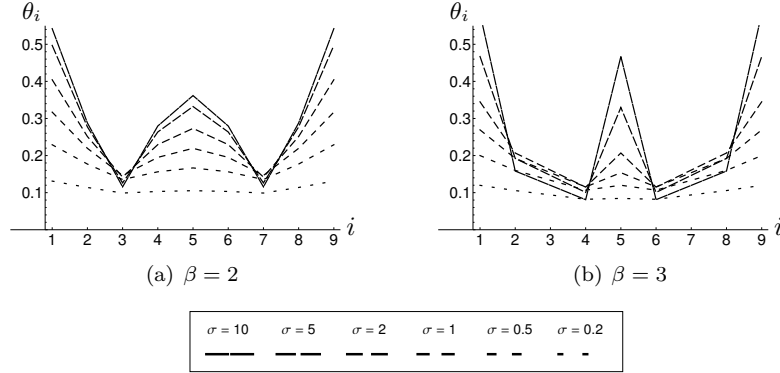


Figure 2: The per-node throughput for $n = 9$ and various values of β and σ

4 Fairness

In this section we present a way to completely remove the unfairness that was addressed in Section 3. In order to do so, we choose node-dependent activation rates ν_i such that all nodes have equal throughput ($\theta_1 = \theta_2 = \dots = \theta_n$). From (4) and (5) we see that in order to meet this objective we have to solve a system of n nonlinear equations. It seems that in general this system cannot be solved directly. We therefore choose a more indirect approach, for which we first consider two special cases that can be solved explicitly. The insight obtained from these exact solutions is then used to guess the general solution to the system of non-linear equations.

The first case is where $\beta = n - 2$, so that all but the two outer nodes will block the entire network.

Proposition 3. *For linear networks with 3 or more nodes, and $\beta = n - 2$, setting $\nu_1 = \nu_n = \alpha$ and $\nu_i = \alpha(1 + \alpha)$ for all other nodes yields equal throughputs*

$$\theta_i = \frac{\alpha}{1 + (n - 1)\alpha}, \quad i = 1, \dots, n. \quad (19)$$

Proof. The expression for the throughput in (5) can be written as

$$\theta_1 = Z_n^{-1} \nu_1 (1 + \nu_n), \quad (20)$$

$$\theta_i = Z_n^{-1} \nu_i, \quad i = 2, 3, \dots, n - 1, \quad (21)$$

$$\theta_n = Z_n^{-1} \nu_n (1 + \nu_1). \quad (22)$$

The inherent symmetry of the model allows us to set $\nu_1 = \nu_n$. Moreover, for the throughput of the other nodes to be equal, we require $\nu_2 = \dots = \nu_{n-1} = \nu_1(1 + \nu_1)$. If we set $\nu_1 = \alpha$, and substitute this into (20)-(22), we get a throughput of

$$\theta_i = Z_n^{-1} \alpha(1 + \alpha). \quad (23)$$

The normalization constant Z_n can be determined by summing over all feasible states:

$$\begin{aligned} Z_n &= 1 + \sum_{i=1}^n \nu_i + \nu_1 \nu_n = 1 + (n - 2)\alpha(1 + \alpha) + 2\alpha + \alpha^2 \\ &= (1 + \alpha)(1 + (n - 1)\alpha). \end{aligned} \quad (24)$$

Substituting (24) into (23) yields (19). \square

The case $n = 5$, $\beta = 3$ of Proposition 3 was considered in [3]. The second special case corresponds to $n = 2(\beta + 1)$, so that a node blocks at least half of the network.

Proposition 4. For linear networks with $n = 2m$ nodes, $m \in \mathbb{N}$, and $\beta = m - 1$, setting $\nu_i = \alpha(1 + \alpha)^{i-1}$ for $i = 1, \dots, m$ yields equal throughputs

$$\theta_i = \frac{\alpha}{1 + m\alpha}, \quad i = 1, \dots, n. \quad (25)$$

Proof. To achieve equal throughputs, we see from (4) and (5) that for the case at hand we should solve the system of equations

$$\begin{aligned} \nu_1 + \nu_1(\nu_{m+1} + \dots + \nu_n) &= \nu_2 + \nu_2(\nu_{m+2} + \dots + \nu_n) \\ &= \nu_3 + \nu_3(\nu_{m+3} + \dots + \nu_n) \\ &\vdots \\ &= \nu_m + \nu_m \nu_n. \end{aligned} \quad (26)$$

Indeed, the throughput of node i can be written as a sum over all states in which node i is active. Using symmetry, (26) can be written as

$$\begin{aligned} \nu_1 + \nu_1(\nu_1 + \dots + \nu_m) &= \nu_2 + \nu_2(\nu_1 + \dots + \nu_{m-1}) \\ &= \nu_3 + \nu_3(\nu_1 + \dots + \nu_{m-2}) \\ &\vdots \\ &= \nu_m + \nu_m \nu_1. \end{aligned} \quad (27)$$

Let $\nu_1 = \alpha > 0$. The solution of (27) is easily seen to be

$$\nu_i = \alpha(1 + \alpha)^{i-1}, \quad i = 1, \dots, m, \quad (28)$$

and hence

$$\theta_i = Z_n^{-1} \alpha(1 + \alpha)^m. \quad (29)$$

Summing over all possible states yields

$$\begin{aligned} Z_n &= 1 + \sum_{i=1}^n \nu_i + \sum_{i=1}^m \nu_i \sum_{j=i+m}^n \nu_j \\ &= 1 + \sum_{i=1}^m \nu_i + \sum_{i=1}^m \nu_i(1 + \nu_1 + \dots + \nu_{m-i-1}) \\ &= 1 + ((1 + \alpha)^m - 1) + m\alpha(1 + \alpha)^m \\ &= (1 + m\alpha)(1 + \alpha)^m. \end{aligned} \quad (30)$$

Substituting (30) into (29) gives (25). \square

It is clear that the complexity of the system of equations governed by (5) reduces considerably for the choices of β discussed in Propositions 3 and 4. For general β this system remains rather complicated, and the direct approach taken in the proofs of Propositions 3 and 4 no longer seems to work.

However, we can use Propositions 3 and 4 to make an educated guess about the general solution. First observe that the fair activation rates in Propositions 3 and 4 only depend on the number of neighbors (nodes within β hops) that each node has. Denote by $\gamma(i)$ the number of neighbors of node i , let $\alpha > 0$, and choose activation rates ν_i^* as in (1). We see that this choice is consistent with the fair activation rates in Propositions 3 and 4. We now show that ν_i^* indeed achieves fairness for all β . To this end, we first show that when the activation rates are chosen according to (1), the recursive relation (6)-(7) for the normalization constant Z_i has a closed-form solution.

Lemma 1. Let $\alpha > 0$ and choose ν_i^* as in (1). Then

$$Z_i = (1 + \alpha)^i, \quad i = 1, 2, \dots, n - \beta. \quad (31)$$

Proof. This can be verified by substituting the solution (31) into the relations (6) and (7). For $i \leq \beta + 1$, (6) gives

$$Z_i = 1 + \alpha + \alpha(1 + \alpha) + \dots + \alpha(1 + \alpha)^{i-1} = (1 + \alpha)^i. \quad (32)$$

From (7) we see that

$$Z_i = (1 + \alpha)^{i-1} + \alpha(1 + \alpha)^\beta(1 + \alpha)^{i-\beta-1} = (1 + \alpha)^i, \quad (33)$$

covering the case $i \geq \beta + 2$. \square

With Lemma 1 we are now in the position to prove our main result.

Theorem 2. Let $\alpha > 0$, $\beta \leq n - 1$ and choose ν_i^* as in (1). Then

$$\theta_i = \frac{\alpha}{1 + (1 + \beta)\alpha}, \quad i = 1, \dots, n. \quad (34)$$

Proof. To prove this result we substitute the normalization constants from Lemma 1 into the expression for the throughput in (8). We distinguish between different values of i .

For $i \geq \beta + 1$ and $i \leq n - \beta$ we see that $\nu_i^* = \alpha(1 + \alpha)^\beta$ and

$$Z_{i-\beta-1} = (1 + \alpha)^{i-\beta-1}, \quad Z_{n-i-\beta} = (1 + \alpha)^{n-i-\beta}. \quad (35)$$

Similarly for $i \geq \beta + 1$ and $i \geq n - \beta + 1$ we have $\nu_i^* = \alpha(1 + \alpha)^{n-i}$ and

$$Z_{i-\beta-1} = (1 + \alpha)^{i-\beta-1}, \quad Z_{n-i-\beta} = 1. \quad (36)$$

For $i \leq \beta$ and $i \leq n - \beta$ we have $\nu_i^* = \alpha(1 + \alpha)^{i-1}$, and

$$Z_{i-\beta-1} = 1, \quad Z_{n-i-\beta} = (1 + \alpha)^{n-i-\beta}, \quad (37)$$

and finally for $i \leq \beta$ and $i \geq n - \beta + 1$ we have $\nu_i^* = \alpha(1 + \alpha)^{n-\beta-1}$ and

$$Z_{i-\beta-1} = 1, \quad Z_{n-i-\beta} = 1. \quad (38)$$

Substituting (35)-(38) into (8) yields

$$\theta_i = Z_n^{-1} \alpha(1 + \alpha)^{n-\beta-1}. \quad (39)$$

We next consider the normalization constant. Choose m such that $n = \beta + m$. Then

$$\begin{aligned} Z_n &= Z_{n-1} + \nu_n^* Z_{n-\beta-1} \\ &= Z_{n-2} + \nu_{n-1}^* Z_{n-\beta-2} + \nu_n^* Z_{n-\beta-1} \\ &\vdots \\ &= Z_{n-\beta} + \sum_{i=1}^{\beta} \nu_{n+1-i}^* Z_{n-\beta-i}. \end{aligned} \quad (40)$$

Substituting (31) into (40) yields

$$\begin{aligned} Z_n &= (1 + \alpha)^{n-\beta} + \sum_{i=1}^{\min\{m, \beta\}} \alpha(1 + \alpha)^{i-1} (1 + \alpha)^{n-\beta-i} + \sum_{i=m+1}^{\beta} \alpha(1 + \alpha)^{n-\beta-i} \\ &= (1 + \alpha)^{n-\beta-1} (1 + (\beta + 1)\alpha). \end{aligned} \quad (41)$$

Combining (41) and (39) leads to (34). \square

To better understand why the rates (1) successfully ensure strict fairness using only the number of neighbors of each node, we study its behavior in the two limiting regimes of light traffic ($\alpha \downarrow 0$) and heavy traffic ($\alpha \rightarrow \infty$). First rewrite (1) as

$$\nu_i^* = \alpha \sum_{j=0}^{\gamma(i)-\gamma(1)} \binom{\gamma(i)-\gamma(1)}{j} \alpha^j, \quad i = 1, \dots, n. \quad (42)$$

When α is small, nodes activate slowly, and few nodes will be active simultaneously. In fact, the limiting distribution is dominated by states with at most one node active, and node interaction (blocking) is negligible. This is reflected in the light-traffic activation rates that follow immediately from (42):

$$\nu_i^* = \alpha + (\gamma(i) - \gamma(1))\alpha^2 + \mathcal{O}(\alpha^3), \quad \alpha \downarrow 0. \quad (43)$$

So the first-order light-traffic activation rate $\nu_i^* = \alpha$ is the same for all nodes. Indeed, if at most one node is active (as is the case for α small), there is no blocking, and therefore no need to differentiate among nodes. As α increases, states with two active nodes are increasingly likely. This introduces node interaction, as nodes may now block their neighbors (all nodes within distance β). This is accounted for in the activation rate by the term $(\gamma(i) - \gamma(1))\alpha^2$, which is linear in the number of neighbors. So in light-traffic, only the number of neighbors is of importance, rather than the structure of the entire network. This reasoning extends to more general networks.

When α is large, we see from (42) that

$$\nu_i^* = \alpha^{\gamma(i)-\gamma(1)+1} + \mathcal{O}(\alpha^{\gamma(i)-\gamma(1)}), \quad \alpha \rightarrow \infty. \quad (44)$$

Nodes activate almost immediately when they get the chance to do so (i.e. when all neighbors are inactive), and thus the system spends almost all the time in maximal independent sets of active nodes. However, as some nodes have a higher activation rate than others, the only states ω that have positive limiting probability for $\alpha \rightarrow \infty$ are those that maximize the sums of the exponents of α over all active nodes, $\sum_{i=1}^n (\gamma(i) - \gamma(1)) I_{\{\omega_i=1\}}$, with I_{Ω} the indicator function. This is in contrast to the heavy-traffic behavior of the system in which all nodes have equal activation rates σ . Here only states with the most active nodes (maximal independent sets) have positive limiting probability when $\sigma \rightarrow \infty$. As (1) provides fairness even in heavy-traffic, the heavy-traffic fair activation rates are such that all nodes are contained an equal number of times in these dominant sets. This result strongly depends on the structure of the network, as the maximal independent sets may change drastically with the addition or the removal of even a single node. So contrary to light-traffic, the simple, locally determined heavy-traffic activation rates (44) may not easily extend beyond linear networks.

5 Network-average throughput

The fair rates ν_i^* in (1) are designed to remove the unfairness that arises when all nodes have equal activation rates σ . In order to compare the two schemes, we want to set their respective parameters α and σ such that the average per-node throughputs are equal. For given equal rates $\sigma > 0$ we write $Z_i(\sigma)$ and $\theta_i(\sigma)$ for the normalization constant of a network with i nodes, and the throughput of node i , respectively. We let $\bar{\theta}_n(\sigma) := n^{-1} \sum_{i=1}^n \theta_i(\sigma)$ denote the average per-node throughput in a network with n nodes, and we denote $\bar{\theta}(\sigma) := \lim_{n \rightarrow \infty} \bar{\theta}_n(\sigma)$.

In Section 4 we showed that all nodes have equal throughput $\alpha/(1 + \alpha(\beta + 1))$ when using the fair activation rates. When all nodes have equal activation rates, a closed-form expression for the throughput does not seem available. However, we can express the average throughput in terms of the normalization constant Z_n .

Proposition 5. *Let $\nu_i = \sigma$, $i = 1, \dots, n$. The average per-node throughput is given by*

$$\bar{\theta}_n(\sigma) = \frac{\sigma}{n} \frac{dZ_n(\sigma)}{d\sigma} / Z_n(\sigma). \quad (45)$$

Proof. We have from (8) with $\nu_i = \sigma$ that

$$\bar{\theta}_n(\sigma) = \frac{\sigma}{nZ_n} \sum_{i=1}^n Z_{i-\beta-1} Z_{n-i-\beta}. \quad (46)$$

We compute, using the definition of Z_i in Theorem 1,

$$\sum_{n=1}^{\infty} \left(\sum_{i=1}^n Z_{i-\beta-1} Z_{n-i-\beta} \right) x^n = x \left(\frac{x^\beta - 1}{x - 1} + x^\beta G_Z(x) \right)^2, \quad (47)$$

where $G_Z(x)$ is the generating function of the Z_i . Then from (66) and some simplifications, we get

$$\sum_{n=1}^{\infty} \left(\sum_{i=1}^n Z_{i-\beta-1} Z_{n-i-\beta} \right) x^n = \frac{x}{(1 - x - \sigma x^{\beta+1})^2}. \quad (48)$$

On the other hand, we compute from the explicit form (66) of $G_Z(x)$ that

$$\frac{d}{d\sigma} [G_Z(x)] = \frac{x}{(1 - x - \sigma x^{\beta+1})^2}, \quad (49)$$

and the result follows. \square

Using Proposition 5 and the partial fraction expansion (11) of Z_i , we can express the average per-node throughput in terms of the roots $\lambda_0, \dots, \lambda_\beta$ of (12).

Proposition 6. *Let $\nu_i = \sigma$, $i = 1, \dots, n$. The average per-node throughput is given by*

$$\bar{\theta}_n(\sigma) = \frac{\sigma P}{n Q}, \quad (50)$$

where

$$P = \sum_{j=0}^{\beta} \frac{\lambda_j^{n+1}}{(\beta+1)\lambda_j - \beta} \left(\frac{n+\beta+1}{(\beta+1)\lambda_j - \beta} - \frac{(\beta+1)\lambda_j}{((\beta+1)\lambda_j - \beta)^2} \right), \quad Q = \sum_{j=0}^{\beta} \frac{\lambda_j^{n+\beta+1}}{(\beta+1)\lambda_j - \beta}. \quad (51)$$

Proof. By (68) and (69) we have

$$Z_n(\sigma) = \sum_{j=0}^{\beta} \frac{\lambda_j^{n+\beta+1}}{(\beta+1)\lambda_j - \beta}, \quad (52)$$

where λ_j are the $(\beta+1)$ roots λ of (12). By implicit differentiation of (12) with respect to σ we find

$$\frac{d\lambda_j}{d\sigma} = \frac{1}{\lambda_j^{\beta-1}} \frac{1}{(\beta+1)\lambda_j - \beta}. \quad (53)$$

Then from (52) and (53) we get

$$\begin{aligned} \frac{dZ_n(\sigma)}{d\sigma} &= \sum_{j=0}^{\beta} \left(\frac{(n+\beta+1)\lambda_j^{n+\beta}}{(\beta+1)\lambda_j - \beta} - \frac{(\beta+1)\lambda_j^{n+\beta+1}}{((\beta+1)\lambda_j - \beta)^2} \right) \frac{d\lambda_j}{d\sigma} \\ &= \sum_{j=0}^{\beta} \frac{\lambda_j^{n+1}}{(\beta+1)\lambda_j - \beta} \left(\frac{n+\beta+1}{(\beta+1)\lambda_j - \beta} - \frac{(\beta+1)\lambda_j}{((\beta+1)\lambda_j - \beta)^2} \right). \end{aligned} \quad (54)$$

The result follows from substituting (52) and (54) into the expression for the average per-node throughput (45). \square

When the network grows large ($n \rightarrow \infty$) the root of largest modulus, λ_0 , becomes dominant, and (50)-(51) simplifies.

Corollary 1. *Let $\nu_i = \sigma$, $i = 1, \dots, n$. The limiting average per-node throughput as $n \rightarrow \infty$ is given by*

$$\bar{\theta}(\sigma) = \frac{\lambda_0 - 1}{(\beta + 1)\lambda_0 - \beta}. \quad (55)$$

Proof. We have, as $n \rightarrow \infty$,

$$P = \frac{\lambda_0^{n+1}}{(\beta + 1)\lambda_0 - \beta} \frac{n}{(\beta + 1)\lambda_0 - \beta} (1 + o(1)), \quad Q = \frac{\lambda_0^{n+\beta+1}}{(\beta + 1)\lambda_0 - \beta} (1 + o(1)). \quad (56)$$

Hence

$$\bar{\theta}_n(\sigma) = \frac{\sigma P}{n Q} = \frac{\sigma \lambda_0^{-\beta}}{(\beta + 1)\lambda_0 - \beta} (1 + o(1)), \quad (57)$$

and the result follows as $\sigma \lambda_0^{-\beta} = \lambda_0 - 1$ by (12). \square

The limiting expression (55) occurs in a variety of contexts in [13, 1, 20, 5, 16]. When $\beta\sigma \rightarrow \infty$, the throughput (55) simplifies even further.

Corollary 2. *Let $\nu_i = \sigma$, $i = 1, \dots, n$ and let $n \rightarrow \infty$. The limiting average throughput as $\beta\sigma \rightarrow \infty$ satisfies*

$$\bar{\theta}(\sigma) = \frac{1}{\beta + 1} (1 + o(1)). \quad (58)$$

Proof. By rewriting (55) we have

$$\bar{\theta}(\sigma) = \frac{1}{\beta + 1} \frac{1}{1 + \frac{1}{(\beta + 1)(\lambda_0 - 1)}}. \quad (59)$$

So for $\bar{\theta}(\sigma) = \frac{1}{\beta + 1} (1 + o(1))$ to hold, it is necessary and sufficient that $(\beta + 1)(\lambda_0 - 1) \rightarrow \infty$. Recall from (12) that λ_0 is such that

$$\lambda_0^\beta (\lambda_0 - 1) = \sigma. \quad (60)$$

Let $M > 0$ be some positive constant, and assume that $\beta\sigma \leq M$. Then

$$\beta \lambda_0^\beta (\lambda_0 - 1) = \beta\sigma \leq M, \quad (61)$$

and so $\beta(\lambda_0 - 1) \leq M$. Conversely, assume that $\beta(\lambda_0 - 1) \leq K$ for some positive constant $K > 0$. Then

$$\beta\sigma = \beta(\lambda_0 - 1)\lambda_0^\beta \leq \beta(\lambda_0 - 1)\exp(\beta(\lambda_0 - 1)) \leq Ke^K. \quad (62)$$

Hence

$$\beta(\lambda_0 - 1) \text{ bounded} \Leftrightarrow \beta\sigma \text{ bounded}. \quad (63)$$

It follows that a sufficient condition for $(\beta + 1)(\lambda_0 - 1) \rightarrow \infty$ is that $\beta\sigma \rightarrow \infty$. \square

Corollary 2 implies for β fixed and $\sigma \rightarrow \infty$ that $\bar{\theta}(\sigma) \rightarrow \frac{1}{\beta + 1}$. So for $n \rightarrow \infty$, both the equal and fair activation rates can achieve the maximum throughput by letting $\sigma \rightarrow \infty$ and $\alpha \rightarrow \infty$, respectively. This can be explained by the observation that for both sets of activation rates, the system spends almost all the time in maximal independent sets of active nodes, thus maximizing spatial reuse.

Next, fix $\sigma > 0$ and search for $\alpha = \alpha_n(\sigma)$ such that

$$\bar{\theta}_n(\sigma) = \frac{\alpha}{1 + \alpha(\beta + 1)},$$

so the network-average throughput is identical for the fair rates and equal rates. For $\alpha(\sigma) := \lim_{n \rightarrow \infty} \alpha_n(\sigma)$ we can make this comparison explicit.

Proposition 7. *We have*

$$\alpha(\sigma) = \lambda_0 - 1. \quad (64)$$

Proof. This follows at once by equating (34) and (55) and solving for α . \square

It is intuitively clear that imposing fairness may compromise the throughput. From (34) it is seen that the fair per-node throughputs are bounded above by $\frac{1}{\beta+1}$, and that this upper bound can be approached by letting $\alpha \rightarrow \infty$. Corollary 1 shows that, as $n \rightarrow \infty$, the average throughputs in the fair case and unfair case are equal when α is taken to be $\lambda_0 - 1$. The maximum activation rate in this limiting case equals

$$\nu_{\max}^* = \alpha(1 + \alpha)^\beta = \nu_i, \quad \beta + 1 \leq i \leq n - \beta \quad (65)$$

as is seen from (1). This maximum grows like σ since $1 + \alpha = \lambda_0 = \sigma^{1/(\beta+1)}(1 + o(1))$, $\sigma \rightarrow \infty$, according to (73). Hence, as $n \rightarrow \infty$, the fair case achieves the same average throughput with activation rates that are of smaller or of the same size as as in the unfair case. The situation is considerably more complicated when we keep n bounded. Then it may well occur that $\bar{\theta}_n(\sigma)$ exceeds $\frac{1}{\beta+1}$, which is the upper bound for the throughput achievable by the fair scheme. In Figure 3 we have plotted curves $\alpha = \alpha_n(\sigma)$, for the case that $n = 10$ and various values of β in the range $0 \leq \sigma \leq 20$. The occurrence of the asymptotes for the curves with $\beta = 2, 5$ shows that $\bar{\theta}_n(\sigma) > \frac{1}{\beta+1}$ when σ is to the right of the asymptote. Remarkably, in the cases that $\beta = 1, 4, 9$, no such asymptotes exist. This can be explained by the observation that a system operating with equal rates $\sigma \rightarrow \infty$ has at all times $\lceil \frac{n}{\beta+1} \rceil$ nodes activated, cf. [3]. When $\frac{n}{\beta+1}$ is integer, the resulting average throughput does not exceed the upper bound $\frac{1}{\beta+1}$ for the throughput in the fair case and so no asymptote occurs.

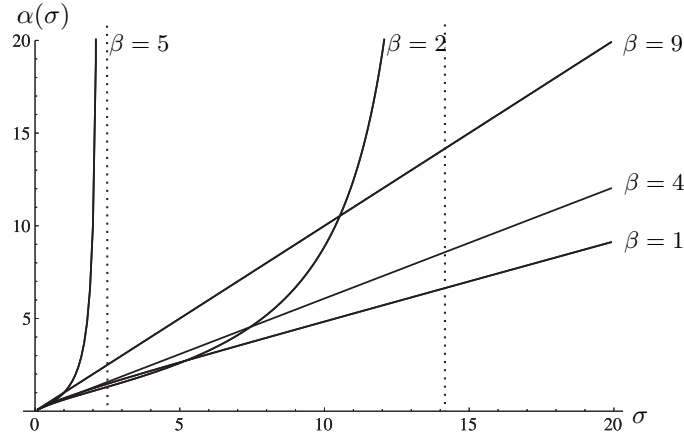


Figure 3: $\alpha_n(\sigma)$ plotted against σ for various values of β and $n = 10$

6 Non-saturated networks

The saturated model described in Section 2 assumes all nodes to have an infinite supply of packets pending for transmission. We now relax this assumption by introducing queueing dynamics. Packets arrive at node 1 according to a Poisson process with rate r , and are routed along nodes $2, \dots, n$. When node n finishes a transmission, the corresponding packet exits the system. All nodes have infinite buffer capacity. These dynamics describe a non-saturated network, where packets are forwarded through the network, using nodes $2, \dots, n$ as relays. As nodes are no longer saturated, they may not always have packets available for transmission. When the backoff timer

of a node expires, and the node is not blocked, it can activate if its buffer contains at least one packet. Otherwise, the node will remain inactive and restart its backoff timer.

To describe the state of the non-saturated network we have to take into account the queue length at each node, as well as node activity. The stochastic process that describes both the activity of nodes and their queue lengths is once again a Markov process, but does not have the appealing product-form limiting distribution encountered in the saturated network. The added complexity of the model is reflected by the lack of available analytical results. When node 1 is saturated (i.e. $r \rightarrow \infty$) and all nodes have equal activation rate $\sigma \rightarrow \infty$ the system simplifies, and in some cases can be analyzed. For $n = 3$, $\beta = 1$ it can be shown (cf. [3]) that both nodes 1 and 2 are unstable, so that the number of packets in their buffers grows without bounds. Using this, [3] shows the throughput of node 3 to be equal to $3/10$. Shneer and Van de Ven [15] demonstrate that for general β and $n = 2\beta + 1$, nodes $1, \dots, \beta + 1$ are unstable, and the throughput of node n is bounded as $1/(2\beta + 2) \leq \theta_n \leq 1/(2\beta + 1)$.

The throughput of node n is of special interest, as it represents the end-to-end throughput of the network, that is, the intensity at which packets leave the network. When $\theta_n = r$, all packets that enter the network eventually leave. If $\theta_n < r$, packets arrive into the system at a higher rate than the network can process, and packets will accumulate at certain bottleneck nodes. The two results for the special case in [3] and [15] paint a particularly bleak picture as, despite the high activation rate, the end-to-end throughput is well below the theoretical upper bound $1/(\beta + 1)$ that can be attained through centralized schedulers. Instead, we investigate the performance of a non-saturated network that uses the fair rates (1).

Figure 4 shows simulation results for the end-to-end throughput in a network using the fair rates, with $n = 5$ and $\beta = 1$, plotted against the arrival rate r . The dashed line corresponds to the network where all nodes have equal activation rate $\sigma = 6$, and the solid line shows the throughput of a network with fair activation rates (1) and $\alpha = 11.68$. The values of σ and α are chosen such that the average per-node throughput is equal for both activation schemes, see Section 5. The network with equal activation rates performs poorly. When the arrival rate grows beyond a certain threshold, node 2 saturates and the throughput drops. The network with fair rates, on the other hand, can sustain higher arrival rates and experiences no throughput drop when in overload. In fact, the end-to-end throughput approaches the per-node throughput in the corresponding saturated network (indicated by the dotted horizontal line). So the network is stable whenever $r < \alpha/(1 + \alpha(\beta + 1))$.

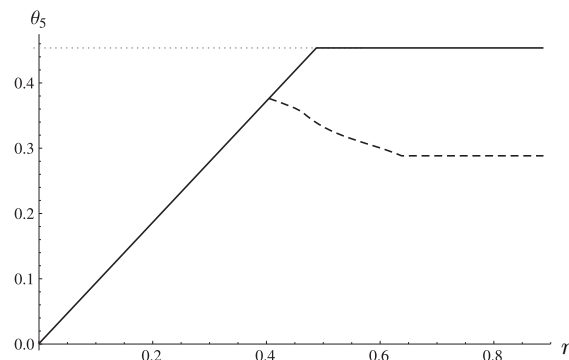


Figure 4: The end-to-end throughput of a network with equal activation rates (dashed) and fair activation rates (solid), plotted against the arrival rate at node 1

We see from Figure 4 that as r increases, the end-to-end throughput eventually reaches the per-node throughput in the saturated network for $\alpha = 11.68$. Sustaining this high end-to-end throughput requires an efficient activation scheme, in which maximal independent sets of active nodes are persistent for a long time, yet all nodes get equal access. Figure 5 shows a typical sample of the activity pattern, for both the equal rates and the fair rates. The arrival rate equals

r equals 0.47, so the network with fair rates is stable. The barcode plots indicate for each point in time whether nodes are active (black) or inactive (white). We see from Figure 5(a) that for equal activation rates, the pattern is very irregular. Node 1 is more active than the other nodes, indicating that the network is in fact unstable, and node 2 is the bottleneck node. The fair activation rates result in a more regular pattern, and maintain certain maximal independent sets of active nodes for a long period of time. Figure 5(b) shows that the maximal independent sets $\{1, 3, 5\}$ and $\{2, 4\}$ are persistent. So instead of sending a single packet before relinquishing access to the medium, nodes send packets in large batches.

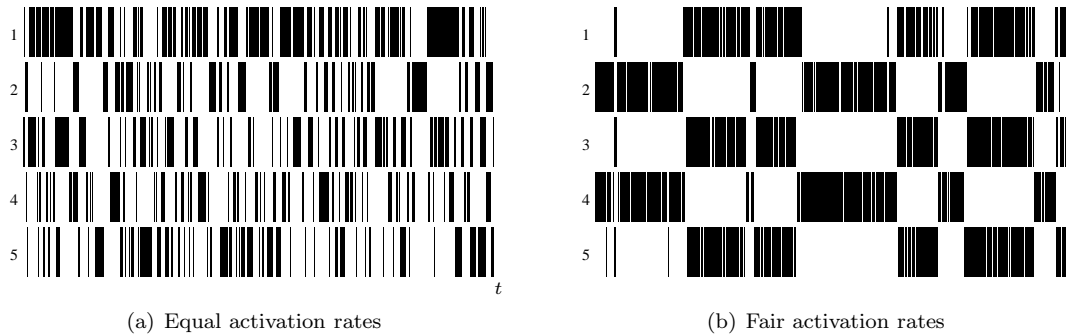


Figure 5: The activity charts for both equal activation rates and fair activation rates for $n = 5$, $\beta = 1$ and $r = 0.47$

The length of the dominant periods of a maximal independent set of active nodes, and therefore the size of the batches, is related to the queue lengths of the nodes in this set. Typically, when one of the queues of the dominant maximal independent set empties, dominance reverts to the other maximal independent set. Figure 6 shows the queue lengths of all nodes over time. It shows large fluctuations in queue length, which are directly related to the large batch sizes. Nodes 4 and 5 in particular empty relatively often, suggesting that these are in fact the nodes that typically break the dominance of their own maximal independent set.

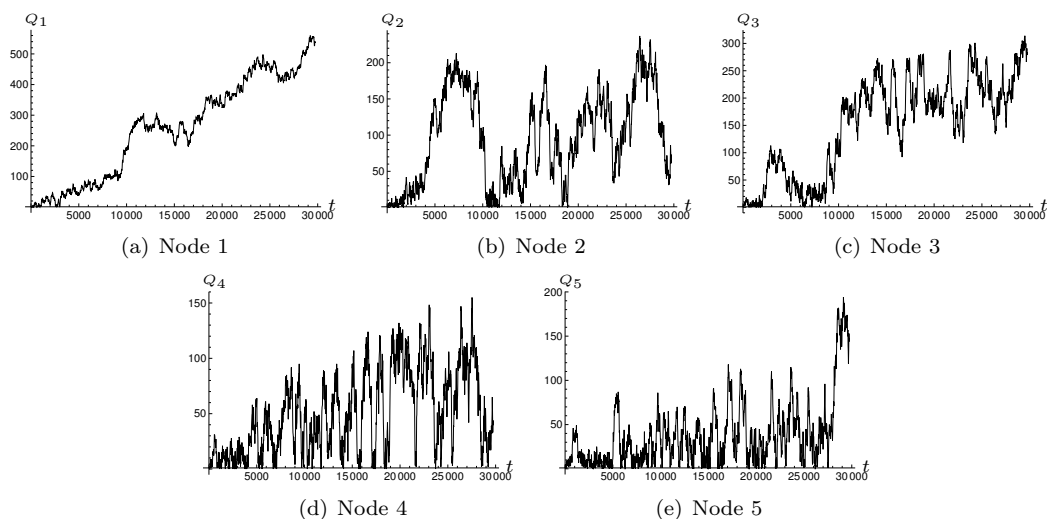


Figure 6: The queue lengths for the various nodes with fair activation rates, $n = 5$, $\beta = 1$ and $r = 0.47$, on a large time scale

In contrast to the saturated network, the non-saturated network may waste capacity when

unblocked nodes cannot activate because of empty buffers. Figure 7 is similar to Figure 6, except for the smaller time scale. It suggests that queues rarely empty, and moreover do not stay empty for a long time, so the wasted capacity is very small, if not negligible.

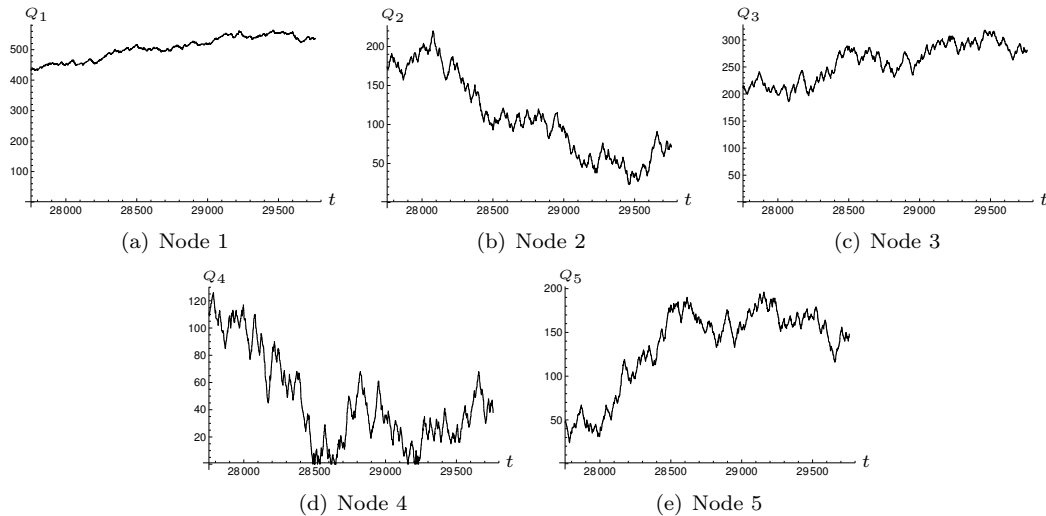


Figure 7: The queue lengths for the various nodes with fair activation rates, $n = 5$, $\beta = 1$ and $r = 0.47$, on a small time scale

The results presented in this section carry over to general n and β . The system with fair activation rates can sustain arrival rates up to the per-node throughput of the corresponding saturated network, and maintains its throughput when r grows even larger. The network with equal activation rates with the same (saturated) average per-node throughput fails to attain this end-to-end throughput. Whenever this system is overloaded, its throughput drops. Recall from (34) that for $\alpha \rightarrow \infty$, the saturated network approaches a per-node throughput of $1/(\beta + 1)$. So the corresponding non-saturated network even approaches the theoretic upper bound on the end-to-end throughput.

The high end-to-end throughput of a non-saturated network using the fair rates, and the poor performance of such a network where all nodes have equal activation rates, can be explained by the result derived for saturated networks in Sections 3 and 4. Indeed, although non-saturated networks are more complex than their saturated counterparts, the basic premise that nodes in the middle of the network have an unfavorable position compared with nodes on the network boundary remains.

7 Conclusions and outlook

In this paper we studied the unfairness in linear multi-access networks. We proposed node-specific activation rates in (1) as a function of the number of neighbors, and showed that these rates provide equal throughput for all nodes. The rates increase with the number of neighbors. Intuitively, this structure can be explained by the observation that, as the number of neighbors increases, a node needs a higher activation rate to retain its throughput. Consequently, the rule in (1), which is exact in linear networks, might serve as a heuristic in more complex networks.

Finding exact expressions for the activation rates that provide strict fairness for networks beyond the linear network is challenging. Kelly [9] obtained results for trees with nearest-neighbor blocking, and finds that rates such as in (1), where nodes on the leaves of the tree have lower rates than those in the stem of the tree, provide strict fairness. For such trees, it seems possible to extend Kelly's result to the β -hop blocking situation. Also other regular topologies such as certain grids appear to admit similar analysis.

Finally, let us mention that regular networks like grids or trees may not always be a good representation of topologies encountered in practice, which in general are less structured. The results obtained in this paper, however, rely heavily on the diagonal structure of the capacity matrix A in (3), which only exists for certain well-structured networks. For more general networks, and hence more general matrices A , the objective of equal throughputs boils down to solving the system of nonlinear equations that follows from (5). In fact, (5) can be described in terms of the mapping (with $\nu = (\nu_1, \dots, \nu_n) \in (0, \infty)^n$)

$$\nu \mapsto \theta(\nu) = Z_n \cdot X \cdot \Pi = \left(\sum_{\substack{\omega \in \Omega \\ \omega_k \neq 0}} \prod_{i=1}^n \nu_i^{\omega_i} \right)_{k=1, \dots, n}$$

with $\theta(\nu) \in (0, \infty)^n$. It can be shown that the mapping θ is globally invertible on $(0, \infty)^n$. Thus, given a vector $c \in (0, \infty)^n$, there is a unique $\nu = \nu(c) \in (0, \infty)^n$ such that $\theta(\nu) = c$. When c has identical entries, this corresponds to all nodes having equal throughput. The full analysis of this fixed-point equation is rather involved and will appear elsewhere.

A Normalization constant

We take a closer look at the normalization constant Z_i defined in (6)-(7), in the regime where all nodes have equal activation rates $\nu_i = \sigma$. This section closely follows Pinsky and Yemini [13] and Janssen *et al.* [16]. All proofs can be found in [16].

The generating function $G_Z(x)$ of the Z_i is given by

$$G_Z(x) = \sum_{i=0}^{\infty} Z_i x^i = \frac{x - 1 + \sigma x^{\beta+1} - \sigma x}{(x-1)(1-x-\sigma x^{\beta+1})}. \quad (66)$$

Let $\lambda_0, \dots, \lambda_\beta$ denote the $\beta + 1$ roots of

$$\lambda^{\beta+1} - \lambda^\beta - \sigma = 0. \quad (67)$$

These roots can be shown to be distinct, and from Rouché's theorem it follows that there exists a unique positive real root λ_0 such that $\lambda_0 > |\lambda_j|$, $j = 1, \dots, \beta$. We can obtain the Z_i from the generating function by applying partial fraction expansion, which gives ([16], Proposition 1)

$$Z_i = \sum_{j=0}^{\beta} c_j \lambda_j^i, \quad i = 0, 1, \dots, \quad (68)$$

where λ_j are the roots of (67), and

$$c_j = \frac{\lambda_j^{\beta+1}}{(\beta+1)\lambda_j - \beta}. \quad (69)$$

When the index grows large ($i \rightarrow \infty$), the normalization constant (68) is dominated by the largest root λ_0 , thus simplifying considerably, i.e.,

$$Z_i = c_0 \lambda_0^i (1 + o(1)), \quad i \rightarrow \infty. \quad (70)$$

The roots $\lambda_0, \dots, \lambda_\beta$ have representations in terms of infinite series expressions. Let $(x)_n = \Gamma(x+n)/\Gamma(x)$ denote the Pochhammer symbol. Depending on the value of σ , either of the following two expansions applies (See [16], Proposition 7 and 8). For small $\sigma > 0$, we have

$$\lambda_0(\sigma) = 1 + \sum_{l=1}^{\infty} \frac{(-1)^{l-1} (\beta l)_{l-1}}{l!} \sigma^l, \quad (71)$$

$$\lambda_j(\sigma) = \sum_{l=1}^{\infty} \frac{(l/\beta)_{l-1}}{l!} w_j^l, \quad j = 1, 2, \dots, \beta, \quad (72)$$

where $w_j = \sigma^{1/\beta} e^{2\pi i(j-1/2)/\beta}$. The series expansions in (71) and (72) converge for

$$0 \leq \sigma \leq \frac{\beta^\beta}{(\beta+1)^{\beta+1}} =: \xi(\beta),$$

and diverge otherwise. For large $\sigma > 0$, we have

$$\lambda_j(\sigma)^{-1} = \sum_{l=1}^{\infty} \frac{\binom{-l}{\beta+1}_{l-1}}{l!} v_j^{-l}, \quad (73)$$

where $v_j = \sigma^{1/(\beta+1)} e^{2\pi i j/(\beta+1)}$. The series expansion in (73) converges for $\sigma \geq \xi(\beta)$, and diverges otherwise.

B Proof of Proposition 1

We first establish an auxiliary result. Define $a(i, l, n)$ as the number of states in which exactly l nodes are active, including node i . For successive nodes, the following relation holds.

Lemma 2. *Let $n \in \mathbb{N}, i \leq \lceil \frac{n}{2} \rceil - 1$. Then:*

$$a(i, l, n) = a(i+1, l, n), \quad l \leq i, \quad (74)$$

$$a(i, l, n) > a(i+1, l, n), \quad i \text{ odd}, i < l \leq \lceil n/2 \rceil, \quad (75)$$

$$a(i, l, n) < a(i+1, l, n), \quad i \text{ even}, i < l \leq \lceil n/2 \rceil. \quad (76)$$

Proof. The proof is by induction on i . Conditioning on activity of node 1 and node n yields the relations

$$a(i, l, n) = a(i-2, l-1, n-2) + a(i-1, l, n-1), \quad (77)$$

$$a(i, l, n) = a(i, l-1, n-2) + a(i, l, n-1), \quad (78)$$

with boundary conditions

$$a(0, l, n) = 0 \text{ for all } n \text{ and } l;$$

$$a(1, l, n) = 1 \text{ for } l > 0 \text{ and all } n;$$

$$a(1, l, n) = 0 \text{ for } l \leq 0 \text{ and all } n.$$

Hence, the initialization step of the induction is

$$a(0, l, n) < a(1, l, n), \quad 0 < l < \lceil n/2 \rceil,$$

$$a(0, l, n) = a(1, l, n), \quad l \leq 0.$$

Consider odd $i \leq \lceil n/2 \rceil - 2$, let $i+1 < l < \lceil n/2 \rceil$, and assume $a(i, l, n) > a(i+1, l, n)$. Using (77) and (78) we get

$$\begin{aligned} a(i+1, l, n) &= a(i+1, l-1, n-2) + a(i+1, l, n-1) \\ &< a(i, l-1, n-2) + a(i+1, l, n-1) = a(i+2, l, n). \end{aligned}$$

This proves assertion (75). Assertions (74) and (76) can be proved in a similar manner. \square

We now use Lemma 2 to prove Proposition 1.

Proof. (Proposition 1) Assertion (i) can be shown by rewriting the throughput as follows:

$$\theta_i = Z_n^{-1} \sum_l a(i, l, n) \sigma^l = Z_n^{-1} \sum_l a(n-i+1, l, n) \sigma^l = \theta_{n-i+1}. \quad (79)$$

To prove assertion(ii) we first show that $(-1)^i(\theta_{i+1} - \theta_i)$ is positive. That is,

$$\begin{aligned} (-1)^i(\theta_{i+1} - \theta_i) &= (-1)^i Z_n^{-1} \sum_l (a(i+1, l, n) - a(i, l, n)) \sigma^l \\ &= (-1)^i Z_n^{-1} \sum_{l=i+1}^{\lfloor n/2 \rfloor} (a(i+1, l, n) - a(i, l, n)) \sigma^l > 0, \end{aligned} \quad (80)$$

where the inequality follows from Lemma 2. Using (80), Proposition 1(ii) follows from

$$\begin{aligned} (-1)^i(\theta_{i+1} - \theta_i) &= (-1)^i \left(\theta_{i+1} - Z_n^{-1} \sum_l a(i, l, n) \sigma^l \right) \\ &= (-1)^i \left(\theta_{i+1} - Z_n^{-1} \sum_l (a(i, l-1, n-2) + a(i, l, n-1)) \sigma^l \right) \\ &> (-1)^i \left(\theta_{i+1} - Z_n^{-1} \sum_l (a(i, l-1, n-2) + a(i+1, l, n-1)) \sigma^l \right) \\ &= (-1)^i \left(\theta_{i+1} - Z_n^{-1} \sum_l a(i+2, l, n) \sigma^l \right) \\ &= (-1)^{i+1} (\theta_{i+2} - \theta_{i+1}). \end{aligned}$$

This completes the proof. \square

References

- [1] Y. Baryshnikov, E. G. Coffman, Jr., and P. Jelenković. Space filling and depletion. *Journal of Applied Probability*, 41(3):691–702, 2004.
- [2] R. R. Boorstyn and A. Kershenbaum. Throughput analysis of multihop packet radio. In *Proc. of ICC*, pages 1361–1366, 1980.
- [3] D. Denteneer, S. C. Borst, P. M. van de Ven, and G. Hiertz. IEEE 802.11s and the philosophers’ problem. *Statistica Neerlandica*, 62(3):283–298, 2008.
- [4] M. Durvy, O. Dousse, and P. Thiran. Modeling the 802.11 protocol under different capture and sensing capabilities. In *Proc. of INFOCOM*, pages 2356–2360, 2007.
- [5] M. Durvy, O. Dousse, and P. Thiran. Self-organization properties of CSMA/CA systems and their consequences on fairness. *IEEE Transactions on Information Theory*, 55(3):931–943, 2009.
- [6] M. Garetto, T. Salonidis, and E. W. Knightly. Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks. In *Proc. of INFOCOM*, pages 1–13, 2006.
- [7] G. Hiertz, Y. Zang, S. Max, T. Junge, E. Wei, D. Denteneer, L. Berlemann, and S. Mangold. IEEE 802.11s - WLAN mesh standardization and high performance extensions. *IEEE Network*, 22(3):12–19, 2008.
- [8] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. In *Proc. of Allerton conference on communication, control, and computing*, pages 1511–1519, 2008.
- [9] F. P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B*, 47(3):379–395, 1985.
- [10] F. P. Kelly. Loss networks. *The Annals of Applied Probability*, 1(3):319–378, 1991.

- [11] L. Kleinrock and F. A. Tobagi. Packet switching in radio channels: part I - carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Transactions on Communications*, 23(12):1400–1416, 1975.
- [12] J. Liu, Y. Yi, A. Proutière, M. Chiang, and H. V. Poor. Convergence and tradeoff of utility-optimal CSMA. <http://arxiv.org/abs/0902.1996> (preprint), 2009.
- [13] E. Pinsky and Y. Yemini. The asymptotic analysis of some packet radio networks. *IEEE Journal on Selected Areas in Communications*, 4(6):938–945, 1986.
- [14] S. Rajagopalan and D. Shah. Distributed algorithm and reversible network. In *Proc. of CISS*, pages 498–502, 2008.
- [15] V. Shneer and P. M. van de Ven. Comparing slotted and non-slotted CSMA. Working paper, 2009.
- [16] P. M. van de Ven, A. J. E. M. Janssen, and J. S. H. van Leeuwen. Optimal tradeoff between exposed and hidden nodes in large wireless networks. submitted, 2009.
- [17] A. Vasan, R. Ramjee, and T. Woo. ECHOS - Enhanced capacity 802.11 hot spots. In *Proc. of INFOCOM*, pages 1562–1572, 2005.
- [18] X. Wang and K. Kar. Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks. In *Proc. of INFOCOM*, pages 23–34, 2005.
- [19] X. Yang and N. Vaidya. On physical carrier sensing in wireless ad hoc networks. In *Proc. of INFOCOM*, pages 2525–2535, 2005.
- [20] M. Zafer and E. Modiano. Blocking probability and channel assignment in wireless networks. *IEEE Transactions on Wireless Communications*, 5(4):869–879, 2006.
- [21] J. Zhu, X. Guo, L. Yang, and W. S. Conner. Leveraging spatial reuse in 802.11 mesh networks with enhanced physical carrier sensing. In *Proc. of ICC*, pages 4004–4011, 2004.