Spatial Fairness in Multi-Channel CSMA Line Networks

Robbe Block, Benny Van Houdt

University of Antwerp - iMinds, Middelheimlaan 1, B2020 - Belgium

Abstract

In this paper we consider a multi-channel random-access carrier-sense multiple access (CSMA) line network with n saturated links, where each link can be active on at most k of the C available channels at any time. Using the product form solution of such a network, we develop fast algorithms to compute the per-link throughputs and use these to study the spatial fairness in such a network.

Recently it was shown that fairness in a single channel CSMA line network can be achieved by means of a simple formula for the activation rates, which depends solely on the number of interfering neighbors. In this paper we show that this formula still achieves fairness in the multi-channel setting under heavy traffic, but no such simple formula seems to exist in general, unless k equals C. On the other hand, numerical experiments show that the fairness index when using the simple single channel formula in the multi-channel setting is close to one and tends to improve as k increases. In other words this simple formula eliminates most of the spatial unfairness in a multi-channel network, especially when k is close to C.

Keywords: CSMA, multi-channel, fairness, mac

1. Introduction

Random-access carrier-sense multiple access (CSMA) networks have received considerable attention over the past few decades and various stochastic models have been developed and studied in great detail, e.g., [6, 10, 9, 18, 19], we refer to [7, 17] for a detailed literature overview. Most of these studies

Email addresses: robbe.block@uantwerpen.be (Robbe Block), benny.vanhoudt@uantwerpen.be (Benny Van Houdt)

Preprint submitted to Performance Evaluation - SI: Valuetools *February 25, 2016*

have focused on the channel throughput, stability, packet delay and/or fairness (between different links) in case of a single channel CSMA network. Studies on multi-channel CSMA networks are far less abundant and include [4, 14], where [4] focuses on throughput optimality and stability and [14] on computing the throughput in large circular and line networks where all the links make use of equal activation rates.

Spatial unfairness in single channel CSMA networks is fairly well understood [9, 20] as links at the border of the network have a restricted neighborhood and thus a higher probability to access the channel. In large line networks these border effects do not propagate inside the network as opposed to more general network topologies. Recently it was shown that spatial unfairness in line networks of limited size can also be eliminated by adapting the activation rates (i.e., mean backoff times) using a simple formula [20]. More specifically, all links achieve the same long run average throughput if the activation rate of link i is of the form $\alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$, for any α, where $\gamma(i)$ is the number of interfering neighbors of link i. Further, these activation rates are the only ones that achieve fairness due to [19].

In this paper, which is an extended version of [1], we study spatial fairness of multi-channel CSMA line networks (of moderate size). We consider a similar network model as in [14], which differs from [3] in the sense that we limit ourselves to line networks, assume that all links have access to all the channels, a link can be active on at most k channels at a time and that interference is the same on each channel. While we can relax some of these assumptions, this might considerably increase the time complexity of the algorithms developed to compute the per-link throughputs.

The following contributions are made in this paper. We start by focusing on the case where each link cannot be active on multiple channels at the same time, that is, $k = 1$. For this case we develop fast algorithms to compute the per-link throughputs, where the time complexity grows linear in the number of links, by exploiting the product form solution of the network. Second, we prove that the simple formula $\alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ to achieve fairness in a single channel network still guarantees fairness in the multi-channel setting under heavy traffic, that is, if α is large. We prove this result first for $k = 1$ and subsequently generalize the result for $1 < k \leq C$. Third, we show that in general a simple formula to achieve spatial fairness in the multi-channel setting which depends only on the number of interfering neighbors does not exist, unless $k = C$ in which case the simple formula for $C = 1$ still achieves fairness. Fourth, by making use of the fast algorithms developed to compute the per-link throughput, we show that while the simple formula $\alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ does not eliminate all spatial unfairness in the multi-channel setting (when $k < C$), it does eliminate most of the unfairness as the Jain's fairness index [12] is typically close to one. Further, this formula results in more fairness as k increases, while the opposite occurs when using equal rates. Finally, we show that the same methodology can be used to compute throughputs in non-equidistant line networks.

Apart from the standard multi-channel CSMA model we also consider a multi-channel CSMA network with channel repacking. Channel repacking indicates that a channel is assigned to a link whenever its backoff timer expires and there is either a channel available or one can be made available by reassigning some of the channels already in use. While this is hard to achieve in practice, especially on a general network topology, we mainly study this variant as we felt that a simple formula to achieve fairness is more likely to exist in this case. The results however indicate that this is not the case and all the findings listed above for the standard CSMA network also apply to the network with repacking.

The paper is structured as follows. In Section 2 we present the model under consideration assuming that a link cannot be simultaneously active on multiple channels, that is, when $k = 1$. The results presented in Sections 3 to 6 are all limited to the setting in which $k = 1$. In these sections we consider systems with and without channel repacking. In Sections 7 and 8 we relax the assumption that $k = 1$ and limit ourselves to the system without repacking. The results in the latter two sections are the main novel contribution with respect to [1] (which was limited to the $k = 1$ case).

2. Model description

Consider a CSMA line network consisting of C channels, n links and an interference range of β, meaning a link cannot be simultaneously active with one of its β left or right neighbors on the same channel. Assume a link can only be active on one channel at a time and packet lengths follow an exponential distribution (with mean 1). It is worth noting that the results presented in this paper remain valid for more general packet length distributions, i.e., phase-type distributions, due to the underlying loss network (see for instance [6, 18, 4] for more details). Backoff timers are assumed to follow an exponential distribution, the average of which is specified further on. Links are assumed to be saturated at all times, that is, each link has at least one packet ready for transmission at any time. We consider two systems: with and without channel repacking.

In case of *channel repacking* we assume that each link maintains a single backoff timer and is assigned a channel when the timer expires provided that a channel is available or *one can be made available* by reassigning some of the already assigned channels.

Without channel repacking we still maintain a single backoff timer per link, but when it expires a channel is selected uniformly at random among the C channels. This channel is assigned in case it is not being used by any of the interfering links, otherwise a new backoff period (of exponential duration) starts. As we assume exponential backoff times, this is equivalent to maintaining C timers, one for each channel.

To emphasize the difference between both systems, assume $C = 2$, $\beta = 1$, link 1 is using channel 1 and link 3 is using channel 2. In this case link 2 can become active with channel repacking (as user 1 simply needs to switch channels), while it cannot without channel repacking.

Due to the exponential nature of the packet lengths and backoff times it is easy to see that the evolution of both systems can be captured by a continuous-time Markov chain. More specifically, for the system *with repacking* all feasible states are given by $\bar{\Omega}_n$ the set of all binary strings of length n such that there are at most C ones in any sequence of $\beta+1$ consecutive bits:

$$\bar{\Omega}_n = \{(\bar{w}_1, \ldots, \bar{w}_n) \in \{0,1\}^n | \sum_{j=k}^{k+\beta} \bar{w}_j \leq C \text{ for } k = 1, \ldots, n-\beta\}.$$

Note that due to repacking it suffices to keep track of the links that are active, meaning there is no need to keep track of the channel ids. Let $\bar{w}_i = 1$ if link i is active on some channel and set $\bar{w}_i = 0$ otherwise. When $\beta < C$, all the links can be active simultaneously in case of repacking and $\bar{\Omega}_n$ is simply the set of all binary strings of length n. Hence, without loss of generality we may assume that $\beta \geq C$ in case of repacking.

Denote ν_i as the activation rate of link i, that is, the mean length of the backoff period equals $1/\nu_i$. With rate 1 the Markov chain makes a transition from a state of the form $(\bar{w}_1, \ldots, \bar{w}_{i-1}, 1, \bar{w}_{i+1}, \ldots, \bar{w}_n)$ to state $(\bar{w}_1, \ldots, \bar{w}_{i-1}, 0, \bar{w}_{i+1}, \ldots, \bar{w}_n)$ as the mean packet length equals 1, while the reverse transition occurs at rate ν_i provided that $(\bar{w}_1, \ldots, \bar{w}_{i-1}, 1, \bar{w}_{i+1}, \ldots, \bar{w}_n)$ belongs to $\bar{\Omega}_n$ as this implies that it is possible to assign a channel to node i (using repacking) when the backoff period ends.

The above Markov chain is reversible and therefore has a product form. To see this, consider a set of n independent $M/M/1/1$ queues, where queue i has arrival rate ν_i and service rate 1. As an M/M/1/1 queue is clearly reversible, so is the union of n such independent queues. The above Markov chain evolves in exactly the same manner as this set of n M/M/1/1 queues, except that the state space is truncated to $\bar{\Omega}_n$. In other words, transitions that would result in leaving the set $\bar{\Omega}_n$ are ignored. The reversibility therefore follows from the fact that a process obtained by truncating the state space of a reversible process is reversible [13]. Moreover, the steady state probabilities of the Markov chain obtained by truncation are identical to the ones of the original process, up to normalization.

Hence, as the steady state probabilities of the i-th M/M/1/1 queue are given by $1/(1+\nu_i)$ and $\nu_i/(1+\nu_i)$, the steady state probability $\bar{\pi}(\bar{w})$ of being in state $\bar{w} = (\bar{w}_1, \ldots, \bar{w}_n) \in \bar{\Omega}_n$ can be expressed as

$$\bar{\pi}(\bar{w}) = \bar{Z}_\nu^{-1} \prod_{i=1}^{n} \nu_i^{\bar{w}_i}, \tag{1}$$

where $\bar{Z}_\nu = \sum_{\bar{w} \in \bar{\Omega}_n} \prod_{i=1}^{n} \nu_i^{\bar{w}_i}$ is the normalizing constant and $\nu = (\nu_1, \ldots, \nu_n)$.

Without repacking we clearly do need to keep track of the ids of the channels in use as they may affect whether a link can become active (as in the example before). Thus the set of all feasible states is given by Ω_n, the set of all strings of length n over the alphabet $\{0, 1, \ldots, C\}$ such that any sequence of $\beta+1$ consecutive symbols does not contain more than one $c > 0$:

$$\begin{aligned}\Omega_n = \{(w_1, \ldots, w_n) \in \{0, 1, \ldots, C\}^n | w_i = 0 \text{ or} \\ w_j \neq w_i \text{ for } j = \max(1, i-\beta), \ldots, \min(n, i+\beta)\}.\end{aligned}$$

Let $C\nu_i$ be the parameter of the exponential distribution of the backoff timer of link i. Using similar arguments as in the system with repacking one finds that this Markov chain is reversible. The difference is that we do not use a simple M/M/1/1 queue as a starting point, but instead consider an M/M/1/1 queue with C types of customers where each customer type has the same arrival rate ν_i and service rate 1. The steady state probability $\pi(w)$ of being in state $w = (w_1, \ldots, w_n) \in \Omega_n$ can therefore be expressed as

$$\pi(w) = Z_\nu^{-1} \prod_{i=1}^{n} \nu_i^{1[w_i > 0]}, \tag{2}$$

where $Z_\nu = \sum_{w\in\Omega_n} \prod_{i=1}^n \nu_i^{1[w_i>0]}$ is the normalizing constant, $\nu = (\nu_1, \ldots, \nu_n)$ and $1[A] = 1$ if A is true and $1[A] = 0$ otherwise.

Throughout the paper we add a bar to a variable or symbol whenever it is related to the system with repacking, unless it concerns a common parameter such as C, β, etc.

3. Matrix expressions for the normalizing constant

In this section we derive a matrix expression for the constants Z_ν and $\bar{Z}_\nu$. Using these expressions we can compute the normalizing constant of the system with repacking in $O(n\binom{\beta+1}{C})$ time and of the system without repacking in $O(n\min(2^\beta, (\beta+1)^C))$ time.

3.1. With Channel Repacking

Theorem 1. *The normalizing constant $\bar{Z}_\nu$ can be written as*

$$\bar{Z}_\nu = \left(\prod_{i=1}^n (1+\nu_i)\right) \bar{P}_n(C, \beta+1, \nu),$$

where $\bar{P}_n(C, \beta+1, \nu)$ is the probability that we have at most C successes in any $\beta+1$ consecutive Bernoulli trials when performing a total of n independent Bernoulli trials where the i-th trial has success probability $p_i = \nu_i/(1+\nu_i)$.

Proof. The result is immediate by noting $\bar{Z}_\nu$ can be written as

$$\bar{Z}_\nu = \frac{\sum_{\bar{w}\in\bar{\Omega}_n} \prod_{i=1}^n \left(\frac{\nu_i}{1+\nu_i}\right)^{1[\bar{w}_i=1]} \left(1-\frac{\nu_i}{1+\nu_i}\right)^{1[\bar{w}_i=0]}}{\prod_{i=1}^n \frac{1}{1+\nu_i}}.$$

□

Probabilities of the type $\bar{P}_n(C, \beta+1, \nu)$ have been studied previously in the area of reliability theory [8], in fact the result in Theorem 2, where $\beta = C$, is equivalent to the method presented in [11] for the so-called consecutive-k-out-of-n:F system. We will generalize this method to any $\beta \geq C$ which implies that our proposed method is also useful to analyze the reliability of a consecutive-k-out-of-m-from-n:F system.

Theorem 2. *When $\beta = C$, we can express $\bar{P}_n(\beta, \beta+1, \nu)$ as*

$$\bar{P}_n(\beta, \beta+1, \nu) = e_1^* \left(\prod_{i=1}^{n} \bar{M}_{\beta,\beta+1}(\nu_i) \right) e,$$

where e_1^* is first row of the size $\beta+1$ identity matrix, e is a column vector of ones and

$$\bar{M}_{\beta,\beta+1}(\nu_i) = \frac{1}{1+\nu_i} \left[\begin{array}{c|ccc} 1 & \nu_i & & \\ \vdots & & \ddots & \\ 1 & & & \nu_i \\ \hline 1 & 0 & \dots & 0 \end{array} \right],$$

for $i = 1, \dots, n$.

Proof. When $\beta = C$ we can have at most β successes in a row. To obtain an expression for $\bar{P}_n(\beta, \beta+1, \nu)$ we construct a time-inhomogeneous Markov chain with $\beta+1$ transient, labeled 0 to β, and one absorbing state. We start in state 0 and the i-th transition corresponds to performing the i-th Bernoulli trail. The $\beta+1$ transient states keep track of the number of consecutive successes, meaning a success increases the state by 1, while a failure resets the state to 0. If a success occurs in state β, meaning we have more than β successes in a row, we move to the absorbing state. The probability $\bar{P}_n(\beta, \beta+1, \nu)$ can be expressed as the probability of being in a transient state at time n. $\square$

This theorem allows us to compute $\bar{Z}_\nu$ in $O(n\beta)$ time when $\beta = C$.

In order to generalize the previous idea, we introduce the matrices $\bar{M}_{C,\beta+1}(\nu_i)$ of size $\sum_{k=0}^{C} \binom{\beta-C+k}{k}$. The rows and columns of $\bar{M}_{C,\beta+1}(\nu_i)$ are labeled by the strings $\bar{w} \in \bar{\Omega}_{C,\beta}$ with

$$\bar{\Omega}_{C,\beta} = \cup_{k=0}^{C} \{\bar{w} \in \{0,1\}^{\beta-C+k} | \sum_i \bar{w}_i = k\}.$$

Note the length of $\bar{w} \in \bar{\Omega}_{C,\beta}$ is limited by β. Let $l(\bar{w})$ be the length of $\bar{w}$ and $z(\bar{w})$ the position of the first zero (which exists for $\beta > C$), e.g., $l((1,1,0,1,0,1)) = 6$ and $z((1,1,0,1,0,1)) = 3$, then

$$(1+\nu_i) \left(\bar{M}_{C,\beta+1}(\nu_i) \right)_{\bar{w},\bar{w}'} = \begin{cases} 1 & \bar{w}' = (\bar{w}_{z(\bar{w})+1}, \dots, \bar{w}_{l(\bar{w})}, 0), \\ \nu_i & l(\bar{w}) < \beta, \bar{w}' = (\bar{w}_1, \dots, \bar{w}_{l(\bar{w})}, 1), \\ 0 & \textit{otherwise}. \end{cases}$$

Theorem 3. *When $\beta > C \geq 1$, we can express $\bar{P}_n(C, \beta+1, \nu)$ as*

$$\bar{P}_n(C, \beta+1, \nu) = e_1^* \left(\prod_{i=1}^{n} \bar{M}_{C,\beta+1}(\nu_i) \right) e,$$

where e_1^* is first row of the identity matrix, e is a column vector of ones. Further, the matrices $\bar{M}_{C,\beta+1}(\nu_i)$ are of size $\binom{\beta+1}{C}$.

Proof. We rely on a time-inhomogeneous Markov chain as before and label the transient states by the strings in $\bar{\Omega}_{C,\beta}$. The binary string $\bar{w} \in \bar{\Omega}_{C,\beta}$ reflects the outcome of all the previous trials that occurred after the $(\beta+1-C)$-last failure. A new success is only allowed if the $(\beta+1-C)$-last failure occurred strictly less than β trials ago and simply adds a 1 to the state. If a failure occurs we can forget about the outcome of all the trials up until and including the first 0 in $\bar{w}$, while adding a 0.

It is easy to see that $|\bar{\Omega}_{C,\beta}| = \binom{\beta+1}{C}$ as

$$\sum_{k=0}^{C} \binom{\beta - C + k}{k} = \sum_{k=0}^{C} \binom{\beta - k}{C - k} = \binom{\beta+1}{C},$$

where the latter equality follows from repeatedly applying $\binom{n+1}{k} = \binom{n}{k} + \binom{n}{k-1}$. $\square$

Note, each row of $\bar{M}_{C,\beta+1}(\nu_i)$ contains at most 2 nonzero entries, meaning that multiplying $\bar{M}_{C,\beta+1}(\nu_i)$ with a column vector requires at most $2|\bar{\Omega}_{C,\beta}|$ floating point operations, which means the time complexity to compute $\bar{Z}_\nu$ is bounded by $O(n\binom{\beta+1}{C})$.

Remark. When $\nu_i = \sigma$, for $i = 1, \ldots, n$, it is also possible to express $\bar{Z}_n$ as

$$\bar{Z}_n = \sum_{j=0}^{n} \bar{B}_n(\beta+1, C, j)\sigma^j, \tag{3}$$

where $\bar{B}_n(m, C, j)$ denotes the number of binary strings of length n with exactly j ones such that no m consecutive bits contain more than C ones. When $\beta = C$, we are thus interested in the number of binary strings of length n with at most C consecutive ones. As shown in [2, Theorem 3.3], $B_n(C+1, C, j)$ can be expressed as

$$\bar{B}_n(C+1, C, j) = \binom{n-j+1}{j}_C,$$

where $\binom{n}{i}_s$ is the generalized binomial coefficient defined by the recursion:

$$\binom{n+1}{i}_s = \sum_{k=0}^{s-1} \binom{n}{i-k}_s,$$

and $\binom{n}{0}_s = 1$. Note, when $s = 1$ these are the usual binomial coefficients, while for $s > 1$ they can be expressed in terms of the usual ones [5, p.19] by

$$\binom{n}{i}_s = \sum_{k=0}^{\lfloor i/s \rfloor} (-1)^k \binom{n}{k} \binom{n+i-sk-1}{n-1}.$$

For $\beta > C$ counting these strings is equivalent to solving the so-called generalized birthday problem. Rather involved closed form expressions for $\bar{B}_n(m, C, j)$ were derived in [16] when $j/2 < C$ and in [15, Theorem 1] for the general case. The latter however are expressed as a large sum of determinants and therefore does not result in an efficient manner to compute $\bar{B}_n(m, C, j)$.

3.2. Without Channel Repacking

Consider a $(C+1)$-sided coin with outcomes $0, 1, \dots, C$ and assume that the probability of having outcome c, for $c \in \{1, \dots, C\}$, equals p, while the outcome 0 has the remaining probability $1 - Cp$, for some $p \in (0, 1/C)$. Let $\mathcal{S}_\beta$ be the set of all binary strings of length β that contain at most C ones. To define the set of matrices $M_{C,\beta+1}(\nu_i)$ of size $|\mathcal{S}_\beta| = \sum_{k=0}^{\min(C,\beta)} \binom{\beta}{k} \leq 2^\beta$, we label the rows and columns of $M_{C,\beta+1}(\nu_i)$ by the strings in $\mathcal{S}_\beta$. For $z \in \mathcal{S}_\beta$, let $n(z)$ be the value of the binary number represented by z, e.g., $n((0,1,0,1)) = 5$, and define

$$(1 + C\nu_i) \left(M_{C,\beta+1}(\nu_i)\right)_{z,z'} = \begin{cases} 1 & n(z) < 2^{\beta-1}, n(z') = 2n(z), \\ 1 & n(z) \geq 2^{\beta-1}, n(z') = 2n(z) - 2^\beta, \\ \nu_i(C-k) & \sum_{i=1}^{\beta} z_i = k, n(z) < 2^{\beta-1}, \\ & n(z') = 2n(z) + 1, \\ \nu_i(C-k) & \sum_{i=1}^{\beta} z_i = k, n(z) \geq 2^{\beta-1}, \\ & n(z') = 2n(z) - 2^\beta + 1, \\ 0 & \textit{otherwise}. \end{cases} \tag{4}$$

The normalization constant Z_ν can be expressed as follows using the matrices $M_{C,\beta+1}(\nu_i)$:

Theorem 4. *The normalizing constant Z_ν can be written as*

$$Z_\nu = \left(\prod_{i=1}^{n} (1 + C\nu_i) \right) P_n(C, \beta + 1, \nu),$$

and

$$P_n(C, \beta + 1, \nu) = e_1^* \left(\prod_{i=1}^{n} M_{C,\beta+1}(\nu_i) \right) e,$$

where e_1^* is first row of the size $|\mathcal{S}_\beta|$ identity matrix, e is a column vector of ones.

Proof. The proof is similar to the proof of Theorem 1 by noting that $P_n(C, \beta + 1, \nu)$ is the probability that when flipping n coins with $C + 1$ sides, where $p = \nu_i/(1 + C\nu_i)$ for coin i, no sequence of $\beta + 1$ consecutive flips results in two or more identical outcomes equal to some $c > 0$.

To express $P_n(C, \beta + 1, \nu)$ we construct a time-inhomogeneous Markov chain (as in the proof of Theorem 3) with $|\mathcal{S}_\beta|$ transient, labeled $z \in \mathcal{S}_\beta$, and one absorbing state. We start in state $(0, \ldots, 0)$ and the i-th transition corresponds to performing the i-th $(C + 1)$-sided coin flip. The transient states keep track of the position of the outcomes $c > 0$ in the last β coin flips. If we are in transient state z and the outcome of coin flip i is 0, we simply shift z to the left, drop the leading bit and add a zero to the right. If the outcome is $c > 0$ and $\sum_{i=1}^{\beta} z_i = k$ there is a probability $(C - k)/C$ that the outcome differs from the k outcomes with $c > 0$ in the last β coin flips. If the outcome differs, we shift z to the left, drop the leading bit and add a one to the right, otherwise we jump to the absorbing state. The probability $P_n(C, \beta + 1, \nu)$ can be expressed as the probability of being in a transient state at time n. □

Note, each row of $M_{C,\beta+1}(\nu_i)$ contains at most 2 nonzero entries, meaning multiplying $M_{C,\beta+1}(\nu_i)$ with a column vector requires at most $2|\mathcal{S}_\beta|$ floating point operations, which means the time complexity to compute Z_ν is bounded by $O(n \min(2^\beta, (\beta + 1)^C))$ as $\sum_{k=0}^{\beta} \binom{\beta}{k} = 2^\beta$ and $\sum_{k=0}^{C} \binom{\beta}{k} \leq (\beta + 1)^C$.

4. Computing Link Throughputs

In case of channel repacking, denote the long-run average throughput of link j as $\bar{\theta}_j(\nu)$. It corresponds to the long-run fraction of time that link j is active on some channel. To express $\bar{\theta}_j(\nu)$ define the matrices $\bar{N}_{C,\beta+1}(\nu_i)$ as

$$(1+\nu_i)\left(\bar{N}_{C,\beta+1}(\nu_i)\right)_{w,w'} = \begin{cases} \nu_i & l(w) < \beta, w' = (w_1, \dots, w_{l(w)}, 1), \\ 0 & otherwise, \end{cases}$$

i.e., they are obtained by setting all the entries of $\bar{M}_{C,\beta+1}(\nu_i)$ that correspond to a failure to zero.

Theorem 5. *The throughput $\bar{\theta}_j(\nu)$ of node j can be computed as*

$$\bar{\theta}_j(\nu) = \frac{\bar{P}_n^{(j)}(C, \beta+1, \nu)}{\bar{P}_n(C, \beta+1, \nu)},$$

where

$$\bar{P}_n^{(j)}(C, \beta+1, \nu) = e_1^* \left(\prod_{i=1}^{j-1} \bar{M}_{C,\beta+1}(\nu_i) \right) \bar{N}_{C,\beta+1}(\nu_j) \left(\prod_{i=j+1}^{C} \bar{M}_{C,\beta+1}(\nu_i) \right) e.$$

Proof. Using the expression for the steady state we get

$$\bar{\theta}_j(\nu) = \bar{Z}_\nu^{-1} \sum_{\bar{w} \in \Omega} \prod_{i=1}^{n} \nu_i^{\bar{w}_i} 1[\bar{w}_j = 1].$$

The result now follows from Theorem 1 and by noting that $\bar{P}_n^{(j)}(C, \beta+1, \nu)$ represents the probability that we have at most C successes in any $\beta+1$ consecutive Bernoulli trials when performing a total of n independent Bernoulli trials where the i-th trial has success probability $p_i = \nu_i/(1+\nu_i)$ and the j-th trial is successful. $\square$

By first computing the vectors $e_1^* \prod_{i=1}^{j-1} \bar{M}_{C,\beta+1}(\nu_i)$ as well as the vectors $\prod_{i=j+1}^{C} \bar{M}_{C,\beta+1}(\nu_i) e$, for $j = 1, \dots, n$, we can compute the vector of throughputs $\bar{\theta}(\nu) = (\bar{\theta}_1(\nu), \dots, \bar{\theta}_n(\nu))$ in $O(n\binom{\beta+1}{C})$ time.

For the system without channel repacking we can proceed in exactly the same way to compute the vector $\theta(\nu) = (\theta_1(\nu), \dots, \theta_n(\nu))$ of channel throughputs, by defining the matrices $N_{C,\beta+1}(\nu_i)$ as

$$(1+C\nu_i)\left(N_{C,\beta+1}(\nu_i)\right)_{z,z'} = \begin{cases} \nu_i(C-k) & \sum_{i=1}^{\beta} z_i = k, n(z) < 2^{\beta-1}, \\ & n(z') = 2n(z)+1, \\ \nu_i(C-k) & \sum_{i=1}^{\beta} z_i = k, n(z) \geq 2^{\beta-1}, \\ & n(z') = 2n(z) - 2^\beta + 1, \\ 0 & otherwise. \end{cases} \quad (5)$$

i.e., they are obtained by setting all the entries of $M_{C,\beta+1}(\nu_i)$ that correspond to outcome 0 to zero. The time complexity to determine the vector of throughputs therefore equals $O(n \min(2^\beta, (\beta+1)^C))$ (by first computing the vectors $e_1^* \prod_{i=1}^{j-1} M_{C,\beta+1}(\nu_i)$ and $\prod_{i=j+1}^{C} M_{C,\beta+1}(\nu_i)e$, for $j = 1, \ldots, n$).

Remark. In [14] the authors also propose the use of a matrix product to compute the throughput $\theta_j(\nu)$ of link j, but they focus on large networks with equal activation rates. Further, the matrices used are considerably larger than the ones used in our approach. For instance, for $C = 3$ and $\beta = 2$ matrices of size 13 are used, while in our case size $\sum_{k=0}^{2} \binom{2}{k} = 4$ suffices.

5. Fairness

Let $\gamma(i)$ be the number of links that interfere with link i. The main result in [20] showed that in case of a single channel, i.e., $C = 1$, fairness can be achieved in a line network consisting of n links if $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$, for $i = 1, \ldots, n$, for any choice of α. The following section indicates that this choice of ν_i still guarantees fairness in case of multiple channels, i.e., $C \geq 1$, under heavy traffic with and without repacking.

5.1. Heavy traffic

We start by considering the case where the number of channels C is at most $\beta + 1$.

Theorem 6. *Let $n > \beta \geq 1$, $C \leq \beta + 1$ and set $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$, then*

$$\lim_{\alpha \to \infty} \theta_j(\nu) = \lim_{\alpha \to \infty} \bar{\theta}_j(\nu) = \frac{C}{\beta + 1},$$

for $j = 1, \ldots, n$

Proof. We restrict ourselves to the system *without* channel repacking. The argument for the system with repacking proceeds similarly. When α becomes large $\nu_i \approx \alpha^{\gamma(i)-\gamma(1)+1}$ and the product form in (2) implies that the Markov chain spends most of its time in the states w that maximize

$$val(w) \stackrel{def}{=} \sum_{i=1}^{n} (\gamma(i) - \gamma(1) + 1) 1[w_i > 0].$$

We will argue that there are $C!\binom{\beta+1}{C}$ states w for which $val(w)$ is maximized and that each $j \in \{1, \dots, n\}$ is active in exactly $C!\binom{\beta}{C-1}$ of these states. This results in a throughput of $\binom{\beta}{C-1}/\binom{\beta+1}{C} = C/(\beta+1)$ for each link.

Define the following subset of Ω_n of size $C!\binom{\beta+1}{C}$:

$$\mathcal{M}_n = \{w \in \Omega_n | \sum_{i=1}^{\beta+1} 1[w_i > 0] = C, w_j = w_{j-(\beta+1)}, j > \beta + 1\}.$$

Note for $w \in \mathcal{M}_n$ any set of $\beta + 1$ consecutive elements contains C *distinct* positive elements. Further, $w_j > 0$ in exactly $C!\binom{\beta}{C-1}$ states $w \in \mathcal{M}_n$, as there are $\binom{\beta}{C-1}$ ways to select the remaining $C-1$ positive elements in the first $\beta+1$ positions. To complete the proof we now show that $val(w) = (n-\beta)C$ for $w \in \mathcal{M}_n$ and $val(w) < (n-\beta)C$ for $w \notin \mathcal{M}_n$.

For $n = \beta + 1$ it is clear that $val(w) = C$ for $w \in \mathcal{M}_n$ as $\gamma(i) = \beta$ for all $i \in \{1, \dots, n\}$ and $val(w)$ is therefore equal to the number of ones in w. If we add a link to a line network of n links, we see that $\gamma(n-\beta+1), \dots, \gamma(n)$ increase by one, while $\gamma(i)$ remains identical for $i \leq n-\beta$ and $\gamma(n+1) = \gamma(1)$. Further, any state w can have at most C positive elements in the last $\beta + 1$ positions, thus

$$val(w_1, \dots, w_{n+1}) \leq val(w_1, \dots, w_n) + C.$$

Hence, $val(w) \leq (n-\beta)C$ for all $w \in \Omega_n$. When $w \in \mathcal{M}_n$ the last $\beta + 1$ positions of $w = (w_1, \dots, w_n)$ contain exactly C positive elements and $w_{n+1} = w_{n-\beta}$, thus each time we add an element to $w \in \Omega_n$ such that $w_{n+1} = w_{n-\beta}$, its value increases by C. This implies that $val(w) = (n-\beta)C$ for $w \in \mathcal{M}_n$.

Assume $w \notin \mathcal{M}_n$ and w contains less than C positive elements in the first $\beta+1$ positions. In this case $val(w) < (n-\beta)C$ as $val(w_1, \dots, w_{\beta+1}) < C$ and adding a single element can only increase the value by C. If w does contain exactly C positive elements in the first $\beta+1$ positions, let j be smallest index such that $w_j \neq w_{j-(\beta+1)}$. In this case we must have $w_j = 0$ and $w_{j-(\beta+1)} > 0$, as we otherwise get $C + 1$ positive elements in $(w_{j-\beta}, \dots, w_j)$ and thus a repetition of the same positive value within a set of $\beta+1$ consecutive values. Thus, when adding w_j, the value of $(w_1, \dots, w_{j-1})$ increases by $C-1$ instead of C, which implies that $val(w)$ must be less than $(n-\beta)C$. $\square$

When $C \geq \beta + 1$ the throughput $\theta_j(\nu)$ approaches one even for $\nu_i = \alpha$ as α tends to infinity, as Ω_n contains states where all the links are active on some channel and these will dominate as α becomes large.

5.2. Intermediate traffic

For intermediate rates and with repacking, setting $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ does not guarantee fairness except when $C = 1$ as will become clear from the following proposition:

Proposition 1. *Let $\beta = n-2$ and let $\phi = \nu_2 = \ldots = \nu_{n-1}$, then fairness is achieved in a system* **with repacking** *if and only if*

$$\nu_1 = \nu_n = \frac{1}{2}\left(\sqrt{(1-\phi\bar{S}_2(\phi)/\bar{S}_1(\phi))^2 + 4\phi} - [1-\phi\bar{S}_2(\phi)/\bar{S}_1(\phi)]\right), \tag{6}$$

with $\bar{S}_k(y) = \sum_{i=0}^{C-k}\binom{n-3}{i}y^i$ for $k \geq C$ and $\bar{S}_k(y) = 0$ for $k > C$.

Proof. Clearly $\nu_1 = \nu_n$ due to the symmetry of the system, while $\phi = \nu_2 = \ldots = \nu_{n-1}$ implies that $\bar{Z}_\nu^{-1}\bar{\theta}_1(\nu) = \bar{Z}_\nu^{-1}\bar{\theta}_i(\nu)$, for $i = 2,\ldots,n-1$, can be written as

$$\nu_1(1+\nu_1)\sum_{i=0}^{C-1}\binom{n-2}{i}\phi^i = \phi\left((1+\nu_1)^2\sum_{i=0}^{C-2}\binom{n-3}{i}\phi^i + \binom{n-3}{C-1}\phi^{C-1}\right),$$

as link 1 can be simultaneously active with link n and at most $C-1$ intermediate links and if an intermediate link i is active with at most $C-2$ other intermediate links both link 1 and n can be active, while they must both be silent if there are $C-1$ other active intermediate links. In other words, ν_1 is the positive solution of a quadratic equation and (6) follows by noting that $\sum_{i=0}^{C-1}\binom{n-2}{i}\phi^i - \phi\sum_{i=0}^{C-2}\binom{n-3}{i}\phi^i = \sum_{i=0}^{C-1}\binom{n-3}{i}\phi^i$. □

The next proposition establishes a similar result for the system without repacking:

Proposition 2. *Let $\beta = n-2$ and let $\phi = \nu_2 = \ldots = \nu_{n-1}$, then fairness is achieved in a system* **without repacking** *if and only if*

$$\nu_1 = \nu_n = \frac{\sqrt{(1+\phi S_2(\phi)/S_1(\phi))^2 + 4\phi} - [1-\phi S_2(\phi)/S_1(\phi)]}{2(1+S_2(\phi)/S_1(\phi))}, \tag{7}$$

with $S_k(y) = \sum_{i=0}^{C-k}\frac{C!}{(C-k-i)!}\binom{n-3}{i}y^i$ for $k \geq C$ and $S_k(y) = 0$ for $k > C$.

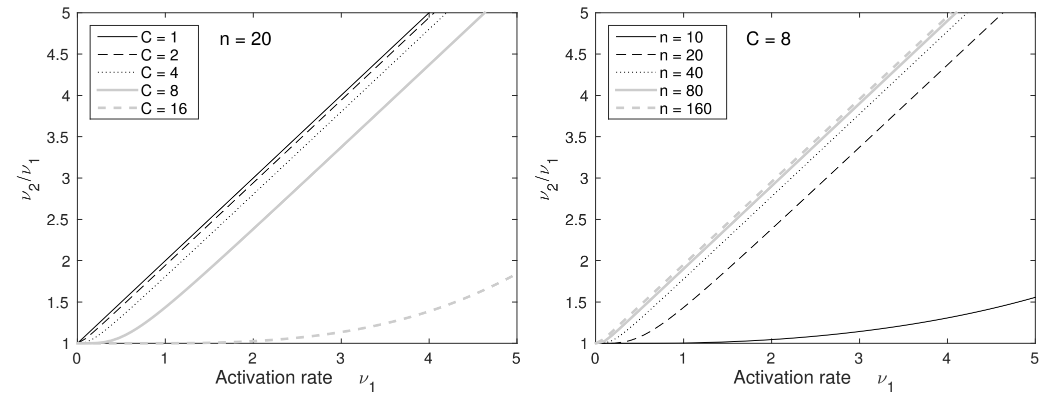

Figure 1: With repacking: ratio of ν_2 and ν_1 to achieve fairness as a function of ν_1 when $\beta = n - 2$. For $C > 1$ channel this ratio is no longer a linear function of ν_1 and depends on n.

Proof. Clearly $\nu_1 = \nu_n$ due to the symmetry of the system, while $\phi = \nu_2 = \ldots = \nu_{n-1}$ implies that $Z_\nu^{-1}\theta_1(\nu) = Z_\nu^{-1}\theta_i(\nu)$, for $i = 2, \ldots, n-1$, can be written as

$$\nu_1(1+\nu_1)C\sum_{i=0}^{C-1} i!\binom{C-1}{i}\binom{n-2}{i}\phi^i +$$
$$\nu_1^2 C(C-1)\sum_{i=0}^{C-2} i!\binom{C-2}{i}\binom{n-2}{i}\phi^i =$$
$$\phi C\sum_{i=0}^{C-1} i!\binom{C-1}{i}\binom{n-3}{i}\phi^i +$$
$$\phi(\nu_1^2+2\nu_1)C(C-1)\sum_{i=0}^{C-2} i!\binom{C-2}{i}\binom{n-3}{i}\phi^i +$$
$$\phi\nu_1^2 C(C-1)(C-2)\sum_{i=0}^{C-3} i!\binom{C-3}{i}\binom{n-3}{i}\phi^i,$$

as link 1 can be simultaneously active with link n and at most $C-1$ or $C-2$ intermediate links depending on whether link 1 and n use the same or a different channel. If a intermediate link i is active with at most $C-2$ other intermediate links both link 1 and n can be active (either on the same or a different channel), while they must both be silent if there are $C-1$

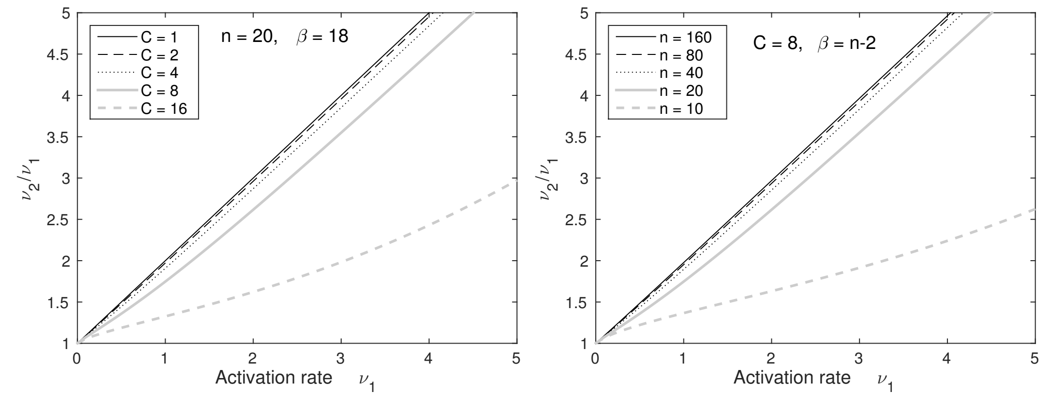

Figure 2: Without repacking: ratio of ν_2 and ν_1 to achieve fairness as a function of ν_1 when $\beta = n-2$. For $C > 1$ channel this ratio is no longer a linear function of ν_1 and depends on n.

other active intermediate links. In other words, ν_1 is the positive solution of a quadratic equation and (7) follows by noting that

$$\sum_{i=0}^{C-k} \frac{C!}{(C-k-i)!}\binom{n-2}{i}\phi^i - \phi \sum_{i=0}^{C-k-1} \frac{C!}{(C-k-1-i)!}\binom{n-3}{i}\phi^i = \sum_{i=0}^{C-k} \frac{C!}{(C-k-i)!}\binom{n-3}{i}\phi^i,$$

due to Pascal's triangle identity. □

When $C = 1$, both results reduce to $\nu_1 = (\sqrt{1+4\phi} - 1)/2$, meaning $\phi = \nu_2 = \nu_1(1+\nu_1)$ and $\nu_2/\nu_1 = (1+\nu_1)$. Figures 1 and 2 indicate that when $C > 1$ the ratio ν_2/ν_1 needed to achieve fairness is no longer linear in ν_1 and this ratio depends on the network size n. The results do seem to indicate that if $n \gg C$ the fairness ratio is close to $(1+\nu_1)$, which is the fairness ratio for $C = 1$.

6. Numerical Results

In this section we investigate the impact of having multiple channels on the fairness in the network. We limit ourselves to the system without channel repacking as this is the most relevant from a practical point of view and numerical experiments not shown here confirm that the main conclusions

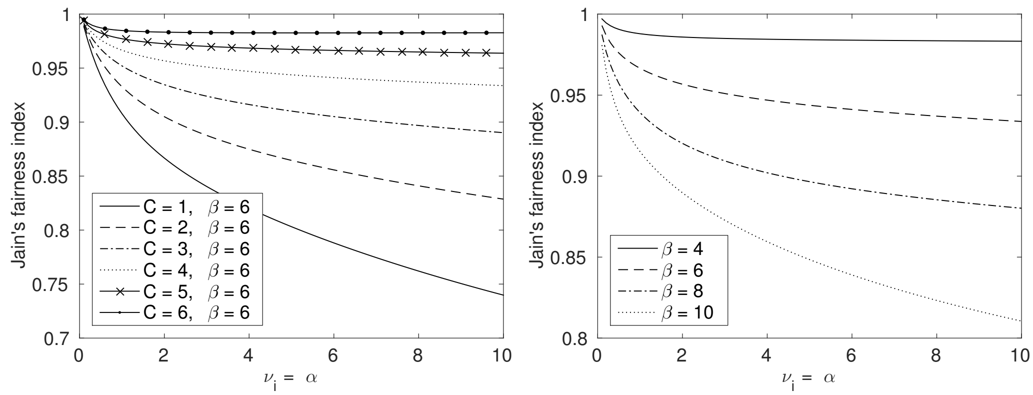

Figure 3: Without repacking: fairness index as a function of the activation rate $\nu_i = \alpha$ with $n = 40$ and either C or β fixed.

for the system with repacking are in fact similar. To express the fairness of the system we make use of Jain's well-known fairness index [12], which is computed as

$$\mathcal{J}(\theta(\nu)) = \frac{\left(\sum_{j=1}^{n} \theta_j(\nu)\right)^2}{n \sum_{j=1}^{n} \theta_j(\nu)^2}.$$

We start by considering the case where all the links make use of the same activation rate, that is, $\nu_i = \alpha$ for $i = 1, \ldots, n$.

Figure 3 depicts the fairness index in a line network consisting of $n = 40$ links as a function of the activation rate α for different combinations of C and β. This figure demonstrates that fairness improves as the number of channels C increases with β fixed, while increasing β with C fixed increases unfairness. This is quite expected as decreasing C or increasing β implies that a link is more severely influenced by the activity of its neighboring links. We also note that the unfairness is quite severe as the index is well below one (unless C is close to β) and worsens as links become more aggressive, i.e., α increases.

We now repeat the same experiment, but instead of using equal rates we set the activation rate $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$, which guarantees fairness in heavy traffic as proven in Theorem 6 and fairness in general when $C = 1$ due to [20].

Figure 4 depicts the fairness index in a line network consisting of $n = 40$ links as a function of the parameter α. The first thing to note is that the index is now very close to one (above 0.995), meaning the activation rates

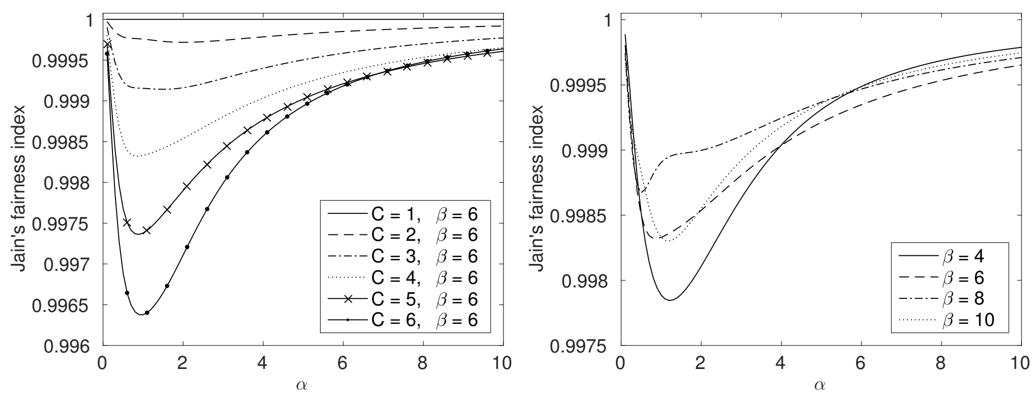

Figure 4: Without repacking: fairness index as a function of α with the activation rate $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ with $n = 40$ and either C or β fixed.

that guarantee fairness in the single channel setup cause only a very limited degree of unfairness in the multi-channel setup. We further note that as opposed to the equal rate case, fairness slightly decreases with the number of available channels C in most cases. Further, when C is fixed, having more or less interference, that is, increasing β, may result in either an increase or a decrease in fairness depending on the value of α.

Figure 5 further demonstrates that setting ν_i equal to $\alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ results in a drastic improvement of the network fairness compared to using fixed activation rates. The fairness index in this particular case increases from 0.8583 to 0.9998. Note that the choice $\nu_i = 0.5(1.5)^6 \approx 5.7$ corresponds to the rate of the links in the middle of the network when $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ and $\alpha = 0.5$.

7. Multi-reception capabilities

In this section we discuss a generalization of the earlier results in which we permit each link to be active on up to $k \leq C$ channels at a time. As in the previous section, we focus on the system without repacking capabilities, but we would like to stress that similar methods can be used to calculate the throughput of a system with multi-reception capabilities and repacking. The method to analyze such networks with repacking is in fact a more straightforward adaptation of the methods presented in Sections 3 and 4.

We first argue that introducing multiple reception capabilities keeps the product form intact and subsequently show how the throughput can be com-

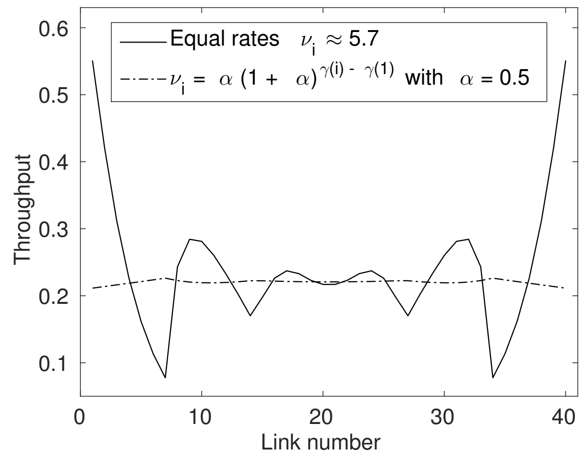

Figure 5: Without repacking: comparison of per-link throughputs $\theta_j(\nu)$ between equal activation rates and setting $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$, when $\alpha = 0.5$, $n = 40$, $C = 2$ and $\beta = 6$.

puted. We assume link i maintains a backoff timer per channel with the same rate ν_i. In Section 2, we noted that this was equivalent to a single backoff timer with rate $C\nu_i$ and selecting a channel uniformly at random. This equivalence only holds when $k = 1$. When $k > 1$, the rate of the single backoff timer should decrease as a function of the number of channels used by the link. Indeed, assuming link i is using $s(i) < k$ channels, it uses $C - s(i)$ exponentially distributed timers, leading to a rate of $(C - s(i))\nu_i$ if we would replace them by a single timer.

7.1. Model description

Let u be a binary $C \times n$-matrix, its elements denoted as $u_{i,j}$. Element $u_{i,j}$ equals one if and only if link j is active on channel i. Define $s(u,j) = \sum_{x=1}^{C} u_{x,j}$, i.e. the number of active channels of link j, and $d(u,i,j) = \sum_{x=j}^{j+\beta} u_{i,x}$, i.e. the number of links that are active on channel i when considering links j until $j+\beta$. The restriction imposed by the interference range results in the following state space:

$$\Omega_{C,k}^{n} = \{u | d(u,i,j) \leq 1, s(u,m) \leq k, 1 \leq i \leq C, 1 \leq j \leq n-\beta, 1 \leq m \leq n\},$$

meaning that in every sub-row of $\beta + 1$ entries, only one entry is allowed to be 1 and all column sums must be less than or equal to k.

The transition rates from state $u \in \Omega^n_{C,k}$ are as follows. If $u_{i,j} = 1$, entry $u_{i,j}$ becomes 0 at rate one. If $u_{i,j} = 0$, entry $u_{i,j}$ becomes 1 with rate ν_i provided that the matrix obtained by setting $u_{i,j} = 1$ in u belongs to $\Omega^n_{C,k}$. The steady state of this Markov chain has a simple product form. This can be noted using a similar argument as in Section 2. Indeed, this particular model can be obtained by truncating the state space of a set of Cn independent $M/M/1/1$ queues organized in a $C \times n$ grid, where channel i is used by node j if queue (i, j) is currently serving a customer. The service rate is one in all of these queues, while there are exactly C queues with arrival rate ν_i. Therefore, we have

$$\pi(u) = Z_\nu^{-1} \prod_{j=1}^{n} \prod_{i=1}^{C} \nu_j{}^{u_{i,j}}. \tag{8}$$

7.2. Throughput calculation

Using a similar approach as in Section 3.2, we could construct matrices $M^{\beta+1}_{C,k}(\nu_j)$ of size $|\Omega^\beta_{C,k}|$ to calculate the normalizing constant. However, as $|\Omega^\beta_{C,k}|$ can be quite large, this approach would be very time and memory consuming. Instead, we argue that a more compact representation of the states can be used to compute the throughput. For this purpose we define the set

$$\hat{\Omega}^n_{C,k} = \{z \in \{0, 1, \dots, k\}^n | \sum_{x=0}^{\beta} z_{j+x} \leq C, j = 1, \dots, n - \beta\}.$$

Given $z = (z_1, \dots, z_n)$, we define $z_{l:m} = (z_l, \dots, z_m)$ and $f(z) = \sum_{i=1}^n z_i$. Furthermore, we assume $z_l = 0$ for $l < 1$ or $l > n$.

If the β left neighbors of link j have $f(z_{j-\beta:j-1})$ channels in use, link j still has $C - f(z_{j-\beta:j-1})$ channels out of which it can choose z_j channels. In other words, if we represent the state of the left neighbors of link j using $z_{j-\beta:j}$, there are $\binom{C-f(z_{j-\beta:j-1})}{z_j}$ different possibilities for link j to be active on z_j channels. This leads to the following definition of the $|\hat{\Omega}^\beta_{C,k}|$ sized matrices

$$(\hat{M}^{\beta+1}_{C,k}(\nu_j))_{z,z'} = \begin{cases} \nu_j^l \binom{C-f(z)}{l} & f(z) + l \leq C, z' = (z_2, \dots, z_\beta, l), \\ 0 & otherwise, \end{cases} \tag{9}$$

with $z, z' \in \hat{\Omega}^\beta_{C,k}$. These matrices can be used to calculate the normalizing constant $Z_{\nu,k}$ as indicated by the following theorem:

Theorem 7. *The normalization constant of the system in which each link is allowed to transmit on up to k channels simultaneously can be expressed as*

$$Z_{\nu,k} = e_1^* \left(\prod_{j=1}^{n} \hat{M}_{C,k}^{\beta+1}(\nu_j) \right) e.$$

Proof. By making use of (8) one finds by induction on i that

$$\left(e_1^* \prod_{j=1}^{i} \hat{M}_{C,k}^{\beta+1}(\nu_j) \right)_z = \sum_{\substack{u \in \Omega_{C,k}^n, (s(u,i-\beta+1),\ldots,s(u,i))=z \\ (s(u,i+1),\ldots,s(u,n))=0}} Z_{\nu,k}\pi(u),$$

for any $z \in \hat{\Omega}_{C,k}^{\beta}$ and $1 \le i \le n$. □

Note that each row of the matrices $\hat{M}_{C,k}^{\beta+1}(\nu_j)$ has at most $k+1$ nonzero elements, thus multiplying it with a column vector takes $(k+1)|\hat{\Omega}_{C,k}^{\beta}|$ time. This yields an overall complexity of $O(n(k+1)|\hat{\Omega}_{C,k}^{\beta}|)$, as opposed to $O(n|\hat{\Omega}_{C,k}^{n}|)$ when summing over the entire state space.

For $k = 1$ the throughput of a link was defined as the long-run fraction of time in which the link was active on some channel. When $k > 1$ the throughput should also take the number of channels on which a link is active into account. As such

$$\hat{\theta}_{k,j}(\nu) = \sum_{u \in \Omega_{C,k}^n} s(u,j)\pi(u).$$

For the actual calculation of the throughput $\hat{\theta}_{k,j}(\nu)$, we define

$$(\hat{N}_{C,k}^{\beta+1}(\nu_j))_{z,z'} = \begin{cases} l\nu_j^l \binom{C-f(z)}{l} & f(z)+l \le C, z' = (z_2,\ldots,z_\beta,l), \\ 0 & \textit{otherwise}, \end{cases} \tag{10}$$

which is used in the following theorem:

Theorem 8. *The throughput $\hat{\theta}_{k,j}(\nu)$ of node j can be computed as*

$$\hat{\theta}_{k,j}(\nu) = \frac{e_1^* \left(\prod_{i=1}^{j-1} \hat{M}_{C,k}^{\beta+1}(\nu_i) \right) \hat{N}_{C,k}^{\beta+1}(\nu_i) \left(\prod_{i=j+1}^{n} \hat{M}_{C,k}^{\beta+1}(\nu_i) \right) e}{e_1^* \left(\prod_{i=1}^{n} \hat{M}_{C,k}^{\beta+1}(\nu_i) \right) e}.$$

Proof. The result is a consequence of Theorem 7 and the fact that we multiply $\pi(u)$ by $s(u,j)$. $\square$

This method can also be used when the throughput does not scale linearly with the number of channels. Indeed, if we define $\vartheta_j(z_j)$ as the instantaneous throughput of link j given that it uses z_j channels, we merely need to use the matrices

$$(\bar{N}_{C,k}^{\beta+1}(\nu_j))_{z,z'} = \begin{cases} \vartheta_j(l)\nu_j^l\binom{C-f(z)}{l} & f(z)+l \leq C, z' = (z_2,\ldots,z_\beta,l), \\ 0 & otherwise, \end{cases}$$

instead of $\hat{N}_{C,k}^{\beta+1}(\nu_j)$ to calculate the throughput. As in the case with $k=1$ the time complexity to compute the vector of throughputs is identical to the time complexity of computing the normalizing constant, that is, $O(n(k+1)|\hat{\Omega}_{C,k}^{\beta}|)$.

7.3. Fairness

In this section, we establish two results: when $k = C$ setting the rates equal to $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$ results in fairness for any α (as proven in [20] for $C = 1$) and when $k < C$ using the same rates results in fairness under heavy traffic (which is a generalization of Theorem 6 for $k = 1$). The structure of the proof of the latter, is similar to the one of Theorem 6.

Theorem 9. *Let $n > \beta \geq 1$, $k = C \leq \beta+1$, $\alpha > 0$ and set $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$, then*

$$\hat{\theta}_{k,j}(\nu) = \frac{\alpha}{1+(1+\beta)\alpha}C,$$

for $j = 1,\ldots,n$

Proof. Since link j uses a separate backoff timer for each channel and $C = k$, this network may be regarded as a network consisting of Cn links with a conflict graph that is composed of C connected components, where each component is identical to the conflict graph of a line network that uses a single channel. When $C = 1$ [20] showed that the rates ν_j result in a throughput of $\frac{\alpha}{1+(1+\beta)\alpha}$ for each link. As the throughput of the multi-channel links is merely the sum of the throughputs of their single channel counterparts, the result follows. $\square$

Theorem 10. *Let $n > \beta \geq 1$, $k \leq C \leq \beta+1$ and set $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$, then*

$$\lim_{\alpha\to\infty} \hat{\theta}_{k,j}(\nu) = \frac{C}{\beta+1},$$

for $j = 1, \ldots, n$

Proof. Using an argument similar to the proof of Theorem 6, one finds that the joint process spends most of its time in states maximizing

$$val_k(\hat{u}) \stackrel{def}{=} \sum_{j=1}^{n} \sum_{i=1}^{C} \hat{u}_{i,j}(\gamma(j) - \gamma(1) + 1).$$

and this value is only reached for states belonging to the set

$$\hat{\mathcal{M}}_{C,k}^{n} = \{\hat{u} \in \Omega_{C,k}^{n} | \sum_{j=1}^{\beta+1} \sum_{i=1}^{C} \hat{u}_{i,j} = C, \hat{u}_{i,j} = \hat{u}_{i,j-(\beta+1)}, j > \beta + 1\}.$$

We now argue that the number of states in which link j is active on x channels, for $x = 1, \ldots, k$, is the same for any j. This can be seen by noting that if a state $\hat{u} \in \hat{\mathcal{M}}_{C,k}^{n}$, we can construct a state $\hat{u}' \in \hat{\mathcal{M}}_{C,k}^{n}$ by a shift to the left, i.e. $\hat{u}'_{i,1} = \hat{u}_{i,2}, \ldots, \hat{u}'_{i,n-1} = \hat{u}_{i,n}, \hat{u}'_{i,n} = \hat{u}_{i,1}$. By repeating this shift β times, we can construct for every given state $\hat{u} \in \hat{\mathcal{M}}_{C,k}^{n}$ in which node j is active on $s(\hat{u}, j)$ channels, a state $\hat{u}_l \in \hat{\mathcal{M}}_{C,k}^{n}$ such that $s(\hat{u}_l, l) = s(\hat{u}, j)$ for any $l \neq j$. Further, as the Markov chain spends on average an equal amount of time in every state belonging to $\hat{\mathcal{M}}_{C,k}^{n}$, we may conclude that every link achieves the same throughput. As we have activity on each of the C channels for any set of $\beta + 1$ consecutive links in $\hat{\mathcal{M}}_{C,k}^{n}$, this leads to a throughput of $\hat{\theta}_{k,j}(\nu) = \frac{C}{\beta+1}$. □

The previous result may seem like a very natural generalization of Theorem 6, but it turns out that this generalization only works if we do not allow repacking. In other words, if we do allow repacking, $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$ does not result in fairness when α tends to infinity as demonstrated by the following example. Consider $C = k = \beta = 2$ and $n = 4$, meaning all links interfere with each other except for link 1 and 4. With repacking one obtains a Markov chain by simply keeping track of the number of channels that each link is occupying. Further, as the system behaves as a set of n $M/M/k/k$ queues with a truncated state space, one can establish the following product form for this Markov chain

$$\grave{\pi}_k(\bar{u}) = \grave{Z}_{\nu,k}^{-1} \prod_{j=1}^{n} {\nu_j}^{\bar{u}_j} \binom{C}{\bar{u}_j},$$

with $\bar{u} = (\bar{u}_1, \ldots, \bar{u}_n)$ and $\bar{u}_j$ denotes the number of acquired channels by link j. Due to this product form, the following six states dominate as α tends to infinity: 1101, 1011, 0110, 2002, 0200 and 0020 (their probability grows as α^4). Due to the presence of the coefficients $\binom{C}{\bar{u}_j}$, the probability of the first two states tends to 8/23, of the third tends to 4/23 and of the last three tends to 1/23 as α tends to infinity. Hence, the first and the last link have a throughput of $18/23 = 0.7826$, the links in the middle have a throughput of $14/23 = 0.6037$. Generalizing the proof of Theorem 6 thus fails because of the factor $\binom{C}{\bar{u}_j}$ in the product form. In our particular example links in the middle have a disadvantage because the probability of being in state 0110 only tends to 4/23. However, if we would consider a network of $n = 6$ links, one finds that all the links have throughput 2/3. It appears that in general, the rates $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$ still achieves heavy traffic fairness in systems with repacking if n is a multiple of $\beta + 1$.

7.4. Numerical results

In Figure 6 we plotted the fairness index as a function of α for $k = 1, \ldots, 4$ when $n = 40$, $C = 4$, $\beta = 6$ and this for $\nu_j = \alpha$ and $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$. The main observation is that k affects the fairness in a very different manner in both scenarios. When $\nu_j = \alpha$ the fairness decreases with increasing k, while when setting $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$ the opposite occurs. When $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$ we have fairness for any α when $k = C$, but for $k < C$, we only have fairness under heavy traffic. As such it can be expected that the larger k values result in more fairness. What is perhaps somewhat surprising is that most of the unfairness already vanishes when we increase k from 1 to 2. When $\nu_j = \alpha$ the decrease in fairness can be understood by noting that for $k = 1$ the unfair advantage of links 1 and n is constrained to the first channel they acquire, leaving the other 3 channels available for the other links. However, if we increase k, these links can occupy more channels, leaving fewer channels for the other links. This can also be seen in the left plot of Figure 7: the throughput of the border links grows considerably while the throughput of its β neighbors decreases somewhat.

Another observation that can be made in Figure 7 is that for $k < C$, the rates $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$ tend to overcompensate, giving the links that benefited in the equal rates regime a clear disadvantage. Finally, we also observe that k has barely any influence on the mean link throughput when $\nu_j = \alpha(1+\alpha)^{\gamma(j)-\gamma(1)}$. On the other hand, increasing k when using equal

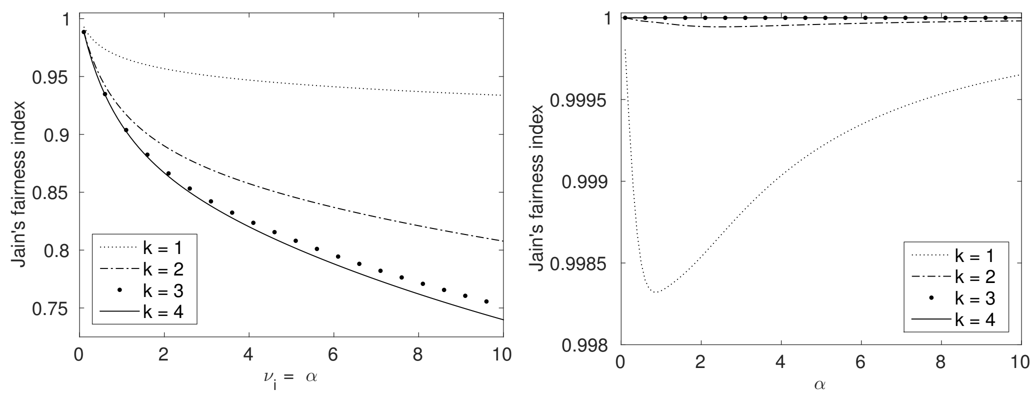

Figure 6: Fairness in terms of load with $n = 40$, $C = 4$, $\beta = 6$, k varying between 1 and 4 with rates $\nu_i = \alpha$ (left) and $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ (right).

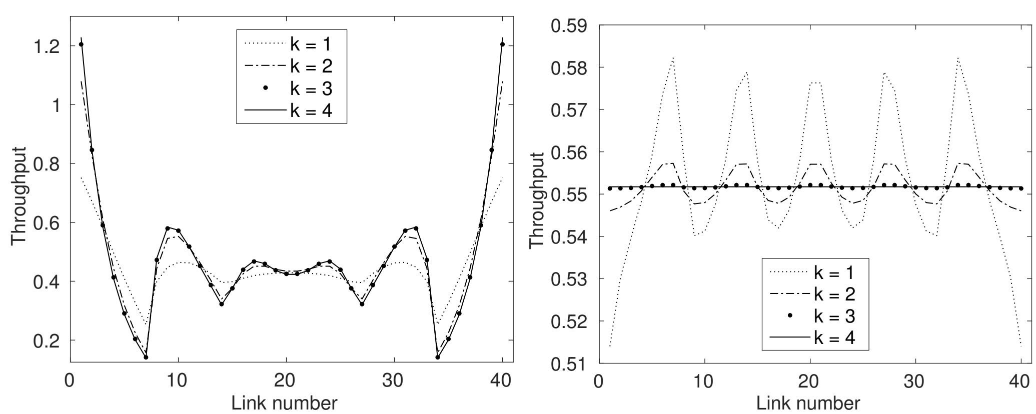

Figure 7: Throughput for each link with $n = 40$, $C = 4$, $\beta = 6$, k varying between 1 and 4 with rates $\nu_i = \alpha = 4$ (left) and $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ (right).

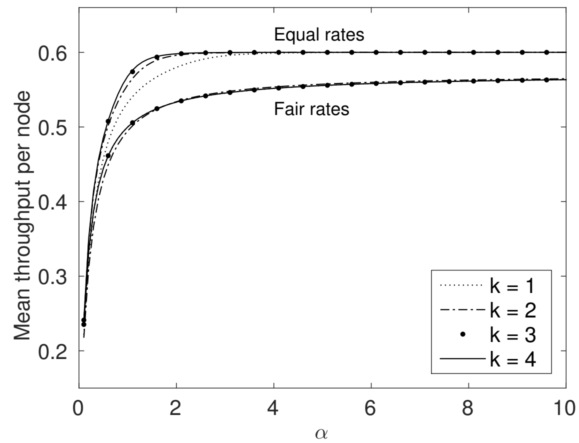

Figure 8: Mean throughput in terms of load with $n = 40$, $C = 4$, $\beta = 6$, k varying between 1 and 4 with rates $\nu_i = \text{avg}(\alpha(1+\alpha)^{\gamma(i)-\gamma(1)})$ (left) and $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ (right).

rates leads to an increased mean link throughput, meaning that the links that benefit gain more than what the remaining links lose.

8. Inhomogeneous line networks

In this section, we further generalize the model of Section 7 by allowing more flexibility in the number of interfering left and right links. Since the methodology is quite similar to the one in Section 7, we simply stress the main differences.

8.1. Model description

Let β_i be the number of left neighbors interfering with link i and define $\mathcal{B} = (\beta_1, \ldots, \beta_n)$. As we assume that interference is symmetric, $\mathcal{B}$ fully characterizes the interference graph. We also remark that $\beta_{i+1} \leq \beta_i + 1$, as we assume that every left interfering link $j \neq i$ of link $i+1$ also interferes with link i. Without loss of generality, we assume that $\beta_i > 0$, for $i = 2, \ldots, n$ (as $\beta_1 = 0$ by definition). If $\beta_i = 0$, there would be no interference between links $1, \ldots, i-1$ and $i, \ldots, n$ and both networks can be analyzed independently.

Similar to Section 7 we can define a Markov chain, the states of which can be represented by means of a $C \times n$-matrix with entries $\bar{u}_{i,j}$. Its state space is however different and is given by

$$\Omega_{C,k}^{\mathcal{B}} = \{\bar{u} | d(\bar{u}, i, j, \mathcal{B}) \leq 1, s(\bar{u}, m) \leq k, 1 \leq i \leq C, 1 \leq j \leq n, 1 \leq m \leq n\},$$

where $d(\bar{u}, i, j, \mathcal{B}) = \sum_{x=j-\beta_j}^{j} \bar{u}_{i,x}$ and $s(\bar{u}, m)$ is defined as before. Furthermore, its steady state probabilities have the same form:

$$\pi(\bar{u}) = \bar{Z}_{\nu,k}^{-1} \prod_{i=1}^{n} \prod_{j=1}^{C} \nu_i^{\bar{u}_{i,j}}.$$

To express the throughput of each link we define the set $\hat{\Omega}_{C,k}^{\mathcal{B}}$ as

$$\hat{\Omega}_{C,k}^{\mathcal{B}} = \{(z_1, \ldots, z_n) \in \{0, 1, \ldots, k\}^n | \sum_{x=0}^{\beta_j} z_{j-x} \leq C, j = 1, \ldots, n\}.$$

Further, let $(\hat{\Omega}_{C,k}^{\mathcal{B}})_{i:j} = \{(z_i, z_{i+1}, \ldots, z_j) | z \in \hat{\Omega}_{C,k}^{\mathcal{B}}\}$. We define the vector

$$\bar{M}_{C,k}^{\beta_1}(\nu_1) = (1, C\nu_1, \binom{C}{2}\nu_1^2, \ldots, \binom{C}{k}\nu_1^k),$$

and the matrices $\bar{M}_{C,k}^{\beta_i}(\nu_i)$, $1 < i \leq n$, each of size $|(\hat{\Omega}_{C,k}^{\mathcal{B}})_{i-1-\beta_{i-1}:i-1}| \times |(\hat{\Omega}_{C,k}^{\mathcal{B}})_{i-\beta_i:i}|$, with entry (z, z') given by

$$(\bar{M}_{C,k}^{\beta_i}(\nu_i))_{z,z'} = \begin{cases} 1 & z' = (z_{\beta_{i-1}-\beta_i+2}, \ldots, z_{\beta_{i-1}+1}, 0), \\ \nu_i^j \binom{C-f(z_{\beta_{i-1}-\beta_i+2:\beta_{i-1}+1})}{j} & f(z_{\beta_{i-1}-\beta_i+2:\beta_{i-1}+1}) + j \leq C, \\ & z' = (z_{\beta_{i-1}-\beta_i+2}, \ldots, z_{\beta_{i-1}+1}, j). \\ 0 & \textit{otherwise}, \end{cases}$$

The idea behind the construction is similar to (9), except that the number of left neighbors varies with i. Similarly, we define

$$\bar{N}_{C,k}^{\beta_1}(\nu_1) = (1, C\nu_1, 2\binom{C}{2}\nu_1^2, \ldots, k\binom{C}{k}\nu_1^k),$$

and

$$(\bar{N}_{C,k}^{\beta_i}(\nu_i))_{z,z'} = \begin{cases} j\nu_i^j \binom{C-f(z_{\beta_{i-1}-\beta_i+2:\beta_{i-1}+1})}{j} & f(z_{\beta_{i-1}-\beta_i+2:\beta_{i-1}+1}) \leq C, \\ & z' = (z_{\beta_{i-1}-\beta_i+2}, \ldots, z_{\beta_{i-1}+1}, j), \\ 0 & \textit{otherwise}. \end{cases}$$

The normalizing constant and the throughput of the links can now be obtained using the following theorems.

Theorem 11. *The normalization constant of an inhomogeneous line network in which each link is allowed to transmit on up to k channels simultaneously, can be expressed as*

$$\bar{Z}_{\nu,k} = \left(\prod_{i=1}^{n} \bar{M}_{C,k}^{\beta_i}(\nu_i)\right) e.$$

Proof. The proof is similar to the proof of Theorem 7. $\square$

Theorem 12. *The throughput $\bar{\theta}_{k,j}(\nu)$ of node j can be computed as*

$$\bar{\theta}_{k,j}(\nu) = \frac{\left(\prod_{i=1}^{j-1} \bar{M}_{C,k}^{\beta_i}(\nu_i)\right) \bar{N}_{C,k}^{\beta_i}(\nu_j) \left(\prod_{i=j+1}^{n} \bar{M}_{C,k}^{\beta_i}(\nu_i)\right) e}{\left(\prod_{i=1}^{n} \bar{M}_{C,k}^{\beta_i}(\nu_i)\right) e}.$$

Proof. The result is a consequence of Theorem 11 and the fact that we multiply $\pi(u)$ by $s(u,j)$. $\square$

8.2. Numerical results

To demonstrate that the rates needed to achieve fairness in a homogeneous line network affect the throughputs in an inhomogeneous network in a rather unpredictable manner, we perform a simple experiment. Consider a homogeneous network of $n = 40$ links, let $C = 4$ and $k = C$. We split this network into two networks of 20 links each in 6 steps by reducing the number of interfering links in the middle during each step. The resulting number of left neighbors for links 20 to 26 after each step is shown in Table 1, β_i remains 5 for the other links. During each step, we calculate the number of interfering neighbors $\gamma(i)$ of each link and use this number to set its rate to $\alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$. The outcome of this experiment is shown in Figure 9 for two different values of α.

Note that as $k = C$, Theorem 9 implies that we have perfect fairness in step 1 and in step 6. However, as expected, it is clear in Figure 9 that fairness is lost when the network is no longer homogeneous. Furthermore, it is hard to predict what happens to the fairness during the intermediate steps. For instance, when $\alpha = 4$, we see that the fairness is the most affected in step 4, while when $\alpha = 100$, fairness worsens with every step, except for the last step, which restores the fairness in a single step.

Step \ i	20	21	22	23	24	25	26
1	5	5	5	5	5	5	5
2	5	4	4	4	4	4	5
3	5	3	3	3	3	4	5
4	5	2	2	2	3	4	5
5	5	1	1	2	3	4	5
6	5	0	1	2	3	4	5

Table 1: Number of left neighbors of links 20 to 26 after each step when splitting a homogeneous line network of 40 links into two networks of size 20.

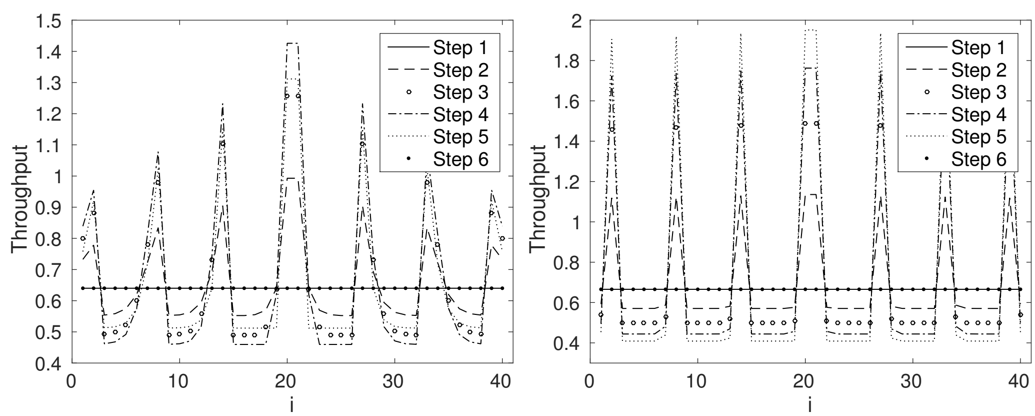

Figure 9: Throughput for each link when stepwise splitting a network of 40 links with $C = k = 4$, $\beta = 5$ and rates $\nu_i = \alpha(1+\alpha)^{\gamma(i)-\gamma(1)}$ with $\alpha = 4$ (left) and $\alpha = 100$ (right).

9. Conclusion

In this paper we considered an idealized multi-channel CSMA line network characterized by the number of links n, the interference range β, the number of channels C and the vector of backoff rates $(\nu_1, \ldots, \nu_n)$. We developed a numerical method to compute the vector of throughputs in a time complexity that is linear in the number of links n (while being exponential in β and C).

Using this method we analyzed the fairness in a multi-channel CSMA line network and found that while the simple formula of [20] for the single channel setting does not generalize to a system with $C > 1$ channels, the degree of unfairness that this formula causes is very limited. We further showed that this unfairness vanishes under heavy traffic conditions, that is, as the backoff rates tend to infinity.

While the initial model is restricted to homogeneous line networks where a link can be active on at most one channel at a time, we generalized the proposed methodology such that both these restrictions can be relaxed.

[1] R. Block and B. Van Houdt. Spatial fairness in multi-channel CSMA line networks. In *8th International Conference on Performance Evaluation Methodologies and Tools, VALUETOOLS 2014, Bratislava, Slovakia, December 9-11, 2014*, 2014.

[2] R. C. Bollinger. Fibonacci k-sequences, Pascal T-triangles, and k-in-a-row problems, 1982.

[3] T. Bonald, S. Borst, N. Hegde, and A. Proutiére. Wireless data performance in multi-cell scenarios. *ACM SIGMETRICS Performance Evaluation Review*, 32(1):378–387, 2004.

[4] T. Bonald and M. Feuillet. Performance of CSMA in multi-channel wireless networks. *Queueing Systems*, 72(1-2):139–160, 2012.

[5] B. A. Bondarenko. Generalized pascal triangles and pyramids: Their fractals, graphs, and applications, 1993.

[6] R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin. Throughput analysis in multihop csma packet radio networks. *IEEE Transactions on Communication*, 35(3):267–274, 1987.

[7] N. Bouman. *Queue-based Random Access in Wireless Networks*. PhD thesis, Technical University Eindhoven, 2013.

[8] M. T. Chao, J. C. Fu, and M. V. Koutras. Survey of reliability studies of consecutive-k-out-of-n:F and related systems. *IEEE Transactions on Reliability*, 44(1):120–127, 1995.

[9] M. Durvy, O. Dousse, and P. Thiran. On the fairness of large CSMA networks. *IEEE Journal on Selected Areas in Communications*, 27(7):1093–1104, 2009.

[10] 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.

[11] J. C. Fu. Reliability of consecutive-k-out-of-n:F systems with (k-1)step Markov dependence. *IEEE Transactions on Reliability*, 35:602–606, 1986.

[12] R. Jain, D.-M. Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical report, DEC-TR-301, Digital Equipment Corporation, 1984.

[13] F. P. Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 2011.

[14] Soung Chang Liew, JiaLiang Zhang, Chi-Kin Chau, and Minghua Chen. Analysis of frequency-agile CSMA wireless networks. *CoRR*, abs/1007.5255, 2010.

[15] J. Naus. Probabilities for a generalized birthday problem. *Journal of the American Statistical Association*, 69(347):810–815, 1974.

[16] B. Saperstein. The generalized birthday problem. *Journal of the American Statistical Association*, 67:425–428, 1972.

[17] P. M. van de Ven. *Stochastic models for resource sharing in wireless networks*. PhD thesis, Technical University Eindhoven, 2011.

[18] P. M. van de Ven, S. C. Borst, J. S. H. van Leeuwaarden, and A. Proutière. Insensitivity and stability of random-access networks. *Performance Evaluation*, 67(11):1230–1242, November 2010.

[19] P. M. van de Ven, A. J. E. M. Janssen, J. S. H. van Leeuwaarden, and S. C. Borst. Achieving target throughputs in random-access networks. *Performance Evaluation*, 68(11):1103–1117, November 2011.

[20] P. M. van de Ven, J. S. H. van Leeuwaarden, D. Denteneer, and A. J. E. M. Janssen. Spatial fairness in linear random-access networks. *Performance Evaluation*, 69(3-4):121–134, March 2012.