Achieving Target Throughputs in Random-Access Networks

Abstract

Random-access algorithms such as CSMA provide a popular mechanism for distributed medium access control in large-scale wireless networks. In recent years, tractable stochastic models have been shown to yield accurate throughput estimates for CSMA networks. We consider a saturated random-access network on a general conflict graph, and prove that for every feasible combination of throughputs, there exists a unique vector of back-off rates that achieves this throughput. This result entails proving global invertibility of the non-linear function that describes the throughputs of all nodes in the network. We present several numerical procedures for calculating this inverse, based on fixed-point iteration and Newton's method. Finally, we provide closed-form results for several special conflict graphs using the theory of Markov random fields.

1 Introduction

Carrier-Sense Multiple-Access (CSMA) is a popular mechanism for distributed medium access control in wireless networks with conflicting nodes. Under CSMA, a node attempts to activate after a certain random back-off time, but freezes its back-off timer whenever it senses activity of an interfering node, until the medium is again sensed idle. The local interactions lead to rather complex behavior on a macroscopic scale, which can be studied via models which assume that the interference constraints can be represented by a general conflict graph, and that the various nodes activate asynchronously after some generally distributed back-off time whenever none of their neighbors are presently active. Such models were first pursued in the context of IEEE 802.11 systems by Wang & Kar [20], and further studied in [4, 5, 6]. These models in fact long pre-date the IEEE 802.11 standard and were already considered in the 1980's [1, 2, 10, 12] for the study of packet radio networks. The model has strong connections with Markov random fields and migration processes, and can under certain assumptions be interpreted as a special instance of a loss network [9, 11, 16, 22].

Our main object of interest is the throughput vector, which expresses for each node the rate at which transmissions are completed. The throughput of a node is obtained by summing the stationary probabilities of all states in which this node is active, and a key feature of this model is that the stationary measure of the

activity state has a closed-form solution, also known as a Gibbs measure. The throughput vector is thus a function of the Gibbs measure, and in fact is a highly non-trivial and non-linear function of the back-off rates. Classical references mostly focus on the analysis of the network behavior, given the back-off rates. In this paper we address the *inverse problem* of determining the back-off rates that achieve some target throughput vector, and for this, we need to study in detail the non-linear function describing the throughputs.

The importance of choosing the back-off rates to achieve some target throughput vector is illustrated by the observation that when all back-off rates are equal, the resulting network throughput is highly unfair [4, 5, 20]. This unfairness is related to the fact that some nodes are in a spatially disadvantageous position (e.g., have many neighbors). Intuition suggests that in order to resolve unfairness, the back-off times of poorly positioned nodes should be shorter than those of high-throughput nodes. However, due to the complex interactions between nodes, it is hard to see how to choose the back-off times exactly, in order to balance the throughputs. This problem has been solved for special conflict graphs: regular trees [10], linear networks [19] and certain rectangular grids [17], which obtain an exact description of the fair vector of back-off rates.

We are also interested in the more general problem of finding back-off rates that yield certain relative throughputs, and we aim to do this for general conflict graphs. To this end we study the non-normalized throughput function, and show that it is *globally invertible*. That is, for every target vector $(s_1, s_2, \ldots, s_n) \in (0, \infty)^n$, with s_i the relative throughput of node i, we can find a unique vector of mean back-off times that achieves this target vector. The special case $s_1 = \ldots = s_n$ corresponds to fairness.

In addition to studying just the relative throughputs, we can also examine the throughput function itself. This throughput function is obtained by dividing the non-normalized throughput by the normalization constant. We may show that the full throughput function is also globally invertible, so for every achievable throughput vector, there exists a unique vector of back-off rates that achieves this vector. The proof of this result relies on the global invertibility of the non-normalized throughput and a monotonicity property of the normalization constant.

Hence, our first contribution is that we prove global invertibility, for both the non-normalized and normalized throughput functions. This result of course raises the question of how to actually compute the inverse, or the vector of back-off rates that will render the throughput targets. Finding the inverse requires solving a set of non-linear equations, which can be written as a fixed-point equation. Closed-form solutions of this fixed-point equation exist only for a few specific choices of the network topology [10, 17, 19] in the case of strict fairness. For the general conflict graph, we present three numerical methods to determine the fixed point: fixed-point iteration, basic Newton iteration, and a continuation method (consisting of a sequence of Newton iteration steps). The formulation of the inverse problem as a fixed-point equation that can be solved numerically is the second contribution of this paper. A third contribution is that we exploit the theory of

Markov random fields to find closed-form solutions to the fixed-point equations of specific conflict graphs, as well as to decompose certain conflicts graphs into separate subgraphs for which the inverse problem can be treated in isolation, reducing its complexity.

Our results pertain to the case where all nodes are saturated, and the back-off rates of nodes are fixed. The case of fixed back-off rates is relevant, since in practice the adaptation of back-off parameters involves a wide range of non-trivial implementation issues (finite-range precision, communication overhead, information exchange), and hence it is important to gain insight in the achievable performance of non-adaptive algorithms.

Related results have been obtained for scenarios where the nodes each have a buffer fed by exogenous arrival processes with certain rates, and the back-off rates can be adapted. For this case, several authors have proposed clever backlog-based algorithms for adapting back-off rates which achieve stability whenever feasible to do so at all, and thus can realize certain throughput vectors [7, 8, 13, 14, 15]. In fact, in this setting, a result similar to Theorem 2 was derived in Section 5.1 of [7] by considering an optimization problem which naturally arises in the context of the variational principle related to Gibbs measures and Markov random fields. This optimization problem was shown to have a solution which corresponds to the vector of back-off rates that will render a service rate vector that precisely matches the arrival rates. It was further established in [7] that the time-varying backlog-dependent back-off rates, which are controlled via a specific distributed algorithm, converge to such a vector of back-off rates. We show that these rates in fact form a unique solution.

The paper is organized as follows. In Section 2 we introduce the model and describe the product-form solution for the stationary distribution and the throughputs. Our main results on global invertibility are presented in Section 3. In Section 4 we describe several numerical methods for determining the inverse throughput function. Section 5 is concerned with results for special conflict graphs. Finally, Section 6 presents some conclusions and a discussion. We denote $\mathbb{R}_+ = (0, \infty)$, and use boldface notation for vectors. We write x_i for the ith element of a column vector $\boldsymbol{x}$. The transpose of a column vector $\boldsymbol{x}$ is a row vector, denoted by $\boldsymbol{x}^T$.

2 Model description

Consider a network of nodes sharing the wireless medium according to a CSMA-type protocol. We say that nodes are either active or inactive, depending on whether they are transmitting or in back-off, respectively. The network structure is characterized by a conflict graph $G(V, E)$, where the set of vertices $V = \{1, 2, \ldots, n\}$ represents the nodes, and the set of (undirected) edges specify the interference constraints. Neighboring nodes are prevented from simultaneous activity, so the independent sets of the graph correspond to the feasible joint activity states. An inactive node is said to be blocked whenever any of its neighbors is active, and

unblocked otherwise.

We assume the network to be saturated, which means that nodes always have packets available for transmission. The transmission times of nodes are independent and generally distributed with unit mean. After finishing a transmission, node i starts a generally distributed, independent back-off period with mean ν_i. If a node in back-off becomes blocked, it freezes its back-off period until all neighboring nodes are inactive, at which point the back-off period is resumed.

Define $\Omega \subseteq \{0,1\}^n$ as the set of all feasible joint activity states of the network, i.e., all independent sets of the conflict graph. Denote the total number of feasible activity states by $K+1$, and write $\Omega = \{\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K\}$. The states are ordered such that $\boldsymbol{\omega}_0 = \mathbf{0}$, the empty state, and $\boldsymbol{\omega}_k = \mathbf{e}_k$, the kth unit vector of $\mathbb{R}^n$, $k = 1, 2 \ldots, n$. Note that the case $K = n$ corresponds to the complete conflict graph, at most one node can be active at any time.

Let $Y(t) \in \Omega$ represent the activity state of the network at time t, with $Y_i(t)$ indicating whether node i is active at time t ($Y_i(t) = 1$) or not ($Y_i(t) = 0$). Denote by $\pi(\boldsymbol{y}) = \lim_{t\to\infty} \mathbb{P}\{Y(t) = \boldsymbol{y}\}$ the limiting probability that the activity state of the network is $\boldsymbol{y} \in \Omega$, which is given by the product form distribution [18]

$$\pi(\boldsymbol{y}) = Z^{-1} \prod_{i=1}^{n} \nu_i^{y_i},$$

with $Z = \sum_{\boldsymbol{y}\in\Omega} \prod_{i=1}^{n} \nu_i^{y_i}$.

For later purposes, it is convenient to explicitly reflect the ordering of the states and the dependence on the back-off rate vector $\boldsymbol{\nu} = (\nu_1, \nu_2, \ldots, \nu_n)^T$ in the notation, and introduce $\pi_k(\boldsymbol{\nu}) = \pi(\boldsymbol{\omega}_k)$:

$$\pi_k(\boldsymbol{\nu}) = \frac{\Lambda_k(\boldsymbol{\nu})}{Z(\boldsymbol{\nu})}, \quad k = 0, 1, \ldots, K,$$

where (with $\boldsymbol{\omega}_k = (\omega_{k,1}, \ldots, \omega_{k,n})^T$)

$$\Lambda_k(\boldsymbol{\nu}) = \prod_{i=1}^{n} \nu_i^{\boldsymbol{\omega}_{k,i}} \tag{1}$$

and $Z(\boldsymbol{\nu}) = \sum_{k=0}^{K} \Lambda_k(\boldsymbol{\nu})$ is the normalization constant.

We are interested in the long-term behavior of the network, characterized by the throughputs of the various nodes. Denote by $\theta_i(\boldsymbol{\nu})$ the rate at which node i finishes packet transmissions, and write $\boldsymbol{\theta}(\boldsymbol{\nu}) = (\theta_1(\boldsymbol{\nu}), \theta_2(\boldsymbol{\nu}), \ldots, \theta_n(\boldsymbol{\nu}))^T$, the throughput vector of the network. As active nodes finish their transmissions at unit rate, we have that

$$\boldsymbol{\theta}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \pi_k(\boldsymbol{\nu}) \boldsymbol{\omega}_k.$$

Let $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_n)^T \in \mathbb{R}_+^n$ belong to the range Γ of the mapping $\boldsymbol{\theta} : \mathbb{R}_+^n \to \Gamma$. In [8] it is shown that Γ is

the interior of the convex hull formed by all states $\boldsymbol{\omega}_0, \boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_K$, i.e.,

$$\Gamma = \Big\{ \sum_{k=0}^{K} \alpha_k \boldsymbol{\omega}_k \;\Big|\; \sum_{k=0}^{K} \alpha_k = 1,\ \alpha_k \geq 0, k = 0, \ldots, K \Big\}.$$

The problem of finding back-off rates that achieve a certain throughput vector can be formalized as finding $\boldsymbol{\nu}_{\boldsymbol{\theta}} = \boldsymbol{\nu}_{\boldsymbol{\theta}}(\boldsymbol{\gamma})$ that solves

$$\boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = \boldsymbol{\gamma}, \tag{2}$$

and hence, we need to study in detail the mapping $\boldsymbol{\theta}$.

3 Global invertibility

We start by considering the non-normalized throughput

$$\boldsymbol{\eta}(\boldsymbol{\nu}) = Z(\boldsymbol{\nu})\boldsymbol{\theta}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \Lambda_k(\boldsymbol{\nu})\boldsymbol{\omega}_k.$$

This function is monotone in $\boldsymbol{\nu}$ and hence easier to handle than the normalized throughput. In fact, we can show the following result.

Theorem 1. *The mapping $\boldsymbol{\eta} : \mathbb{R}_+^n \to \mathbb{R}_+^n$ is globally invertible on its range $\mathbb{R}_+^n$.*

The proof of Theorem 1 is presented in Appendix A.1. Theorem 1 says that the range of $\boldsymbol{\eta}$ is $\mathbb{R}_+^n$, and that for any $\boldsymbol{\gamma}$ within the range of $\boldsymbol{\eta}$, we can find a unique $\boldsymbol{\nu}_{\boldsymbol{\eta}} = \boldsymbol{\nu}_{\boldsymbol{\eta}}(\boldsymbol{\gamma})$ that solves

$$\boldsymbol{\eta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}) = \boldsymbol{\gamma}. \tag{3}$$

In some cases, it might be beneficial from a computational point of view to invert $\boldsymbol{\eta}$ rather than $\boldsymbol{\theta}$. Although $\boldsymbol{\eta}$ only represents the non-normalized throughput, this is sufficient when interested solely in the throughputs ratios (for instance, when aiming for fairness).

The crucial difference between $\boldsymbol{\eta}$ and $\boldsymbol{\theta}$ is the normalization constant $Z(\boldsymbol{\nu})$, for which we have the following result:

Lemma 1. *Let $\boldsymbol{c} \in \mathbb{R}_+^n$, $s > 0$ and write $\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})$ for the unique $\boldsymbol{\nu} \in \mathbb{R}_+^n$ such that $\boldsymbol{\eta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) = s\boldsymbol{c}$. Then, the function $f_{\boldsymbol{c}}(s) = s/Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))$ is injective.*

The proof of Lemma 1 is presented in Appendix A.2. Lemma 1 suggests that we can control the throughput $\boldsymbol{\theta}$ via the non-normalized throughput $\boldsymbol{\eta}$, and indeed, it turns out to be a crucial ingredient in the proof of the following result:

Theorem 2. *The mapping* $\boldsymbol{\theta} : \mathbb{R}_+^n \to \Gamma$ *is globally invertible on* Γ.

Proof. It suffices to show that $\boldsymbol{\theta}$ is injective. Let $\boldsymbol{\nu}_1, \boldsymbol{\nu}_2 \in \mathbb{R}_+^n$ be such that $\boldsymbol{\theta}(\boldsymbol{\nu}_1) = \boldsymbol{\theta}(\boldsymbol{\nu}_2)$. Then we have

$$\boldsymbol{\eta}(\boldsymbol{\nu}_1) = Z(\boldsymbol{\nu}_1)\boldsymbol{\theta}(\boldsymbol{\nu}_1), \quad \boldsymbol{\eta}(\boldsymbol{\nu}_2) = Z(\boldsymbol{\nu}_2)\boldsymbol{\theta}(\boldsymbol{\nu}_2). \tag{4}$$

With $\boldsymbol{c} = \boldsymbol{\theta}(\boldsymbol{\nu}_1) = \boldsymbol{\theta}(\boldsymbol{\nu}_2) \in \mathbb{R}_+^n$, we consider the trajectory $\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})$, $s > 0$, for which we have

$$\boldsymbol{\eta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) = s\boldsymbol{c} = Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))\boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})). \tag{5}$$

With $s_1 = Z(\boldsymbol{\nu}_1)$, $s_2 = Z(\boldsymbol{\nu}_2)$, it follows from Theorem 1 and (4), (5) that $\boldsymbol{\nu}_{\boldsymbol{\eta}}(s_1\boldsymbol{c}) = \boldsymbol{\nu}_1$, $\boldsymbol{\nu}_{\boldsymbol{\eta}}(s_2\boldsymbol{c}) = \boldsymbol{\nu}_2$, and that

$$\frac{1}{s_1} Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s_1\boldsymbol{c})) = \frac{1}{s_2} Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s_2\boldsymbol{c})).$$

Hence, by injectivity of $f_{\boldsymbol{c}}(s)$, it follows that $s_1 = s_2$, so that $\boldsymbol{\nu}_1 = \boldsymbol{\nu}_2$. □

Theorem 2 thus says that for any $\boldsymbol{\gamma}$ within the range of $\boldsymbol{\theta}$, there is a unique vector $\boldsymbol{\nu}_{\boldsymbol{\theta}} = \boldsymbol{\nu}_{\boldsymbol{\theta}}(\boldsymbol{\gamma})$ that solves (2). The proofs of Theorems 1 and 2 require the description of the entire network structure, which appears at odds with the distributed nature of CSMA. However, in actual implementations, the back-off times only have to be determined once, after which the nodes can operate fully autonomously. Thus, if the network structure is fixed, or if the time scale on which it changes is slower than that of the network operations, we retain a fully distributed CSMA protocol, while achieving the throughput targets.

4 Inversion methods

In Section 3 we established that both the non-normalized throughput $\boldsymbol{\eta}$ and the normalized throughput $\boldsymbol{\theta}$ are globally invertible on their respective ranges. In this section we present several numerical procedures to compute the inverse of a given (normalized) throughput vector, as well as a light-traffic approximation of the throughput inverse.

4.1 Fixed-point iteration

A first numerical procedure to determine the inverse vector is fixed-point iteration. This procedure follows naturally from rewriting the system of non-linear equations (2) as fixed-point equations. We distinguish between normalized throughput and non-normalized throughput.

4.1.1 Non-normalized throughput

Write

$$\theta_i(\boldsymbol{\nu}) = \sum_{k=0}^{n} \pi_k(\boldsymbol{\nu})\omega_{k,i} + \sum_{k=n+1}^{K} \pi_k(\boldsymbol{\nu})\omega_{k,i} = \nu_i \frac{1 + G_i(\boldsymbol{\nu})}{Z(\boldsymbol{\nu})}, \tag{6}$$

with

$$G_i(\boldsymbol{\nu}) = \frac{1}{\nu_i} \sum_{k=n+1}^{K} \Lambda_k(\boldsymbol{\nu})\omega_{k,i}. \tag{7}$$

We can thus write (3) as

$$\boldsymbol{\nu_\eta} = \mathbf{H}(\boldsymbol{\nu_\eta}), \quad \text{where} \quad \mathbf{H}(\boldsymbol{\nu}) = \Big(\frac{\gamma_i}{1 + G_i(\boldsymbol{\nu})}\Big)_{i=1,\dots,n},$$

and G_i as in (7). Note that $\mathbf{H} : [\mathbf{0}, \boldsymbol{\gamma}] \to [\mathbf{0}, \boldsymbol{\gamma}]$, where we denote $[\mathbf{0}, \boldsymbol{\gamma}] = [0, \gamma_1] \times \cdots \times [0, \gamma_n]$. By global invertibility of $\boldsymbol{\eta}$, we know that $\boldsymbol{\nu_\eta}$ is the unique fixed point that solves $\boldsymbol{\nu_\eta} = \mathbf{H}(\boldsymbol{\nu_\eta})$. Alternatively, since $\mathbf{H}$ is continuous, the existence of a fixed point also follows directly from Brouwer's fixed point theorem.

The fixed-point iteration is defined as

$$\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(0)} = \mathbf{0}, \quad \boldsymbol{\nu}_{\boldsymbol{\eta}}^{(l+1)} = \mathbf{H}(\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(l)}), \quad l = 0, 1, \dots. \tag{8}$$

We next show that the iterands obtained through (8) approach the fixed point in a monotone fashion.

Proposition 1. *Assume that the conflict graph has no fully connected nodes (i.e., nodes that are connected to all the other nodes). Then for $i = 1, 2, \dots, n$ and $l = 1, 2, \dots$ it holds that*

$$0 = \nu_{\eta,i}^{(0)} < \nu_{\eta,i}^{(2)} < \cdots < \nu_{\eta,i}^{(2l-2)} < \nu_{\eta,i} < \nu_{\eta,i}^{(2l-1)} < \nu_{\eta,i}^{(2l-3)} < \cdots < \nu_{\eta,i}^{(3)} < \nu_{\eta,i}^{(1)} = \gamma_i. \tag{9}$$

Proof. We have $\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(0)} = \mathbf{0}$ by definition, $\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(1)} = \boldsymbol{\gamma}$ since $G_i(\mathbf{0}) = 0$, $i = 1, \dots, n$, and $0 < \nu_{\eta,i} < \gamma_i$, $i = 1, \dots, n$. Now let l be such that (9) holds for all $i = 1, \dots, n$. Then $\nu_{\eta,i} < \nu_{\eta,i}^{(2l-1)}$, $i = 1, \dots, n$, and by the exclusion of fully connected nodes we have that

$$G_i(\boldsymbol{\nu_\eta}) < G_i(\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(2l-1)}), \quad i = 1, \dots, n,$$

and so

$$H_i(\boldsymbol{\nu_\eta}) > H_i(\boldsymbol{\nu}_{\boldsymbol{\eta}}^{(2l-1)}) = \nu_{\eta,i}^{(2l)}, \quad i = 1, \dots, n,$$

i.e.,

$$\nu_{\eta,i} > \nu_{\eta,i}^{(2l)}, \quad i = 1, \dots, n. \tag{10}$$

In a similar fashion it follows from (10) that

$$\nu_{\eta,i} < \nu_{\eta,i}^{(2l+1)}, \quad i = 1, \ldots, n.$$

The proof follows by induction. □

Proposition 1 shows that the iteration scheme in (8) approaches the fixed-point ever more closely, although it does not necessarily imply convergence.

4.1.2 Normalized throughput

We now present a similar fixed-point iteration scheme for $\boldsymbol{\nu_\theta}(\boldsymbol{\gamma})$. Setting $\boldsymbol{\theta}(\boldsymbol{\nu_\theta}) = \boldsymbol{\gamma}$ and rewriting (6) yields $\boldsymbol{\nu_\theta} = \mathbf{K}(\boldsymbol{\nu_\theta})$ with

$$\mathbf{K}(\boldsymbol{\nu}) = \left(\frac{\gamma_i Z(\boldsymbol{\nu})}{1 + G_i(\boldsymbol{\nu})} \right)_{i=1,\ldots,n}. \tag{11}$$

We have thus established that $\boldsymbol{\nu_\theta}(\boldsymbol{\gamma})$ is the unique solution to the fixed-point equation (11), and we can again try to find $\boldsymbol{\nu_\theta}(\boldsymbol{\gamma})$ by iteration. That is, we let $\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(0)} = \mathbf{0}$ and define recursively

$$\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l+1)} = \mathbf{K}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)}), \quad l = 0, 1, \ldots. \tag{12}$$

To gain some insight into this fixed-point iteration, we give below two special cases for which we can prove convergence to the fixed point.

Example 1. (complete conflict graph) Assume that only one node may be active at any time. Let $\boldsymbol{\gamma} = (\gamma, \ldots, \gamma)^T$, $\gamma \in \mathbb{R}_+$. By symmetry, both the solution $\boldsymbol{\nu_\theta}(\boldsymbol{\gamma})$ as well as the iterands $\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)}$, $l = 0, 1, \ldots$ have identical components. Thus $Z(\boldsymbol{\nu_\theta}) = 1 + n\nu_{\theta,1}$ and $G_i(\boldsymbol{\nu_\theta}) = 0$. Iterating according to (12), gives for all $i = 1, 2, \ldots, n$,

$$\nu_{\theta,i}^{(l)} = \gamma(1 + n\gamma + \cdots + (n\gamma)^{l-1}), \quad l = 0, 1, \ldots,$$

and $\nu_{\theta,i}(\boldsymbol{\gamma}) = \lim_{l\to\infty} \nu_{\theta,i}^{(l)} = \frac{\gamma}{1-n\gamma}$ for $\gamma < \frac{1}{n}$. The requirement for convergence $\gamma < \frac{1}{n}$ is equivalent to $\boldsymbol{\gamma} \in \Gamma$.

In this particular example, $\boldsymbol{\nu_\theta}(\boldsymbol{\gamma})$ can also be determined analytically. Noting that at most one node can be active at a time, and assuming all nodes to have the same back-off rate, it was shown in [18] that the throughput of node i equals $\theta_i = \nu_{\theta,1}/(1 + n\nu_{\theta,1})$. Solving $\gamma = \theta_i$ for $\boldsymbol{\nu_\theta}$ then gives the same result as the fixed-point iteration.

Example 2. (disconnected graph) Assume that all nodes are isolated ($E = \emptyset$). As nodes do not interact with each other at all, the throughput of node i equals $\nu_{\theta,i}/(1 + \nu_{\theta,i})$, and thus the choice $\nu_{\theta,i} = \gamma_i/(1 - \gamma_i)$, yields per-node throughputs γ, $\gamma < 1$.

The same result can be obtained by fixed-point iteration. Let $\boldsymbol{\gamma} = (\gamma, \ldots, \gamma)^T$, $\gamma \in \mathbb{R}_+$, so the target vector, solution and iterands have identical components. We have

$$Z(\boldsymbol{\nu_\theta}) = \binom{n}{0} + \binom{n}{1}\nu_{\theta,1} + \binom{n}{2}\nu_{\theta,1}^2 + \cdots + \binom{n}{n}\nu_{\theta,1}^n = (1+\nu_{\theta,1})^n$$

and

$$G_i(\boldsymbol{\nu_\theta}) = \binom{n-1}{1}\nu_{\theta,1} + \binom{n-1}{2}\nu_{\theta,1}^2 + \cdots + \binom{n-1}{n-1}\nu_{\theta,1}^{n-1} = (1+\nu_{\theta,1})^{n-1} - 1,$$

so that

$$K_i(\boldsymbol{\nu_\theta}) = \gamma(1+\nu_{\theta,1}). \tag{13}$$

By iterating (13), we get for all $i = 1, 2, \ldots, n$,

$$\nu_{\theta,i}^{(l)} = \gamma + \gamma^2 + \cdots + \gamma^l, \quad l = 0, 1, \ldots.$$

Thus $\nu_{\theta,i}(\boldsymbol{\gamma}) = \lim_{l\to\infty} \nu_{\theta,i}^{(l)} = \gamma/(1-\gamma)$, as expected.

Due to the inclusion of the normalization constant, the fixed point iteration for the normalized throughput becomes theoretically less tractable than for the non-normalized throughput, and a counterpart to Proposition 1 remains elusive. In applying the iteration, though, we encountered no convergence issues for the fixed-point iteration. See Section 4.3 for an example. In fact, for both the non-normalized and normalized throughputs the fixed-point iterations seems to work equally well.

4.2 Newton-based methods

A second numerical method for inverting the throughput function is Newton iteration. We present two versions: classical Newton iteration, and a continuation method. The latter method consists of a sequence of Newton iteration steps. Since there is no essential difference in these methods between the non-normalized and normalized case, we present the numerical procedures only for the normalized throughput $\boldsymbol{\theta}$.

4.2.1 Classical Newton iteration

Recall from basic Newton iteration that one selects an initial vector $\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(0)} \in \Gamma$, and iterates according to

$$\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l+1)} = \boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)} - \Big(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})\Big)^{-1}(\boldsymbol{\gamma} - \boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})), \quad l = 0, 1, \ldots,$$

where

$$\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}} = \Big(\frac{\partial \theta_i}{\partial \nu_j}\Big)_{i=1,\ldots,n;j=1,\ldots,n}$$

is the functional matrix, which also plays a crucial role in the proof of Theorem 1 (see (26) and further).

4.2.2 Continuation method

Let $\boldsymbol{\gamma} = s\boldsymbol{c}$, with $\boldsymbol{c} \in \mathbb{R}_+^n$ and $s > 0$. The general idea behind the continuation method is to compute a sequence of back-off rates $\boldsymbol{\nu}_{\boldsymbol{\theta}}(\boldsymbol{\gamma}^{(l)})$, with $\boldsymbol{\gamma}^{(l)} = l\delta\boldsymbol{c}$, and δ the step size such that s/δ is integer. Successive iterands are computed by performing a single Newton iteration step:

$$\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l+1)} = \boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)} - \Big(\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})\Big)^{-1}(l\delta\boldsymbol{c} - \boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)})), \quad l = 0, \ldots, s/\delta.$$

The step size affects the accuracy of the resulting approximation, as well as the computation time. In addition to finding $\boldsymbol{\nu}_{\boldsymbol{\theta}}(\boldsymbol{\gamma})$, the continuation method approximates the entire path $\boldsymbol{\nu}_{\boldsymbol{\theta}}(\boldsymbol{\gamma}^{(l)})$, $l = 0, 1, \ldots, s/\delta$. Similar to the fixed-point iteration, both Newton-based methods can be modified to work for $\boldsymbol{\eta}$ as well. This is done in both cases by replacing the functional matrix $\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}$ by $\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}}$ (see (18) and further).

4.3 Comparison of inversion methods

To illustrate the inversion methods we consider a linear network with $n = 15$ nodes in which an active node blocks all nodes within 5 hops. Note that for this particular network, we have a closed-form expression for the target back-off rates, see [19] and Section 5. We set as target throughput $\boldsymbol{\gamma} = (1/5, \ldots, 1/5) \in \Gamma$. Each of the three inversions methods (fixed-point iteration, Newton iteration and continuation method) for 30 iterations, and compare in each step the back-off rates and throughputs to their respective target values. We measure the error of the iterands by the Euclidean norm. The results are shown in Figure 1. Figure 1(a) plots the error in the back-off rates, and Figure 1(b) shows the error in the corresponding throughputs. Both figures show convergence of all three methods. Note that the error of the throughput estimate by the continuation method decreases linearly in the number of steps, which is related to the incremental nature of this method. More precisely, in its lth iteration this method produces an estimate for $\nu_{\boldsymbol{\theta}}(l\delta\boldsymbol{\gamma})$. If this estimate is accurate (i.e., the estimate $\boldsymbol{\nu}_{\boldsymbol{\theta}}^{(l)}$ yields a throughput vector close to $l\delta\boldsymbol{\gamma}$), the throughput error indeed decreases linearly in l.

In general, it is difficult to compare the various numerical methods, since the fixed-point method and the Newton-based methods have different computational bottlenecks. For the Newton-based methods, the initialization stage is the most cumbersome, in particular the computation of the matrix $\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}$ (or $\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}}$). The

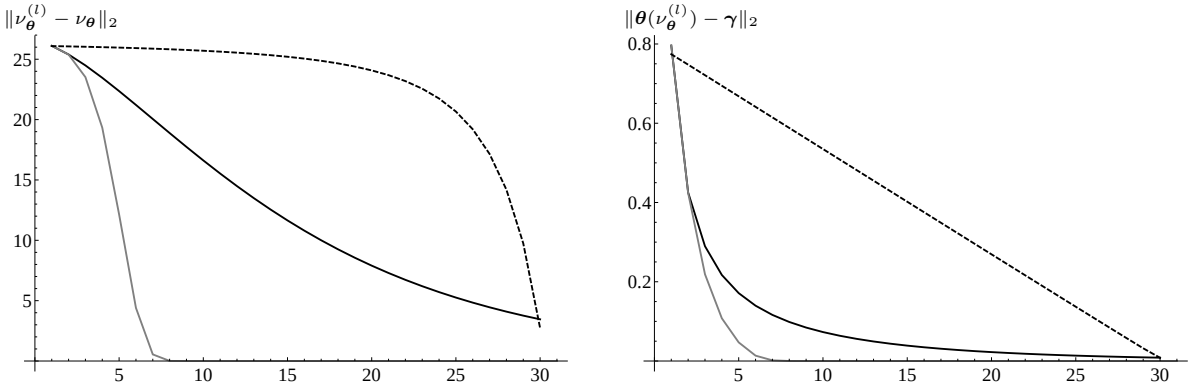

(a) Back-off rates

(b) Throughputs

Figure 1: The relative difference with the target back-off rates and throughputs for the iterands of the fixed-point iteration (black), Newton iteration (gray) and continuation method (dashed).

iteration itself has a relatively low complexity. In contrast, the fixed-point method barely requires any initialization, but its iteration stage typically takes longer than that of the Newton-based methods (when aiming for equal accuracy). Thus, either method may be best, depending on the conflict graph, target throughput and required accuracy.

4.4 Light-traffic approximation

Starting from the fixed-point equation (11) we derive an approximation for the inverse $\boldsymbol{\nu}_{\boldsymbol{\theta}}(\boldsymbol{\gamma})$ that is accurate when the elements of the normalized throughput vector $\boldsymbol{\gamma}$ are relatively small (a similar result can be obtained for the non-normalized throughput):

Proposition 2. *Let $\boldsymbol{\gamma} \in \Gamma$ and denote by $\mathcal{N}_i = \{j : \{i,j\} \in E\}$ the set of neighbors of node i. Then, as $\|\boldsymbol{\gamma}\| \downarrow 0$,*

$$\nu_{\theta,i}(\boldsymbol{\gamma}) = \gamma_i(1 + \gamma_i + \sum_{j \in \mathcal{N}_i} \gamma_j) + O\big(\|\boldsymbol{\gamma}\|^3\big), \quad i = 1, 2, \ldots, n.$$

Proof. For $\|\boldsymbol{\nu}\| \downarrow 0$, we know that $G_i(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = O\big(\|\boldsymbol{\nu}_{\boldsymbol{\theta}}\|\big)$, and $Z(\boldsymbol{\nu}_{\boldsymbol{\theta}}) = 1 + O\big(\|\boldsymbol{\nu}_{\boldsymbol{\theta}}\|\big)$. Thus we obtain from (11)

$$\nu_{\theta,i} = \gamma_i + O\big(\|\boldsymbol{\gamma}\|^2\big),$$

and hence $\boldsymbol{\nu}_{\boldsymbol{\theta}} = \boldsymbol{\gamma} + O\big(\|\boldsymbol{\gamma}\|^2\big)$. Substituting this into (11) once more, and noting that $G_i(\boldsymbol{\gamma}) = \sum_{j \neq i, j \notin \mathcal{N}_i} \gamma_j +$

$O(\|\boldsymbol{\gamma}\|^2)$, we deduce

$$\nu_{\theta,i} = \frac{\gamma_i\big(1+\sum_{j=1}^n \gamma_j + O(\|\boldsymbol{\gamma}\|^2)\big)}{1+G_i(\boldsymbol{\gamma})+O(\|\boldsymbol{\gamma}\|^2)} = \gamma_i(1+\sum_{j=1}^n \gamma_j) - \gamma_i G_i(\boldsymbol{\gamma}) + O(\|\boldsymbol{\gamma}\|^3)$$
$$= \gamma_i(1+\gamma_i+\sum_{j\in\mathcal{N}_i}\gamma_j) + O(\|\boldsymbol{\gamma}\|^3), \quad \|\boldsymbol{\gamma}\| \downarrow 0,$$

as required. □

The approximation in Proposition 2 may be interpreted by observing that a fraction of the time $\gamma_i + \sum_{j\in\mathcal{N}_i}\gamma_j + O(\|\boldsymbol{\gamma}\|^2)$ node i is prevented from activating due to either its own activity or that of one of its neighbors. The first-order approximation can be extended to an infinite-series expression by using the Lagrange-Good theorem [3]. However, the Lagrange-Good theorem in general requires cumbersome computations, and the result provides little insight.

5 Special conflict graphs

For certain specific conflict graphs, we can either find an exact expression for the fixed point $\boldsymbol{\nu_\theta}(\boldsymbol{\gamma})$, or we can decompose the graph into several components in order to reduce the complexity of the inversion methods. For this we shall exploit the fact that our model is a Markov random field. The crucial property of Markov random fields that we will use is that for any subset $S \subseteq V$, the distribution of S is determined by the state of its boundary, and is independent of all other nodes. That is, for a general conflict graph G, for any $\boldsymbol{y} \in \Omega$ and subset $S \subset V$, we have

$$\mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{V\backslash S} = \boldsymbol{y}^{V\backslash S}\} = \mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{\partial S} = \boldsymbol{y}^{\partial S}\}, \tag{14}$$

with $Y^S = (Y_i)_{i\in S}$ the components of the vector Y indexed by S, and $\partial S = \{j \in V\backslash S : \{i,j\} \in E \text{ for some } i \in S\}$, the boundary of S.

The next proposition identifies the 'fair' back-off rates that render equal throughputs for all nodes in a linear topology. In particular, we consider n nodes on a line with a k-hop interference range, i.e., each node interferes with up to k adjacent nodes to the left and to the right, $n \geq k$.

Proposition 3. *Consider the conflict graph that arises from the linear network described above, and let* $\boldsymbol{\gamma} = (\gamma, \ldots, \gamma)$ *with* $\gamma < 1/(k+1)$. *Then*

$$\nu_{\theta,i}(\boldsymbol{\gamma}) = \frac{\gamma(1-k\gamma)^{h_i-1}}{(1-(k+1)\gamma)^{h_i}},$$

with

$$h_i := \begin{cases} i, & i = 1, \ldots, k, \\ k+1, & i = k+1, \ldots, n-k, \\ n-i+1, & i = n-k+1, \ldots, n, \end{cases}$$

the number of interferers of node i minus $k-1$.

Note that $\gamma \to 1/(k+1)$ as $\nu_{\theta,i}(\boldsymbol{\gamma}) \to \infty$, i.e., the throughput approaches the maximum achievable fair throughput as the back-off rates tend to infinity. Proposition 3 was originally proven in [19], and in Appendix A.3 we provide an alternative proof of Proposition 3 based on the Markov random field representation of the stationary distribution of the joint activity state, extending the approach in Kelly [10].

Before proceeding, we first introduce some additional notation. For any subset $S \subseteq V$, we may consider a modified version of the system with the nodes in $V \setminus S$ removed, or equivalently, a system associated with a conflict graph that is the subgraph of G induced by the nodes in S and the same back-off rates. For brevity, we will call such a modified version the *system induced by S*. Denote by $Y(S)$ a random variable with the stationary distribution of the activity process in the system induced by S and by $\theta(S)$ the associated throughput vector. Moreover, for any $S \subseteq V$, $W \subseteq V \setminus S$, let

$$\Delta(S; \boldsymbol{y}^W) = S \setminus \bigcup_{i \in W : y_i = 1} \mathcal{N}_i$$

be the set of those nodes in S that are not blocked by nodes in W that are active under $\boldsymbol{y}^W$. By the Markov random field property (14), we have that $\Delta(S; \boldsymbol{y}^{V \setminus S}) := \Delta(S; \boldsymbol{y}^{\partial S})$. Finally, let us denote by Ω^S the state space restricted to S.

Recall that $\mathcal{N}_i = \{j : \{i,j\} \in E\}$ is the set of neighbors of node i in the conflict graph G. We will now apply the property (14) to show that the problem of finding the stationary distribution of S can be reduced to finding the stationary distribution of several smaller systems, by conditioning on the state of ∂S.

Proposition 4. *For any conflict graph $G = (V, E)$, $S \subseteq V$, and $\boldsymbol{y}^S \in \{0,1\}^{|S|}$,*

$$\mathbb{P}\{Y^S = \boldsymbol{y}^S\} = \sum_{\boldsymbol{y}^{\partial S} \in \Omega^{\partial S}} \mathbb{P}\{Y(\Delta(S; \boldsymbol{y}^{\partial S})) = \boldsymbol{y}^{\Delta(S; \boldsymbol{y}^{\partial S})}\} \cdot \mathrm{I}_{\{\sum_{i \in \partial S} \sum_{j \in \mathcal{N}_i} y_i y_j = 0\}} \mathbb{P}\{Y^{\partial S} = \boldsymbol{y}^{\partial S}\}.$$

The proof of Proposition 4 is given in Appendix A.4.

Proposition 4 may seem convoluted, but can very useful in certain conflict graphs for reducing the complexity of solving the above inversion problems, by choosing the set S in a smart way. For example, consider the conflict graph in Figure 2. In this case, the node set can be partitioned into two subsets V_1 and V_2 and a single node v, so $V = V_1 \cup V_2 \cup \{v\}$. The sets V_1 and V_2 are not connected, and v shares edges with nodes

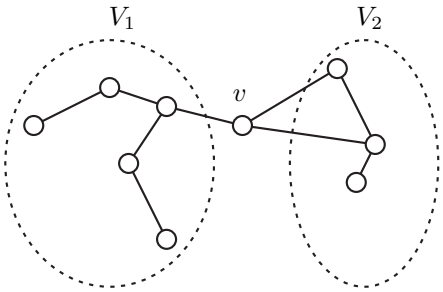

Figure 2: A decomposable graph.

in both subgraphs. We can decompose the graphs V_1 and V_2 as follows.

Corollary 1. *For any* $\boldsymbol{y} \in \{0,1\}^{|V_1|}$*, we have*

$$\mathbb{P}\{Y^{V_1} = \boldsymbol{y}^{V_1}\} = \mathbb{P}\{Y(V_1) = \boldsymbol{y}^{V_1}\}(1-\theta_v(V))\mathbb{P}\{Y(V_1 \setminus \mathcal{N}_v) = \boldsymbol{y}^{V_1 \setminus \mathcal{N}_v}\}\mathrm{I}_{\{\boldsymbol{y}^{\mathcal{N}_v}=0\}}\theta_v(V).$$

In particular, for any $i \in V_1 \cup \mathcal{N}_v$*,*

$$\theta_i(V) = \theta_i(V_1)(1-\theta_v(V)), \tag{15}$$

and for any $i \in V_1 \setminus \mathcal{N}_v$*,*

$$\theta_i(V) = \theta_i(V_1)(1-\theta_v(V)) + \theta_i(V_1 \setminus \mathcal{N}_v)\theta_v(V). \tag{16}$$

The proof of Corollary 1 is presented in Appendix A.5.

If we now substitute $\theta_v(V) = \gamma_v$ into (15) and (16), then we see that the resulting inverse problem for finding $\nu_{\theta,i}$ only depends on the nodes in V_1, and no longer requires knowledge about any node in V_2. This allows us to solve the inversion problems for V_1 and V_2 separately. Doing so considerably reduces complexity, as the number of feasible states of the induced subgraph on V_1 is much smaller than that of the entire graph. The result in Corollary 1 can also be applied when v is replaced by a clique of nodes. Naturally, when the conflict graph is disconnected, each of the components can also be handled separately.

6 Conclusion and outlook

In this paper we have established *global invertibility* of both the non-normalized and normalized throughput function for wireless random-access networks with general conflict graphs. This fundamental result, presented in Section 3, says that we can match any achievable throughput vector with a unique vector of back-off rates that will lead to that target throughput vector. This result allows us, for example, to compute the back-off

rates that give equal throughputs, or instead to create various user classes by designing the back-off rates to give certain nodes higher throughput than others.

We presented in Section 4 several algorithms for determining the (inverse) back-off rates. The implementation of the inversion algorithms involves the computation of the normalizing constant $Z(\boldsymbol{\nu})$, the (inverse of) the functional matrix $\frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\nu}}$, and the functions G_i in (7). These require the enumeration of the entire state space Ω, which essentially boils down to counting all independent sets of the conflict graph, a problem which is known to be computationally cumbersome for large graphs. An important task for future research is to find ways of dealing with this curse of dimensionality. One possible approach is to exploit the structure of the conflict graphs and using the theory of Markov random fields, as was done in Section 5.

Acknowledgment

In the course of this project we have benefitted from discussions with many colleagues, including Dee Denteneer, Alexandre Proutière, Devravat Shah, Patrick Thiran and Jean Walrand.

References

[1] R.R. Boorstyn and A. Kershenbaum. Throughput analysis of multihop packet radio. In *Proc. ICC*, pages 1361–1366, 1980.

[2] R.R. Boorstyn, A. Kershenbaum, B. Maglaris, and V. Sahin. Throughput analysis in multihop CSMA packet radio networks. *IEEE Trans. Commun.*, 35(3):267–274, 1987.

[3] N.G. de Bruijn. The Lagrange-Good inversion formula and its application to integral equations. *Journal of Mathematical Analysis and Applications*, 92(2):397–409, 1983.

[4] M. Durvy, O. Dousse, and P. Thiran. Modeling the 802.11 protocol under different capture and sensing capabilities. In *Proc. Infocom*, pages 2356–2360, Anchorage AK, May 6–12 2007.

[5] M. Durvy, O. Dousse, and P. Thiran. Self-organization properties of CSMA/CA systems and their consequences on fairness. *IEEE Trans. Inform. Theory*, 55(3):931–943, 2009.

[6] M. Durvy and P. Thiran. A packing approach to compare slotted and non-slotted medium access control. In *Proc. Infocom*, Barcelona, Spain, 2006.

[7] L. Jiang, D. Shah, J. Shin, and J. Walrand. Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inform. Theory*, 56(12):6182–6207, 2010.

[8] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Trans. Netw.*, 18(3):960–972, 2010.

[9] F.P. Kelly. *Reversibility and Stochastic Networks*. Wiley, Chichester York, 1979.

[10] F.P. Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B*, 47(3):379–395, 1985.

[11] F.P. Kelly. Loss networks. *Ann. Appl. Prob.*, 1(3):319–378, 1991.

[12] A. Kershenbaum, R.R. Boorstyn, and M. Chen. An algorithm for evaluation of throughput in multihop packet radio. *IEEE J. Sel. Areas Commun.*, 5(6):1003–1012, 1987.

[13] J. Liu, Y. Yi, A. Proutière, M. Chiang, and H.V. Poor. Maximizing utility via random access without message passing. Technical Report 2008-128, Microsoft Research, 2008.

[14] P. Marbach and A. Eryilmaz. A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks. In *Proc. 46th Annual Allerton Conf. Commun., Control and Comp.*, 2008.

[15] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: an efficient randomized protocol for content resolution. In *Proc. ACM SIGMETRICS/Performance*, pages 133–144, 2009.

[16] Y. Suhov and U.A. Rozikov. A hard-core model on a Cayley tree: An example of a loss network. *Queueing Systems*, 46(1–2):197–212, 2004.

[17] P.M. van de Ven, S.C. Borst, D. Denteneer, A.J.E.M Janssen, and J.S.H. van Leeuwaarden. Equalizing throughputs in random-access networks. 2010. Performance Evaluation Review, Special issue - Proc. MAMA (to appear).

[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, 2010.

[19] 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 (to appear)*, 2011.

[20] X. Wang and K. Kar. Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks. In *Proc. Infocom*, pages 23–34, 2005.

[21] F.F. Wu and C.A. Desoer. Global inverse function theorem. *IEEE Trans. Circuit Theory*, 19(2):199–201, 1972.

[22] S. Zachary and I. Ziedins. Loss networks and Markov random fields. *J. Appl. Prob.*, 36(2):403–414, 1999.

A Remaining proofs

A.1 Proof of Theorem 1

Rather than showing invertibility of $\boldsymbol{\eta}$ itself, we consider the mapping

$$\mathbf{f}(\boldsymbol{x}) = \ln \boldsymbol{\eta}(\mathrm{e}^{\boldsymbol{x}}), \quad \boldsymbol{x} = (x_1, x_2, \ldots, x_n)^T \in \mathbb{R}^n,$$

with $\mathrm{e}^{\boldsymbol{x}} = (\mathrm{e}^{x_1}, \mathrm{e}^{x_2}, \ldots, \mathrm{e}^{x_n})^T$ and $\ln \mathbf{y} = (\ln y_1, \ldots, \ln y_n)^T$ for $\mathbf{y} = (y_1, y_2, \ldots, y_n)^T \in \mathbb{R}^n$. Because ln and exp are invertible, global invertibility of $\mathbf{f}$ and $\boldsymbol{\eta}$ is equivalent.

By the main result in [21] we have that $\mathbf{f}$ is globally invertible if and only if (i) $\mathbf{f}$ is locally invertible and (ii) $\max_i |f_i(\boldsymbol{x})| \to \infty$ as $\max_i |x_i| \to \infty$.

To show that condition (i) holds, it suffices to show that the functional matrix

$$\frac{\partial \mathbf{f}}{\partial \boldsymbol{x}} = \Big(\frac{\partial f_i}{\partial x_j}\Big)_{\substack{i=1,\ldots,n \\ j=1,\ldots,n}},$$

is non-singular at any point $\boldsymbol{x} \in \mathbb{R}^n$. Observe that

$$\frac{\partial f_i}{\partial x_j} = \frac{1}{\eta_i(\mathrm{e}^{\boldsymbol{x}})} \frac{\partial \eta_i(\mathrm{e}^{\boldsymbol{x}})}{\partial \mathrm{e}^{x_j}} \frac{\partial \boldsymbol{x}}{\partial x_j} = \frac{1}{\eta_i(\boldsymbol{\nu})} \frac{\partial \eta_i(\boldsymbol{\nu})}{\partial \nu_j} \nu_j.$$

Thus

$$\frac{\partial \mathbf{f}}{\partial \boldsymbol{x}} = \mathrm{diag}\Big(\frac{1}{\eta_1(\boldsymbol{\nu})}, \ldots, \frac{1}{\eta_n(\boldsymbol{\nu})}\Big) \frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}} \mathrm{diag}(\nu_1, \ldots \nu_n), \tag{17}$$

with

$$\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}} = \big(\frac{\partial \eta_i}{\partial \nu_j}\big)_{\substack{i=1,\ldots,n \\ j=1,\ldots,n}}.$$

Because $\nu_1, \ldots, \nu_n > 0$, both diagonal matrices in (17) are non-singular, and we only have to verify that $\partial \boldsymbol{\eta} / \partial \boldsymbol{\nu}$ is non-singular as well.

By taking the derivative of $\Lambda_k(\boldsymbol{\nu})$, see (1), with respect to ν_j, we get

$$\frac{\partial \Lambda_k(\boldsymbol{\nu})}{\partial \nu_j} = \frac{1}{\nu_j} \Lambda_k(\boldsymbol{\nu}) \mathrm{I}_{\{\omega_{k,j}=1\}}, \quad k = 0, 1, \ldots, K, \ j = 1, \ldots, n.$$

Consequently,

$$\frac{\partial \eta_i}{\partial \nu_j} = \frac{1}{\nu_j} \sum_{k=0}^{K} \omega_{k,i} \omega_{k,j} \Lambda_k(\boldsymbol{\nu}), \quad i,j = 1,2,\ldots,n. \tag{18}$$

Thus the functional matrix $\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}}$ may be written as

$$\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}} = \mathbf{P}(\boldsymbol{\nu}) \mathbf{D}(\boldsymbol{\nu}),$$

with

$$\mathbf{P}(\boldsymbol{\nu}) = \sum_{k=0}^{K} \Lambda_k(\boldsymbol{\nu}) \boldsymbol{\omega}_k \boldsymbol{\omega}_k^T,$$

and

$$\mathbf{D}(\boldsymbol{\nu}) = \operatorname{diag}(\nu_1^{-1}, \ldots, \nu_n^{-1}). \tag{19}$$

The matrix $\mathbf{P}$ is positive definite since $\Lambda_k(\boldsymbol{\nu}) > 0$, $\boldsymbol{\omega}_k = \mathbf{e}_k$, $k = 1,2,\ldots,n$. Therefore, $\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}}$ is non-singular, as required.

In order to verify condition (ii), we write $\boldsymbol{\eta}(\mathrm{e}^{\boldsymbol{x}})$ as

$$\eta_i(\mathrm{e}^{\boldsymbol{x}}) = \mathrm{e}^{x_i} \Big(1 + \mathrm{e}^{-x_i} \sum_{k=n+1}^{K} \Lambda_k(\mathrm{e}^{\boldsymbol{x}}) \omega_{k,i} \Big). \tag{20}$$

Let

$$m = \min_i \mathrm{e}^{x_i}, \quad M = \max_i \mathrm{e}^{x_i},$$
$$a = -\min_i x_i, \quad b = \max_i x_i.$$

It is seen from (1) and (20) that

$$\max_i \eta_i(\mathrm{e}^{\boldsymbol{x}}) \geq M = \mathrm{e}^b, \tag{21}$$

$$\min_i \eta_i(\mathrm{e}^{\boldsymbol{x}}) \leq m(1 + (K-n)M^{n-1}) = \mathrm{e}^{-a}(1 + (K-n)\mathrm{e}^{(n-1)b}). \tag{22}$$

Assume that $\max_i |x_i| = \max\{a,b\} \to \infty$. We need to show that $\max_i |f_i(\boldsymbol{x})| \to \infty$ as well.

When $b \geq a$ we have

$$\max_i |f_i(\mathrm{e}^{\boldsymbol{x}})| \geq b = \max\{a,b\}. \tag{23}$$

When $b \leq a$, we see from (21) and (22) that

$$\max_i |f_i(\boldsymbol{x})| \geq \max\{b, a - \ln(1 + (K-n)\mathrm{e}^{(n-1)b})\} \geq \max\{b, a - A - Bb\},$$

for some $A, B > 0$ only depending on K, n. Now

$$\min_{0 \leq b^* \leq a} \max\{b^*, a - A - Bb^*\} \geq b(a),$$

with $b(a)$ the solution of $b = a - A - Bb$, i.e., $b(a) = \frac{a-A}{B+1}$. Hence, when $a \geq b$

$$\max_i |f_i(\mathrm{e}^{\boldsymbol{x}})| \geq \frac{a-A}{B+1} = \frac{\max\{a.b\} - A}{B+1}. \tag{24}$$

From (23) and (24) we see that

$$\max_i |f_i(\mathrm{e}^{\boldsymbol{x}})| \to \infty$$

as $\max\{a, b\} \to \infty$, and the proof is complete.

A.2 Proof of Lemma 1

In order to prove this lemma, we compute some derivatives. We have, compare (18),

$$\frac{\partial Z(\boldsymbol{\nu})}{\partial \nu_j} = \sum_{k=0}^{K} \omega_{k,j} \frac{1}{\nu_j} \Lambda_k(\boldsymbol{\nu}) = \frac{1}{\nu_j} \eta_j(\boldsymbol{\nu}), \quad j = 1, \ldots, n.$$

Recall from (18) that

$$\frac{\partial \eta_i}{\partial \nu_j} = \frac{1}{\nu_j} \sum_{k=0}^{K} \omega_{k,i} \omega_{k,j} \Lambda_k(\boldsymbol{\nu}), \quad i, j = 1, 2, \ldots, n.$$

Differentiating $\boldsymbol{\eta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) = s\boldsymbol{c}$ with respect to s, we see that

$$\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) \cdot \boldsymbol{\nu}'_{\boldsymbol{\eta}}(s\boldsymbol{c}) = \boldsymbol{c},$$

i.e., that

$$\boldsymbol{\nu}'_{\boldsymbol{\eta}}(s\boldsymbol{c}) = \Big(\frac{\mathrm{d}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))_1}{\mathrm{d}s}, \ldots, \frac{\mathrm{d}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))_n}{\mathrm{d}s} \Big)^T = \Big(\frac{\partial \boldsymbol{\eta}}{\partial \boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) \Big)^{-1} \boldsymbol{c}. \tag{25}$$

Moreover, we have

$$\frac{\partial\theta_i}{\partial\nu_j} = \frac{\partial}{\partial\nu_j}\Big(\frac{\eta_i(\boldsymbol{\nu})}{Z(\boldsymbol{\nu})}\Big) = \frac{1}{Z^2(\boldsymbol{\nu})}\Big(\frac{\partial\eta_i(\boldsymbol{\nu})}{\partial\nu_j}Z(\boldsymbol{\nu}) - \eta_i(\boldsymbol{\nu})\frac{\partial Z(\boldsymbol{\nu})}{\partial\nu_j}\Big)$$
$$= \frac{1}{\nu_j}\Big(\sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\omega_{k,i}\omega_{k,j} - \sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\omega_{k,i}\sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\omega_{k,j}\Big).$$

Note that $\sum_{k=0}^{K}\pi_k(\boldsymbol{\nu}) = 1$ and that $\boldsymbol{\theta}(\boldsymbol{\nu}) = \sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\boldsymbol{\omega}_k$. Hence we have

$$\frac{\partial\theta_i}{\partial\nu_j} = \sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\big(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu})\big)_i\big(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu})\big)^T_j\frac{1}{\nu_j}. \tag{26}$$

So

$$\frac{\partial\boldsymbol{\theta}}{\partial\boldsymbol{\nu}} = \mathbf{Q}(\boldsymbol{\nu})\mathbf{D}(\boldsymbol{\nu}), \tag{27}$$

with D the diagonal matrix in (19) and

$$\mathbf{Q}(\boldsymbol{\nu}) = \sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\big(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu})\big)\big(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu})\big)^T. \tag{28}$$

We shall show below that $\mathbf{Q}(\boldsymbol{\nu})$ is positive definite. Assuming this, we compute from $\boldsymbol{\theta}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})) = \frac{s}{Z(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))}\boldsymbol{c}$, for any $s > 0$

$$\frac{\partial\boldsymbol{\theta}}{\partial\boldsymbol{\nu}}(\boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c}))\boldsymbol{\nu}'_{\boldsymbol{\eta}}(s\boldsymbol{c}) = f'_{\boldsymbol{c}}(s)\boldsymbol{c}. \tag{29}$$

By (25) we have that $\boldsymbol{\nu}'_{\boldsymbol{\eta}}(s\boldsymbol{c}) \neq \mathbf{0}$ and by the fact that $\mathbf{Q}(\boldsymbol{\nu})$ is positive definite and (27) we have that $\frac{\partial\boldsymbol{\theta}}{\partial\boldsymbol{\nu}}$ is non-singular at $\boldsymbol{\nu} = \boldsymbol{\nu}_{\boldsymbol{\eta}}(s\boldsymbol{c})$. Hence, the left-hand side of (29) is a non-zero vector and so $f'_{\boldsymbol{c}}(s)\boldsymbol{c} \neq \mathbf{0}$. Hence $f'_{\boldsymbol{c}}(s) \neq 0$ for any $s > 0$. Since $f_{\boldsymbol{c}}(0) = 0$, $f_{\boldsymbol{c}}(s) > 0$ for $s > 0$, the claim follows.

It remains to show that $\mathbf{Q}(\boldsymbol{\nu})$ is positive definite. Assume $\boldsymbol{y} \in \mathbb{R}^n$ is such that $\mathbf{Q}(\boldsymbol{\nu})\boldsymbol{y} = \mathbf{0}$. Then

$$0 = \boldsymbol{y}^T\mathbf{Q}(\boldsymbol{\nu})\boldsymbol{y} = \sum_{k=0}^{K}\pi_k(\boldsymbol{\nu})\big|(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu}))^T\boldsymbol{y}\big|^2,$$

and so, as $\pi_k(\boldsymbol{\nu}) > 0$, $k = 0, 1, \ldots, K$, we have

$$(\boldsymbol{\omega}_k - \boldsymbol{\theta}(\boldsymbol{\nu}))^T\boldsymbol{y} = 0, \quad k = 0, 1, \ldots, K,$$

i.e.,

$$\boldsymbol{\omega}_k^T\boldsymbol{y} = \boldsymbol{\theta}(\boldsymbol{\nu})^T\boldsymbol{y}, \quad k = 0, 1, \ldots, K. \tag{30}$$

Since $\boldsymbol{\omega}_0 = \mathbf{0}$, we get $\boldsymbol{\theta}(\boldsymbol{\nu})^T\boldsymbol{y} = 0$ from (30) with $k = 0$. Then, for $k = 1, \ldots, n$, it follows from $\boldsymbol{\omega}_k = \mathbf{e}_k$

and (30) that

$$y_k = \boldsymbol{\omega}_k^T \boldsymbol{y} = \boldsymbol{\theta}(\boldsymbol{\nu})^T \boldsymbol{y} = 0.$$

Hence $\boldsymbol{y} = \mathbf{0}$. We conclude that $\mathbf{Q}(\boldsymbol{\nu})$ is non-singular, and then from (28) it is seen that $\mathbf{Q}(\boldsymbol{\nu})$ is positive definite.

A.3 Proof of Proposition 3

For conciseness, denote

$$\psi_i := \mathbb{P}\{Y_{i-k}, \ldots, Y_{i-1} = 0\}, \quad i = k+1, \ldots, n+1,$$

and

$$a_i := \mathbb{P}\{Y_i = 0 | Y_{i-k}, \ldots, Y_{i-1} = 0\}, \quad i = k+1, \ldots, n.$$

By definition,

$$\begin{aligned}\theta_i &= \mathbb{P}\{Y_i = 1\} = \mathbb{P}\{Y_i = 1, Y_{i-k}, \ldots, Y_{i-1} = 0\} \\ &= \mathbb{P}\{Y_i = 1 | Y_{i-k}, \ldots, Y_{i-1} = 0\}\mathbb{P}\{Y_{i-k}, \ldots, Y_{i-1} = 0\} = (1 - a_i)\psi_i, \end{aligned} \tag{31}$$

for all $i = k+1, \ldots, n$.

The idea of the proof is to consider probabilities of the form $\mathbb{P}\{Y_i = 1, Y_j = 0, j \in \mathcal{N}_i\}$ and $\mathbb{P}\{Y_i = 1, Y_j = 1, j \in \mathcal{N}_i\}$ and use two different relationships between these in order to obtain a set of equations for the coefficients a_i.

First of all, it follows from the product form of the stationary distribution (or the local balance property) that

$$\mathbb{P}\{Y_i = 1, Y_j = 0, j \in \mathcal{N}_i\} = \nu_i \mathbb{P}\{Y_i = 0, Y_j = 0, j \in \mathcal{N}_i\},$$

for all $i = 1, \ldots, n$.

The second relationship between these two probabilities follows from the Markov random field representation of the stationary distribution.

Specifically, for all $i = 1, \ldots, k$, we may write

$$\begin{aligned}\mathbb{P}\{Y_i = 0, Y_j = 0, j \in \mathcal{N}_i\} =& \mathbb{P}\{Y_1, \ldots, Y_k = 0\} \prod_{l=k+1}^{i} \mathbb{P}\{Y_i = 0 | Y_1, \ldots, Y_{i-1} = 0\} \\ &= \psi_{k+1} \prod_{l=k+1}^{i} \mathbb{P}\{Y_i = 0 | Y_{i-k}, \ldots, Y_{i-1} = 0\} = \psi_{k+1} \prod_{l=k+1}^{i} a_i.\end{aligned}$$

For all $i = k+1, \ldots, n$, we may write

$$\mathbb{P}\{Y_i = 0, Y_j = 0, j \in \mathcal{N}_i\} = \mathbb{P}\{Y_{i-k}, \ldots, Y_{i-1} = 0\} \prod_{l=i}^{\min\{i+k,n\}} \mathbb{P}\{Y_l = 0 | Y_{i-k}, \ldots, Y_{l-1} = 0\}$$
$$= \psi_i \prod_{l=i}^{\min\{i+k,n\}} \mathbb{P}\{Y_l = 0 | Y_{l-k}, \ldots, Y_{l-1} = 0\} = \psi_i \prod_{l=i}^{\min\{i+k,n\}} a_i,$$

and

$$\mathbb{P}\{Y_i = 1, Y_j = 0, j \in \mathcal{N}_i\} = \mathbb{P}\{Y_{i-k}, \ldots, Y_{i-1} = 0\} \mathbb{P}\{Y_i = 1 | Y_j = 0, j \in \mathcal{N}_i^-\}$$
$$\prod_{l=i+1}^{\min\{i+k,n\}} \mathbb{P}\{Y_l = 0 | Y_i = 1, Y_{i-k}, \ldots, Y_{i-1}, Y_{i+1}, \ldots, Y_{l-1} = 0\} = \psi_i (1 - a_i),$$

yielding

$$\frac{\mathbb{P}\{Y_i = 1, Y_j = 0, j \in \mathcal{N}_i\}}{1 - a_i} = \frac{\mathbb{P}\{Y_i = 0, Y_j = 0, j \in \mathcal{N}_i\}}{a_i \prod_{j=i+1}^{\min\{i+k,n\}} a_j}.$$

Now observe that $\psi_i + \sum_{j=i-k}^{i-1} \theta_j = 1$ for all $i = k+1, \ldots, n$, and in particular $\psi_{k+1} + \sum_{j=1}^{k} \theta_j = 1$. Combining the above two sets of equations, we obtain

$$\theta_i = \nu_i (1 - \sum_{j=1}^{k} \theta_j) \prod_{j=k+1}^{i+k} a_j, \tag{32}$$

for $i = 1, \ldots, k$, while

$$1 - a_i = \nu_i a_i \prod_{j=i+1}^{\min\{i+k,n\}} a_j, \tag{33}$$

for all $i = k+1, \ldots, n$.

A solution to (33) is provided by $a_i = a$ and $\nu_i = (1-a)a^{-h_i}$, or equivalently $\nu_i = \alpha(1+\alpha)^{h_i - 1}$, with $\alpha = 1/a - 1 > 0$. Taking $\nu_i = (1-a)a^{-h_i}$, $i = 1, \ldots, k$, we obtain from (32)

$$\theta_1 = \cdots = \theta_k = \theta = (1-a)(1-k\theta),$$

i.e.,

$$\theta_1 = \cdots = \theta_k = \theta = \frac{1-a}{1+k(1-a)} = \frac{\alpha}{1+(k+1)\alpha},$$

and (31) then yields

$$\theta = \frac{1-a}{1+k(1-a)} = \frac{\alpha}{1+(k+1)\alpha} \quad \text{for all} \quad i = k+1, \ldots, n.$$

Noting that

$$a = \frac{1-(k+1)\theta}{1-k\theta} \quad \text{or} \quad \alpha = \frac{\theta}{1-(k+1)\theta}$$

then completes the proof.

A.4 Proof of Proposition 4

The product form of the stationary distribution implies

$$\begin{aligned}
&\mathbb{P}\{Y^S = \boldsymbol{y}^S \mid Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\} \\
&= \frac{\mathbb{P}\{Y^S = \boldsymbol{y}^S, Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\}}{\mathbb{P}\{Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\}} = \frac{\mathbb{P}\{Y = \boldsymbol{y}\}}{\sum_{(\boldsymbol{x}^S, \boldsymbol{y}^{V\setminus S})\in\Omega} \mathbb{P}\{Y^S = \boldsymbol{x}^S, Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\}} \\
&= \frac{Z^{-1}\prod_{j=1}^{n} \nu_j^{y_j}}{Z^{-1}\sum_{(\boldsymbol{x}^S, \boldsymbol{y}^{V\setminus S})\in\Omega} \prod_{i\in S} \nu_i^{x_i} \prod_{i\notin S} \nu_j^{y_j}} = K^{-1}(S; \boldsymbol{y}^{V\setminus S}) \prod_{j\in S} \nu_j^{y_j},
\end{aligned}$$

for any $\boldsymbol{y} \in \Omega$, with

$$K(S; \boldsymbol{y}^{V\setminus S}) = \sum_{\boldsymbol{x}^S : (\boldsymbol{x}^S, \boldsymbol{y}^{V\setminus S})\in\Omega} \prod_{i\in S} \nu_i^{x_i}.$$

Note that

$$\prod_{j\in S} \nu_j^{y_j} = \prod_{j\in\Delta(S;\boldsymbol{y}^{\partial S})} \nu_j^{y_j}$$

for any $y \in \Omega$. Likewise,

$$\prod_{i\in S} \nu_i^{x_i} = \prod_{i\in\Delta(S;\boldsymbol{y}^{\partial S})} \nu_i^{x_i}$$

for any $(\boldsymbol{x}^S, \boldsymbol{y}^{V\setminus S}) \in \Omega$.

Let $\Omega(S; \boldsymbol{y}^{V\setminus S})$ be the collection of independent sets in the subgraph of G induced by $\Delta(S; \boldsymbol{y}^{\partial S})$.

It is easily verified that $(\boldsymbol{x}^S, \boldsymbol{y}^{V\setminus S}) \in \Omega$ if and only if it holds that $\boldsymbol{x}^{\Delta(S;\boldsymbol{y}^{\partial S})} \in \Omega(S; \boldsymbol{y}^{\partial S})$. It follows that the coefficient

$$K(S; \boldsymbol{y}^{V\setminus S}) = \sum_{\boldsymbol{x}^{\Delta(S;\boldsymbol{y}^{\partial S})}\in\Omega(S;\boldsymbol{y}^{\partial S})} \prod_{i\in\Delta(S;\boldsymbol{y}^{\partial S})} \nu_i^{x_i},$$

and thus corresponds to the normalization constant of the system induced by $\Delta(S; \boldsymbol{y}^{\partial S})$.

We conclude that

$$\mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\} = \frac{\prod_{j\in\Delta(S;\boldsymbol{y}^{\partial S})} \nu_j^{y_j}}{\sum_{\boldsymbol{x}^{\Delta(S;\boldsymbol{y}^{\partial S})}\in\Omega(S;\boldsymbol{y}^{\partial S})} \prod_{i\in\Delta(S;\boldsymbol{y}^{\partial S})} \nu_i^{x_i}}$$
$$= \mathbb{P}\{Y(\Delta(S;\boldsymbol{y}^{\partial S})) = \boldsymbol{y}^{\Delta(S;\boldsymbol{y}^{\partial S})}\},$$

for any $\boldsymbol{y} \in \Omega$. Informally speaking, the distribution of the activity state of the nodes in S in the original system, conditional on the activity states of the remaining nodes, equals the stationary distribution of the system induced by $\Delta(S;\boldsymbol{y}^{\partial S})$. Since $\Delta(S;\boldsymbol{y}^{\partial S})$ only depends on $\boldsymbol{y}^{V\setminus S}$ through $\boldsymbol{y}^{\partial S}$, it further follows that $\mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\} = \mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{\partial S} = \boldsymbol{y}^{\partial S}\}$. This corroborates the fact that the stationary distribution is a Markov random field with a neighborhood structure defined by the conflict graph G.

Now observe that

$$\mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{V\setminus S} = \boldsymbol{y}^{V\setminus S}\} = \mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{\partial S} = \boldsymbol{y}^{\partial S}\} = 0$$

unless

$$\sum_{i\in\partial S} \sum_{j\in\mathcal{N}_i} y_i y_j = 0.$$

Thus we may write

$$\mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{\partial S} = \boldsymbol{y}^{\partial S}\} = \mathbb{P}\{Y(\Delta(S;\boldsymbol{y}^{\partial S})) = \boldsymbol{y}^{\Delta(S;\boldsymbol{y}^{\partial S})}\} \mathrm{I}_{\{\sum_{i\in\partial S}\sum_{j\in\mathcal{N}_i} y_i y_j = 0\}},$$

for all $\boldsymbol{y} \in \{0,1\}^V$, rather than just $\boldsymbol{y} \in \Omega$.

We deduce that

$$\mathbb{P}\{Y^S = \boldsymbol{y}^S\} = \sum_{\boldsymbol{y}^{\partial S}} \mathbb{P}\{Y^S = \boldsymbol{y}^S | Y^{\partial S} = \boldsymbol{y}^{\partial S}\} \mathbb{P}\{Y^{\partial S} = \boldsymbol{y}^{\partial S}\}$$
$$= \sum_{y^{\partial S}} \mathbb{P}\{Y(\Delta(S;\boldsymbol{y}^{\partial S})) = \boldsymbol{y}^{\Delta(S;\boldsymbol{y}^{\partial S})}\} \mathrm{I}_{\{\sum_{i\in\partial S}\sum_{j\in\mathcal{N}_i} y_i y_j = 0\}} \mathbb{P}\{Y^{\partial S} = \boldsymbol{y}^{\partial S}\}. \tag{34}$$

A.5 Proof of Corollary 1

For the specific graph under consideration, apply Proposition 4 with $S = V_1$ so that $\partial S = \{v\}$. We have

$$\begin{aligned}\mathbb{P}\{Y^{V_1} = \boldsymbol{y}^{V_1} | Y^{V\setminus V_1} = \boldsymbol{y}^{V\setminus V_1}\} &= \mathbb{P}\{Y^{V_1} = \boldsymbol{y}^{V_1} | Y_v = y_v\} \mathrm{I}_{\{y_v \sum_{j\in\mathcal{N}_v} y_j = 0\}} \\ &= \begin{cases} \mathbb{P}\{Y(V_1) = \boldsymbol{y}^{V_1}\}, & y_v = 0, \\ \mathbb{P}\{Y(V_1 \setminus \mathcal{N}_v) = \boldsymbol{y}^{V_1\setminus\mathcal{N}_v}\} \mathrm{I}_{\{\boldsymbol{y}^{\mathcal{N}_v} = 0\}}, & y_v = 1, \end{cases}\end{aligned}$$

with $\mathbb{P}\{Y_v = 1\} = 1 - \mathbb{P}\{Y_v = 0\} = \theta_v(V)$, so that (34) reduces to

$$\mathbb{P}\{Y^{V_1} = \boldsymbol{y}^{V_1}\} = \mathbb{P}\{Y(V_1) = \boldsymbol{y}^{V_1}\}(1 - \theta_v(V)) + \mathbb{P}\{Y(V_1 \setminus \mathcal{N}_v) = \boldsymbol{y}^{V_1\setminus\mathcal{N}_v}\} \mathrm{I}_{\{\boldsymbol{y}^{\mathcal{N}_v} = 0\}} \theta_v(V).$$

In particular, we obtain

$$\begin{aligned}\theta_i(V) &= \sum_{\boldsymbol{y}^{V_1} : y_i = 1} \mathbb{P}\{Y^{V_1} = \boldsymbol{y}^{V_1}\} \\ &= \sum_{\boldsymbol{y}^{V_1} : y_i = 1} \left[\mathbb{P}\{Y(V_1) = \boldsymbol{y}^{V_1}\}(1 - \theta_v(V)) + \mathbb{P}\{Y(V_1 \setminus \mathcal{N}_v) = \boldsymbol{y}^{V_1\setminus\mathcal{N}_v}\} \mathrm{I}_{\{\boldsymbol{y}^{\mathcal{N}_v} = 0\}} \theta_v(V)\right] \\ &= \begin{cases} \theta_i(V_1)(1 - \theta_v(V)), & i \in \mathcal{N}_v, \\ \theta_i(V_1)(1 - \theta_v(V)) + \theta_i(V_1 \setminus \mathcal{N}_v)\theta_v(V), & i \notin \mathcal{N}_v. \end{cases}\end{aligned}$$