



# Temporal starvation in multi-channel CSMA networks: an analytical framework

Alessandro Zocca<sup>1</sup> 

Received: 10 November 2018  
© Springer Science+Business Media, LLC, part of Springer Nature 2019

## Abstract

In this paper, we consider a stochastic model for a frequency-agile CSMA protocol for wireless networks where multiple orthogonal frequency channels are available. Even when the possible interference on the different channels is described by different conflict graphs, we show that the network dynamics can be equivalently described as that of a single-channel CSMA algorithm on an appropriate virtual network. Our focus is on the asymptotic regime in which the network nodes try to activate aggressively in order to achieve maximum throughput. Of particular interest is the scenario where the number of available channels is not sufficient for all nodes of the network to be simultaneously active and the well-studied temporal starvation issues of the single-channel CSMA dynamics persist. For most networks, we expect that a larger number of available channels should alleviate these temporal starvation issues. However, we prove that the aggregate throughput is a non-increasing function of the number of available channels. To investigate this trade-off that emerges between aggregate throughput and temporal starvation phenomena, we propose an analytic framework to study the transient dynamics of multi-channel CSMA networks by means of first hitting times. Our analysis further reveals that the mixing time of the activity process does not always correctly characterize the temporal starvation in the multi-channel scenario and often leads to pessimistic performance estimates.

**Keywords** Random-access networks · Performance evaluation · Throughput analysis · Markov process · Hitting times · Mixing times · Partial  $q$ -coloring

**Mathematics Subject Classification** 60J27 · 60K35 · 68M10 · 68M20 · 90B15 · 90B18

---

✉ Alessandro Zocca  
azocca@caltech.edu

<sup>1</sup> California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA

## 1 Introduction

The carrier-sense multiple-access (CSMA) algorithm is a popular distributed medium-access control mechanism, various incarnations of which are currently implemented in IEEE 802.11 WiFi networks. CSMA is a random-access algorithm, in the sense that it relies on randomness to both avoid simultaneous transmissions and share the medium in the most efficient way. Being intrinsically a randomized scheme, many stochastic models have been developed in the literature to study its performance in terms of throughput, stability, delay, and spatial fairness (see, for example, [27] and references therein for a detailed overview).

It is well known that the delay performance of CSMA algorithms can be rather poor and much worse than for other mechanisms, such as MaxWeight. One of the root causes for these poor delay performances has been identified in the *temporal starvation phenomenon*: Even in scenarios where all nodes have a good long-term throughput, they may individually experience long sequences of transmissions in rapid succession, interspersed with extended periods of starvation. These starvation effects are particularly pronounced when the nodes become more aggressive in trying to activate, which is the regime in which the network should operate to achieve maximum throughput.

Most of the literature focuses on single-channel CSMA algorithms, in which all the transmissions occur on the same frequency. In this paper, we consider the natural generalization in which multiple orthogonal frequency channels are available; various variations of this multi-channel CSMA model have been studied in [1,2,4,5,20,22,23]. Our ultimate goal is understanding whether temporal starvation effects can be effectively mitigated by the usage of multiple channels, as suggested in [20].

In order to make a fair comparison, we do not assume that additional channels can be added to the existing one, but rather that the total available wireless spectrum can be divided into  $C$  non-overlapping channels with smaller capacity on which nearby nodes can transmit without interfering with each other. In this way, at the cost of a potentially lower throughput, the starvation effects in the network can be alleviated or even eliminated and in this way obtain substantial improvements in delay performance. The aim of the present paper is to investigate this non-trivial trade-off, and for this reason we analyze the aggregate throughput of multi-channel CSMA networks and develop an analytic framework for quantifying temporal starvation effects for these networks.

Temporal starvation in the context of multi-channel CSMA networks has been mostly studied by means of mixing times [20]. The large majority of these results, however, assumes either that the activation rate is very small or that there are enough channels for all the nodes to be active simultaneously, so that existing results for multi-color Glauber dynamics can be exploited. The first crucial difference of our approach is that we are mostly interested in the scenario in which the number of available channels does not allow simultaneous activity of all the nodes. This creates complex network dynamics in which nodes still compete for transmission, and phenomena like temporal starvation and spatial unfairness persist.

Mixing times have often been successfully used to study performances of single-channel CSMA-like algorithms [7,8,11,12,14–16,20,24,25,28]. However, in

the present paper, instead of focusing on mixing times, inspired by [18], we study temporal starvation by means of hitting times of the Markov process describing the CSMA dynamics, which, as we will show later, enable a more accurate performance evaluation analysis than mixing times in the multi-channel scenario.

We focus on the regime in which the activation rate grows large, in which the CSMA dynamics favor the activity states with a maximum number of active nodes, to which we will refer as *dominant states*. These activity states play a crucial role in our analysis, since the timescales at which transitions between them occur are intimately related to the magnitude of the temporal starvation effects.

Our contributions can be summarized as follows:

- We show that the multi-channel CSMA dynamics can be represented as single-channel CSMA dynamics on an appropriate virtual network, even when the possible interference on the different channels is described by different conflict graphs. This equivalent representation allows us to immediately derive the stationary distribution of multi-channel CSMA dynamics and to prove its insensitivity to back-off and transmission time distributions.
- We prove various properties of the asymptotic aggregate throughput and, in particular, that is a non-increasing function of the number  $C$  of available channels.
- We show how temporal starvation can be evaluated in the high-activation limit by studying the expected transition times between the dominant activity states. Furthermore, we characterize the timescale of temporal starvation phenomena in terms of the structure of the state space of the Markov process describing the CSMA dynamics.
- We analyze the mixing time of multi-channel CSMA dynamics using the same framework built to study the transition times between dominant states. Our analysis suggests that in the high-activation regime the mixing time does not always correctly characterize the temporal starvation timescale in multi-channel CSMA networks, often leading to pessimistic performance estimates.
- By means of various counterexamples, we show that many desirable properties and performance indices are not monotone in the number  $C$  of available channels, revealing the difficulty of finding analytically the best trade-off between throughput and temporal starvation effects.

The rest of the paper is organized as follows: The CSMA network models for both the single-channel and multiple-channel case are described in Sect. 2. The equivalent representation for multi-channel CSMA networks is presented in Sect. 3. In Sect. 4, we introduce the notion of dominant activity states and derive the properties of the aggregate throughput. We show in Sect. 5 how we characterize temporal starvation by looking at asymptotic results for transition times between dominant states and relate it to structural properties of the underlying state space. Section 6 is entirely devoted to mixing times and to their relation with the worst temporal starvation timescale. Lastly, in Sect. 7 we discuss some generalizations of our model and show how our framework extends to them.

## 2 Model description

In this section, we introduce the stochastic models that describe the dynamics of random-access networks operating according to CSMA-like algorithms. We present first the model for the case of the single-channel CSMA algorithm and later its generalization to the multi-channel scenario. The two models are presented separately for notational convenience, especially since in the next section we will show how the multi-channel model has an equivalent representation as single-channel model.

### 2.1 Single-channel CSMA network

We consider a network of transmitter–receiver pairs sharing a wireless medium according to a CSMA-type algorithm. A *node* indicates potential data transmission between a transmitter and a receiver.

Every node can be either active or inactive, depending on whether the data transmission is ongoing or not. We assume that the network consists of  $N$  such nodes, so that the network activity state can then be described by an  $N$ -dimensional vector  $x$ , where  $x_i = 1$  if node  $i$  is active and  $x_i = 0$  otherwise.

We assume that the network structure and interference conditions can be described by means of an undirected finite graph  $G = (V, E)$ , called the *conflict graph*, where the set of vertices  $V = \{1, \dots, N\}$  represents the nodes of the network and the set of edges  $E \subseteq V \times V$  indicate which pairs of nodes cannot be active simultaneously. Therefore, neighboring nodes in the conflict graph are prevented from simultaneous activity by the carrier-sensing mechanism.

We focus on the scenario where nodes are saturated, which means that nodes always have packets available for transmission; this particularly relevant in high-load regimes. The transmission times of node  $i$  are independent and exponentially distributed with mean  $1/\mu_i$ . When the transmission of a packet is completed, node  $i$  *deactivates* (i.e., releases the medium) and starts a back-off period. The back-off periods of node  $i$  are independent and exponentially distributed with mean  $1/\nu_i$ . At the end of each back-off period, node  $i$  *activates* (i.e., starts the transmission of a new packet) if and only if all its neighboring nodes in  $G$  are currently inactive.

Let  $\mathcal{X} \subseteq \{0, 1\}^N$  be the set of all feasible joint activity states of the network. Since the interference is modeled by the conflict graph  $G$ , the set  $\mathcal{X}$  consists of the incidence vectors of all independent sets of the conflict graph  $G$ :

$$\mathcal{X} := \left\{ x \in \{0, 1\}^N : x_i x_j = 0 \ \forall (i, j) \in E \right\}.$$

If we let  $X(t) \in \mathcal{X}$  denote the network activity state at time  $t$ , then the CSMA dynamics are described by a continuous-time Markov process  $\{X(t)\}_{t \geq 0}$  on the state space  $\mathcal{X}$  with transition rates between  $x, y \in \mathcal{X}$  given by

$$q(x, y) := \begin{cases} \nu_i & \text{if } y = x + e_i \in \mathcal{X}, \\ \mu_i & \text{if } y = x - e_i \in \mathcal{X}, \\ 0 & \text{otherwise,} \end{cases}$$

where  $e_i \in \{0, 1\}^N$  is the vector with all zeros except for a 1 in position  $i$ . The Markov process  $\{X(t)\}_{t \geq 0}$  is reversible [6] and has a product-form stationary distribution

$$\pi(x) := Z^{-1} \prod_{i=1}^N \left( \frac{v_i}{\mu_i} \right)^{x_i}, \quad x \in \mathcal{X}, \tag{1}$$

where  $Z$  is the normalizing constant

$$Z := \sum_{x \in \mathcal{X}} \prod_{i=1}^N \left( \frac{v_i}{\mu_i} \right)^{x_i}.$$

Stationary distribution (1) is insensitive to the distributions of back-off periods and transmission times, in the sense that it depends on these only through their averages  $1/v_i$  and  $1/\mu_i$ , as proved in [26]. Hence, (1) holds in fact for general back-off and transmission time distributions.

### 2.2 Multi-channel CSMA network

In this paper, we consider a generalization of the saturated single-channel CSMA model described in Sect. 2.1 where each node can sense the interference and transmit on any of the  $C$  available channels, but at most on one at a time. We assume that for every  $c = 1, \dots, C$  all possible conflicts between nodes on channel  $c$  are described by a conflict graph  $G_c = (V, E_c)$ . Node  $i$  in the network has a different back-off timer for each of the  $C$  available channels and we model these timers as  $C$  independent Poissonian clocks, ticking at rates  $v_{i,1}, \dots, v_{i,C}$ . When the first of these  $C$  clocks rings for an inactive node, the node senses the corresponding channel, say  $c$ , and starts transmitting on it if and only if the neighboring nodes of  $i$  in  $G_c$  are not active on the same channel. The transmission times of node  $i$  on channel  $c$  are independent and exponentially distributed with mean  $1/\mu_{i,c}$ .

A network activity state is described by a vector  $x \in \{0, 1, \dots, C\}^N$ , where  $x_i = 0$  if node  $i$  is inactive and  $x_i = c$  if node  $i$  is active on channel  $c$ , with  $1 \leq c \leq C$ . Let  $\mathcal{X}_C$  be the collection of feasible network activity states, which consists of all the vectors  $x \in \{0, 1, \dots, C\}^N$  such that, for every  $c = 1, \dots, C$ , two neighboring nodes in  $G_c$  are not simultaneously active on channel  $c$ , i.e.,

$$\mathcal{X}_C := \{x \in \{0, 1, \dots, C\}^N : \forall c \forall (i, j) \in E_c, x_i x_j = 0 \text{ or } x_i \neq x_j\}. \tag{2}$$

If the vector  $X_C(t) \in \mathcal{X}_C$  describes the activity state of the network at time  $t$ , then  $\{X_C(t)\}_{t \geq 0}$  is a Markov process on the state space  $\mathcal{X}_C$  with transition rates between  $x, y \in \mathcal{X}_C$  given by

$$q_C(x, y) = \begin{cases} v_{i,c} & \text{if } y = x + c \cdot e_i \in \mathcal{X}_C, \\ \mu_{i,c} & \text{if } y = x - c \cdot e_i \in \mathcal{X}_C, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Later, we will briefly discuss a generalization of this multi-channel CSMA model in which each node can possibly be simultaneously active on more than one channel and discuss which of our results extend to this setting (see Sect. 7).

### 3 An equivalent representation for multi-channel CSMA networks

In this section, we argue that the evolution of a network using a multi-channel CSMA algorithm can be equivalently described as a *virtual network* operating under the single-channel CSMA algorithm on a modified conflict graph  $G^*$ .

Consider the undirected graph  $G^* = (V^*, E^*)$  with vertex set  $V^* := V \times \{1, \dots, C\}$ , where two nodes  $(i, c)$  and  $(i', c')$  are adjacent if and only if

$$\begin{cases} c \neq c', \\ i = i', \end{cases} \quad \text{or} \quad \begin{cases} c = c', \\ (i, i') \in E_c. \end{cases}$$

To avoid confusion with the original nodes, we refer to the nodes of  $G^*$  as *virtual nodes*. The activity state of the virtual network is then represented by a 0-1 vector of length  $C \cdot N$  and we denote by  $\mathcal{X}^* \subset \{0, 1\}^{C \cdot N}$  the set of admissible virtual activity states on  $G^*$ .

Define the function  $\mathcal{V} : \mathcal{X}^* \rightarrow \mathcal{X}_C$  that maps a virtual network state  $\mathbf{x} \in \mathcal{X}^*$  to the network state  $\mathcal{V}(\mathbf{x}) \in \mathcal{X}_C$  such that, for every  $i = 1, \dots, N$ ,

$$\mathcal{V}(\mathbf{x})_i = \begin{cases} c & \text{if } \exists c \text{ s.t. } \mathbf{x}_{(i,c)} = 1, \\ 0 & \text{otherwise.} \end{cases}$$

The function  $\mathcal{V}$  is well defined, since for every  $c \neq c'$  the virtual nodes  $(i, c)$  and  $(i, c')$  are neighbors in  $G^*$ , and thus only one of them can be active. Furthermore, it is easy to check that  $\mathcal{V}$  is a bijection, from which the following lemma immediately follows.

**Lemma 1** *The collection  $\mathcal{X}^*$  of activity states on the virtual network  $G^*$  is in one-to-one correspondence with the collection  $\mathcal{X}_C$  of multi-channel activity states on  $G$  introduced in (2).*

Suppose further that the  $C \cdot N$  virtual nodes in  $G^*$  operate using the single-channel CSMA algorithm described in Sect. 2.1, assuming that the virtual node  $(i, c)$  has back-off periods that are exponentially distributed with mean  $1/\nu_{i,c}$  and transmission periods that are exponentially distributed with mean  $1/\mu_{i,c}$ . The virtual network activity then evolves as a continuous-time Markov process on  $\mathcal{X}^*$ , which we denote by  $\{\tilde{X}(t)\}_{t \geq 0}$ .

It is easy to check that the transition rate between any two virtual network states  $\mathbf{x}, \mathbf{y} \in \mathcal{X}^*$  is the same as the transition rate given in (3) between the multi-channel activity states  $\mathcal{V}(\mathbf{x}), \mathcal{V}(\mathbf{y}) \in \mathcal{X}_C$ . Hence, the processes  $\{X_C(t)\}_{t \geq 0}$  and  $\{\tilde{X}(t)\}_{t \geq 0}$  are two representations of the same Markovian dynamics, and if  $X_C(0) = \mathcal{V}(\tilde{X}(0))$ , then

$$X_C(t) \stackrel{d}{=} \mathcal{V}(\tilde{X}(t)), \quad \forall t \geq 0.$$

This equivalent representation can be used in combination with (1) to obtain the stationary distribution  $\pi_C$  of the multi-channel activity process  $\{X_C(t)\}_{t \geq 0}$ , as illustrated by the following proposition.

**Proposition 1** *The multi-channel network activity process  $\{X_C(t)\}_{t \geq 0}$  has product-form stationary distribution*

$$\pi_C(x) = Z^{-1} \prod_{i=1}^N \prod_{c=1}^C \left( \frac{\nu_{i,c}}{\mu_{i,c}} \right)^{\mathbb{1}_{\{x_i=c\}}}, \quad x \in \mathcal{X}_C, \tag{4}$$

where  $Z$  is the appropriate normalizing constant.

We remark that multi-channel CSMA dynamics were already shown in [5] to have a product-form distribution and, in the special case where the interference is described by the same conflict graph on every channel  $E_1 = \dots = E_C$ , also in [23]. Our proof approach based on this equivalent representation recovers the same result in the most general case and immediately shows that the insensitivity property carries over to the stationary distribution (4) of the multi-channel CSMA dynamics, which thus holds for general back-off and transmission time distributions.

### 4 Dominant states and aggregate throughput

In this section, we study how the number of channels  $C$  affects the performance in terms of aggregate throughput. While increasing the number of channels evidently provides greater transmission opportunities, the net impact on throughput performance is non-obvious since the transmission capacity per channel is inversely proportional to the number of channels.

Consider a multi-channel CSMA network as described in Sect. 2.2 and assume further that each of these channels has capacity  $1/C$  and that all the nodes have homogeneous activation and transmission rates, namely

$$\nu_{i,c} \equiv \nu \text{ and } \mu_{i,c} \equiv \mu \quad \forall c = 1, \dots, C, \quad \forall i = 1, \dots, N.$$

Without loss of generality, we henceforth also assume that  $\mu = 1$ . The stationary distribution (4) of the activity process  $\{X_C(t)\}_{t \geq 0}$  then reads

$$\pi_C(x) = Z^{-1} \prod_{i=1}^N \nu^{\mathbb{1}_{\{x_i \neq 0\}}} = Z^{-1} \nu^{a(x)}, \quad x \in \mathcal{X}_C, \tag{5}$$

where  $a(x) := \sum_{i=1}^N \mathbb{1}_{\{x_i \neq 0\}}$  is the number  $a(x)$  of active nodes in  $x$ , regardless of the channels they are active on. As (5) shows, when  $\nu > 1$ , the stationary probability of an activity state  $x \in \mathcal{X}_C$  increases with its cardinality in an exponential fashion.

Define  $\mathcal{A}(C)$  to be the maximum number of active nodes in the conflict graph  $G$  when  $C$  channels are available, i.e.,

$$\mathcal{A}(C) := \max_{x \in \mathcal{X}_C} \sum_{i \in V} \mathbb{1}_{\{x_i \neq 0\}} = \max_{x \in \mathcal{X}_C} a(x).$$

In this regime, the stationary distribution (5) favors the activity states with a maximum number of active nodes. We refer to such activity states as *dominant (activity) states* and denote by  $\mathcal{D}_C$  their collection, i.e.,

$$\mathcal{D}_C := \arg \max_{x \in \mathcal{X}_C} a(x) = \{x \in \mathcal{X}_C : a(x) = \mathcal{A}(C)\}.$$

In the rest of the section, we assume further that the same conflict graph  $G = (V, E)$  describes the interference on all  $C$  channels, allowing us to draw a parallel between network activity states and colorings of the graph  $G$ .

In order to characterize the dominant states of a multi-channel CSMA network, we need some further definitions. A *vertex coloring* of the graph  $G$  is a labeling of the graph's vertices with colors such that no two vertices sharing the same edge have the same color. A (*proper*)  $C$ -*coloring* is a vertex coloring using at most  $C$  colors (see, for example, [3,9]). The smallest number of colors needed to color a graph  $G$  is called its *chromatic number* and is denoted by  $\chi(G)$ .

If  $C = 1$ , we are back in the single-channel scenario and the activity states are in one-to-one correspondence with the independent sets of the graph  $G$ . In particular,  $\mathcal{A}(1)$  is equal to the cardinality  $\alpha(G)$  of the maximum independent set on  $G$ , often referred to as the *independence number* of the graph  $G$ .

For  $C \geq 2$ , we showed in Sect. 3 that the multi-channel activity on  $G$  is equivalent to the single-channel activity on the virtual network  $G^*$ . Since we assumed  $E_1 = \dots = E_C$ , the virtual network  $G^*$  is the Cartesian product of  $G$  and the complete graph with  $C$  nodes. In this scenario  $\mathcal{A}(C) = \alpha(G^*)$ , where  $\alpha(G^*)$  is the independence number of the graph  $G^*$ , as proved in [3, Lemma 1].

If the number of available channels in a CSMA network is larger than or equal to the chromatic number of the corresponding conflict graph, i.e.,  $C \geq \chi(G)$ , then every proper  $C$ -coloring of the graph  $G$  corresponds to a dominant state for the network dynamics and, in particular,  $\mathcal{A}(C) = N$ .

If instead  $C < \chi(G)$ , by the definition of the chromatic number, a proper  $C$ -coloring of the graph  $G$  does not exist and, in particular, there are no admissible activity states where all nodes are active. All admissible activity states correspond to *partial C-colorings* of the graph  $G$ . However, the problem of finding the maximum number of nodes that can be active simultaneously in a general conflict graph  $G$  with  $C$  available channels is non-trivial, since it is at least as hard as finding the maximum independent sets of  $G^*$ .

Since nodes are saturated, the *throughput*  $\theta_i(C)$  of node  $i$  is proportional to the long-run fraction of time that node  $i$  is active and can be expressed as

$$\theta_i(C) := \frac{1}{C} \sum_{x \in \mathcal{X}_C} \pi_C(x) \mathbb{1}_{\{x_i \neq 0\}}, \tag{6}$$

where the constant  $1/C$  accounts for the fact that each channel has capacity  $1/C$ . Being a function of the stationary distribution, the throughput is also insensitive to the distributions of the back-off and transmission times.

We analyze the *aggregate throughput*  $\Theta(C)$  of a multi-channel CSMA network with  $C$  available channels in the asymptotic regime where the activation rate  $\nu$  grows large, defined as



$$\Theta(C) := \lim_{\nu \rightarrow \infty} \sum_{i \in V} \theta_i(C) = \lim_{\nu \rightarrow \infty} \sum_{i \in V} \frac{1}{C} \sum_{x \in \mathcal{X}_C} \pi_C(x) \mathbb{1}_{\{x_i \neq 0\}}.$$

Even if it may be infeasible to calculate  $\mathcal{A}(C)$  for a given random-access network using  $C$  channels, some conclusions for its aggregate throughput  $\Theta(C)$  can still be drawn, as illustrated by the following theorem.

**Theorem 1** (Aggregate throughput) *The following statements hold for the aggregate throughput  $\Theta(C)$  of a multi-channel CSMA network where the same conflict graph  $G = (V, E)$  describes the interference on all  $C$  channels:*

- (i)  $\Theta(C) = \mathcal{A}(C)/C$  for every  $C \geq 1$ ;
- (ii)  $\Theta(C) = \alpha(G)$  for every  $1 \leq C \leq C^*(G)$ , where  $C^*(G)$  is the number of disjoint maximum independent sets of the graph  $G$ ;
- (iii)  $\Theta(C) = N/C$  for every  $C \geq \chi(G)$ ;
- (iv)  $\Theta(C)$  is a non-increasing function of the number of channels  $C$ .

**Proof** The definition of a dominant state and (5) together yield that  $\sum_{x \in \mathcal{D}_C} \pi_C(x) \rightarrow 1$  as  $\nu \rightarrow \infty$ , and thus,

$$\begin{aligned} \Theta(C) &= \lim_{\nu \rightarrow \infty} \sum_{i \in V} \frac{1}{C} \sum_{x \in \mathcal{X}_C} \pi_C(x) \mathbb{1}_{\{x_i \neq 0\}} \\ &= \frac{1}{C} \lim_{\nu \rightarrow \infty} \sum_{x \in \mathcal{X}_C} \pi_C(x) \sum_{i \in V} \mathbb{1}_{\{x_i \neq 0\}} \\ &= \frac{\mathcal{A}(C)}{C} \lim_{\nu \rightarrow \infty} \sum_{x \in \mathcal{D}_C} \pi_C(x) \\ &= \frac{\mathcal{A}(C)}{C}. \end{aligned}$$

As mentioned earlier, when  $C \geq \chi(G)$  there exists a proper  $C$ -coloring of the graph  $G$ , which means all nodes can be simultaneously active and thus  $\mathcal{A}(C) = N$ . Furthermore, it is easy to prove by induction that  $\mathcal{A}(C) = C \cdot \alpha(G)$  for every  $C \leq C^*(G)$ . To prove (iv), we will show that the following inequality holds:

$$\frac{\mathcal{A}(C + 1)}{C + 1} \leq \frac{\mathcal{A}(C)}{C}, \quad \forall C \geq 1.$$

The inequality trivially holds when  $C \geq \chi(G)$ , since  $\mathcal{A}(C) = N$ . The proof in the case  $C \leq \chi(G)$  readily follows by considering a  $(C + 1)$ -coloring that yields  $\mathcal{A}(C + 1)$ ; pick the color that is used the least and discolor the corresponding nodes, which are at most  $\lfloor \frac{\mathcal{A}(C+1)}{C+1} \rfloor$ , obtaining a  $C$ -coloring with  $\mathcal{A}(C + 1) - \lfloor \frac{\mathcal{A}(C+1)}{C+1} \rfloor \geq \frac{C}{C+1} \mathcal{A}(C + 1)$  colored nodes. Using this newly obtained  $C$ -coloring, we immediately get that  $\mathcal{A}(C) \geq \frac{C}{C+1} \mathcal{A}(C + 1)$ .  $\square$

In view of the fact that  $\Theta(C)$  is a non-increasing function of the number of channels  $C$ , it seems that the network should be operated using a single channel, but this is also




Fig. 1 Example of a network and its dominant states with  $C = 1$  and  $C = 2$  channels available, respectively

the scenario where temporal starvation phenomena and spatial unfairness are often more pronounced.

Jain’s fairness index [13] is often used as a fairness measure in the context of random-access algorithms, and we consider here its asymptotic counterpart  $J(C)$  in the high-activation limit  $\nu \rightarrow \infty$ , i.e.,

$$J(C) := \lim_{\nu \rightarrow \infty} \frac{\left(\sum_{i=1}^N \theta_i(C)\right)^2}{N \sum_{i=1}^N \theta_i(C)^2}.$$

One may conjecture that a larger number of available channels should always reduce the spatial unfairness in the network. This is indeed true for many conflict graphs, but does not hold in general. Figure 1 shows a particular example of a network in which by adding an additional channel, the spatial unfairness worsens as  $J(2) < J(1)$ .

As far as temporal starvation is concerned, we expect that by exploiting more channels the temporal starvation effects can be mitigated. However, making this dependence more explicit is a challenging task, since the structure of the conflict graph plays a crucial role via the parametric family  $\mathcal{D}_C$  of dominant activity states. In general, a non-trivial trade-off emerges between the magnitude of temporal starvation effects and the aggregate throughput that the network can achieve, and a balance can be found using a number of channels in the range between 1 and  $\chi(G)$ .

### 5 Temporal starvation evaluation via transition times

In this section, we introduce a framework to study and quantify temporal starvation effects in multi-channel CSMA networks in the high-activation regime  $\nu \rightarrow \infty$ .

When the activation rate grows large, the network spends roughly the same fraction of time in each dominant state. However, it takes a long time for the activity process to move from one dominant state to the other, since such a transition involves the occurrence of rare events. Indeed, intuitively, the activity process must follow a transition path through some activity states with fewer active nodes which are highly unlikely in view of (5) and the time to reach such activity states is correspondingly long.

We study the transitions between dominant states in terms of first hitting times of the activity process  $\{X_C(t)\}_{t \geq 0}$ . In the limit  $\nu \rightarrow \infty$ , the asymptotic behavior of the transition time between any pair of dominant states can be characterized exactly on a logarithmic scale by studying structural properties of the state space  $\mathcal{X}_C$ .

The main idea of our method is inspired by the concept of “traps” of the state space introduced in the context of single-channel CSMA networks in [17,18]. Traps are essentially subsets of activity states where the activity process remains for a long time before leaving. By focusing on the asymptotic regime  $\nu \rightarrow \infty$ , we can restrict our analysis only to the traps corresponding to dominant states and, moreover, characterize the timescale of the transitions between them without the need of a detailed description of the structure of these traps.

In this section, we consider the multi-channel CSMA dynamics with homogeneous rates as in Sect. 4, but here we do *not* need to assume that the interference structure on different channels is described by the same conflict graph  $G$ . Note that the assumption of homogeneous rates across different nodes and channels is not crucial, and, in fact, it can be relaxed, but to keep the exposition simple we do not present here the general case and the interested reader is referred to Sect. 7, where we illustrate also other generalizations of our model for which this framework is still valid.

Let  $\tau_{x,A}(\nu) := \inf\{t > 0 : X_C(t) \in A\}$  be the *first hitting time* of the subset  $A \subset \mathcal{X}_C$  for the Markov process  $\{X_C(t)\}_{t \geq 0}$  starting in  $x$  at  $t = 0$ . For any  $x, y \in \mathcal{X}_C$ , we say that  $\omega \subset \mathcal{X}_C$  is a *path* from  $x$  to  $y$  in  $\mathcal{X}_C$  and denote it by  $\omega : x \rightarrow y$ , if  $\omega$  is a finite sequence of states  $\omega_1, \dots, \omega_n \in \mathcal{X}_C$  such that  $\omega_1 = x, \omega_n = y$  and the CSMA dynamics allow the step from  $\omega_i$  to  $\omega_{i+1}$  for every  $i = 1, \dots, n - 1$ . The *communication height* between two states  $x, y \in \mathcal{X}_C$  is the minimum number of nodes (with respect to that of any dominant state) that need to become simultaneously inactive at some point along any path from  $x$  to  $y$ , i.e.,

$$\Delta(x, y) := \min_{\omega: x \rightarrow y} \max_{z \in \omega} (\mathcal{A}(C) - a(z)). \tag{7}$$

Both the notions of a path and communication height naturally extend to the case of two subsets  $A, B \subset \mathcal{X}_C$ .

The communication height is a crucial quantity in our analysis since in the regime  $\nu \rightarrow \infty$  the most likely trajectories from  $x$  to  $y$  are precisely those that achieve the min–max in (7). Indeed, all the other trajectories  $\omega : x \rightarrow y$  will visit at some point an activity state with fewer active nodes which is thus correspondingly more rare in view of (5). The communication height  $\Delta(\cdot, \cdot)$  defines an *ultrametric* on  $\mathcal{X}_C$ , since

1.  $\Delta(x, y) \geq 0$  and equality holds if and only if  $x = y$ ;
2.  $\Delta(x, y) = \Delta(y, x)$ ;
3.  $\Delta(x, y) \leq \max\{\Delta(x, z), \Delta(z, y)\}$ .

The second statement in property (1) holds since if  $x \neq y$  at every step of the CSMA dynamics, the number of active nodes must either increase or decrease by 1 and cannot stay constant along any path from  $x$  to  $y$  and, hence,  $\Delta(x, y) > 0$ .

For any  $a \in \mathbb{R}^+$ , we denote by  $\log_\nu a$  the logarithm in base  $\nu$ . Since most of the results are in the limit  $\nu \rightarrow \infty$ , we will henceforth assume that  $\nu > 1$ , so that  $\log_\nu(\cdot)$  is a strictly increasing function from  $\mathbb{R}^+$  to  $\mathbb{R}$ . The next theorem shows that the communication height  $\Delta(s, D)$  between a dominant state  $s \in \mathcal{D}_C$  and a subset  $D \subseteq \mathcal{D}_C \setminus \{s\}$  completely characterizes on a logarithmic scale the expected transition time from  $s$  to  $D$  in the limit  $\nu \rightarrow \infty$ .

**Theorem 2** (Timescale of transitions between dominant states) *For every dominant state  $s \in \mathcal{D}_C$  and target subset  $D \subseteq \mathcal{D}_C \setminus \{s\}$  of dominant states different from  $s$ , the following asymptotic equality holds:*

$$\lim_{\nu \rightarrow \infty} \log_\nu \mathbb{E}\tau_{s,D}(\nu) = \Delta(s, D) - 1.$$

The proof of this result is presented in Sect. 5.1. We now illustrate how Theorem 2 allows us to identify the longest timescale for which each node of the network starves. For each node  $i \in V$ , let  $\mathcal{S}_C(i) := \{s \in \mathcal{D}_C : s_i \neq 0\}$  be the subset of dominant states in which node  $i$  is active. It is natural to study the temporal starvation only for the nodes which are active in at least one dominant state, i.e.,  $\mathcal{S}_C(i) \neq \emptyset$ , and that are not active in every dominant state, i.e.,  $\mathcal{S}_C(i) \neq \mathcal{D}_C$ . For every such node  $i$ , we define its *starvation index* as

$$\Upsilon_i(C) := \max_{s \in \mathcal{D}_C \setminus \mathcal{S}_C(i)} \min_{s' \in \mathcal{S}_C(i)} \Delta(s, s'). \tag{8}$$

From property (1) of the communication heights, it immediately follows that  $\Upsilon_i(C) \geq 1$ . We choose not to define the starvation index of the nodes  $i \in V$  such that  $\mathcal{S}_C(i) = \emptyset$ , since these nodes do not properly suffer from *temporal* starvation, since they permanently starve in the limit  $\nu \rightarrow \infty$ , having asymptotically zero throughput. Similarly,  $\Upsilon_i(C)$  is not defined for any node  $i$  that is active in every dominant state of the network, i.e.,  $\mathcal{S}_C(i) = \mathcal{D}_C$ , since any such node  $i$  does not starve in view of the fact that  $\sum_{x \in \mathcal{X}_C} \pi_C(x) \mathbb{1}_{\{x_i \neq 0\}} \rightarrow 1$  as  $\nu \rightarrow \infty$ .


The starvation index  $\Upsilon_i(C)$  captures the largest timescale at which node  $i$  suffers from temporal starvation at stationarity in the high-activation regime  $\nu \rightarrow \infty$ . Indeed,  $\Upsilon_i(C)$  characterizes on a logarithmic scale how long it takes the activity process to reach the subset  $\mathcal{S}_C(i)$  of dominant states where node  $i$  is active when starting in the worst possible dominant state in  $\mathcal{D}_C \setminus \mathcal{S}_C(i)$ , as illustrated by the following identity:

$$\lim_{\nu \rightarrow \infty} \log_\nu \left( \max_{s \in \mathcal{D}_C \setminus \mathcal{S}_C(i)} \mathbb{E}\tau_{s,\mathcal{S}_C(i)}(\nu) \right) = \Upsilon_i(C) - 1.$$

Lastly, we define the *starvation index of the network* as the worst starvation index of its nodes, i.e.,

$$\Upsilon(C) := \max_{i \in V: \mathcal{S}_C(i) \neq \emptyset, \mathcal{D}_C} \Upsilon_i(C). \tag{9}$$

The index  $\Upsilon(C)$  is a non-increasing function of the number of available channels for all networks of which we analyzed numerically the corresponding state spaces



**Fig. 2** Example of a network with the unique dominant state  $s^{(1)} \in \mathcal{D}_1$  (on the left) and two dominant states  $s^{(2)}, s^{(3)} \in \mathcal{D}_2$  (in the middle and on the right). In particular, the dominant state  $s^{(2)}$  is not a superset of the dominant state  $s^{(1)}$

$\mathcal{X}_1, \dots, \mathcal{X}_{\chi(G)-1}$ . However, proving this claim analytically is challenging in view of the intricate nature of the dominant states. Indeed, dominant states in  $\mathcal{D}_{C+1}$  are not necessarily obtained from dominant states in  $\mathcal{D}_C$  by activating a maximum number of inactive nodes on the  $(C + 1)$ -th channel, as illustrated for the network in Fig. 2. Hence, there is no simple relation between the sets  $\mathcal{D}_C$  and  $\mathcal{D}_{C+1}$  nor any other “monotonicity property” that one could leverage. In particular, this means that the maximum in (9) may have to possibly be taken over different and/or non-nested subsets of nodes when increasing  $C$ .

**5.1 Asymptotics for transition times between dominant states**

This subsection is devoted to the proof of Theorem 2. We first introduce the uniformized discrete-time version of the activity process, which will allow us to use classical results for expected hitting times of reversible Markov chains and their analogy with electrical networks (see, for example, [10, Section 2.1] and references therein).

In this subsection, we will almost entirely suppress the subscript  $C$  for notational compactness, but the reader should keep in mind that the results hold for general multi-channel CSMA dynamics as described in Sect. 2.2. We consider the uniformized discrete-time version of the activity process  $\{X_C(t)\}_{t \geq 0}$ . In more detail, we construct a discrete-time Markov chain  $\{\tilde{X}(t)\}_{t \in \mathbb{N}}$  starting from the continuous-time Markov process  $\{X_C(t)\}_{t \geq 0}$  by means of uniformization at rate  $q_{\max} := \max_{x \in \mathcal{X}_C} \sum_{y \neq x} q(x, y) = CN\nu$ . The transition probabilities of this Markov chain read

$$P(x, y) = \begin{cases} \frac{1}{CN} \nu^{-[a(x)-a(y)]^+} & \text{if } d(x, y) = 1, \\ 0 & \text{if } d(x, y) > 1, \\ 1 - \sum_{z \in \mathcal{X}, z \neq x} P(x, z) & \text{if } x = y, \end{cases}$$

where  $d(x, y)$  is the distance function on  $\mathcal{X} \times \mathcal{X}$  defined as

$$d(x, y) := |\{i \in V : x_i \neq y_i\}| + |\{i \in V : x_i \neq y_i \text{ and } x_i y_i \neq 0\}|.$$

The quantity  $d(x, y)$  represents the minimum number of steps that the CSMA dynamics need to go from state  $x$  to state  $y$ ; the second term on the right-hand side accounts for the fact that if a node is active on a channel  $c$ , the immediate activation on another frequency  $c' \neq c$  is not allowed and requires first the node to become inactive.

This distance function on  $\mathcal{X}$  reflects the main difference between the multi-channel CSMA dynamics and the classical multi-color Glauber dynamics, in which a node can change channel/color in a single step (see [19, Section IV.C] for more details). The two corresponding state spaces, while still having the same states, have thus fundamentally different structures and, in particular, there are many more possible trajectories between dominant states in the one corresponding to the Glauber dynamics. This fact means that using Glauber dynamics to study temporal starvation phenomena may lead to overly optimistic performance bounds.

It is easy to check that (5) is the stationary distribution also for the uniformized chain  $\{\tilde{X}(t)\}_{t \in \mathbb{N}}$ . Let  $T_{x,A}(\nu) := \inf\{t \in \mathbb{N} : \tilde{X}(t) \in A\}$  be the discrete-time counterpart of the first hitting time  $\tau_{x,A}(\nu)$ . These hitting times are closely related as

$$\tau_{x,A}(\nu) \stackrel{d}{=} \sum_{i=1}^{T_{x,A}(\nu)} \mathcal{E}_i(\nu),$$

where  $\{\mathcal{E}_i(\nu)\}_{i \in \mathbb{N}}$  are i.i.d. exponential random variables with mean  $q_{\max}^{-1} = (CN\nu)^{-1}$ . In particular, it holds that

$$\mathbb{E}\tau_{x,A}(\nu) = (CN\nu)^{-1} \mathbb{E}T_{x,A}(\nu). \tag{10}$$

This equality shows that we can study the expected transition times between dominant states in the discrete-time setting where the theory is richer and then immediately translate them back for the continuous-time activity process  $\{X_C(t)\}_{t \geq 0}$ . We are now ready to present the proof of the main result of this section.

**Proof of Theorem 2** We first introduce the following notation to asymptotically compare two functions  $f, g : \mathbb{R} \rightarrow \mathbb{R}$ : We write  $f(\nu) \prec g(\nu)$  if  $f(\nu) = o(g(\nu))$ ,  $f(\nu) \leq g(\nu)$  if  $f(\nu) = O(g(\nu))$  as  $\nu \rightarrow \infty$  and  $f(\nu) \asymp g(\nu)$  if both  $f(\nu) \leq g(\nu)$  and  $g(\nu) \leq f(\nu)$ .

Furthermore, we will need some classical notions from the theory of reversible Markov chains on finite state spaces. For every pair of states  $x, y \in \mathcal{X}$  for which  $P(x, y) \neq 0$ , we define the *resistance* of the edge  $e = (x, y)$  as  $r(e) = r(x, y) = (\pi(x) P(x, y))^{-1}$ . For any state  $x \in \mathcal{X}$  and subset  $A \subset \mathcal{X} \setminus \{x\}$ , define then the *effective resistance*  $R(x \leftrightarrow A) := \pi(x)^{-1} \mathbb{P}(T_{x,A} < T_{x,x}^+)$  and the *critical resistance*  $\Psi(x, A) := \min_{\omega: x \rightarrow A} \max_{e \in \omega} r(e)$ . As shown in [10, Proposition 2.2], effective resistance and critical resistance of the same pair of states are intimately related, namely there exists a constant  $k \geq 1$  independent of  $\nu$  such that

$$\frac{1}{k} \Psi(x, A) \leq R(x \leftrightarrow A) \leq k \Psi(x, A). \tag{11}$$

As illustrated in [10, Section 2.1], the following identity holds for the expected hitting time of the target subset  $A$  starting from a state  $x \notin A$

$$\mathbb{E}T_{x,A} = \pi(x) R(x \leftrightarrow A) \sum_{z \in \mathcal{X}} \frac{\pi(z)}{\pi(x)} \mathbb{P}(T_{z,x} < T_{z,A}). \tag{12}$$

Consider now the special case relevant for this theorem in which the starting state is a dominant state, i.e.,  $x = s \in \mathcal{D}_C$ , and the target subset  $D$  is such that  $D \subseteq \mathcal{D}_C \setminus \{x\}$ . From (5), it immediately follows that

$$\pi(s) > \pi(z) \quad \forall z \in \mathcal{X} \setminus \mathcal{D}_C \quad \text{and} \quad \pi(s) \geq \pi(z) \quad \forall z \in \mathcal{D}_C.$$

Using this fact and identity (12), we deduce that  $\mathbb{E}T_{s,D} \asymp \pi(s) R(s \leftrightarrow D)$  as  $\nu \rightarrow \infty$ . In view of (11), it then follows that

$$\mathbb{E}T_{s,D} \asymp \pi(s) \Psi(s, D), \quad \text{as } \nu \rightarrow \infty,$$

and thus

$$\lim_{\nu \rightarrow \infty} \log_\nu \mathbb{E}T_{s,D} = \lim_{\nu \rightarrow \infty} \log_\nu \pi(s) \Psi(s, D). \tag{13}$$

Consider the right-hand side of the latter identity. Using the definition of critical resistance, it follows that

$$\begin{aligned} \log_\nu \pi(s) \Psi(s, D) &= \log_\nu \min_{\omega: s \rightarrow D} \max_{e \in \omega} \pi(s) r(e) \\ &= \log_\nu \min_{\omega: s \rightarrow D} \max_{(x,y) \in \omega} \frac{\pi(s)}{\pi(x) P(x, y)} \\ &= \min_{\omega: s \rightarrow D} \max_{(x,y) \in \omega} \log_\nu \frac{CN_\nu^{\mathcal{A}(C)}}{\nu^{a(x)\nu - [a(x) - a(y)]^+}} \\ &= \log_\nu CN + \min_{\omega: s \rightarrow D} \max_{(x,y) \in \omega} \mathcal{A}(C) - a(x) + [a(x) - a(y)]^+ \\ &= \log_\nu CN + \min_{\omega: s \rightarrow D} \max_{(x,y) \in \omega} \mathcal{A}(C) - (a(x) \wedge a(y)) \\ &= \log_\nu CN + \min_{\omega: s \rightarrow D} \max_{x \in \omega} \mathcal{A}(C) - a(x) \\ &= \log_\nu CN + \Delta(s, D), \end{aligned} \tag{14}$$

where the third last passage above follows from the identity  $-a(x) + [a(x) - a(y)]^+ = -(a(x) \wedge a(y))$ , while the second last one is a consequence of the fact that we take the maximum over *all the configurations visited by the path  $\omega$* , rather than the edges crossed by the same.

Taking the limit in (14), we obtain  $\lim_{\nu \rightarrow \infty} \log_\nu \pi(s) \Psi(s, D) = \Delta(s, D)$ , which, combined with identity (13), yields

$$\lim_{\nu \rightarrow \infty} \log_\nu \mathbb{E}T_{s,D}(\nu) = \Delta(s, D).$$

In view of (10), we have

$$\begin{aligned} \lim_{\nu \rightarrow \infty} \log_{\nu} \mathbb{E} \tau_{s,D}(\nu) &= \lim_{\nu \rightarrow \infty} \log_{\nu} (CN\nu)^{-1} \mathbb{E} T_{s,D}(\nu) \\ &= \lim_{\nu \rightarrow \infty} \log_{\nu} \mathbb{E} T_{s,D}(\nu) + \lim_{\nu \rightarrow \infty} \log_{\nu} (CN\nu)^{-1} \\ &= \Delta(s, D) - 1, \end{aligned}$$

and this concludes the proof. □

### 6 Mixing time asymptotics for multi-channel CSMA networks

In this section, we turn attention to the long-run behavior of the Markov process  $\{X_C(t)\}_{t \geq 0}$  and in particular examine the rate of convergence to the stationary distribution. We measure the rate of convergence in terms of the total variation distance and the so-called *mixing time*, which describes the time required for the distance to stationarity to become small.

The mixing time becomes particularly relevant when the network has two or more dominant states which together attract the entire probability mass in the limit as  $\nu \rightarrow \infty$ . Indeed, in this case, the mixing time provides an indication of how long it takes the activity process to reach a certain level of fairness among the dominant components.

For any fixed activation rate  $\nu$ , the maximal distance over  $x \in \mathcal{X}_C$ , measured in terms of total variation, between the distribution at time  $t$  of the activity process and the stationary distribution  $\pi_C$  is defined as

$$d_t(\nu) := \max_{x \in \mathcal{X}_C} \|P_{\nu}^t(x, \cdot) - \pi_C\|_{TV},$$

where  $P_{\nu}^t(x, \cdot)$  is the distribution at time  $t$  of the Markov process  $\{X_C(t)\}_{t \geq 0}$  started at time 0 in  $x$ . The *mixing time* of the process  $\{X_C(t)\}_{t \geq 0}$  is then defined as

$$t_{\text{mix}}(\varepsilon, \nu) := \inf\{t \geq 0 : d_t(\nu) \leq \varepsilon\}.$$

The next result shows that the worst communication height between any two dominant states provides a lower bound for the timescale at which the activity process mixes.

**Theorem 3** (Mixing time timescale) *The mixing time of the multi-channel activity process  $\{X_C(t)\}_{t \geq 0}$  satisfies*


$$\lim_{\nu \rightarrow \infty} \log_{\nu} t_{\text{mix}}(\varepsilon, \nu) \geq \Gamma(C) - 1, \tag{15}$$

where  $\Gamma(C) \geq 1$  is the worst communication height between a pair of dominant states, i.e.,  $\Gamma(C) := \max_{s,s' \in \mathcal{D}_C} \Delta(s, s')$ .


The proof of this theorem is based on a conductance argument that leverages the bottlenecks of the state space  $\mathcal{X}_C$  and is presented later in Sect. 6.1.



**Fig. 3** The four dominant states  $\mathcal{D}_1 = \{s^{(1)}, s^{(2)}, s^{(3)}, s^{(4)}\}$  on a network with  $C = 1$  channels available



**Fig. 4** The four dominant states  $\mathcal{D}_2 = \{s^{(5)}, s^{(6)}, s^{(7)}, s^{(8)}\}$  on a network with  $C = 2$  channels available



**Table 1** The values of the communication height  $\Delta(\cdot, \cdot)$  between all pairs of dominant states in  $\mathcal{D}_1$  (on the left, cf. Fig. 3) and in  $\mathcal{D}_2$  (on the right, cf. Fig. 4)

	$s^{(1)}$	$s^{(2)}$	$s^{(3)}$	$s^{(4)}$	$s^{(5)}$	$s^{(6)}$	$s^{(7)}$	$s^{(8)}$	
$s^{(1)}$	0	2	2	2	$s^{(5)}$	0	1	3	3
$s^{(2)}$	2	0	1	1	$s^{(6)}$	1	0	3	3
$s^{(3)}$	2	1	0	1	$s^{(7)}$	3	3	0	1
$s^{(4)}$	2	1	1	0	$s^{(8)}$	3	3	1	0

The first consideration is that for some networks the mixing time worsens when the number of available channels increases. Consider the network depicted in Figs. 3 and 4, which illustrate its dominant states for  $C = 1$  and  $C = 2$ , respectively. For such a network, it is easy to check that (15) holds with equality when  $C = 1$  and to compute that  $\Gamma(1) = 2$ , while  $\Gamma(2) = 3$  (see Table 1).

We now compare this new exponent  $\Gamma(C)$  characterizing the mixing time in the regime  $\nu \rightarrow \infty$  with the starvation index  $\Upsilon(C)$  introduced in Sect. 5. The inequality  $\Upsilon_i(C) \leq \Gamma(C)$  holds for every node  $i \in V$  with  $\mathcal{S}_C(i) \neq \emptyset$  and  $\mathcal{S}_C(i) \neq \mathcal{D}_C$ , since

$$\gamma_i(C) = \max_{s \in \mathcal{D}_C \setminus \mathcal{S}_C(i)} \min_{s' \in \mathcal{S}_C(i)} \Delta(s, s') \leq \max_{s, s' \in \mathcal{D}_C} \Delta(s, s') = \Gamma(C),$$

and trivially holds also for all the other nodes for which we set  $\gamma_i(C) = 0$ , as  $\Gamma(C) \geq 1$ . Therefore, the starvation index of the network  $\Upsilon(C)$  is also not larger than  $\Gamma(C)$ , i.e.,

$$\Upsilon(C) \leq \Gamma(C). \tag{16}$$

These two indices can be equal, for instance in the case of a conflict graph with a grid structure and an even number of nodes. As illustrated in [30], the CSMA dynamics with  $C = 1$  on this bipartite conflict graph have exactly two dominant states and each node is active in only one of the two, and therefore  $\Upsilon(1) = \Gamma(1)$ .

However, for most conflict graphs and choices of the number  $C$  of available channels, inequality (16) is strict. The network presented in Fig. 4 is an example for which  $\Upsilon(2) < \Gamma(2)$ . We already calculated earlier that  $\Gamma(2) = 3$ . As far as the starvation index is concerned, the only two nodes  $v, w$  for which it is well defined are the two topmost ones of each dominant state in Fig. 4, and again using Table 1, we can deduce that  $\Upsilon_v(2) = \Upsilon_w(2) = 1$ , and, thus,  $\Upsilon(2) = 1$ .

The fact that inequality (16) is often strict suggests that the mixing time may not be the best metric to study temporal starvation, at least in the high-activation regime  $\nu \rightarrow \infty$ , since it may yield a pessimistic performance assessment by overestimating the duration of the starvation effects. Indeed, while it may take a long time for the activity process  $\{X_C(t)\}_{t \geq 0}$  to visit all the dominant states and thus “mix” on the state space  $\mathcal{X}_C$ , the temporal starvation of the nodes often occurs at shorter timescales even in the worst case, that is, the one captured by the network starvation index  $\Upsilon(C)$ .

We notice that determining precisely for which class of conflict graphs the strict inequality  $\Upsilon(C) < \Gamma(C)$  holds is probably a very difficult combinatorial problem, whose solution is beyond the scope of this paper.

### 6.1 Asymptotic lower bound for the mixing time

This subsection is entirely devoted to the proof of Theorem 3. The proof relies on a classical conductance argument, which is often used in the literature to obtain lower bounds for the mixing time of Markov chains with a finite state space (see [21, Theorem 7.3]).

For any subset  $S \subset \mathcal{X}_C$ , the *conductance* of  $S$  is defined as

$$\Phi(S) := \frac{Q(S, S^c)}{\pi_C(S)},$$

where  $Q(S, S^c)$  is *probability flow out of  $S$* , i.e.,

$$Q(S, S^c) := \sum_{x \in S, y \in S^c} \pi_C(x) q(x, y).$$

In order to bound the mixing time of the activity process  $\{X_C(t)\}_{t \geq 0}$ , we will make use of the continuous-time counterpart of that theorem, which is summarized in the following lemma (see [29, Lemma 7.1] for a proof).

**Lemma 2** For any  $\varepsilon \in (0, 1/2)$ ,

$$t_{\text{mix}}(\varepsilon, \nu) \geq \frac{1/2 - \varepsilon}{\min_{S \subset \mathcal{X}_C : \pi_C(S) \leq 1/2} \Phi(S)}. \tag{17}$$

To get a sharp bound for the conductance and hence a sharp lower bound for the mixing time, we need to identify a subset  $S$  with low conductance. Let  $s, s' \in \mathcal{D}_C$  be two dominant states such that  $\Delta(s, s') = \Gamma(C)$  and consider the subset  $S$  of network states that can be reached by the CSMA dynamics starting in  $s$  while keeping no less than  $\mathcal{A}(C) - \Gamma(C) + 1$  nodes active, i.e.,

$$S = \{x \in \mathcal{X}_C : \Delta(s, x) < \Gamma(C)\}.$$

Note that  $s \in S$  and  $s' \in S^c$ , which implies that both subsets  $S$  and  $S^c$  have nonvanishing stationary probabilities containing at least one dominant set each. Without loss of generality, let us assume that  $\pi_C(S) \leq 1/2$  for all  $\nu$  sufficiently large.

By the construction of  $S$ , all states  $x \in S$  such that there exists  $y \in S^c$  such that  $d(x, y) = 1$  must have exactly  $\mathcal{A}(C) - \Gamma(C) + 1$  active nodes. Indeed, they clearly cannot have fewer active nodes, otherwise  $\Delta(s, x) \geq \Gamma(C)$ , and if they had strictly more than  $\mathcal{A}(C) - \Gamma(C) + 1$ , it would be impossible to exit in a single step from  $S$ , since the CSMA dynamics prescribe the number of active nodes to increase or decrease at most by one. In view of this fact, the probability flow of  $S$  can be rewritten as

$$Q(S, S^c) = \sum_{x \in S, y \in S^c} \pi_C(x) q(x, y) = Z^{-1} \nu^{\mathcal{A}(C) - \Gamma(C) + 1} B(S),$$

where the quantity  $B(S) := |\{(x, y) \in S \times S^c : d(x, y) = 1\}|$  counts the number of possible way of exiting from  $S$ , which is independent of  $\nu$  and depends only on the structure of the state space  $\mathcal{X}_C$ . Then,

$$\begin{aligned} \lim_{\nu \rightarrow \infty} \log_\nu Q(S, S^c) &= \lim_{\nu \rightarrow \infty} \log_\nu (Z^{-1} \nu^{\mathcal{A}(C) - \Gamma(C) + 1} B(S)) \\ &= \lim_{\nu \rightarrow \infty} \left( \log_\nu (Z^{-1} \nu^{\mathcal{A}(C) - \Gamma(C) + 1}) + \log_\nu B(S) \right) \\ &= \lim_{\nu \rightarrow \infty} \log_\nu (Z^{-1} \nu^{\mathcal{A}(C) - \Gamma(C) + 1}) \\ &= \mathcal{A}(C) - \Gamma(C) + 1 - \mathcal{A}(C) = 1 - \Gamma(C). \end{aligned}$$

Using the conductance of the subset  $S$  in (17) yields the lower bound  $t_{\text{mix}}(\varepsilon, \nu) \geq (1/2 - \varepsilon) / \Phi(S)$ . Therefore,

$$\begin{aligned}
 \lim_{\nu \rightarrow \infty} \log_{\nu} t_{\text{mix}}(\varepsilon, \nu) &\geq \lim_{\nu \rightarrow \infty} \log_{\nu} \frac{1/2 - \varepsilon}{\Phi(S)} \\
 &= - \lim_{\nu \rightarrow \infty} \log_{\nu} \Phi(S) \\
 &= \lim_{\nu \rightarrow \infty} \log_{\nu} \pi(S) - \log_{\nu} Q(S, S^c) \\
 &= \Gamma(C) - 1,
 \end{aligned}$$

and this concludes the proof of Theorem 3.

### 7 Model generalizations

In this section, we discuss some generalizations of our model and show to what extent our results remain valid.

Our framework for transition and mixing times does allow for more general heterogeneous activation rates of the form  $\nu_{i,c} = f_{i,c}(\nu)$ , where  $f_{i,c}(\cdot)$  are nonnegative real functions with a common scaling parameter  $\nu$ . These more general rates can be used to model interesting scenarios where nodes have preferred channels or frequency channels with different features. Allowing for this more general parametric family of activation rates does not affect the state space  $\mathcal{X}_C$ , but does change the stationary distribution  $\pi_C$  and, ultimately, the dominant states of the process. The asymptotic results derived in Sects. 5 and 6 in the regime  $\nu \rightarrow \infty$  still hold for this more general setting. However, it is not enough to keep track only of the number of active nodes and one needs to work with the more involved functions  $\tilde{a}(x) := \sum_{i=1}^N \sum_{c=1}^C w_{i,c} \mathbb{1}_{\{x_i=c\}}$  and  $\tilde{\mathcal{A}}(C) := \max_{x \in \mathcal{X}_C} \tilde{a}(x)$ , where the weights  $w_{i,c}$  are defined as  $w_{i,c} := \lim_{\nu \rightarrow \infty} \log_{\nu} f_{i,c}(\nu)$  (assuming that such limits exist and are finite).

Another natural generalization of our model is the one where each node  $i \in V$  of the network possibly has more than one radio interface and thus can simultaneously transmit on  $c_i$  channels, with  $1 \leq c_i \leq C$ , having an independent back-off timer for each one of them. In this setting, there is no natural way to represent the multi-channel dynamics as single-channel dynamics on a virtual network as we did in Sect. 3, but the resulting Markov process still has a product-form stationary distribution, as proved in [23]. Despite the fact that this assumption results in a much larger and fundamentally different state space, the asymptotic results for transition and mixing times proved in Sects. 5 and 6 remain valid, after having opportunely modified the definition of the functions  $a(x)$  and  $\mathcal{A}(C)$ . However, the identification of the dominant states becomes more involved, especially in the case where a node cannot be active with different radio interfaces on the same frequency channel.

As far as the results of Sect. 4 are concerned, none of the statements of Theorem 1 but the first one hold in general for the two aforementioned model extensions or in the case where the interference on different channels is described by different conflict graphs. Indeed, its proof used in a crucial way the equivalent description of a network activity state as an (unweighted) partial  $C$ -coloring of the conflict graph, which is not valid anymore in the more general settings described above.

The only exception is the scenario in which every node has exactly  $C$  radio interfaces, each one dedicated to precisely one of the  $C$  frequency channels. The dynamics on different channels become in this way decoupled and the network dynamics can be described by means of  $C$ -dimensional Markov process, each component being an independent single-channel CSMA Markov process. It is easy to show that in this case the aggregate throughput  $\Theta(C)$  is constant in  $C$ , meaning that increasing the number  $C$  of available channels may alleviate the starvation effects without negatively affecting the throughput.

## 8 Conclusions

In this paper, we investigate the performance of random-access networks in which  $C$  non-overlapping channels are available for transmission. We consider a Markovian model for the multi-channel network dynamics evolving according to a CSMA-like algorithm, aiming to understand how the network throughput performance and starvation issues depend on the number of available channels.

The most relevant scenario for our analysis is the one where the number  $C$  of available channels is not sufficient to allow all nodes to be active simultaneously, and for this reason, temporal starvation phenomena persist. Focusing on the high-activation regime, we show how the activity states with a maximum number of active nodes play a crucial role in determining both the aggregate throughput and the temporal starvation effects in the network.

We then characterize the asymptotic behavior of transition times between dominant states by leveraging specific structural properties of the state space where the activity process evolves. This analysis allows us to infer the precise timescales at which starvation phenomena occur in the network, paving the way for a detailed analysis of the delay performance of multi-channel CSMA networks.

The same analytic framework that we develop for expected hitting times also yields an asymptotic lower bound for the mixing time of the process, which we show tends to overestimate the magnitude of the temporal starvation phenomena.

**Acknowledgements** The author acknowledges the NWO Rubicon Grant 680.50.1529 and the Resnick Sustainability Institute at Caltech for the support and is grateful to Sem Borst for the discussions at the early stages of this work.

## References

1. Andrews, M., Zhang, L.: Utility optimization in heterogeneous networks via CSMA-based algorithms. In: 2013 WiOpt Proceedings, pp. 372–379 (2013)
2. Andrews, M., Zhang, L.: Utility optimization in heterogeneous networks via CSMA-based algorithms. *Wirel. Netw.* **23**(1), 219–232 (2017). <https://doi.org/10.1007/s11276-015-1149-z>
3. Berge, C.: Minimax relations for the partial  $q$ -colorings of a graph. *Discrete Math.* **74**(1–2), 3–14 (1989). [https://doi.org/10.1016/0012-365X\(89\)90193-3](https://doi.org/10.1016/0012-365X(89)90193-3)
4. Block, R., Van Houdt, B.: Spatial fairness in multi-channel CSMA line networks. *Perform. Eval.* **103**, 69–85 (2016). <https://doi.org/10.1016/j.peva.2016.06.002>

5. Bonald, T., Feuillet, M.: Performance of CSMA in multi-channel wireless networks. *Queueing Syst.* **72**(1–2), 139–160 (2012). <https://doi.org/10.1007/s11134-012-9300-3>
6. Boorstyn, R., Kershbaum, A., Maglaris, B., Sahin, V.: Throughput analysis in multihop CSMA packet radio networks. *IEEE Trans. Commun.* **35**(3), 267–274 (1987). <https://doi.org/10.1109/TCOM.1987.1096769>
7. Bouman, N., Borst, S., van Leeuwen, J.: Delays and mixing times in random-access networks. In: *Proceedings of the ACM SIGMETRICS* (2013). <https://doi.org/10.1145/2465529.2465759>
8. Bouman, N., Borst, S., van Leeuwen, J.: Delay performance in random-access networks. *Queueing Syst.* **77**(2), 211–242 (2014). <https://doi.org/10.1007/s11134-014-9408-8>
9. Cameron, K.: A min-max relation for the partial  $q$ -colourings of a graph. Part II Box Perfect. *Discrete Math.* **74**(1–2), 15–27 (1989). [https://doi.org/10.1016/0012-365X\(89\)90194-5](https://doi.org/10.1016/0012-365X(89)90194-5)
10. den Hollander, F., Nardi, F., Taati, S.: Metastability of hard-core dynamics on bipartite graphs. *Electron. J. Probab.* **23**, paper no. 97 (2018). <https://doi.org/10.1214/18-EJP210>
11. Ghaderi, J., Srikant, R.: On the design of efficient CSMA algorithms for wireless networks. In: *49th IEEE Conference on Decision and Control (CDC)*, pp. 954–959 (2010). <https://doi.org/10.1109/CDC.2010.5717965>
12. Ghaderi, J., Srikant, R.: The impact of access probabilities on the delay performance of Q-CSMA algorithms in wireless networks. *IEEE/ACM Trans. Netw.* **21**(4), 1063–1075 (2013). <https://doi.org/10.1109/TNET.2012.2215964>
13. Jain, R., Chiu, D.M., Hawe, W.: *A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer System*. Eastern Research Laboratory, Digital Equipment Corporation, Hudson (1984)
14. Jiang, L., Leconte, M., Ni, J., Srikant, R., Walrand, J.: Fast mixing of parallel Glauber dynamics and low-delay CSMA scheduling. In: *2011 Proceedings INFOCOM*, pp. 371–375. *IEEE* (2011). <https://doi.org/10.1109/INFCOM.2011.5935185>
15. Jiang, L., Ni, J., Srikant, R., Walrand, J.: Performance bounds of distributed CSMA scheduling. *Inf. Theory Appl. Workshop (ITA)* **4**, 1–6 (2010). <https://doi.org/10.1109/ITA.2010.5454123>
16. Jiang, L., Shah, D., Shin, J., Walrand, J.: Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inf. Theory* **56**(12), 6182–6207 (2010). <https://doi.org/10.1109/TIT.2010.2081490>
17. Kai, C., Liew, S.C.: Temporal Starvation in CSMA Wireless Networks. In: *2011 IEEE International Conference on Communications (ICC)*, pp. 1–6. *IEEE* (2011). <https://doi.org/10.1109/icc.2011.5963320>
18. Kai, C., Liew, S.C.: Temporal starvation in CSMA wireless networks. *IEEE Trans. Mob. Comput.* **14**(7), 1515–1529 (2015). <https://doi.org/10.1109/TMC.2014.2352256>
19. Lam, K.K., Chau, C.K., Chen, M., Liew, S.C.: Mixing time and temporal starvation of general CSMA networks with multiple frequency agility. *Tech. rep.* (2011)
20. Lam, K.K., Chau, C.K., Chen, M., Liew, S.C.: Mixing time and temporal starvation of general CSMA networks with multiple frequency agility. In: *2012 IEEE International Symposium on Information Theory Proceedings*, pp. 2676–2680. *IEEE* (2012). <https://doi.org/10.1109/ISIT.2012.6284005>
21. Levin, D., Peres, Y., Wilmer, E.: *Markov Chains and Mixing Times*. AMS, Providence, Rhode Island (2009). <https://doi.org/10.1090/mbk/058>
22. Liew, S.C., Zhang, J., Chau, C.K., Chen, M.: Analysis of frequency-agile CSMA wireless networks. Preprint at [arXiv:1007.5255](https://arxiv.org/abs/1007.5255) (2010)
23. Proutière, A., Yi, Y., Lan, T., Chiang, M.: Resource allocation over network dynamics without timescale separation. In: *2010 Proceedings, INFOCOM*, pp. 1–5. *IEEE* (2010). <https://doi.org/10.1109/INFCOM.2010.5462201>
24. Rajagopalan, S., Shah, D., Shin, J.: Network adiabatic theorem: an efficient randomized protocol for contention resolution. *ACM Sigmetrics Perform. Eval. Rev.* **37**(1), 133–144 (2009). <https://doi.org/10.1145/2492101.1555365>
25. Shah, D., Shin, J., Tetali, P.: Medium access using queues. In: *52nd Annual Symposium on Foundations of Computer Science*, pp. 698–707. *IEEE* (2011). <https://doi.org/10.1109/FOCS.2011.99>
26. van de Ven, P., Borst, S., van Leeuwen, J., Proutière, A.: Insensitivity and stability of random-access networks. *Perform. Eval.* **67**(11), 1230–1242 (2010). <https://doi.org/10.1016/j.peva.2010.08.011>

27. Yun, S.Y., Yi, Y., Shin, J., Eun, D.: Optimal CSMA: A survey. In: 2012 IEEE International Conference on Communication Systems, ICCS 2012 pp. 199–204 (2012). <https://doi.org/10.1109/ICCS.2012.6406138>
28. Zocca, A., Borst, S., van Leeuwaarden, J.: Mixing properties of CSMA networks on partite graphs. In: Proceedings of VALUETOOLS 2012, pp. 117–126. IEEE (2012). <https://doi.org/10.4108/valuetools.2012.250264>
29. Zocca, A., Borst, S., van Leeuwaarden, J.: Slow transitions and starvation in dense random-access networks. *Stoch. Models* **31**(3), 361–402 (2015). <https://doi.org/10.1080/15326349.2015.1018441>
30. Zocca, A., Borst, S., van Leeuwaarden, J., Nardi, F.: Delay performance in random-access grid networks. *Perform. Eval.* **70**(10), 900–915 (2013). <https://doi.org/10.1016/j.peva.2013.08.019>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.