



HAL
open science

Distributed Link Scheduling in Wireless Networks

Jean-Claude Bermond, Dorian Mazauric, Vishal Misra, Philippe Nain

► **To cite this version:**

Jean-Claude Bermond, Dorian Mazauric, Vishal Misra, Philippe Nain. Distributed Link Scheduling in Wireless Networks. Discrete Mathematics, Algorithms and Applications, 2020. hal-01977266v1

HAL Id: hal-01977266

<https://inria.hal.science/hal-01977266v1>

Submitted on 10 Jan 2019 (v1), last revised 6 Apr 2020 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed Link Scheduling in Wireless Networks

Jean-Claude Bermond · Dorian Mazaauric · Vishal Misra ·
Philippe Nain

Received: date / Accepted: date

Abstract This work investigates distributed transmission scheduling in wireless networks. Due to interference constraints, “neighboring links” cannot be simultaneously activated, otherwise transmissions will fail. Here, we consider any binary model of interference. We use the model described by Bui, Sanghavi, and Srikant in [BSS09, SBS07]. We assume that time is slotted and during each slot there are two phases: one control phase which determines what links will be activated and a data phase in which data are sent. We assume random arrivals on each link during each slot, so that a queue is associated to each link. Since nodes do not have a global knowledge of the network, our aim (like in [BSS09, SBS07]) is to design for the control phase a distributed algorithm which determines a set of non-interfering links. To be efficient the control phase should be as short as possible; this is done by exchanging control messages during a constant number of mini-slots (constant overhead).

In this paper, we design the first fully distributed local algorithm with the following properties: it works for any arbitrary binary interference model; it has a constant overhead (independent of the size of the network and the values of the queues), and it does not require any knowledge of the queue-lengths. We prove that this algorithm gives a maximal set of active links, where in each interference set there is at least one active link. We also establish sufficient conditions for stability under general Markovian assumptions. Finally, the performance of our algorithm (throughput, stability) is investigated and compared via simulations to that of previously proposed schemes.

Keywords Wireless network · Transmission scheduling · Interference · Distributed algorithm · Stability

1 Introduction

Link scheduling is a key issue in communication networks as it drastically impacts the overall performance of the system (throughput, delay, etc.). It has received a lot of attention for both wired and wireless networks (radio, ad hoc network, sensor network, etc.) - see Section 3 for prior influential work in this area. In a wireless network, an additional difficulty arises from the existence of physical interference, as only transmissions which do not interfere with each other over a given period of time should be scheduled. There are three main classes of interference models: protocols and geometric models [BB10, GK00], in which a geometric notion of interference is used,

Jean-Claude Bermond
Université Côte d’Azur-CNRS-Inria-I3S, 2004 Route des Lucioles, B.P. 93, F-06902, Sophia Antipolis, France
E-mail: jean-claude.bermond@inria.fr

Dorian Mazaauric
Université Côte d’Azur-Inria, 2004 Route des Lucioles, B.P. 93, F-06902, Sophia Antipolis, France
E-mail: dorian.mazaauric@inria.fr

Vishal Misra
Dept. Computer Science, Columbia University, New York, NY, USA
E-mail: misra@cs.columbia.edu

Philippe Nain
Inria, ENS de Lyon/LIP, 46 avenue d’Italie, 69364, Lyon, France
E-mail: philippe.nain@inria.fr

SINR-based models [CXW09], in which correct message reception at a receiver is driven by the experienced SINR (Signal-to-Interference-plus-Noise-Ratio) value, and graph-based models.

In this paper, built on a preliminary work by the authors in [BMMN10], we use the graph-based *binary interference model* (see e.g. [BKK⁺10, BSS09, GLS09, Wan09]). The transmission network is modeled by a graph (undirected or directed), where transmissions may only occur between two neighboring nodes. We give a binary symmetric relation between the links of the network (edges of the undirected graph or arcs of the directed graph) which indicates which pairs of links interfere.

Time is slotted and during a time-slot (or simply slot) transmissions only occur on some links called *active* links; no concurrent transmissions on interfering links are allowed (see Section 2 for accurate definitions).

A well-known example of binary interference is the *primary node interference model* (see [MSS06]), where only links that do not share a common node can be active at the same time. A *matching* is a set of links in a graph without common nodes. In particular, in the primary node interference model, links can be simultaneously activated in the same slot only if they form a matching of the corresponding graph. Consider, for example, the four-node square grid network represented in Figure 1. Here, in each slot, we can simultaneously activate either the two vertical links of the matching 1 (Figure 1(a)) or the two horizontal links of the matching 2 (Figure 1(b)). Note that it is also possible to only activate a single link in this network; however, we will only be interested by maximal sets of active links. Figure 2(a) shows a matching in a larger grid.

The purpose of a link scheduling algorithm is to identify sets of active non-interfering links and, among these sets, to activate the links in the set that will yield the better performance in terms of throughput and delay. In general, links with the largest backlogs should be activated. In the primary node interference model, this corresponds to finding a maximum matching or a large matching. Centralized algorithms have been proposed to solve this problem both for random [TE92, Tas97, SW12b, SW12a] and deterministic arrivals [KMP08].

In this paper, we assume that the message arrival process to each link of the network is stochastic, but with characteristics not necessarily known to the network designers. As a result, a scheduling algorithm has to determine the set(s) of active links as a function of current and past arrivals.

In a decentralized setting, only distributed algorithms with limited local knowledge can be used. This is in contrast with centralized algorithms, which make their decisions based on a total knowledge of the network (i.e. current link backlogs). Devising a distributed algorithm based on a total information about the system state is very difficult to do because of the interference. Actually, acquiring this information is at least as much complicated as devising a decentralized link scheduling algorithm, which is the objective of this paper. In addition, exchanging messages with a central unit in a distributed setting will be both energy consuming and will generate high communication overheads, which will in turn affect the performance.

There are two different classes of scheduling algorithms for wireless networks: contention or random algorithms and collision-free (deterministic) algorithms. Algorithms in the former class may achieve throughput optimality/near throughput optimality in terms of optimizing the wireless network resources (a precise mathematical is given by Tassiulas et al. [TE92] in the context of wireline networks) both without any local information [JW11, SST11, SS12] or with partial local information [RSS09, JW08, JSSW10] but at the expense of some requirements (e.g. perfect carrier sense information) not available in practice. Throughput optimality is typically obtained at the expense of high latencies due, primarily, to the back-off mechanisms that are used to establish the optimality [SBB13]. Algorithms in the second class will in general induce lower latencies than algorithms in the former class, in particular because they are collision-free, but it is extremely challenging to come up with any throughput optimality result for them. In this paper, we will focus on the class of deterministic collision-free algorithms.

Many (collision-free) distributed algorithms which have been proposed yield communication overheads which increase with the size of the network (see e.g. [BZM06, EAM07, MSZ06]). In particular, [GLS09] presents a distributed algorithm valid for any binary interference model but at the expense of a non-constant overhead (which increases with the size of the network). The need for distributed algorithms with a small constant overhead has

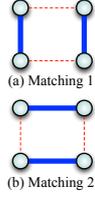


Fig. 1 Scheduling algorithm for grid for $d = 0$.

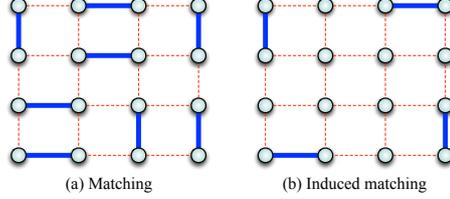


Fig. 2 Two admissible schedules (selected active links for transmission represented by solid bold blue lines).

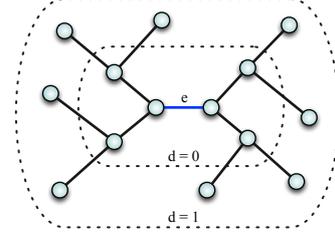


Fig. 3 Interference sets of a link e for $d = 0$ and $d = 1$.

been emphasized in [BSS09,SBS07], where a distributed algorithm with a constant overhead depending on the quality of the desired approximation is described; however, it is only valid for the primary node interference model. An overview of the algorithms in [BSS09,?] is given in Section 3.

Our algorithm can handle any binary symmetric interference model; however, in the examples given in the paper only the so-called d -interference model ([BKK⁺10]) will be considered. In this model, two links interfere if their distance is at most d (a precise definition is given in Section 2 (1)). The case $d = 0$ corresponds to the primary node interference model. At this point, it is worth noting that, already for $d = 1$, the problem of finding a set of links fulfilling the interference constraints and maximizing the sum of the weights (i.e. the total backlog on the links) is *NP-Complete*.

The case $d = 1$ is known as the 802.11 interference model or distance-2 matching problem (see [Wan09,BKK⁺10, BBK⁺04,KMPS04]): here two links interfere if the first one contains a node adjacent to one of the second link. An *induced matching* M in a graph G is a matching where no two edges of M are joined by an edge of G . In the case $d = 1$, a set of non-interfering transmissions form an induced matching (see examples in Figure 2(b)).

The problem of finding a maximum induced matching in a graph (case $d = 1$ in the d -interference model) is *NP-Complete* [SV82,Cam89] and it remains *NP-Complete* even for very special graphs (for 3-regular planar graphs for instance). Furthermore it is proved in [OFGZ08] that the problem of finding a maximum induced matching in a graph cannot be approximated in polynomial time within a factor of $n^{\frac{1}{2}-c}$ for any constant $c > 0$ (unless $P = NP$), where n is the number of vertices in the graph (representing the network).

Under these considerations, we make the following contributions to the general problem:

- we propose, to the best of our knowledge, the first distributed transmission algorithm – called *AlgoLog* – which holds for *any* binary interference model and which has a communication overhead *independent* of the size of the network (Section 4);
- we formally prove some salient properties of *AlgoLog* and compute design tradeoffs. In particular, we show that the overhead grows logarithmically with the *maximum node degree*¹ and requires *no explicit knowledge* of the queue size information between nodes (Section 4);
- under general Markovian assumptions for the arrivals, we identify (Section 5) a set of admissible arrival rates at the links for which the system is stable, a result directly related to the achievable throughput.
- we investigate the performance of *AlgoLog* via simulations (Section 6).

The “link model” above, which is commonly used in the literature (see e.g. [BSS09,CKLS08]) is idealized since in practice only nodes, and not links, can run a link scheduling algorithm. Indeed, a wireless link is immaterial and does not have any computing or sensing capability. The link model is often preferred to the node model as it allows the obtention of generic results as the latter model is much more dependent on the technology (e.g. type of antennas) than the former model. In addition to the contributions in Section 4-6, we show in Section 7 how *AlgoLog* can be emulated if buffers are located in the nodes. First, notation and model are introduced in Section 2 and algorithms in [GLS09,BSS09] are discussed in Section 3.

¹ From a practical standpoint, this presents a more attractive design tradeoff since node degree is bounded by physical layer characteristics whereas the network size can be unbounded.

2 Definitions, notation, network and interference models

Throughout the paper $\mathbb{N} := \{0, 1, \dots\}$ is the set of nonnegative integers and $\mathbb{N}^* := \mathbb{N} - \{0\}$.

The network is modeled as a graph $G = (V, E)$, where a link exists between two nodes (or vertices) $u \in V$ and $v \in V$ if both nodes are within transmission range of one another. If the graph is undirected the link will correspond to the edge $\{u, v\}$. If the graph is directed the link will correspond to the arc (u, v) if u is the sender and v is the receiver. In practice the networks are symmetric digraphs where if there exists an arc (u, v) , then there exists also the arc (v, u) .

Due to interference, not all links can successfully transmit messages in the same slot. Here, we use the so-called Binary Symmetric Interference model. In such a model, to each link e is associated a set of interfering links, denoted by $I(e)$. Note that due to symmetry, if e' belongs to $I(e)$, then e belongs to $I(e')$. We assume that e does not belong to $I(e)$.

The (distributed) link scheduling algorithm proposed in Section 4 applies to any kind of transmission graph (directed or undirected) and interference sets $I(e)$, $e \in E$. However, for the sake of both simplicity and concreteness in the examples, we will mainly consider undirected graphs with the d -interference model (with d fixed in \mathbb{N}) defined below. Here $d(u, v)$ is the distance in G between u and v , that is the length of a shortest path between these two nodes.

$$\text{For every link } e = \{u_1, u_2\} \in E, I(e) = \{\{v_1, v_2\} \in E \setminus \{e\}, \exists i, j \in \{1, 2\}, d(u_i, v_j) \leq d\}. \quad (1)$$

In other words, two links interfere if one node of the first link is located at distance at most d in G from a node of the second link. As said in the introduction, two particular cases have been mainly studied in the literature. The case $d = 0$ is known as the primary node (or node exclusive) interference model. In this model, two links interfere if they are incident, so that an admissible schedule forms a *matching* of the transmission graph G , namely, a set of links without common nodes (Figure 2(a)). The more realistic model with $d = 1$ is known as the *802.11* interference model. In this model, two nodes may communicate only if their neighbors are not involved in a communication. Said otherwise, an admissible schedule is what is called in graph theory an *induced matching* of G , namely a matching where no two edges of M are joined by an edge of G (Figure 2(b)). Two instances of binary interference sets are represented in Figure 3 for these two models: when $d = 0$ (resp. $d = 1$) messages are successfully transmitted on link e during a slot if all links other than link e within the first circle (resp. second circle) remain silent in that slot.

We assume that time is slotted and denote by slot t the time interval $[t, t + 1)$, $t \in \mathbb{N}^*$. Transmissions on different links are synchronized and the unit of information to be transmitted between two nodes is called a *message*. We denote by $c(e) \in \mathbb{N}^*$ the capacity of link $e \in E$, defined as the maximum number of messages that can be transmitted on link e in a slot.

Link $e \in E$ is equipped with a buffer of size $B(e) \in \mathbb{N}^*$ messages called *buffer e* . We denote by $q_t(e)$ the *weight* of link e , defined as the number of messages in buffer e , at the beginning of slot t , waiting to be transmitted on link e . If link e is allowed to transmit messages in slot t , it will (successfully) transmit $\min(c(e), q_t(e))$ messages in that slot. A link $e \in E$ is said to be *busy* in slot t if buffer e is non-empty at the beginning of slot t or, equivalently, if $q_t(e) > 0$.

The purpose of a link scheduling algorithm is to identify an *admissible schedule* in each slot, namely, a set of non-interfering busy links. Furthermore we want to insure the stability of the system. At this time, we do not need to specify the queue length of the weights $\{q_t(e)\}_t$, $e \in E$. Such a model will be introduced in Section 5 when we study the stability of the distributed link scheduling proposed in Section 4.

3 Related works

In this section, we briefly review two algorithms proposed in [GLS09] and [BSS09], respectively. They are both distributed but the algorithm in [GLS09] does not admit a constant overhead, whereas the one described in [BSS09] is valid only for the primary node interference model ($d = 0$).

In the algorithm described in [GLS09], for each time-slot $t \geq 1$ there are a control phase and a data phase, like in our modeling (Section 2). Before each control phase, composed of T mini-slots, each link $e \in E$ is undetermined, and chooses a backoff value $t(e)$, $1 \leq t(e) \leq T$. In this context undetermined means that the link does not know if it will be active or inactive during the data phase of the current slot. If link e receives a message (from a link in

$I(e)$	Links interfering with link e , excluding e
d	d - interference model defined by (1)
$d = 0$	Primary node interference model
$d = 1$	802.11 node interference model
$c(e)$	Capacity of link e (in number of messages/slot)
$B(e)$	Size of buffer e (in number of messages; $B(e) \in [1, \infty]$)
$q_t(e)$	Weight of link e (= number of messages in buffer e) at beginning of slot t
C	Number of colors
$w_t(e)$	Virtual weight of link of e at beginning of slot t function of $q_t(e)$ and C (8)
$\mathbf{v}_t(e) = (v_t(e, 1), \dots, v_t(e, T))$	Control vector of link e in slot t
S	Number of control sub-phases in a slot
$T + 1$	Number of mini-slots in a sub-phase
$K \geq 2$	Number of possible intervals for the virtual weights
L	Initial value of the last interval

Table 1 Glossary of main notation.

$I(e)$) during a mini-slot t , $1 \leq t \leq t(e) - 1$, then e becomes inactive. Otherwise e sends a control message during mini-slot $t(e)$ and if it does not receive a message, then e becomes active (e is inactive otherwise). At the end, a valid set of active links is computed, allowed to send messages during the data phase. This simple algorithm is valid for any interference set, but the choice of the backoff value $t(e)$, for any link $e \in E$, is a function of the weights of links located in its interference set $I(e)$ and function of the weights of links located in interference set of each $e' \in I(e)$. More precisely $t(e)$ depends on the weights of links belonging to the following set $S = \{e', e' \in I(e)\} \cup \{e'', e'' \in I(e'), e' \in I(e)\}$. Thus at each time slot $t \geq 1$, each link has to update the weights of links located in its $2d$ -neighborhood (links at distance at most $2d$ from e), if we have an interference model based on distance d . Therefore the overhead is not constant and furthermore one has to obtain this information, which, due to interference, is a problem as difficult as the transmission scheduling problem.

The algorithm described in [BSS09], *Augmenting Paths Algorithm*, has a constant overhead but it is specific to the primary node interference model ($d = 0$). It uses the “augmenting path tool” developed for matching theory [LP86] and which is used to find polynomial centralized algorithms to determine maximum matchings. In that case, at each slot $t \geq 1$, a valid set of active links is a matching of the transmission graph $G = (V, E)$. The main idea of *Augmenting Paths Algorithm*, is to compute, at a slot $t + 1 > 1$, a matching of G from the matching of G found at slot t . In [BSS09], it is proved that, if $2k + 1$ is the maximum length of the augmenting paths, the algorithm needs a constant overhead of order $4k + 2$ and achieves $\frac{k}{k+2}$ of the capacity region (we can improve via a careful analysis this value to $\frac{k}{k+1}$).

Moreover at the beginning of each slot $t \geq 1$, each node $v \in V$ becomes *seed* with a constant probability p . A seed is a node allowed to start an alternating path. In [BSS09], it is not described how to compute p analytically. Finally it is necessary to precise that the augmenting path technique works only for matchings in graphs and so this algorithm is specific to the primary node interference model ($d = 0$), not the more realistic one. Recall that the problem of determining a maximum matching (if the interference model is defined by $d = 0$) can be done in polynomial time with a centralized algorithm, but for interference models defined by $d \geq 1$, determining a maximum valid set of links is an NP-complete problem. In Section 6.2 we compare via simulations the performance of *Augmenting Paths Algorithm* to the performance of our distributed algorithm, proposed in Section 4, for a square grid composed of 121 nodes.

4 Presentation of AlgoLog

4.1 Aim of the control phase

In this section, we propose a distributed link scheduling algorithm - called *AlgoLog* - with a constant communication overhead and which is valid for any binary interference model. (Recall that for the sake of simplicity, we will only consider binary d -interference sets as defined Equation (1).) This is in contrast with the algorithm proposed in [GLS09] whose communication overhead increases with the size of the network, and with the algorithm proposed in [BSS09, SBS07] which works only for the primary node interference model ($d = 0$, see Section 2).

Link scheduling algorithms proposed in the literature like those in [BSS09, GLS09] require exchanges of information between nodes (typically, any node needs to know the number of pending messages of its neighbors) which is in itself a difficult task due to interference. Our algorithm does not require any exchange of information between nodes and is therefore fully distributed.

In `ALgoLog` the time is slotted, with slot t referring to the time interval $[t, t+1)$. A slot is composed of two different phases, a *control* phase followed by a *data* phase. The aim of the control phase is to determine all links – called **active links** – which will be allowed to send messages in the data phase of slot t . At slot t , a link $e \in E$ is said *busy* if $q_t(e) > 0$. Links with $q_t(e) = 0$ will not participate in the control phase. Our first goal is to ensure that all active links form a maximal admissible schedule, i.e. to determine a maximal set E' of non-interfering busy links (a set is maximal if a busy link in $E - E'$ cannot be activated without interfering with a link in E'). To do that, `ALgoLog` will use interference as information. Indeed, when a signal is sent on a link $e \in E$, every link $e' \in I(e)$ hears it.

In addition to finding a maximal admissible schedule, `ALgoLog` activates all links whose virtual weights are local maximum (see just after the definition of a virtual weight which is related to the values of $q_t(e)/c(e)$).

4.2 Virtual weights and their binary representation

Throughout the paper W is a fixed nonnegative integer whose role and value will be discussed later on. To each link $e \in E$, we associate an integer $w_t(e)$ at the beginning of slot $t \geq 1$ – called the *virtual weight* of link e at time t – defined precisely in Section 4.5 (see (9)).

The mapping $e \rightarrow w_t(e)$ will satisfy the following properties: for all $t \geq 1$,

$$w_t(e) \in \{0, 1, \dots, W\},$$

$$w_t(e) = 0 \text{ if and only if } q_t(e) = 0, \quad (2)$$

and two interfering busy links at time t will have different virtual weights namely,

$$\text{if } q_t(e) > 0 \text{ and } q_t(e') > 0 \text{ then } w_t(e') \neq w_t(e) \text{ for all } e' \in I(e), e \in E. \quad (3)$$

There are two main reasons for working with virtual weights, and not with the actual weights: one is to impose that two interfering busy links at time t will have different (virtual) weights, a feature that `ALgoLog` will use, and the other one is to parametrize (through W) the duration of the control phase, typically to limit its duration - see Section 4.5.

Since the virtual weights $\{w_t(e), e \in E\}$ always lie in the set $\{0, 1, \dots, W\}$, we may (and will) identify $w_t(e)$ with its binary representation the (so-called) *control vector* $\mathbf{v}_t(e)$ defined by

$$\mathbf{v}_t(e) = (v_t(e, 1), \dots, v_t(e, T)) \in \{0, 1\}^T, \quad (4)$$

where binary numbers $v_t(e, 1), \dots, v_t(e, T)$ enter the decomposition of $w_t(e)$ in base 2, namely,

$$w_t(e) = \sum_{i=1}^T 2^{T-i} v_t(e, i), \quad (5)$$

with $T := \min\{q \in \mathbb{N}^* : W \leq 2^q - 1\}$. Alternatively,

$$T = \lceil \log_2(W + 1) \rceil, \quad (6)$$

with $\lceil x \rceil$ the smallest integer greater than or equal to x . Table 3 shows the control vectors associated to values of $w_t(e)$ between 1 and 15 (here $T = 4$).

Require: $w_t(e)$.

Ensure: return *active* ($s(e) = A$) or *inactive* ($s(e) = I$).

```

1:  $e$  computes  $v_t(e, i)$ ,  $i = 1, \dots, T$ 
2:  $s(e) = I$  if  $w_t(e) = 0$  otherwise  $s(e) = U$  ( $w_t(e) > 0$ )
3: for  $j = 1, \dots, S$  do
4:   for  $i = 1, \dots, T$  do
5:     if  $s(e) = U$  and  $v_t(e, i) = 1$  then
6:        $e$  sends a signal (heard by links in  $I(e)$ )
7:       if  $e$  does not hear any signal then
8:          $s(e) = A$ 
9:       if  $s(e) = U$ ,  $v_t(e, i) = 0$ , and  $e$  hears a signal then
10:         $s(e) = PI$ 
11:     if  $s(e) = A$  then
12:        $e$  sends a signal (of synchronization)
13:     if  $s(e) = PI$  and  $e$  does not receive a signal then
14:        $s(e) = U$ 
15:     if  $s(e) = PI$  and  $e$  receives a signal then
16:        $s(e) = I$ 
17: return  $s(e)$ 

```

Table 2 Control phase of AlgoLog at link $e \in E$ in slot t .

4.3 The algorithm AlgoLog

The control phase is composed of $S \geq 1$ successive sub-phases labeled $j = 1, \dots, S$ with a mini-slot of synchronization between two consecutive sub-phases. We will see later that we can choose $S = T$, so that the parameter T , or equivalently W from (6), controls the duration of the control phase.

At the end of the control phase a link will be in one of the two final states: **active** (state A) or **inactive** (state I). However, during the running of the algorithm a link can also be in one of the two other temporary states: **undetermined** (state U) or **potentially inactive** (state PI). The possible transitions between the states are described in Figure 4. During a sub-phase an undetermined link may become either definitely active or potentially inactive. The objective of the mini-slot of synchronization between two sub-phases is to determine which potentially inactive links will become definitely inactive; the others will become undetermined and will participate in the next sub-phase.

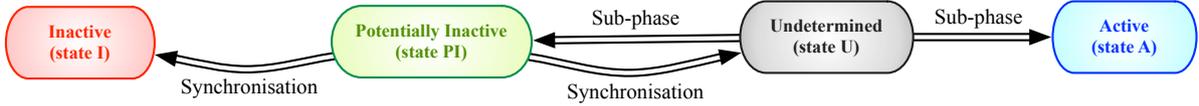


Fig. 4 The four different possible states and transitions for links in AlgoLog.

We can now describe precisely the algorithm which is also represented in pseudocode in Table 2, where $s(e) \in \{A, I, U, PI\}$ gives the state of link e . The reader can also follow the running of the algorithm in the examples in Figures 6-8. In Figure 6 the graph consists of a cycle with 9 links (with $d = 0$); to make the figure more readable we have repeated the first node u at the end. In Figure 7 the graph is a grid with 16 nodes and 24 links (with $d = 0$), and in Figure 8 the graph is a random graph with 48 nodes and 91 links (with $d = 1$). The random graph has been generated using the method described in Section 6.3. Active links are indicated with thick bold blue lines, inactive links with thin dashed red lines, undetermined links with thin black lines, and potentially inactive links with dotted black lines.

At the beginning of the algorithm (**line 2**) link e marks itself as inactive if $w_t(e) = 0$ and as undetermined otherwise ($w_t(e) > 0$). Each sub-phase is divided itself into T (so-called) *mini-slots*, labeled $i = 1, \dots, T$. Each link e undetermined at the beginning of the mini-slot i does the following in the mini-slot i :

- (**lines 5-8**) if $v_t(e, i) = 1$, link e sends a signal (which will be heard by all links in $I(e)$). If link e does not hear anything from other links during mini-slot i , it marks itself as active, otherwise it stays undetermined;
- (**lines 9-10**) if $v_t(e, i) = 0$, link e does not transmit in mini-slot i . If link e does not hear anything from other links during mini-slot i , then it stays undetermined; otherwise it marks itself as potentially inactive.

Figure 5 describes the four transitions described before.

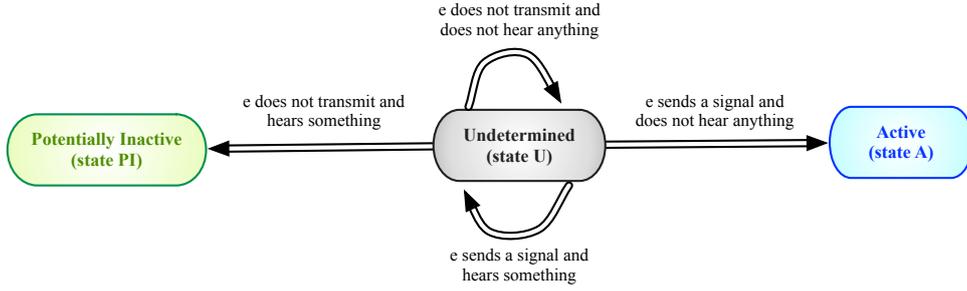


Fig. 5 The four transitions and the three states in one sub-phase of `AlgoLog` for a link e .

In Figure 6 at the beginning of the algorithm the last link is inactive (thin dotted red line) as $q_t(e) = 0$ and all the other links are undetermined (solid black lines). Here $T = 4$ and the 4 mini-slots of sub-phase 1 are indicated in lines (d)-(g). At the end of the 4th mini-slot only link 5 is active (represented by a solid bold blue line), the others being potentially inactive or inactive (for the last link).

`AlgoLog` could stop after the end of the first sub-phase (**line 10** with $j = 1$ and $i = T$) since, by construction, all active links form an admissible schedule. It is however possible that this admissible schedule is not maximal. Such a situation is depicted in Figure 6 (for $d = 0$) where at the end of mini-slot T (in this example $T = 4$) only link 5 is active and the others are potentially inactive (or inactive). But we will see that in this example, at the end of the S subphases of `AlgoLog`, links 1, 3, 5, 7 are active yielding a maximal admissible schedule (Figure 6 (n)).

To enforce an admissible schedule to be maximal, the following is done: at the end of the first sub-phase each active link transmits a signal, called a *synchronization signal* (**lines 11-12**). All potentially inactive links that do not hear anything during this additional mini-slot mark themselves as undetermined (**lines 13-14**); the other potentially inactive links (i.e. those that do hear something from neighbor(s)) are marked inactive (**lines 15-16**). `AlgoLog` then enters the second sub-phase (**line 4** with $j = 2$, $i = 1$) and the previous process is repeated until the end of the S -th sub-phase (**line 17** with $j = S$ and $i = T$).

We will show in Proposition 2 that property (3) together with the choice for $S = T$ implies that `AlgoLog` always generates a maximal admissible schedule.

In Figure 6 (h), during the mini-slot of synchronization links 4 and 6 become inactive the other links 1, 2, 3, 7, 8 becoming undetermined. After the first four mini-slots of sub-phase 2 (Figure 6 (i)-(l)), links 3 and 7 become active links and 1, 2, 8 are potentially inactive. During the mini-slot of synchronization (Figure 6 (m)) links 2, 8 become inactive and link 1 undetermined. At the end of the next sub-phase (Figure 6 (n)) link 1 becomes active.

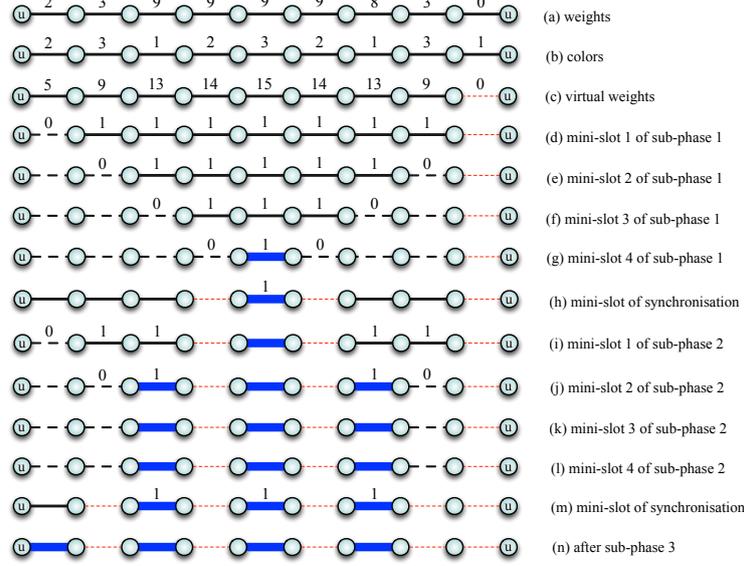


Fig. 6 Output of AlgoLog for a cycle composed of 9 links for $d = 0$, $C = 3$, $K = 5$, $L = 4$, $W = 15$, and $T = 4$. Active links are indicated with thick bold blue lines, inactive links with thin dashed red lines, undetermined links with thin black lines, and potentially inactive links with dotted black lines.

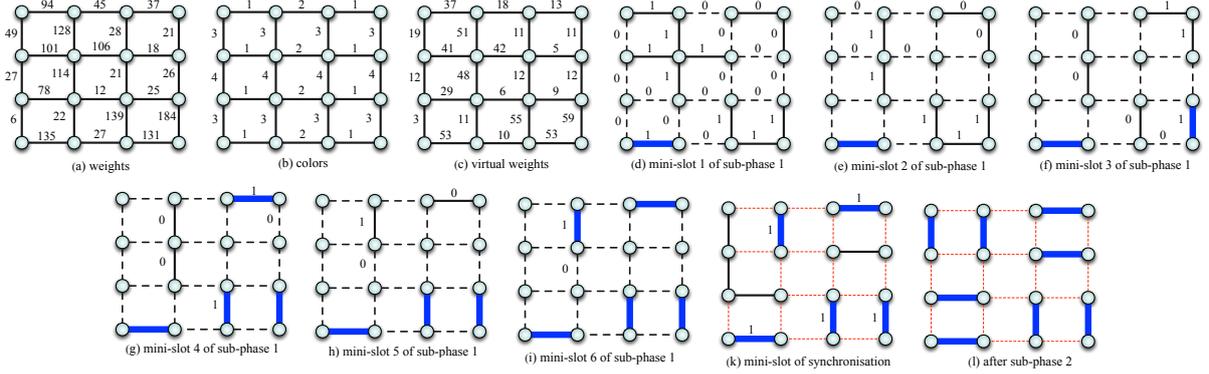


Fig. 7 Output of AlgoLog for a grid composed of 24 links for $d = 0$, $C = 4$, $K = 15$, $L = 140$, $W = 60$, and $T = 6$.

4.4 Data phase

All active links at the end of the control phase of a slot transmit during the data phase of that slot. More specifically, if link e is active in slot t as the result of the control phase of AlgoLog , then it will transmit $\min(q_t(e), c(e))$ messages during the data phase of that slot.

4.5 Definition of the virtual weights

Recall that we want that two interfering links have different virtual weights. Furthermore, to enhance the performance of AlgoLog (stability, delay, ...), we want as much as possible to schedule links with the highest backlogs or, by fairness to links with small bandwidth, links with the highest ratios $q_t(e)/c(e)$ in slot t ($t \geq 1$).

To overcome the problem that interfering links may have the same virtual weights, we will use an *admissible coloring* of the links, which consists in associating to each link e a color $\gamma(e)$ so that any two interfering links have different colors. Let C be the number of colors of such an admissible coloring. We have $\gamma(e) \in \{1, 2, \dots, C\}$. We will see in Section 4.7 how to compute or approximate the minimum value for C in the d -interference model.

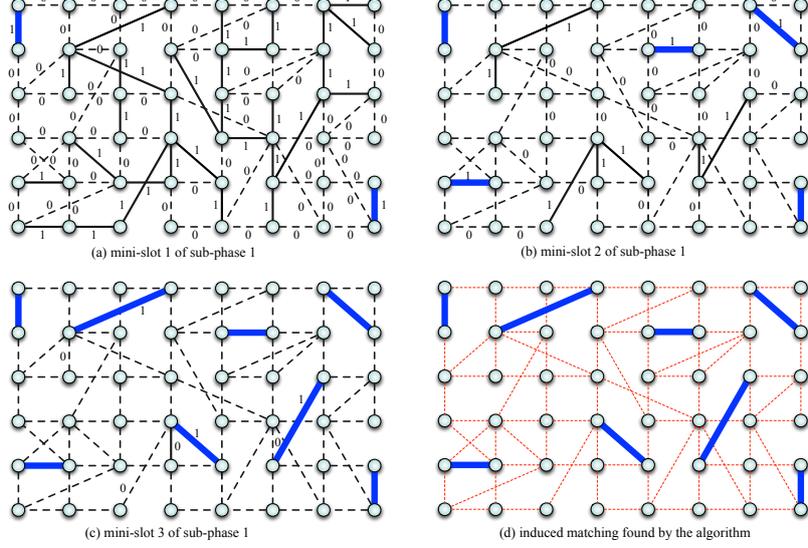


Fig. 8 Output of AlgoLog for a graph composed of 91 links for $d = 1$.

Figure 6 (b) shows an admissible coloring of a cycle graph with 9 links with $C = 3$ and Figure 7 (b) an admissible coloring of the grid with $C = 4$ colors, both within the primary node interference model ($d = 0$).

In fact, to ensure fairness between the links we will use a *time-dependent coloring*, which associates to each link e in slot $t \geq 1$ a color $\gamma_t(e) \in \{1, 2, \dots, C\}$ defined by:

$$\gamma_t(e) = (\gamma(e) + t - 1) \bmod(C) \quad e \in E, t \geq 1, \quad (7)$$

where, here, $p \bmod(C) = p - C \lfloor (p-1)/C \rfloor$ for any integer $p \geq 1$, so that $p \bmod(C) \in \{1, 2, \dots, C\}$.

As the initial coloring $\gamma_1 := (\gamma(e), e \in E)$ is admissible, $\gamma_t := (\gamma_t(e), e \in E)$ is also an admissible coloring for all $t \geq 2$. We state this fact as a lemma:

Lemma 1 *If e and e' are interfering links, then $\gamma_t(e) \neq \gamma_t(e')$ for all $t \geq 1$.*

From now on we will work under the following assumption.

Assumption 1 (Initial coloring) *The initial coloring $\gamma(e), e \in E$, is known and fixed throughout.*

To insure that links with the highest ratios $q_t(e)/c(e)$ will have more chance to be scheduled in slot t , to each link e with a positive weight (i.e. $q_t(e) > 0$) we will assign the *virtual weight* $w_t(e)$ defined as

$$w_t(e) = \begin{cases} Cg(q_t(e)/c(e)) + \gamma_t(e) & \text{if } q_t(e) > 0 \\ 0 & \text{if } q_t(e) = 0. \end{cases} \quad (8)$$

The mapping g is defined as follows: $[0, \infty)$ is partitioned into $K \geq 2$ (necessarily) disjoint sets I_0, \dots, I_{K-1} , where $I_0 \cup \dots \cup I_{K-2} = \{x, x \leq L\}$ and $I_{K-1} = \{x, x > L\}$ with $L \geq 0$. Then,

$$g(x) = k \quad \text{if } x \in I_k \quad \text{for } k = 0, 1, \dots, K-1.$$

In summary, either $q_t(e) = 0$ and then $w_t(e) = 0$, or $q_t(e) > 0$ and

$$w_t(e) = C \sum_{k=0}^{K-1} k \mathbf{1}\{q_t(e)/c(e) \in I_k\} + \gamma_t(e) \in \{1, \dots, CK\}. \quad (9)$$

Note that there is no need to select $L \geq \max_{e \in E} B(e)/c(e)$ since $\max_{e \in E} q_t(e)/c(e) \leq \max_{e \in E} B(e)/c(e)$ for all t . Also notice that, as $k \leq K-1$ then $w_t(e) \leq C(K-1) + C = CK$ for all t , so that the maximum virtual weight W satisfies

$$W = CK. \quad (10)$$

$q_t(e)/c(e)$	$\gamma_t(e)$	$w_t(e)$	$v_t(e,1)$	$v_t(e,2)$	$v_t(e,3)$	$v_t(e,4)$
(0,1]	1	1	0	0	0	1
(0,1]	2	2	0	0	1	0
(0,1]	3	3	0	0	1	1
(1,2]	1	4	0	1	0	0
(1,2]	2	5	0	1	0	1
(1,2]	3	6	0	1	1	0
(2,3]	1	7	0	1	1	1
(2,3]	2	8	1	0	0	0
(2,3]	3	9	1	0	0	1
(3,4]	1	10	1	0	1	0
(3,4]	2	11	1	0	1	1
(3,4]	3	12	1	1	0	0
> 4	1	13	1	1	0	1
> 4	2	14	1	1	1	0
> 4	3	15	1	1	1	1

Table 3 ($w_t(e), v_t(e)$) for all entries $(q_t(e)/c(e), \gamma_t(e))$ for $C = 3, K = 5, L = 4$, and $W = 15$.

From (6) and (10), we obtain

$$T = \log_2(\lceil CK + 1 \rceil). \quad (11)$$

For the first example of the cycle (cf. Figure 6 and Table 3), we choose $L = 4, K = 5$, intervals $I_k = (k, k + 1]$ for $k = 0, 1, 2, 3$ and $I_4 = (4, \infty)$. For the example of the grid (Figure 7) we choose $L = 140, K = 15$ and $I_k = (10k, 10k + 10]$ for $0 \leq k \leq 13$ and $I_{14} = (140, \infty)$.

The following lemma shows that property in Eq. (3) holds.

Lemma 2 *Let $e, e' \in E$ be two busy interfering links in slot $t \geq 1$. Then, $w_t(e) \neq w_t(e')$.*

Proof. Fix $t \geq 1$. There exist j and k such that $\lfloor q_t(e)/c(e) \rfloor \in I_j$ and $\lfloor q_t(e')/c(e') \rfloor \in I_k$. Assume that $j \neq k$. Then (cf. (8))

$$\begin{aligned} |w_t(e) - w_t(e')| &= |C(j - k) - (\gamma_t(e') - \gamma_t(e))| \\ &\geq C|j - k| - |\gamma_t(e') - \gamma_t(e)| \quad (\text{Hint: } |a - b| \geq |a| - |b|) \\ &\geq C|j - k| - (C - 1) \\ &\geq 1, \end{aligned} \quad (12)$$

where (12) follows from the fact that $1 \leq \gamma_t(e), \gamma_t(e') \leq C$. Hence, $w_t(e) \neq w_t(e')$. Assume that $j = k$. We have (cf. (8)) $w_t(e) - w_t(e') = \gamma_t(e) - \gamma_t(e')$. Therefore, as $\gamma_t(e) \neq \gamma_t(e')$ from Lemma 1, we get that $w_t(e) \neq w_t(e')$. ■

4.6 AlgoLog generates a maximal schedule

We will first show that AlgoLog activates all links whose virtual weights are local maximum, and then that AlgoLog generates a maximal admissible schedule. We say that a busy link e is a *local maximum* in slot $t \geq 1$, if $w_t(e) > w_t(e')$ for all $e' \in I(e)$. Note that, by Lemma 2, $w_t(e) \neq w_t(e')$ for $e' \in I(e)$, which justifies the strict inequality in the definition of a local maximum.

Proposition 1 *A local maximum link is always active at the end of a control phase of AlgoLog. More precisely, it is already active at the end of sub-phase 1.*

Proof.

Assume that link e is local maximum and that it is not active at the end of mini-slot T of sub-phase 1. So it is either PI (potentially inactive) or U (undertermined) at the end of mini-slot T of sub-phase 1. If link e has become PI in mini-slot $i, 1 \leq i \leq T$, that is due to the existence of a link $e' \in I(e)$ such that link e' is U at the beginning of mini-slot i and coordinate i of their vector controls satisfy $v_t(e', i) = 1$ and $v_t(e, i) = 0$. If link e is U, let i be

the last mini-slot where e has sent a message; as it stayed U there exists a link $e' \in I(e)$ such that link e' is U at the beginning of mini-slot i and coordinate i of their vector controls satisfy $v_t(e', i) = 1$ and $v_t(e, i) = 1$. But, as $w_t(e) \geq w_t(e')$, in both cases there exists an $\ell < i$ such that $v_t(e, \ell) = 1$ and $v_t(e', \ell) = 0$ and so at mini-slot ℓ , e' should have become PI, yielding a contradiction. ■

In essence, Proposition 1 says that links with ‘large’ weights (from a local point of view) will be more likely to be scheduled for transmission since the virtual weight of a link is a non-decreasing function of its weight (see (8) and (9)).

The statement in Proposition 1 should however be interpreted with care. In particular, it does not imply that, in any slot t , `AlgoLog` always finds a maximum schedule that is one for which the sum of the virtual weights of the activated links is maximized. To illustrate this point, consider a path composed of three links: e_1 , e_2 , and e_3 . Let us suppose we have the primary node interference model ($d = 0$); so $C = 2$. Assume that the virtual weights satisfy: $w_t(e_2) > w_t(e_1) = w_t(e_3)$ but $w_t(e_1) + w_t(e_3) > w_t(e_2)$. As a result, `AlgoLog` will only activate link e_2 (which is a local maximum) in slot t (note that this schedule is maximal). Note, however, that a maximum schedule consists of link e_1 and e_3 as $w_t(e_1) + w_t(e_3) > w_t(e_2)$, thereby showing that `AlgoLog` does not always find a maximum schedule. Note also that the situation might be worse with the real weights. Indeed we can have, by equation 8, $w_t(e) > w_t(e')$, if $\gamma_t(e) > \gamma_t(e')$ and $q_t(e)/c(e)$ and $q_t(e')/c(e')$ belong to the same interval although $q_t(e')/c(e')$ might be considerably larger than $q_t(e)/c(e)$.

The above example shows that `AlgoLog` will not allow the network to operate at its maximum throughput (defined as the average number of transmissions per slot) since it cannot ensure that no power transmission will be lost. Indeed, in the example above, it would be more efficient to activate links e_1 and e_3 instead of link e_2 (the throughput can be the double if $q_t(e_2) = q_t(e_1) = q_t(e_3)$ and $\gamma_t(e_2) = 2$ and $\gamma_t(e_1) = \gamma_t(e_3) = 1$). This is of course not surprising as this is the price to pay for using a decentralized scheduling algorithm like `AlgoLog`.

We now prove that `AlgoLog` finds a maximal schedule when $S \geq T$, namely, a schedule with the property that each inactive busy link in E is such that there is (at least) one active link in its interference set.

Proposition 2 *Let $S \geq T = \log_2(\lceil CK + 1 \rceil)$. For each slot $t \geq 1$ and for every $e \in E$ such that $q_t(e) > 0$, there exists at least one link $e' \in I(e) \cup \{e\}$ such that e' is active at the end of the control phase of slot t .*

Proof. Consider an arbitrary slot t and an arbitrary link e such that $w_t(e) > 0$ (i.e. e is busy).

- If e is active at the end of the control phase of t , then the result is proved by taking $e' = e$.
- If e is inactive at the end of the control phase of t , that is due to the fact that, when e became inactive (during a mini-slot of synchronization), there was some link $e' \in I(e)$ active, and the proposition is proved with e' .
- If e is neither active nor inactive at the end of the control phase of slot t , then it is U (undetermined) at the end of *all* sub-phases (after the mini-slot of synchronization) in which case e is either U (Undetermined) or PI (Potentially Inactive) at the end of mini-slot T of any sub-phase. We will see that if the number S of sub-phases satisfies $S \geq T$ this case cannot happen.

More precisely, we will prove that in any sub-phase j , $j = 1, \dots, S$, there exists a link which became PI no sooner than in mini-slot $2 + S - j$. In particular, at the end of sub-phase 1, there exists a link which became PI no sooner than in mini-slot $S + 1$, which yields a contradiction when $S \geq T$.

First, let us see how a link f can be PI or U at the end of mini-slot T of any sub-phase. The reader can follow the proof on the example of Figure 9, where `AlgoLog` is applied to a path with 5 links, labeled e_1, e_2, e_3, e_4, e_5 , $d = 0, W = 8$, and $T = 3$.

Case 1: let f be PI at the end of mini-slot T of a sub-phase and let $i(f)$ be the mini-slot of this sub-phase in which f became PI. The latter is due to the existence of a link $f' \in I(f)$ which, at mini-slot $i(f)$, was U with the $i(f)$ coordinate of its control vector given by $v_t(f', i(f)) = 1$ (while $v_t(f, i(f)) = 0$). Hence, f' will be PI, U or A at the end of the sub-phase. Note that, if $f = e$, then f' cannot be A, as otherwise e would become inactive during the mini-slot of synchronization, implying that it would be inactive at the end of the control phase of t , thereby contradicting the assumption that it is U or PI at the end of the control phase of t .

In the example of Figure 9 it is the case in the sub-phase 1 for link e_1 (resp. e_3, e_4) which became PI at mini-slot 1 (resp. 2, 3) or in sub-phase 2 for links e_1 and e_2 .

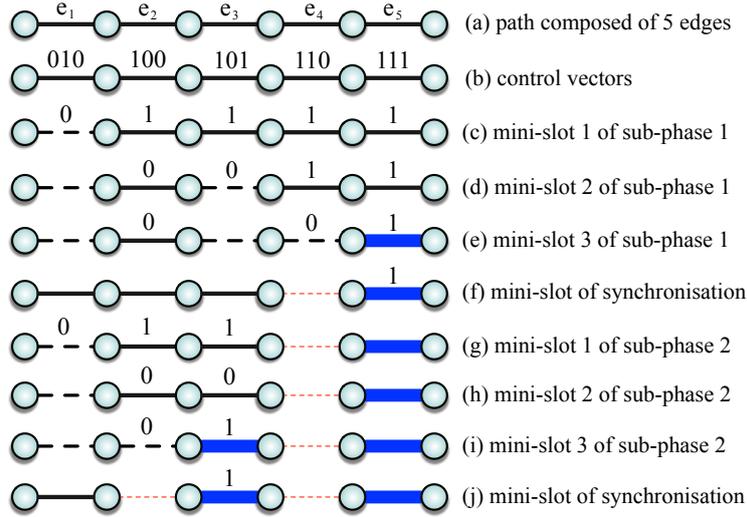


Fig. 9 Example showing the necessity of $S \geq T$ in Proposition 2. Active links are indicated with thick bold blue lines, inactive links with thin dashed red lines, undetermined links with thin black lines, and potentially inactive links with dotted black lines.

Case 2: let f be U at the end of mini-slot T of some sub-phase and let $i(f)$ be the last mini-slot of this sub-phase where link f has sent a signal. We claim that it is due to the existence of a link $f' \in I(f)$ which became PI in a mini-slot $i(f') > i(f)$. Indeed, if link f is U at the end of mini-slot T then it was U at mini-slot $i(f)$. But the latter is due to another undetermined link $f' \in I(f)$, which also sent a signal during mini-slot $i(f)$. Let $v_t(f)$ and $v_t(f')$ be the control vectors of f and f' , respectively.

First, note they have the same coordinates for $i = 1, \dots, i(f)$. To see why, suppose that for some $i < i(f)$, $v_t(f, i)$ and $v_t(f', i)$ differ; then, as both links f and f' were U, the one with the coordinate 0 should have become PI due to the other one, and so could not have been U at mini-slot $i(f)$. Now, since $w_t(f) \neq w_t(f')$ from Lemma 2, there is a mini-slot $l > i(f)$, where $v_t(f, l)$ and $v_t(f', l)$ differ. However, as $i(f)$ was the last mini-slot where f sent a signal, we have $v_t(f, l) = 0$ and so $v_t(f', l) = 1$. But f being U at the end of mini-slot T , this is possible only if, at mini slot l , f' is PI. Therefore, f' became PI at some mini-slot $i(f') > i(f)$ as, at mini-slot $i(f)$, it was U with coordinate 1.

In the example of Figure 9, this situation shows up for link $f = e_2$ at the end of sub-phase 1. Indeed, e_2 sent its last signal in mini-slot 1 and stayed U in this mini-slot due to link $f' = e_3$. Here, $v_t(e_2, 3) = 0$ and $v_t(e_3, 3) = 1$; but e_3 is F and so e_2 stays U at mini-slot 3.

Recall that we suppose e is U or PI at the end of mini-slot T of all sub-phases. Consider a sub-phase j .

- If e is U at the end of mini-slot T of sub-phase j and if i_1^j is the last mini-slot of this sub-phase j where link e has sent a signal, then by case 2 there exists a link $e_2^j \in I(e)$ which became PI in a mini-slot $i_2^j > i_1^j \geq 1$, and as a result $i_2^j \geq 2$.
- If e is PI at the end of mini-slot T of sub-phase j and i_1^j denotes the mini-slot in which e became PI, then, by case 1, there exists a link $e_2^j \in I(e)$ which is PI or U at the end of mini-slot T of sub-phase j , with the i_1^j -th coordinate of its control vector equal to 1 (or, equivalently, e_2^j has sent a signal at mini-slot i_1^j). As we noted in case 1, e_2^j cannot be A, as otherwise e would be inactive at the end of the sub-phase j (due to the synchronization slot), and so will be Inactive at the end of control phase of t .
- If e_2^j is U and i_2^j is the last mini-slot of this sub-phase j where link e_2^j has sent a signal, then, by case 2 there exists a link $e_3^j \in I(e_2^j)$ which became PI at a mini-slot $i_3^j > i_2^j \geq i_1^j$, so that $i_3^j \geq 2$.
- If e_2^j is PI, as the i_1^j -th coordinate of its control vector is equal to 1, it became PI at a mini-slot $i_2^j > i_1^j \geq 1$, and so $i_2^j \geq 2$. Furthermore, there exists by case 1 a link $e_3^j \in I(e_2^j)$ which is PI or U or A at the end of mini-slot T of sub-phase j .
- If e_3^j is A, we stop with the sequence e, e_2^j, e_3^j . That is the case of sub-phase 2 in the example of Figure 9, with the sequence e_1, e_2, e_3 respectively PI, PI, A.

Otherwise, as e_3^j is PI (resp. U), there exists by case 1 (resp. 2) a link $e_4^j \in I(e_3^j)$ which is A, PI or U. If e_4^j is A, we stop with the sequence e, e_2^j, e_3^j, e_4^j ; otherwise we continue the process with e_4^j until we find an active link. In the sub-phase 1 in the example of Figure 9, we get the sequence e_1, e_2, e_3, e_4, e_5 respectively PI, U, PI, PI, A.

Summarizing the above construction, we have created at the end of mini-slot T of any sub-phase j a sequence of links $\mathbf{S}^j = (e = e_1^j, e_2^j, \dots, e_{k_j}^j, e_{k_j+1}^j)$ which satisfies the following properties:

- (i) link e_m^j is U or PI for $m = 1, \dots, k_j - 1$; link $e_{k_j}^j$ is PI and link $e_{k_j+1}^j$ is A.
- (ii) link $e_{m+1}^j \in I(e_m^j)$ for $m = 1, \dots, k_j$.
- (iii) If e_m^j is U and i_m^j is the last mini-slot where e_m^j has sent a signal, then e_{m+1}^j is PI and became PI at a mini-slot $i_{m+1}^j > i_m^j$.
- (iv) If e_m^j is PI and became PI at the mini-slot i_m^j , then e_{m+1}^j has sent a signal in mini-slot i_m^j . Properties (iii) and (iv) imply the following property
- (v) Suppose e_m^j and $e_{m'}^j$ became PI respectively at mini-slots i_m^j and $i_{m'}^j$; then if $m' > m$, $i_{m'}^j > i_m^j$. (Said otherwise, a link becomes PI at a mini-slot strictly greater than the mini-slot where the preceding PI link in the sequence became PI).

It suffices to prove (v) for two consecutive PI links in the sequence. Suppose e_m^j became PI at mini-slot i_m^j . If the next PI link is e_{m+1}^j , then by property (iv), as e_{m+1}^j has sent a signal in mini-slot i_m^j , it becomes PI at a mini-slot $i_{m+1}^j > i_m^j$. If the link e_{m+1}^j is U then e_{m+2}^j is PI and is the first PI link after e_m^j ; by property (iv) e_{m+1}^j has sent a signal in mini-slot i_m^j and so has sent its last signal in the mini-slot $i_{m+1}^j \geq i_m^j$. By property (iii) applied to e_{m+1}^j we have $i_{m+2}^j > i_{m+1}^j \geq i_m^j$.

We now have all the ingredients to prove by induction on $S - j$ the following claim, which will imply that there exists a link which became PI in sub-phase 1 at a mini-slot at least $S + 1$, giving a contradiction when $S \geq T$.

Claim At sub-phase j , there exists a sequence \mathbf{S}^j as described above whose last PI link $e_{k_j}^j$ became PI at a mini-slot $i_{k_j}^j \geq 2 + S - j$.

Proof. The claim is true for $j = S$ ($S - j = 0$). In fact, for any j , $i_{k_j}^j \geq 2$; indeed we have seen that either e_2^j is PI and then $i(e_2^j) \geq 2$ or e_2^j is U and then $i(e_3^j) \geq 2$. By property (v) it follows that the last link became PI at a mini-slot ≥ 2 .

Assume that the claim is true for j and let us prove that it is still true for $(j - 1)$ (induction on $S - j$). The induction assumption implies that there exists a sequence $\mathbf{S}^j = (e = e_1^j, e_2^j, \dots, e_{k_j}^j, e_{k_j+1}^j)$ whose last PI link $e_{k_j}^j$ became PI at a mini-slot $i_{k_j}^j \geq 2 + S - j$. Consider the same sequence at sub-phase $j - 1$. Link $e_{k_j+1}^j$ (which was A at the end of sub-phase j) cannot be A at the end of the sub-phase $j - 1$, otherwise during the synchronization slot of sub-phase $j - 1$, link $e_{k_j}^j$ would have become Inactive and would have not participated in sub-phase j . If link $e_{k_j+1}^j$ is PI at the end of mini-slot T of sub-phase $j - 1$, then by property (v) it became PI at a mini-slot $i_{k_j+1}^{j-1} > i_{k_j}^j \geq 2 + S - j$. If link $e_{k_j+1}^j$ is U at the end of mini-slot T of sub-phase $j - 1$, then by case 2, there exists a link $e_{k_j+2}^{j-1}$ which became PI at a mini-slot $i_{k_j+2}^{j-1} > i_{k_j+1}^j \geq 2 + S - j$.

The above construction generates a sequence $\mathbf{S}^{j-1} = (e = e_1^j, e_2^j, \dots, e_{k_j}^j, e_{k_j+1}^j, \dots, e_{k_{j-1}+1}^{j-1})$ satisfying all the properties (i)-(iv), and which contains a link (either $e_{k_j+1}^j$ or $e_{k_j+2}^j$) which became PI at a slot $> 2 + S - j$. Therefore, by property (v), the last link of \mathbf{S}^{j-1} also became PI at a mini-slot $i_{k_{j-1}}^{j-1} \geq 2 + S - (j - 1)$. ■

We are now in position to conclude the proof of the proposition. By the claim above, there exists at sub-phase 1 a PI link $e_{k_1}^1$ which became PI at a mini-slot $i_{k_1}^1 \geq S + 1$. However, in a sub-phase there are exactly T mini-slots, which yields a contradiction when $S \geq T$. This completes the proof.

In the example of Figure 9 one can see that we need $S \geq 3$; indeed, after sub-phase 2 there is no active link in $I(e_1) \cup \{e_1\}$. ■

4.7 Setting the parameters of AlgoLog

Let us now address the communication overhead of AlgoLog, denoted by O , defined as the number of mini-slots contained in the control phase of a slot. To obtain an admissible maximal schedule we have to choose, by Proposition 2, $S \geq T$. The best is to choose $S = T$. Therefore, there are $S (= T)$ sub-phases in the control phase plus $S - 1$ mini-slots of synchronization; furthermore each sub-phase is composed of T mini-slots. Hence,

$$O = T^2 + T - 1 \quad \text{with} \quad T = \log_2(\lceil CK + 1 \rceil). \quad (13)$$

Parameters K and C directly impact the overhead of AlgoLog since the overhead is an increasing function of the product $C \cdot K$, as shown in (13). We choose for C the smallest possible value (see next paragraph). The choice of K is more delicate as, on one hand, we want a small K to reduce the overhead but, on the other hand, we want a large K to ensure the finest possible discrimination between the ratios $q_t(e)/c(e)$, a key feature for ensuring that (non-interfering) links with the highest backlogs are scheduled first. One possible way to choose K consists in determining an admissible threshold for the overhead, which in turns implies an upper-bound K_0 on K . Then, if the ratio $q_t(e)/c(e)$ is bounded by some constant B , we can choose $K = K_0$ and $L = \frac{B(K_0-1)}{K_0}$. Finally, we divide the interval (O, L) into $K_0 - 1$ intervals of the same length.

Computation of C . Recall that computing the minimum value of C consists in determining the minimum number of colors needed to color the links in such a way that two interfering links have different colors. In general, that is an NP-complete problem (in fact it corresponds to a coloring of the vertices of the so called link-interference graph of G whose vertices represent the links of G , two vertices being joined if the corresponding links interfere). In the d -interference model, there exists a simple greedy algorithm to do this in linear time using $2\Delta(G)^{d+1}$ different colors [Die97], where $\Delta(G)$ is the maximum degree in G . In the d -interference model, if G is a *path* composed of at least $d + 2$ links, then we can use only $d + 2$ colors. In particular, if $d = 0$ and G is a path then only two colors are needed.

A particular interesting case is when the binary interference model is the primary interference model (i.e. $d = 0$). In this case, the minimum value of C is the *edge chromatic number* of G (see [FW77]). Vizing's theorem [FW77] establishes that for any graph G , this number is either $\Delta(G)$ or $\Delta(G) + 1$. A constructive proof of Vizing's theorem can be found in [MG92]. There also exists a polynomial time algorithm to compute a valid coloring of the links by using at most $\Delta(G) + 1$ colors, therefore giving a $+1$ approximation. Finally, if G is bipartite then the edge chromatic number of G is exactly $\Delta(G)$; therefore when $d = 0$, we have $C = 2$ for any path and $C = 4$ for any grid.

5 Stability of AlgoLog

In this section, we address the stability (defined below) of AlgoLog when each link is equipped with an infinite capacity buffer. This is done both for single-hop and multi-hop communications.

5.1 Single-hop communications

To make the model as general as possible, we assume that arrivals to links are modulated by a discrete-time, aperiodic, irreducible, and homogeneous Markov chain $\mathbf{Z} := \{z_t, t \geq 1\}$, taking values in a countable and finite set R . Let $\pi = (\pi(z), z \in R)$ be the stationary distribution of \mathbf{Z} , and let $\mathbf{P} = [P(z, z')]_{z, z' \in R}$, with $P(z, z') := P(z_{t+1} = z' | z_t = z)$, be its probability transition matrix.

When $z_t = z$, the arrivals to links in slot t are given by the stochastic sequence $\{A_t(e, z), e \in E\}$, with $A_t(e, z) \in \{0, 1, \dots\}$ the arrivals to link e .

Let $q_t(e)$ be the number of pending messages in the buffer (of link) e at the beginning of slot $t \geq 1$. We denote by $x_t(e)$ the number of messages transmitted on link e in slot t , with $x_t(e) = 0$ if link e is inactive in slot t and $x_t(e) = \min(q_t(e), c(e))$ if link e is active in slot t (note that in the latter case $q_t(e) > 0$ as otherwise link e cannot be active). The sequence $\{q_t(e), t \geq 1\}$ satisfies the following recursion: for each $t \geq 1$ and each $e \in E$,

$$q_{t+1}(e) = q_t(e) + A_t(e, z) - x_t(e) \quad (14)$$

if $z_t = z$. Notice that the definition of $x_t(e)$ implies that new messages joining buffer e in slot t are not transmitted in this slot. We recall that the decision to activate or not a link in slot t is made by `AlgoLog`, and that this decision is a deterministic function of the virtual weight $(w_t(e), e \in E)$ at the beginning of slot t , which, by (9), only depends on the queue-lengths vector $\mathbf{q}_t := (q_t(e), e \in E)$ and on the coloring vector $\gamma_t = (\gamma_t(e), e \in E)$.

We assume that:

- (a) for every $e \in E$ and $z \in R$, $\{A_t(e, z), t \geq 1\}$ is a renewal sequence, i.e. the random variables $A_1(e, z), A_2(e, z), \dots$ are independent and identically distributed (iid). Define

$$a(e, z) = E[A_t(e, z)], \text{ for all } e \in E, z \in R.$$

Moreover, for each e , the $|R|$ renewal sequences $\{A_t(e, z), t \geq 1\}$, $z \in R$, are mutually independent. Last, we assume that the sequences $\{A_t(e, z), e \in E, z \in R\}$ and $\{A_{t'}(e, z), e \in E, z \in R\}$ are independent for $t \neq t'$;

- (b) for each $t \geq 1$,

$$m_0 := \max_{e \in E, z \in R} a(e, z) < \infty \quad (15)$$

and

$$\sigma_0 := \max_{e \in E, f \in I(e) \cup \{e\}, z \in R} E[A_t(e, z)A_t(f, z)] < \infty. \quad (16)$$

Observe that (a) does not rule out correlated arrivals at links in the same slot.

To show the versatility of the arrival processes defined above, let us specialize the Markov chain \mathbf{Z} and assumption (a) above as follows. Assume that $R = \times_{e \in E} R(e)$, and that $\mathbf{Z} = \{\mathbf{Z}(e), e \in E\}$ with $\mathbf{Z}(e) = \{z_t(e), t \geq 1\}$, $z_t(e) \in R(e)$, are $|E|$ mutually independent, aperiodic irreducible, homogeneous, and discrete-time Markov chains. Assume further that $A_t(z) = A_t(e, z(e))$ if $z_t = z$ with $z = (z(e), e \in E)$, and that $\{A_t(e, k), t \geq 1, k \in R(e)\}$ are $|E|$ mutually independent renewal sequences. Then, for each $e \in E$, $\{A_t(e), t \geq 1\}$ is the so-called Discrete-time Batch Markov Arrival Process (D-BMAP) and these D-BMAPs are mutually independent. If in addition we assume that, for each $e \in E$, the set $R(e)$ is a singleton, then these arrival processes are all independent renewal sequences.

Under assumption (a), the process $\mathbf{Y} := \{\mathbf{y}_t := (\mathbf{q}_t, z_t), t \geq 2\}$ is a discrete-time, homogeneous, Markov chain on $\mathbb{N}^{|E|} \times R$. We further assume that (see Remark 1):

- (c) the Markov chain \mathbf{Y} is irreducible on $\mathbb{N}^{|E|} \times R$. This will hold, in particular, if $P(A_t(e, z) = k) > 0$ for all $k \geq 0$, $e \in E, z \in R$.

`AlgoLog` is stable if the Markov chain \mathbf{Y} is positive recurrent [MT93]. In words, queues do not “build up” when `AlgoLog` is stable. Observe that the system is necessarily stable when all buffers have finite capacity since in this case the state space is finite.

Let $a(e)$ be the expected number of arrivals in link e in a slot. We have

$$a(e) = \sum_{z \in R} a(e, z)\pi(z).$$

The following lemma will be used in the proof of Proposition 3 and, more precisely, in the definition of the Lyapounov function used to investigate the stability of the queue-lengths process $\{\mathbf{q}_t, t \geq 1\}$.

For any \mathbf{v} row vector we denote by \mathbf{v}^{Tr} its transpose.

Lemma 3 For each $e \in E$, let $\mathbf{u}(e) = (u(e, z), z \in R)^{Tr}$ be the unique solution of the equation

$$(\mathbf{I} - \mathbf{P})\mathbf{u}(e) = a(e)\mathbf{1} - \mathbf{D}(e)\mathbf{1},$$

such that $\sum_{z \in R} u(e, z) = 1$, with $\mathbf{1} = (1, \dots, 1)^{Tr}$ the unit vector of dimension $|R|$, and $\mathbf{D}(e)$ the $|R| \times |R|$ diagonal matrix with diagonal elements $(a(e, z), z \in R)$. Equivalently, $u(e, z)$ satisfies

$$u(e, z) = \sum_{z' \in R} P(z, z')u(e, z') + a(e) - a(e, z), \quad (17)$$

for all $e \in E, z \in R$.

Proof. Fix $e \in E$. Since \mathbf{P} is a stochastic matrix, the matrix $\mathbf{I} - \mathbf{P}$ has rank $|R| - 1$ and therefore the system of linear equations $(\mathbf{I} - \mathbf{P})\mathbf{u}(e) = a(e)\mathbf{e} - \mathbf{D}(e)\mathbf{e}$, $\mathbf{u}(e) \cdot \mathbf{1} = 1$, has a unique solution. \blacksquare

The following sufficient stability condition holds:

Proposition 3 (Sufficient stability condition for single-hop communications)

Under assumptions (a)-(d) above, the arrival rate vector $(a(e), e \in E)$ stabilizes *AlgoLog* if

$$\sum_{f \in I(e) \cup \{e\}} \frac{a(f)}{c(f)} < 1, \quad \forall e \in E, \quad (18)$$

when $c(e) = 1$ for all $e \in E$, and if

$$\sum_{f \in I(e) \cup \{e\}} \frac{a(f)}{c(f)} < \min(1, L), \quad \forall e \in E, \quad (19)$$

when $c(e) > 1$ for at least one $e \in E$.

Recall that $L \geq 0$ is a free parameter of *AlgoLog*- see Section 4.5. We see from (19) that selecting $L \geq 1$ will improve the stability condition when at least one $c(e)$ is equal to 1.

Proof of Proposition 3. Consider the Markov chain \mathbf{Y} with the Lyapounov function

$$V(\mathbf{y}) := \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q(f)}{c(f)} - 2 \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{u(f, z)}{c(f)}, \quad (20)$$

with $\mathbf{y} = (\mathbf{q}, z) \in \mathbb{N}^{|E|} \times R$ where $\mathbf{q} = (q(e) \in E)$ and the mappings $\{u(f, z), f \in E, z \in R\}$ are defined in Lemma 3. Notice that the first double summation in the right hand side of (20) is the Lyapounov function considered in [WSP06, Section II-A].

We will show that Foster's criterion [MT93, Fos53] applies to $V(\mathbf{y})$, i.e. that there exist a finite set $\mathcal{S} \subset \mathbb{N}^{|E|} \times R$ and a constant $\varepsilon > 0$ such that $\Delta(\mathbf{y}) := E[V(\mathbf{y}_{t+1}) - V(\mathbf{y}_t) | \mathbf{y}_t = \mathbf{y}] < -\varepsilon$ for all $\mathbf{y} \notin \mathcal{S}$ and $\sup_{\mathbf{y} \in \mathcal{S}} \Delta(\mathbf{y}) < \infty$. Observe that Foster's criterion applies to $V(\mathbf{y})$ as this function is bounded from below.

With $\mathbf{y}_{t+1} = (q_{t+1}, z_{t+1})$ and $\mathbf{y}_t = (q, z)$, we have

$$\begin{aligned} \Delta(\mathbf{y}) &= E \left[\sum_{e \in E} \frac{q_{t+1}(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q_{t+1}(f)}{c(f)} - \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q(f)}{c(f)} \mid \mathbf{y}_t = (\mathbf{q}, z) \right] \\ &\quad - 2E \left[\sum_{e \in E} \frac{q_{t+1}(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{u(f, z_{t+1})}{c(f)} - \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{u(f, z)}{c(f)} \mid \mathbf{y}_t = (\mathbf{q}, z) \right] \\ &= E \left[\sum_{e \in E} \frac{q_{t+1}(e) - q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q_{t+1}(f) - q(f)}{c(f)} - 2 \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q(f)}{c(f)} \right. \\ &\quad \left. + \sum_{e \in E} \frac{q_{t+1}(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q(f)}{c(f)} + \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q_{t+1}(f)}{c(f)} \mid \mathbf{y}_t = (\mathbf{q}, z) \right] \\ &\quad - 2E \left[\sum_{e \in E} \frac{q_{t+1}(e) - q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{u(f, z_{t+1})}{c(f)} + \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{u(f, z_{t+1}) - u(f, z)}{c(f)} \mid \mathbf{y}_t = (\mathbf{q}, z) \right]. \end{aligned}$$

From the set equality $\{(e, f) \in E^2, f \in I(e) \cup \{e\}\} = \{(e, f) \in E^2, e \in I(f) \cup \{f\}\}$ which holds from the symmetric conflict assumption, we see that

$$\sum_{e \in E} \frac{q_{t+1}(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q(f)}{c(f)} = \sum_{e \in E} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \frac{q_{t+1}(f)}{c(f)},$$

which yields

$$\Delta(\mathbf{y}) = \Delta_1(\mathbf{y}) + \Delta_2(\mathbf{y})$$

with

$$\begin{aligned} \Delta_1(\mathbf{y}) &:= E \left[\sum_{e \in E} \frac{q_{t+1}(e) - q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \left(\frac{q_{t+1}(f) - q(f)}{c(f)} - 2 \frac{u(f, z_{t+1})}{c(f)} \right) \mid \mathbf{y}_t = (\mathbf{q}, z) \right] \\ &\quad + 2E \left[\sum_{e \in E: q(e)/c(e) \leq L} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \left(\frac{q_{t+1}(f) - q(f)}{c(f)} - \frac{u(f, z_{t+1}) - u(f, z)}{c(f)} \right) \mid \mathbf{y}_t = (\mathbf{q}, z) \right] \\ \Delta_2(\mathbf{y}) &:= 2E \left[\sum_{e \in E: q(e)/c(e) > L} \frac{q(e)}{c(e)} \sum_{f \in I(e) \cup \{e\}} \left(\frac{q_{t+1}(f) - q(f)}{c(f)} - \frac{u(f, z_{t+1}) - u(f, z)}{c(f)} \right) \mid \mathbf{y}_t = (\mathbf{q}, z) \right]. \end{aligned}$$

We first show that there exists a constant b_0 such that $\Delta_1(\mathbf{y}) < b_0$. By using (14) (recall that $A_t(e) = A_t(e, z)$ when $z_t = z$, as assumed in (a) above) and $0 \leq x_t(e) \leq c(e)$, we find

$$\begin{aligned} \Delta_1(\mathbf{y}) &\leq \sum_{e \in E} \sum_{f \in I(e) \cup \{e\}} \left(\frac{E[A_t(e, z)A_t(f, z)]}{c(e)c(f)} + \frac{a(e, z)}{c(e)} + \frac{a(f, z)}{c(f)} + 1 \right) \\ &\quad + 2 \max_{f \in E, z \in \mathcal{R}} |u(f, z)| \sum_{e \in E} \left(\frac{a(e, z)}{c(e)} + 1 \right) \sum_{f \in E} \frac{1}{c(f)} \\ &\quad + 2L \sum_{e \in E} \left(\sum_{f \in E} \left(\frac{a(f, z)}{c(f)} + 1 \right) + 2 \max_{f \in E, z \in \mathcal{R}} |u(f, z)| \sum_{f \in E} \frac{1}{c(f)} \right) \\ &\leq |E|^2 \left(\frac{\sigma_0}{c_0^2} + 2m_0c_0 + 1 + 2L \left(\frac{m_0}{c_0} + 1 \right) \right) + 4 \frac{|E|^2}{c_0} \left(\frac{m_0}{c_0} + 1 + L \right) \max_{f \in E, z \in \mathcal{R}} |u(f, z)| := b_0, \quad (21) \end{aligned}$$

where we have set $c_0 = \min_{e \in E} c(e) > 0$. The bound b_0 in (21) is finite as $|E|, L, m_0$ and σ_0 are finite quantities (see (15) and (16)), and $\max_{f \in E, z \in \mathcal{R}} |u(f, z)| < \infty$ thanks to the constraints $\sum_{z \in \mathcal{R}} u(e, z) = 1$ for all $e \in E$ (see Lemma 3).

Let us now focus on $\Delta_2(\mathbf{y})$. With $E[u(f, z_{t+1}) \mid \mathbf{y}_t = (\mathbf{q}, z)] = \sum_{z' \in \mathcal{R}} P(z, z') u(f, z')$ and (14), $\Delta_2(\mathbf{y})$ can be written as

$$\begin{aligned} \Delta_2(\mathbf{y}) &= 2 \sum_{e \in E: q(e)/c(e) > L} \frac{q(e)}{c(e)} \left(\sum_{f \in I(e) \cup \{e\}} \frac{1}{c(f)} \left(a(f, z) - \sum_{z' \in \mathcal{R}} P(z, z') u(f, z') + u(f, z) \right) \right. \\ &\quad \left. - \sum_{f \in I(e) \cup \{e\}} \frac{E[x_t(f) \mid \mathbf{q}_t = \mathbf{q}]}{c(f)} \right) \\ &= 2 \sum_{e \in E: q(e)/c(e) > L} \frac{q(e)}{c(e)} \left(\sum_{f \in I(e) \cup \{e\}} \frac{a(f)}{c(f)} - \sum_{f \in I(e) \cup \{e\}} \frac{E[x_t(f) \mid \mathbf{q}_t = \mathbf{q}]}{c(f)} \right), \quad (22) \end{aligned}$$

by using (17).

Define $\alpha = 1$ if $c(e) = 1$ for all $e \in E$ and $\alpha = \min(1, L)$ if $c(e) < 1$ for at least one $e \in E$. By (18)-(19) there exists $\varepsilon' > 0$ such that $\sum_{f \in I(e) \cup \{e\}} a(f)/c(f) < \alpha - \varepsilon'$. Therefore, by (22),

$$\Delta_2(\mathbf{y}) \leq 2 \sum_{e \in E: q(e)/c(e) > L} \frac{q(e)}{c(e)} \left(\alpha - \varepsilon' - \sum_{f \in I(e) \cup \{e\}} \frac{E[x_t(f) \mid \mathbf{q}_t = \mathbf{q}]}{c(f)} \right). \quad (23)$$

Let us show that when $q(e)/c(e) > L$

$$\sum_{f \in I(e) \cup \{e\}} \frac{E[x_t(f) \mid \mathbf{q}_t = \mathbf{q}]}{c(f)} \geq \alpha. \quad (24)$$

Assume first that $c(e) = 1$ for all $e \in E$. This implies that $q(e) \geq 1$ (recall $q(e)$ is an integer). Proposition 2 then says that at least one link in $I(e) \cup \{e\}$ is active, namely, there exists $f \in I(e) \cup \{e\}$ such that $q(f) \geq 1$ and

$x_t(f) = \min(q(f), c(f)) = 1$, which proves (24) with $\alpha = 1$.

Assume now that $c(f) > 1$ for some $f \in E$. Because $q(e)/c(e) > L$, we may again invoke Proposition 2 to conclude that at least one link in $I(e) \cup \{e\}$ is active. If e is active then $x_t(e) = \min(c(e), q(e)) \geq \min(c(e), c(e)L) = c(e) \min(1, L) = \alpha c(e)$ which proves (24). If e is not active then there exists an active link in $I(e)$, say f , with virtual weight, $w_t(f)$, larger than the virtual weight of e . If $w_t(f) > w_t(e)$, then $q(f)/c(f) \in I_{K-1}$ by the definition of the virtual weights or, equivalently, $q(f)/c(f) > L$ and again (24) holds.

In summary, cf. (23)-(24),

$$\Delta(\mathbf{y}) \leq -2\varepsilon' \sum_{e \in E: q(e)/c(e) > L} \frac{q(e)}{c(e)} + b_0 \leq -2\varepsilon' \sum_{e \in E} \frac{q(e)}{c(e)} + 2\varepsilon' L|E| + b_0. \quad (25)$$

Take $\varepsilon > 0$ and let M be any constant such that $M > \frac{2\varepsilon' L|E| + b_0 + \varepsilon}{2\varepsilon'}$. Let $\mathcal{S} = \{\mathbf{q} : \sum_{e \in E} q(e)/c(e) \leq M\} \times \mathcal{R}$; we conclude from (25) that $\Delta(\mathbf{y}) \leq -\varepsilon$, for all $\mathbf{y} \notin \mathcal{S}$ so that Foster's criterion applies with the finite set \mathcal{S} , which completes the proof. \blacksquare

Remark 1 Due to the generality of the network topology, the interference sets, the statistical assumptions, and the behavior of `AlgoLog` it may be difficult to guarantee the irreducibility of the Markov chain \mathbf{Y} . If so, we can adopt the definition of stability in [TE92]. Theorem 3.1 in [TE92] ensures that Proposition 3 holds for this alternative definition.

5.2 Multi-hop communications

We will refer to a result obtained in [WSP06]. There are M distinct flows, defined by their source and destination links. All messages belonging to a flow follow the same (deterministic) route in the network. Let $B_t(s_m, d_m)$ be the number of messages of flow m entering the network in slot t with s_m (resp. d_m) the source (resp. destination) link. We assume that $\{B_t(s_1, d_1), t = 1, \dots\}, \dots, \{B_t(s_m, d_m), t = 1, \dots\}$ are mutually iid sequences. Let $b_m = E[B_t(s_m, d_m)]$ and assume that $E[B_t(s_m, d_m)^2] < \infty$ for all m .

To each flow visiting a link is associated a so-called λ -regulator, whose role is to regulate the output of this flow from the link. If the number of messages of a flow, say flow m , at link e exceeds its capacity $c(e)$, then with probability $\lambda_m/c(e)$ (λ_m is defined below) $c(e)$ messages are injected into the queue dedicated to flow m , and no messages are injected otherwise. If link e is the k -th hop of flow m then $\lambda_m = b_m + (k-1)\varepsilon$, with $\varepsilon > 0$ small enough. The following Theorem 2 in [WSP06] applies to `AlgoLog`: the network is stable for any rate vector (b_1, \dots, b_m) satisfying

$$\sum_{f \in I(e) \cup \{e\}} \frac{b_m I(f, m)}{c(f)} < 1, \quad (26)$$

for every $m = 1, \dots, M$ and $e \in E$, where $I(f, m) = 1$ if flow m visits link f and $I(f, m) = 0$ otherwise.

This result can be extended to the more general arrival processes considered in Section 5.1.

In this context, `AlgoLog` can be implemented by letting $q_t(e) = \sum_{m=1}^M q_t(e, m) I(f, m)$, with $q_t(e, m)$ the number of messages of flow m (if any) ready to be transmitted in slot t on link e , where a message is ready for transmission if it has left the regulator. Assume without loss of generality that $q_1(e, m) = 0$ for every $m = 1, \dots, M$ so that either $q_t(e, m) = 0$ or $q_t(e, m) \geq c(e)$ at the beginning of slot t , thanks to the regulator. If link e is activated in slot t then transmit $c(e)$ messages from flow $m' = \arg \max_{1 \leq m \leq M} q_t(e, m)$ in that slot, for every $e \in E$.

6 Numerical results

In this section, we investigate via simulations the performance of `AlgoLog`. More precisely, we study the time evolution of the queue lengths over a certain time-window (number of slots). The simulations have been carried out for path networks (Section 6.1), square grid networks (Section 6.2), and random networks (Section 6.3). For

each topology we investigate both the evolution of the largest queue length and of the average queue length over tens of thousands of slots. In Section 6.1, for path networks and $d = 0$, we compare the total sum of the weights of the links of the matching found by `AlgoLog` in a given slot to the total sum of the weights of an optimal matching. In Section 6.2, we compare the performance of `AlgoLog` to that of the `Augmenting Paths Algorithm` for a grid topology.

6.1 Path network

In this section, we consider a path network and the primary node interference model ($d = 0$). We first study the ratio between the total sum of the weights (or queue lengths) of the links present in the matching computed by `AlgoLog` and the total sum of the weights of the links present in an optimal matching for centralized wireline network). Figure 10 displays this ratio for a path $G = (V, E)$ composed of $|E| = 50$ links for 1 000 tests. At each test, we assign a weight at every $e \in E$ as follows: $P(q_t(e) = i) = 1/51$ for $i = 0, 1, \dots, 50$. Observe that, for each test, `AlgoLog` finds a matching yielding a ratio larger than 0.80 and, in most experiments, this ratio is larger than 0.95.

We now investigate the evolution of the queue length at each buffer. We consider a path composed of $|E| = 100$ links that we observe during 100 000 consecutive time slots. Since $d = 0$, we can take $C = 2$ (see discussion at the end of Section 4.7). With the initial coloring $(\gamma_1(1), \gamma_1(2), \gamma_1(3), \dots, \gamma_1(100)) = (1, 2, 1, \dots, 2)$, we have from (7) that $(\gamma_t(1), \gamma_t(2), \gamma_t(3), \dots, \gamma_t(100))$ is equal to $(1, 2, 1, \dots, 2)$ (resp. $(2, 1, 2, \dots, 1)$) if t is an odd (resp. even) integer. We take $K = 1000$, $L = 999$, and the capacity of every link $e \in E$ is $c(e) = 18$. The arrival process is a renewal process defined as follows: with probability $p(e)$, $c(e)$ messages join link e in a slot and with probability $1 - p(e)$ no message joins link e , so that $a(e) = p(e)c(e)$. We assume that $a(e) = a_1$ (resp. $a(e) = a_2$) for $e = 1, 3, 5, \dots, 99$ (resp. for $e = 2, 4, 6, \dots, 100$). Figures 12(a), 12(b), and 12(c) display the largest queue length in each slot for pairs $(a_1, a_2) = (16, 1)$, $(a_1, a_2) = (12, 4)$, and $(a_1, a_2) = (8, 8)$, respectively.

In these figures, we observe that the size of the largest queue is always smaller than 400, 180, and 140, respectively. It is also worth noting that the network is stable in Figures 12(a)-12(c) in the sense that the largest queue length does not increase as a function of t , despite the fact that the (sufficient) stability condition (19) in Proposition 3 is not satisfied for all $e \in E$ (it is only satisfied for links 1 and 100).

We finally investigate the impact of the value of K on the largest and average size of the queues for a path composed of 100 links, with $d = 0$, $C = 2$, $c(e) = 1$ for all $e \in E$. The arrival process is similar to that in Figures 12(a)-12(c) with $a(e) = 0.75$ (i.e. $p(e) = 0.75$) for $e = 1, 3, 5, \dots, 99$ and $a(e) = 0.2$ (i.e. $p(e) = 0.2$) for $e = 2, 4, 6, \dots, 100$. Figure 11 displays the largest and the average queue-length after 50 000 slots as a function of K for $K = 2^i$ for $i = 0, 1, \dots, 7$ and $L = K - 1$. The plots illustrate the performance-overhead tradeoff of K , a higher K yielding a lower largest and average queue lengths; however, most of the gains come with relatively small values of K . Here too, the system appears to be stable although condition (18) in Proposition 3 is not satisfied as $\sum_{f \in I(e) \cup \{e\}} \frac{a(f)}{c(f)} = 1.15$ when e is an odd number and $\sum_{f \in I(e) \cup \{e\}} \frac{a(f)}{c(f)} = 1.7$ when e an even number.

6.2 Square grid network

In this section, we compare the performance of `AlgoLog` to that of the `Augmenting Paths Algorithm` proposed in [BSS09]. The simulations have been done for a square grid $G = (V, E)$ composed of $|V| = 121$ nodes and $|E| = 220$ links (see Figure 13). For every link $e \in E$, the capacity of e is $c(e) = 1$.

Figure 14 displays the evolution of the largest and average queue-lengths over 100 000 time-slots under `AlgoLog` for $d = 1$, with $C = 4$, $K = 1 000$, and $L = 999$. The arrival rates $(a(e), e \in E)$ are reported in Figure 13 (the arrival process is the same as in Section 6.1, namely, with probability $p(e)$, 1 message joins link e and with probability $1 - p(e)$ no message joins link e in a slot; here $p(e) = 0.7$ or $p(e) = 0.2$ as indicated in Figure 13). The numerical results indicate that the network appears to be stable under these arrival rates even if conditions (18) are all violated when $d = 1$.

Figures 15(a)-15(c) display the evolution of the largest and average queue-lengths over 50 000 – 200 000 time-slots under `AlgoLog` for $C = 4$, $K = 1 000$, and $L = 999$, and under the `Augmenting Paths Algorithm`. As described in Section 3, the `Augmenting Paths Algorithm` has two input parameters: k and p . Like in [BSS09],

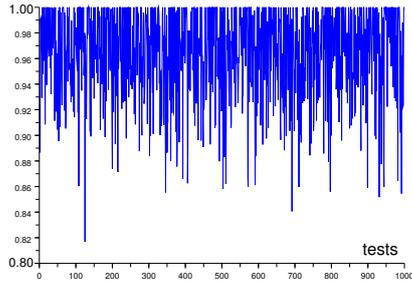


Fig. 10 Ratio between the sum of the weights of the links of the matching found by `AlgoLog` and the sum of the weights of the links of an optimal matching for a path composed of 50 links for 1 000 tests, $d = 0$, and $P(q_t(e) = i) = 1/51$ for $0 \leq i \leq 50$, $e \in E$.

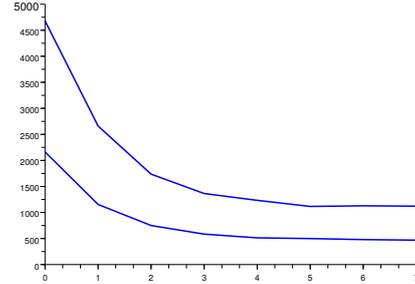


Fig. 11 Largest and average queue length as function of $K = 2^i$ for $i = 0, 1, \dots, 7$ after 10 000 slots for a path composed of 100 links for $d = 0$, $C = 2$, $L = K - 1$, $c(e) = 1$, $a(e) = 0.75$ for odd links and $a(e) = 0.2$ for even links.

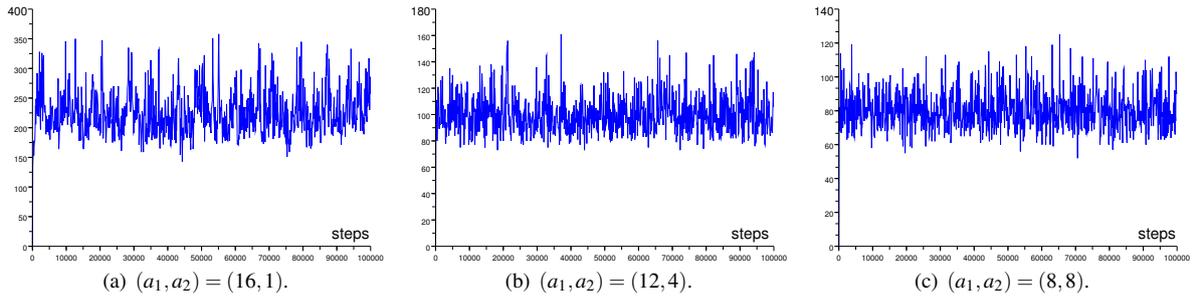


Fig. 12 Evolution of the largest queue-length for a path network composed of 100 links for `AlgoLog` during 50 000 slots for $d = 0$, $C = 2$, $K = 1 000$, $L = 999$, $c(e) = 18$, $a(e) = a_1$ for odd links and $a(e) = a_2$ for even links.

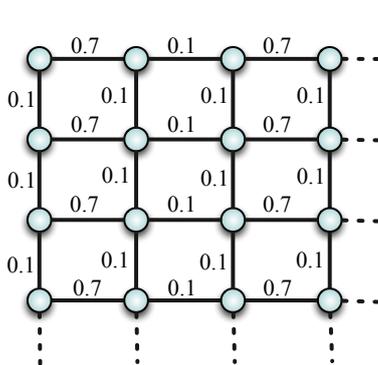


Fig. 13 Grid network. Values near the links give $a(e)$ for all $e \in E$.

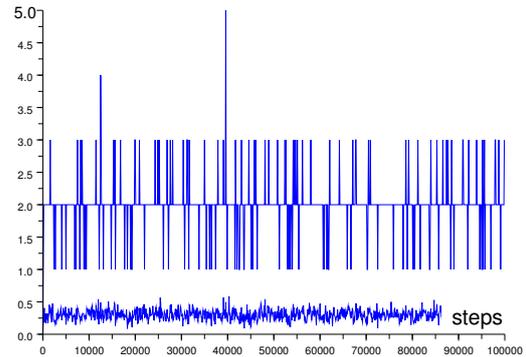


Fig. 14 Evolution of the largest and average queue-lengths for the grid network in Fig 13 under `AlgoLog` with $d = 1$, $C = 4$, $K = 1 000$, $L = 999$.

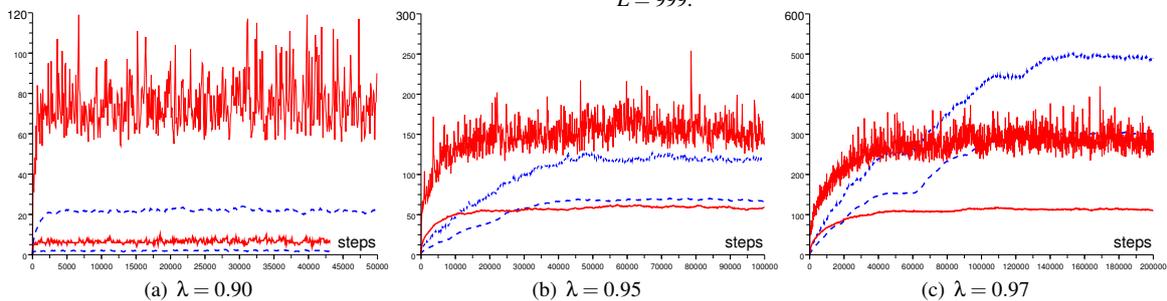


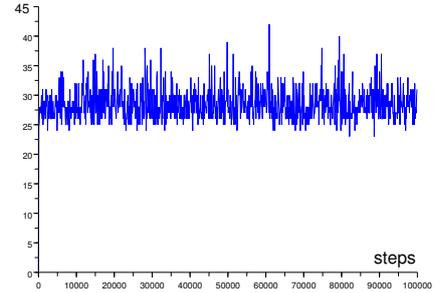
Fig. 15 Evolution of the largest and average queue-lengths for a grid network composed of 121 nodes and 200 links under either the Augmenting Paths Algorithm (with $k = 2$, $p = 0.2$) or `AlgoLog` (with $d = 0$, $C = 4$, $K = 1 000$, $L = 999$) with $(c(e) = 1, e \in E)$ and $(a(e), e \in E)$ given in Fig. 13. Arrival rates are $(\lambda a(e), e \in E)$ with $\lambda = 0.90, 0.95, 0.97$ in Fig. 15(a)-15(c), resp., where $(a(e), e \in E)$ are given in Fig. 14. Red straight lines represent Augmenting Paths Algorithm and blue dotted lines represent `AlgoLog`.

the simulations have been carried out in Figures 15(a)-15(c) for $\{k = 2, p = 0.2\}$, $\{k = 3, p = 0.2\}$, and $\{k = 3, p = 0.1\}$, respectively. In Figure 15(a)-(c), $d = 0$ as the `Augmenting Paths Algorithm` is only defined for the primary interference model. In Figure 15(a) (resp. Figure 15(b), Figure 15(c)) the arrival rates are $(\lambda a(e), e \in E)$ with $\lambda = 0.90$ (resp. $\lambda = 0.95$, $\lambda = 0.97$), where $(a(e), e \in E)$ are given in Figure 14. Again, it is worth noting that the arrival rates in Figure 15(a)-15(c) stabilize the grid network although condition (18) in Proposition 3 is not satisfied. We observe that the flexibility of `AlgoLog` - in terms of working for arbitrary values of d as opposed to the `Augmenting Paths Algorithm` which only works for $d = 0$ - comes at a price for high arrival rates, as can be seen in Figure 15(c) where the `Augmenting Paths Algorithm` performs better than `AlgoLog`.

6.3 Random networks

To conclude, we apply `AlgoLog` to a random network, generated as follows: given a square grid $G' = (V, E')$, we first delete link e with probability p for all $e \in E'$, and then add a link with probability q between any pair of nodes located within a given Euclidean distance. This process yields a new (random) transmission graph $G = (V, E)$. Figure 8 shows an instance of a random graph composed of $|V| = 48$ nodes and $|E| = 91$ links.

The figure on the right displays the evolution of the largest queue-length under `AlgoLog` for a random graph composed of 200 nodes and 424 links. The parameters of `AlgoLog` are $d = 0$, $C = 20$, $K = 10\,000$, $L = 9\,999$. For every $e \in E$, $c(e) = 20$. The arrivals per slot on link $e \in E$ are iid and follow a uniform distribution in $[0, \delta(e)]$, with $\delta(e) \leq 40$. The parameters $(\delta(e), e \in E)$ have been chosen so that the sufficient stability conditions (19) hold.



7 `AlgoLog`: from link scheduling to node scheduling

Until now we have focused on the link version of `AlgoLog` - referred to as the link model - following in that the models used in the literature. On a theoretical point of view, the link model is natural as ultimately arrivals and buffers are on the links. In practice, however, things are different as all decisions are made by the nodes and not by the links which are immaterial entities. There are many ways to go from a link model to a node model depending on the technology that is used, but most of the node models use symmetric directed transmission graph as there is a need for some acknowledgment at least in the transmission phase.

In the control phase of `AlgoLog`, we have used the possibility for a link e to send a control signal and the fact that this signal is heard by all links interfering with e , namely, by all links in $I(e)$ (cf. line 6 of the pseudo-code of `AlgoLog` in Table 2) and decisions were taken in each sub-phase according to that. In a node model, now sending or hearing should be done in the nodes. We can design some new algorithms adapted to this model, but we will have to redo all the analysis. We prefer to indicate how to emulate the link model, so that we can reuse all the preceding results. To this end, we will emulate a sub-phase of `AlgoLog` (cf. lines 5-10 in Table 2). However, this has to be done in a distributed manner and with the smallest possible overhead. To emulate the fact that link (u, v) send a signal when $s(e) = U$ and $v_{t,e}(i) = 1$, we will let node u send a signal to node v . But now we have to be able to differentiate among all arcs with origin u those which are sending.

In summary, emulation of the link model will be possible if a node v can determine in some manner if it is concerned or not by the signal emitted by node u . This determination will depend on the physical model. It can always be done by using the initial coloring (Assumption 1); indeed it suffices to replace a mini-slot by C mini-slots. If the arc (u, v) wants to send a signal, then u sends a signal in the mini-slots corresponding to the color of the arc (u, v) . But the overhead can be significant. However there are also cases where we are able to differentiate the direction of the signal; for example, if we use a model with unidirectional antennas where a signal is sent only in a specific direction. We show now how, with this hypothesis, we can emulate in the primary node model a mini-slot i of sending on the arc $e = (u, v)$ and hearing only with two mini-slots (that is how to emulate the behavior of lines 5-10 of `AlgoLog`). Therefore the overhead will be only multiplied by 2. The node version of `AlgoLog` is denoted

AlgoLogNodes.

In a nutshell, in the first mini-slot, if arc $e = (u, v)$ is sending a signal in AlgoLog, then u sends a signal to v . At the end of this mini slot any node u will be able to know if there is an arc f incident to u implied in a signal transmission; then u will take decisions on the state of some arcs according to the answer (see Cases (a) and (b) below). In the second mini-slot a node v informs (by sending a signal) all its neighbors except when it is not implied in a sending and has heard a signal from an unique node say u , in which case it sends a signal to all its neighbors except u . At the end of this mini-slot any node u will be able to know for each of its neighbors if there is an arc, f incident to it, implied in a transmission; then u will take decisions on the state of the arcs incident to it according to the answer (see Cases (c) and (d) after).

In the following, we describe precisely the two mini-slots of AlgoLogNodes (this corresponds to one mini-slot of AlgoLog lines 5-6) and the decisions taken by the nodes in order to update the states of the arcs (lines 7-10 of AlgoLog). We illustrate AlgoLogNodes by describing its execution for the oriented grid network depicted in Figure 16.

We recall that, $s(f)$, the state of arc f , takes value in $\{U, PI, I, A\}$, where U (resp. PI, I, A) stands for undetermined (resp. potentially inactive, inactive, active).

In AlgoLogNodes, each mini-slot $i = 1, \dots, T$ of AlgoLog (lines 5-6) is replaced by two mini-slots called mini-slot i_1 and mini-slot i_2 with a slight abuse of notation and defined as follows.

- mini-slot i_1 : If $s(e = (u, v)) = U$ and $v_{t,e}(i) = 1$, then u sends a signal in the direction of node v (signal only heard by v).
- mini-slot i_2 :
 - Case 1: if in mini-slot i_1 node v has sent a signal, then v sends a signal (busy) to all of its neighbors;
 - Case 2: if in mini-slot i_1 node v has not sent a signal, but has heard a signal from at least two of its neighbors, then v sends a signal (busy) to all of its neighbors;
 - Case 3: if in mini-slot i_1 node v has not sent any signal and has heard a signal from exactly one of its neighbors (say node u), then v sends a signal (busy) to all of its neighbors, except to node u .

Let us describe the two mini-slots for the oriented grid depicted in Figure 16. Here we suppose that at the beginning of mini-slot i , all the arcs are undetermined and five of them satisfy $v_{t,e}(i) = 1$ namely (u_1, u_2) , (u_4, u_7) , (u_5, u_4) , (u_8, u_7) , and (u_8, u_9) . In mini-slot i_1 (Figure 16(b)), u_1 sends a signal to u_2 (repectively u_4 to u_7 , u_5 to u_4 and u_8 both to u_7 and u_9). In mini-slot i_2 (Figure 16(c)), according to Case (1), u_1, u_4, u_5, u_8 being implied in one or 2 transmissions send a signal (busy) to all their neighbors; according to Case (2), u_7 sends a signal (busy) to all its neighbors; according to Case (3), u_2 (resp u_9) sends a signal to its neighbors except u_1 (resp u_8); u_3 and u_6 do nothing. The three cases correspond to arcs labeled busy $i \in \{1, 2, 3\}$ in Figure 16 (c).

According to what happens in mini-slots i_1 and i_2 , the nodes take decisions which exactly emulate the decisions taken in the link model (lines 7-10). There decisions are the following:

- Case a: if in mini-slot i_1 node u hears a signal, then:
 - it changes the state of the arcs f with origin u such that $s(f) = U$ and $v_t(f, i) = 0$ from U to PI .
 - it leaves the state of the arcs f with origin u such that $s(f) = U$ and $v_t(f, i) = 1$ at U ;
- Case b: if in mini-slot i_1 node u has not heard a signal from any of its neighbors, but has sent signals to at least two of its neighbors then:
 - it changes the state of the arcs f with origin u such that $s(f) = U$ and $v_t(f, i) = 0$ from U to PI .
 - it leaves the state of the arcs f with origin u such that $s(f) = U$ and $v_t(f, i) = 1$ at U ;
- Case c: if in mini-slot i_1 node u has not heard a signal from any of its neighbors and has sent a signal only to node v , then:
 - u changes the state of all the arcs f different than (u, v) with origin u which necessarily satisfy $s(f) = U$ and $v_t(f, i) = 0$ from U to PI ;
 - if in mini-slot i_2 node u does not hear a signal from node v , then it changes the state of arc $e = (u, v)$ from U to A ;
 - if in mini-slot i_2 node u hears a signal sent by node v it leaves the state of arc $e = (u, v)$ at U .
- Case d: if in mini-slot i_1 node u has not sent a signal and has not heard a signal from any of its neighbors, then

- if in mini-slot i_2 node u hears a signal sent by node w and if arc $f = (u, w)$ satisfies $s(f) = U$ (necessarily $v_t(f, i) = 0$), then node u changes the state of arc $f = (u, w)$ from U to PI ;
- if in mini-slot i_2 node u has not heard a signal from node w and if arc $f = (u, w)$ satisfies $s(f) = U$ (necessarily $v_t(f, i) = 0$), it leaves the state of arc f to U .

Proposition 4 proves the equivalence between `AlgoLog` and `AlgoLogNodes`. In the proof, we describe the decisions taken by the nodes after the two mini-slots, for the oriented grid depicted in Figure 16. The state of an arc is represented by a dash line for PI (Potentially Inactive), a solid line for U (Undetermined) and Blue bold for A (Active). The four cases correspond to arcs labeled case $x \in \{a, b, c, d\}$ in Figure 16(d).

Proposition 4 *Under the primary node interference model, decisions taken by `AlgoLogNodes` after the two mini-slots i_1 and i_2 emulate the decisions taken by `AlgoLog` after mini-slot i (lines 7-10).*

Proof. We have to check that under `AlgoLogNodes` links behave as under `AlgoLog` (see the description of `AlgoLog` in Section 4.1 or in Figure 5). Throughout the proof, we focus on mini-slot $i = 1, \dots, T$, of slot t . Recall that under `AlgoLogNodes` mini-slot i is composed of the two (sub-) mini-slots i_1 and i_2 .

First, consider an arc $e = (u, v)$ such that $s(e) = U$ and $v_{t,e}(i) = 1$. Under `AlgoLog`, e sends a signal in the mini-slot i and from there two things may happen: (i) e hears a signal from an arc e' in $I(e)$ and it remains in state U or (ii) e does not hear any signal and its state is changed to A .

When (i) occurs, we need to investigate four cases according to the origin or destination of e' (i.e. either e' has destination/origin u or v):

- if there is an arc $e' = (w', u)$ with destination u satisfying $s(e') = U$ and $v_{t,e'}(i) = 1$ then in mini-slot i_1 u hears a signal from w' and, from Case (a), u leaves arc e undetermined.
In Example 16, this is the case of u_4 which in mini-slot i_1 hears a signal from u_5 , and so leaves (u_4, u_7) in state U .
- if there is an arc e' with origin u satisfying $s(e') = U$ and $v_{t,e'}(i) = 1$; then u is implied in at least two transmissions and, from Case (b), u leaves arc e undetermined.
In Example 16, this is the case of u_8 (implied in two transmissions) which leaves arcs (u_8, u_7) and (u_8, u_9) in the state U .
- if there is an arc $e' = (v, w')$ with origin v satisfying $s(e') = U$ and $v_{t,e'}(i) = 1$ then v has sent a signal in mini-slot i_1 and, from Case (1), in mini-slot i_2 , v has sent a signal to u so that, from Case (c), u leaves e undetermined.
In Example 16, this is the case of u_5 whose neighbor u_4 has sent a signal “busy” in mini-slot i_2 , and so it leaves (u_5, u_4) in state U .
- if there is an arc $e' = (w', v)$ ($e' \neq e$, i.e. $w' \neq u$) with destination v such that $s(e') = U$ and $v_{t,e'}(i) = 1$ then, from Case (2), in mini-slot i_2 , u hears a signal from v , and so from Case (c) u leaves e undetermined.
There is no such a case in Example 16 but it would have been the case of u_4 which in mini-slot i_2 hears a signal from u_7 (and so will have left (u_4, u_7) in state U) if u_5 did not sent to u_4 in mini-slot i_1 . It would also have been the case of u_8 which leaves (u_8, u_7) in state U if u_8 has not sent to u_9 in mini-slot i_1 .

When (ii) occurs, we are in Case (c) above and u changes the state of e from U to A . In the example, it is the case of u_1 which has sent only to u_2 in mini-slot i_1 and has heard nothing in mini-slot i_2 (Case (3)). So the state of the arc (u_1, u_2) is changed to A .

Second, consider an arc $f = (u, w)$ such that $s(f) = U$ and $v_{t,f}(i) = 0$. In `AlgoLog`, f does not send any signal in the mini-slot i and from there two things may happen: (i') f hears a signal from an arc f' in $I(f)$ and it changes its state from U to PI or (ii') f does not hear any signal and its state remains in U .

When (i') occurs, we need to investigate four cases according to the origin or destination of f' (i.e. either f' has destination/origin u or w):

- If there is an arc $f' = (w', u)$ with destination u satisfying $s(f') = U$ and $v_{t,f'}(i) = 1$, then u hears in mini-slot i_1 a signal from w' and so by Case (a) u changes the state of f from U to PI .

In Example 16, this is the case of all the arcs f with origin u_2, u_4, u_7 and u_9 satisfying $s(f) = U$, and $v_{t,f}(i) = 0$; so, the states of $(u_2, u_1), (u_2, u_3), (u_2, u_5), (u_4, u_1), (u_4, u_5), (u_7, u_4), (u_7, u_8), (u_9, u_6)$, and (u_9, u_8) are changed to PI .

- If there is an arc f' with origin u satisfying $s(f') = U$, and $v_{t,f'}(i) = 1$; u is implied in at least one transmission and so by Case (b) or (c) u changes the state of f from U to PI .

In Example 16, this is the case of all the arcs f with origin u_1, u_4, u_5 , and u_8 satisfying $s(f) = U$, and $v_{t,f}(i) = 0$ and not already ruled out; so, $(u_1, u_4), (u_5, u_2), (u_5, u_6), (u_5, u_8)$ and (u_8, u_5) are changed to PI .

- If there is an arc $f' = (w, w')$ with origin w satisfying $s(f') = U$ and $v_{t,f'}(i) = 1$, then, by Case (1), w has sent in mini slot i_2 a signal to u and so, from Case (d), u changes the state of f from U to PI .

In Example 16, this is the case of all the arcs f with destination u_1, u_4, u_5 and u_8 satisfying $s(f) = U$, and $v_{t,f}(i) = 0$ and not already ruled out; so, the state of (u_6, u_5) is changed to PI .

- If there is an arc $f' = (w', w)$ with destination w satisfying $s(f') = U$, and $v_{t,f'}(i) = 1$, then u hears in mini-slot 2 a signal from w (by Case (2) or (3) and so by Case (d)), u changes the state of f from U to PI .

In Example 16, this is the case of all arcs f with destinations u_2, u_4, u_7 , and u_9 satisfying $s(f) = U$ and $v_{t,f}(i) = 0$ and not already ruled out; so, states of (u_3, u_2) and (u_6, u_9) are changed to PI .

When (ii') occurs, we are in Case (d) above and u leaves the state of f to U .

In Example 16, this is the case of arcs (u_3, u_6) and (u_6, u_3) whose state remains in U . ■

8 Conclusion

We have devised the first (to the best of our knowledge) distributed algorithm - referred to as `AlgoLog`- for the transmission scheduling problem in wireless networks with constant overhead and arbitrary binary interference. We have proved that the set of active links at each slot is maximal. We have proposed sufficient stability conditions and have investigated performance of `AlgoLog` via simulations. We have also proposed extensions of `AlgoLog` in the multi-hop case and shown how the algorithm can be emulated at the node level. In terms of future work, it would be interesting to establish a better stability condition and, ideally, to characterize the entire stability region of `AlgoLog`.

References

- [BB10] F. Baccelli and B. Błaszczyszyn. *Stochastic Geometry and Wireless Networks Volume II: Applications*. Now Publishers, 2010.
- [BBK⁺04] H. Balakrishnan, C.L. Barrett, V.S.A. Kumar, M.V. Marathe, and S. Thite. The distance-2 matching problem and its relationship to the mac-layer capacity of ad hoc wireless networks. *IEEE, J. Selected Areas in Communication*, 22(6):1069–1079, 2004.
- [BKK⁺10] V. Bonifaci, R. Klasing, P. Korteweg, L. Stougie, and A. Marchetti-Spaccamela. Data gathering in wireless networks. In Arie Koster and Xavier Muñoz, editors, *Graphs and Algorithms in Communication Networks, Studies in Broadband, Optical, Wireless and Ad Hoc Networks*, pages 357–377. Springer-Verlag, 2010.
- [BMMN10] J-C. Bermond, D. Mazauric, V. Misra, and P. Nain. A distributed scheduling algorithm for wireless networks with constant overhead and arbitrary binary interference. In *ACM SIGMETRICS Performance Evaluation Review*, volume 38, pages 345–346, New York, NY, USA, June 14-18 2010.
- [BSS09] L.X. Bui, S. Sanghavi, and R. Srikant. Distributed link scheduling with constant overhead. *IEEE/ACM Transactions on Networking*, 17(5):1467–1480, 2009.
- [BZM06] A. Brzezinski, G. Zussman, and E. Modiano. Enabling distributed throughput maximization in wireless mesh networks: a partitioning approach. In *Proc. ACM MobiCom*, pages 26–37, Los Angeles, CA, USA, September 24-29, 2006.
- [Cam89] K. Cameron. Induced matchings. *Discrete Applied Mathematics*, 24(1-3):97–102, 1989.
- [CKLS08] P. Chaporkar, K. Kar, X. Luo, and S. Sarkar. Throughput and fairness guarantees through maximal scheduling in wireless networks. *IEEE Transactions on Information Theory*, 54(2):572–594, 2008.
- [CXW09] H. Chen, X. Xie, and H. Wu. A queue-aware scheduling algorithm for multihop relay wireless cellular networks. In *Proc. IEEE Mobile WiMAX Symposium*, pages 63–68, Napa Valley, CA, USA, July 9-10, 2009.
- [Die97] R. Diestel. *Graph Theory (Graduate Texts in Mathematics)*, 1997.
- [EAM07] A. Eryilmaz, O. Asuman, and E. Modiano. Polynomial complexity algorithms for full utilization of multi-hop wireless networks. In *Proc. IEEE INFOCOM*, pages 499–507, Anchorage, AK, USA, May 6-12, 2007.
- [Fos53] F.G. Foster. On the stochastic matrices associated with certain queuing processes. *Ann. Math. Statistics*, 24:355–360, 1953.
- [FW77] S. Fiorini and R.J. Wilson. *Edge-colourings of graphs*, volume 16 of *Research Notes in Mathematics*. Pitman, 1977.
- [GK00] P. Gupta and P.R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, March 2000.
- [GLS09] A. Gupta, X. Lin, and R. Srikant. Low-complexity distributed scheduling algorithms for wireless networks. *IEEE/ACM Transactions on Networking*, 17(6):1846–1859, December 2009.

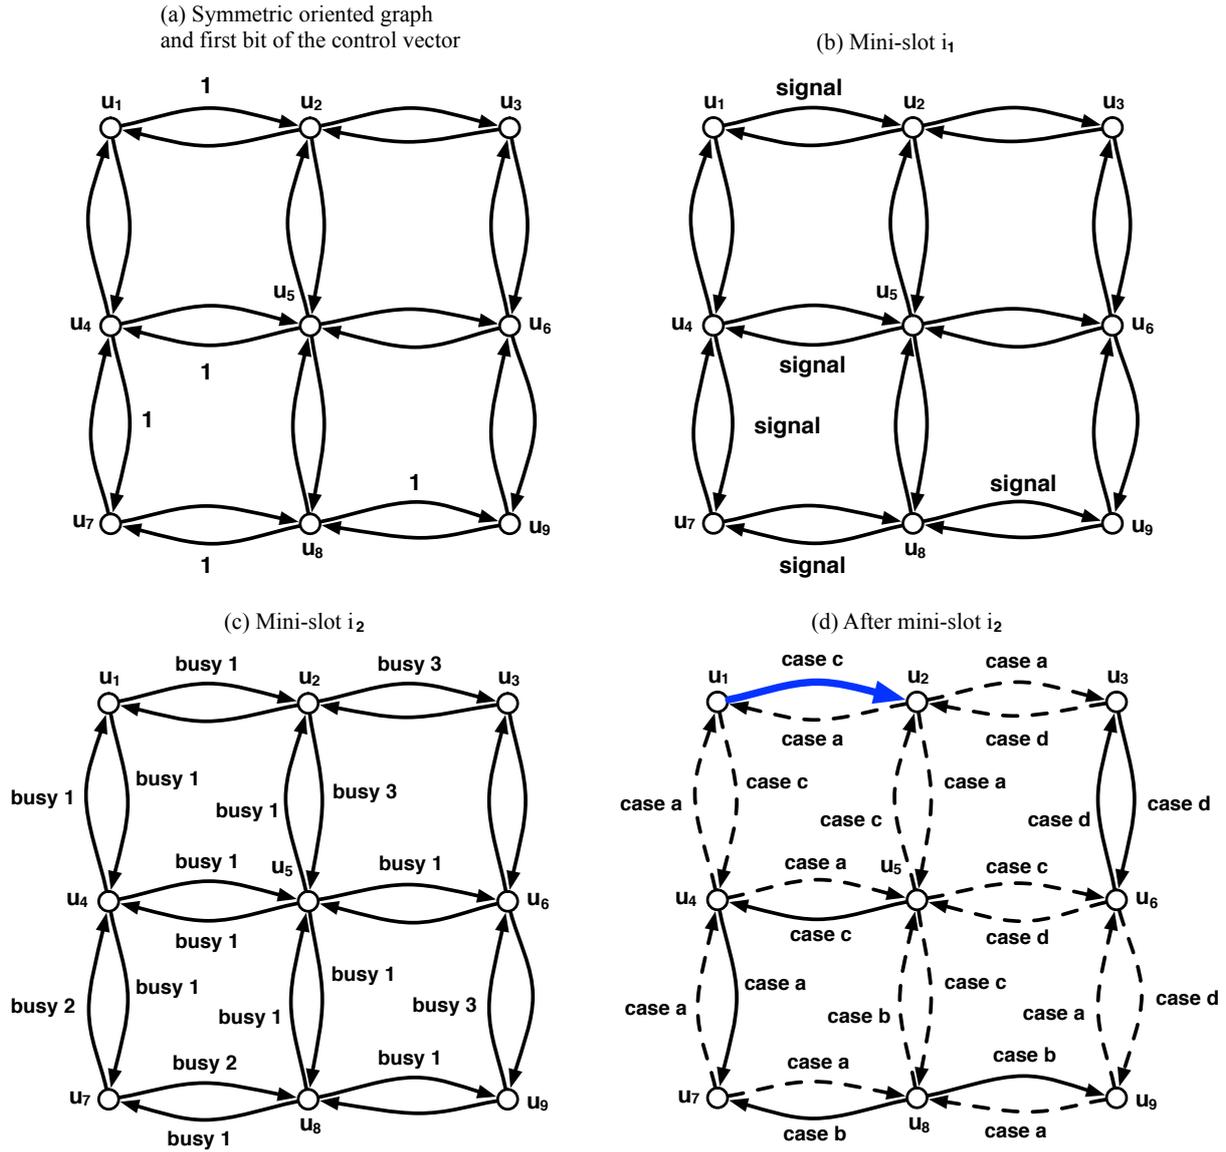


Fig. 16 Two mini-slots of AlgoLogNodes (that emulates AlgoLog) for an oriented grid network.

- [JSSW10] L. Jiang, D. Shah, J. Shin, and J. Walrand. Distributed random access algorithm: scheduling and congestion control. *IEEE Transactions on Information Theory*, 56(12):6182–6207, December 2010.
- [JW08] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in a wireless networks. In *Proc. 46th Allerton Conference on Communication, Control, and Computing*, Urbana-Champaign, IL, USA, 2008.
- [JW11] L. Jiang and J. Walrand. Approaching throughput-optimality in distributed CSMA scheduling algorithms with collisions. *IEEE/ACM Transactions on Networking*, 19(3):816–829, 2011.
- [KMP08] R. Klasing, N. Morales, and S. Pérennes. On the complexity of bandwidth allocation in radio networks. *Theoretical Computer Science*, 406(3):225 – 239, 2008.
- [KMPS04] V.S.A. Kumar, M. Marathe, S. Parthasarathy, and A. Srinivasan. End-to-end packet-scheduling in wireless ad-hoc networks. In *Proc. ACM-SIAM SODA*, pages 1021–1030, New Orleans, LO, USA, January 11-13, 2004.
- [LP86] L. Lovász and M.D. Plummer. *Matching Theory*, volume 29 of *Annals of Discrete Mathematics*. North-Holland, 1986.
- [MG92] J. Misra and D. Gries. A constructive proof of vizing’s theorem. *Information Processing Letters*, 41, 1992.
- [MSS06] R. Mazumdar, G. Sharma, and N. Shroff. Maximum weighted matching with interference constraints. in *Proc. FAWN, Pisa, Italy*, March 2006.
- [MSZ06] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless networks via gossiping. In *Proc. ACM SIGMETRICS - IFIP PERFORMANCE*, volume 34, pages 27–38, Saint Malo, France, June 26-30, 2006.
- [MT93] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, 1993.
- [OFGZ08] Y. Orlovich, G. Finke, V. Gordon, and I. Zverovich. Approximability results for the maximum and minimum maximal induced matching problems. *Discrete Optimization*, 5(3):584 – 593, 2008.

- [RSS09] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: an efficient randomized protocol for contention resolution. In *Proc. ACM SIGMETRICS - IFIP PERFORMANCE*, volume 37, pages 133–144, Seattle, WA, USA, June 15-19, 2009.
- [SBB13] F. Simatos, N. Bouman, and S. Borst. Linger issues in distributed scheduling. In *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '13, pages 141–152, New York, NY, USA, 2013. ACM.
- [SBS07] S. Sanghavi, L. Bui, and R. Srikant. Distributed link scheduling with constant overhead. In *Proc. ACM SIGMETRICS*, pages 313–324, San Diego, CA, USA, June 12-16, 2007.
- [SS12] D. Shah and J. Shin. Randomized scheduling algorithm for queueing networks. *Ann. Appl. Probab.*, 22(1):128–171, February 2012.
- [SST11] D. Shah, J. Shin, and P. Tetali. Medium access using queues. In *Proc. IEEE 52nd Annual Symposium on Foundations of Computer Sciences (FOCS)*, Palm Springs, CA, USA, October 2011.
- [SV82] L.J. Stockmeyer and V.V. Vazirani. NP-completeness of some generalizations of the maximum matching problem. *Information Processing Letters*, 15(1):14–19, 1982.
- [SW12a] D. Shad and D. Wischik. Log-weight scheduling in switched networks. *Queueing Systems (QUESTA)*, 71(1–2):97–136, June 2012.
- [SW12b] D. Shah and D. Wischik. Switched networks with maximum weight policies: Fluid approximation and multiplicative state space collapse. *Ann. Appl. Probab.*, 22(1):70–127, February 2012.
- [Tas97] L. Tassiulas. Scheduling and performance limits of networks with constantly changing topology. *IEEE Transactions on Information Theory*, 43(3):1067–1073, 1997.
- [TE92] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37(12), December 1992.
- [Wan09] P.-J. Wan. Multiflows in multihop wireless networks. In *Proc. ACM MobiHoc*, pages 85–94, New Orleans, LO, USA, May 18-21, 2009.
- [WSP06] X. Wu, R. Srikant, and J.R. Perkins. Queue length stability of maximal greedy schedules in wireless networks. In *Proc. of Information Theory and Applications Inaugural Workshop*, San Diego, CA, USA, February 6-10, 2006.