Modelling IEEE 802.11 CSMA/CA RTS/CTS with stochastic bigraphs with sharing

Muffy Calder and Michele Sevegnani

School of Computing Science, University of Glasgow, Lilybank Gardens, Glasgow G12 8QQ, UK

Abstract. Stochastic bigraphical reactive systems (SBRS) is a recent formalism for modelling systems that evolve in time and space. However, the underlying spatial model is based on sets of trees and thus cannot represent spatial locations that are shared among several entities in a simple or intuitive way. We adopt an extension of the formalism, *SBRS with sharing*, in which the topology is modelled by a directed acyclic graph structure. We give an overview of SBRS with sharing, we extend it with rule priorities, and then use it to develop a model of the 802.11 CSMA/CA RTS/CTS protocol with exponential backoff, for an *arbitrary* network topology with possibly overlapping signals. The model uses sharing to model overlapping connectedness areas, instantaneous prioritised rules for deterministic computations, and stochastic rules with exponential reaction rates to model constant and uniformly distributed timeouts and constant transmission times. Equivalence classes of model states modulo instantaneous reactions yield states in a CTMC that can be analysed using the model checker PRISM. We illustrate the model on a simple example wireless network with three overlapping signals and we present some example quantitative properties.

Keywords: 802.11 protocol; Formal verification; Quantitative analysis; Reactive systems; Stochastic bigraphs

1. Introduction

Systems of mobile devices and wireless networks evolve in both time and space. Stochastic bigraphical reactive systems (SBRS) [Mil09, KMT08] is a recent formalism for modelling the temporal and spatial evolution of computation: it should be ideally suited to this domain. However, SBRS assume an underlying model of space that is not overlapping, that is, spaces can be nested and connectivity is defined by a tree structure. A crucial aspect of wireless networks is signal interference: a device may be in the range of more than one signal. In other words, wireless networks are *overlapping*. Therefore, we adopt SBRS *with sharing* [SC10], a novel extension of the standard formalism whose spatial model is based on Directed Acyclic Graphs (DAG), thus allowing for a native representation of overlapping space. Additionally, we introduce rule priorities to allow for partial ordering of reaction rules. The result, *PSBRS with sharing*, allows us to model signal interference and in this paper we use it to model IEEE 802.11 CSMA/CA RTS/CTS, a protocol designed to handle signal interference.

Fig. 1. The use of virtual channel sensing using CSMA/CA

The main contributions of this paper are the following:

- extension of SBRS to *PSBRS with sharing*;
- an application of PSBRS with sharing to the 802.11 CSMA/CA RTS/CTS protocol for *any* wireless network topology, including interference;
- analysis of example properties for quantitative analysis.

Some aspects of the CSMA protocol have been modelled previously: for example collision detection on Ethernet is modelled by a MDP (Markov Decision Process) in [DFH⁺05]. A similar approach was taken in [KNS02] where probabilistic timed automata are used to model the basic two-way handshake mechanism¹ of the 802.11 protocol. The authors assume a fixed network topology consisting of two senders and two receivers. Furthermore, in their model there is exactly one shared signal, and thus each station can sense any other station. Properties of the system are specified in CSL (Continuous Stochastic Logic) [ASSB96] and automatically verified using probabilistic model checker PRISM [KNP11]. The model we present here differs in the following significant ways: support for arbitrary network topologies, and explicit representation of potentially overlapping wireless signals for all the stations in the network. These features are essential to represent networks in which two or more stations transmit to the same receiver and they cannot sense each other, thus causing a transmission collision. This is generally known in the literature as the *hidden node problem*.

The paper is organised as follows. The protocol is described informally in the next section and in Sect. 3 we give a brief overview of SBRS with sharing. Section 4 describes the bigraphical model for WLANs with arbitrary network topology. In Sect. 5 we present the graphical form of the reaction rules used to model the protocol. The evolutions of an example WLAN of three stations is given in Sect. 6. Section 7 outlines the CTMC encoding and some analysis results. Conclusions and directions for future work are in Sect. 8.

2. The protocol: IEEE 802.11 CSMA/CA with RTS/CTS handshake

We now describe informally the functioning of the protocol. Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) is the basic access mechanism in the 802.11 protocol [IEE05]. CSMA/CA adopts a slotted binary exponential backoff scheme to reduce collisions due to stations transmitting simultaneously. It defines two access mechanisms: the default, two-way handshaking technique called *basic access* and the optional four-way handshaking RTS/CTS reservation scheme. We focus on the latter here. Observe that this mechanism is more complicated because an extra handshake is introduced to reduce the collisions caused by the hidden node problem.

¹ Note that this protocol is different from RTS/CTS.

Table 1. Parameters of the CSMA/CA protocol

Parameter	Symbol	Typical value
DIFS	λ_D	50 μs
SIFS	λ_S	10 μs
Time slot	τ	20 μs
Time-out	λ_T	30 μs
Min contention window	t_{min}	15
Max contention window	t_{max}	1023
RTS	λ_R	160 μs
CTS	λ_C	112 μs
ACK	λ_A	112 μs

To illustrate the protocol, assume stations A, B, C, D where B and C are in the range of A, and B is in the range of D. In CSMA/CA a transaction is defined to be an exchange of packets between two stations that terminates with the transmission of an *acknowledgement* ACK. A is willing to transmit a packet to B and so transmits a *request-to-send* RTS, which includes the source, destination, and the duration of the following transaction (i.e. transmission of a *clear-to-send* CTS followed by a the data packet and the corresponding ACK). B responds (if the medium is free) with a CTS. Upon successful receipt of the CTS, A is allowed to transmit the actual data packet. Station B checks whether the received packet is corrupted by performing a cyclic redundancy check (CRC) and sends back an ACK. Receipt of the acknowledgement indicates that no collision occurred. Note that CTS, ACK and the data packet are sent only if the medium is idle for a *short inter-frame space* (SIFS) period. If A does not receive the CTS or the ACK packets, then it will retransmit until it gets acknowledged or reaches a given number of transmissions. All other stations receiving either the RTS and/or the CTS, e.g. C and D, set their virtual carrier sense timer (called NAV) with the duration value contained in those packets, and so will not attempt transmission until the NAV reaches 0. Figure 1 illustrates the scenario of A transmitting to B.

The exponential backoff scheme is executed each time the medium is sensed as busy or when a retransmission occurs. The station defers its transmission and the scheme initialises its random backoff timer. This value is obtained by uniformly choosing a value in the interval $[0, t]$, where $t_{min} \leq t \leq t_{max}$ is the current *contention window* (CW) size, and multiplying it for the duration of a *slot time* denoted by τ . At each timer initialisation, the interval is updated as follows: $[0, t']$, with $t' = 2t + 1$. Upon successful transmission t is reset to t_{min} . The parameters of the protocol specified by the standard are summarised in Table 1. The values used correspond to the times defined for a physical layer (PHY) with a data-rate of 1 Mbps. For example, λ_C is the time taken to transmit a CTS packet (with length 14 bytes = 112 bits) at 1 Mbps.

3. Stochastic BRS with sharing

Bigraphical reactive systems (BRS) is a recent formalism for modelling the temporal and spatial evolution of computation. It was initially introduced by Milner [Mil09] to provide a fully graphical model capable of representing both connectivity and locality. A BRS consists of a set of *bigraphs* and a set of *reaction rules*, which defines the dynamic evolution of the system by specifying how the set of bigraphs can be reconfigured. In this section we define informally BRS and its extensions (e.g. sharing, stochastic semantics and rule priorities), with enough detail to support the models adopted in our application.

Constituents of bigraphs and graphical notation

Some example bigraphs are depicted in Table 2 (left column). Dashed rectangles denote *regions*. Their rôle is to describe parts of the system that are not necessarily adjacent. The ovals and circles are *nodes*, which can represent physical or logical components within the system. Each node has a type, called *control*, denoted here by the labels A to D. The set of controls of a bigraphs is called *signature*. Each node can have zero, one or many *ports*, indicated by bullets, which represent possible connections. Actual connections are represented as *links*, depicted by solid (green) lines, which may connect ports and *names*. In the examples, they are ranged over by x, y, z . They can be thought of as links (or potential links) to other bigraphs. Gray squares are called *sites*. They encode parts of the model that have been abstracted away. Summarising, nodes represent the spatial placement of agents while links represent their communication capabilities.

Table 2. Operations on bigraphs

Operation	Algebraic	Graphical
Parallel product	$A_{xy} \parallel B_{yz}$
Merge product	$A_{xy} \mid B_{yz}$
Nesting	$A_{xy}.B_{xz}$
Name closure/New name	$/z A_{xz} \parallel y$
Sharing	share $A.1 \parallel B.1$ by ϕ in $C \mid D$ $\phi = [\{0\}, \{0, 1\}]$

Interfaces and sorting

The capabilities of bigraph B to interact with the external environment are recorded in its *interface*. For example, we write $B : 1 \rightarrow \langle 2, \{x, y\} \rangle$ to indicate that B has one site, two regions and its names are x and y . We sometimes use ϵ to denote 0 and the pair $\langle 0, \emptyset \rangle$.

Controls and links in a bigraph can be classified by means of *sorts*. A sorting discipline is a triple $\Sigma = (\Theta, \mathcal{K}, \Phi)$ where Θ is a non empty set of sorts, $\mathcal{K}$ is a signature and Φ is a *formation rule*. Sorts are ranged over by $a, b, \dots$. An example of sorting is given in Tables 3 and 4. A formation rule can be thought of as a set of properties a bigraph has to satisfy. For examples, it can specify that nodes of sort a may only contain b -nodes or that a -nodes may only be linked to b -nodes. Disjunctive sorts are written as $\widehat{ab}$, meaning that a node can either be of sort a or sort b . The interface of a sorted bigraph is expressed as follow: $B : a \rightarrow \langle b, \{z\} \rangle$. The notation indicates that B 's site is of sort a , its region is of sort b and its name is z . From now on, all bigraphs are assumed sorted.

Fig. 2. Example bigraph with sharing (a), corresponding DAG (b) and alternative graphical representation highlighting placing $\phi = [\{0\}, \{0, 1\}]$ (c)

Bigraphs with sharing

In [SC10] we have extended standard bigraphs so that a node or a site may be placed in the intersection of other nodes or regions. This implies a locality model based on DAGs instead of forests. An example bigraph with sharing and the DAG encoding its node hierarchy are shown in Fig. 2a, and b, respectively. Our extension yields several advantages. First, its completeness allows the direct representation of *any* topology with overlapping places. Second, the modelling phase is more natural and immediate, because no additional links, copies of nodes and special controls have to be introduced to encode intersections. Third, the DAG structure of the spatial model appears to have many similarities with the notions from category theory used in the standard definition of bigraphs. In more detail, bigraphs with sharing can be defined in the categorical setting of Milner’s s-categories and an axiomatisation of algebraic terms can be defined by a bialgebra over palce graphs.

For brevity, in the remainder of this paper, we use the term ‘bigraphs’ to refer to ‘bigraphs with sharing’.

Algebraic definition

The structure of a bigraph can also be formulated in an algebraic form closing resembling traditional process calculi. This is done by combining *elementary bigraphs* via the operations listed in Table 2. Note that, the algebraic form is equivalent to the graphical notation.

Parallel product $F \parallel G$ expresses a bigraphical term obtained by juxtaposing bigraphs F and G and merging their common names. Similarly, merge product $F \mid G$ denotes the juxtaposition of bigraphs F and G which is then placed inside a single region. Nesting operation $F.G$ allows us to place bigraph G inside F . Also in this case, common names are merged. Name closure $/x F$ is used to disallow connections on name x in bigraph F . This means that all F ’s links to x are broken and x removed. In example $/z A_{xz} \parallel y$ given in Table 2, new name introduction is simply indicated by y . The new name y is not linked to any node. Sharing operation $\text{share } F \text{ by } \phi \text{ in } G$ is a specialised version of nesting in which the regions of bigraph F can be placed inside the sites of bigraph G . The association between F ’s regions and G ’s sites is specified by *placing* ϕ . This allows the expression of shared nodes, i.e. nodes situated in the intersection of other nodes. Numbering of regions and sites proceeds from left to right starting from zero. Therefore, placings can be expressed by a vector of sets indicating the sites associated to each region. For instance, placing $\phi = [\{0\}, \{0, 1\}]$ indicates that the first F ’s region is placed in the first G ’s site while the second F ’s region is shared between the first and the second G ’s site. This is depicted in Fig. 2.

The elementary bigraphs most commonly used in our application are *identities* and 1. An example identity is $\text{id}_{2,ab}$. It indicates a bigraph with two regions each one containing a site, and two separate links a and b . Bigraph $1 : 0 \rightarrow 1$ consists of one single region.

Fig. 3. Bigraphical representation of a station M with two packets P^l and W^l . Station M has sent a data packet P^l and another packet W^l is stored in its queue (a). Simplified representation in which nodes of control A and Q are encoded by special diamond-shaped ports (b)

Bigraphical reactive systems

A *Bigraphical Reactive System* (BRS) consists of a set of bigraphs representing the state of the system, and a set of *reaction rules*, defining how the system can reconfigure itself. A reaction rule R is a pair (R, R') , where R and R' are bigraphs that can be inserted into the same host bigraph (i.e. they have the same interface). We sometimes indicate a rule as $R = R \rightarrow R'$. The evolution of a system S_t is derived by checking if R is an occurrence in S_t (this is also called *bigraph matching*) and by substituting it with R' to obtain a new system S_{t+1} . Such a *reaction* is indicated with $S_t \rightarrow_R S_{t+1}$. We also use $\rightarrow^*$ to indicate zero or more applications of a reaction.

In a *Stochastic BRS* (SBRS) [KMT08], rates are associated with reaction rules; the resulting transition system can be translated into a Continuous Time Markov Chain (CTMC) for quantitative analysis. We adopt the following notational conventions: $R = R \xrightarrow{\alpha} R'$ indicates a stochastic reaction rule with rate $\alpha \in \mathbb{R}^+$, while transition $S_t \xrightarrow{\rho} S_{t+1}$ denotes the update of bigraph S_t via R with rate ρ . Instantaneous reactions are characterised by rate $\alpha = \infty$. From now on, a reaction is deemed as instantaneous when its rate is omitted in the notation.

Priority BRS and priority SBRS

A *Priority BRS* (PBRS) is a BRS with *rule priorities* in the style of [BBKW89], i.e. by introducing a partial ordering on the rules of the reactive system. A reaction rule of lower priority can be applied only if no rule of higher priority is applicable. We write $R < R'$ to indicate that reaction rule R' has higher priority than reaction rule R . A priority class $\mathfrak{P}$ is a set of reaction rules with the same priority. By an abuse of notation, we write $\mathfrak{P} < \mathfrak{P}'$ when, for any two rules $R \in \mathfrak{P}$ and $R' \in \mathfrak{P}'$, we have $R < R'$. We also say that class $\mathfrak{P}$ has lower priority than class $\mathfrak{P}'$.

Similarly, a *Priority SBRS* (PSBRS) consists of an SBRS augmented with rule priorities. We require all the rules in a priority class to be instantaneous, if it contains an instantaneous rule.

4. Bigraphical model of wireless network topology

In this section we define the class of bigraphs used to represent WLANs in our model. Informally, nodes represent the main entities present in the network, namely stations, wireless signals and packets. Their relationships are expressed by links between them and by node sharing. Overlapping signals are modelled by a node representing a station in the intersection of two or more nodes representing the signals. This modelling strategy is explained in greater detail later on. The graphical form of an example bigraph encoding a station is shown in Fig. 3a. Packets to be sent are indicated by triangle-shaped nodes. Another example is given in Fig. 4b where the graphical form for the encoding of three stations with overlapping wireless signals is drawn. Before proceeding with a formal description of the bigraphical model, we observe that the oval shape used for nodes representing signals does not correspond to the actual shape of an area covered by a signal in a real WLAN. This is because the nesting of nodes represents a topological space and not an Euclidean space. Therefore, distances between places and their shapes are not expressible in bigraphs.

Table 3. Node controls for the bigraphical model

Description	Control	Arity	Sort	Notation
Signal idle	S^l	1	s	—
Signal locked	S_L^l	1	s	—
Signal clear	S_C^l	1	s	—
Signal error	S_E	1	s	—
Station idle	M	2	m	—
Station locked	M_L	2	m	—
Station deferred	M_D	2	m	—
Station potential	M_P	2	m	—
Station backoff	M_B	2	m	—
Address	A	1	a	—
Queue	Q	1	q	—
Packet waiting	W^l	2	p	▲
RTS packet	RTS^l	2	p	▲
CTS packet	CTS^l	2	p	▲
Data packet	P^l	2	p	▲

The parameters are $t \in T$ and $272 \mu s \leq l \leq 18768 \mu s$

The fifteen controls listed in Table 3 are used to encode the entities forming a WLAN. These are organised into sorts as follows. Controls describing wireless signals are grouped together into sort $s = \{S^l, S_L^l, S_C^l, S_E\}$. Parameter t records the contention window size of the station associated to the signal. Its values are drawn from set

$$T = \{t_{min} = 15, 31, 63, 127, 255, 511, 1023 = t_{max}\},$$

when parameters are defined as in Table 1. This corresponds to allowing seven transmission attempts for every data packet. Sort $m = \{M, M_L, M_D, M_P, M_B\}$ is used to represent nodes encoding stations, whilst a packet is indicated with a node of sort $p = \{W^l, RTS^l, CTS^l, P^l\}$. Here, parameter l is used to store the time employed to transmit a packet. Observe that the various controls in each sort are required to represent faithfully the WLAN throughout the different stages of the protocol. For instance, a node of control M_L indicates a *locked* station while a M_B -node expresses a station in *backoff* state. The meaning of each control is explained in Sect. 5, where the reaction rules of the SBRS modelling the protocol are introduced. Two additional sorts are defined: $a = \{A\}$ and $q = \{Q\}$. They represent the address of a node and a queue of packets respectively: virtual entities introduced only to simplify the modelling process. Specifically, they allow one to distinguish their links from the links of m and s -nodes. The set of sorts is written as $\Theta = \{s, m, p, a, q\}$. The signature is given by $\mathcal{K} = \bigcup_{s \in \Theta} s$.

The sorting discipline ensures that only bigraphs with a meaningful structure are constructed. For example, it forces packets to be always located inside a machine and a machine to be always surrounded by its signal. This is formalised in formation rule Φ with conditions Φ_i , $1 \leq i \leq 10$, given in Table 4. We briefly comment on each condition. Condition Φ_1 states that dummy nodes of control Q and A are atomic (i.e. they contain nothing). Conditions Φ_2 – Φ_4 specify a hierarchic structure on the placing of the nodes. In particular, s -nodes are the outermost nodes in the model. They contain m -nodes, which in turn contain $\widehat{p}a$ -nodes. Conditions Φ_5 , Φ_6 specify the placing described in the example in Fig. 3a. Condition Φ_7 ensures that an address A is always connected to a unique name and to one ancestor signal node. This allows the discrimination of signals according to the machine, and thus the address they belong to. Conditions Φ_8 , Φ_9 describe the queue-like structure formed by packets which sit within a station as in Fig. 3a. Condition Φ_{10} states that an address may be linked to several packets. This models the destination field in each packet. Finally, we refer the sorting used in the model as $\Sigma_{802.11} = (\mathcal{K}, \Theta, \Phi)$.

Our approach to modelling network topologies is explained through an example. Consider for instance the WLAN of three stations shown in Fig. 4a. Assume that both stations A and C are willing to transmit to station B. The bigraphical model needs to record that the system is composed of three stations, each one with its own wireless signal. Then, it also has to satisfy the following requirements:

- Req1** A and C sense B,
- Req2** B senses A and C,
- Req3** A does not sense C and C does not sense A.

Table 4. Conditions of formation rule Φ

Φ_1	all $\widehat{qa}$ -nodes are atomic
Φ_2	all children of a θ -root have sort θ , where $\theta \in \{s, m\}$
Φ_3	all children of an s -node have sort m
Φ_4	all children of an m -node have sort $\widehat{pa}$
Φ_5	an m -node has one a -child
Φ_6	a p -node has one q -child
Φ_7	an a -node is always linked to one grandparent of sort s and a name
Φ_8	in an m -node, one port may be linked to one m -node and the other may be linked to one p -node
Φ_9	in a p -node, one port may be linked to an a -node and the other is always linked to one $\widehat{mq}$ -node (see Fig. 3a)
Φ_{10}	an a -node may be linked to many p -nodes

Fig. 4. Modelling of network topologies. A network diagram of a WLAN of three stations (a) and the corresponding bigraphical encoding $N_0 : \epsilon \rightarrow \langle s, \{a_A, a_B, a_C, \} \rangle$ (b). Stations A and C are both willing to send to station B

This wireless network is modelled by bigraph $N_0 : \epsilon \rightarrow \langle s, \{a_A, a_B, a_C, \} \rangle$ given in Fig. 4b. For purposes of aesthetic clarity we adopt a simplified graphical notation in which $\widehat{aq}$ -nodes are encoded by special diamond ports as shown in Fig. 3b. Stations A, B, and C are encoded by the three M -nodes linked to names a_A , a_B and a_C , respectively. Their signals are indicated by nodes S^{t_A} , S^{t_B} , S^{t_C} , respectively, and they are linked to the corresponding name. Observe that each station sits within its own signal node. **Req1** is satisfied by station B being in the intersection of the two signals S^{t_A} and S^{t_C} . **Req2** is satisfied because the two m -nodes for stations A and C are contained by node S^{t_B} . Finally, **Req3** holds because station A is outside S^{t_C} and vice versa. An equivalent algebraic form of N_0 is

$$N_0 = \text{share}(M_A \parallel M_B \parallel M_C) \text{ by } \psi \text{ in } (\text{id}_{a_A a_B a_C} \mid S_{a_A}^{t_A} \mid S_{a_B}^{t_B} \mid S_{a_C}^{t_C}),$$

where terms

$$\begin{aligned} M_A &= (\text{id}_{1, a_A a_B} \parallel /x \parallel /r \parallel /q) (M_{rx}.(W_A \mid A_{a_A}.1)) \\ M_B &= (\text{id}_{1, a_B} \parallel /x \parallel /r) (M_{rx}.A_{a_B}.1) \\ M_C &= (\text{id}_{1, a_C a_B} \parallel /x \parallel /r \parallel /q) (M_{rx}.(W_C \mid A_{a_C}.1)), \end{aligned}$$

indicate stations A, B and C, respectively, and terms

$$W_A = W_{x a_B}^l.Q_q.1 \quad W_C = W_{x a_B}^k.Q_q.1$$

encode A's and C's packets (i.e. the blue triangles in Fig. 4b), respectively. The network topology is specified by placing

$$\psi = [\{0, 1\}, \{0, 1, 2\}, \{1, 2\}].$$

This simple example shows how our modelling strategy can provide adequate expressive power for the representation of arbitrary network topologies. Moreover, bigraphs with sharing allow us to express succinctly complex network topologies.

Table 5. Reaction rules

	Reaction rule	Rate	Description	Graphical form
1	$R_{RTS}(t, l)$	$\rho_1(t)$	machine is idle and willing to send $\rightarrow$ machine is locked and an RTS sent.	Fig. 5
2	$R_{CTS}(t, t', l)$	ρ_2	sender has sent an RTS $\rightarrow$ sender and receiver locked and linked, CTS sent.	Fig. 6
3	$R_{DATA}(l)$	$\rho_3(l)$	sender has sent a CTS $\rightarrow$ sender has sent its data packet.	Fig. 7
4	$R_{ACK}(t, t', l)$	ρ_4	data packet sent $\rightarrow$ ACK sent and locks released.	Fig. 8
5	$R_{BACK1}(t, t', l)$	ρ_5	sender (not the most recent) has to back off and receiver has detected a collision $\rightarrow$ locks are released and the sender is timed out.	Fig. 10
6	$R_{BACK2}(t, t', l)$	ρ_5	most recent sender has to back off and receiver has detected a collision $\rightarrow$ locks are released and the sender is timed out.	Fig. 11
7	$R_D(t)$	∞	station receives an RTS or a CTS $\rightarrow$ station defers.	Fig. 12
8	$R_P(t, t')$	∞	station receives an RTS from more than one station $\rightarrow$ station defers and detects conflict.	Fig. 13
9	$R_B(l)$	∞	sender cannot receive a CTS from the receiver $\rightarrow$ sender has to back off.	Fig. 14
10	$R_{UD}(t)$	∞	the NAV of a deferred station expires $\rightarrow$ lock released.	Fig. 15
11	$R_{UP}(t)$	∞	the NAV of a station in a conflict state expires $\rightarrow$ lock released.	Fig. 16
12	$R_{UC}(t)$	∞	sender or receiver terminated a transmission $\rightarrow$ lock released.	Fig. 17

We now list our modelling assumptions:

- All signals have equal power. This implies that if a station A is within the range of another station B, then B is within the range of A.
- The range of broadcast packets is the same of unicast packets.
- The range of small packets is the same of large packets.
- No packet is lost during transmission, *i.e.* ideal channel conditions.
- Static topology.

We remark that these assumptions derive from those made in the specification of the 802.11 protocol [IEEE05].

5. Stochastic reaction rules modelling the protocol

Before presenting the bigraphical reactive system, we introduce some notation and conventions. The stochastic reaction rules in our model have the form $R_L(\pi) = R \xrightarrow{\rho} R'$, where subscript L is a label describing the rule semantics and π is a list of parameters. All the rules respect the sorting $\Sigma_{802.11}$. We call the interface of the rule the interface of bigraphs R and R' . Reaction rates are indicated by $\rho(\pi)$. Names used as addresses range over $a, a', a_A, a_B, \dots$. Destination addresses are denoted by $d, d', \dots$. When defining reaction rules, name x is used to denote a link between two m-nodes while name q represents a link between a p-node and a node of control Q.

The 802.11 CSMA/CA RTS/CTS protocol is modelled as a PSBRS in which WLAN's configurations are encoded by $\Sigma_{802.11}$ sorted bigraphs and system updates are expressed by reaction rules. These are described in the remainder of this section. In order to obtain a general model of the protocol, each reaction rule is indexed by parameters such as the values of timers (indicated with t and t') and packet sizes (denoted by l). In the following a graphical description of each reaction rule is presented. Equivalent algebraic definitions are reported in Appendix A. A summary of the rules is in Table 5 using the following convention: each description takes the form $\langle \text{description of bigraph on lhs} \rangle \rightarrow \langle \text{description of bigraph on rhs} \rangle$.

The first rule $R_{RTS}(t, l)$ models the initial phase of the CSMA/CA protocol: whenever a sender senses the channel free, it is allowed to initiate a communication. A graphical representation of the reaction is given in Fig. 5. The station willing to transmit is encoded on the left-hand side by a node of control M containing a triangle of control W^l . On the right-hand side, a communication is instantiated. This is shown by the RTS^l triangle, the M_L and the S_L^l nodes. The reaction rate is defined as follows:

$$\rho_1(t) = (\lambda_D + \tau \frac{t}{2} + \lambda_R)^{-1}.$$

Fig. 5. $R_{\text{RTS}}(t, l)$: machine is idle and willing to send $\rightarrow$ machine is locked and an RTS sent

Fig. 6. $R_{\text{CTS}}(t, t', l)$: sender has sent an RTS $\rightarrow$ sender and receiver locked and linked, CTS sent

Note that it depends on t , i.e. the contention window size of the sending station. Term $\tau \frac{t}{2}$ is a constant indicating the average time the sender waits idle before its contention window of size t is fully consumed². In algebraic terms, this reaction rule is defined as

$$R_{\text{RTS}}(t, l) = \text{share}(\text{id} \parallel M_w) \text{ by } \psi_1 \text{ in } (\text{id}_{1,adq} \parallel S_a^t) \xrightarrow{\rho_1(t)} \text{share}(\text{id} \parallel M_l) \text{ by } \psi_1 \text{ in } (\text{id}_{1,adq} \parallel S_{La}^t)$$

where

$$M_w = (\text{id}_{1,adq} \parallel /x \parallel /r)(M_{rx} \cdot (\text{id} \mid W_{xd}^l \cdot Q_q \cdot 1 \mid A_a \cdot 1)) \quad \psi_1 = [\{1\}, \{0, 1\}]$$

$$M_l = (\text{id}_{1,adq} \parallel /x \parallel /r)(M_{Lrx} \cdot (\text{id} \mid \text{RTS}_{xd}^l \cdot Q_q \cdot 1 \mid A_a \cdot 1)).$$

The interface of the reaction rule is given by: $R_{\text{RTS}}(t, l) : \widehat{\text{mpa}} \rightarrow \langle \text{sm}, \{a, d, q\} \rangle$.

The second rule $R_{\text{CTS}}(t, t', l)$ describes the transmission of a CTS packet from the sender to the receiver. Its graphical representation is given in Fig. 6. The reaction can be triggered only when the sender and the receiver sense each other and the receiver is available for a transmission. These two preconditions are encoded in the left-hand side by two nodes of control M_L and M_D being in the same intersection of signals and by the node of control $S^{t'}$ containing the receiver, respectively. Moreover, M_D can be identified as the receiver because it is linked to the sender's RTS^l triangle. Note also that both the receiver and the sender can be in the range of other stations. On the right-hand side the receiver is locked (nodes of control $S_L^{t'}$ and M_L) and a link between the two machines is established. The reaction rate is $\rho_2 = (\lambda_S + \lambda_C)^{-1}$.

² Formally, term $\tau \frac{t}{2}$ is the mean of a uniformly distributed random variable on the interval $[0, \tau t]$.

Fig. 7. $R_{DATA}(l)$: sender has sent a CTS $\rightarrow$ sender has sent its data packet

Fig. 8. $R_{ACK}(t, t', l)$: data packet sent $\rightarrow$ ACK sent and locks released

The third rule $R_{DATA}(l)$ models the transmission of a data packet from the receiver to the sender. Its graphical form is depicted in Fig. 7. On the left-hand side, the sender's triangle is of control CTS^l , while on the right-hand side it is of control P^l . The rule can only be applied when a communication between two stations has already been initiated. This is encoded by the edge connecting the nodes of control M_L . Observe that the sender and the receiver can sense two different sets of stations because the M_L nodes are placed in two different regions. The reaction rate depends on the size of the data packet to be transmitted: $\rho_3(l) = (\lambda_S + l)^{-1}$.

The fourth rule $R_{ACK}(t, t', l)$ specifies the last phase of the protocol: upon successfully transmitting a data packet, the receiver sends back to the sender an ACK control packet and stops transmitting. A graphical representation of the rule is given in Fig. 8. On the left-hand side the two stations are engaged in a communication. This is encoded by the intersecting nodes of control S_L and by the link connecting the two M_L s. The triangle of the sender is of control P^l , meaning that a data packet has been transmitted. On the right-hand side locks on both the stations are released and the sender's backoff timer is reset to t_{min} . This is expressed by the nodes of control M , $S_C^{t_{min}}$ and S_C^t . Moreover, the link between the two machines is removed. Also the triangle is removed. The reaction rate is $\rho_4 = (\lambda_S + \lambda_A)^{-1}$.

A pictorial representation of how the previous four stochastic reaction rules: $R_{RTS}(t, l)$, $R_{CTS}(t, t', l)$, $R_{DATA}(l)$, and $R_{ACK}(t, t', l)$, are used to encode the packet exchange between two stations as specified by the protocol is given in Fig. 9. From this figure, it is possible to infer the contribution of the relevant constants to the associated rates (summarised in Table 5). For example, rate $\rho_1(t)$, the rate for $R_{RTS}(t, l)$ rule, includes the sum of λ_D , the constant for DIFS, $\tau \frac{t}{2}$, the constant denoting the average duration of the sender's contention window of size t and λ_R , the constant indicating the transmission time of an RTS packet.

Fig. 9. Packet exchange between two stations and corresponding reaction rules

Fig. 10. $R_{BACK1}(t, t', l)$: sender (not the most recent) has to back off and receiver has detected a collision $\rightarrow$ locks are released and the sender is timed out

Fig. 11. $R_{BACK2}(t, t', l)$: most recent sender has to back off and receiver has detected a collision $\rightarrow$ locks are released and the sender is timed out

The fifth rule $R_{BACK1}(t, t', l)$ models collisions when two stations simultaneously try to transmit an RTS packet. It is expressed graphically in Fig. 10. On the left-hand side, triangle of control RTS is placed within a node of control M_B . This represents the fact that the sender has to back off. The receiver is in a potential collision state, indicated by control M_P . On the right-hand side locks are released and the sender's contention window is increased. This is encoded by the node of control S_C^{2t+1} , the triangle of control W^l and nodes of control M . The modification of parameter t models the exponential back-off procedure. The reaction rate is $\rho_5 = \lambda_T^{-1}$, i.e the sender's time-out expires before a CTS packet is received.

The sixth rule $R_{BACK2}(t, t', l)$ is very similar to the previous one. It is applied when the receiver has control M_D instead of M_P . This happens when all the other conflicting stations have already backed off. A graphical representation of the rule is given in Fig. 11. In both the previous two rules, nodes of control S_C^{2t+1} are replaced by S_E when $t \geq t_{max}$. This models the protocol when the maximum number of transmission attempts is reached. Rates associated to each of the previous reaction rules are summarised in Table 5.

The six stochastic reaction rules described so far encode behaviours of stations as specified by the CSMA/CA RTS/CTS protocol. In the remainder of this section, a set of six instantaneous reaction rules is introduced. These rules are used to enforce conditions on the system that assure the rules presented above can only be fired in a meaningful order. More formally, they force some invariants to be satisfied before one of the stochastic rule is triggered. For example, any station receiving more than one RTS has to be marked as M_P . This allows the correct choice of whether rule $R_{CTS}(t, t', l)$ or rule $R_{BACK1}(t, t', l)$ has to be applied. We explain the interplay between stochastic and instantaneous rules in more detail when rule priorities are assigned (at the end of this section).

An overview of the instantaneous rules is the following. The first three rules deal with the detection and marking of stations involved in a communication. Rule $R_D(t)$ detects and marks stations that have to defer. This is depicted in Fig. 12. On the left-hand side, an idle station has received an RTS from some other station. This is modelled by a node of control M placed within a node of control S_L^t . Moreover, links indicating addresses are not joined together i.e. there are two distinct links to names a and a' . On the right-hand side, the station is marked by assigning control M_D to it.

Fig. 12. $R_D(t)$: station receives an RTS or a CTS $\rightarrow$ station defers

Fig. 13. $R_P(t, t')$: station receives an RTS from more than one station $\rightarrow$ station defers and detects conflict

Rule $R_P(t, t')$ detects and marks stations that are in a potential conflict state. The graphical representation is given in Fig. 13. In the left-hand side, a station receives an RTS from more than one station. This is shown by the node of control M_D placed within the intersection of two nodes of control S_L^t . Observe that the station can be in the range of other machines. This is modelled by the extra region surrounding M_D . On the right-hand side, the station detecting the conflict is marked by assigning control M_P to it.

Rule $R_B(t)$ detects and marks stations that have to backoff when a conflict occurs. This is drawn in Fig. 14. On the left-hand side, the sender is trying to communicate with a receiver in a conflict state. This is shown by nodes of control M_L , RTS^t and M_P . Note that the triangle and the receiver are linked together on name d . On the right-hand side, the sender is marked by assigning control M_B to it.

The last three instantaneous reactions are used to release the locks when a communication is terminated. Rule $R_{UD}(t)$ releases the lock of a deferred station when its NAV reaches 0. This is modelled by the node of control M_D enclosed by a node of control S_C^t on the left-hand side. On the right-hand side, the lock is released by assigning control M to the station. A graphical representation of the reaction rule is given in Fig. 15.

Fig. 14. $R_B(t)$: sender cannot receive a CTS from the receiver $\rightarrow$ sender has to back off

Fig. 15. $R_{UD}(t)$: the NAV of a deferred station expires $\rightarrow$ lock released

Fig. 16. $R_{UP}(t)$: the NAV of a station in a conflict state expires $\rightarrow$ lock released

Similar to the previous rule $R_{UP}(t)$ is used to release the lock of a station in a potential conflict state. This is encoded by the node of control M_P on the left-hand side. The rule is depicted in Fig. 16.

Finally, rule $R_{UC}(t)$ models the release of the lock at the end of a transmission. The graphical representation is in Fig. 17. On the left-hand side, the station is idle and ready to release the lock. This is encoded by the node of control S_C^t linked to its child node of control M . On the right-hand side, the final stage of the protocol is performed by substituting S_C^t with S^t . Note that this rule is not in conflict with rules $R_{UD}(t)$ and $R_{UP}(t)$ because it can only be applied to a sender or a receiver.

Reaction rules are organised into six priority classes as follows:

$$\begin{aligned} \mathfrak{P}_5 &= \{R_{UD}(t), R_{UP}(t)\} & \mathfrak{P}_4 &= \{R_{UC}(t)\} & \mathfrak{P}_3 &= \{R_D(t)\} & \mathfrak{P}_2 &= \{R_P(t, t')\} & \mathfrak{P}_1 &= \{R_B(l)\} \\ \mathfrak{P}_0 &= \{R_{RTS}(t, l), R_{CTS}(t, t', l), R_{DATA}(l), R_{ACK}(t, t', l), R_{BACK1}(t, t', l), R_{BACK2}(t, t', l)\} \end{aligned}$$

Fig. 17. $R_{UC}(t)$: sender or receiver terminated a transmission $\rightarrow$ lock released

All stochastic rules belong to $\mathfrak{P}_0$, i.e. the class with lowest priority. Therefore, any stochastic rule can be applied only when no instantaneous rule can be applied. A brief description of the priority classes is the following.

$\mathfrak{P}_5$ is the class with highest priority. Rules assigned to this class mark with control M an idle station in a clear signal (i.e. of control S_C^I). Class $\mathfrak{P}_4$ has one rule that marks with control M a sender or receiver inside its clear signal, the signal is marked as S^I . Since $\mathfrak{P}_4 < \mathfrak{P}_5$, all idle stations must be marked as M before the clear signal containing them is marked as S^I . Similarly, rules in priority classes $\mathfrak{P}_3, \mathfrak{P}_2, \mathfrak{P}_1$ mark a station inside a locked signal (S_L^I) as deferred (M_D), in a potential conflict state (M_P) and backing off (M_B), respectively. This is necessary because rules with higher priorities may have marked a station as M even if it is contained in a locked signal. Finally, stochastic rules in $\mathfrak{P}_0$ can be applied safely since all the preconditions are enforced by the prior application of the rules with higher priorities.

Discussion

A difficulty with rewriting rules in general, and therefore a problem for us here, is how to deal with situations where we want to rewrite because a condition does *not* apply. For instance, we want to apply rule $R_{RTS}(t, l)$ only when the sender is not in a locked signal. In our approach, this issue has been solved by introducing priorities on reaction rules as described above. However, there are two modelling choices other than using PSBRs.

One option is to parameterise the reactions on the topology of the network. For instance, the first rule (RTS packet transmission) can be modified so all the idle signals sensed by the sender are explicitly listed. Since only idle signals are included and no other signals are allowed via shared regions, then the rule can only be applied when the sender is not within a locked signal. An additional parameter n indicating the number of signals in the sender's range is required, i.e. $R_{RTS}(t, l, n)$. A major drawback of this approach is that a different rule has to be instantiated for every network topology. Moreover, the specification of the rule becomes cumbersome and difficult to read. An advantage is that less controls are required for the modelling of the protocol's phases and no instantaneous reaction rule is introduced.

The other option involves extending the notion of reaction rule to include what may be called *conditional reaction rules*, having the form (φ, R, R', ρ) , where φ is a **BiLog** [CMS05] predicate expressing a condition imposed upon the bigraphical context in which left-hand side R occurs. Also with this option, less controls are required and instantaneous reaction rules can be avoided. A disadvantage is that it is often necessary to specify predicates involving universal quantifiers or expressing the absence of a pattern. These may be difficult to verify especially when the size of the system grows.

Our approach allows the specification of concise rules that can be applied to any network topology. This conciseness is important not only because it allows for a clearer description of the model, as we showed in this section, but also because the size of the left-hand sides of our reaction rules is always constant. Therefore, there is no performance deterioration when a reaction rule is applied to a bigraph with many nodes.

A common issue with instantaneous rules is the generation of intermediate states that may not represent meaningful network configurations. In order to minimise the number of states stored during the execution of a model, we treat instantaneous rules like rewriting within an equivalence class, and only store a canonical form, after applying all possible instantaneous rules. This is possible because in our model instantaneous rules within a priority class are confluent, thus they always yield a unique fixed point. We explain this procedure in Sect. 7.

We note that we have defined exponential rates from time constants, i.e. if we have a constant k waiting (or transmission, or sum of constants) time, then we use as CTMC rate k^{-1} .

Finally, we remark that although the protocol assumes a static topology, we could easily define rules for station movement in and out of a signal range, i.e. protocol behaviour in a dynamic topology.

6. Execution of an example network with three stations

In this section we describe two example executions of our model starting from initial state N_0 , the bigraph encoding of an example WLAN with three stations given in Fig. 4. Recall that all the possible evolutions of a system are obtained by successive applications of reaction rules chosen among the ones described above.

The first execution encodes the successful transmission of a packet from station A to station B, with no collisions occurring. It is the result of the application of reaction rules $R_{RTS}(t_A, l)$, $R_D(t_B)$, $R_{CTS}(t_A, t_B, l)$, $R_D(t_C)$, $R_{DATA}(l)$, $R_{ACK}(t_A, t_B, l)$, $R_{UD}(t_C)$, $R_{UC}(t_A)$ and $R_{UC}(t_B)$, in that order. The dynamic evolution of the network is represented pictorially in Fig. 18. Observe that in state N_6 (second diagram in the third row), the sender's contention windows size is reset: $t'_A = t_{min}$. An alternative execution can be obtained by changing the order in which the locks are released, i.e. by applying $R_{UC}(t_B)$ before $R_{UC}(t_A)$. The resulting state is again N_9 . Execution $N_0 \xrightarrow{\rho_1(t_A)} \triangleright_{RTS} \dots \triangleright_{UC} N_9$ can further proceed in an analogous way by transmitting the remaining packet from C to B. The symmetrical evolution in which station C sends its packet to B can be obtained from initial state N_0 by applying reaction rules $R_{RTS}(t_C, k)$, $R_D(t_B)$, $R_{CTS}(t_C, t_B, l)$, $R_D(t_A)$, $R_{DATA}(k)$, $R_{ACK}(t_C, t_B, k)$, $R_{UD}(t_A)$, $R_{UC}(t_C)$ and $R_{UC}(t_B)$, in that order.

The second execution encodes a transmission in which a collision occurs. It is shown in Fig. 19. The initial state is bigraph N_2 obtained in the trace described above. As can be seen, in this evolution station C sends its RTS before B transmits its CTS. Therefore, stations C and A have to defer and the backoff scheme is executed. This is encoded by the application of rules $R_{BACK1}(t_A, t_B, l)$ and $R_{BACK2}(t_C, t_B, k)$. Station C can try to retransmit (by applying rule $R_{RTS}(t'_C, k)$) only after its timeout has expired. The contention windows of senders A and C are exponentially increased in states N'_4 and N'_7 , respectively. Therefore, $t'_i = 2t_i + 1$ with $i \in \{A, C\}$. An alternative execution can be computed by inverting the order of application of rules $R_B(l)$ and $R_B(k)$. This leads to state N'_3 . Another interleaving is obtained when $R_{BACK1}(t_C, t_B, k)$ is applied before $R_{BACK2}(t_A, t_B, l)$. Also in this case, the final state is still N'_8 .

The two evolutions above show how instantaneous rules must be applied before stochastic rules as required by the priority hierarchy of our model. For instance, in the second trace, rules $R_P(t_A, t_C)$, $R_B(l)$ and $R_B(k)$ are applied before rule $R_{BACK1}(t_A, t_B, l)$.

7. CTMC analysis

We first explain how to compute the CTMC capturing the semantics of a PSBRS. In particular, we describe our approach to reduce the number of intermediate states generated by the application of instantaneous reaction rules. Whenever a reconfiguration S' of a state S (i.e. $S \xrightarrow{\rho} S'$) is obtained, the unique fixed point S^* is computed by iteratively applying to S' all the instantaneous rules. Then, state S^* is stored in memory in place of S' as a reconfiguration of S (i.e. $S \xrightarrow{\rho} S^*$). A formal representation of the procedure is given by

$$S \xrightarrow{\rho} S' \xrightarrow{\mathfrak{P}_i} S'' \xrightarrow{\mathfrak{P}_{i-1}} S''' \dots \xrightarrow{\mathfrak{P}_j} S^* = S \xrightarrow{\rho} S^*$$

with $j < i - 1$. All the priority classes, $\mathfrak{P}_i < \dots < \mathfrak{P}_j$, contain only instantaneous reaction rules. Note that all the rule applications respect the priority hierarchy and that all the intermediate states S' , S'' , S''' , $\dots$ obtained during the computation of S^* are discarded. Furthermore, by definition of fixed point, no instantaneous rule can be applied to S^* . This always forces the first reaction to have rate $\rho \neq \infty$. This approach can only be used if each priority class whose instantaneous reactions are ignored has a unique fixed point (as is the case in our model).

The procedures described above can be explained by showing how to compute the CTMC starting from bigraph N_0 , the example WLAN of three stations used in the example in the previous section. The parameters of the model are packet transmission times $l = 8464 \mu s$ and $k = 4368 \mu s$.³ Contention windows for stations A, B and C are $t_A, t_B, t_C = 15$. In the PSBRS, iterated applications of instantaneous reaction rules belonging to the same priority class always yield a fixed point. Hence, intermediate states generated by $\mathfrak{P}_5, \dots, \mathfrak{P}_1$ can be ignored. Another property of our model is that any reaction $S \xrightarrow{\rho} S'$ can only be obtained by a single application of a stochastic reaction rule R . This may not hold with more complex topologies.

³ Parameters l and k correspond to the transmission times of data-packets with payloads of 512 bytes and 1024 bytes, respectively. The values are obtained by adding to the payload lengths 30 bytes for the MAC header and 4 bytes for the frame check sequence (FCS).

Fig. 18. Evolution $N_0 \xrightarrow{\rho_1(t_A)} \text{RTS} \cdots \xrightarrow{\text{UC}} N_9$ showing the successful transmission of a packet from sender A to destination B. Sender's contention window size is reset after the application of rule $R_{\text{ACK}}(t_A, t_B, l)$ (i.e. $t'_A = t_{\min}$)

Fig. 19. Evolution $N_2 \xrightarrow{\rho_1(t_C)}_{RTS} N'_0 \dots \xrightarrow{UC} N'_8$ showing the collision occurring when senders A and C try to transmit to the same destination B. The contention windows of the senders are exponentially increased: $t'_i = 2t_i + 1$ with $i \in \{A, C\}$

The resulting CTMC (indicated by $\mathcal{M}$) is given in Fig. 20. Reactions corresponding to actions on station C are indicated with left facing arrows and those on station A by right facing arrows. The diagram is organised into seven diamond sub-parts each one representing a different transmission attempt after a collision. Each diamond is indicated by $\mathcal{M}^i$ with $1 \leq i \leq 7$. Note that only the first two diamonds are shown in Fig. 20. The central node of a diamond (indicated with $\circ$) is the top vertex of the successive one as for instance state N_8' in $\mathcal{M}^1$ and $\mathcal{M}^2$. The bottom node of $\mathcal{M}^1$ (indicated with $\diamond$) is shared between all the seven diamonds. It represents the state in which both packets are successfully transmitted to B. It is shared because the contention windows of stations A and C are reset to $t_{min} = 15$ in every diamond after every transmission of the ACK control packet. The reactions going to the centre of the diamond encode transmission collisions. A side of a diamond is formed from four reactions, encoding the successful transmission of a packet. For instance, side $N_0 \rightarrow^* N_9$ encodes transmission of A's packet before C's packet at the first attempt. Dashed arrows correspond to the encoding of the two executions described in Sect. 6. The CTMC has two deadlock states. One is node $\circ$ in $\mathcal{M}^7$ representing stations A and C when the maximum number of transmission attempts is reached, i.e. both signals are S_E . The other is node $\diamond$ described above.

As can be seen in the diagram, the CTMC does not contain the intermediate states generated by the application of instantaneous rules. Take for instance reaction $N_0 \xrightarrow{\rho_1(t_A)} N_2$ in $\mathcal{M}^1$. By inspecting the execution in Fig. 18, we see that it corresponds to reactions

$$N_0 \xrightarrow{\rho_1(t_A)} N_1 \xrightarrow[\text{RTS}]{\text{D}} N_2.$$

In this case, intermediate bigraph N_1 is not used as a CTMC state. Also note that no instantaneous reaction rule can be applied to N_1 . Analogously, reaction $N_2 \xrightarrow{\rho_1(t_C)} N_3'$ in $\mathcal{M}^1$ corresponds to

as shown in Fig. 19. Here, intermediate bigraphs N_0' , N_1' , N_2' and $\overline{N_2'}$ are discarded. Note that only one between N_2' and $\overline{N_2'}$ is computed by the rewriting engine, because only one path is needed to compute fix point N_3' . In both the reactions considered, the stochastic rate is the rate of the stochastic reaction rule applied at the beginning of the corresponding sequence.

Finally, we note that the CTMCs generated from our model, for any finite topology, will be finite because there are a finite number of transmissions, a finite number of packets, and a finite number of stations.

7.1. Analysis of quantitative properties

It is possible to express quantitative properties about the behaviour of our model by combining **BiLog** and **CSL** (Continuous Stochastic Logic) [ASSB96]. In more detail, we use **BiLog** to express predicates that are used as atomic propositions in **CSL**. In [CS12], it was shown that this kind of predicate can be reduced to one or more instances of bigraph matching. Therefore, the predicates in our analysis are represented as bigraphs. For example, we define pattern $P_{\varphi_{error}} = S_E$ to encode predicate φ_{error} , which is true whenever a station reaches the maximum number of transmission attempts, namely there is a signal node in error state. This predicate can then be used in any **CSL** formula such as $\psi_1 = \mathbf{P}_{<0.001}(\mathbf{F} \varphi_{error})$, i.e. eventually a station is in error state with probability less than 0.001. Note that the only state satisfying φ_{error} is node $\circ$ in $\mathcal{M}^7$.

Quantitative properties can be evaluated with standard tools for CTMC model checking. Here we use the probabilistic model checker **PRISM** [KNP11] on some simple **CSL** formulae expressing properties on $\mathcal{M}$. In particular, we take advantage of the *explicit model import* feature in **PRISM** which allows the importation of a CTMC's rate matrix and its labelling function.

It is also possible to specify properties that evaluate to a numerical value. Let us analyse some examples.

Fig. 20. Structure of CTMC $\mathcal{M}$ derived from the PSBRS on initial state N_0 (see Fig. 4b). Dashed arrows correspond to the two executions described in Figs. 18 and 19

Example 7.1 The probability of a station being in an error state is given by formula $\mathbf{P}_{=?}(\mathbf{F} \varphi_{\text{error}}) = 6.08 \times 10^{-7}$, which implies that $\mathcal{M} \models \psi_1$. This small value is explained by the fact that only two stations compete for transmission in our example topology. When more overlapping stations are present in the network, iterated collisions are more likely, thus this value increases.

Example 7.2 The probability that a collision occurs is expressed by $\mathbf{P}_{=?}(\mathbf{F} \varphi_{\text{collision}}) = 0.36$. Predicate $\varphi_{\text{collision}}$ is encoded by the matching of pattern $P_{\varphi_{\text{collision}}} = \mathbf{M}_{\text{Bxy}} \parallel \mathbf{M}_{\text{Bzw}}$. It corresponds to state N'_3 in $\mathcal{M}^1$ and the states in the same position in $\mathcal{M}^i$ with $1 < i \leq 7$.

Example 7.3 The probability of station A successfully transmitting its packet in $1 \leq n \leq 7$ transmission attempts is expressed by $\mathbf{P}_{=?}(\varphi_A^n)$. The cumulative plot is given in Fig. 21. Predicate φ_A^n corresponds to pattern

$$\text{share}(\text{id} \parallel \text{rx} \parallel \mathbf{M}_{\text{Lrx}} \cdot (\text{id} \mid \mathbf{P}_{\text{xd}}^l \mid \mathbf{A}_{\text{aA}} \cdot 1) \text{ by } [\{0\}, \{0, 1\}] \text{ in } (\mathbf{S}_{\text{LdA}}^t \parallel \text{id}_{1, \text{draA}}))$$

with $t = 2^{n+3} - 1$. Therefore, the labelling function marks the nodes indicated with $\star$ in $\mathcal{M}$. Observe that the predicate cannot be defined to simply match state $\diamond$ because information about the number of transmission attempts is lost after every application of reaction rule $\xrightarrow{\rho_4} \text{ACK}$ when the contention window is reset.

Discussion

Previous studies of the IEEE 802.11 standard have focussed mainly on the quantitative analysis of the basic access mechanism. Two example works that consider the RTS/CTS handshake are [Bia00] and [CG97]. The results reported by those authors are used to evaluate the performance of the protocol, e.g. system throughput and congestion rate, and are obtained by means of simulation. In both cases, the network topology is not encoded directly, but a probability of collision p is assigned to each station. Hence, higher values of p allow one to encode “denser” topologies. Another example of analysis is the analytic approach presented in [AA10], but it is difficult to compare with our analysis, since again the network topology is modelled implicitly by a collision probability.

Fig. 21. Cumulative probability of sending A's packet against the number of transmission attempts

Fig. 22. Example topology not meeting the assumptions of the protocol: B is in the range of C but C is not in the range of B

In general, the model checking approach has some benefits over simulation [Cla08]. Probably, the most relevant advantage in the context of wireless networks modelling is that all possible states are explored. In particular, model checking can explore those states having a reachability probability that is too small to be discoverable by simulation. This is fundamental to obtain precise and trustworthy results especially when analysing safety-critical systems and border-line or unusual system behaviours. Consider for instance Example 7.1 in which the probability to reach an error state is computed, i.e. $\mathbf{P}_{=?}(\mathbf{F} \varphi_{\text{error}}) = 6.08 \times 10^{-7}$. In this case, simulation based approaches could erroneously conclude that the system can never be in an error state (thus giving a wrong answer) because the probability to discover such states is too low to be found in a reasonable number of simulation runs.

The quantitative properties we presented earlier in this section serve mainly to show how SBRS can effectively be analysed and to highlight the expressive power of both bigraphs with sharing and **BiLog**. Following these examples, more advanced analysis can be carried out by researchers in wireless networking and practitioners in the field of performance evaluation. For instance, network throughput can be studied by varying the parameters of the protocol and the exponential increase of the contention window size can be compared with different back-off schemes. Another interesting investigation is to analyse the effects of the RTS/CTS mechanism on different network topologies (e.g. sparse or dense) to find out when the mechanism improves network throughput and when it introduces overhead.

Our bigraphical model can be extended easily to analyse more realistic systems, i.e. systems in which the assumptions made in the 802.11 protocol are not met. For example, formation rule Φ of sorting $\Sigma_{802.11}$ given in Table 4 allows for the specification of networks in which signals do not have equal power as it is the case in the WLAN encoded by the bigraph in Fig. 22. The diagram shows that receiver B is within the signal range of station C but C is not in the range of B. Therefore, station C cannot receive a CTS packet from B and reaction rule $\mathbf{R}_{\text{CTS}}(t, l)$ can never be applied. In order to model this kind of system, stochastic reaction rule $\mathbf{R}_{\text{T}}(t, l)$ given in Fig. 23 needs to be introduced. It encodes the behaviour of a sender timing out before a CTS packet is received. It can be thought as the opposite of $\mathbf{R}_{\text{RTS}}(t, l)$ described in Sect. 5 where the size of the contention window is increased in the reactum as in reaction rules $\mathbf{R}_{\text{BACK1}}(t, t', l)$ and $\mathbf{R}_{\text{BACK2}}(t, t', l)$. Also in this case, node of control

Fig. 23. $R_T(t, l)$: machine is locked and an RTS sent $\rightarrow$ lock is released and the machine is timed out

S_C^{2t+1} is replaced by S_E when $t \geq t_{max}$. With the addition of this new reaction rule, the system can evolve to a state in which C's signal has control S_E after seven transmission attempts.

Another modification for a more realistic model is to represent a constant transmission time τ with an Erlang distribution of shape k and rate $\frac{k}{\tau}$. In a CTMC, this can be achieved by replacing transitions in the form $N \xrightarrow{\frac{1}{\tau}} N'$ with $N = N_1 \xrightarrow{\frac{k}{\tau}} N_2 \xrightarrow{\frac{k}{\tau}} \dots N' = N_{k+1}$. Note that when $k = 1$ we have an exponential distribution. A drawback of this approach is that $k - 1$ states are added to the CTMC and that an additional priority class needs be added in order to avoid interleaving with the other reaction rules of the model. We observed that in most cases $k \approx 4$ is sufficient for an accurate approximation.

Finally, the model can be further refined by adding reaction rules encoding the movement of stations in and out of signals. They would encode both the spatial movement of mobile stations and the variation of signal coverage due to frequency interference. These reaction rules would usually be slower than the other stochastic reaction rules.

8. Conclusions and future work

We have adopted SBRS with sharing to exploit a locality model based on DAGs and we illustrated the extended formalism with a model of the IEEE 802.11 CSMA/CA RTS/CTS protocol. The graphical form of the rules illustrates the spatial aspects of the protocol in a precise and intuitive way. Reaction rules are then organised into priority classes to form a PSBRS allowing the specification of concise rules that can be efficiently applied on any arbitrary network topology. The simple WLAN of three stations with overlapping signals provides an example of a practical application of stochastic bigraphs and one that could not be modelled using the original formalism. The CTMC encoding offers the possibility of quantitative analysis, allowing us to quantify probabilities such as a packet is received, or two stations are in an error state, or how the former depends on the number of retransmissions. Future work includes further extensions to the tools for bigraphs such as support for automatic CTMC generation and translation of **BiLog** predicates.

Appendix

A. Algebraic definitions of reaction rules

In the following, we give an algebraic definition for each reaction rule summarised in Table 5.

Bigraphs indicating machine nodes at various stages of the protocol are given by:

$$\begin{aligned} M_w &= (\text{id}_{1,aqd} \parallel /x \parallel /r)(M_{rx}.(\text{id} \mid W_{xd}^l.Q_q.1 \mid A_a.1)) \\ M_l &= (\text{id}_{1,aqd} \parallel /x \parallel /r)(M_{Lrx}.(\text{id} \mid \text{RTS}_{xd}^l.Q_q.1 \mid A_a.1)) \\ M_b &= (\text{id}_{1,aqd} \parallel /x \parallel /r)(M_{Brx}.(\text{id} \mid \text{RTS}_{xd}^l.Q_q.1 \mid A_a.1)). \end{aligned}$$

We begin with the algebraic description of the six stochastic reaction rules:

1. $R_{\text{RTS}}(t, l) : \widehat{\text{mpa}} \rightarrow \langle \text{sm}, \{a, d, q\} \rangle$

$$R_{\text{RTS}}(t, l) = \text{share}(\text{id} \parallel M_w) \text{ by } \psi_1 \text{ in } (\text{id}_{1, \text{adq}} \parallel S_a^t) \xrightarrow{\rho_1(t)} \text{share}(\text{id} \parallel M_l) \text{ by } \psi_1 \text{ in } (\text{id}_{1, \text{adq}} \parallel S_{La}^t)$$

$$\psi_1 = [\{1\}, \{0, 1\}]$$

2. $R_{\text{CTS}}(t, t', l) : \widehat{\text{papamm}} \rightarrow \langle \text{ssmm}, \{x, a, d, q\} \rangle$

$$R_{\text{CTS}}(t, t', l) = \text{share}(M_l \parallel M_{\text{Drx}}.(\text{id} \mid A_d.1) \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{La}^t \mid S_d^{t'}) \parallel /r)$$

$$\xrightarrow{\rho_2}$$

$$\text{share}(S \parallel R \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{La}^t \mid S_{Ld}^{t'}) \parallel /r)$$

$$S = (\text{id}_{1, \text{aqdr}} \parallel /x)(M_{\text{Lrx}}.(\text{id} \mid \text{CTS}_{\text{xd}}^l.Q_q.1 \mid A_a.1))$$

$$R = M_{\text{Lrx}}.(\text{id} \mid A_d.1)$$

$$\psi_2 = [\{0, 2, 3\}, \{1, 2, 3\}, \{2\}, \{3\}]$$

3. $R_{\text{DATA}}(l) : \widehat{\text{papa}} \rightarrow \langle \text{mm}, \{x, a, d, q\} \rangle$

$$R_{\text{DATA}}(l) = (\text{id}_{2, \text{xdaq}} \parallel /r)(S \parallel R) \xrightarrow{\rho_3(l)} (\text{id}_{2, \text{xdaq}} \parallel /r)(S' \parallel R)$$

$$S' = (\text{id}_{1, \text{aqdr}} \parallel /x)(M_{\text{Lrx}}.(\text{id} \mid P_{\text{xd}}^l.Q_q.1 \mid A_a.1))$$

4. $R_{\text{ACK}}(t, t', l) : \widehat{\text{papamm}} \rightarrow \langle \text{ssmm}, \{x, a, d, q\} \rangle$

$$R_{\text{ACK}}(t, t', l) = \text{share}(S' \parallel R \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{La}^t \mid S_{Ld}^{t'}) \parallel /r)$$

$$\xrightarrow{\rho_4}$$

$$\text{share}(S'' \parallel R' \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{Ca}^{t_{\min}} \mid S_{Cd}^{t'}))$$

$$S'' = (\text{id}_{1, \text{aq}} \parallel /r)(M_{\text{rq}}.(\text{id} \mid A_a.1))$$

$$R' = (\text{id}_{1, \text{dx}} \parallel /r)(M_{\text{rx}}.(\text{id} \mid A_d.1))$$

5. $R_{\text{BACK1}} : \widehat{\text{papamm}} \rightarrow \langle \text{ssmm}, \{x, a, d, q\} \rangle$

$$R_{\text{BACK1}}(t, t', l) = \text{share}(M_b \parallel M_{\text{Prx}}.(\text{id} \mid A_d.1) \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{La}^t \mid S_d^{t'}) \parallel /r) \xrightarrow{\rho_5} F$$

$$F = \text{share}(M_w \parallel M_{\text{rx}}.(\text{id} \mid A_d.1) \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{Ca}^{2t+1} \mid S_d^{t'}) \parallel /r)$$

6. $R_{\text{BACK2}}(t, t', l) : \widehat{\text{papamm}} \rightarrow \langle \text{ssmm}, \{x, a, d, q\} \rangle$

$$R_{\text{BACK2}}(t, t', l) = \text{share}(M_b \parallel M_{\text{Drx}}.(\text{id} \mid A_d.1) \parallel \text{id}_2) \text{ by } \psi_2 \text{ in } (\text{id}_{2, \text{xadq}} \parallel (S_{La}^t \mid S_d^{t'}) \parallel /r) \xrightarrow{\rho_5} F$$

Finally, instantaneous reaction rules are defined as follows:

7. $R_{\text{D}}(t) : \widehat{\text{mpa}} \rightarrow \langle \text{ms}, \{a', a, x\} \rangle$

$$R_{\text{D}}(t) = \text{share}(\text{id} \parallel M_{\text{rx}}.(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1, \text{ax}} \parallel S_{La'}^t \parallel /r)$$

$$\longrightarrow$$

$$\text{share}(\text{id} \parallel M_{\text{Drx}}.(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1, \text{ax}} \parallel S_{La'}^t \parallel /r)$$

$$8. R_P(t, t') : \widehat{mmpa} \rightarrow \langle ms, \{a', a'', a, x\} \rangle$$

$$\begin{aligned} R_P(t, t') &= \text{share}(\text{id}_2 \parallel M_{D_{rx}}(\text{id} \mid A_a.1)) \text{ by } \psi_3 \text{ in } (\text{id}_{1,ax} \parallel (S'_{La'} \mid S'_{La''}) \parallel /r) \\ &\longrightarrow \\ &\quad \text{share}(\text{id}_2 \parallel M_{P_{rx}}(\text{id} \mid A_a.1)) \text{ by } \psi_3 \text{ in } (\text{id}_{1,ax} \parallel (S'_{La'} \mid S'_{La''}) \parallel /r) \\ \psi_3 &= [\{1\}, \{2\}, \{0, 1, 2\}] \end{aligned}$$

$$9. R_B(l) : \widehat{papa} \rightarrow \langle mm, \{xaqd\} \rangle$$

$$R_B(l) = M_l \parallel (\text{id}_{1,xd} \parallel /r)(M_{P_{rx}}(\text{id} \mid A_d.1)) \longrightarrow M_b \parallel (\text{id}_{1,xd} \parallel /r)(M_{P_{rx}}(\text{id} \mid A_d.1))$$

$$10. R_{UD}(t) : \widehat{mpa} \rightarrow \langle ms, \{a', a, x\} \rangle$$

$$\begin{aligned} R_{UD}(t) &= \text{share}(\text{id} \parallel M_{D_{rx}}(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1,ax} \parallel S'_{Ca'} \parallel /r) \\ &\longrightarrow \\ &\quad \text{share}(\text{id} \parallel M_{rx}(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1,ax} \parallel S'_{Ca'} \parallel /r) \end{aligned}$$

$$11. R_{UP}(t) : \widehat{mpa} \rightarrow \langle ms, \{a', a, x\} \rangle$$

$$\begin{aligned} R_{UP}(t) &= \text{share}(\text{id} \parallel M_{P_{rx}}(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1,ax} \parallel S'_{Ca'} \parallel /r) \\ &\longrightarrow \\ &\quad \text{share}(\text{id} \parallel M_{rx}(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1,ax} \parallel S'_{Ca'} \parallel /r) \end{aligned}$$

$$12. R_{UC}(t) : \widehat{mpa} \rightarrow \langle ms, \{a', a, x\} \rangle$$

$$\begin{aligned} R_{UC}(t) &= \text{share}(\text{id} \parallel M_{rx}(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1,ax} \parallel S'_{Ca} \parallel /r) \\ &\longrightarrow \\ &\quad \text{share}(\text{id} \parallel M_{rx}(\text{id} \mid A_a.1)) \text{ by } \psi_1 \text{ in } (\text{id}_{1,ax} \parallel S'_a \parallel /r) \end{aligned}$$

References

- [AA10] Alshanyour A, Agarwal A (2010) Performance of IEEE 802.11 RTS/CTS with finite buffer and load in imperfect channels: modeling and analysis. In: GLOBECOM, pp 1–6
- [ASSB96] Aziz A, Sanwal K, Singhal V, Brayton R (1996) Verifying continuous time Markov chains. In: Alur R, Henzinger T (eds) Computer aided verification. Lecture notes in computer science, vol 1102. Springer, Berlin-Heidelberg, pp 269–276
- [BBKW89] Baeten JCM, Bergstra JA, Klop JW, Weijland WP (1989) Term-rewriting systems with rule priorities. Theor Comput Sci 67:283–301
- [Bia00] Bianchi G (2000) Performance analysis of the IEEE 802.11 distributed coordination function. IEEE J Selected Areas Commun 18(3):535–547
- [CG97] Chhaya HS, Gupta S (1997) Performance modeling of asynchronous data transfer methods of IEEE 802.11 MAC protocol. Wireless Networks 3(3):217–234
- [Cla08] Clarke E (2008) The birth of model checking. In: Grumberg O, Veith H (eds) 25 Years of model checking. Lecture notes in computer science, vol 5000. Springer, Berlin-Heidelberg, pp 1–26
- [CMS05] Conforti G, Macedonio D, Sassone V (2005) Spatial logics for bigraphs. In: Caires L, Italiano GF, Monteiro L, Palamidessi C, Yung M (eds) Proceedings of the 32th international colloquium on automata, languages and programming (ICALP'05). Lecture notes in computer science, vol 3580. Springer, Berlin, pp 766–778
- [CS12] Calder M, Sevegnani M (2012) Matching and rewriting for bigraphs with sharing. Submitted for publication
- [DFH⁺05] Duflot M, Fribourg L, Herault T, Lassaigle R, Magniette F, Messika S, Peyronnet S, Picaronny C (2005) Probabilistic model checking of the CSMA/CD protocol using prism and APMC. Electron Notes Theor Comput Sci 128(6):195–214
- [IEE05] IEEE. (2005) Std. 802.11e-2005, Part 11: wireless LAN medium access control (MAC) and physical layer (PHY) specifications, Amendment 8: medium access control (MAC) quality of service enhancements
- [KMT08] Krivine J, Milner R, Troina A (2008) Stochastic bigraphs. Electron Notes Theor Comput Sci 218:73–96
- [KNP11] Kwiatkowska M, Norman G, Parker D (2011) PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan G, Qadeer S (eds) Proceedings of 23rd international conference on computer aided verification (CAV'11). Lecture notes in computer science, vol 6806. Springer, Berlin, pp 585–591

- [KNS02] Kwiatkowska M, Norman G, Sproston J (2002) Probabilistic model checking of the IEEE 802.11 wireless local area network protocol. In: Hermanns H, Segala R (eds) Proceedings of 2nd joint international workshop on process algebra and probabilistic methods, performance modeling and verification (PAPM/PROBMIV'02). Lecture notes in computer science, vol 2399. Springer, Berlin, pp 169–187
- [Mil09] Milner R (2009) The Space and motion of communicating agents. Cambridge University Press, Cambridge
- [SC10] Sevegnani M, Calder M (2010) Bigraphs with sharing. Technical report TR-2010-310, University of Glasgow

Received 23 July 2012

Revised 16 November 2012

Accepted 17 November 2012 by Joachim Parrow

Published online 8 January 2013