

A versatile infinite-state Markov reward model to study bottlenecks in 2-hop ad hoc networks

Anne Remke, Boudewijn R. Haverkort, Lucia Cloth
University of Twente

[anne, brh, lucia]@cs.utwente.nl

Abstract

In a 2-hop IEEE 802.11-based wireless LAN, the distributed coordination function (DCF) tends to equally share the available capacity among the contending stations. Recently alternative capacity sharing strategies have been made possible. We propose a versatile infinite-state Markov reward model to study the bottleneck node in a 2-hop IEEE 802.11-based ad hoc network for different adaptive capacity sharing strategies. We use infinite-state stochastic Petri nets (iSPNs) to specify our model, from which the underlying QBD-type Markov-reward models are automatically derived. The impact of the different capacity sharing strategies is analyzed by CSRL model checking of the underlying infinite-state QBD, for which we provide new techniques. Our modeling approach helps in deciding under which circumstances which adaptive capacity sharing strategy is most appropriate.

1 Introduction

The availability of cheap yet powerful wireless access technology, most notably the IEEE 802.11 (wireless LAN), has given an impulse to the development of wireless ad hoc networks. In such a network, the stations (also called nodes) that are in reach of each other, facilitate connectivity by forwarding traffic, e.g., to obtain access to the fixed internet. In an 802.11 ad hoc network, the stations that are in mutual reach, and that help each other in obtaining and maintaining connectivity, are at the same time also competitors, as they all contend for the same resource, i.e., the shared ether as transmission medium. The medium access control of 802.11 has been based on CSMA/CA [10, 16] and is commonly referred to as the distributed coordination function (DCF) [10, 16]. Research has shown that, effectively, the DCF tends to equally share the capacity among the contending stations [3, 12]. Although this appears to be a nice fairness property, this fairness may lead to undesirable situations in case one of the nodes happens to function as a bridge toward either another group of nodes, or to the fixed internet, because in that case a clear bottleneck situation arises. In this paper we will show how model checking can be used to analyze alternative and more flexible capac-

ity sharing approaches. Earlier work on the performance of IEEE 802.11 ad hoc networks considers a variety of scenarios, [17] categorizes this earlier work nicely. However, none of the analytical papers mentioned in [17], explicitly addresses the delays or through-puts in a *multihop* ad hoc network. In [17], a two-hop ad hoc network is considered, where the second hop has to forward the traffic of many sources (the first hops), thus forming a bottleneck, since all active stations have to share the transmission capacity. This study yields explicit (closed-form) equations for the expected overall delay and the expected delay at the bottleneck, by translating the model at hand into a generalized processor sharing model as studied extensively by Cohen [5]. Although the analysis is approximate, good results are obtained, as confirmed by simulations. However, this evaluation approach is limited in that it only allows for an equal sharing of transmission capacity between active stations.

In this paper, we follow the same line of modeling as in [17], however, we do allow for alternative capacity sharing strategies as well; such strategies are made possible through the recent QoS-extension of the IEEE 802.11 standard, e.g., through the EDCA (“E”) version [9]. In doing so, we can study the impact of adaptive capacity sharing strategies that recognize potential bottlenecks and adapt accordingly, as detailed in Section 2. Since we want to have flexibility in modeling a variety of adaptive capacity sharing strategies and at the same time more modeling convenience, we have chosen to specify our models as infinite-state stochastic Petri nets (iSPNs) and not at state level. The underlying infinite-state Markov chain, which can be automatically generated from the iSPN, obeys a quasi-birth-death (QBD) structure, for which we showed that CSL model checking is feasible [15]. In this paper we equip QBDs with rewards and use the logic CSRL (continuous stochastic reward logic) that has been proposed to reason about time *and* rewards in finite Markov chains [1] for model checking infinite state *Markov reward models*. To do so, we adapt techniques for CSRL model checking finite Markov chains to the infinite case. The rest of this paper is organized as follows. In Section 2 we present three different strategies for dividing the available radio capacity among the competing stations; we also provide a concise introduction into iSPNs and QBDs. Section 3 presents two different adaptive capacity sharing strategies. Then, Section 4 presents a set




Figure 1. Quasi-Birth-Death Process

of steady-state and transient performance measures of interest as expressions in the continuous and stochastic reward logic (CSRL) [7]. We explain how the CSRL operators can be evaluated and which measures we calculate with each of them. We will analyze the expected buffer occupancy, the expected number of active sources, the work done in a given time by the sources and by the bridge as well as some other measures. The results are presented and interpreted in 5. We provide conclusions in Section 6.

2 Modeling approach and basic model

In this section we first describe Quasi-Birth-Death models and their high-level specification as infinite-state stochastic Petri nets, which we will use to model the bottleneck in 2-hop ad hoc networks. We provide a slightly more detailed system description, before we model the system as iSPN and transform it to the underlying QBD.

2.1 QBDs and iSPNs

The infinite state space of a QBD can be viewed as a two-dimensional strip, which is finite in one dimension and infinite in the other. Formally, a *labeled QBD* Q with rewards of order (N_0, N_r) (with $N_0, N_r \in \mathbb{N}^+$) is a labeled infinite-state CTMC, where the reward function $\rho : S \rightarrow \mathbf{R}_{\geq 0}$ assigns a reward $\rho(s)$ to each state s . Figure 1 gives a graphical representation of a QBD. The set of states is composed as $S = \{0, \dots, N_0 - 1\} \times \{0\} \cup \{0, \dots, N_r - 1\} \times \mathbb{N}^+$, where the first part represents the boundary level with N_0 states, and the second part the infinite number of repeating levels, each with N_r states. The steady-state probabilities of a QBD can be calculated in a level-wise fashion, using e.g., *matrix-geometric methods* [11, 13], which exploit the repetitive structure in the generator matrix. To compute transient state probabilities for the infinite-state QBDs, we developed a uniformization-based approach [15]. iSPNs [14] can be used as high-level description of QBDs. They have one unbounded place that corresponds directly to the infinite dimension of the QBD.

2.2 System description

We address a wireless ad hoc network in which the individual nodes communicate with each other through the IEEE 802.11 CSMA/CA media access control protocol. We do not discuss or model the DCF explicitly; this has been




Figure 2. Wireless ad hoc network




Figure 3. High level specification as iSPN

done in the past and by now this access mechanism is, as such, well understood [3, 12]. The scenario under study has N active nodes, the so-called *sources*, that are all within reach of each other. Additionally there is a special node, referred to as *bridge* or *bottleneck node* B . This node, which can be reached by the N sources, is the only node that connects to yet another (set of) external node(s) E through which, for instance, the fixed internet can be reached. Thus, the bridge B forms a natural bottleneck: all traffic originating from the N sources, as well as the traffic passing through the bridge has to share the same wireless transmission capacity. Figure 2 illustrates this scenario.

Assuming the classical WLAN capacity of 11 Mbps, earlier work has shown that the effective capacity is about 60% of this value. This effective capacity is about equally shared among the competing nodes. For an average packet size of 1500 bytes, we obtain an effective capacity of 500 packets per second, that is denoted C . An inactive source becomes active after a negative exponentially distributed amount of time (with rate λ) and immediately instantiate a flow. During such a *flow time*, the source remains active and continues to transmit packets to the bottleneck station B . Following the assumptions made in [17], the amount of work put forward by each flow (the amount of packets comprising the flow) on average is $F = 500$ packets. The assumed values for the system parameters are summarized in Table 1.

Table 1. Values for the system parameters

parameter	
source arrival rate	$\lambda \in [0.1, 0.5] \frac{1}{\text{sec}}$
average flow size	$E[F] = 500$ packets
overall radio capacity	$C = 500$ packets/s
maximum of active sources	$K = 10$

When modeling the bridge in 2-hop ad hoc networks as an iSPN, clearly the unbounded place has to be the buffer of the bottleneck B . We limit the maximum number of active

sources to some finite number $K = 10$, as iSPNs only allow for one unbounded place. This is a reasonable restriction, as the number of active sources in an ad-hoc network cannot be arbitrarily high.

2.3 Basic model

Since the active sources as well as the bridge share the same wireless capacity, the amount of time to complete the work put forward by a single flow depends on the activity of other flows. In the basic model with N active sources ($N \in \{1, \dots, K\}$), the bottleneck as well as each of the active sources receive a fraction $\frac{C}{N+1}$ of the overall effective capacity. Hence, the sources in total receive a fraction $\frac{N \cdot C}{N+1}$ of the transmission capacity, and the bridge just $\frac{C}{N+1}$. However, all the transmission capacity devoted to the sources leads immediately to traffic arriving at the bridge, that is, all the packets the sources are allowed to send result in packets to be queued and processed at B . This makes that the bridge B will be overloaded and unstable as soon as more than one source becomes active. In case just one source is active, that source and the bridge will equally share the capacity, such that the source cannot generate more traffic than the bridge can accommodate. A backlog of packets at the bridge can only be cut back whenever none of the sources is active. From the interaction between the regular sources and the bridge, we note that, the effective rate with which the bridge can complete its packet transmissions depends on the number of active sources (flows) and that the time it takes to complete a flow depends on the number of active flows, that is the number of active sources. As we do not need to distinguish between individual active sources, we can model the number of active sources as shown in the left part of the Petri net in Figure III. Sources become active after a negative exponentially distributed amount of time and start transmitting a flow that has an exponentially distributed length with parameter $\frac{1}{F}$. However, the duration of a flow does not only depend on its size but also on the radio capacity a source can use to transmit the flow, as can be seen in Table 2: the rate of transition *source departure* is $\frac{1}{F} \cdot \frac{N \cdot C}{N+1}$, where N is the number of active sources. The right part of Figure 2.2 that contains the unbounded place *buffer* actually models the bridge of the system. Transition *input* models the total arrival stream of packets from all active sources and transition *output* models the service of packets at B . Both rates depend on the number of active sources and the amount of radio capacity that is distributed to each source and the bottleneck node. The rates for the individual transitions as used in the basic model are given in Table 2. To obtain the underlying QBD, we specify the iSPN in CSPL. Using the CSPL implementation [2], we obtain the underlying finite CTMC that consists of the finite boundary part and the repeating levels. The resulting QBD for the basic model is given in Figure 4. Every level consists of $K+1$ states, modeling the number of active sources. Whenever at least one source is present, packets can arrive and whenever at least one packet is present, this packet can




Figure 4. Underlying QBD of the basic model

be served. In the basic model the bottleneck node and every individual source always receive the same amount of radio capacity; this results in a very homogeneous QBD, where only the leftmost K states belong to the boundary level.

Table 2. Rates for Two-hop models

	input	output	source arrival	source departure
Basic	$\frac{N \cdot C}{N+1}$	$\frac{C}{N+1}$	λ	$\frac{1}{F} \cdot \frac{C \cdot N}{N+1}$
BRT				
- low occupancy	$\frac{N \cdot C}{N+1}$	$\frac{C}{N+1}$	λ	$\frac{1}{F} \cdot \frac{C \cdot N}{N+1}$
- high occupancy	$\frac{1}{2} \cdot C$	$\frac{1}{2} \cdot C$	λ	$\frac{1}{F} \cdot \frac{1}{2} \cdot C$
SRT				
- startup	$\frac{N \cdot C}{N+1}$	$\frac{C}{N+1}$	λ	$\frac{1}{F} \cdot \frac{C \cdot N}{N+1}$
- run	$\frac{N \cdot C}{N+m}$	$\frac{C \cdot m}{N+m}$	λ	$\frac{1}{F} \cdot \frac{C \cdot N}{N+m}$
- clearance	$\frac{N \cdot C}{N+m}$	$\frac{C \cdot m}{N+m}$	λ	$\frac{1}{F} \cdot \frac{C \cdot N}{N+m}$

3 Adaptive capacity sharing

We propose two different approaches for adaptive capacity sharing, namely the *buffer-related threshold* model (BRT) and the *source-related threshold* model (SRT). Both approaches are first modeled as iSPNs and then transformed to the underlying QBD. The SRT model can be implemented by appropriately setting the differentiation parameters according to [9], whereas the BRT model cannot be realized yet, due to the relation to the buffer-content. However the BRT model describes a sensible way to allocate radio capacity, as it directly refers to the place where the problem lies, that is, to the buffer in the bottleneck station. Using this type of model-based analysis allows us to evaluate a strategy like BRT even before it can be implemented.

3.1 Buffer-related threshold

In the basic setting we allocate most of the radio capacity to the sources. Clearly, this benefits the sources and can lead to very high buffer occupancy from which the bridge may not be able to recover. To improve the performance of the bridge, we allocate more radio capacity to the bridge as

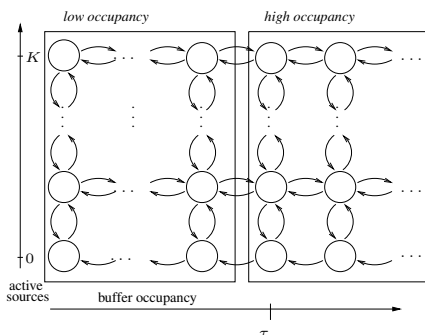


Figure 5. Underlying QBD of the BRT model

soon as the buffer occupancy becomes a given threshold τ . The buffer-related threshold model (BRT) then consists of two phases: a *low occupancy* phase and a *high occupancy* phase. In the *low occupancy* phase we allocate the same amount of radio capacity to every source and to the bridge (as in the basic model). As soon as the buffer occupancy becomes τ we switch to the *high occupancy* phase. In this second phase, half of the server capacity is allocated to the bridge and the other half is distributed equally between the sources. The structure of the iSPN for the BRT model is exactly the same as for the basic model, only the rates have to be adapted according to Table 2. The rates now depend on the buffer occupancy, that is, the content of the unbounded place, however, from threshold τ onwards they remain unchanged. This only enlarges the boundary part of the underlying QBD from the leftmost K states to the leftmost $K \cdot \tau$ states. The two phases of the model are located next to each other in 'the infinite dimension' as shown in Figure 5.

3.2 Source-related threshold

Whereas the BRT model defines its threshold related to the buffer occupancy, the phases of the source-related threshold model (SRT) depend on the number of active sources. Whenever the system starts, it has zero active sources. Up to a small number of active sources we allow for the same distribution of radio capacity as in the basic model (*start up* phase). However, as soon as the number of active sources reaches a given number m , the model enters the *run* phase, in which we assign m times as much radio capacity to the bridge than to each single source. Whenever the number of active sources falls again below m , the system switches to its third phase: *clearance*. To allow the bottleneck node to first work off the complete buffer content, in *clearance* the bottleneck node keeps the higher share of radio capacity as in *run* until the buffer is empty. The model then switches back to *startup*. In the SRT model we have to keep track in which phase the model currently is. This is done by adding an extra part to the iSPN that consists of one place for each phase. The rates for the SRT model are also given in Table 2. Note, that the SRT model for $m = 1$ equals the basic model. Each level of the underlying QBD contains states that belong to the three different phases of the SRT model, as shown in Figure 6. The number of states

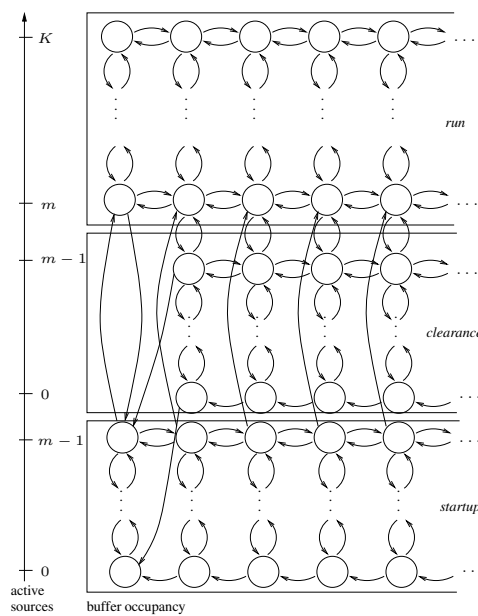


Figure 6. Underlying QBD of the SRT model

per level in the SRT model is $2m + K - m + 1 = K + m + 1$, that is, $K - m + 1$ states for *run*, m states for *startup*, and m states for *clearance*. In the SRT model the 'leftmost' K states belong to the boundary level.

4 Measures of Interest

In the following we will analyze how the strategies for adaptive capacity sharing influence the system behavior. This is done by means of model checking. With model checking we first formally specify the measures of interest as formulas of a given logic. Second, model checking provides algorithms for automatically evaluating these formulas on a given model. In this section we will first describe the logic CSRL, before we describe which measures we are going to evaluate, how these measures can be expressed in CSRL, and how we have to assign rewards and atomic propositions to the QBD, to do so. Furthermore, we discuss how to validate the CSRL operators, that are needed in this study. Note, that we are the first to do CSRL model checking on infinite-state CTMCs.

4.1 The logic CSRL

The logic CSRL (continuous stochastic reward logic) has been proposed to reason about time and rewards in finite Markov chains [1]. It is an extension of the logic CSL (continuous stochastic logic), which we recently applied to infinite-state Markov chains, in particular to QBDs [15]. As such, CSRL is a specification formalism for performability measures over CTMCs extended with a reward structure (Markov reward models) [7, 1]. The syntax of a CSRL state

formula is defined as follows:

$$\Phi ::= tt \mid ap \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathcal{S}_{\triangleright p}(\Phi) \mid \mathcal{P}_{\triangleright p}(\phi) \mid \mathcal{E}_{\triangleright r}(\Phi),$$

where ϕ is a path formula constructed by

$$\phi ::= \mathcal{X}_{\leq r}^{\leq t} \Phi \mid \Phi \mathcal{U}_{\leq r}^{\leq t} \Phi. \quad (1)$$

The key difference to CSL is that the path-operators are equipped with two parameters. The additional parameter r represents a bound on the accumulated reward. The steady-state operator $\mathcal{S}_{\triangleright p}(\Phi)$ denotes that the steady-state probability for a Φ -state meets the bound p . $\mathcal{P}_{\triangleright p}(\phi)$ asserts that the probability measure of the paths satisfying ϕ meets the bound p . The expected reward operator $\mathcal{E}_{\triangleright r}(\Phi)$ expresses that the steady-state reward rate (that is, $\sum_{s' \in \text{Sat}(\Phi)} \pi(s, s') \rho(s') \leq r$ with $\pi(s, s')$ the steady-state probability to be in state s' , having started in s) meets the bound r . The next operator $\mathcal{X}_{\leq r}^{\leq t} \Phi$ asserts that a transition to a Φ -state is made at some time instant smaller or equal to t with an accumulated reward smaller or equal to r . The until operator $\Phi \mathcal{U}_{\leq r}^{\leq t} \Psi$ asserts that Ψ is satisfied at some time instant smaller or equal to t and that at all preceding time instants Φ holds; furthermore, the accumulated reward up to time instant t should be at most r . For a CSRL formula Φ , the satisfaction set contains all states that fulfill Φ .

4.2 Buffer occupancy and active sources

As discussed earlier, the bridge B is the bottleneck of the 2 hop ad-hoc network. We therefore study the expected buffer occupancy and the expected number of active sources under different capacity sharing strategies. The expected buffer occupancy can be expressed using the expected reward operator $\mathcal{E}_{\leq r}(\Phi)$, where Φ is just *true*. We choose the reward to be the number of packets currently in the buffer, which is just the level-index in the QBD, so that, we have a level-dependent reward.

The number of expected active sources can also be expressed with the expected reward operator (again with $\Phi = \text{true}$), where we now assign to each state the number of sources that is currently active as reward. In the models the number of active sources is the same in corresponding states of different levels, so that we have level-independent rewards in this case. The expected reward operator requires the computation of a sum over all elements of the satisfaction set $\text{Sat}(\Phi)$. However, for both of the above measures $\Phi = \text{true}$, so that the satisfaction set contains an infinite number of elements. In case of level-independent rewards (expected number of active sources), the QBD structure of the Markov chain allows us to compute, in closed form, the sum of the steady-state probabilities for corresponding states over all levels as $\underline{z}_0 + \underline{z}_1(\mathbf{I} - \mathbf{R})^{-1}$, where \underline{z}_0 and \underline{z}_1 are the steady-state probabilities for the boundary level and the first repeating level. Multiplying these accumulated steady-state probabilities with the corresponding level-independent rewards, allows us to evaluate the expected reward operator. For level-dependent rewards (expected buffer occupancy) the expected reward operator can

be evaluated by exploiting a property of the geometric series, leading to a closed-form expression for the expected reward of the form $\underline{z}_1(\mathbf{I} - \mathbf{R})^{-2}\mathbf{1}$ for the expected buffer occupancy [6]. Following a similar procedure, closed-form expressions can also be derived for the higher moments.

4.3 Time the system spends in a given phase

We now analyze how the steady-state fraction of time the system spends in a given phase changes for the adaptive capacity sharing strategies. As we only allow for a finite maximum of active sources K , an arriving source is “lost” whenever K sources are already active. Thus, we are interested in the fraction of time the system spends in states which correspond to K active sources. Measures concerning the fraction of time the system spends in a given set of states can be expressed by the steady-state operator. It can be evaluated as described in [15] in detail. We just assign the atomic proposition *loss* to all these states that correspond to K active sources and evaluate $\mathcal{S}_{\triangleright p}(\text{loss})$ for the loss probability. To measure the fraction of time each model spends in its different phases, we assign to every state the name of the corresponding system phase as an atomic proposition. In the BRT model we then have *low* in all states of the boundary part, and *high* in all states of the repeating levels. In the SRT model we have *start*, *run* and *clearance* in corresponding states of all levels. As can be seen easily, we have level-independent atomic propositions in all cases.

4.4 Transient analysis of start up

In both adaptive capacity sharing approaches we allow the sources a start up phase, called *low occupancy* in the BRT model and *start up* in the SRT model. In this phase the sources can use most of the available radio capacity. To analyze how long this phase is, we use the time-bounded until operator. Again, we assign the same atomic proposition to all states that belong to the same phase and evaluate $\mathcal{P}_{\triangleright p}(\text{low occupancy } \mathcal{U}_{\leq t}^{\leq t} \text{high occupancy})$ for the BRT model and $\mathcal{P}_{\triangleright p}(\text{start up } \mathcal{U}_{\leq t}^{\leq t} \text{run})$ for the SRT model. These measures express for varying p and t , the distribution of time spend in the first phase. The until operator with only a time constraint can be validated as for CSL, as proposed in [15]. In this study we will consider only the empty state (no sources active and no packets present) as starting state.

4.5 Work done in a given time

We want to analyze how the adaptive capacity sharing strategies influence the amount of work that is done in a given time by the bridge and the sources, respectively. This can be done with the until operator with time and reward bounds. To analyze the work done by the bridge, we assign to each state the amount of radio capacity it currently uses as a reward. In the SRT model the server rate changes with the number of active sources, however, in corresponding states of different levels the server has the same amount

of radio capacity, therefore the reward is level independent. In the BRT model the rate changes with the buffer occupancy, however, the states of the boundary level all have the same reward and the states of the repeating levels also have the same reward. To analyze the work done by the sources, we assign to each state the amount of radio capacity the sources currently occupy. Again this provides us with level-independent rewards for both models.

For the time- and reward-bounded until, $\Phi U_{\leq_r}^t \Psi$, we consider how to check the formula for only one starting state (the all empty state). This allows us to use well-known algorithms for CSRL model checking [7, 1]; since we address a finite time horizon, these uniformization-based algorithms will consider only a finite number of steps, so that only a finite portion of the underlying QBD needs to be taken into account, cf. [15]; in particular, we use the so-called Markovian approximation [8, 7, 4].

4.6 Other CSRL operators

Even though we do not use the full range of CSRL operators, we want to sketch how the *until operator with only a reward bound* and the *instantaneous reward operator* can be validated for QBDs. For checking the until operator with only a reward constraint, e.g., $\Phi U_{\leq_r} \Psi$, for finite CTMCs the *duality theorem* can be used. This theorem states that the progress of time can be regarded as the earning of reward and vice versa [1]. Formulas with only a reward constraint can thus be checked as formulas with just a time constraint on a transformed CTMCs. On QBDs, the Duality Theorem is applicable only in case of level-independent rewards (since the transition rates are rescaled by the reward rates, the QBD structure would be destroyed otherwise). In this case, the QBD can then be checked as stated in [7].

The *instantaneous reward* ($\mathcal{E}_{\leq_r}^t(\Phi)$) operator gives the expected reward at time t . To calculate the transient probabilities in a QBD we always consider only a finite number of steps. That is, the instantaneous reward operator can always be checked for a single starting state s , regardless of the reward structure. To calculate the satisfaction set we distinguish between level-independent and level-dependent rewards. We will eventually find a level from which onwards the transient probabilities do not change anymore. With level-independent rewards the validity of the instantaneous reward operator does not change anymore from this level onwards. We do not know how to check the instantaneous reward operator with level-dependent rewards and an infinite satisfaction set in all cases as the reward modifies the transient probabilities.

To compute, in general, the whole satisfaction set for time- and reward-bounded until properties, i.e., $\Phi U_{\leq_r}^t \Psi$, we have to distinguish between two different reward types. In the case of *level-independent rewards*, the satisfaction set is potentially of infinite size. However, there will eventually be a level from which onwards the validity of the formula will be the same in all corresponding states of the repeating levels. This is just a straightforward extension of the ideas

presented in [15]. In case of *level-dependent rewards*, the satisfaction set will always be finite whenever the reward rates are increasing in the level index.

5 Results

5.1 Buffer occupancy and active sources

All active sources and the bridge have to share a fixed amount of radio capacity. We analyze how the different adaptive capacity sharing strategies influence the expected buffer occupancy and the expected number of active sources. Figure 7 shows five curves representing five different values for the buffer related threshold τ for the BRT model. The curves range over λ from 0.1 to 0.5. For all values of λ , the buffer occupancy is larger for higher thresholds τ . This is due to the fact that for larger values of τ the model stays longer in *low occupancy*, where the sources have exactly the same amount of radio capacity as the bridge and thus have more time to fill the buffer.




Figure 7. Expected buffer occupancy in BRT

The lowest buffer occupancy is achieved for $\tau = 2$, as this setting benefits the bridge most. However, the buffer occupancy cannot be judged without the number of expected sources. In Figure 8 we show the expected number of active sources for five different values of τ , for λ ranging from 0.1 to 0.5. No difference in the number of active sources can be seen, as the differences between the five values of τ are too small. Recapitulating Figure 7, we can say that the buffer related threshold model benefits the bridge and that there is no notable impact of τ on the expected number of active sources.

For the SRT model we show the expected buffer filling for four different values of m over the same range of λ in Figure 9. Note that we use a logarithmic scale on the y -axis, as we observe a much higher buffer occupancy than in the BRT model. Recall that the basic model is included in the SRT model for $m = 1$. Looking more closely, we see that the four depicted curves cross each other. This happens because the parameter m plays a double role: On the one hand, m serves as a source-related threshold for switching between the three phases of the SRT model. On the other




Figure 8. Expected number of active sources in the BRT model

hand, m times as much capacity is assigned to the bridge than to each of the sources. Thus, for larger m the buffer occupancy drops as more radio capacity is assigned to the bridge, however, for larger m the bridge obtains its larger share of radio capacity later. Thus, the backlog of packets can only be removed after m sources have been active at the same time. The probability for m sources being active at the same time depends on λ , therefore different values of m might be optimal for different λ .




Figure 9. Expected buffer occupancy in SRT

In Figure 10 we show the expected number of active sources, again for four different values of m . The number of expected active sources is lower than in the BRT model. Again, we see that the curves cross each other due to the trade-off between the two effects that play a role due to m . Comparing these curves to the BRT model in Figure 8, the expected number of sources in the SRT model differs, because the amount of radio capacity the sources can use as soon as the model leaves *start up*, heavily depends on m . For small values of λ , the system stays longer in *start up*, where the sources get the same share of radio capacity, independent of m , so that the curves are still close to each other. For higher values of λ the probability to be in the *run phase* or *clearance phase* rises and the curves start to diverge.




Figure 10. Expected number of active sources in the SRT model

To achieve a reasonable buffer occupancy with the SRT model, λ should stay below ≤ 0.3 , however, the sources will then be able to transmit much faster as in the BRT model and the expected number of active sources in this setting is below one. Concluding we can say, that the SRT model benefits the sources much more than the bridge. This behavior is due to the double role that the parameter m plays. We did not expect this a priori. A variant of the SRT model with two parameters, one for the threshold and one for the adaption of the capacity distribution, should be considered and studied in the future. Comparing the buffer occupancy and the number of active sources for the BRT and the SRT model, the BRT model benefits the server, whereas the choice of τ has no impact on the number of active sources, while the SRT model benefits the sources and the choice of m has a huge impact on the buffer occupancy.

5.2 Loss probability

To analyze the restriction to a maximum of K active sources, we compute the probability that an arriving source is lost. In Figure 11 we show the probability that all sources are active over the range for λ from 0.1 to 0.5 for both models. The upper curve shows the loss probability in the BRT model and the lower four curves show the loss probability in the SRT model. Again the curves for the BRT model do not differ for the different values of τ . For small values of λ the loss probability is zero, as the source arrival rate is too small. Only when the models leave *low occupancy* and *start up* the loss probability increases as the sources then have less radio capacity and the expected number of active sources grows. In essence, we see the same behavior as for the expected number of active sources, as these measures are closely related. As the loss probability is low for the depicted setting, we can conclude that the restriction to a maximum of K sources is not a severe one. It can easily be seen, that the SRT model benefits the sources more than the BRT model does.




Figure 11. Probability for 10 active sources

5.3 Work done in a given time

To further analyze the influence of the capacity sharing strategies we calculate the accumulated work done at time t for given $\lambda = 0.4$, with the empty state (no sources active and no packets present) as starting state. In Figure 12 we show the probability that the accumulated reward (the work done by the bridge) at time t is at most 10 packets, for four different values of τ in the BRT model. For $\tau = 2$ the probability for the accumulated reward to be at most 10 packets is always highest, which means that the accumulated reward is smaller than for higher τ . In case of higher thresholds τ the probability for the reward to be at most 10 decreases and for $\tau \geq 8$ the curves are almost the same. Thus, for higher thresholds the accumulated reward is higher. Clearly, the bridge can only forward packets which it received beforehand from the sources. For a low threshold the sources do not have much time to fill the buffer in the *low occupancy* phase, so that, when the bridge gets half of the available radio capacity it sometimes cannot use it, as no packets are present. So, in order not to waste radio capacity, the threshold τ should be high enough, so that the sources can fill the buffer well in the *low capacity* phase. In Figure 13 we observe again, that the behavior of the sources in the BRT model does not change for the different thresholds. Thus the work done at a given time t by the sources does not influence the choice of τ in the BRT model.

Figure 14 shows the probability that the accumulated work done by the bridge in the first 10 seconds after the system has been empty is at most 10 packets in the SRT model. This probability heavily depends on the threshold m . For $m = 8$ the bridge has almost no work to do, as the probability for 8 sources to become active within the first 10 seconds is quite small. The smaller the value of m is, the more work can be done in the first 10 seconds. Figure 15 shows the same probability for the work done by the sources. We observe that the behavior of the bridge and the sources is pretty much the same. The accumulated amount of work done by the sources decreases for higher thresholds m , as the probability for the accumulated work to be below 10 increases.




Figure 12. Probability that the accumulated work done by the bridge at time t is at most 10 in BRT.




Figure 13. Probability that the accumulated work done by the sources at time t is at most 10 in the BRT model.




Figure 14. Probability that the accumulated work done by the bridge at time t is at most 10 in SRT.




Figure 15. Probability that the accumulated work done by the sources at time t is at most 10 in the SRT model.




Figure 16. Fraction of time BRT model spends in *low occupancy*




Figure 17. Fraction of time SRT model spends in *start up*

5.4 Time the system spends in a given phase

We show the steady-state fraction of time the BRT model spends in *low occupancy* in Figure 16 for five different values of τ . The curves are straightforward, for higher values of τ the probability to be in *low occupancy* is higher and for growing values of λ the probability declines, as a higher arrival rate for sources induces a higher buffer occupancy. Again the curves stay close to each other. In Figure 17 we show the steady-state fraction of time the SRT model spends in *start up*. The basic behavior of the curves resembles Figure 16, however, the curves are spread more widely due to the influence of m .

5.5 Transient analysis of the start up phase

Figure 18 depicts the transient probability to reach *high occupancy* within time t for different values of τ and a fixed value for $\lambda = 0.4$ in the BRT model. The starting point for this transient analysis is the empty system with no active sources and no packets present. For lower thresholds τ the probability to reach the next phase is highest as the threshold is crossed sooner. For $\tau \geq 16$ the probability to reach the next phase in ten seconds is still very low. The same can be observed in Figure 19, where we show the transient probability to reach *run* within time t for different values of m in the SRT model. The low probabilities, especially for high thresholds, show that the phases *low occupancy* in the BRT model and *start up* in the SRT model last quite a while, which is useful for the sources, as they can start transmitting with a greater share of radio capacity.




Figure 18. Transient probability to reach *high occupancy* in the BRT model

6 Conclusions

Recently alternative capacity sharing strategies in IEEE 802.11 (wireless LAN) have been made possible [9]. In this paper, we build a versatile infinite-state Markov reward model to study such adaptive capacity sharing strategies.




Figure 19. Transient probability to reach run in the SRT model

They are modeled as iSPNs and then automatically transformed into the underlying QBD. The behavior of the different capacity sharing strategies is then analyzed by means of CSRL model checking. We adapt techniques for CSRL model checking of finite CTMC with rewards to the infinite case and explain how CSRL model checking is done on such infinite-state QBDs with rewards. Throughout this paper we have assumed at various points that certain time durations obey a negative exponential distribution. This has been done for modeling convenience only, and is not a fundamental necessity of the approach. Indeed, we can use phase-type distributions instead, at all times. The modeling and solution procedure will not change for that, but just become more time consuming. Note that we cannot analyze source-based behavior, as *flow transfer time*, as we cannot distinguish between sources.

For the transient measures we choose the empty system as a starting point, however, it is also possible to look at different starting states with the same techniques. For steady-state measures the starting point does not matter. The analysis of the two adaptive capacity sharing strategies clearly shows that the buffer-related threshold benefits the bridge and the source-related threshold benefits the sources. We found that the parameter m of the SRT model plays two different roles, on the one hand it serves as a threshold and on the other hand as a multiplier for the capacity that is assigned to the server. This leads to a complex behavior in the SRT model. Our modeling approach helps in deciding under which circumstances which adaptive capacity sharing strategy is most appropriate. In the future we will study a variant of the SRT model with two parameters (instead of just m), and will present, in a separate paper, all the details of the employed CSRL model checking algorithms for QBDs.

Acknowledgments

We thank J.L. van den Berg and his co-authors for providing us with an early version of their paper [17] as well as for discussing their model with us. Thanks also go

to Alexander Bell for his help with the state-space generation. Anne Remke is supported through the MC=MC project (612.000.311) and Lucia Cloth is supported through the MOCS project (642.000.505), both financed by the Netherlands Organization for Scientific Research (NWO). We thank the reviewers for their constructive and helpful criticism.

References

- [1] C. Baier, B. R. Haverkort, H. Hermanns, and J.-P. Katoen. On the logical characterisation of performability properties. In *ICALP '00*, pages 780–792, London, UK, 2000. Springer-Verlag.
- [2] A. Bell. *Distributed evaluation of stochastic Petri nets*. PhD thesis, RWTH Aachen, 2004.
- [3] G. Bianchi. Performance analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications*, 18:535–547, 2000.
- [4] L. Cloth. *Model checking algorithms for Markov Reward Models*. PhD thesis, University of Twente, 2005.
- [5] J. Cohen. The multiple phase service network with Generalized Processor Sharing. *Acta informatica*, 12:245–284, 1979.
- [6] B. Haverkort. *Performance of Computer Communication Systems*. John Wiley & Sons, 1998.
- [7] B. Haverkort, L. Cloth, H. Hermanns, J. Katoen, and C. Baier. Model checking performability properties. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN'02)*, pages 102–112. IEEE Computer Society, 2002.
- [8] B. R. Haverkort, H. Hermanns, J.-P. Katoen, and C. Baier. Model checking CSRL-specified performability properties. In *Proceedings of the 5th International Workshop on Performability Modeling of Computer and Communications Systems (PMCCS'01)*, pages 105–109, 2001.
- [9] IEEE 802.11e/D13.0. Wireless LAN Medium Access Control (MAC) and Physical layer (PHY) specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS). *draft supp. to IEEE 802.11 std.*, 2005.
- [10] J. Kurose and K. Ross. *Computer Networking*. Addison-Wesley, 2005.
- [11] G. Latouche and V. Ramaswami. *Introduction to Matrix Analytic Methods in Stochastic Modeling*. ASA-SIAM, 1999.
- [12] R. Litjens, R. Roijers, J. van den Berg, R. Boucherie, and M. Fleuren. Analysis of flow transfer times in IEEE 802.11 wireless lans. *Annals of Telecommunications*, 59:1407–1432, 2004.
- [13] M. Neuts. *Matrix Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Johns Hopkins University Press, 1981.
- [14] A. Ost. *Performance of communication Systems*. Springer, 2001.
- [15] A. Remke, B. Haverkort, and L. Cloth. Model checking infinite-state Markov chains. In *TACAS 2005, LNCS 3440 proceedings*, pages 237–252. Springer, 2005.
- [16] J. Schiller. *Mobile Communications*. Addison-Wesley, 2003.
- [17] J. van den Berg, M. Mandjes, and F. Roijers. Performance modeling of a bottleneck node in an IEEE 802.11 ad-hoc network. In *To appear in: Proceedings 5th International Conference on AD-HOC Networks and Wireless*, 2006.