

Modeling Media Access in Embedded Two-Flow Topologies of Multi-hop Wireless Networks

Michele Garetto, Jingpu Shi, and Edward W. Knightly
Department of Electrical and Computer Engineering
Rice University
Houston, TX 77005
{garetto,jingpu,knightly}@rice.edu*

ABSTRACT

In this paper, we decompose a large- or small-scale multi-hop wireless network into embedded subgraphs, each consisting of four nodes and two flow pairs. We systematically study all twelve possible topologies that arise according to whether the different nodes are in radio range of each other. We show that under both a random spatial distribution of nodes and random waypoint mobility with shortest-path routing, a critical and highly probable scenario is a class in which the channel state shared by the two flows is not only incomplete (i.e., the graph is not fully connected), but there is also *asymmetry* in the state between the two flows. We develop an accurate analytical model validated by simulations to characterize the long-term unfairness that naturally arises when CSMA with two- or four-way handshake is employed as a random access protocol. Moreover, we show that another key class of topologies consists of incomplete but *symmetric* shared state. We show via modeling and simulations that in this case, the system achieves long-term fairness, yet endures significant durations in which one flow dominates channel access with many repeated transmissions before relinquishing the channel. The model predicts the time-scales of this unfairness as a function of system parameters such as the maximum retransmission limit.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Design, Performance Evaluation

Keywords

CSMA, CSMA/CA, Fairness, Analytical Modeling

*This research is supported by NSF ITR Grants ANI-0331620 and ANI-0325971 and by Intel Corporation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'05, August 28–September 2, 2005, Cologne, Germany.
Copyright 2005 ACM 1-59593-020-5/05/0008 ...\$5.00.

1. INTRODUCTION

To better understand the performance characteristics of large- or small-scale multi-hop wireless networks, we begin with the basic building blocks of pairs of contending one-hop flows. If all four nodes of the two flows are within radio range, or if the two senders are within radio range, a CSMA protocol with two- or four-way handshake obtains fair performance that can be accurately predicted using existing models [4].

Unfortunately, all remaining two-flow topologies encounter performance problems, including severe long-term unfairness and short-term unfairness on timescales of up to seconds. While the existence of some problematic topologies is well known (see [2, 11, 12, 17] for example), this paper is the first to systematically and comprehensively identify and study via analytical modeling and simulations all twelve possible configurations of two flows in a random access environment. We compute the likelihood of each scenario occurring in a random topology, develop models that accurately predict each flow's short- and long-term performance in each scenario where senders are not in range, and predict the impact of key system parameters such as the maximum retransmission limit.

In particular, our contributions are as follows. First, we enumerate the twelve distinct scenarios and classify them into three groups based on their geometric properties: Senders Connected (SC), Asymmetric Incomplete State (AIS), and Symmetric Incomplete State (SIS). We develop a spatial model that predicts the likelihood of the scenarios and groups when nodes are placed randomly. Moreover, we use simulations with random waypoint mobility and routing protocols to further characterize scenario likelihood. We show that the problematic AIS and SIS groups should not be viewed as rare “corner cases,” but rather characterize topologies that occur with high frequency.

Second, we perform an extensive set of simulations to characterize the short- and long-term performance in each of the twelve scenarios. The experiments quantify the severe long-term unfairness of the AIS class and the seconds time-scale short-term unfairness of the SIS class. We describe the MAC mechanisms that lead to each class' performance limits. Moreover, we show how random waypoint mobility allows flows to alternate among classes, resulting in improved long-term performance but velocity-controlled epochs of poor performance in the AIS and SIS classes.

Third, we develop an accurate analytical model for the AIS class. We begin with a general “decoupling model” of IEEE 802.11 which provides key temporal embedding points that are critical for development of a Markov model for topologies in which nodes are not fully connected and performance is asymmetric. We generalize the model to incorporate flows that are not continuously backlogged and use simulations to validate the model's accuracy.

Finally, we develop a model to characterize the short-term unfairness that arises in the SIS class. Because this class achieves long-term fairness, the steady state distribution is not of interest. Instead, we derive an expression for the expected time for the system to switch between one flow dominating and the other flow dominating. The model incorporates effects of key system parameters such as the maximum retransmission limit on the timescales of unfairness. Simulation results indicate that the model is able to accurately predict transition times ranging from 10's of msec to 1 sec, with the range accounted for by both the maximum retransmission limit and whether or not RTS/CTS is used.

The remainder of this paper is organized as follows. In Section 2 we identify the twelve topologies for the flow pairs, classify them into three groups, and conduct a spatial analysis to determine their likelihood in a random network. In Section 3, we study via simulations the impact of protocols on the fairness of the three classes and describe the protocol-driven origins of this behavior. In Section 4 and Section 5, we model the AIS and SIS classes respectively. Finally, in Section 6 we discuss related work and in Section 7 we conclude.

2. TOPOLOGY & NETWORK GEOMETRY

In this section, we first identify all feasible topologies in which exactly two directional flows are communicating. Next, using a grouping of these topologies, we perform a spatial analysis that characterizes the likelihood of each of these topologies occurring under random node placement. Finally, we compare scenario likelihood via the model and simulations that incorporate mobility.

2.1 Two Flow Topologies

We consider four stations that are communicating pairwise, where two senders are transmitting a one-way data flow to their two respective receivers. When two stations are within radio range of each other (i.e., the received SNR is above the carrier sense threshold) we refer to them as having a connection or link between them. Depending on the distances and propagation paths between the four stations, a link may or may not be established between the two flows.

Denote station A as the sender for flow A and station B as the sender for flow B , and stations a and b as their respective receivers. Links are named by using the names of the stations that they interconnect. There are four possible inter-flow links: AB , ab , aB and Ab . In a general topology, each one of these links may be present or not, yielding $2^4 = 16$ different scenarios. Notice that links Aa and Bb are always present, given that senders are connected with their respective receivers. However, scenarios where link Ab exists and link aB does not are symmetric to those where aB exists and Ab does not. After ruling out four such cases, we have twelve distinct scenarios depicted in Figure 1. For example, Scenario 1 in Figure 1 depicts the case in which the two flows are out of radio range and hence operate independently.

Omitting Scenario 1 which is trivial to analyze, we classify the remaining eleven scenarios into three groups as follows.

- Senders Connected (SC) - Scenarios 2-7, in which link AB is present.
- Asymmetric Incomplete State (AIS) - Scenarios 11 and 12, in which senders are disconnected, and only one of links Ab or aB is present (we assume it is always aB), resulting in asymmetric connection between the two flows.
- Symmetric Incomplete State (SIS) - Scenario 8, 9 and 10, in which senders are disconnected, and either both Ab and




Figure 1: Twelve scenarios for two flows sharing a wireless channel.

aB are present or neither is present (in the latter case ab is present) resulting in symmetric connection between the two flows.

We will show that the flow pairs obtain dramatically different performance profiles according to which of the three groups represent their topology. We demonstrate this using simulations for both two-way and four-way handshakes in Section 3 and via analytical models in Sections 4 and 5.

2.2 Spatial Analysis

Now we develop a model to compute the probability that each scenario in Figure 1 occurs in a random graph. Our approach is to view a scenario as three joint events, and to conduct a spatial analysis to compute the probability of each of these events occurring. We assume the four stations are uniformly distributed in the network and that the size of the network is large enough so that border effects are negligible. We describe in detail only the derivation for Scenario 11. The same approach can be applied to predict the occurrence probabilities of the other scenarios.

We consider a simplified propagation model in which r is the radio range of a station.¹ Let s denote the physical area (size) of the network and $T(\cdot)$ represent a region of the plane satisfying certain geometric conditions. In particular, $T(A)$ indicates the region within radio range of station A and $T(\bar{A} \cap B)$ represents the region within radio range of B but outside that of A . The distance between station A and a is denoted by d_1 . Similarly, the distance between station B and b is denoted by d_2 . The coordinates of station z are denoted by x_z and y_z , respectively, where $z \in \{A, a, B, b\}$.

Node placement for Scenario 11 is shown in Figure 2, where the angle β indicates whether station b is within $T(A \cup a)$. Scenario 11 can be decomposed into the following three events: (i) $d_1 \leq r$; (ii) station B is within $T(\bar{A} \cap a)$ given that the first event occurs; (iii) $d_2 \leq r$ and station b is within $T(\bar{A} \cap \bar{a})$ given that the first two events occur. Note that due to the symmetry of region $T(\bar{A} \cap a)$, we only need to consider cases where $y_B \geq 0$ and then apply a multiplicative factor of 2.

We now compute the probability of each of the three events occurring. In event (i), the probability of station a falling into a doughnut area comprised in the interval $[d_1, d_1 + \Delta d_1]$ with A

¹For simplicity, we assume that the transmission range is equal to the sensing range for all stations. We remove this assumption later.

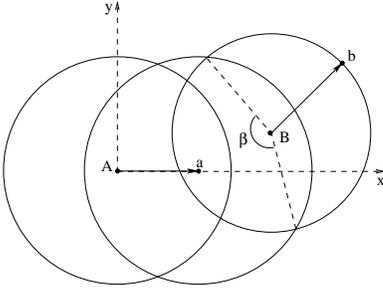


Figure 2: Node placement and topology of Scenario 11. The two circles centered at A and a indicate respective transmission range of station A and a . The circle centered at B indicates possible positions of station b given the distance between B and b is fixed.

at the center is given by

$$\begin{aligned} \lambda_1 &= \frac{2\pi(d_1 + \Delta d_1)^2 - 2\pi d_1^2}{s} \\ &= \frac{2\pi\Delta d_1^2 + 4\pi d_1\Delta d_1}{s}. \end{aligned} \quad (1)$$

When $\Delta d_1 \rightarrow 0$ we can neglect the second order term, obtaining

$$\lambda_1 = \frac{4\pi d_1\Delta d_1}{s}. \quad (2)$$

In event (ii), the probability of station B falling into a small square region in $T(\bar{A} \cap a)$ defined by the interval $[x_B, x_B + \Delta x_B]$ on the x axis and $[y_B, y_B + \Delta y_B]$ on the y axis is given by

$$\lambda_2 = \frac{\Delta x_B \Delta y_B}{s}. \quad (3)$$

In event (iii), the joint probability that the distance between station B and b is within the interval $[d_2, d_2 + \Delta d_2]$ and station b is within $T(\bar{A} \cap \bar{a})$, is given by

$$\begin{aligned} \lambda_3 &= \frac{1 - \beta}{2\pi} \times \frac{2\pi(d_2 + \Delta d_2)^2 - 2\pi d_2^2}{s} \\ &= \frac{1 - \beta}{2\pi} \times \frac{2\pi\Delta d_2^2 + 4\pi d_2\Delta d_2}{s}, \end{aligned} \quad (4)$$

where β is the angle shown in Figure 2. Neglecting the second order term, this equation becomes

$$\lambda_3 = \frac{1 - \beta}{2\pi} \times \frac{4\pi d_2\Delta d_2}{s}, \quad (5)$$

when $\Delta d_2 \rightarrow 0$.

From Equations (2), (3) and (5), the probability that Scenario 11 occurs given x_B, y_B, d_1 and d_2 is

$$\begin{aligned} p'_{11} &= \lambda_1 \times \lambda_2 \times \lambda_3 \\ &= \frac{4\pi d_1\Delta d_1}{s} \times \frac{\Delta x_B \Delta y_B}{s} \times \frac{(1 - \beta)}{2\pi} \times \frac{4\pi d_2\Delta d_2}{s} \\ &= 8\pi \frac{d_1\Delta d_1\Delta x_B\Delta y_B d_2\Delta d_2(1 - \beta)}{s^3}. \end{aligned} \quad (6)$$

Finally the probability of Scenario 11 occurring is

$$\begin{aligned} p_{11} &= \int_0^r \int_0^r \int_{\frac{d_1}{2}}^{r+d_1} \int_{f_1(d_1, x_B)}^{f_2(d_1, x_B)} \\ &2 \times p'_{11} d(y_B) d(x_B) g(d_2) d(d_2) g(d_1) d(d_1), \end{aligned} \quad (7)$$

where $g(\cdot)$ is a given probability density function of the distance between a transmitter and a receiver. The x coordinate of any point

within $T(\bar{A} \cap B)$ lies in the interval $[\frac{d_1}{2}, r + d_1]$, which explains the bounds of the third integral.

To solve Equation (7), we need to compute β , $f_1(d_1, x_B)$ and $f_2(d_1, x_B)$. As shown in Figures 3 and 4, computations for β , $f_1(d_1, x_B)$ and $f_2(d_1, x_B)$ are different when station B falls into different regions of $T(\bar{A} \cap a)$. Therefore, we need to compute p_{11} in different regions and sum the results.

To compute β , we consider Figure 3 which indicates that if d_2 is fixed, $T(\bar{A} \cap a)$ can be divided by the dashed circle into two areas, denoted by $R1$ and $R2$ respectively. β is computed differently in $R1$ than in $R2$ because the circle centered at B intersects with other circles differently in different regions. However, the computation for β is trivial in either $R1$ or $R2$ although its expression is tedious, so we omit it here.




Figure 3: Different locations of station B lead to different computations for β .

To determine $f_1(d_1, x_B)$ and $f_2(d_1, x_B)$, we further divide region $R1$ and $R2$ in Figure 3 into smaller regions. Let Γ denote such a value that when $d_2 = \Gamma$, $x_Q = d_1$, where x_Q is the x axis of the intersection Q shown in Figure 4. When $d_2 \leq \Gamma$, $T(\bar{A} \cap a)$ is divided into five areas labeled by I, II, III, IV and V as shown in Figure 4(a). When $d_2 > \Gamma$, $T(\bar{A} \cap a)$ is divided into five different areas shown in Figure 4(b). $f_1(d_1, x_B)$ and $f_2(d_1, x_B)$ can be determined within each of the five regions and Equation 7 can be solved numerically. Similar to β , $f_1(d_1, x_B)$ and $f_2(d_1, x_B)$ are easy but tedious to obtain, thus we omit their expressions here.

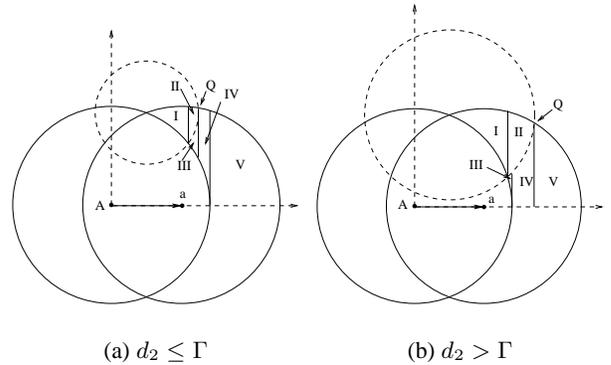


Figure 4: Different value of d_2 decides how to further divide $T(\bar{A} \cap a)$ into different regions.

A similar analysis can be conducted for computing the probabil-

ity that each of the other 10 scenarios in Figure 1 occurs, denoted by p_i , where $i \in \{2, 3, 4, \dots, 12\}$.

2.3 Comparing Scenario Likelihood

We now compare the likelihood of the three classes occurring. In the context of a multi-hop wireless network, it is especially interesting to evaluate the probability of each of the three classes occurring between direct neighbors as a function of hop distance. We assume that the distance between transmitters and receivers is constant, and we compute p_i as a function of d , where $d = d_1 = d_2$.

We can obtain results independent of flow density by computing the conditional probability that a particular Scenario i occurs, given that some connection exists between the two flows, i.e., they are not isolated. This is then the probability of Scenario i occurring conditioned on the event that any of the scenarios except Scenario 1 occurs. Thus it is given by p_i/p , where $p = \sum_{i=2}^{12} p_i$. Results validated by Monte Carlo simulations are shown in Figure 5 as a function of the normalized hop distance, which is the actual distance between a sender and a receiver divided by the radio range. Simulations are run dropping two pairs of nodes at a given distance from each other uniformly at random in a square area with wrap-around (to avoid border effects), and seeing which case they map into (excluding Scenario 1). By so doing, results are insensitive of area size and node density.

We observe that Scenario 11 (belonging to the AIS class and one with problematic performance) emerges as the dominating scenario when the distance between sender and receiver increases.

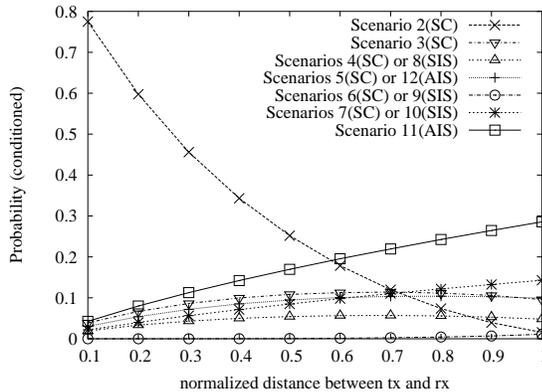


Figure 5: Probability of each of the eleven scenarios occurring as a function of normalized hop distance.

Denote the respective probabilities of the Senders Connected (SC), Asymmetric Incomplete State (AIS), and Symmetric Incomplete State (SIS) classes occurring as ϵ_{SC} , ϵ_{AIS} and ϵ_{SIS} . We have $\epsilon_{SC} = \{p_2 + p_3 + p_4 + p_5 + p_6 + p_7\}/p$, $\epsilon_{SIS} = \{p_8 + p_9 + p_{10}\}/p$, and $\epsilon_{AIS} = \{p_{11} + p_{12}\}/p$. The three probabilities above are reported in Figure 6 as a function of the normalized hop distance.

Results indicate that the AIS and SIS classes account for a significant fraction of all possible scenarios. Increasing the hop distance, the likelihood of class AIS approaches that of the SC class despite accounting for two vs. six scenarios. The ratio between AIS and SIS probabilities is about 2.

In multi-hop wireless networks, the distribution of hop distances depends on the routing protocol deployed. From Figure 6, it is clear that routing protocols can have a significant impact on the probability of each class occurring. To evaluate this impact, we conducted simulation experiments to measure the hop distance dis-




Figure 6: Probabilities of the two groups, assuming identical transmission and sensing range.

tribution resulting from the operation of current routing protocols in large-scale networks with mobility. In our simulations, we consider 300 stations randomly deployed in a $2000\text{ m} \times 2000\text{ m}$ area. A random waypoint model [22] is used to simulate mobility, and connections are randomly established among the nodes. The routing algorithm considered is Distance Sequence Distance Vector (DSDV) (we tested other routing algorithms and obtain similar results). To compute the normalized hop distance distribution, we calculate the hop distance of every link traversed by each packet and divide it by the maximum transmission range. After averaging the results of several simulation runs, we obtained the distribution of normalized hop distances reported in Figure 7. Compared to the results in




Figure 7: Hop distance distribution in a multi-hop wireless network.

[14], in which the impact of routing protocol is not considered, our simulation experiments suggest that in a multi-hop wireless network, about 40 percent of normalized distances are within the $[0.9, 1]$ interval. This can be explained by the fact that routing protocols select minimum-hop paths to reach the destination. Combined with the results in Figure 6, this means that current routing protocols make the scenarios of classes AIS and SIS more likely to occur in a random network, by favoring larger hop distances when choosing next hops. The occurrence probabilities of classes SC, AIS and SIS under routing can be computed from Figures 6 and 7, assuming random distribution of two-hop flows.

In the previous analysis, we assumed identical transmission and sensing range for all stations. In a real network, however, sens-

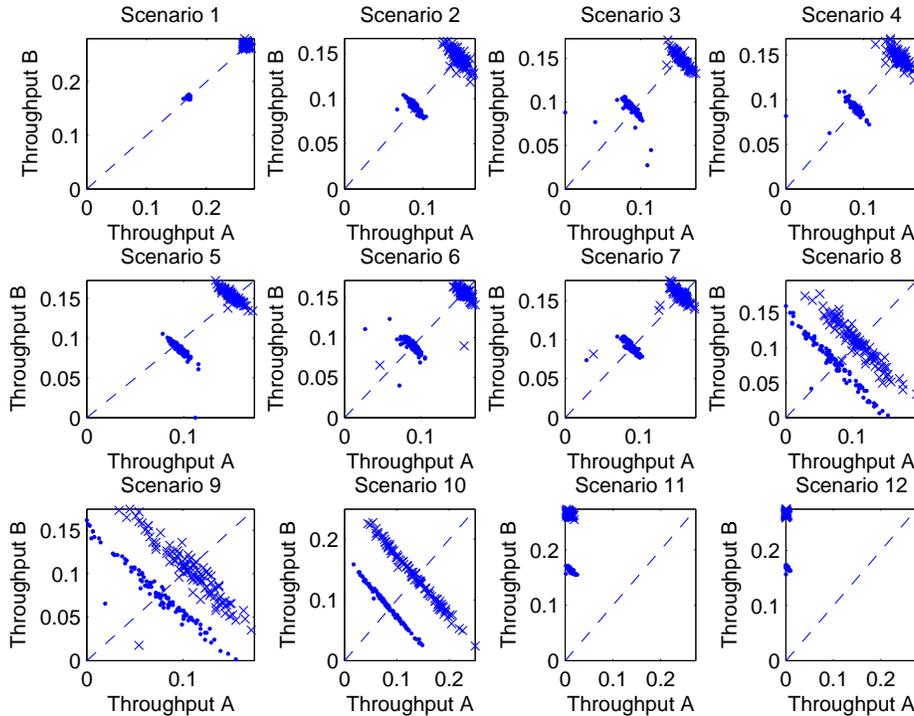


Figure 9: Short-term normalized throughput for the twelve scenarios. The ‘x’ marks correspond to two-way handshake simulations; the dot marks correspond to four-way handshake simulations.

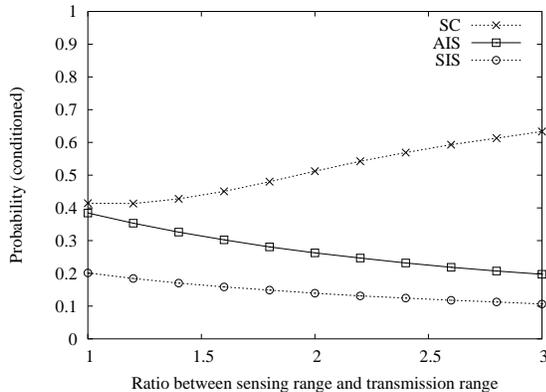


Figure 8: Probabilities of the two groups, in case of different transmission and sensing range.

ing range and transmission range of a station are usually different. In order to examine the impact of this factor, we vary the sensing range of the nodes, while keeping the transmission range fixed. We assume the distance between senders and receivers is equal to the transmission range. Results are shown in Figure 8, as a function of the ratio between sensing range and transmission range. Although the occurrence probabilities of the AIS and SIS classes decrease when the ratio between sensing range and transmission range increases, the probability of class AIS occurring is significant over the whole range of values encountered in practice. The ratio between AIS and SIS is always around 2.

3. IMPACT OF MAC PROTOCOLS

In this section we qualitatively assess by simulation experiments the performance of IEEE 802.11 in the two-flow subgraphs identified by the spatial analysis of Section 2. Our objective is to identify the critical performance issues that can arise when a typical CSMA-based MAC protocol is employed to arbitrate channel access between the two flows.

The main channel access mechanisms implemented in the IEEE 802.11 standard [7] are the “basic access” two-way handshake without RTS/CTS and the four-way handshake with RTS/CTS. We study how both access methods perform in each of the twelve possible scenarios comprising two flows.

We consider a data rate of 11 Mbps, and a fixed packet size of 500 bytes. Both flows are continuously backlogged with UDP traffic. We measure the average throughput of each flow during consecutive periods of 0.4 seconds. Each simulation experiment lasts 20 seconds, thus we collect 50 throughput samples for each flow. In Figure 9, we plot the normalized throughput of flow B versus the normalized throughput of flow A in each measurement interval. The “x” marks correspond to the two-way handshake access method; the dot marks correspond to the four-way handshake. Dashed lines represent ideal fairness, since they consist of the points where the throughput of the flows are equal.

Omitting Scenario 1 in which the flows are isolated due to spatial reuse, we make the following observations on the other 11 scenarios.

- For the Senders Connected (SC) class consisting of Scenarios 2-7, most of the throughput points reside close to the dashed line, indicating that neither short-term nor long-term fairness problems exists.

- For the Asymmetric Incomplete State (AIS) class consisting of Scenarios 11 and 12, the variance of the throughput points is small. However, these points largely deviate from the fairness line, indicating severe unfairness at all time scales and flow starvation.
- For the Symmetric Incomplete State (SIS) class consisting of Scenarios 8, 9 and 10, the throughput points are symmetrically scattered around the dashed line, indicating short-term unfairness, but in the long term, the throughputs of the two flows are the same.

Since the SC class does not encounter fairness problems and because its performance can be analyzed with existing techniques [4], we will not further consider it. We focus instead on the problematic AIS and SIS classes. In Sections 4 and 5 we will develop a detailed analysis of these two classes. Here we provide a qualitative explanation of the behavior observed in the simple experiments of Figure 9.

Origins of AIS Long-Term Unfairness. Figure 10 shows example topologies of Scenarios 11 and 12 of the AIS class. The core property of the AIS class is the asymmetric view of the channel state possessed by the two flows. When transmitters are not in range of each other, channel state information is necessarily incomplete because transmitters cannot sense when the other flow is transmitting. This lack of information affects the two AIS flows in very different ways because of the asymmetry of the topology. In particular, flow *A* lacks the necessary information to compete fairly with flow *B*, while flow *B* does not suffer from the incomplete channel state information. This disparity is due to the fact that sender *A* does not sense any packets belonging to flow *B*, and consequently, completely ignores the activity of the other flow. On the other hand, sender *B* can hear the control packets sent by node *a* (CTS and/or ACK), and hence can detect the activity of the other flow. While sender *B* knows exactly when to start contending for the channel, sender *A* has to discover an available time-slot randomly, without any coordination with sender *B*. This fact results in many transmission attempts of sender *A* without any response back from receiver *a*, most often because *A* attempts to transmit in the middle of a transmission of flow *B*, when receiver *a* cannot receive correctly the packets sent by *A*, or is not able to reply. Consequently, sender *A* is forced to timeout and to repeatedly double its contention window. As a result, the probability of flow *A* capturing the channel is significantly smaller than that of flow *B*.




Figure 10: Example topologies of the AIS class.

Origins of SIS Short-Term Unfairness. Figure 11(a) reports an example topology of Scenario 8 of the SIS class. The classic “hidden terminal” problem [2, 12], depicted in Figure 11(b), is a special case of Scenario 8 when the two flows have the same receiver. The core property of the SIS class is the symmetric view

of the channel state possessed by the two flows while the channel state information is incomplete. This property results in short-term unfairness but long-term fairness. The origin of the short-term unfairness lies on the binary exponential backoff mechanism coupled with the large packet loss probability that characterizes all scenarios of the SIS class. The large packet loss probability is due to incomplete channel state information. In particular, a sender does not sense the other sender’s activity, thus it can start transmitting the first packet of the two-way or four-way handshake while the other sender is also attempting to transmit. Indeed, a sender does not stop decrementing the backoff counter as soon as the other sender starts transmitting. This fact significantly increases the collision probability of the flows. After experiencing a collision, a source doubles its contention window, thus reducing the chances of attempting a new transmission in the next available slot. On the contrary, a source resets its contention window to the minimum value after a successful transmission, increasing the likelihood of a new transmission attempt in the near future. Thus, in all scenarios of the SIS group, the system endures significant durations in which one flow dominates channel access with many repeated transmissions, while the other flow is forced to repeatedly double its contention window significantly reducing the chance to seize the channel. However, this problem affects the two flows equally, because the geometric relationships in the scenarios of the SIS group are symmetric. The two flows alternate capturing the channel and dominating over the other flow. Therefore, flow pairs belonging to the SIS group do not suffer from long-term unfairness (essentially because of symmetry). Instead, simulations show significant short-term unfairness, as illustrated in Figure 9, which is clearly undesirable as it can adversely affect delay-sensitive applications, such as voice.

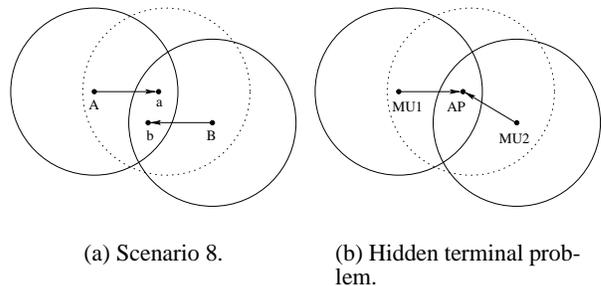


Figure 11: Example topologies of the SIS class. (b) is a special case of (a), when two Mobile Units (MU1 and MU2) are sending packets simultaneously to the same Access Point (AP) in infrastructure mode, as encountered in the classic “hidden terminal” problem.

Impact of Mobility. In a network where stations move randomly, a flow is expected to belong to different local subgraphs constantly changing over time. Will this mobility alleviate the unfairness problem observed in Figure 9? We conducted a simulation experiment where 40 stations move according to the random waypoint model in a 1000x1000 region with a speed uniformly distributed in a [7, 15] m/s interval. 20 connections are established between randomly chosen pair of stations. Figure 12 shows that although mobility averages out unfairness over time-scales of 120 seconds, severe unbalance is observed during time windows of 10 seconds, in which two dominating flows starve all of the other flows. This means that severe fairness problems still exist over time scales associated with the speed of the nodes.




Figure 12: Flow throughput comparison between a 10 second snapshot and a 120 second snapshot.

4. ASYMMETRIC INCOMPLETE STATE

In this section we develop an analytical model to study the behavior of flow pairs belonging to the Asymmetric Incomplete State (AIS) group, comprising Scenarios 11 and 12 of Figure 1. Examples of topologies that satisfy the geometric properties of these two scenarios are illustrated in Figure 10(a) and 10(b). In both cases, flow *B* achieves a significantly higher throughput as compared to flow *A* for CSMA both with and without RTS/CTS (see Figure 9). The only difference between the two cases is that in scenario 12 the receivers are in radio range of each other, whereas this is not true in scenario 11. Below, we show how this topology difference affects performance.

Our objective is to analytically compute the throughput of the two flows to both characterize the root cause of the starvation of flow *A* and to evaluate the impact of key system parameters on the extent of the starvation.

The remainder of the section is organized as follows. In Section 4.1 we introduce a general model of the behavior of a backlogged source employing the standard 802.11 DCF. This model will be applied in Section 4.2 to the particular scenarios of the AIS group. In Section 4.3 the analysis is extended to the case of non-backlogged sources in order to assess the impact of the starvation problem in more general network scenarios. Finally, numerical results and simulations are presented in Section 4.4.

4.1 General Decoupling Model of an 802.11 Station

Our modeling framework for the AIS group contrasts with existing techniques (e.g., [4]) in that we account for the fact that, in a general topology, the channel state as perceived by a station can be different from node to node. In [4], all stations are assumed to be in range of each other so that they share a common view of the channel. In contrast, we build a model representing the channel state as seen by each individual source, instead of the channel state shared by all nodes. Yet, in the scenarios of the AIS group, the behavior of each station can still be decoupled from that of the other stations, as done in [4]. This property significantly simplifies the analysis of the interaction among the two flows, as we show below.

Channel State as Seen by a Single Source. The behavior of an arbitrary station employing a CSMA protocol such as the DCF function of 802.11 can be abstractly represented by a temporal diagram such as the one illustrated in Figure 13. We identify 4 different states : (i) idle channel; (ii) channel occupied by successful transmission of the station; (iii) channel occupied by a collision of

the station; (iv) busy channel due to activity of other nodes, detected by means of either physical or virtual carrier sensing (i.e., NAV).




Figure 13: A station's channel view and embedded discrete time renewal process.

The durations of the time intervals during which the channel remains in the four states above are denoted by σ , T_s , T_c , and T_b , respectively. In Figure 13, the time instants of a possible state change are pointed to by arrows placed below the temporal axis. While σ is a constant equal to one 802.11 time slot, the duration of the other intervals can be variable (with general distribution) depending on the access mechanism (basic access or RTS/CTS), the frame size, and the sending rate of the transmitting node(s).

Analysis of the Behavior of a Single Source. In order to analyze the behavior of a station, we make the fundamental assumption that the channel evolution can be described by a renewal process: at each switching time the next state does not depend on the current state, and the four states occur with fixed probabilities Π_s , Π_c , Π_σ and Π_b . The resulting process is thus semi-Markov. Notice in Figure 13 that the durations of T_s , T_c and T_b comprise an idle slot at the end of the interval, which occurs deterministically, so that it is not considered as an individual event of the overall stochastic process (hence these special slots are marked with a dotted arrow in Figure 13).

At the end of an idle slot, the station decrements its backoff counter, and starts transmitting in the next interval if the counter reaches zero. Let τ be the probability that the station sends out a packet after an idle slot, under the assumption that it is always backlogged (we will remove this assumption later in the section). Let p be the probability that a transmission of the station is not successful. The probability p is usually referred to as the *conditional* packet loss probability [4]. We also introduce b , as the probability that the channel becomes busy after an idle slot due to activity of other nodes (assuming that the station does not start transmitting). Using these probabilities, we can specify the occurrence probability of each of the four channel states at the switching instants as follows,

$$\begin{aligned}
 \Pi_s &= \tau(1-p), \\
 \Pi_c &= \tau p, \\
 \Pi_\sigma &= (1-\tau)(1-b), \\
 \Pi_b &= (1-\tau)b.
 \end{aligned} \tag{8}$$

Computation of the throughput. Using renewal theory, the throughput of a station (expressed in packet/s), is given by

$$T_P = \frac{\Pi_s}{\Delta}, \tag{9}$$

where Δ is the average duration of a channel state (in seconds). The final expression for the throughput of a station is then given by

$$T_P = \frac{\tau(1-p)}{\tau(1-p)\bar{T}_s + \tau p\bar{T}_c + (1-\tau)(1-b)\sigma + (1-\tau)b\bar{T}_b}. \tag{10}$$

Now, the probability τ is a deterministic function of p , which depends only on backoff parameters such as the window size, the number of backoff stages, etc. For 802.11, the expression of τ as a function of p has been first computed in [4]. More recently, it has

been shown that one can easily write similar expressions of τ as a function of p for a large class of backoff mechanisms, employing arbitrary window distributions and backoff multipliers [13].

The complete expression of τ for 802.11, which takes into account the maximum retransmission limit jointly with the maximum window size, is given by

$$\tau = \frac{2q(1 - p^{m+1})}{q(1 - p^{m+1}) + W_0[1 - p - p(2p)^{m'}(1 + p^{m-m'}q)]}, \quad (11)$$

where $q = 1 - 2p$, W_0 is the minimum window size, m is the maximum retry limit, and m' is the backoff stage at which the window size reaches its maximum value, $m' \leq m$.

The average durations \bar{T}_s and \bar{T}_c of a successful transmission or of a collision in which the station is involved are also known *a priori* (see [4]). It turns out that the only unknown variables are the occurrence probability b of a busy period, its average duration \bar{T}_b , and p , the conditional packet loss probability. These quantities are specific to each station, and their values derive from the interaction of the station with the rest of network.

In the next section, we apply the above modeling technique to the study of the AIS scenarios. However, we remark that this methodology has general applicability to modeling long-term throughput of flows in arbitrary networks with any number of nodes. In complex topologies comprising more than two flows, evaluation of the variables b and \bar{T}_b for each node turns out to be the most difficult task, whereas their computation is quite simple in scenarios comprising only two flows. While our study in this paper is restricted to two-flow scenarios, computation of b and \bar{T}_b in more general network scenarios can be found in [9]. The analysis of all possible combinations of flow pairs provides instead the basis to evaluate the packet loss probability p (the other fundamental variable that we need to compute the throughput) of a node in an arbitrary topology. That is, two-flow scenarios are the building blocks that can be used to evaluate the collision probability of a transmitter in any network topology: a transmission on a link is successful if it does not collide with any other transmission on neighboring links. Therefore, the careful analysis of all flow pairs presented in this paper is the necessary first step toward the throughput analysis in arbitrary networks.

4.2 Analysis of AIS Flows

Now we apply the general model introduced in Section 4.1 to *independently* study the behavior of the two transmitting nodes A and B in the scenarios of the AIS group. We add an index A or B within brackets to all quantities defined in Section 4.1 to distinguish between the values of the two stations. For example, $\tau(A)$ is the transmission probability of node A . For simplicity, we assume that the payload size of all data frames is constant. The analysis can be extended to the case of variable payload sizes.

We start by considering the behavior of flow A . As described in Section 3, sender A does not detect any activity on the channel produced by flow B , neither by means of physical nor virtual carrier sensing. As a consequence, $b(A) = 0$. It turns out that the only parameter that we need to compute for flow A is the conditional collision probability $p(A)$, or its complement $s(A)$, the conditional success probability.

The transmission attempts of sender A are not coordinated with those of sender B , and occur at random points in time according to the backoff process of A . Our approximation is to assume that each transmission attempt of A is an independent random look at the activity of flow B . Hence we need to characterize the activity of flow B . A fundamental property of flow B is that all transmis-

sion attempts of sender B are successful, i.e., $p(B) = 0$. Indeed, transmissions of B could only collide with the control packets sent by a , but the probability that this happens is negligible because of the lack of synchronization: almost always B or a avoid collisions by sensing the channel busy and refraining from transmitting.²

Under saturated traffic, the activity of flow B is a sequence of successful transmissions, separated by a random number of backoff slots uniformly distributed in the minimum window size W_0 (the window size of sender B is never increased, since $p(B) = 0$). Occasionally, sender B receives a *CTS* or *ACK* packet from node a , freezing its backoff counter for the remaining part of the successful packet exchange of flow A .

The only chance flow A has to successfully transmit is when the initial packet of the two-way or four-way handshake (a *DATA* frame or an *RTS* frame) happens to arrive during those short gaps in which sender B is in the backoff phase. More precisely, we have to examine the channel occupation state as perceived by the receiving node a while node A is trying to initiate a new data transfer, which is illustrated in Figure 14 for both scenarios of the AIS group. Notice that we remove the amount of time in which receiver a is actively transmitting or receiving from sender A , because we are considering the channel around a conditioned on the fact that A starts transmitting a new packet.




Figure 14: Channel occupation state as perceived by node a while sender A attempts to initiate a new data transfer. The top diagram (a) refers to Scenario 12, the bottom (b) to Scenario 11.

The temporal evolution of the channel as perceived by node a can be divided into cycles of variable duration C . Each cycle comprises a successful data transfer of flow B , of duration T_s , and a variable number i of slots σ corresponding to the backoff phase of sender B . The “gap” G during which a packet originated by node A can be received by node a is also reported in Figure 14. We observe that this gap comprises also the DIFS space at the end of T_s . In Scenario 11, node a does not receive the ACK of node b , thus the gap is enlarged by the duration of an ACK and a SIFS. A key observation is that for flow A to be successful, not only node A must start transmitting during the gap, but the entire packet that node A places on the channel must fit into the same gap. This explains why we do not consider as potential gaps the short inter-frame spaces in between the packets of flow B , nor even, in Scenario 11, the space corresponding to the (unheard) CTS of node b , because $RTS > CTS + 2 \text{ SIFS}$.

If the RTS/CTS mechanism is not used, the situation is very similar to that represented in Figure 14, with the only difference that there are no RTS and CTS packets. Also, notice that in this case an entire *DATA* packet must fit into the gap in order to be successfully received by node b .

²A collision would occur only if the time instants at which nodes a and B start placing a packet on the channel are separated in time by less than the propagation delay between the two nodes, which is a rare event.

Table 1: The value of D to be used in (12).

RTS/CTS - scenario 11	ACK + DIFS - RTS - SIFS
RTS/CTS - scenario 12	DIFS - RTS
Basic Access - scenario 11	ACK + DIFS - DATA - SIFS
Basic Access - scenario 12	DIFS - DATA

In order to compute the conditional packet loss probability $p(A)$, we make the simplifying assumption that the initial packet sent by node A (RTS or DATA packet) arrives at an arbitrary point in time during a cycle C . Also, we assume that all transmission attempts of node A randomly and independently sample a point within a generic cycle C . Although this is only an approximation, the analytical predictions produced by the resulting model are quite accurate (see Section 4.4).

We observe that since the duration of a cycle is variable (the number i of slots is randomly chosen by the transmitting node B), we must consider the fact from renewal theory that the probability of arriving within a cycle of duration C is proportional to C .

The final expression of the packet loss probability of flow A is as follows

$$p(A) = 1 - \frac{2}{W_0[2T_s + (W_0 - 1)\sigma]} \sum_{i=0}^{W_0-1} \max(0, D + i\sigma), \quad (12)$$

where D is a parameter that depends on the access mechanism used (basic access or RTS/CTS) and on the considered scenario. Table 1 provides the value of D for all combinations of cases. Notice that D can take negative values, which explains the max operator in Equation (12). If D is positive, Equation (12) simplifies to

$$p(A) = \frac{2(T_s - D)}{2T_s + (W_0 - 1)\sigma}, \quad D > 0. \quad (13)$$

We observe that $p(A)$ can be directly computed as a function of all known system parameters. Thus, we can already compute the throughput of flow A by first obtaining $\tau(A)$ from Equation (11), and then substituting both $\tau(A)$ and $p(A)$ into Equation (10) (recall that $b(A) = 0$).

We now turn to the analysis of flow B . We have already seen that $p(B) = 0$, thus we can obtain from Equation (11) $\tau(B) = 2/(W_0 + 1)$. The only unknown variable of flow B is $b(B)$, the probability that node B , after an idle slot during the backoff phase, starts receiving a control packet from a (CTS or ACK), after which B sets the NAV and suspends its activity, allowing the packet exchange of flow A to complete successfully. The duration \bar{T}_b of this suspension is equal to \bar{T}_s minus the duration of the first packet (RTS or DATA) sent by A , which is not heard by B .

Since we have already independently computed the throughput of flow A , we know the rate at which sender B has to suspend its activity during the backoff phase. Indeed, the following equation, similar to Equation (10), has to be satisfied

$$T_P(A) = \frac{[1 - \tau(B)]x}{\tau(B)\bar{T}_s + [1 - \tau(B)](1 - x)\sigma + [1 - \tau(B)]x\bar{T}_b}, \quad (14)$$

from which one can obtain the unknown variable $x = b(B)$, to be used into the expression of the throughput of flow B .

We remark that as a result of our analysis, the throughputs of both flows are available in closed form expressions. This is made possible by the hypothesis that both flows are backlogged. In the next section, the analysis is extended to the case of non-backlogged sources.

4.3 Non-Continuously-Backlogged Flows

The starvation problem observed in the scenarios of the AIS group is particularly severe when flow B is continuously backlogged and transmits at the maximum achievable rate, occupying the largest possible fraction of channel time with its own transmissions, and leaving few gaps to be discovered by flow A . Therefore it is important to model these scenarios under more general assumptions, i.e., when flow B does not utilize all of the available bandwidth. This can happen because flow B represents a variable rate flow that empties its transmission queue, or if the sender or receiver of flow B are deferring elsewhere, e.g., if flow B senses the activity of other flows in the network by means of either physical or virtual carrier sensing.

In our analysis, we assume that the maximum throughput achievable by a station is known. For example, if the transmission queue of a source is fed by an arrival process of data packets from the upper protocol layers at rate λ , the achieved rate clearly cannot exceed this value. If the queue is backlogged but the station senses the activity of other flows in its neighborhood, the achievable throughput will be limited by the resulting share of the channel capacity. In this paper we limit ourselves to the analysis of two-flow scenarios, therefore we do not deal with the problem of solving the interaction of many flows in arbitrary topologies. Regardless, we can model the behavior of two-flow scenarios embedded in a large topology by considering each scenario in isolation from the network, and assuming that the transmission queues of the senders are fed by a given arrival rate of packets λ , that can either represent the actual data rate offered by the upper protocol layers, or the maximum rate resulting from the interaction with the rest of the network.

Moreover, we assume the the actual throughput T achieved by a station as a function of λ is equal to the input rate λ up to a saturation value T_{sat} , after which it remains constant and equal to T_{sat} . Our analysis in the previous sections has actually computed the saturation throughput values $T_{\text{sat}}(A)$ and $T_{\text{sat}}(B)$ when both flows are backlogged.

Now we extend the analysis to the case in which senders are fed by arbitrary input rates $\lambda(A)$ and $\lambda(B)$. We add a new probability e to the description of the behavior of a single station introduced in Section 4.1, which is the conditional probability that the transmission queue of the station is empty, given that the station can potentially start a new transmission (i.e., when its backoff counter is zero and the channel has been sensed idle for a time slot). The occurrence probability of each of the four channel states at the switching instants are modified as follows,

$$\begin{aligned} \Pi_s &= \tau(1 - p)(1 - e), \\ \Pi_c &= \tau p(1 - e), \\ \Pi_\sigma &= [(1 - \tau) + \tau e](1 - b), \\ \Pi_b &= [(1 - \tau) + \tau e]b, \end{aligned} \quad (15)$$

and the throughput expression in Equation (9) must be changed accordingly. As a result, the only unknown variable to compute is the probability e . This can be easily done by considering that $T = \min(\lambda, T_{\text{sat}})$ as described above, and by assuming that all other variables are known: if $T|_{e=0} < \lambda$, then $e = 0$ (the source is saturated); otherwise e is equal to the value e^* such that $T|_{e^*} = \lambda$, which is easily obtained by inverting the throughput formula.

The solution of the scenarios of the AIS group under arbitrary input rates requires an iterative approach: each flow is studied assuming that the throughput of the other flow is given. The independent analysis of each flow allows to compute a new estimate of its throughput, to be used in the analysis of the other flow in the next step of the iteration. After a few iterations we obtain the fixed-point solution.

SIFS	10 μ s
DIFS	50 μ s
EIFS	364 μ s
σ	20 μ s
BasicRate	2 Mbps
DataRate	11 Mbps
PLCP length	192 bits @ 1 Mbps
MAC header (RTS,CTS,ACK,DATA)	(20,14,14,28) bytes @ BasicRate
(CW_{min} : CW_{max})	(31,1023)
Short Retry Limit	7
Long Retry Limit	4

Table 2: Parameters setting for the MAC and physical layers.

For flow A , we only need to compute the collision probability $p(A)$, or its complement $s(A)$. To do so, we model the activity of flow B as perceived by a while sender A attempts to transmit as an alternating on-off process. The *on* period has a fixed duration T_{ON} , equal to the portion of the cycle in Figure 14 not occupied by the gap G , which depends only on the access method and the specific scenario. The *off* period is the gap available for flow A , that now can contain also periods of time in which the queue of sender B is empty. The average duration \bar{T}_{OFF} can be computed from the following expression,

$$\frac{1}{T_{ON} + \bar{T}_{OFF}} = \frac{T(B)}{1 - T(A)\bar{T}_b}, \quad (16)$$

which states that the rate at which *on* periods occur (the inverse of the average duration of a cycle) must be equal to the throughput of flow B , normalized by the fraction of time in which the channel is not occupied by successful transmissions of flow A . We assume that the duration of the *off* period is exponentially distributed, and that the transmission attempts of sender A arrive at random point in time. We obtain the following formula for the successful probability of flow A ,

$$s(A) = \frac{\bar{T}_{OFF}}{T_{ON} + \bar{T}_{OFF}} e^{-\frac{d}{\bar{T}_{OFF}}}, \quad (17)$$

where d is the duration of the first packet sent by A (RTS or DATA). Equation (17) states that a transmission attempt of flow A is successful if the first packet arrives during the *off* period of flow B and if it is fully received by a before the beginning of the next *on* period. A new estimate of the throughput of flow A can now be derived using the collision probability $p(A) = 1 - s(A)$.

For flow B , the only unknown is the probability $b(B)$, which is updated using the same reasoning applied at the end of Section 4.2.

4.4 Simulations and Model Validation

Here we validate our analysis of the AIS group and compare the analytical predictions of the throughput of the two flows with simulation results obtained with *ns*. We consider stations operating according to the 802.11b standard, with a data rate of 11 Mb/s. The common parameters at the MAC and physical layers for all of the experiments of this paper are reported in Table 4.4.

We first consider the case of continuously backlogged sources. In Figure 15 we compare the throughput achieved by flow A (in pkt/s) in the two scenarios of the AIS group. We vary the data payload size from 100 bytes to 1500 bytes, and we use the RTS/CTS access mechanism to transfer all data frames, irrespective of their size. Figure 16 reports the corresponding results for flow B .

We observe that despite a number of approximations, the model's predictions provide an excellent match with simulation results for all payload sizes and in both scenarios. As expected, the throughput of Flow B is significantly larger than the throughput of flow A (notice the different scales on the y axis). As predicted by the




Figure 15: Throughput of Flow A vs. data payload size (with RTS/CTS).




Figure 16: Throughput of Flow B vs. data payload size (with RTS/CTS).

model, starvation is partially alleviated in scenario 11, when nodes a and b are not in range of each other. As already explained, this is due to the fact that in this case, the gap in which node A can successfully send to a is enlarged by the duration of an ACK (see Figure 14).

In Figure 17 we compare the throughput achieved by flow A in scenario 11, considering the basic access method and assuming that the maximum retry limit is equal to either the Short Retry Limit or the Long Retry Limit as specified in the 802.11 standard, irrespective of the data payload size (notice that the Short Retry Limit corresponds in the model to $m = 6$, while the Long Retry Limit corresponds in the model to $m = 3$).




Figure 17: Throughput of Flow A vs. data payload size (basic access).

We observe that flow A achieves very different throughput depending on the packet payload size used. Flow A is completely starved (i.e. achieves zero throughput) when the size of a DATA

packet exceeds the maximum possible gap left free by flow B . For very small payload sizes, the throughput of flow A can actually be higher than the throughput obtained employing the RTS/CTS mechanism due to the reduced MAC overhead.

We also observe that flow A achieves significantly higher throughput when the Long Retransmission Limit ($m = 3$) is used, with respect to the case in which the Short Retransmission Limit ($m = 6$) is used. This is due to the fact that when $m = 3$, sender A spends less time in backoff, because after reaching the maximum retransmission limit, the window is reset to the minimum value W_0 . This increases the aggressiveness of sender A , which is more likely to find an available gap in the activity of flow B .

Finally, we consider flows that are not continuously backlogged. The most interesting case to analyze is when we limit the rate of flow B so as to leave more time available to flow A . In Figures 18 and 19 we plot the throughput achieved by the flows as a function of the input rate $\lambda(B)$, while keeping flow A always backlogged. Figure 18 refers to scenario 12 with the RTS/CTS mechanism, whereas 18 refers to scenario 11 with basic access. The data payload size is constant and equal to 1000 bytes, and $m = 6$. We also report on the plots the sum of the throughput of the two flows.




Figure 18: Throughput of Flow A vs. arrival rate of flow B (with RTS/CTS, Scenario 12).




Figure 19: Throughput of Flow A vs. arrival rate of flow B (basic access, Scenario 11).

We observe that the asymmetry between the flows results in nearly strict priority of flow B over flow A . Indeed, flow B achieves a throughput exactly equal to the input rate up to a sharp saturation point, whereas flow A , even if backlogged, only gets a fraction of the remaining channel capacity. Notice that there is a significant loss in aggregate throughput when the two flows obtain similar throughput. This is due to the time wasted by flow A during backoff, even when flow B does not have packets to transmit, leaving the channel idle for a large fraction of time. By limiting the rate of

flow B , it is possible to give the flows the same rate, at the expense of a loss in the aggregate throughput. This illustrates the trade-off between fairness and capacity (aggregate throughput) in networks as it is realized with an unfair access mechanism.

5. SYMMETRIC INCOMPLETE STATE AND SHORT TERM UNFAIRNESS

In this section we develop an analytical model to study the behavior of flow pairs belonging to the Symmetric Incomplete State (SIS) group, comprising Scenarios 8, 9 and 10 in Figure 1. We explore the resulting short-term unfairness by means of analysis and examine the impact of various protocol parameters.

5.1 Analytical Model

The main difficulty in analyzing the scenarios of the SIS group resides in the fact that the behavior of the two flows is tightly correlated: when one flow starts dominating over the other, the states of the flows clearly cannot be considered to be independent, and we therefore cannot employ the decoupling technique adopted in Section 4. In order to correctly analyze the system, it is necessary to consider the joint behavior of the two flows.

Thus, we represent the system state as the pair (S_A, S_B) , where S_A and S_B represent the current backoff stages of transmitters A and B , $0 \leq S_A, S_B \leq m$, respectively. Recall that $m + 1$ is equal to the maximum retransmission limit, which plays a fundamental role in the behavior of the flows as we show in Section 5.2. The total number of states is $(m + 1)^2$ yielding a computationally efficient solution.

Using our bi-dimensional state description, we build a discrete time Markov Chain embedded over continuous time at the time instants in which both senders can (potentially) start transmitting the first packet of a new data exchange (either the RTS or the DATA packet), provided that their backoff counter is equal to zero. We use the same channel view as depicted in Figure 13, but consider only time epochs at which both transmitters can attempt a new transmission.

We assume that the backoff counter of a station is geometrically distributed, instead of uniformly distributed, over the current window. By doing so, we can exploit the memoryless property of the geometric distribution and avoid explicitly incorporating in the state description the remaining number of backoff slots of each station. Our simulations indicate that this approximation does not compromise the model's accuracy. The parameter γ_i of the geometric distribution that characterizes the backoff counter at stage i ($0 \leq i \leq m$) is given by $\gamma_i = \frac{2}{W_i - 1}$, where W_i is the window size of backoff stage i . Consequently, at each time epoch a station in stage i attempts a new transmission with probability γ_i .

5.1.1 Transition Probabilities

We consider the behavior of the two flows in Scenario 8 (or equivalently, Scenario 9). The key point to analyze the system dynamics is the computation of the (conditional) packet collision probability p . We observe that sender A (for example) does not know if sender B has started accessing the channel until a packet sent by B has been fully received at b , triggering the transmission of a CTS or ACK that can be immediately sensed by A . This leads to the following typical situation: sender A starts transmitting a new packet, but before it is fully received by a , sender B also starts transmitting a packet, resulting in a collision at the receivers in which both packet are destroyed. A transmission from sender A is successful only if sender B does not attempt to transmit in all transmission opportunities that occur during the duration of the first

packet (either the RTS or the DATA packet) sent by A . The number of such transmission opportunities is given by the duration of the first packet (RTS or DATA) expressed in the number of backoff slots, and denoted by f . The (conditional) successful probability of sender A in state (i, j) is then given by $(1 - \gamma_j)^f$ (recall that j is the backoff stage of sender B , and γ_j the corresponding transmission probability).

from state	to state	probability
i, j	i, j	$(1 - \gamma_i)(1 - \gamma_j)$
i, j	$0, j$	$\gamma_i (1 - \gamma_j)^f$
i, j	$i, 0$	$(1 - \gamma_i)^f \gamma_j$
i, j	$(i + 1) \bmod m, (j + 1) \bmod m$	otherwise

Table 3: Transition probabilities of the Markov Model.

The transition probabilities stemming from the generic state (i, j) are summarized in Table 3. The first row of the table is the self transition corresponding to the case in which both flows do not start transmitting a new packet. The second and third rows refer to successful transmissions from sender A or B , respectively. Notice that in this case the backoff stage of the station successfully transmitting is reset to the initial value (stage 0). Finally, the last row corresponds to a collision event for both flows, with the consequent increase of the backoff stage – if the backoff stage reaches the maximum retransmission limit, it is reset to 0, which explains the modulus operator.

5.1.2 Performance Metrics

By numerically solving the Markov Chain, which is ergodic for any choice of parameters, we obtain the stationary distribution $\pi = \{\pi_{i,j}\}, \forall i, j$. An example of such a distribution is reported in Figure 20, obtained using the set of parameters in Table 4.4 and considering the RTS/CTS access mechanism. The plot clearly suggests a bi-stable behavior, in which the most likely states are those in which one flow maintains a small value of window size (thus obtaining high throughput), while the other falls into deeper and deeper backoff stages (thus obtaining low throughput). Only when the poor flow reaches the maximum retransmission limit, it resets its window and competes equally with the rich flow.




Figure 20: Steady state probabilities of the Markov model in case of RTS/CTS mechanism, default 802.11b parameters.

Long-term performance metrics can be obtained directly from the solution of the Markov Chain. From renewal-reward theory,

the throughput of either flow is given by

$$T = \frac{\sum_{i,j} \pi_{i,j} \gamma_i (1 - \gamma_j)^f}{\Delta},$$

where Δ is the average duration of a step. The duration of a successful transmission is equal to T_s , as defined in Section 4. The idle slot is σ , while a collision has an approximate average duration of $T_c + \sigma f/2$, assuming that the colliding packet starts on average in the middle of the packet that is transmitted first. Δ is computed as the average of the duration of all possible events in all states, weighted by their respective probabilities.

The average conditional collision probability p for each flow can be computed as

$$p = \frac{\sum_{i,j} \pi_{i,j} P_c(i, j)}{\sum_{i,j} \pi_{i,j} [P_c(i, j) + \gamma_i (1 - \gamma_j)^f]},$$

where $P_c(i, j)$ is the collision probability at state (i, j) (fourth row of Table 3).

5.1.3 Transient Analysis

As SIS scenarios are long-term fair, we are most interested in evaluating the short-term unfairness of these cases. In particular, it would be desirable to have an indication of the average amount of time during which one flow experiences poor throughput while the other gets most of the available channel capacity. To this purpose, we consider the system states $(m, 0)$ and $(0, m)$, where one flow has reached the last backoff stage, while the other is at stage 0. These two states are expected to be reached in the two symmetric conditions in which one flow strongly dominates over the other. The average amount of time necessary to transition from one of these states to the other provides a good estimate of the duration for the system to switch from one equilibrium point to the other, i.e., the system's time-scale of unfairness.

To compute this, we exploit the symmetry of our system, and reduce the problem to finding the average time necessary to re-enter one of the above states (for example $(m, 0)$) after having left it. To simplify notation, let i be state $(m, 0)$.

We proceed as follows. First, we remove all self-transitions of state i , since we want to consider re-entries to this state only after a change of state. To do so, transitions leaving state i are re-normalized to sum up to one.³ The time actually spent in state i in between two successive visits to this state will be considered separately at the end of the computation.

After applying this modification to the transition probabilities, we recompute the stationary distribution of the model obtaining a different state vector π' . From the renewal-reward theorem applied to cycles defined by visits to state i , we know that the average number of visits to state j between returns to state i is given by $E(V_{ij}) = \pi'_j / \pi'_i$. Moreover, the average number of transitions of type $j \rightarrow k$ between returns to state i is given by [21]

$$E(V_{ijk}) = \pi'_j p_{jk} / \pi'_i, \quad \forall i, j, k. \quad (18)$$

Using this result, we can compute the average time to return to state i after leaving it by summing all durations associated with transitions $j \rightarrow k$, weighted by their expected average number of occurrences given by Equation (18). Finally, we add the average time spent in state i before leaving it, which was initially removed. This can be easily done by considering that the average number of steps spent in a state is geometrically distributed, with parameter equal to the exit probability from the state.

³Indeed, following all transitions leaving i , the poor flow (flow A) resets its window to the minimum value, exiting the starvation condition and starting to compete fairly with the rich flow.

5.2 Simulations and Model Validation

In this section we validate the analytical model of the behavior of the two flows in Figure 11(a) and investigate several properties of this scenario with a focus on short-term unfairness. Our results are summarized in Table 4, in which we compare analytical predictions with ns simulations in four different cases: (C1) RTS/CTS access, $m = 6$, $CW_{\max} = 1024$; (C2) RTS/CTS access, $m = 8$, $CW_{\max} = \infty$; (C3) basic access, $m = 3$, $CW_{\max} = 1024$; (C4) basic access, $m = 6$, $CW_{\max} = 1024$.

case	model			ns			SC
	T	p	Δt	T	p	Δt	T
C1	218	0.25	235	216	0.25	223	250
C2	229	0.11	982	230	0.09	1156	250
C3	125	0.69	15	107	0.75	15	337
C4	222	0.37	59	220	0.38	60	337

Table 4: Average throughput per flow T (in pkt/s), conditional packet loss probability p , and average transient time Δt (in ms) obtained in Scenario 8 for four different settings of parameters.

Case C1 corresponds to the default parameters of 802.11 when RTS/CTS is used, having the maximum retransmission limit equal to 7, the *Short Retry Limit*. Case C2 also employs RTS/CTS, but increases the maximum retransmission limit to 9 and does not bound the maximum window size. Case C3 corresponds to the default parameters of 802.11 with basic access, having maximum retransmission limit 4, and the *Long Retry Limit*. Case C4 differs from C3 in that the maximum retransmission limit is set to 7, the *Short Retry Limit*.

We compare both long-term performance metrics (average throughput and collision probability) and our characterization of the short-term unfairness by means of the average time Δt to transition from state $(m, 0)$ to $(0, m)$. This latter quantity is computed analytically using the approach described above and has also been measured in simulation via inspection of the backoff stage of the two flows. The last column of Table 4 reports the per-flow throughput (in pkts/s) achieved in each case when senders are connected (SC).

In addition to the excellent agreement between analysis and simulations in all cases, we make the following observations. All configurations achieve similar throughput (around 220 pkt/s) except case C3, for which we observe a severe penalty (62%) mainly due to the large packet loss probability combined with the small retransmission limit, which leads to many packets being dropped by the MAC. In the other cases the throughput loss with respect to the case in which senders are connected is not significant using RTS/CTS (around 10%, cases C1 and C2), while it is important in case C4, in which basic access is employed (34%).

The maximum throughput is achieved in case C2, where we have increased both the maximum number of retransmissions and the maximum window size with respect to standard values. The small gain in throughput comes at the expense of a severe short term unfairness, as the average transition time from one equilibrium point to the other approaches 1 second. This confirms the significant impact of m on the time-scales of short-term unfairness.

Surprisingly, the basic/access mechanism with $m = 6$ achieves the smallest value of Δt while obtaining a throughput comparable to that obtained using the RTS/CTS mechanism. This is remarkable, as the RTS/CTS mechanism was actually proposed to overcome the hidden terminal problem in infrastructure mode (see Figure 11(b)). Unfortunately, it leads to further short-term unfairness than that obtained when the basic access mechanism is used.

As a further validation of our conclusions, we present in Figure




Figure 21: Packets successfully sent by one flow every 100 ms, in all four cases.

21 simulation results showing the average number of packets successfully sent by one of the two flows every 100 ms, in the window $[50s - 60s]$, for each of the four cases. We also show the average number of packets obtained when senders are connected, for comparison. We observe that the throughput oscillates dramatically in the cases having a large value of Δt , enduring prolonged intervals in which the achieved rate is either very high or very low.

6. RELATED WORK

Problematic Topologies. Variants of some of the scenarios we investigated are known to incur poor performance. Among them are the classic hidden terminal and exposed terminal problems examined in [2, 12, 18]. Likewise, AIS scenarios are described in [2, 3, 11, 12, 20]. Yet no prior studies comprehensively analyze all scenarios within a single analytical framework nor predict scenario likelihood in a random graph.

Modeling Throughput. The fully connected topology and Senders Connected (SC) group can be accurately modeled using [4] and hence we do not consider them here. In contrast, we model topologies in which senders are disconnected, which leads us to develop new modeling techniques for both the AIS and SIS groups.

In [15, 16], a queuing analysis of Scenario 11 in Figure 1 is developed under (unrealistic) assumptions that (i) the time between retransmissions (i.e. the backoff delay) is negligible, and (ii) the maximum number of retransmissions allowed for each packet is unlimited. Consequently, the model matches simulations well only under light traffic conditions and with very large packets transmitted at low data rate. In contrast, our analysis incorporates all details of 802.11 DCF, applies to both saturated and non-saturated conditions, incorporates variable packet sizes, and shows the significant impact of the maximum retransmission limit and access method (two- or four-way handshake). In [10] the authors present a throughput analysis that incorporates topology dependent relations, and report preliminary results for the ring topology neglecting the impact of the binary exponential backoff.

Modeling Short-Term Unfairness. While a number of studies have analyzed short-term unfairness due to various aspects of the wireless medium ranging from channel errors to contention, (see [1, 5, 6] for example), none accurately characterize the time scale in which 802.11 stations alternate domination in SIS scenarios. In particular, in our analysis of the SIS group, we model short-term

unfairness of IEEE 802.11 with and without RTS/CTS, derive the behavior of the flows from the collision probability, analyze the transient behavior of the system, and accurately predict the time-scale of this unfairness as a function of key system parameters.

Receiver Oriented Media Access. A number of receiver-based access mechanisms, in which channel contention is initiated by the receiver, not the sender, have been proposed to improve fairness and performance in various scenarios [2, 8, 19]. The Request-for-Request-to-Send (RRTS) solution is an example receiver-oriented mechanism originally proposed in MACAW [2]: whenever a station receives an RTS to which it cannot respond (due to deferral), it contends during the next contention period on behalf of the sender, reserving the channel by means of an additional control packet called RRTS. The RRTS message solicits the sender to immediately send a new RTS.




Figure 22: Comparison of the number of packets sent by one flow during intervals of 100 ms in scenario 10, with or without RRTS

We have investigated the effectiveness of this solution in both the AIS and SIS group, implementing the RRTS packet exchange in the *ns* simulator. We have found that this solution has a fundamental limitation in the fact that it requires the receiver of a flow to correctly decode the RTS packet sent by the sender. Unfortunately, this occurs extremely rarely when a transmission attempt is not successful. Consider the scenarios of the AIS group: the RTS packet sent by *A* typically arrives at *a* during a transmission of flow *B*, or it is destroyed by a new packet sent by *B* before it is fully received by *a*. Simulation results confirm that the RRTS mechanism does not help in the scenarios of the AIS group.

In the SIS group, the RRTS mechanism is also not useful in scenarios 8 and 9, because during a collision event both RTS packets are destroyed at the receivers. The only exception is scenario 10: in this case, the two flows are connected only through the receivers; the receiver of a flow is not in range of the transmitter of the other flow, and correctly receives the RTS sent by the sender. This scenario leads to short-term unfairness, and can be analyzed in a similar way as described in Section 5. Figure 22 shows that the RRTS mechanism partially mitigates the observed unfairness. Notice that the occurrence probability of Scenario 10 is only about 10%, as reported in Figure 5.

7. CONCLUSIONS

We have systematically and comprehensively investigated the foundational scenarios of multi-hop wireless networks. We identified all possible topologies consisting of four nodes and two flows, classified them into three geometric groups, and computed their likelihood under random node placement. In each case, we showed how fundamental properties of two- and four-way handshake CSMA protocols yield short-term unfairness in one group, and long-term unfairness in another. Because these problematic scenarios are highly likely to occur in a random graph, we developed an analytical model that can accurately predict the performance in each of

these scenarios. For example, we are able to precisely predict each flow's throughput in scenarios with long-term unfairness and the time scale that flows alternate domination in scenarios with short-term unfairness.

8. REFERENCES

- [1] G. Berger-Sabbatel, A. Duda, M. Heusse, and F. Rousseau. Short-term fairness of 802.11 networks with several hosts. In *Proceedings of the 6th IFIP/IEEE International Conference on Mobile and Wireless Communication Networks*, Paris, France, October 2004.
- [2] V. Bharghavan, S. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LANs. In *Proceedings of ACM SIGCOMM '94*, London, UK, August 1994.
- [3] V. Bharghavan. Performance evaluation of algorithms for wireless medium access. In *Proceedings of IEEE International Computer Performance and Dependability Symposium (IPDS '98)*, Durham, NC, March 1998.
- [4] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, March 2000.
- [5] M. Carvalho and J.J. Garcia-Luna-Aceves. Delay analysis of IEEE 802.11 in single-hop networks. In *Proceedings of IEEE ICNP '03: 11th IEEE International Conference on Network Protocols*, Atlanta, Georgia, November 2003.
- [6] C. Chaudet, D. Dhoutaut, and I. G. Lassous. Experiments of some performance issues with IEEE 802.11b in ad hoc networks. In *Second Annual Conference on Wireless On-demand Network Systems and Services (WONS'05)*, St. Moritz, Switzerland, January 2005.
- [7] B. P. Crow, I. K. Widjaja, G. Jeong, and P. T. Sakai. IEEE 802.11 wireless local area networks. *IEEE Communications Magazine*, 35(9):116–126, September 1997.
- [8] J.J. Garcia-Luna-Aceves and A. Tzamaloukas. Reversing the collision-avoidance handshake in wireless networks. In *Proceedings of ACM MobiCom '99*, Seattle, Washington, August 1999.
- [9] M. Garetto, T. Salonidis, and E. Knightly. Starvation in CSMA-Based Multi-hop Wireless Networks: Model and Solutions. Technical report, Rice University, Houston, TX, 2005.
- [10] N. Gupta and P. R. Kumar. A Performance Analysis of the 802.11 Wireless LAN Medium Access Control. *Communications in Information and Systems*, 3(4):279–304, 2004.
- [11] V. Kanodia, C. Li, A. Sabharwal, B. Sadeghi, and E. Knightly. Ordered packet scheduling in wireless ad hoc networks: Mechanisms and performance analysis. In *Proceedings of ACM MobiHoc*, Lausanne, Switzerland, June 2002.
- [12] P. Karn. MACA: A new channel access method for packet radio. In *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, April 1990.
- [13] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New insights from a fixed point analysis of single cell IEEE 802.11 WLANs. In *Proceedings of IEEE INFOCOM '05*, Miami, FL, 2005.
- [14] L. E. Miller. Distribution of link distances in a wireless network. *Journal of Research of the National Institute of Standards and Technology*, 106(2):401–412, 2001.

- [15] S. Ray, J. Carruthers, and D. Starobinski. Evaluation of the masked node problem in ad-hoc wireless LANs. *IEEE Transactions on Mobile Computing*, in press.
- [16] S. Ray, D. Starobinski, and J. Carruthers. Performance of wireless networks with hidden nodes: A queueing-theoretic analysis. *Journal of Computer Communications (Special Issue on the Performance of Wireless LANs, PANs, and Ad-Hoc Networks)*, in press.
- [17] Y. Sun, X. Gao, E. M. Belding-Royer, and J. Kempf. Model-based resource prediction for multi-hop wireless networks. In *Proceedings of the 1st IEEE International Conference on Mobile Ad hoc and Sensor Systems (MASS)*, Ft. Lauderdale, FL, October 2004.
- [18] F. A. Tobagi and L. Kleinrock. Packet switching in radio channels: Part II - the hidden terminal problem in carrier sense multiple access and the busy tone solution. *IEEE Transactions on Communications*, 23(12):1417–1433, 1975.
- [19] Y. Wang and J.J. Garcia-Luna-Aceves. A new hybrid channel access scheme for ad hoc networks. In *Proceedings of Med-Hoc-Net*, Sardegna, Italy, September 2002.
- [20] Y. Wang and J.J. Garcia-Luna-Aceves. Channel sharing of competing flows in ad hoc networks. In *IEEE Symposium on Computers and Communications (ISCC '03)*, Kemer-Antalya, Turkey, July 2003.
- [21] R. W. Wolff. *Stochastic Modeling and the Theory of Queues*. Prentice Hall, 1989.
- [22] J. Yoon, M. Liu, and B. Noble. Sound mobility models. In *Proceedings of ACM MobiCom '03*, San Diego, CA, September 2003.