Multi-interface Multi-channel Wireless Mesh Networks

by

Mohammad Ahmad Munawar

A thesis
presented to the University of Waterloo
in fulfilment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2004

Author's Declaration For Electronic Submission of a Dissertation

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

ABSTRACT

In this thesis we propose a multi-channel wireless network based on nodes that use multiple 802.11 radio interfaces. The proposed system is singular, as it does not require new hardware or a new MAC, but instead leverages commodity 802.11-based products. With this system, we target scenarios where the nodes are stationary and where their location can often be controlled. We evaluate the performance in this setup using an ad-hoc network approach whereby nodes generate as well as forward data. We also present and appraise a purely-wireless multi-channel infrastructure, which operates like the WLAN infrastructure-based networks in existence today, but without any fixed-line support. In such an infrastructure nodes dedicated for routing purposes provide wireless connectivity to users. We show that a multi-interface system provide significantly higher capacity in many scenarios. Our work puts forward various challenges, points to various anomalies in the operation of the 802.11 MAC protocol, and shows the need to tackle unfairness issues. Our experiments demonstrate that the mere use of more dual-interface nodes does not necessarily create higher capacity. We also show that traffic differentiation significantly increases aggregate throughput in realistic scenarios. Finally, we provide an example of how simple channel-allocation algorithms in controlled random topologies can allow us to take advantage of a multi-interface system.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Paul A.S. Ward, for his precious guidance and continuous support throughout this work. I also extend my gratitude to all members of our group and members of the Shoshin Distributed System Group for their help and for making the working environment enjoyable as well as instructive. Finally, I am thankful to the International Council for Canadian Studies (ICCS) for their financial support through the Commonwealth Scholars program.

To my parents...

CONTENTS

Figures

TABLES

1 Introduction

In today's information age access to information has become essential irrespective of time and location. Fulfilling this need entails the use of technologies that do not constrain users to specific locations. In this respect wireless networking is transforming this need for ubiquitous information access into a reality. Wireless networks enable us to connect devices without wires, which are not only disruptive but also very costly. Not surprisingly, we see wireless-capable devices becoming increasingly used in our daily lives. For example, cell phones, text-messaging devices like Blackberry, PDAs, and laptops are essential tools for many of us. The increasing popularity of wireless networking is witnessed by the soaring sales of related products. As a result of the perceived benefits and the resulting high demand, there is a great deal of research work under way in the area to meet ever-increasing expectations from users.

Present wireless services include voice calls, text messages, email access, small-file exchange, *etc*. Services that can be provided are presently constrained by the capacity of the underlying wireless networks. The vision of pervasive computing suggests many more services, such as permanent connection, video streaming, *etc*. It is anticipated that we will become more dependent on wireless connectivity, and more so with the increasing use of embedded devices. For example, Radio Frequency Identification (*RFId*) is expected to revolutionize management of products from their production to the end of their lifetime. The research community is endeavouring to improve wireless networks in order to meet these expectations. Motivated by the various benefits of wireless networks, the research community has explored many avenues to address numerous challenges in the area. While some researchers are working at improving the physical layer, others are exploring enhancements at higher layers. Even though a lot of progress has been made, researchers are still striving to make wireless networks a viable alternative to their wired-line counterparts.

Wireless cellular technologies such as GSM and CDMA [41] were primarily developed for voice applications. Others, such as 802.11 [19] and HiperLAN2 [9], have specifically been designed for data networks, with higher bandwidth as the key requirement. The IEEE 802.11 has become the most popular standard for communication in a wireless local area network. Since the 802.11 standard was specified, 802.11-based products have seen a tremendous growth in sales. The standard has allowed for better inter-operability between products from different vendors. This has, in turn, lead to a considerable reduction in the cost of these products. Most of the present deployments of 802.11-based wireless networks consist of coverage at fixed locations that are supported by wired-line backbones. Such setups are called infrastructure-based wireless networks. The required infrastructure necessitates a significant initial investment and tends to be intrusive. Large deployments of such networks are currently restricted to business intranets or public access networks, many of which are funded for marketing purposes.

In the past few years, the idea of infrastructure-less networks has been much publicized, mostly driven by applications in the military. These instantaneous mobile networks are termed as *ad-hoc* networks. For example, during a battle, soldiers in range of each other can form, on the fly, a temporary network. Much research has gone into *ad-hoc* networks, as they present many challenges because of the lack of a stable core and their highly-dynamic nature. *Ad-hoc* networks, as deployed and proposed today, typically employ a single radio channel for communication. Due to the shared nature of the wireless channel, the performance of such networks is poor. *Ad-hoc* networking is also plagued with other problems that have hindered its widespread adoption. One such problem is the lack of an economic model to support its operation. In multi-hop *ad-hoc* networks, there is currently no widely-accepted technique to compensate users for their forwarding services. Nevertheless, the various advantages of *ad-hoc* networks have motivated us to look for alternatives that provide similar benefits, but that do not suffer from some of their constraints.

In this work we target some specific scenarios that have many applications. To better understand these scenarios, let us consider the following examples: First, a village without any network infrastructure consisting of a number of houses and a village administration office that is connected to the Internet through a satellite link. Let us assume that the village council wishes to make this connection available to the local habitants. Second, a case where we need to organize a one-day gathering in a park for a large number of users. Let us assume that we need to provide wireless connectivity to these users. In the first case the cost of putting in place a wired network may be prohibitive for the village. Similarly, in the second scenario, installing a wired back-haul or an infrastructure-based wireless network that relies on wires to connect the end-points would be difficult, costly, and inappropriate. Both these scenarios can benefit from a purely-wireless network because it would cost far less and would be less intrusive. In order for such a wireless network to operate the nodes need to route packets from each other in a multi-hop fashion. In the village, each house would have a wireless node that could connect to the administration office via other nodes in other houses. Likewise, organizers of the gathering could spread out wireless nodes that would use each other to cover the area of interest.

Two common characteristics in the two scenarios above are the possibly large number of users that need to be supported and the stationary nature of the core network. However, the two scenarios differ as to how much control the users have on the wireless nodes. While in the village, the nodes are controlled by the users and route both user-data and data from other nodes, users in the park do not need to route other users' data because this function is carried out by dedicated wireless nodes.

Any solution to meet the requirements in the two scenarios needs to provide for sufficient capacity in the core network. Given the stationary component in such scenarios, many features of an *ad-hoc* network may not be required. For example, we may not frequently encounter situations where wireless nodes move away. Moreover, many improvement can be made that would not necessarily work under assumptions of mobility. For example, in our gathering scenario, we can control the placement of the network nodes to enhance available capacity and to guarantee coverage. However, the village scenario does not allow for a controlled placement of wireless nodes. Thus, in this case other techniques are needed to improve capacity.

Our intent is to develop a system that targets scenarios similar to the ones described above. Essentially, we assume that the core of the network is stationary. The main challenge here is to provide high capacity in the the backbone such that a reasonable number of users can adequately be serviced. To this effect, we opt to take advantage of the availability of the few non-overlapping channels available for products that are 802.11-compliant and that use Direct Sequence Spread Spectrum (DSSS) at the physical layer. This approach differs from common *ad-hoc* networking proposals, which use only a single channel. Given that a radio channel is a shared resource, the use of more channels would reduce contention if we distribute users and/or backbone-nodes on different channels.

In this work, we emphasize scenarios where we have some control on the network nodes. We present a system that draws characteristics from both infrastructure-based and *ad-hoc* wireless networks. As such, we look at scenarios where users rely on a pre-existing core network, as in an infrastructure-based network, but without the need for a wired backbone. By doing so, the intent is to allow for deployment flexibility in order for the network to be easily reconfigured and relocated. Once deployed, the backbone in the system remains stationary until it is no longer needed and moved. The core is thus akin to an *ad-hoc* network, but without the full mobility support required by the latter. Such a system can be deployed at any time and anywhere, providing the advantages that one can possibly obtain from *ad-hoc* networks. To deploy this infrastructure, the system owner will intelligently place the nodes to provide adequate coverage and sufficient capacity. Ideally, the nodes will be small devices that can be plugged into power outlets, or else operate on batteries or solar energy. In the proposed system, mobility of users can be handled using techniques developed for roaming in a wireless local area network (WLAN) infrastructure.

In our gathering scenario above, while an *ad-hoc* network may meet the requirements, it presumes that users are spread around such that the desired connectivity is achieved. Moreover, it requires that all users have the required software and are willing to share each others' resources, and in particular their batteries. In contrast, using a wireless infrastructure, the organizers will provide the needed equipment. The core nodes will be appropriately placed to support the desired coverage. These nodes may rely on batteries or any alternative source of power supply for their operation. Once placed, the nodes will self-configure for operation. Users can then benefit from the provided connectivity without the need to modify their software. Once the event has finished, the system nodes can be collected and re-used in other situations.

1.1 MOTIVATION

At present, many real-life deployments of purely wireless networks are based on *ad-hoc* networking, which does not require any pre-existing infrastructure. However, such networks not only present capabilities that are not always needed, but also suffer from various shortcomings. *Ad-hoc* networks employ a single radio channel, and thus allocate this shared resource among competing devices. Such a share may be too small for many applications. *Ad-hoc* networks also need either a practical business model or a reward mechanism in order to become viable and emerge as a useful technology. Furthermore, standardization is required to provide off-the-box support in the most commonly-used operating systems.

Ad-hoc networks based on the 802.11 standard mostly utilize a single shared channel. As such, the bandwidth is divided between the nodes trying to communicate. The throughput per node decreases as a function of the active node density in a particular area. Gupta and Kumar [16] show that per-user capacity is an inverse function of the square-root of the number of users, assuming that nodes are identical and are optimally located. In related work by Li *et al.* [24], it was shown that in a very simple setup *i.e.* a chain topology of length 8, the throughput achieved can be as low as $1/7$ of the raw bandwidth. The fundamental issue here relate to physical limitations of the wireless channel. One way to circumvent such limitations is to use more than one independent wireless channel. For example, when 802.11 is used with the Direct Sequence Spread Spectrum (DSSS) at the physical layer, there exist three non-overlapping channels. Instead of using only one channel, if we leveraged all of the available resources, we could achieve higher performance. In fact, all the three available channels are often used in a WLAN infrastructure to enhance capacity. Such setup, however, present several disadvantages. They require a wired back-haul, which may be expensive to put in place and does not provide flexibility as far as reconfiguration is concerned. But, unlike a multi-hop wireless network, access points (AP) in a WLAN infrastructure do not need to communicate with each other over the radio. This makes frequency re-use easier, as channel can be separated enough to minimize interference.

Ad-hoc networks are not practical yet. Even if the bandwidth issue is addressed, there exist other issues that undermine their widespread use. One such issue is the readiness of mobile systems to access the network. Most systems supporting wireless connectivity are supplied with software to connect to a WLAN infrastructure. In contrast, *ad-hoc* routing requires installation and configuration of various routing modules. This also presumes that a specific routing protocol has been agreed upon by the users and is used throughout the system. Presently, operating systems do not come with support for these routing protocols. Besides, there are still issues in integrating *ad-hoc* routing protocols with IP routing [48]. It is thus our intent to leverage this advantage of a WLAN infrastructure insofar as users do not need any special software to access network resources.

Another major hurdle in the widespread acceptance of *ad-hoc* networking is the lack of a business model to support it. This approach requires that every node participate by forwarding packets from others. Presently, no compensation is given to any user for its relaying services. On the contrary, nodes that forward data from others can be disadvantaged. As such, the approach suits scenarios where benefits of the overall system override individual considerations. Other approaches are necessary to create an economically viable wireless system.

We believe that our proposal of using multiple radio interfaces on a node can address important weaknesses of *ad-hoc* networks. Such nodes can either be used in an *ad-hoc* fashion such as in a community wireless network, or they can be employed in an infrastructure-like setup to meet requirements of a temporary wireless network. At present, very little is known about performance improvements that can be achieved using a multi-interface multi-channel 802.11 DSSS system. It is thus necessary to evaluate the behaviour of such systems through analysis and simulation before an eventual implementation.

We do not intend to make any changes to the existing 802.11 MAC or hardware. However, we will point out weaknesses of the current protocol and suggest areas for improvement. In addition, our study of a multi-interface system shows significant improvement, which we think will motivate current users of 802.11 products to experiment with multi-channel multi-hop networks for various applications.

Despite the lack of knowledge about the performance of multi-channel multi-hop wireless networks, some products are already available. Pure wireless networks in mesh-like topology have begun to be deployed, especially in community wireless networks [5] and are being proposed as a cost-effective solution for Metropolitan Area Networks (MAN). The related products have preceded a thorough study of the potential benefits of using multi-channel solutions. Our work aims to bridge this gap by providing a better understanding of performance in such systems.

1.2 CONTRIBUTIONS

This thesis provides the following contributions towards a better understanding of 802.11-based multi-channel wireless networks:

1. We design and develop a multi-interface multi-channel simulation testbed using the *ns-2* simulator.

2. We show significant improvement in capacity by leveraging multi-interface nodes that use a limited number of orthogonal channels.

3. We pinpoint various anomalies inherent in the present 802.11 MAC that have an important impact on achievable aggregate throughput.

4. We demonstrate that using more interfaces does not automatically improve capacity.

5. We show that unfairness is an important issue in multi-interface wireless networks if an *ad-hoc*-like approach is followed without any consideration for traffic control.

6. We study a multi-hop wireless infrastructure whereby routing is carried out by dedicated stationary nodes and users need not route each other's traffic. This infrastructure can meet the requirements of many scenarios targeted by the *ad-hoc* networking research community.

7. We show that traffic differentiation can significantly improve capacity in such an infrastructure.

8. We argue that a very simple channel allocation algorithm without any location information or signal strength assessment may be used to leverage benefits of using multiple channels.

1.3 Approach

This thesis evaluates potential benefits of using of multi-interface multi-channel wireless nodes. We first carried out a number of experiments with a single-channel network in order to confirm problems described by other researchers. Based on this work, we designed a mechanism to coordinate multiple radio-interfaces on one node and modified the *ns-2* [14] simulator to emulate a network with such nodes. Our changes relate to the routing protocol as well as some lower-level layers. We then appraised the behaviour of the new system. We dwelt into details to identify various issues and consequently, we propose some simple modifications to improve performance.

1.4 Thesis scope

In this work, we assume the use of IEEE 802.11 MAC with DSSS at 2 Mbps at the physical layer. We only report results of experiments with the DSDV routing protocol. Although, some simulations were done with AODV, we did not get any new insight in our area of interest. We believe that other, potentially more efficient, routing protocols need to be designed. A number of issues have deliberately been left out of this work, because they represent significant research problems on their own. As such, we do not specify any channel assignment protocol for the backbone in our system. The channel allocation problem is akin to the well-known graph colouring problem that has been extensively studied. Also, we do not implement any mobility support in our simulations. Despite observing unfairness in our system, we do not propose solutions to this problem. We also do not discuss means to support QoS in such networks. Interested readers are referred to [20, 32, 46, 49, 52] for fairness and [10, 43, 51] for QoS. Furthermore, we do not attempt to leverage multi-path routing [23, 28, 35] for fault-tolerance and/or load-balancing.

In summary, our intent to is to evaluate and analyze the performance of a system that can be readily implemented today and identify its weaknesses. Nevertheless, we aim to point out ways in which the system may be improved.

2 BACKGROUND

In this chapter we review the basics of wireless data networks with an emphasis on the IEEE 802.11 standard, *ad-hoc* routing, and multi-channel transmission protocols. We also briefly describe some variants of *ad-hoc* networks proposed to meet some real-world application requirements.

2.1 WIRELESS NETWORKS

Wireless networks consist of stations that are not connected with wires or fiber, but communicate through other media such as radio signals or infra-red light. These networks can generally be classified into two broad categories: infrastructure-based and *ad-hoc*.

A wireless network with infrastructure consists of fixed base-stations at specified locations that provide wireless connectivity to devices within their coverage area. Examples include cellular networks such as the Global System for Mobile Communication (GSM), Code Division Multiple Access (CDMA) (See [41] for an overview), and 802.11 WLAN. Infrastructure includes equipment needed for communicating with end-users, the backbone to interconnect base-stations, hardware to bridge with other networks, *etc*.

In contrast, *ad-hoc* networks are wireless networks without pre-established infrastructure. Such networks are instantaneously formed when interested nodes come within each other's reach. *Ad-hoc* networks can be very useful in situations where a there is no need for an infrastructure or where its creation would be too costly. Applications of these networks include providing connectivity in a disaster-relief situation, (*e.g.* for keeping contact between rescue-team members), for communication between soldiers in battlefields , for sharing traffic-related information along congested roads, *etc.*

There exist many different standards and technologies for wireless communication. For data-intensive traffic, specifications such as IEEE 802.11, HiperLAN2, and Bluetooth have been developed. While 802.11b/g and Bluetooth operate on the 2.4 GHz band, HiperLAN2 and 802.11a use the 5 GHz band. Today, the 802.11 standard is the most popular standard that draws support from major players in the wireless networking industry. Given that this work is exclusively based on the 802.11 DSSS standard, we first briefly describe IEEE 802.11 MAC specification. We then discuss *ad-hoc* routing protocols, which are essential in providing self-configuration and self-healing capabilities in *ad-hoc* networks.

2.2 IEEE 802.11

In this section we provide an overview of essential topics in the IEEE 802.11 standard that are subsequently referred to in this thesis. A more detailed description is available in the actual specification [19].

2.2.1 IEEE 802.11 ARCHITECTURE

The IEEE 802.11 standard specifies three primary setup and two operational modes. A station is a component that connects to the wireless medium.The simplest setup is a Basic Service Set (BSS), which comprises a number of stations that communicate with each other. When these stations are not connected to a wired network, they are referred to as an Independent Basic Service Set (IBSS). They operate in the *ad-hoc* mode such that each station is able to directly communicate with another within its reach.

In comparison, an Infrastructure Basic Service Set is a BSS, which has a base-station (BS). When a BSS works in the infrastructure mode, each station in the BSS goes through the BS for any communication with other stations. A BS is an access point (AP) if it is connected to a wired network.

When different BSS are connected via their BSs using a network backbone (also called distribution system (DS)), an extended service set (ESS) is formed. An ESS acts as one MAC-layer network. The specification does not mandate any particular technology for the DS.

In the infrastructure mode, a station needs to join a BSS to communicate. It obtains synchronization information from periodic beacons from the base station. It can either obtain this information by requesting it from the BS (active probing), or else it can wait for the periodic beacon from the BS. Before being able to send and receive data, the station has to go through an authentication and association process. The roaming function is not defined in the standard, but logical services have been described for this purpose.

The IEEE 802.11 standard not only defines a Medium Access Protocol (MAC), but also the related management protocols and services, and the physical layer. In this work, we will exclusively deal with the Direct Sequence Spread Spectrum (DSSS) physical layer operating at 2 Mbps on the 2.4 GHz band. The standard defines 11 channels for the US and Canada region, of which only three are non-overlapping.

Before introducing the medium-access mechanism, we describe some important timing intervals prescribed by the standard.

- *Short inter-frame space (SIFS)*: It is the shortest time interval. For DSSS at 2Mbps it is 10 μs. It is used between a frame and its acknowledgment. It is long enough for the sender to switch to the receive mode.

- *Slot time (Slot)*: a little longer than SIFS, it is the basic time unit for the binary exponential back-off algorithm spelled out in the standard.

- *Priority inter-frame space (PIFS)*: it is equal to SIFS + one Slot. It is used by the Point Coordinator to get higher priority in accessing the medium.

- *Distributed inter-frame space (DIFS)*: it is equal to SIFS + two Slot. It is used before starting a new transmission.

2.2.2 Distributed Coordination Function (DCF)

The DCF protocol allows stations to access the medium in a distributed manner. There is no central entity mediating use of the shared channel. Two access mechanisms are spelled out for the DCF: the Basic Access and RTS/CTS.

2.2.3 Basic Access Mechanism

The Basic Access scheme is a carrier sense multiple access with collision avoidance (CSMA/CA). When the MAC needs to transmit a frame, it physically senses the medium to check its status. If the medium is free, the station waits for an interval of DIFS to check that the medium remains free. If it is still free, the station sends its frame. Otherwise, the MAC selects a back-off value randomly selected from a contention window. This scheme is depicted in 2.1. The back-off value is decremented each time the medium is free for one slot time. If a collision happens, the contention window is set to twice its size and a back-off value is chosen from the new interval. After a successful transmission, the contention window is reset to a pre-set minimum value. The random back-off is also called after each successful transmission and each retransmission to reduce probability of collisions. The 802.11 MAC uses a positive acknowledgment scheme to detect collisions. Each unicast frame sent by the MAC has to be acknowledged by the receiver. If an acknowledgment is not received, the frame is retransmitted by the MAC layer. Broadcast packets are not acknowledged. Also, retransmissions are limited to a maximum number of tries, after which they are dropped.

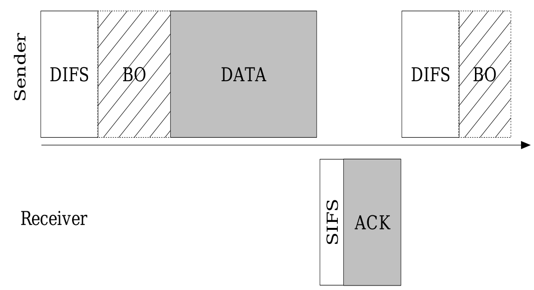

Figure 2.1: 802.11 DCF Basic Access Mechanism

2.2.4 The RTS/CTS Access Mechanism

In a wireless medium, the sender is not able to detect collision because it occurs at the receiver. If a packet collides at the receiver, the whole packet still needs to be transmitted and then re-transmitted when an acknowledgment is not received. In addition, stations in the receiver's surrounding may not sense a transmission from the sender. If any of these stations transmits, there will be a collision at the receiver. This is referred to as the Hidden Node Terminal problem. To circumvent the above problem and enable faster collision detection, the 802.11 MAC specifies a prior hand-shake. Whenever a station has data to sent, it first sends a Request to Send (RTS) frame. The destination replies with a Clear to Send frame (CTS). These two frames contain duration information of the forthcoming data frame. All neighbouring stations hearing these frames set a variable called Network Allocation Vector (NAV) to keep track of the availability of the medium. Checking the NAV before a transmission is also called a Virtual Carrier Sense mechanism. This protocol is shown in Fig. 2.2. The use of the RTS/CTS mechanism is enabled by the the attribute *dot11RTSThreshold* in the management information base (MIB) that specifies the minimum size of the frame requiring a RTS/CTS exchange.

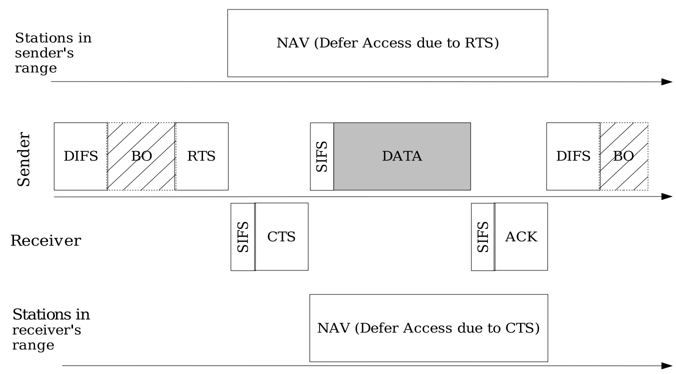

Figure 2.2: 802.11 DCF RTS/CTS Access Mechanism

2.2.5 Point Coordination Function (PCF)

Besides DCF, the 802.11 Specification also defines an optional polling-based access mechanism. PCF is a medium-access protocol built on top of DCF, which is controlled by a Point Coordinator (PC) and is usually implemented in an access point. Stations register themselves with the PC in order to access the medium. The PC periodically polls the stations for data and also delivers data destined to them. PCF alternates the contention-free period (CFP) with a contention period where normal DCF rules apply. At the beginning of the CFP, the PC sends a beacon using DCF. The beacon contains information about the expected duration of the CFP, which allows listening stations to set their NAV. During CFP, the PC uses PIFS as its delay period to prevent stations that did not hear the beacon or are not registered from gaining access to the medium. At the end of the CFP, the PC sends a contention-free end (CF-End) frame so that stations can reset their NAV. Implementation of PCF is optional and is not commonly implemented in commercial products. We thus do not study PCF further in this work.

2.3 ROUTING

This section introduces *ad-hoc* routing and provides some details on the Destination Sequenced Distance Vector (DSDV) and the Ad-hoc On-demand Distance Vector (AODV) routing protocols. Some understanding of DSDV is required to understand experiments that are explained later in this thesis.

2.3.1 AD-HOC ROUTING PROTOCOLS

Routing in *ad-hoc* networks has drawn considerable attention from the research community. The dynamic nature of such networks raises many challenges such as synchronization, adaptivity, fault-tolerance, *etc*. Today, management and operation of some of these protocols, notably Dynamic Source Routing (DSR), AODV, and Optimized Link State Routing (OLSR) have been formalized in IETF experimental Request for Comments (RFC) [11].

Ad-hoc routing protocols can generally be classified into two categories: proactive and reactive protocols. Proactive routing protocols collect information in advance such that it is available when need arises. Mechanisms such as periodic updates may be used to maintain fresh information. Examples include DSDV and OLSR. Such protocols are suitable for networks in which nodes have a low degree of mobility.

Reactive routing protocols on the other hand look for information only when it is required. For example, when a node needs to reach another node, routes are dynamically created as a result. These protocols are very useful for scenarios with high mobility Examples include DSR and AODV.

Below, we further describe DSDV, which is the principal routing protocol employed in our work.

2.3.1.1 Destination-Sequenced Distance Vector (DSDV)

DSDV [36] was one of the first routing protocols devised for *ad-hoc* networks. It is a table-driven proactive routing protocol, whereby each node keeps track of the next hop and the associated cost to all other nodes in the network. This cost represents the number of hops from the node to another one. In addition, a sequence number is kept to avoid the formation of routing loops.

Each node, periodically or when it detects a change in its neighbourhood, broadcasts a routing-table-update packet. This packet comprises a monotonically increasing sequence number generated by the source and also includes a routing cost initialized to one. Every other node receiving this packet updates its table, if necessary, increments the routing cost (*i.e.* number of hops from the source), and further broadcast the packet. When a node updates one of its route, it sets the source of the routing packet as the next hop on that route. DSDV assumes that all links are bidirectional. If a node receives the same routing update packet again, it will only consider it if it has a lower routing cost. Also, only routing update packets with a current or fresher sequence number are considered. If such a packet provides a better or fresher route, the last sender of the packet becomes the next hop for the node that generated the update packet. Each node is required to periodically send a routing-table-update packet reflecting its entire table. When a node detects that a neighbour has moved away or is no more active, it immediately broadcasts an update packet with a route metric of infinity for the neighbour. This causes nodes to purge related entries from their routing tables. To reduce unnecessary routing traffic due to time variation in the delivery of packets (*e.g.* forwarding of routes with higher hop counts), a settling time table is used. The settling time of a destination is a running weighted average over the most recent relevant route updates. The settling time indicates to a node when it can advertise a route *i.e.* when the settling time expires.

Due to the broadcast overhead, DSDV is only suitable for scenarios where mobility is very low [3]. For our purposes, however, this is suitable.

2.3.1.2 Ad-hoc On-Demand Distance Vector Routing (AODV)

Besides DSDV, we also experimented with Ad-hoc On-demand Distance Vector Routing (AODV) [37]. AODV is a reactive routing protocol, that creates routes on-demand when a packet needs to be sent to a new destination.

AODV draws certain features from DSDV such as sequence numbers and per-hop routing. Each node keeps information about each active destination, which includes the next hop, a sequence number for the destination, route cost, and the set neighbours that are using the node as the next hop to the destination. When a node needs a route to another node, it broadcasts a Route Request (RREQ) packet. RREQ packets create a reverse route for replies to return to the source. RREQ packets are matched with replies by the destination or any other node that has a fresher or better route than that specified in the request. The source of the request chooses, among route reply (RREP) packets, one that is fresher or provides a better metric with a valid sequence number. Routes at a node may timeout if no packet is sent through them. If a neighbour of a node is no more reachable, the node sends a RREP packet with a route metric of infinity to all neighbours for which the node is the next hop to the unreachable neighbour.

2.3.2 Routing in WLAN infrastructure

As with *ad-hoc* networks, in an infrastructure-based wireless network packets need to be routed between users who may be attached to different access points. Presently, most 802.11 WLAN infrastructure networks use a wired network (usually Ethernet) to connect the access points. For 802.11 networks, no routing protocol is specified for routing in the distribution system. Most networks rely on IP routing for this purpose. Wireless stations need to obtain a valid IP address a priori, typically by using DHCP [12]. A combination of MobileIP [13] and DHCP can be used to obtain seamless roaming. To support functions such as roaming, the IEEE has proposed a tentative recommended practice for inter-access point (802.11f)[50]. It defines information that needs to be shared between access points in the system as well as with upper layers to support several distribution system functions.

access points in a WLAN can also be connected via a wireless network. Such a wireless distribution system allows for data frames to be routed between access points at the MAC layer. This is possible if the four MAC addresses specified by the IEEE 802.11 standard are used. However, solutions tend to be vendor-specific and thus it is often recommended to purchase equipment from the same vendor to implement a wireless distribution system. Alternatively, routing protocols developed for *ad-hoc* networks may be employed to create a higher-layer distribution system, provided an appropriate tuning of various parameters is used.

2.4 Capacity issues with one-channel 802.11 MAC

Many researcher have studied performance of *ad-hoc* networks and have reported their poor scaling, theoretically and experimentally. Li *et al.* [24] show that in a chain of nodes separated by 200m and where the transmission range is 250 m, we only achieve $1/7$ of the maximum throughput. The authors identify three causes for the poor utilization. A channel utilization of 1/3 should be possible if we assume that a node cannot interfere with another not in its transmission range. In reality, however, a node can affect another one that is well beyond its transmission reach. For example, the transmission range of Lucent ORiNOCO with a data rate of 2 Mbps is 400 m whereas its carrier sensing range in an open space environment can be as large as 670m [54]. Similarly, in *ns-2*, which models a 914 MHz Lucent WAVELAN DSSS radio, a node has a default transmission range of 250 m while the carrier sense range is about 550 m. Because the 802.11 MAC is based on a CSMA protocol, the large carrier sense range leads to a further decrease in throughput.

Another reason for the low is the unbalanced contention experienced by the nodes in a chain. If the first node in a chain is the sender, it only suffers contention from one side. It is able to send packets at a higher rate than can be processed by middle nodes in the chain. Excess packets are dropped, leading to wastage of channel resources.

Li *et al.* also mention the anomalous behaviour of the exponential back-off algorithm in 802.11 MAC. When small packets, such as RTS and CTS, compete for medium with long data packets, the contention window of the nodes trying to send the control packets increases exponentially. This results in excessive back-off and loss of transmission opportunity when the medium eventually becomes free. The authors point out the inability of 802.11 MAC to discover a scheduling that maximizes capacity. Similar problems are also mentioned by Bharghavan *et al.* [2].

Xu *et al.* [55] evaluate the performance of the 802.11 MAC when using TCP traffic over multi-hop *ad-hoc* networks. They describe performance implications of the exposed node problem, which occurs when a node does not transmit because it can sense transmission of the sender, even though it would not interfere with the packet being received at the destination. This particular problem causes intermediate nodes in a chain to experience frequent route failure, especially when trying to send control packets. Route failures in turn leads to TCP timeout, which results in bursty TCP behaviour. The problem is due, as described earlier, to the 802.11 binary exponential back-off algorithm. The authors also report unfairness when there are competing TCP flows caused by the exposed node phenomenon when using 802.11 MAC. The MAC operates in such a way that the last node that is able to send a packet is favoured for the next transmission. A node experiencing high level of contention and trying to initiate communication will often drop packets and report route failures. Repeated failures will severely hamper throughput of a TCP flow going through this node. Some solutions to alleviate the exposed node problem have been proposed in the context of a power-controlled MAC. For further details, the reader is referred to [4, 31].

Tschudin *et al.* [48] claim that TCP-based applications may not be usable on multi-hop *ad-hoc* networks using IEEE 802.11 with paths of three hops or more. This phenomenon is explained by the constant congestion assumed by TCP because of packet drops.

To better understand performance implication of the 802.11 MAC protocol, we need to differentiate between transmission, carrier-sense, and interference range. Transmission Range(R_t) is the maximum distance within which a packet can be correctly exchanged between two nodes, assuming there is no interference. It is determined by the transmission power and the radio propagation characteristics of the environment. Carrier Sense Range (R_{cs}) is the maximum distance at which a signal can be sensed by a receiving node. This value depends on the antenna sensitivity. Finally, Interference Range (R_i) is the distance between a receiver and arbitrary node whose transmission will corrupt the packet being heard at the receiver.

A packet is successfully captured if it's signal to noise ratio (SNR) exceeds a certain threshold. The *ns-2* simulator uses the Two-way Ground model to simulate signal propagation in open space. In the immediate zone around the transmitter (Freznel zone), the received power is inversely proportional to d^2, where d is the separation between the sender and the receiver. Outside the Freznel zone, the received power is inversely proportional to d^4. The power of the signal received at a station is given by:

$$P_r = P_t G_t G_r (h_t^2 h_r^2)/d^4 \tag{2.1}$$

P_t is the power of transmission, G_t and G_r are the antenna gains of the transmitter and the receiver respectively, and h_t and h_r are the antenna heights.

Assuming that there is no noise, to successfully capture a packet, the ratio of the power received from an interfering node and the power received from the source should exceed $(capturethreshold)^{1/4}$. The value 4 is the signal-attenuation coefficient for the Two-way ground model. This capture threshold is usually set to 10.

Xu *et al.* [54] show that a large interference range decreases performance of an *ad-hoc* network. If the distance, d, between the sender and the receiver is larger than $0.56R_t$, a node farther than R_t can interfere with the packet exchange. The power needed to corrupt an ongoing transmission is less than what is needed to successfully transmit. They also point out that when R_{cs} is larger than $d + R_t$, the RTS/CTS handshake is not useful. Finally, they confirm that a large carrier sense range reduces network performance by preventing channel reuse, because some nodes in the R_{cs} range will defer despite the fact that their transmission would not cause any interference.

Ye *et al.* [57] study spatial reuse properties of 802.11 MAC. They show that RTS/CTS cannot completely address the hidden node issue. This is the case when the interference range is larger than the transmission range. When the interference range is smaller than the transmission range, RTS/CTS over-reserves space. Some nodes receiving RTS or CTS packets are prevented from transmitting, even though their transmissions would not affect the transmission in question.

Other papers have identified various additional problems in the 802.11 MAC. For example, in [39], the authors describe the blocking and false blocking problem due to the RTS/CTS exchange. A node is blocked when it cannot transmit a packet. For example, when a node receives an RTS, it sets its NAV to the duration of the expected data packet and it will not reply to any packet irrespective of the state of the medium during this period. If no data packet follows the RTS, a node may repeatedly send RTS packets to the blocked node, assume high channel contention, and unnecessarily back-off.

Tay and Chua [47] developed a model to evaluate performance of the 802.11 basic access mechanism. The authors make the observation that decreasing throughout not only results from collisions, but also happens when nodes waste too much time in back-off. They refer to the phenomenon as *back-off thrashing*.

As we can see, several factors contribute to the poor performance of 802.11-based multi-hop networks. First of all, because the channel is a shared medium, throughput per node decreases as there are more users per area-unit. In addition, the CSMA protocol and the exponential back-off algorithm can lead to under-utilization of the medium, which in turn results in high variability in throughput. While various solutions have been proposed to alleviate these problems, many of these would require changes to the MAC algorithm. A segment of the research community has investigated the use of multi-channel protocols to increase capacity of wireless networks and *ad-hoc* networks in particular. In the next section, we briefly describe the main protocols resulting from this research.

2.5 Multi-channel protocols

A number of multi-channel access protocols have been proposed for *ad-hoc* networks. In this section, we describe some the main protocols. The intent of these protocols is to increase capacity by enabling maximal spatial re-use of available channels in a distributed manner.

Dynamic Private Channel (DPC) [18] is a connection-oriented multiple channel MAC. Stations require multiple radio ports, one of which permanently listens to the control channel (CCH). Access to the CCH is contention-based. A station wanting to send data to another makes a request (RTS) on the CCH indicating duration of the connection. The latter is restricted by a maximum value. The request includes the channel code of a free data channel. The receiver replies with a Reply to RTS (RRTS) packet indicating whether it accepts the data channel proposed. If not, it proposes a new data channel. When the negotiation ends, both stations shift to the data channel and the receiver confirms the reservation of the radio port by sending a CTS frame. With this protocol, the improvement in throughput is restricted due to blocking.

Nasipuri *et al.* [33] propose a protocol whereby the available bandwidth is divided into N channels (frequency domain (FDMA) or code domain (CDMA)). Each station continuously monitors all channels (e.g. when a receiver uses multiple CDMA codes). Channels are marked as free if the received signal strength does not exceed the sensing threshold. When a free channel is available, the station waits for a long inter-frame space to see if the channel remains free. If it is free, the sender transmits. If the station successfully transmitted earlier, it reuses the same channel. In this scheme, even though a node can simultaneously sense multiple channels, it can only send one data frame at a time. The authors report better throughput than using a single channel, especially when exploiting multiple receivers. In a stationary network of $n \times n$ area with n^2 nodes, when the available frequency band is divided into 20 channels (assuming no guard-bands used), the authors show more than a two-fold throughput improvement compared to the use of one channel. The authors also show better delay behaviour for medium and high offered load.

Samir *et al.* [40] extended the work of Nasipuri *et al.* [33] proposing an RTS/CTS exchange. The sender includes its list of free channels in the RTS packet. Upon receiving the RTS, the receiver generates its own list of free channels. It compares the two lists, and picks the best channel w.r.t. the received signal power. It sends back this information in the CTS packet. If no such channel is available, the receiver does not reply. The sender times out and tries again after a back-off. If the RTS/CTS exchange takes place, the data frame and its acknowledgment are send on the selected data channel. In this protocol, while nodes can receive packets on the different channels in parallel, they can only transmit on one channel at any time. The authors report a throughput improvement in a stationary network versus a single channel network that uses all the available bandwidth. The increase in throughput is limited due to the requirement of one transmission at any time and the comparatively reduced capacity per channel. Also, at medium and high offered load, the average packet delay is smaller than the single channel case because there are fewer collisions and fewer retransmissions. The proposed protocol is said to work best with 4 to 7 channels,

Wu *et al.* [53] propose a protocol similar to the one designed by Samir *et al* [40] called Dynamic Channel Assignment (DCA) where the channel is randomly selected from the list of available free channels. In addition to the RTS/CTS exchange, the sender confirms the channel reservation by sending a RES packet including the channel identifier. This allows the sender's neighbours to update the status of the to be used data-channel.

While the solutions described above necessitate the use of multiple radios, there exist other approaches where a single interface. In [45], the authors put forward the Hop Reservation Multiple Access (HRMA) protocol for slow frequency hopping spread spectrum (FHSS) systems. The sender sends an RTS packet to the destination on a particular frequency hop. If the receiver replies with a CTS, the two stations stay on the same frequency for data transmission, while the other nodes continue to hop according to a predetermined sequence. This allows other nodes to communicate on the other channels. With this protocol, performance may be poor if other devices were to use a different hopping sequence [18]. Data transmissions may be long, and thus subject to interference from these other devices.

Recently, a new MAC protocol (MMAC) was presented by So and Vaidya [42], which requires only one radio interface. No separate control channel is necessary, but all stations need to be synchronized. The proposed scheme uses a similar approach as the power saving scheme (PSM) of IEEE 802.11. Every station periodically sends out a beacon that allows other stations to synchronize. In the beginning of a beacon interval, an ATIM (Ad-hoc traffic indication message) interval is defined during which all nodes are active. If a station has to receive data during the beacon interval, it is informed of this fact through an ATIM packet. The receiver acknowledges by sending back an ATIM-ACK. Both stations will thus stay awake during the interval to exchange the data packet. A station that is not to receive any data may go in sleep mode. During the ATIM window, all stations tune to the default channel. Channel selection takes place during this period so as to allow other neighbouring stations to keep track of the availability of channels. When the sender sends the ATIM packet it includes its preferred channel list. The receiver compares it to its own list and includes the chosen channel in the ATIM-ACK packet. The sender then replies with an ATIM-RES packet confirming the channel, and thus allows neighbours to update channel status. If the sender is unable to select the channel proposed by the receiver, it will wait to next beacon to negotiate another one. It is also possible for the sender to shift to the channel of the receiver and transmit the packet when the channel becomes idle. For multi-hop networks, MMAC show slightly better throughput than DCA.

A number of MAC protocols based on the Code Division Multiple Access (CDMA) using spread spectrum have been put forward for ad-hoc network. CDMA allows multiple concurrent conversations to take place at the same time by spreading the corresponding signals using pseudo-random noise codes. However, because we are focusing on 802.11 DSSS, we do not discuss such protocols further in this work.

The schemes described above suffer from various weaknesses. Most of these schemes require special purpose hardware or the use of technologies such as CDMA. With current 802.11 hardware, some of these schemes may not be very efficient. For example, with the availability of only three non-overlapping channels, if one is reserved for control traffic, only $2/3$ of the bandwidth is available for data traffic. Besides, with present hardware, the channel switching time is large, making it difficult to implement a dynamic protocol . In addition, the optimum choice of the division of the available bandwidth between control and data channels depends on the level of contention and the distribution of packet length. As pointed out by Xue *et al.* [56], dynamic channel division is not practical with hardware available today.

Schemes that divide the available frequency band into smaller channels opt for reduced bandwidth per channel, and thus reduced throughput. While some of this reduction is compensated by the gain in spatial re-use of the channels, smaller channels lead to a higher level of blocking. Such problems arise because transmissions can take longer, resulting in more waiting time before a sender and a receiver can undertake a packet exchange. Assuming that a node can undertake only one data exchange at a time, another node willing to communicate with a sender or a receiver has to wait longer for the data exchange to complete.

In parallel to our study, Adya *et al.* at Microsoft Research have worked on multi-interface solutions for community wireless mesh networks. They propose the Multi-radio Unification Protocol (MUP) [1] at the link-layer in order to achieve better utilization of the spectrum. Their goal is similar to ours (*i.e.* leverage commodity 802.11-compliant products and not require any changes to higher-layer protocols). MUP abstracts multiple interfaces by presenting a single virtual MAC address to higher layers. At initialization, the available interfaces are tuned to non-interfering channels. This assignment is fixed for the lifetime of the network. For each neighbour a node keeps track of the MAC address and channel-related information. If the neighbour has multiple interfaces and supports MUP, the protocol proceeds to discover additional information. For any packet transmission, MUP selects the interface with the highest channel quality. This quality is a measure of the roundtriptime of probe packets sent to that neighbour. Each node periodically appraises the need to shift channel to communicate with an MUP-capable node. Because probes need to be sent and acknowledged immediately, they require support for prioritized traffic, which is expected to be implemented in forthcoming 802.11e-compliant products. The authors evaluate the system through simulations in *ns-2* and show important performance gains when compared to a single channel network. They simulate TCP traffic in the presence of a varying background UDP load. Their experiments show that simple striping on multiple interfaces outperforms MUP if all nodes are MUP-capable. However, with a mix of legacy nodes, MUP performs better than simple striping. Striping performs poorly in the latter setup poorly due the imbalance of the load on the interfaces. This in turn causes TCP reordering and thus affects the TCP congestion window.

2.6 Hybrid Ad-hoc Networks

Several researchers have proposed variants of *ad-hoc* networking optimized for specific scenarios. In this section, we discuss some of these initiatives whereby *ad-hoc* networks are either used in a non-mobile context or are mixed with stationary nodes. Two projects described in this section are based on actual deployment and experimentation.

2.6.1 Infrastructured ad-hoc network

Mobile nodes have limited power supply and users are not rewarded for relaying other users' packets. Instead, if the node is close to a frequently used resource (*e.g.* a file server or a gateway), its power will be unfairly consumed. Based on these observations, Lundgren *et al.* [25] propose to mix some fixed stable wireless nodes with normal mobile nodes. The stable nodes have more power available (*e.g.* they may be connected to a power outlet or may draw power from a car) and thus would make such *ad-hoc* networks more appealing and more efficient for mobile devices. Power saving would be achieved by giving a preference to routing through the resource-rich nodes. The special nodes referred to as peudo-base-stations (PBS) run an *ad-hoc* routing protocol as do other mobile nodes. The proposed extension to AODV called ISAIAH, uses a cost function that takes mobility, power availability, and number of hops into account. These factors are weighted so as to give PBS priority in routing packets. With ISAIAH, the route chosen may not have the least number of hops, but will prefer low mobility high-power availability nodes.

2.6.2 K-Hop Hybrid Network

Miller *et al.* [29] present the idea of combining an *ad-hoc* network with a WLAN infrastructure. Such a hybrid network consists of mobile devices, relay nodes, and access points. The main goal is to extend the infrastructure (*e.g.* providing connectivity around a shopping mall). access points can be connected via a wired or wireless network. The authors develop a routing protocol based on AODV with proactive and reactive features. An access point keeps track of reachable mobile nodes using a beaconing and a join process, while nodes reactively discover routes to other nodes they want to contact. For efficiency reasons, a mobile node can only contact another one if it is less than K hops away from it. Longer routes (up to 2K) are possible, by going through two access points. The value of K can be changed to adapt to prevailing conditions.

Similar ideas have been proposed for increasing spatial re-use and decreasing power consumed in cellular networks (See [17, 26] for some related work).

2.6.3 MIT Ad-hoc Grid Networking

While the two previous studies suggest the use of variants of *ad-hoc* networks, the latter have begun to be used without modification in stationary deployments. This the case with the implementation of wireless Community Area Networks. The intent of such deployments is to provide connectivity by freely sharing resources in areas where there is no infrastructure or where access cost is high. Little information is available on the performance of these implementations.

At MIT, the Grid Ad-hoc Networking project [34] employs *ad-hoc* routing protocols in a fixed indoor and outdoor wireless network setup. This experimental testbed has been set up to study a production-quality 802.11b multi-hop network. Despite having a well-connected network, the authors report poor performance using DSDV [7]. The results obtained were below the expectations taking the shared nature of the radio media into consideration. They point out the inability of DSDV to find routes with high delivery rates. DSDV tries to find the shortest path, and if two paths have identical route metrics, one is arbitrarily chosen. Interestingly, the shortest path may not necessarily provide the highest throughput. The study shows that there exist a wide distribution of delivery rates. Links with medium delivery rates tend to be asymmetric. As such, such link may relay routing packets, but may be unable to route data packets.

In the context of the same project at MIT, cost metrics were studied in order to find high-quality paths in a network where the link-level loss varies in time [6]. The authors point out that minimizing the hop metric leads to the selection of links with longer distance. This in turn amounts to choosing links with smaller signal strength and thus higher loss ratio. As an alternative, the authors propose the ETX metric, which finds the fewest expected number of transmissions for a packet to reach its destination. The number of transmissions is estimated from the link-layer loss ratio in both directions. The link behaviour is evaluated by broadcasting a set small-sized packet and measuring the results at the receiver end. A two-fold improvement is reported over the minimum hop-count for multi-hop paths.

All the work mentioned in this section suggest that there exist applications where a hybrid or fixed wireless network would be of a great value. We also see that the large amount of work carried out for routing in the *ad-hoc* networks can be leveraged in semi-ad-hoc networks. Nevertheless, the above work is still based on single channel networks, and as such, can only yield limited improvement. Even though such networks have non-mobile components, no attempt is made to take advantage of the existence of multiple channels. In addition, users require special software to benefit from such systems.

3 Using Multiple Interfaces

The widespread adoption of 802.11-based products has led to a drastic reduction in their cost. This low cost allows us to experiment with configurations that were deemed costly before, but that have the potential to address capacity issues in multi-hop wireless networks. In a traditional multi-hop wireless network, all nodes operate on the same channel. As such, when a specific node uses the medium, all others in its interference zone are affected insofar as they need to defer their transmission. For example, as shown in Fig. 3.1(a), when node 2 is transmitting a packet to node 1, node 3 cannot initiate a send because the medium is busy. Similarly, node 0 cannot send a packet, as it would collide with the packet being received at node 1, causing the latter to be dropped.

Our work proposes a solution that can potentially reduce the level of contention on any particular channel. We investigate a simple approach that involves leveraging multiple interfaces on a single node such that the node can communicate simultaneously on different channels.

By using multiple interfaces tuned to different channels, the level of communication parallelism is increased and contention for the shared medium decreases. More communication can take place at the same time with less interference between nodes, which results in increased throughput. For example, in Figure 3.1(b), communication between node 0 and node 1, node 1 and node 2, and node 2 and node 3 can take place at the same time, giving a three-fold increase in capacity. However, there are only three non-overlapping channels available for 802.11b/g-based networks. In addition, the carrier-sense range of a transmitting station is significantly larger than the transmission range. When these channels are reused, interference is not completely eliminated, and thus reduces throughput. We also study alternate configurations whereby we try to reduce interference by re-using channels at larger distances. An example is depicted in Figure 3.1(c). A channel is used for two hops (*e.g.* channel 1 between nodes 0, 1, and 2). This has the effect of spreading channels further apart, ideally to a distance larger than the interference zone. In this configuration if the distance between the nodes is 200 m and we have three channels, a channel is reused at a distance of 800 m. The level of parallelism achieved in the three configurations is shown in Fig. 3.2.

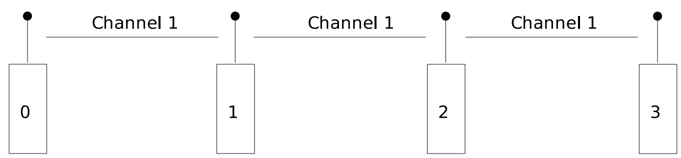

(a) One channel network

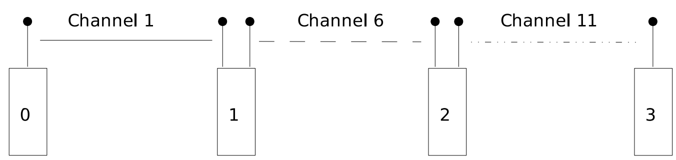

(b) Dual-interface nodes at each hop

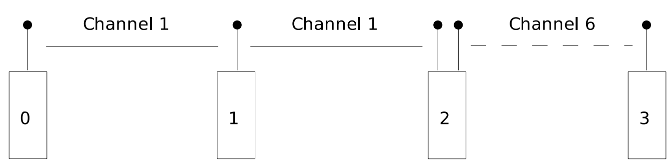

(c) Dual-interface nodes at every two hops

Figure 3.1: One channel vs. multi-channel setup

In this work, we investigate nodes with a maximum of two interfaces. Even with

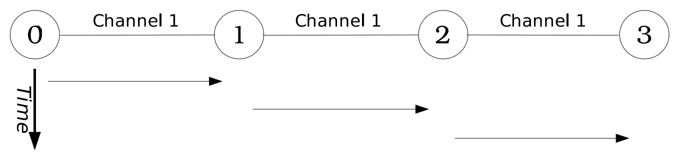

(a) One channel network

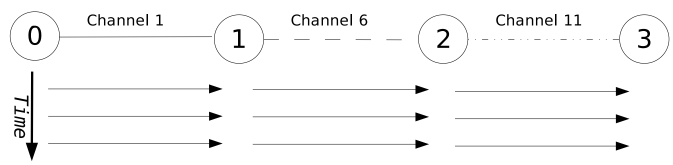

(b) Dual-interface nodes at each hop

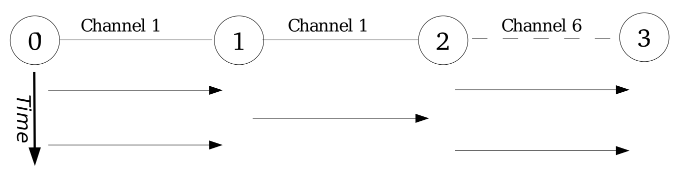

(c) Dual-interface nodes at every two hops

Figure 3.2: Parallelism

two-interface nodes, the approach raises a number of concerns that need to be addressed. The existence of only three non-overlapping channels entails careful allocation, if we are to get maximum spatial re-use.

Hereafter, in this document, we distinguish between an interface and a node. A node is a computing device, such as a laptop, that can have multiple radio interfaces. Unless stated otherwise, nodes with two interfaces operate on two different non-interfering channels.

3.1 Design and Implementation

When two or more interfaces are placed on a node, we are faced with a number of choices as to the way they are interconnected. Two main alternatives are bridging at the link layer and connecting the interfaces at the network layer. While bridging at the link-layer may be acceptable in a wired network, in its wireless counterpart traffic unnecessarily relayed can significantly degrade performance. In this work, we adopt a more generic solution by connecting the interfaces at the routing level. The individual interfaces act as separate logical entities with different addresses for routing purposes.

Because the interfaces are tuned to different channels, they need to be integrated to ensure proper operation. An implementation of such a system requires changes at two main places. First, we need to adapt the ARP module such that an interface on a multi-interface node can reply to a request for the MAC address of one of its counterparts. This change is necessary because the other interfaces are on different channels and thus cannot reply to such a request.

Second, we need to integrate forwarding services of the available interfaces. A packet from one interface may need to be forwarded on one or more of the other interfaces. This is normally achieved by developing a single routing module that associates routing knowledge with interface and channel information. However, rather than develop a new routing module, we opt to use separate routing modules for our design that are appropriately linked. With this approach we can readily use existing routing modules that have already been implemented (See [38] for an example). As we will see, very few changes are required to integrate such routing modules.

In the following section, we further expand on this design and describe our implementation of a dual-interface node in *ns-2*. We assume the use of the DSDV routing protocol. However, any other routing protocol can be adapted for such operation.

3.1.1 NS-2 IMPLEMENTATION

Our implementation is based on the wireless *mobilenode* implementation in *ns-2* [15]. In *ns-2*, a node is organized as a set of independent layers. These layers are implemented as C++ objects, which are combined using the TCL scripting language. These components simulate behaviour of the corresponding actual protocols. Essentially, a simple wireless node comprise a radio interface, a MAC layer, an interface queue, a link-layer, and a routing agent. A network interface can be set to use any channel object of choice. We assume that there is no interference arising from the antennas being close to each other. In reality, however, signal power leakage may give rise to interference despite the use of independent channels [1]. Port and address classifiers allow packets to be directed to the correct destination application component. Packets travel through these components following a set of chosen protocols and algorithms. Transport agents (*e.g.* UDP or TCP) and application components sit on top of the basic *mobilenode* to form an active wireless node.

The organization of these components can be seen in Fig. 3.3 (adapted from [15]). As shown in Fig. 3.3, a dual-interface node consists of two independent radio adapters with separate protocol stacks up to and including the routing modules. While separate, the two interfaces are coordinated in order for the node to appear as a single entity. We associate two individual basic *mobilenodes* by linking them at the routing level and the data link layer (indicated by the bold arrows in Fig. 3.3). These links represent the passing of information between the two sets of components. At the ARP level, the IP-to-MAC mapping information has to be shared. At the routing level, packets need to be forwarded between the peer routing agents. In our implementation two basic nodes are superimposed (*i.e.* they have the same location). We assume that there is a layer that abstracts the two interfaces below the transport and the application layers. As such, an application always binds to one of the two interfaces. Alternatively, at the cost of

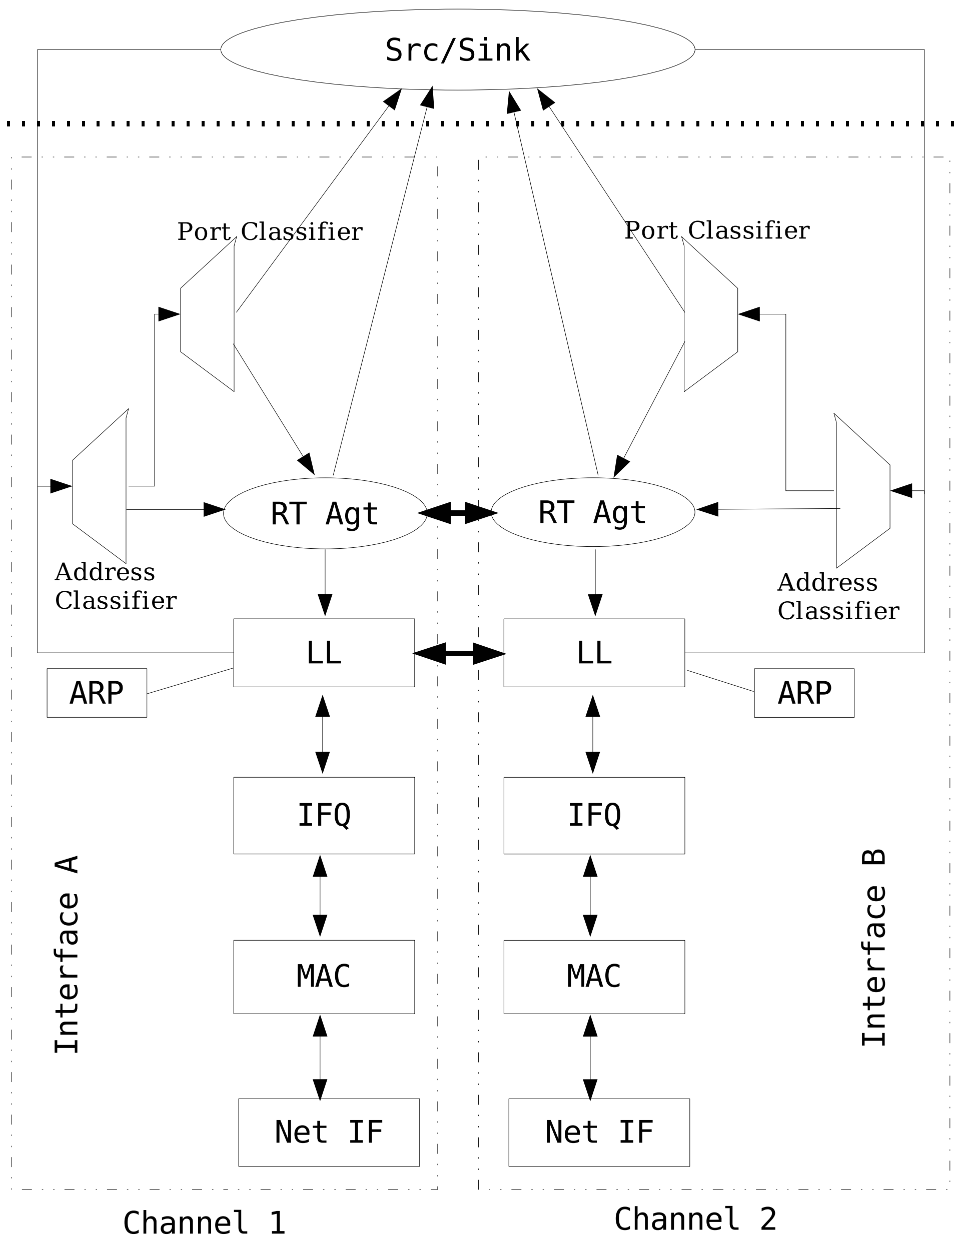

Figure 3.3: Two interconnected ns-2 mobilenodes

some complexity, it is possible to implement a striping mechanism to provide some load-balancing. However, striping may introduce new problems, such as packet reordering in TCP flows.

During operation, when an interface generates a routing-update packet, it not only sends it on the medium, but also passes it to the other interface on the same node. The latter processes the update like a normal routing-update packet. However, unlike when it receives such a packet from a different node, the interface does not increment the hop count for routes found in the packet. This allows the dual-interface node to count as a single hop. An interface that finds a better route from the packet received from its counterpart will record the fact that the route originated from the other interface. When an interface receives a data packet not destined to the node, the routing table entry for the destination is retrieved from both routing modules. The entry with the highest sequence number is chosen. If both have the same sequence number, the one with the smallest hop count to the destination is selected. If both sequence number and hop count are equal, the interface making the routing decision checks where the route originated. If it came from the other interface, the data packet is passed over to the latter. If not, the interface uses the route in its own table to forward the packet.

We also designed an infrastructure-mode in which user nodes do not relay each others' packets. Instead, user nodes only communicate with others via their access points. The access points form a backbone that caters for inter-access point routing. When an access point has two radio interfaces, both of them act as *i.e.* service providers as well as participate in the backbone. For such a scheme to work, we assume that access points can attach to other access points as clients while still being able to service user nodes.

In our design, when a user node associates with an access point, the latter immediately schedules a routing update packet to other access points in the network. The same procedure is applied when a node disassociates from an access point. For such a scheme to work, various parameters of the routing protocol need to be tuned to avoid stale entries and mis-routing.

4 Evaluation and Analysis

In this section we evaluate the multi-interface system proposed by running and analyzing a variety of experiments. We test different characteristics of our system in various configurations in order to ascertain benefits and potential problems. For each experiment, we compare behaviour using both the Basic Access and the RTS/CTS Access schemes. Our motivation is that RTS/CTS handshake can easily be controlled by a user. Our presentation should not be seen as an attempt to devalue its use, as it is very effective during high contention periods provided that any hidden node can be heard by receiving nodes. This means all nodes are within two hops of each other.

4.1 Simulator Setup

We carried out a number of simulations using *ns-2* [14] with the Monarch project wireless extension [30]. We used the version 2.1b8a of the simulator. *ns-2* emulates the operation of 914 MHz Lucent's WaveLAN radio using DSSS at the physical layer with a data rate of 2Mbps. For all our simulations, we use wireless nodes that use a prioritized interface-queue with a buffer of 50 packets. All nodes use an omni directional antenna. In *ns-2*, signal propagation follows a combination of the Free Space and the Two-ray Ground model [15]. The default parameters yield a transmission range of 250 m and a carrier-sensing range of approximately 550 m. We simulate both the 802.11 Basic Access and RTS/CTS access methods. Given the amount of processing involved, unless otherwise stated, each experiment is repeated at least 5 times. Measurements are taken after the routing information has settled. We also reduce the rate of DSDV periodic updates such that it does not affect our measurements. This tuning is reasonable because in our system, nodes do not move. Each measurement is obtained after execution of the simulation for 300 seconds. Unless stated otherwise, the distance between the nodes is set to 200 m. We use constant bit rate (CBR) over UDP as source traffic with packets of length 1500 bytes. In our simulations, all measurements are taken at the application level.

4.2 CHAIN TOPOLOGY

A chain topology is a network setup in which nodes are linearly aligned. An example of such a setup is a Roadside Information Network, which can be built by mounting wireless nodes on light poles. Such a topology is easier to analyze and thus helps understand more complex scenarios, which we experiment with later in this thesis.

4.2.1 THROUGHPUT

The three configurations tested are shown in Fig. 4.1 where each link labeled with the channel number. In Fig. 4.1(a), all nodes are on the same channel. This is the traditional *ad-hoc* network approach and is used as a baseline for comparison. In Fig. 4.1(b), we alternate the three non-overlapping channels available at each hop. Lastly, in 4.1(c), we use a channel on two successive links, and then alternate.

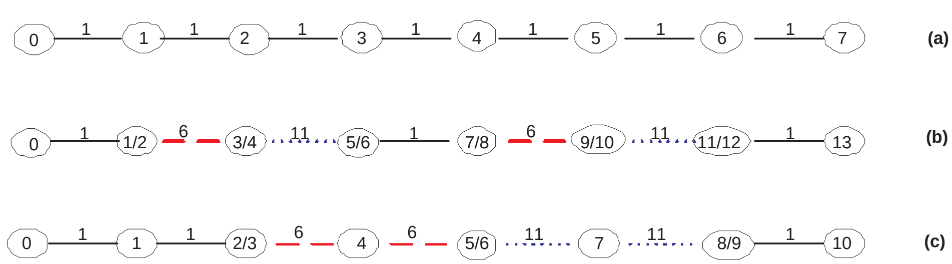

Figure 4.1: Chain topology (a) One channel (b) One hop/channel (c) Two hops/channel

We first analyze throughput as a function of the number of hops in the chain in the three configurations. In all experiments in this section, the source is the last node in the chain, while the destination is the first. The offered load at the source is 2 Mbps, which exceeds the available capacity, but allows us to simulate nodes with backlogged data.

The RTS/CTS mechanism has been devised to avoid the hidden node terminal problem and also to detect the possibility of collision faster. When a collision occurs, a long data packet still needs to be completely transmitted because the the sender is unaware of the corruption at the receiver. The collision assessment time is reduced if an RTS/CTS handshake is used. We thus start with the RTS/CTS access mechanism and compare the three configurations. We then evaluate results with the Basic CSMA/CA mechanism specified in the 802.11 protocol.

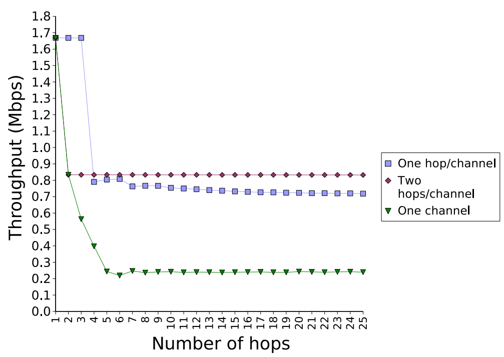

(a) RTS/CTS

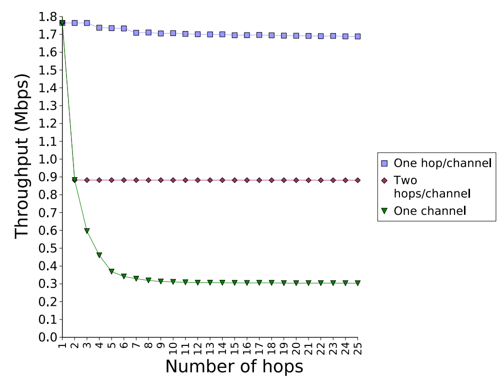

(b) Basic Acess

Figure 4.2: Throughput vs. number of hops

Figure 4.2 shows the results of our experiments. As we expected, both multi-channel configurations offer much higher throughput than the single-channel setup. In the one hop/channel with RTS/CTS (Fig. 4.2(a)), the throughput remains high for up to three hops, as each link uses a different channel. However, when four hops or more are required, the re-use of channels leads to a dramatic drop in throughput because the carrier-sense range is larger than 400 m (the distance between two links on the same channel). Surprisingly, the one hop/channel configuration offers slightly worse performance than the two hops/channel configuration for chains of length four or more. With more radio-interfaces in the one hop/channel configuration, we would expect to achieve higher performance in the one hop/channel case. The reasons underlying the lower throughput of the one hop/channel chain is explained in Section 4.2.3.

Throughput for the two hops/channel stabilizes at around 0.83 Mbps, and remains there for any number of hops. This stability is due to the reduced interference in this setup, as channels that are re-used are far apart, at a far longer distance than the carrier-sense range. However, given that each channel is used for two links, throughput is shared among them. In general with RTS/CTS, for two hops/channel with a chain of five hops or more, we are able to get greater than a three-fold throughput gain when compared to the single-channel chain. The same gain can be seen for one/hop per channel chain, its throughput decreases as we add more hops. In this configuration, the longer the chain, the higher the probability of packet collisions along the path. If we have a single flow, throughput in the chain topology is equivalent to the throughput of the slowest link in the chain. This is the link that suffers from the highest level of contention. In the one hop/channel chain, this corresponds to a middle link that is interfered with on both sides.

Figure 4.2(b) illustrates the results obtained when the RTS/CTS handshake is disabled. In a sparse network, if we ignore collisions, the removal of the RTS/CTS exchange should yield a slight improvement that is proportional to the size of RTS and CTS packets. While we do see a small improvement in the single-channel case and in the two hops/channel configuration (approximately 50 Kbps more than with RTS/CTS), throughput in the one hop/channel configuration for four hops and more almost doubles when compared to the case where RTS/CTS is used. From the two figures, it is clear that the RTS/CTS handshake introduces anomalies that limit performance of a setup with dual interfaces at each hop. Later in this chapter, we dwell into simulation details to find the causes of this anomaly.

4.2.2 Delay

We now look at the end-to-end delay behaviour in the chain experiments described above. Fig. 4.3 shows the total delay experienced by each packet that is not dropped by the network from the point it is sent by the source application to the point where it is received by the destination application. Fig. 4.4 presents the MAC component of the total delay (*i.e.* the time spent by each successively received packet in the MAC layer along the path). The MAC delay includes the propagation delay, back-off time, and time spent for retransmissions. At any node, we compute the time elapsed between the moment when a packet enters the queue and when it leaves the queue. We sum the time thus obtained for all nodes in the chain and subtract this value from the total delay.

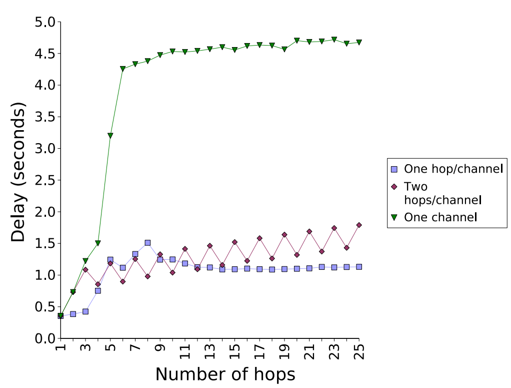

(a) RTS/CTS

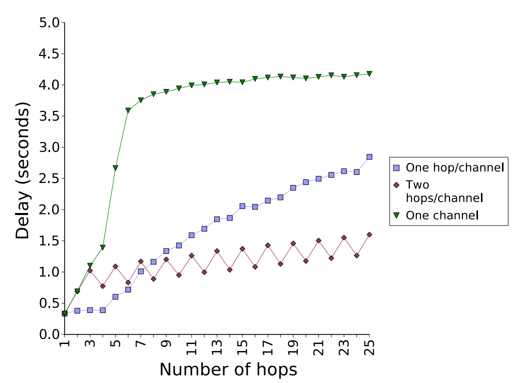

(b) Basic Acess

Figure 4.3: End-to-end delay vs. number of hops

We first observe in Fig. 4.3 that in the single-channel chain, the total delay can become very large as we increase the length of the chain. This effectively makes such a

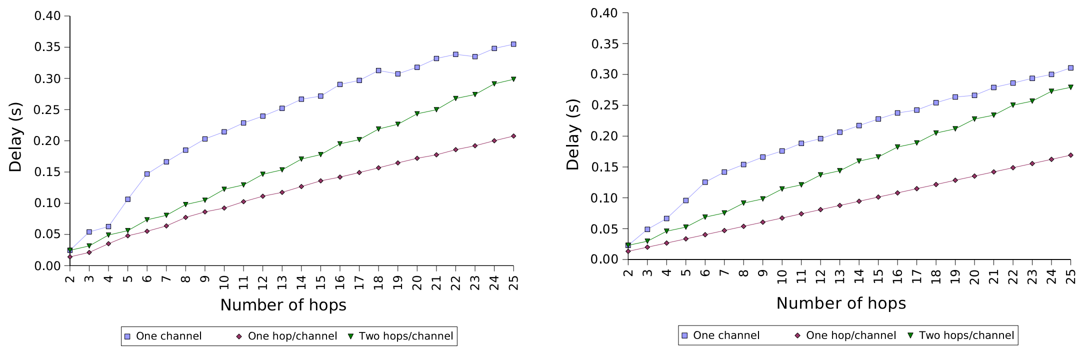

(a) RTS/CTS (b) Basic Acess

Figure 4.4: End-to-end MAC-layer delay vs. number of hops

network unusable for time-sensitive or interactive applications. We note that both multi-channel chains offer much lower delay. When taking both Fig. 4.4 and Fig. 4.3 into consideration, we see that an important largest of this delay can be attributed to queueing. However, it should be noted that queueing delay is affected by the MAC delay of earlier packets.

The step behaviour in the two hops/channel case is due to a mismatch of link-bandwidth when a newly created link uses a channel that is different from the previous link (*e.g.* when adding node 10 to the chain shown in in 4.1(c)). On such a link, the added node is able to send at a much higher rate than can be handled by the next hop. Queue build-up thus occurs both at the last node (because of saturation) and the next-to-last node. Packets need to go through both queues, hence the increased delay. If one more node is added to the chain with the same channel as the last link (*e.g.* adding node 11 to the chain shown in 4.1(c)), load exceeding available bandwidth is dropped at the source. In this case, queue build-up happens at one place only. With reference to Fig. 4.4, we observe that adding two new links on the same channel increases delay at the MAC level too because a packet needs to traverse two competing links. With both Basic and RTS/CTS access, the MAC-layer delay behaviour for the two/hops per channel configuration is similar and steady.

In the one hop/channel case with Basic Access, total delay increases as more hops are added (Fig. 4.3(b)). The same applies for the MAC component of this delay. For chains longer than 8 hops, two/hops per channel appears to offer better total delay. However, we see that one hop/channel provides lower end-to-end delay in the MAC layer. This shows that for longer chains queue build-up occurs more in the one hop/channel configuration. Queues of relaying-nodes along the chain tend to fill up from the source node up to a node where the packet arrival rate is lower than the link capacity. In the chain topology, this point is at interface 7 (See Fig. 4.1(b)). The one hop/channel setup offers the best total delay behaviour of the multi-channel configurations when the RTS/CTS mechanism is used, but the associated throughput is the lowest of the multi-channel scenarios. The delay stabilizes after 10 hops because most excess data packets are dropped at the first interface. From this point, the network would only let in a load that it can handle. When we have fewer than 10 hops, the delay increases steeply up to 8 hops, with a drop at hop 6. This uneven behaviour was traced to queue build-up at different places that are in turn caused by capacity mismatch between various links. Fig. 4.4 shows a steady increase in delay for one hop/channel, and thus confirms that the high variation in the total delay is principally due to queueing.

In the one hop/channel setup, when using the RTS/CTS access scheme, the link capacity mismatch that creates queue build-up is mainly due to short term unfairness introduced when small control packets compete with longer data packets and contention window reset when a packet is successfully transmitted. Our throughput results represent average values over a 300 seconds period. However, there is a high variation of throughput per link in the short run. Figure 4.5 shows data sent by interface 1 and 7 at the MAC level during one instance of the simulation where we have four hops. The throughput is sampled at 4-second intervals. We observe that the link-layer throughput of interface 3 and 5 closely follow the same behaviour as interface 7. This shows that access to the medium is very bursty and is a manifestation of unfairness in the short run.

4.2.3 Burstiness Analysis

We now explain the cause of this burstiness in throughput. We consider a chain of four hops where each intermediate node uses two interfaces. This scenario is depicted in Fig. 4.6.

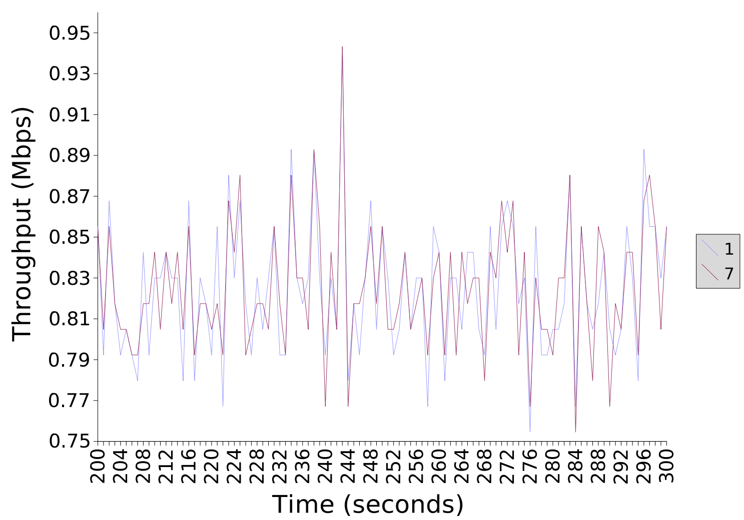

Figure 4.5: Throughput burstiness in a chain of a length 7

We first point out that the closest distance between nodes on the same channel is 400 m in the one hop/channel case. As stated previously, the carrier-sense range of 802.11b wireless adaptors is about twice their transmission range. In *ns-2*, the carrier-sense range is about 550 m whereas a transmission can reach a distance of 250 m. In our one hop/channel setup, if a node A sends a packet to another node B, which is 200 m away, external interference can only come from another node on the same channel that is 400 m away.

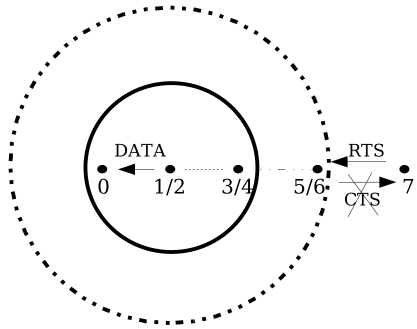

Figure 4.6: Anomaly when using RTS/CTS

In *ns-2*, the interference model is such that if a neighbour sends a packet to a node, the packet will not be interfered with by any other transmission from a node at 400 m if the transmission of the neighbour starts first. Otherwise, a collision occurs at the receiving node. Collisions are not the principal cause of the performance degradation. The main reason for the low throughput when using RTS/CTS lies in the algorithm of the 802.11 MAC, and in particular the exponential back-off mechanism. In the scenario shown in Figure 4.6, consider the situation where interface 1 is sending a data packet to interface 0. It should be noted that only interfaces 0, 1, 6, and 7 are on the same channel. Let our data packet be 1500 bytes long including various overheads (UDP, IP, MAC headers, and overhead from lower layers). In comparison, RTS and CTS are only 40 and 39 bytes long, respectively. The data packet is much larger than these control packets. When interface 1 is sending a data packet to 0, interface 6 would sense the medium as busy and would not transmit a packet during this time. This implies that for any RTS packet sent by interface 7 and correctly received by 6, the former will not receive a corresponding CTS back. Interface 7, not receiving a CTS, assumes contention. It doubles its congestion window, and selects a new random slot for a retry. Several of these back-offs lead to an increasingly larger back-off window. In contrast, a node that successfully transmits uses the minimum contention window for the back-off selection. This gives the latter higher priority in the short-term. If the retry threshold is exceeded the packet is dropped. In addition, larger back-off windows result in a potential under-usage of the channel. When the channel eventually becomes free such that interface 6 is able to send back a CTS, interface 7 may be in back-off mode. The capacity of a chain of length four with three channels (as shown in Fig. 4.6) is determined by the capacity of the link 7-6. Its capacity in turn depends on the interaction of the data on this link with that on the link 1-0. The RTS/CTS problem generates a very bursty throughput per link as shown in Fig. 4.5. The variation is explained by alternating periods of time where link 1-0 monopolizes the medium because it is sending a long data packet. Node 7 periodically enters a back-off mode and remains there until the retry threshold is reached and the packet dropped or until its RTS gets a corresponding CTS from interface 6. The latter is able to do so during those short periods of time when the medium is free because interface 1 is in a small back-off because it successfully sent its last packet. This particular problem is reported by Li *et al.* [24] in the context of a single channel. We observe that the two hops/channel configuration keeps channel reuse far enough apart to avoid this problem.

4.2.4 Performance as a function of load

We selected the 7-hop case to further investigate performance and problems observed. In doing so, our aim was also to derive the maximum throughput achievable with each configuration, even though not at a coarse granularity. We use a step increase of 100 Kbps in the offered load. All parameters are same as the previous chain experiments, with the exception of offered load, which varies from 200 Kbps to 1800 Kbps.

In Figure 4.7(a), once more we see a clear improvement when using multi-interface

	RTS/CTS	Basic Access
One channel	0.30	0.40
One hop/channel	0.80	1.71
Two hops/channel	0.83	0.88

Table 4.1: Maximum throughput per chain configuration (Mbps)

nodes. The peak throughput for the different configurations is given in Table 4.1.

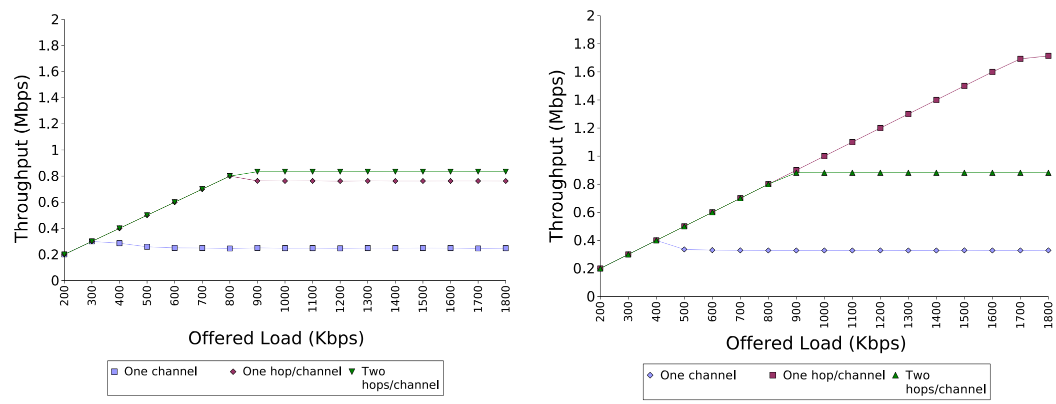

(a) RTS/CTS

(b) Basic Access

Figure 4.7: Throughput vs. Offered load

Under saturation load, the two hops/channel configuration offer higher capacity than the one hop/channel method when using RTS/CTS even though the former employs fewer interfaces. This the same observation we made for experiments shown in Fig. 4.2. The low throughput is the result of the unfairness problem mentioned earlier. Fig. 4.7(b) show the average throughput obtained with the three configurations when only the Basic Access mechanism is used. As pointed out in our first experiments, throughput of one hop/channel is higher than the other two options when RTS/CTS is not used. Unlike CTS, ACK packets are sent after SIFS without doing a carrier sense on the medium (Section 9.2.8 of the IEEE 802.11 Specification [19]). This explains the much higher throughput in the Basic Access case. In Fig. 4.6, consider the case where interface 1 is sending a data packet to 0. Interface 6 is able to capture a packet from interface 7 if this packet's transmission begins before any other packet's interference. After reception, interface 6 will not perform a carrier sense on the medium before sending back the ACK. Node 7 can thus proceed to the next packet transmission. This leads to a higher input rate of data packets in the network.

We now analyze the total end-to end packet delay at the transport layer and the time spent at the MAC layer to get further insight. From Fig. 4.8, we can see that in the one channel case, end-to-end total packet delay can become very high as we increase load whereas in the multi-channel chains, total delay is better controlled. We should note, however, that, as shown in Fig. 4.9, queueing delay accounts for most of the total end-to-end delay.

In all cases, before reaching saturation, the average delay is low and the MAC-level delay contributes to most of the delay. As we reach and pass the saturation point, queueing delay becomes the major component of the total delay. In the one channel case, saturation is reached very quickly (*i.e.* as the offered load reaches 300 Kbps).

For the multi-channel setup, above the saturation point, two hops/channel has a lower total delay when RTS/CTS is used. However, we can see in Fig. 4.9(a) that at the MAC level, once the queueing delay is discarded, one hop/channel has a lower delay. One hop/channel thus suffers more from queueing. In the one hop/channel setup, the sudden increase in delay for the case where RTS/CTS is used is due to a sudden queue build-up at interface 13 and 7 (see Fig. 4.1) when reaching saturation. RTS/CTS and the MAC backoff algorithm help the one hop/channel chain to control the arrival rate of packets in the chain. In one hop/channel below an offered load of 800 Kbps, the MAC is able to find a good scheduling whereby there is little queue build-up.

In the case of Basic Access, whereas in two hops/channel total delay increases abruptly

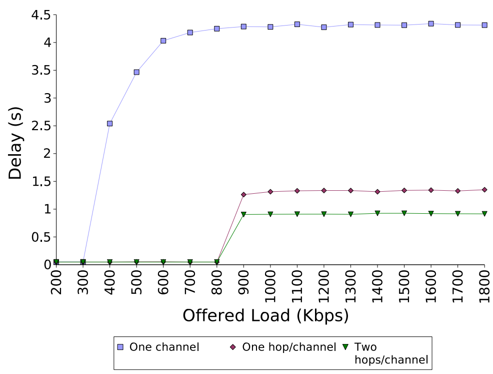

(a) RTS/CTS

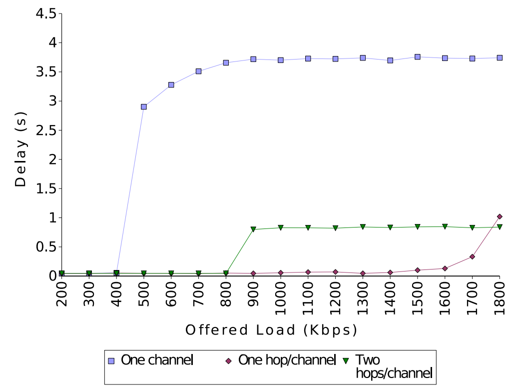

(b) Basic Access

Figure 4.8: End-to-end delay vs. Offered load

as we reach the saturation point, in the one hop/channel the increase is smoother. This is explained by the fact that queue build-up can occur at many more places in one hop/channel and is subject to the scheduling of link transmissions on the same channel. As we reach the saturation point in the one hop/channel configuration, we can see the total delay steeply increases even though the MAC-level delay stays relatively stable. There is a increasingly larger queue build-up at various places up to interface 7 that contribute to the higher packet delay. We see that the one hop/channel is particularly vulnerable to queue build-up at high load.

Fig. 4.10 shows the delivery ratio, which we measure as the percentage of transmitted data packets that are successfully received. As expected the delivery ratio dropped after link-saturation point, as packets are dropped at the source due to a full queue. The multi-channel configurations are able to maintain high delivery ratio for higher loads. With RTS/CTS, two hops/channel offers the best delivery ratio because of the limited loss of packets due to interference between links on the same channel. From Fig. 4.10, we can see that overdriving the network reduces the throughput. Optimal capacity can be

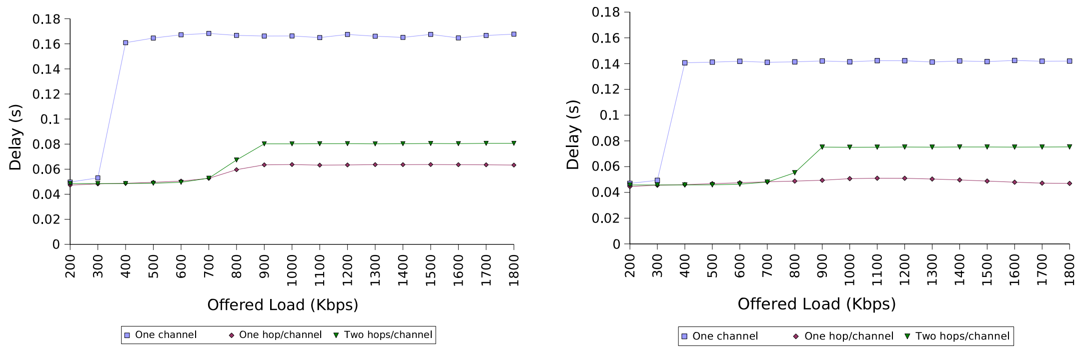

(a) RTS/CTS

(b) Basic Acess

Figure 4.9: End-to-end MAC-layer delay vs. Offered load

achieved by limiting the offered load to a level below the saturation point.

4.2.5 Summary of chain experiments

The chain experiments show that using multi-interface nodes improves both throughput and the end-to-end delay. We also note that the combination of RTS/CTS and the exponential back-off algorithm are problematic in our multi-channel system. This problem is likely to compound in more complex scenarios. We also observe that using more interfaces does not necessarily imply better performance. Thus it is essential to carefully parametrize the MAC protocol to achieve best performance.

4.3 Grid Topology

While the chain topology is simple to analyze, its applicability is very limited. In many scenarios, wireless nodes are likely to be connected to more than two nodes in a mesh-like topology. A mesh network can be used stand-alone to provide wireless access in a particular area, or it may extend a wired network to facilitate last-mile access. In

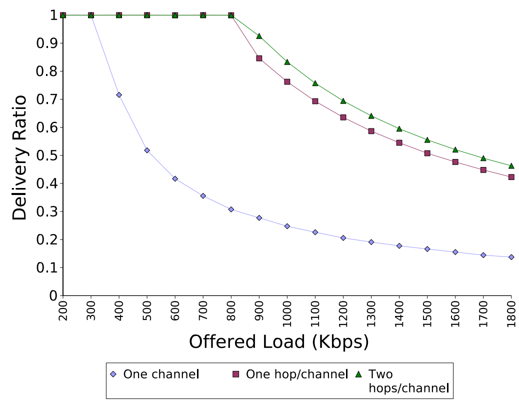

(a) RTS/CTS

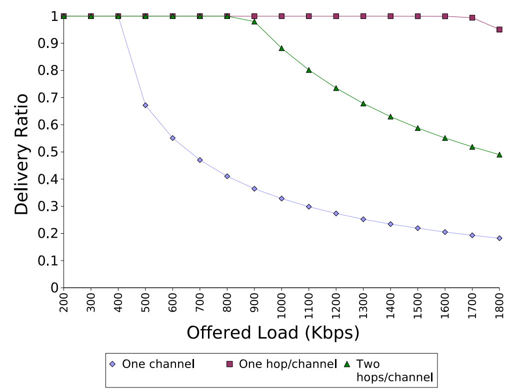

(b) Basic Access

Figure 4.10: Delivery ratio vs. Offered load

this section, we evaluate the performance of our multi-interface setup in a regular grid topology. Compared to the chain topology, the task of analyzing behaviour in a grid topology is harder. Two main factors contributing to this difficulty are multiple sources of interference and unbalanced contention. Nevertheless, we use a range of scenarios to study and explain the behaviour of networks with multi-interface nodes.

Experiments described in this section involve a regular grid-like network. All directly-connected nodes are separated by 200 m. We assume line-of-sight between neighbouring nodes. We compare three grid configurations that are extensions of the chain topology that we elaborated upon in Section 4.2. It should be noted that given the limited number of channels it is not possible to strictly implement the one hop/channel and two/hops per channel configurations both horizontally and vertically. However, we maintain such patterns in the rows of the grid. These three grids are as follows:

- **Grid 1**: This grid uses a single channel and has the same layout as the grids shown in Fig. 4.11.

- **Grid 2** (Fig. 4.11(a)): This grid is based on the one hop/channel pattern.
- **Grid 3** (Fig. 4.11(b)): This grid is based on the two hops/channel pattern.

For simplicity, in Fig. 4.11, we label nodes instead of individual interfaces, and nodes with a cross represent dual-interface nodes, whereas nodes with a circle have one interface.

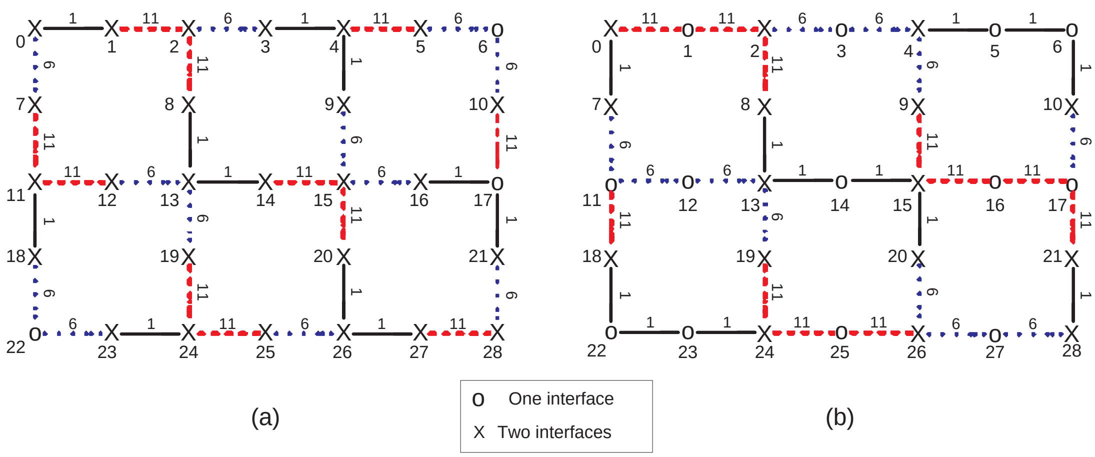

Figure 4.11: Multi-interface grids: (a) Grid 2 (26 two-interface nodes) (b) Grid 3 (16 two-interface nodes)

In our first experiment, we randomly selected 10 sources and 10 destinations on the grid itself. Data is generated at a high rate such that each source always has data to send (2 Mbps in these experiments). The grid-nodes in this case not only generate data, but also relay traffic from other nodes. In Fig. 4.12, we show results for 10 such scenarios whereby each result represents an average of 5 runs.

Our first observation is that in most cases the multi-interface grids yield better throughput than the single-channel grid. The improvement in throughput varies greatly and depends on the selection of the source-destination pairs. In scenarios where we have many

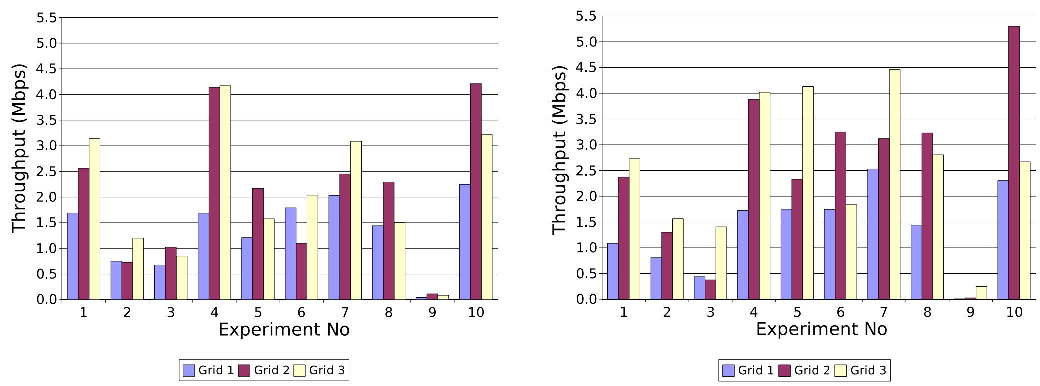

(a) RTS/CTS (b) Basic Access

Figure 4.12: Aggregate throughput of 10 the random flows on the grid

links that are shared, relatively long (more than 3 hops), and highly contended, the multi-channel grids offer a very small improvement or worse throughput when compared to the one-channel grid. This is the case with Experiment 2, 3, and 6 in Fig. 4.12. In scenarios where the flows are spread out on the grid, we obtain higher throughput in the multi-channel grids.

Another observation is that Grid 3 has better aggregate throughput in half of the scenarios when RTS/CTS is used and 7 out of 10 scenarios when only Basic Access is used. This happens despite the fact that in Grid 2, we use more radio interfaces than in Grid 3. Part of the explanation lies in the higher level of interference that exists in Grid 2 when compared to Grid 3, as channel are re-used closer to each other. This causes a higher level of collisions that in turn reduces throughput.

4.3.1 Starvation

To gain more insight, we studied the detailed traces of the simulations. We noted that there are many flows that are not able to enjoy any throughput at all. The results derived are presented in Fig. 4.13. For each scenario in Fig. 4.12, we show the average number of flows that are completely starved.

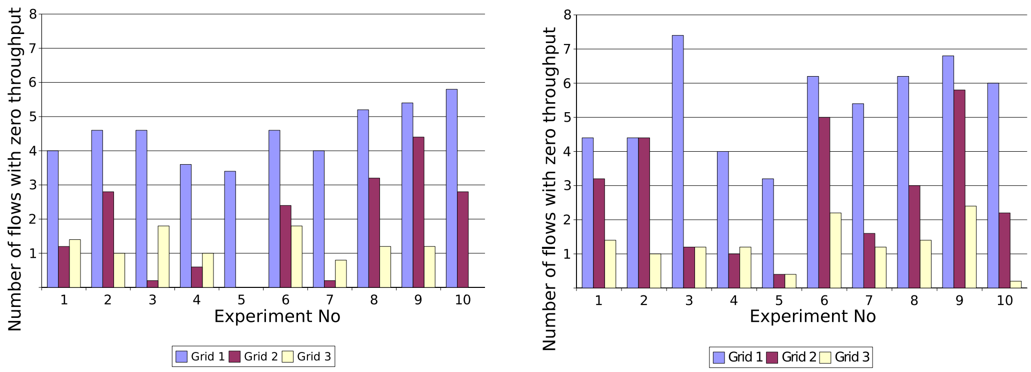

(a) RTS/CTS (b) Basic Access

Figure 4.13: Average no. of flows with zero throughput

We see in Fig. 4.13 that, on average, 4 out of 10 flows in the case of RTS/CTS, and 5 out of 10 flows in the case of Basic Access, completely starve in Grid 1. This indicates a high level of unfairness in this set of experiments. We observed that flows that have shorter routes are those that survive. In our scenario, a short flow is a stream of packets that follows a path of three hops or less. The main problem stems from the fact that flows being relayed by a node are unable to compete with the traffic generated at that node. In a multi-hop network the interface queue of a node either contain packets generated at that node or packets received from other nodes, which need to be forwarded. Any neighbouring node needs to first access the medium for transmitting a packet and then, provided successful reception occurred, this packet may be queued at the receiver. A packet that is generated by a node itself and scheduled for transmission does not undergo such competitive access. It is thus easier for a node to fill up its own queue with its own packets. Therefore at a high load, the probability that a packet belonging to another node, which first needs to be relayed, is queued is significantly smaller than that of a node's own packet. This is the main cause to the high level of unfairness in Grid 1. This problem has been studied in the context of single-channel wireless mesh networks by Jangeun and Sichitiu [21], albeit for simple scenarios. It should be noted that unfairness as discussed here is different from the MAC-layer described in Section 4.2. In this case, unfairness occurs at a higher layer when packets from different nodes compete for a place in the send queue.

Flow starvation also occurs in the multi-interface grids, even though with lower severity. There are many more flows with zero throughput in Grid 2 than Grid 3. In Grid 2, the high throughput that can achieved along a path, can have a similar effect to the one where a node generates and relays data at the same time. We observe in Fig. 4.13 that Basic Access significantly increases starvation the case of Grid 2. As we saw in the chain topology, if channels alternate after each hop along a route, we achieve high throughput especially when Basic Access is used. This implies that unfairness may not only occur at a node generating and relaying traffic, but also happen one or more links ahead from that node. If, at a particular node, the transmission opportunity is being contended by multiple flows, a flow with higher packet arrival rate will obtain an equivalent share of relaying opportunities.

4.3.2 Unfairness

Fig. 4.13 only shows the number of flows that starve. However, there are many flows that have very low throughput. Below, we investigate flow throughput in one of the scenarios in Grid 3 when using RTS/CTS (Experiment No.1 in Fig. 4.12). Fig. 4.14 depicts the 10 flows with the corresponding routes on the grid in one instance of the simulations. Table 4.2 shows the throughput achieved for each flow in this particular simulation.

We see that flows 8-20, 12-5, 7-9, 3-23, and 2-6 have zero or very little throughput.

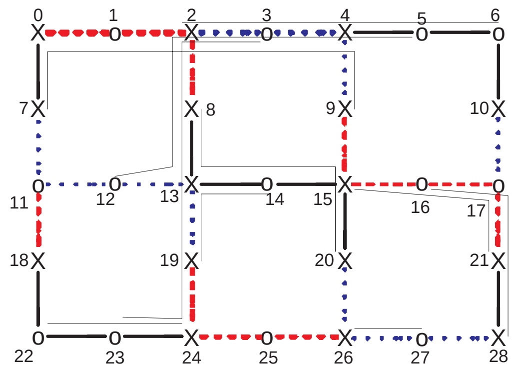

Figure 4.14: Grid 3 with 10 random flows

Flows 8-20 and 12-5 share the highly-loaded node 13. Flow 8-20 starves because node 14 overloads the link to 13. Flow 12-5 not only has to pass through node 13, thus competing with other flows for relaying, but also needs to go through the link 2-3, which is shared by 3 other flows. Similarly, flow 7-9 also competes for the medium at node 2, 3, and 4. Similarly, flow 2-6 competes for relaying opportunity at node 2, 3, 4 and 5. Given that node 2 is generating data, it would normally overload link 2-3 with its own traffic. However, in our scenario, node 2 is not able to do so, as node 3 also generates data destined to node 23. Both 2 and 3 are on the same channel and thus share access to the medium. Flow 3-23, however, suffers from a high level of contention throughout its path. As mentioned earlier, when multiple flows compete at a node, the throughput is a factor of the arrival rate at that node. The higher the arrival rate of packets from a particular flow at the node, the higher the likelihood of a packet from the flow being forwarded. Flow 27-26 achieves the highest throughput as it suffers from very low contention (from node 19). Throughput of flow 22-24 is lower than half of the maximum throughput, as the flow spans over two links on the same channel. In addition, it shares link 23-24 with the flow 3-23. Flow 15-21 uses three consecutive links on the same channel shares the 16-17 and 17-21 links with the 28-16 flow. Flow 28-16 attains higher throughput than

Flow	Throughput (Mbps)
8-20	0.0
15-21	0.24496
27-26	1.66768
22-24	0.66664
2-6	0.0222
12-5	0.00128
28-16	0.47316
14-19	0.02432
3-23	0.0114
7-9	0.00208

Table 4.2: Average per-flow throughput

flow 15-21 because one of the links it uses is on a different channel, whereas the two others use the same channel.

This brief description hints at the complexity involved in analyzing throughput in a grid. Various factors are involved in estimating per-flow throughput. The main factors are the individual link bandwidth, which depends on the level of interference around both ends of the link, the packet arrival rate (external or internal) at each node, and the route of each flow.

We observe a high level of unfairness in the grid as exemplified by the above scenario. To further investigate this fact, we selected one set of 10 randoms flows (same flows as shown in Fig. 4.14) and study the resulting performance with increasing offered load. We repeat each experiment 10 times and present the average results.

Fig. 4.15 shows the aggregate throughput in each experiment as we increase the load at the source. We can see that the aggregate throughput increases as we increase the offered load. With RTS/CTS, aggregate throughput achieved in Grid 1 stabilizes around 1.6 Mbps. With Basic Access, throughput in Grid 1 drops when offered load is increased after a certain point. A close examination of the simulation traces indicates that this

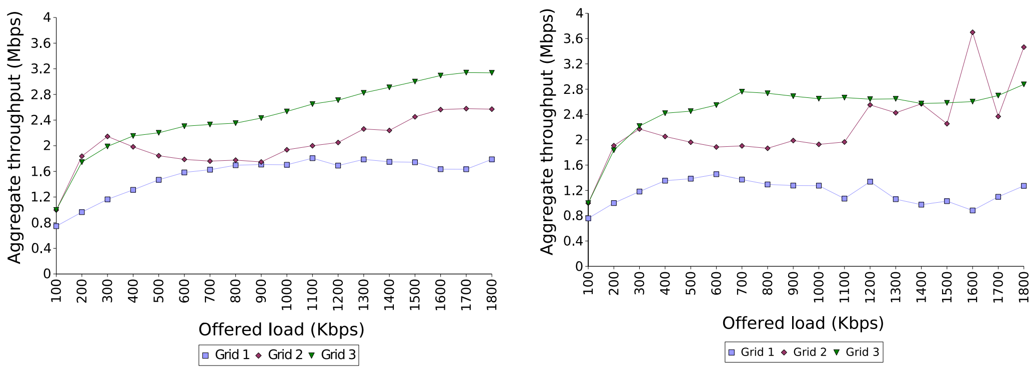

(a) RTS/CTS (b) Basic Access

Figure 4.15: Aggregate throughput of 10 random flows with increasing offered load

drop is due to a bias in the route selection. The drop is due to the selection of different routes compared to those chosen at lower load in Fig. 4.15. We note that there exist two different alternative paths for flows 12-5 and 7-9. If the same routes are chosen for each run, we do not expect such a drop. At higher loads in Grid 1, a few flows account for the total throughput due to starvation of most other flows.

In Grid 2, the throughput first peaks at around 300 Kbps. This increase is achieved by a gain in throughput by most flows. As such, at this level of offered load, there is no flow with lower than 100 Kbps throughput. Beyond this level, the level of interference between nodes and competition among flows at various points on the grid increase and throughput starts to drop for most flows. However, some short flows are still able to deliver high throughput. As the offered load increases, these flows are able to obtain extra throughput, offsetting the loss of throughput of other flows that experience high contention. At higher offered loads, more throughput is achieved because a few flows (4 flows in Grid 2) monopolize the medium and thus contribute to the increase in the aggregate throughput. The routes taken by these flows are shorter than for the flows having low or no throughput. The behaviour is similar both with RTS/CTS and with Basic Access. However, throughput achieved is higher with the Basic Access.

In Grid 3, the overall throughput increases as a function of offered load. We observe the same unfairness phenomenon, whereby at higher offered load, a few flows contribute to most of the overall throughput. With RTS/CTS, Grid 3 provides highest throughput when compared to both Grid 1 and Grid 2. When using the Basic Access scheme, Grid 3 is able to reach higher throughput level at higher levels of load. Grid 3 does not have a drop in aggregate throughput like Grid 2 because, compared to Grid 3, more flows suffer from less interference. More flows are able to increase their throughput, thus offsetting the loss seen by others either due to unfairness and collisions. In Grid 3, Basic Access offers better performance below an offered load of 1200 Kbps, but above this point the the system is unable to find a better scheduling than the one provided by RTS/CTS. This is the result of more collisions without RTS/CTS. An example of such a case in our scenario is collisions at node 16 caused by flows 15-21 and 28-16, which share two links in opposite directions.

For the same set of experiments, we measured the level of starvation. Fig. 4.16 shows the average number of flows that starve at each level of offered load. We can see that as the load increases at the source, we have an increasing number of flows that starve. Starvation is most severe in Grid 1. As explained earlier, when one node on the grid generates data, it reduces the probability of packets from another node being relayed. If the packet arrival rate from another node is high, some packets may be relayed. For flows with long paths, even after one hop, the throughput is likely to significantly decrease to a level that is insufficient to compete with the data generated at an intermediate node. There are far fewer flows with no throughput at all for the multi-interface grids.

With RTS/CTS, Grid 3 experiences comparatively early starvation because we reach saturation level for many links at a lower offered load. Above the saturation point at

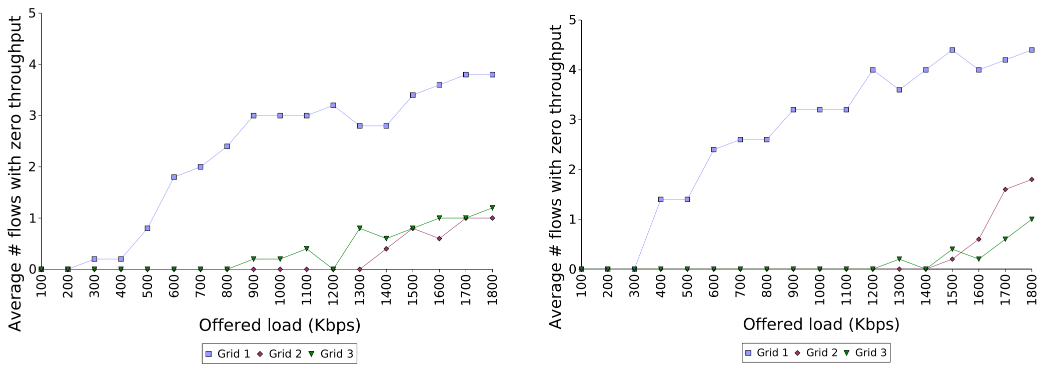

(a) RTS/CTS (b) Basic Access

Figure 4.16: Average no. of zero-throughput flows out of 10 random flows with increasing offered load

any node, starvation depends on the packet arrival rate of the flows sharing the link from that node. At high loads with RTS/CTS, starvation increases more steeply in Grid 2. As described earlier, Grid 2 is able to provide higher throughput paths that can cause unfairness well after the point at which data is generated. With Basic Access, starvation appears earlier and is more severe in Grid 2. This is again explained by the fact that some flows can have a high throughput along their paths. Such high-throughput path are more common in Grid 2.

The number of zero-throughput flows does not provide a good indication of the fairness in the distribution of the network capacity among the competing flows. We now compute the per-flow throughput variance for the same set of experiments. These results are presented in Fig. 4.17. The higher the variance, the larger the variation in throughput of the different flows on the grid. We consider the system as totally fair when the flow variance is zero. We observe that the unfairness problem is far larger in the multi-interface grids than in the single-channel grid. As we increase the source load, unfairness

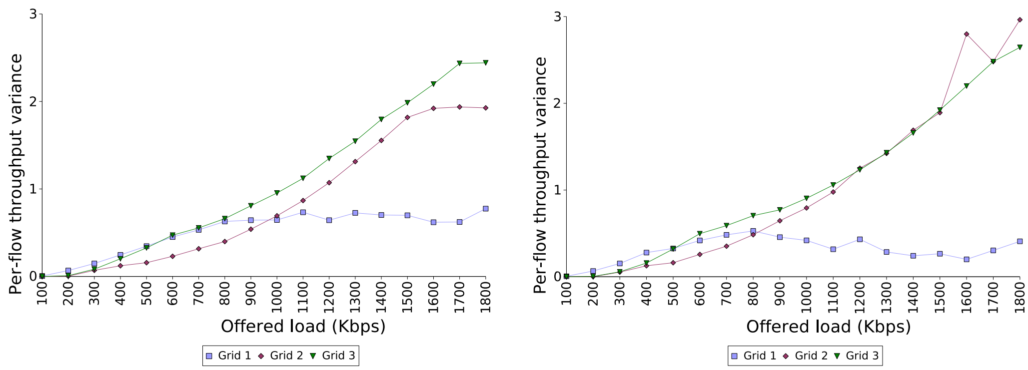

(a) RTS/CTS (b) Basic Access

Figure 4.17: Per-flow throughput variance with increasing offered load

constantly increases. Unfairness in Grid 1 with RTS/CTS steeply increases, but then stabilizes because links cannot carry more data to change the throughput distribution among the flows. In Grid 1 with Basic Access, increasing load unfairness seems to decrease after reaching a first peak at 800 Kbps. This is again due to a bias in the simulation that selects different routes. If we look at Fig. 4.15(b), we can see that this corresponds to a decrease in the overall throughput. Unfairness is relatively higher in Grid 1 with RTS/CTS. This in turn is due to the RTS/CTS problem described in Section 4.2. These problems are more severe in a grid topology, as contention for the medium at a particular location can be imbalanced, giving unfair share to nodes that are less affected by contention.

With RTS/CTS, unfairness is higher in Grid 3. First, we note that because in the rows of Grid 3, we have two consecutive links using the same channel. This implies that these links will saturate earlier. Second, RTS/CTS is less effective on longer single channel paths. In Grid 3, we see many instances of three or more consecutive links using the same channel. This causes more collisions and increased back-off, hence reduced throughput. As we reach saturation, flows with a higher packet arrival rate at the congested link will get an unfair share of the medium. In Grid 2, however, single-channel segments do not exceed two hops. As such RTS/CTS is very effective. The problems of back-off and especially collisions exists in Grid 2 also, but the overall negative effect is less severe than in Grid 3.

When using the Basic Access scheme, unfairness is higher in Grid 3 below an offered load of 1200 Kbps. This is explained by the same reasons as mentioned above. At higher load, the per-flow throughput variance is higher in Grid 2 whereby a few flows are able to take advantage of the availability of high capacity paths.

This first series of experiments with the Grid topology indicate that nodes in a network should not be allowed to generate and relay data at the same time if no traffic discipline is implemented. The problem stems from the fact that a user is easily able to overload a link because its own traffic does not suffer from contention before being queued for transmission. This in turn not only wastes bandwidth, as the node's link-layer share of the medium remains the same, but also reduces the probability of other users' traffic being queued at that node. Our experiments also show that using more interfaces with different channels does not automatically guarantee higher capacity. We have seen that Grid 3 proves to be a better configuration even though it uses fewer radio interfaces than Grid 2. Because, the number of orthogonal channels is very small for 802.11 DSSS and that interference range is larger than the transmission range, channel allocation plays a crucial role in providing adequate gains in performance.

We do not evaluate techniques to alleviate the problems discussed above. Instead, we propose a wireless infrastructure whereby such problems are considerably mitigated. We note, however, that many community area networks are being built such that all nodes can generate data and have to relay other users' traffic. A user that overloads the medium, is likely to severely hinder the network performance and possibly prevent other users from enjoying any usable throughput. Many studies have been carried out in the context of fairness in wireless networks that may help address unfairness problems discussed above.

4.3.3 Purely Wireless Infrastructure

Our main proposal is to provide a wireless backbone, which users can use to connect to each other and to access external networks as well. This approach represents a natural means for traffic control, as each user needs to contend for the medium before enqueueing packets for transmission. user nodes, in this approach, do not relay other users' data. They will instead compete for relaying opportunities at nodes dedicated for routing. While overhead may be higher with this approach, as every communication has to go through an access point, we believe that its benefits outweigh such inefficiency. Based on this idea, we now investigate the performance of our multi-interface system with the grid topology whereby user nodes are randomly spread out in the area covered by the grids. Nodes in the grid act as access points. Each user node associates with its closest access point, which relays all user node's traffic. In such scenarios, it is harder for any one user node to monopolize the medium. We are thus less likely to see a queue at an access point filled up with packets from a single user node.

For these experiments, we randomly place a varying number of user nodes under the area covered by the grids described earlier (See Fig. 4.11. This represents a coverage of an area of 1000 m x 1400 m. Figures 4.12(a) and 4.12(b) show the aggregate throughput of an increasing number of user nodes. Half of these user nodes act as source and the other half as destination. Each source generates traffic at the rate of 100 Kbps. We repeat each experiment five times and present the average results.

We first note that when the number of nodes is low (*e.g.* 10, 20 or 30 based on Fig. 4.18), there is excess capacity in the multi-interface because the offered load is low. As such, a direct comparison with Grid 1 is not informative. In most other ex-

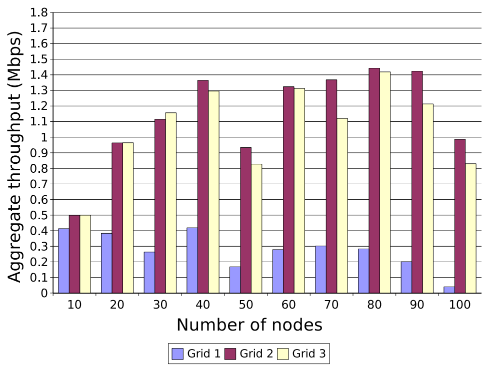

(a) RTS/CTS

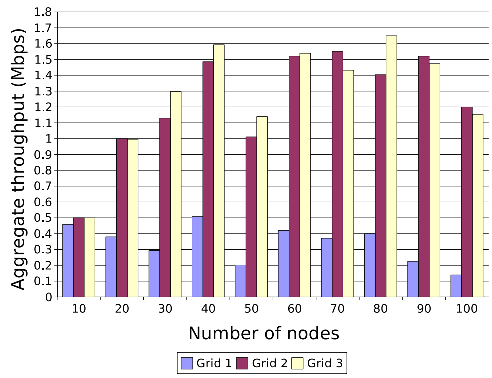

(b) Basic Access

Figure 4.18: Throughput vs. No. of nodes ($nodes/2$ random flows between user nodes)

periments shown in Fig. 4.18, we can see that the multi-interface grids outperform the single-channel grid. In most cases, there is more than a three-fold gain when compared to the single- channel case. Grid 2, which is based on the one hop/channel configuration, offers the highest aggregate throughput when RTS/CTS is employed. This contrasts with the results we obtained earlier where nodes in the grid generated data. In those experiments, Grid 3 is a better setup. When we disable RTS/CTS, performance in Grid 3 is comparable to Grid 2, even better in many cases. With RTS/CTS, the lower performance of Grid 3 is mainly due to the interference caused by user nodes, which adversely affects the already reduced capacity of the backbone. When interference from user nodes, adds to this self-induced backbone interference, the capacity of the network is significantly reduced. Also, because backbone links are shared over several hops, each access point waits before winning the medium. When an access point is waiting, the user nodes it services are unable to communicate and back-off with larger contention windows, if they could not hear the communication taking place. This problem is less severe in Grid 2 because we do not have more than 2 consecutive links that use the same channel. We also observe that Grid 3 gains more by disabling RTS/CTS than Grid 3 because of longer consecutive single-channel links. RTS/CTS is more effective for shorter links, thus the comparatively small improvement for Grid 2.

4.3.3.1 GATEWAYS

When deploying wireless mesh networks, it is most probable that certain points in the network will experience higher load. These locations will correspond to popular resources such as Internet gateways, file servers, *etc.*. To evaluate performance in such situations, we carry out a set of experiments with four hot spots placed at various places in the grids (depicted in Fig. 4.19). Hereafter, we abstract these popular resources as gateways. As with normal user nodes, the gateways associate to the closest access points.

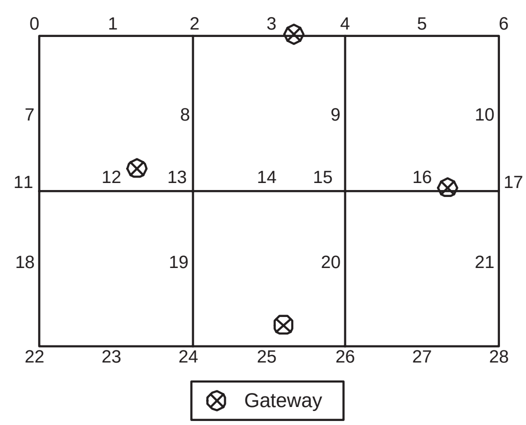

Figure 4.19: Grid with 4 gateways

In this set of experiments, we evaluate the effect of a significant amount of traffic destined to, and originating from, the gateways on the different grid configurations. We use the same topology as the ones used for experiments in Fig. 4.18, but with the addition of four gateways. Half of the user nodes that are randomly placed under the coverage of the grid send data to the gateways. Each of these user nodes sends data to its closest gateway. The gateways send data to the other half of the user nodes. Again, the closest gateway sends data to each of these user node. All nodes generate traffic at the rate of 100 Kbps. Results from this set of experiments are shown in Fig. 4.20.

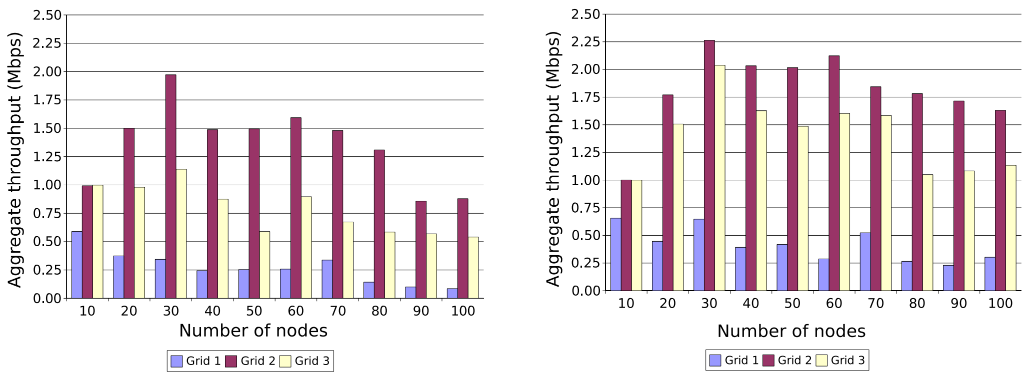

(a) RTS/CTS (b) Basic Access

Figure 4.20: Throughput vs. No. of nodes (random flows between user nodes and 4 gateways)

Fig. 4.20 illustrates again the higher capacity in Grid 2 when we deploy the wireless infrastructure. There is greater than a three-fold improvement in most cases. In these experiments, small clouds of high contention form around the gateways. With both access schemes, we achieve higher aggregate throughput in Grid 2. First, there is some physical separation between the areas with high-contention. Second, we have shorter single-channel links in the backbone. Both these factors lead to the better throughput in Grid 2. We observe, however, that there is a significant difference in the aggregate throughput among the two access schemes. The Basic Access mechanism brings about important gains in throughput. The lower throughput in the case of RTS/CTS is mainly due to access points suffering from the high contention around the gateways. This causes the access points to suffer from large contention windows as a result of the exponential back-off algorithm. In turn, the victim access points cause user nodes attached to them to back-off. The second cause for the poor performance with RTS/CTS lies in the existence of overlapping cells. As we can see in the grids (Fig. 4.11), in most places there are more than one consecutive links on the same channel. Each end-point of a link is an access point to which user nodes can attach. access points that are neighbours are likely to cover each other's users. These users attach to the access point that is closest to them (assuming we do not have any external interference). An RTS handshake initiated between a user node and its access point prevents user nodes and other access points in the neighbourhood of both nodes from transmitting during the transmission of a data packet. However, in general in our simulations, the average distance between a user node and its access point is small. As such, in many cases two users-nodes could simultaneously communicate with their respective access points without interfering with each other, provided they are sufficiently separated.

4.3.3.2 Prioritizing backbone traffic

There exist various ways in which the overall performance in the grid topology can be improved. For improving throughput in particular, it is possible to prioritize the backbone traffic (*i.e.* traffic between access points). This idea is based on the observation that once a packet is accepted by the network, it should be given precedence over packets that have not yet entered the network. Also, the delivery of a packet from the backbone to a user node is given priority over a packet from the user node to the network. This would reduce the wastage of bandwidth due to packet-drops after a certain amount of bandwidth has been consumed. In this section, we evaluate gain in performance using this simple idea.

There is exist various ways to prioritize traffic. For our purposes, we select a subset of ideas proposed for Extended DCF (EDCF) in the 802.11e draft, as described by Mangold *et al.* [27]. With EDCF, higher priority is achieved by a combination of different inter-frame space periods, minimum and maximum contention window, and variations of the back-off algorithm by means of a multiplier variable. In our simulations, for the traffic between access points and from the access points to user nodes, we opted to use PIFS as the inter-frame space period, a minimum contention window of 7 (31 for user nodes), and a maximum contention window of 127 (1023 for user nodes). The use of PIFS amounts to a change to the MAC, but we should point out that prioritizing is mostly the result of using a smaller minimum contention window.

We repeat the same set of experiments as carried out in the previous section. In the first set, we simulate an increasing number of user nodes serviced by the grid. Half of the user nodes are the source and the other half are the destination. The offered load at the source-nodes is 100 Kbps.

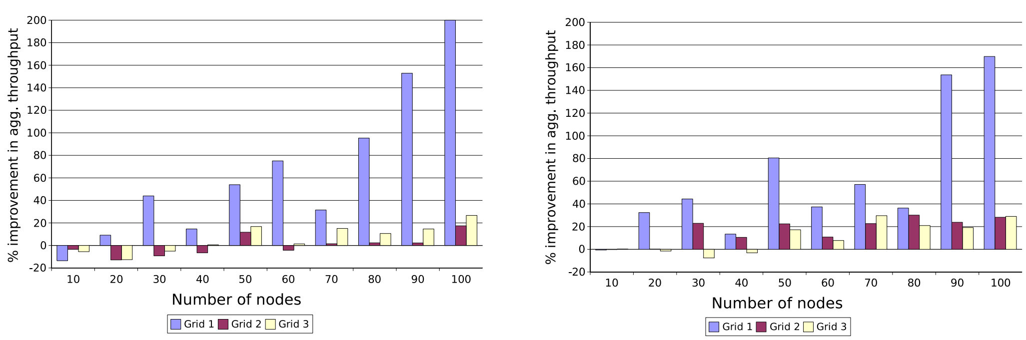

(a) RTS/CTS

(b) Basic Access

Figure 4.21: Throughput vs. No. of nodes with prioritized backbone ($nodes/2$ random flows between user nodes)

The results obtained from these experiments are shown in 4.21 and represent the percentage improvement in the throughput. Any negative value amounts to a decrease in throughput when compared to the base-case scenarios where all nodes use a similarly tuned MAC. Improvement for 100 user nodes when using RTS/CTS in Grid 1 is about six-fold, but this fact is not shown in order to improve readability of the chart.

This set of experiments allow us to roughly estimate the effect of user traffic on the backbone (*i.e.* the access points), which reduces the aggregate throughput. By prioritizing the backbone, we reduce this effect, and the level of gain reflects the magnitude of the problem.

We first observe in the RTS/CTS case that there is a significant gain for the one-channel grid. When a user node uses RTS/CTS to reserve the medium, it will in most cases cause a neighbouring access point to freeze. While in some cases this may be necessary, in others the reservation will be too conservative. As we described in the background, the interference range depends on the distance between the sender and its destination. When a user node is close to its access point, blocking a neighbouring access point would be unnecessary. In addition, binary exponential back-off algorithm also has a negative effect on the backbone. A user node sending a long data packet to its access point will cause other access points, which have the user node within their carrier-sense range, to repeatedly back-off. Prioritizing somewhat alleviates these problems by protecting the access points.

In half of the experiments with RTS/CTS, prioritizing reduces aggregate throughput, especially for the multi-interface configurations. We should note that problems due to RTS/CTS and the back-off algorithm do not disappear with prioritization. Traffic in the backbone has the same priority, and thus self-induced interference still exists. Furthermore, with the backbone traffic having higher priority, user nodes spend more time in back-off and thus are unable to take advantage of transmission opportunities when the medium becomes idle. In many cases, however, the throughput loss due to these problems is offset by the gain achieved from prioritizing. This contrasts with the experiments that do not use RTS/CTS, in which prioritizing the backbone traffic brings about consistent improvement in throughput. With Basic Access, there is increased gain in throughput as the user node density increases. As contention increases around an access point, protection of packets in the backbone becomes more important.

We now repeat our earlier experiments with four gateways, but this time, prioritizing the backbone traffic. In these experiments, half of the user nodes send to the their closest gateway. The gateways send to the other half of the user nodes, each one selecting the ones that are nearest. The offered load at each source is 100 Kbps. The results of these experiments are presented in Fig 4.22.

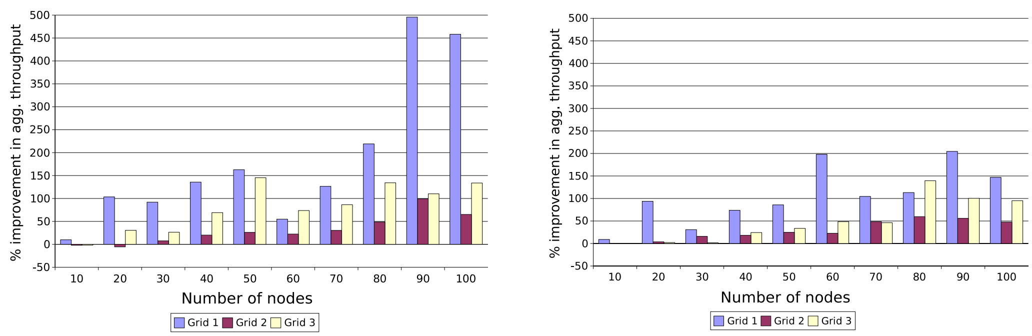

(a) RTS/CTS (b) Basic Access

Figure 4.22: Throughput vs. No. of nodes with prioritized backbone (random flows between user nodes and 4 gateways)

Interestingly, in Fig. 4.22, we see that prioritizing backbone traffic considerably improves throughput when RTS/CTS is used. In Fig. 4.20, we saw that RTS/CTS provides noticeably lower throughput than the Basic Access scheme. By prioritizing backbone traffic, we reduce the likelihood of access points being unfairly treated in the presence of high contention. This explains the large improvement in Fig. 4.22(a). We should also note that even the scenario with Basic Access benefits from the prioritizing of the backbone, as less bandwidth is wasted on packets that are dropped along the way to their destination.

The previous experiments have shown that the idea of a wireless infrastructure not only improves fairness, but also has the potential to provide increased capacity when multiple channels are utilized. We have also shown that RTS/CTS problems have a significant impact on throughput in a wireless infrastructure setup. It is thus necessary to tackle such problems if we are to provide a purely wireless infrastructure using the limited number of channels as available for 802.11 DSSS PHY. We also demonstrated that such problems can be contained by employing traffic differentiation.

4.4 Random Topology

After studying the chain and the grid topology, we now look at the performance of a multi-interface system in a random topology. In contrast to the regular setup studied earlier, here we randomly place nodes that form the backbone. The layout we tested is constrained by the minimum distance between two nodes in the backbone, which cannot be smaller than a threshold. In a real-life deployment, the objective of system engineers is to maximize quality coverage with as few access points as possible. We assume that placing access points closer than a certain threshold does not yield any advantage in terms of coverage. On the contrary, this reduces spatial re-use of the available frequencies.

We further assume that once the nodes are turned on, they operate on a common channel for a short period of time. During this lapse of time, the nodes execute a protocol that allows them to select the channels to use. We devised a simple channel-assignment protocol, which does not attempt to optimize channel allocation. The protocol does not cater for the incremental addition of nodes to the network, even though it can easily be extended to this effect. Our protocol is presented here as a proof of concept for the possibility of using more fine-tuned channel selection algorithms.

Examples of such random topologies are depicted in Fig. 4.23.

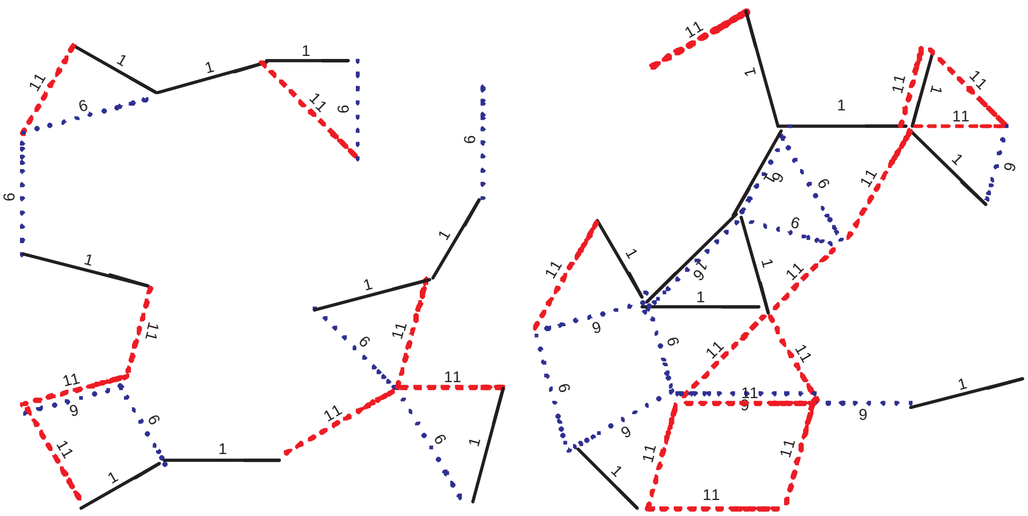

(a) Example 1 (b) Example 2

Figure 4.23: Sample random networks with 20 access points

4.4.1 Channel Selection Algorithm

We hereby present a simple neighbourhood-based distributed channel-selection algorithm. The principle is to use the channel that is least used in the 2-hop neighbourhood of nodes at both ends of a link. The protocol operates as follows: a token is broadcast by one node and passed over by its neighbours to their neighbours, until all nodes have processed the token. A node only acts upon the token once in order to select the channels to use. When a node first receives the packet containing a token, it contacts each of its neighbours in turn to choose the channels. It exchanges channel-usage information with its neighbour and, as a result, picks the one that is least used by both itself and its neighbour. Each node is allowed to choose a maximum of two channels. If both the node and its neighbour have a free radio-interface, both are assigned the selected channel. If one of the them has both its interfaces already assigned channels, the one with a free interface chooses the channel of the other that is least used by both nodes' neighbours. In case two neighbours are connected via two non-overlapping channels, the channels chosen for communication will depend on the interface that is first advertised. For DSDV, the channel selected may change as a result of periodic or triggered updates. In our simulations, the channel used remains the same because we disable periodic updates in order to eliminate the routing overhead. In actual deployments, however, the channel selected used may change over time. Fig. 4.23 shows the results of running the above algorithm on a random graph of 20 nodes. Each pattern in the figure represents a different channel.

In the simulations described below, we used an area of 1100 x 1100 m where the infrastructure network consists of 20 access points. We ensure a minimum distance of 150 m between access points. We vary the number of user nodes from 10 to 160, randomly placed within the area covered by the access points. Half of the user nodes send data to the other half at a rate of 100Kbps.

Results of these simulations are presented in Fig. 4.24. We can see that with both access schemes, we have significant throughput gains using multiple interfaces in comparison to the single-channel network. With the exceptions of a few experiments that have a small number of users, the multi-channel configuration offers more than a three-fold improvement in performance. These experiments are based on a single topology for the core network, but we observed the same benefits in other random topologies.

Despite the important benefits of using multi-interface nodes, we believe that throughput in random topologies can be improved in various ways. A combination of better channel selection algorithms and power-aware transmission or routing protocols ([8, 22, 31, 44]), for example, can further improve advantages put forward by our experiments.

This last set of experiments show that a our wireless infrastructure based on a backbone using multi-interface nodes can be a valuable solution to creating high-bandwidth wireless mesh networks. In turn, such a wireless mesh has many applications varying from commercial movable networks to community networks for sharing costly Internet

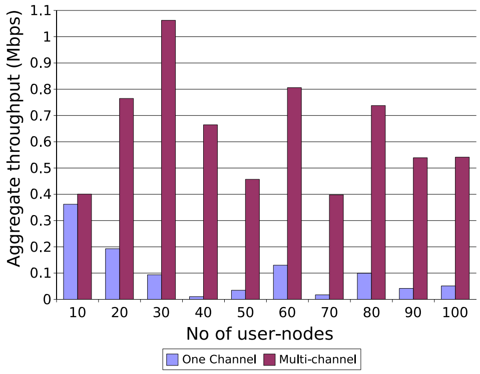

(a) RTS/CTS

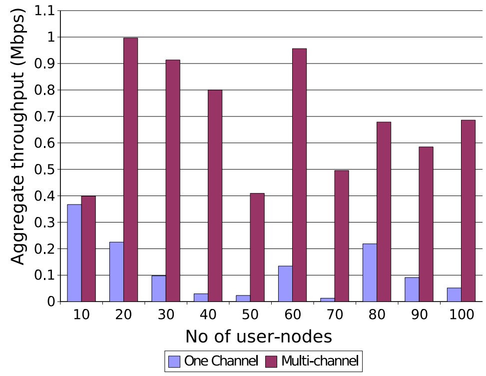

(b) Basic Access

Figure 4.24: Single-channel vs. Multi-channel random mesh

access. However, our study demonstrates that it is essential to address issues related to the RTS/CTS access scheme in order to obtain good performance. In addition, we show that service differentiation can yield valuable capacity improvement.

5 Conclusions and Future Work

In this thesis we evaluate the performance of 802.11-based wireless networks with multi-interface multi-channel nodes. In order to deal with fundamental limitations of single-frequency wireless networks, we propose a multi-channel solution that leverages more than one radio-interface on each node. With this approach, one node can potentially simultaneously communicate with two neighbours. An important requirement of our work is to use the low-cost commodity 802.11-compliant wireless cards, avoiding any hardware changes. We made simple modifications to routing modules in order to co-ordinate multiple interfaces on the same node. We not only study the behaviour of the proposed system when nodes produce and relay packets as in an ad-hoc network, but we also study the use of a wireless infrastructure as an alternative in scenarios where high mobility support is not needed. In such an infrastructure, user nodes do not route packets from each other, but connect to a dedicated multi-channel wireless backbone. We show, through experiments, that using more than one interface tuned to orthogonal channels can generate a valuable increase in capacity of wireless networks. Nevertheless, we single out scenarios where using more multi-interface nodes does not necessarily improve the system. We report serious unfairness issues when a node can schedule its own data and relay others' data without any traffic control. We have shown how the small number of non-overlapping channels available for 802.11 DSSS Specification requires careful channel allocation. We show the anomalous behaviour of the RTS/CTS scheme, due to binary exponential back-off algorithm and the exposed terminal problem, in the multi-channel stationary *ad-hoc* type of network as well as in wireless infrastructure. Finally, we demonstrate that important improvement in capacity can also be achieved in random topologies with simple neighbourhood-based channel allocation algorithms.

As a result of our study, we believe that it is possible to implement very cost-effective multi-channel wireless networks with several-fold more capacity than what is available in single-channel *ad-hoc* networks today. No changes to either the available 802.11 DSSS hardware or the 802.11 MAC are needed to create our proposed system. With the experience accumulated throughout this work, we realize that there exist many avenues for enhancing the proposed system.

There are several aspects of this work that can be extended either to support the main system or to improve its performance. First, power-aware transmission protocols can significantly increase spatial re-use and alleviate the limitation in the number of non-interfering channels. This in turn should yield an increase in the capacity of the system. In our system, channel allocation is only carried out initially. However, dynamic configuration that reflects the load on the network may bring about more enhancement in throughput. During our simulations we also observed that it is possible to improve throughput by controlling traffic scheduled by a node. Schemes developed in the context of fairness and QoS in wireless networks can not only improve the distribution of throughput, but also potentially increase the overall capacity. In this work, we limited ourselves to the DSDV routing protocol. Alternative protocols such as OLSR, based on network clustering, combined with a better routing metric can potentially present far better efficiency in terms of broadcast and routing overhead. Finally, channel allocation can be optimized in various ways. Better algorithms that take location, power, and/or neighbourhood information into account can be used to optimize channel assignment. As an immediate followup it would be very valuable to implement this system in order to ascertain the benefits put forward by our study.

A List of Abbreviations

AODV	-	Ad-hoc On-Demand Distance Vector
ARP	-	Address Resolution Protocol
BS	-	Base Station
BSS	-	Basic Service Set
CDMA	-	Code Division Multiple Access
CTS	-	Clear to Send
DCF	-	Distributed Coordination Function
DIFS	-	Distributed Inter-frame Space
DSDV	-	Destination Sequenced Distance Vector
DSR	-	Dynamic Source Routing
DSSS	-	Direct Sequence Spread Spectrum
EDCF	-	Extended Distributed Coordination Function
ESS	-	Extended Service Set
GSM	-	Global System for Mobile Communications
IBSS	-	Independent Basic Service Set
IETF	-	Internet Engineering Task Force
IP	-	Internet Protocol
MAC	-	Medium Access Control
NAV	-	Network Allocation Vector
OLSR	-	Optimized Link State Routing

PCF - Point Coordination Function

PIFS - Priority Inter-frame Space

RTS - Request to Send

SIFS - Short Inter-frame Space

TCP - Transmission Control Protocol

UDP - User Datagram Protocol

WLAN - Wireless Local Area Network

REFERENCES

[1] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for IEEE 802.11 wireless networks. Technical Report TR-2003-41, Microsoft Research, June 2003.

[2] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LANs. In *Proceedings of ACM SIGCOMM*, pages 212–25, Aug/Sept 1994.

[3] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad-hoc network routing protocols. In *Mobile Computing and Networking*, pages 85–97, 1998.

[4] Yeh Chi-Hsiang, Zhou Helen, Ho Pin-Han, and Mouftah Hussein T. A variable-radius multi-channel MAC protocol for high-throughput low-power heterogeneous ad-hoc networking. 2003.

[5] Wireless Communities. Available at http://www.personaltelco.net/index.cgi/-WirelessCommunities.

[6] Douglas S. J. De Couto, Daniel Aguayo, John Bicket, and Robert Morris. A high-throughput path metric for multi-hop wireless routing. In *Proceedings of the 9th ACM International Conference on Mobile Computing and Networking (MobiCom ’03)*, September 2003.

[7] Douglas S. J. De Couto, Benjamin A. Chambers Daniel Aguayo, and Robert Morris. Performance of multihop wireless networks: Shortest path is not enough. In

Proceedings of the First Workshop on Hot Topics in Networking (HotNets-I), October 2002.

[8] S. Narayanaswamy et al. Power control in ad-hoc networks: Theory, architecture, algorithm and implementation of the COMPOW protocol. In *European Wireless Conference*, 2002.

[9] The HiperLAN2 Global Forum. Available at http://www.hiperlan2.com/.

[10] J. A. Garcia-Macias, Franck Rousseau, Gilles Berger-Sabbatel, Leyla Toumi, and Andrzej Duda. Quality of service and mobility for the wireless Internet. In *ACM Wireless Networks*, volume 9, pages 341–352, July 2003.

[11] MANET Working Group. Mobile Ad-hoc Networking. Available at http://www.-ietf.org/html.charters/manet-charter.html.

[12] The DHC Working Group. Dynamic Host Configuration (DHC), Draft Standard RFC2131 and RFC2132 (for IPv4). Availble at http://www.ietf.org/html.charters/-dhc-charter.html.

[13] The Mobile IP Working Group. IP Routing for Wireless/Mobile Hosts (MobileIP). Available at http://www.ietf.org/html.charters/mobileip-charter.html.

[14] VINT Group. UCB/LBNL/VINT Network simulator (ns) - version 2. Available at http://www.isi.edu/nsnam/ns.

[15] VINT Group. The NS Manual, Available at http://www.isi.edu/nsnam/ns/doc/-index.html.

[16] P.R. Gupta, P.and Kumar. The capacity of wireless networks. In *IEEE Transactions on Information Theory*, volume 46-2, pages 388–404, March 2000.

[17] H.Y. Hsieh and R. Sivakumar. On using peer-to-peer communication in cellular wireless data networks. In *IEEE Transactions on Mobile Computing*, volume 3-1, pages 57–72, 2004.

[18] Wing-Chung Hung, K.L. Eddie Law, and A. Leon-Garcia. A dynamic multi-channel MAC for ad-hoc LAN. In *Proceedings of the 21st Biennial Symposium on Communications*, pages 31–35, June 2002.

[19] IEEE 802.11 WG. ISO/IEC 8802-11:1999(E). International Standard [for] Information Technology—Telecommunications and Information Exchange between Systems—Local and Metropolitan Area Networks—Specific Requirements —Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, 1999.

[20] Jangeun Jun and Mihail L. Sichitiu. Fairness and QoS in multihop wireless networks. In *Proc. of the IEEE Vehicular Technology Conference (VTC 2003)*, pages 6–9, 2003.

[21] Jangeun Jun and Mihail L. Sichitiu. The nominal capacity of wireless mesh networks. In *IEEE Wireless Communications*, October 2003.

[22] Eun-Sun Jung and Nitin H. Vaidya. A power control MAC protocol for ad-hoc networks. In *Proceedings of the 8th annual international conference on Mobile computing and networking*, pages 36–47, Atlanta, Georgia, USA, 2002. ACM Press.

[23] S.J. Lee and M. Gerla. Split multipath routing with maximally disjoint paths in ad-hoc networks. In *Proceedings of the IEEE ICC*, pages 3201–3205, 2001.

[24] J. Li, C. Blake, D. S. De Couto, H. I. Lee, and R. Morris. Capacity of ad-hoc wireless networks. In *Proceedings of ACM MOBICOM*, pages 61–69, July 2001.

[25] Anders Lindgren and Olov Schelen. Infrastructured ad-hoc networks. In *Proceedings of the 2002 International Conference on Parallel Processing Workshops (International Workshop on Ad-Hoc Networking (IWAHN 2002))*, 2002.

[26] Haiyun Luo, Ramachandran Ramjee, Prasun Sinha, Li Li, and Songwu Lu. UCAN: A unified cellular and ad-hoc network architecture. In *Proceedings of ACM MOBICOMM*, Sep 2003.

[27] S. Mangold, S. Choi, P. May, O. Klein, G. Hiertz, and L. Stibor. IEEE 802.11e wireless LAN for quality of service (invited paper). volume 1, pages 32–39, Florence, Italy, Feb 2002.

[28] M. Marina and S. Das. On-demand multipath distance vector routing in ad-hoc networks. In *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, pages 14–23, 2001.

[29] M. J. Miller, W. D. List, and N. H. Vaidya. A hybrid network implementation to extend infrastructure reach. Technical report, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, January 2003.

[30] CMU Monarch. The CMU Monarch Project's Wireless and Mobility Extensions to NS, 1998. Available from http://www.monarch.cs.cmu.edu/.

[31] J. P. Monks, V. Bharghavan, and W.-M.W. Hwu. A power control multiple access protocol for wireless packet networks. In *IEEE INFOCOMM*, pages 219–228, Feb 2001.

[32] Thyagarajan Nandagopal, Tae-eun Kim, Xia Gao, and Vaduvur Bharghavan. Achieving MAC-layer fairness in wireless packet networks. In *Proc. of The 6th*

Annual International Conference on Mobile Computing and Networking, pages 87–98, 2000.

[33] A. Nasipuri, J. Zhuang, and S. Das. A multichannel CSMA MAC protocol for multihop wireless networks. In *Proc. of IEEE Wireless Communications and Networking Conference (WCNC)*, September 1999.

[34] MIT LCS's Parallel and Distributed Operating Systems group. The Grid Ad-hoc Networking Project. Available at http://www.pdos.lcs.mit.edu/grid/.

[35] M. Pearlman, Z. Haas, P. Sholander, and S. Tabrizi. The impact of alternate path routing for load balancing in mobile ad-hoc networks. In *Proceedings of the ACM MobiHoc*, pages 3–10, 2000.

[36] Charles Perkins and Pravin Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, pages 234–244, 1994.

[37] Charles Perkins and Elisabeth Royer. Ad-hoc on-demand distance vector (AODV). In *Proceedings of IEEE WMCSA*, Feb 1999.

[38] Grid Ad-Hoc Networking Project. DSDV and DSR software. Available at http://www.pdos.lcs.mit.edu/grid/software.html.

[39] Saikat Ray, Jeffrey Carruthers, and David Starobinski. RTS/CTS-induced congestion in ad-hoc wireless LANs. In *Proceedings of the IEEE Wireless Communications and Networking Conference 2003*, pages 1516–1521, March 2003.

[40] Nitin Jain Samir, Samir R. Das, and Asis Nasipuri. A multichannel CSMA MAC protocol with receiver-based channel selection for multihop wireless networks. In

Proceedings of the 9th Int. Conf. on Computer Communications and Networks, October 2001.

[41] Jochen Schiller. *Mobile communications*. Addison-Wesley, 2000.

[42] Jungmin So and Nitin H. Vaidya. Multi-channel MAC protocol for ad-hoc wireless networks. Technical report, Coordinated Science Laboratory, University of Illinois at Urbana-Chamapaign, January 2003.

[43] J. L. Sobrinho and A. S. Krishnakumar. Quality-of-service in ad-hoc carrier sense multiple access wireless networks. In *IEEE Journal on Selected Areas in Communications*, volume 17, pages 1353–1368, August 1999.

[44] M. W. Subbarao. Dynamic power-conscious routing for MANETs: An initial approach. In *Proc. of IEEE Vehicular Technology Conference*, pages 1232–1237, 1999.

[45] Z. Tang and J. J. Garcia-Luna-Aceves. Hop-reservation multiple access (HRMA) for ad-hoc networks. In *Proceedings of INFOCOM*, 1999.

[46] L. Tassiulas and S. Sarkar. Max-min fair scheduling in wireless networks. In *Proc. of INFOCOM (3)*, volume 10, pages 320–328, June 2002.

[47] Y.C. Tay and K.C. Chua. A capacity analysis for the IEEE 802.11 MAC protocol. In *ACM Wireless Networks*, volume 7, pages 159–171, March 2001.

[48] Christian Tschudin, Henrik Lundgren, and Erik Nordstrm. Embedding manets in the real world. In *Personal Wireless Communications (PWC) 2003,Venice ITALY*.

[49] Yu Wang and J. J. Garcia-Luna-Aceves. Channel sharing of competing flows in ad-hoc networks. In *Proc. of The IEEE Symposium on Computers and Communications (ISCC ’03)*, pages 87–98, Kemer - Antalya, Turkey, July 2003.

[50] IEEE 802.11 WG. IEEE 802.11f/D5.0, Draft Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution System Supporting IEEE 802.11 Operation, (January 2003).

[51] IEEE 802.11 WG. Draft Supplement for Telecommunications and Information Exchange between Systems—LAN/MAN Specific Requirements — Part II: Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications: MAC Enhancements for Quality of Service (QoS) - IEEE 802.11e/D2.0, Nov 2001.

[52] Alec Woo and David E. Culler. A transmission control scheme for media access in sensor networks. In *Proc. of the 7th Annual Internation Confererence on Mobile Computing and Networking*, pages 221–235, 2001.

[53] S. Wu, C. Lin, Y. Tseng, and J. Sheu. A new multi-channel MAC protocol with on-demand channel assignment for multihop mobile ad-hoc networks. In *Proceedings of the International Symposium on Parallel Architectures, Algorithms, and Networks(I-SPAN 2000)*, pages 232–237, 2000.

[54] Kaixin Xu, M. Gerla, and Sang Bae. How effective is the IEEE 802.11 RTS/CTS handshake in ad-hoc networks. In *Global Telecommunications Conference (IEEE GLOBECOM '02)*, pages 72–76, November 2002.

[55] Shugong Xu and Tarek Saadawi. Does the IEEE 802.11 MAC protocol work well in multihop wireless ad-hoc networks? *IEEE Communications Magazine*, pages 130–137, June 2001.

[56] Xue Yang and N. H. Vaidya. Pipelined packet scheduling in wireless LANs, 2002. UIUC Technical Report.

[57] Fengji Ye, S. Yi, and Biplab Sikdar. Improving spatial reuse of IEEE 802.11-based ad-hoc networks. In *Proceedings of IEEE GLOBECOM*, December 2003.