Are Two Interfaces Better than One?

Mohammad Ahmad Munawar and Paul A.S. Ward
Shoshin Distributed Systems Group
Department of Electrical and Computer Engineering
University of Waterloo, Waterloo, Ontario N2L 3G1, Canada
E-mail: {mamunawa,pasward}@shoshin.uwaterloo.ca

Abstract

Many of the community-area networks use commodity 802.11 hardware to form small wireless networks. Generally organized as a mesh, employing a single channel, and having a few gateways for wider-area access, they tend to offer poor bandwidth to end users. To increase bandwidth, the idea of leveraging multiple interfaces operating on different, non-overlapping, channels has been put forward recently. In this paper, we examine the performance of community wireless networks based on such multi-interface nodes. Our experiments demonstrate that the mere use of more dual-interface nodes does not necessarily create higher capacity. Indeed, in a number of cases we show that the throughput is lower than cases where fewer interfaces are used. We identify three causes for this throughput limitation: channel load, RTS/CTS and exposed nodes, and unfairness due to local traffic. Furthermore, we show that in random topologies, it is very often hard to achieve adequate throughput gain.

Index-terms: Mobile Networking, Home and Ubiquitous Networks, Multi-channel 802.11

I. Introduction

The use of wireless data networks has become widespread, making it easier to access information anytime anywhere. Thanks to standardization, interoperability between vendors has increased and the cost such networks has drastically decreased. It is thus not surprising that in the past few years we have witnessed an ever growing movement towards community networking, which promotes the idea of freely sharing information resources with each other. In these networks users are expected to collaborate by routing each others' packets. Community wireless networks have been deployed in a wide variety of cases to overcome various constraints (*e.g.* absence of infrastructure in rural areas). We are now seeing their implementation in places where there is no compelling reason other than a simple desire to form a community. We believe that such networks will become more widespread and will play an important role in enabling ubiquitous wireless access.

Community wireless networks mostly use 802.11-based products, as they are inexpensive. These networks are generally stationary, organized in a mesh topology, and employ a single channel. Several researchers have shown that throughput of single-channel networks is very limited. For example, Gupta and Kumar [1] show that capacity per user is an inverse function of the square-root of the number of users, assuming that nodes are identical and are optimally located. Li *et al.* [2] have shown that the best achievable throughput for a single flow, with no competing traffic, is theoretically one quarter of the raw bandwidth, and in practice, one seventh, unless the communication is mostly local. The cause of the performance degradation is usually poor sharing of the wireless medium. Researchers have tried to address this issue by improving channel spatial reuse (*e.g.* using directional antennae, power-control, *etc.*) and by dividing the wireless medium into multiple non-interfering channels. These techniques typically presume that the 802.11 systems can be changed as required.

Most multihop 802.11-based wireless systems, typically ad-hoc networks, use one of the channels available even if more channels are available. The Direct Sequence Spread Spectrum (DSSS) technology used in 802.11b/g, for example, provides three non-interfering channels in the North American region. Therefore, a simple solution to improve the capacity that does not require changes to the 802.11 MAC or the hardware is to use multiple interfaces tuned to independent channels on nodes that require high capacity.

The purpose of this paper to evaluate how well a system with dual-interface nodes can perform. We study its performance in various topologies and explain the observed behaviour. The contributions of this paper are several folds. We first compare performance of alternative multi-interface configurations to the single-channel multihop network. We next show bandwidth improvement as well as important problems arising from the use of multi-interface nodes. A distinguishing feature of our work is that we concentrate on the link-layer factors affecting performance.

The remainder of this paper is as follows: Section II, provides an overview of the technologies involved . In Section III, we motivate the need for using multi-interface nodes in wireless networks. In Section IV, we present a number of experiments and analyze their results.

II. Background

A. IEEE 802.11 MAC

The Basic access method of the Distributed Coordination Function (DCF) of IEEE 802.11 is based on the Carrier Sense Multiple Access (CSMA) mechanism. Time is slotted for the purpose of medium access. When nodes sense the medium as free for the Distributed Interframe Space (DIFS) period, they are allowed to transmit. A collision avoidance mechanism is specified to reduce probability of collisions. If the medium is busy, each node waits a random number of slots before accessing the medium. This value is selected from a contention window, that is adaptively adjusted following a binary exponential backoff algorithm. The resulting contention window is reset upon a successful transmission. The 802.11 MAC uses a positive acknowledgment scheme whereby all unicast data packets have to be acknowledged. To avoid the hidden terminal problem, the 802.11 Standard specifies an optional virtual carrier sense mechanism that uses an RTS/CTS handshake. Nodes hearing an RTS or a CTS packet set an internal variable, the Network Allocation Vector (NAV), to the time included in these control packets. Thus, nodes wanting to transmit not only do a physical carrier sense, but also check the NAV before transmitting.

B. Destination-Sequenced Distance Vector

In this work we use Destination-Sequenced Distance Vector (DSDV) [3], which is 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 to all other nodes in the network. The table is updated periodically, or when a station detects change in its neighbours. Routing update packets are broadcast. A node receiving an update packet processes the information, increments the hops metric, and re-broadcasts it. Processing involves modifying the routing table if the update has a fresher sequence number or a better hop count when the sequence numbers are the same. Due to broadcast overhead, DSDV is only suitable for scenarios where mobility is very low [4]. For community wireless networks, such a restriction is acceptable.

C. Related Work

Multi-channel medium access protocols have been proposed in [5], [6], [7], [8], [9], [10]. Most of these schemes propose the separation of the bandwidth into control and data channels. All nodes negotiate use of data channels on the common control channel, and then shift to the chosen frequency for data transmission. A number of these schemes require the use of special hardware. Most of them entail either a new or a significantly different MAC. There is also the presumption in some of this work that there exists either an infinite number of non-interfering channels or that the frequency band can be dynamically allocated. Hardware with dynamic frequency division is not typically available and such allocation needs to consider workload. Moreover, with only three non-overlapping channels available when using 802.11b/g, using one such channel for control purposes may be very inefficient. Others have assumed short channel switching time (about 1 us), even though with 802.11 hardware it takes much longer. The multi-channel schemes mentioned have been developed with mobile ad-hoc networks in mind. The dynamic nature of such networks asks for adaptive channel assignment, which introduces important overhead and reduces capacity. For fairly static network topologies, simpler schemes that alleviate channel alocation overhead can be used.

The capacity of a multi-channel MAC protocol has been studied in [11], where scaling laws for the capacity of multi-channel wireless network in various topologies are derived. In [12], the authors study fairness and capacity in wireless mesh networks that use the same channel. The authors point out fairness problem in these networks when there are few gateways to external networks such as the Internet. Other researchers have studied performance of actual community-type wireless mesh networks. The Roofnet project [13] at MIT experiments with stationary testbeds that use a single channel. The work has exposed shortcomings of the hop-count routing metric. As far as the industry is concerned, several providers offer wireless-mesh network solutions (*e.g.*, [14], [15], [16]). Nortel [17] and Intel [18] are also involved in active research and development in this area. Products offered are based on technologies that vary from using new modulation techniques to mixing available standard equipment such as 802.11b/g/a.

As far as multi-interface solutions are concerned, researchers at Microsoft Research [19] propose a virtual MAC layer that abstracts multiple interfaces, which are assigned non-overlapping channels. A node, which supports their system may communicate with any other node, provided both have interfaces that share a common channel. The channel with best link quality is used for this purpose. This system does not require a channel allocation algorithm, but fails to make full spatial re-use of the channels. Their system is similar to ours in that they target community networks. However, they require the use of products capable of traffic prioritization (*e.g.* 802.11e-compliant) in order to accurately measure quality of links. Recent work by Raniwala and Chiueh [20] also propose a multi-interface system to support a complete wireless access network. In contrast to our work, they use 802.11a in the backbone and 802.11b/g for user access. By using 802.11a, even though they have more non-interfering channels, they cannot leverage the large customer base that already uses 802.11b/g. 4GSystems [21] and Locust World [22] sell multi-interface products. However, there are no studies showing how networks with such products perform.

III. Using Multiple Interfaces

In 802.11 multihop wireless network using one channel, when a node is transmitting, others in its interference zone need to defer transmission. For example, in Fig. 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. A simple technique to alleviate the above issue, is to employ multiple interfaces on a single node such that it can communicate simultaneously on different channels. Advantages of this approach include reduction of contention per area unit and increase in communication parallelism. For example, in Fig. 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 802.11 DSSS physical layer. Furthermore, the carrier-sense range is generally much larger than the transmission range. When channels are reused, interference thus caused reduces throughput. It is possible to minimize this interference by spacing out links that use the same channel. However, this is only possible at the cost of either coverage or lower capacity. For example, in Fig. 1(c) a channel is used for two hops (*e.g.* channel 1 between nodes 0, 1, and 2). If the distance between the nodes is 200 m, and we have three channels, a channel is reused at 800 m. But, capacity per link is reduced, as two consecutive segments use the same frequency.

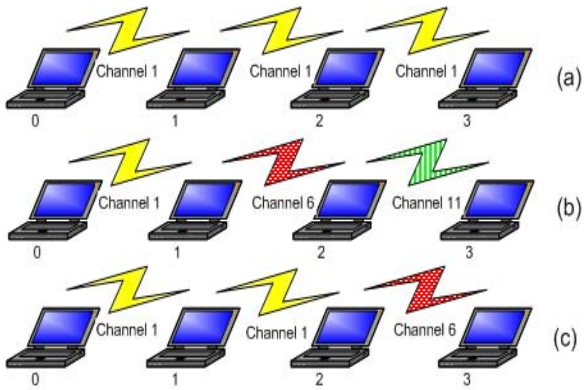

Fig. 1. (a) One channel network (b) Multi-channel network: dual-interface node at each hop(c) Multi-channel network: dual-interface nodes at every two hops

Hereafter, we distinguish between an interface and a node. A node can have multiple radio interfaces. In this paper, we investigate nodes that have at most two interfaces. In our model, the two interfaces are connected at the routing level. Thus the routing protocol not only needs to keep track of routes and associated cost, but also the interfaces from which these routes can be taken. While such a routing protocol can be implemented, for simulation purposes, we chose a design whereby each interface has an independent routing module with different IP addresses, but which are connected. Each module advertises the presence of its counterpart on its channel. We assume that passing packets between the two modules incurs negligible delay. We also modify the ARP module in order for each interface to respond on behalf of the other.

IV. Simulation Results and Analysis

In this section, we present results of our experiments showing performance of multi-interface multi-channel wireless networks. This paper concentrates on throughput, which is an important performance criterion for applications that are not delay sensitive.

A. Simulator Setup

We use ns-2 [23] for our experimentation, which emulates the operation of Lucent's 914 MHz WAVELAN radio with a data rate of 2 Mbps. The modeled transmission range is 250 m whereas the carrier-sensing range is approximately 550 m. For each interface the queue length is set to 50 packets. Since nodes do not move, we disable DSDV periodic updates for more accurate measurement. All measurements are taken at the application layer after the routing information has settled. We use constant bit rate (CBR) over UDP as source traffic with packets of length 1500 bytes. Each measurement is obtained after execution of the simulation for 300 seconds. Unless otherwise stated, each experiment is run five times. We do not experiment with TCP to avoid additional complexity and to concentrate on lower-level system performance.

B. Chain Topology

We begin experimentation with a simple chain topology where nodes are separated by 200 meters. We use three configurations: (a) all nodes are on the same channel (b) we alternate the three non-overlapping channels available on each link (c) we use a channel on two adjacent links, and then alternate every two links.

In these experiments the source is the last node in the chain, while the destination is node 0. The offered load at the source is 2 Mbps. Results in Fig. 2 and Fig. 3 show that both multi-channel configurations offer much higher throughput than the single-channel chain. Surprisingly, one hop/channel with RTS/CTS offers slightly worse performance than the two hops/channel for chains of length four or more. The reuse of channels leads to a dramatic drop in throughput. Because nodes still interfere with others in the carrier-sense range, the longer the chain, the more contention we have along the chain. However, this does not explain alone the steep drop in throughput. More important reasons underlying the lower throughput in the one hop/channel chain are described later in this section.

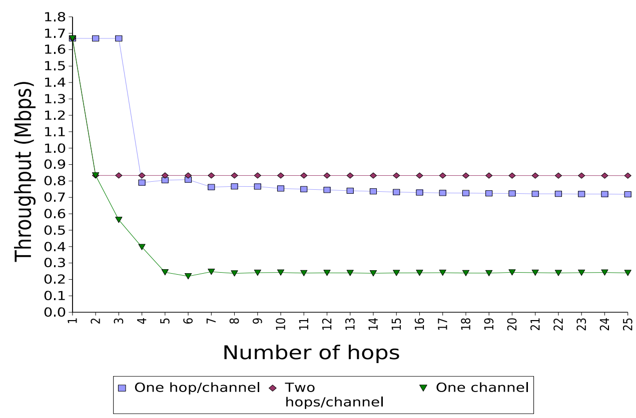

Fig. 2. Throughput vs. Number of hops (with RTS/CTS)

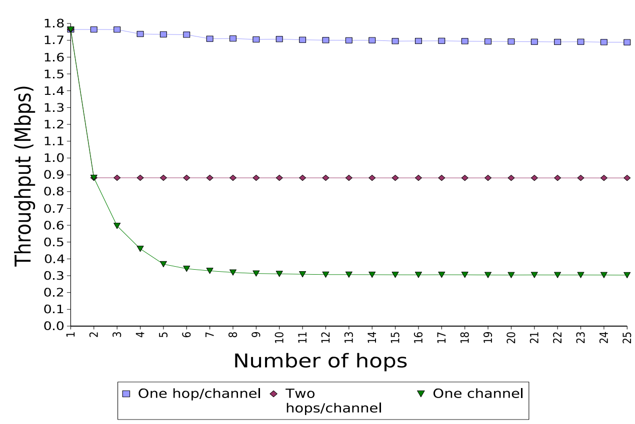

Fig. 3. Throughput vs. Number of hops (Basic)

Throughput for the two hops/channel with RTS/CTS is stable, but reaches 0.83 Mbps only. This stability is due to the reduced interference, as channels are reused farther away than the carrier-sense range. However, given that each channel is used for two adjacent links, throughput is apportioned accordingly.

When RTS/CTS is disabled (See Fig. 3), throughput in the one hop/channel configuration for four hops and more almost doubles when compared to the same experiment with RTS/CTS. It is clear that the RTS/CTS handshake introduces anomalies that limit performance of the one hop/channel setup. RTS/CTS, however, was devised to reduce collisions and detect probability of collision faster for long data packets.

To understand this anomaly, consider a chain of four hops where each intermediate node uses two interfaces. This scenario is depicted in Fig. 4. Interfaces 0, 1, 6, and 7 are on the same channel. While in *ns-2* with the default parameters, a node

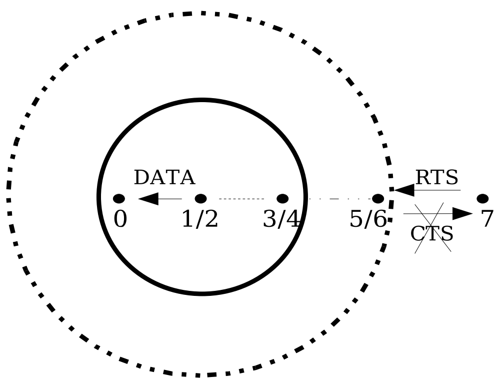

Fig. 4. Example of RTS/CTS backoff anomaly (nodes are 200m away)

at 400 m can still collide with a packet being captured at a receiver, collisions are not the principal cause of the performance degradation. The main reason for the low throughput when using RTS/CTS lies in the large carrier-sense range, which prevents completion of RTS/CTS handshake. In Fig. 4 consider the case where interface 1 is sending a data packet to interface 0. Let our data packet be 1500 bytes long including various overheads. In comparison, RTS and CTS are only 40 and 39 bytes long, respectively. For each packet sent by interface 1 to interface 0, interface 6 would sense the medium as busy and not transmit a packet during this time. Thus any RTS correctly received by interface 6 would remain unanswered. This would in fact halve the throughput, because the two links are operating in lock-steps. Furthermore, interface 7, not receiving a CTS, assumes contention. It doubles its congestion window, and waits for a new random value. Several of these backoff lead to an increasingly larger contention windows and possible drop of the packet, if the retry threshold is exceeded. When the channel eventually becomes free, interface 7 may be in backoff, thus loose an opportunity to transmit. This leads to further decrease in throughput. Unlike CTS, ACK packets are sent after SIFS without doing a carrier sense [24]. This explains why the Basic access scheme is not affected by the large carrier-sense range. As such, link 1-0 and link 7-6 can support a high data rate in parallel.

Our throughput results represent average values over a 300 seconds period. In the short run with RTS/CTS, one hop/channel chain suffers from high variation of throughput per link. Fig. 5 shows the bursty MAC-level transmission rate of interface 1 and 7 during one instance of the simulation along a four-hop chain. This burstiness is a manifestation of the backoff aloritihm, exacerbated by the RTS/CTS versus long data packet problem.

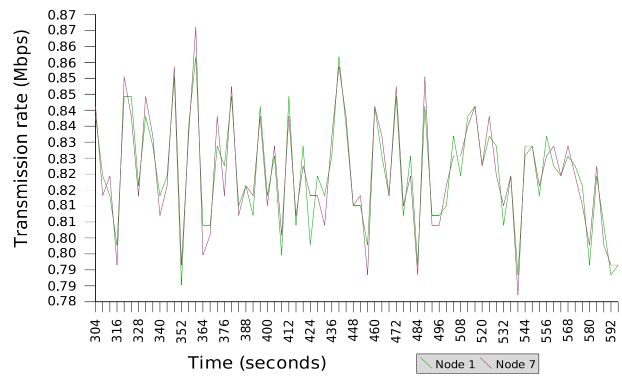

Fig. 5. Transmission burstiness in a chain of length 4

A node that successfully transmits uses the minimum contention window for the next backoff selection. This gives the latter unfair access to the medium, as it is able to hold on to the medium for a longer period by repeatedly accessing the medium with smaller contention windows. The variation in throughput, depicted in Fig. 5, is a manifestation of the binary exponential backoff algorithm on node 7. The excessive backoff problem is also reported by Li *et al.* [2] in the context of a single channel chain. We observe that the two hops/channel configuration keeps channel reuse far enough apart to avoid this problem.

As we can see, multi-interface nodes improve throughput in a chain provided we are careful in enabling RTS/CTS. The Basic access scheme is potentially less power efficient, but in our context it is reasonable to assume a stable supply of power. We now move on to more complex scenarios where the scheduling problem due to RTS/CTS and a large carrier sense range together with excessive backoff are likely to compound.

C. Grid Topology

In this section we study regular grid-like topologies that are based on the chain topology described in Section IV-B. Nodes in the grid operate in an ad-hoc networking fashion (*i.e.* they are producer/consumer of traffic as well as routers).

To experiment with this topology, we use the three grid configurations. Grid 1 uses a single channel and has the same layout as the grid shown in Fig. 6. 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. Grid 2 depicted in Fig. 6 is based on the one hop/channel pattern, and Grid 3 shown in Fig. 7 is based on the two hops/channel pattern. All directly-connected nodes are separated by 200 m. For simplicity, in Fig. 6 and Fig. 7 we label nodes instead of individual interfaces. Nodes with a cross represent dual-interface nodes, whereas nodes with a circle have one interface.

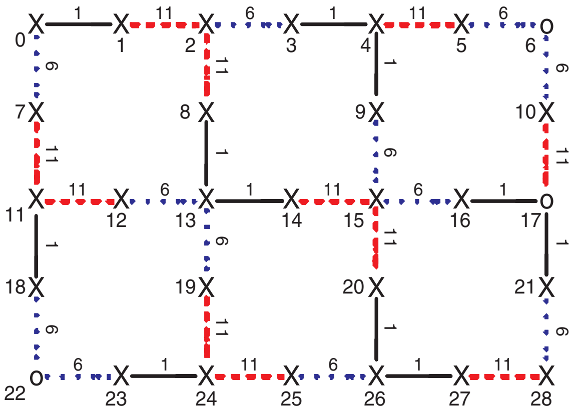

Fig. 6. Grid 2 (26 dual-interface nodes)

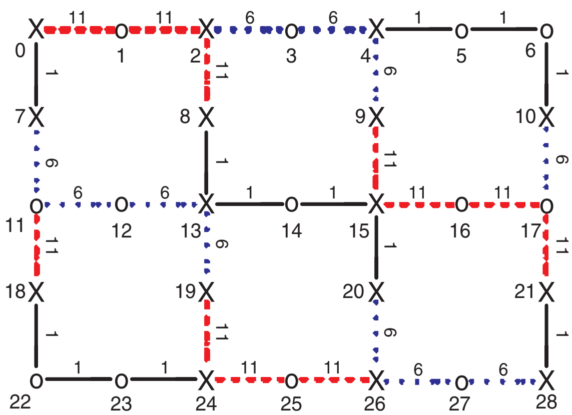

Fig. 7. Grid 3 (16 dual-interface nodes)

In our first experiment we randomly selected 10 unique source and destination pairs on the grid itself. In order for each source to always have data to send, the rate is set to 2 Mbps. In Fig. 8 we present results for 10 such scenarios, when RTS/CTS is used.

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 and depends on the the source-destination pairs selected. In scenarios where many links are shared, where end-to-end flows have long routes (more than 3 hops), and where we have highly contended intermediate nodes, the multi-channel grids offer worse or slightly better throughput than Grid 1. This is the case with Experiment 2, 6, and 8 in Fig. 8. In scenarios where flows take different and relatively short routes, 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 we use fewer interfaces in Grid 3 than Grid 2. 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.

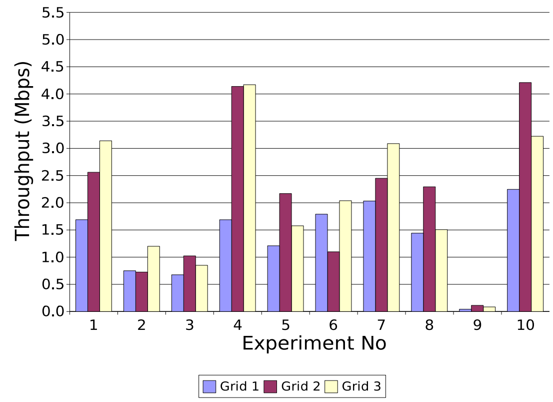

Fig. 8. Aggregate throughput of 10 flows (with RTS/CTS)

In the same experiments we noted that there many flows have no throughput at all. These results are presented in Fig. 9. For each scenario in Fig. 8, we show the average number of flows that are completely starved.

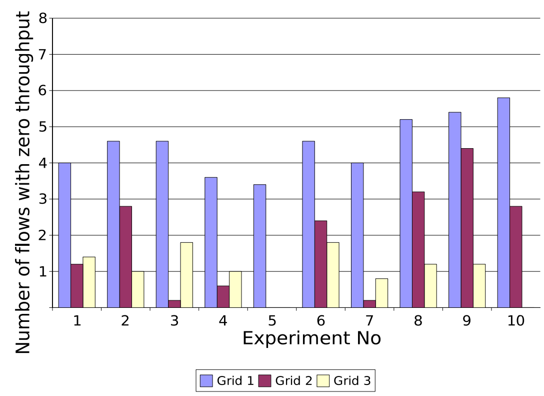

Fig. 9. Average no. of starved flows (with RTS/CTS)

We see in Fig. 9 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. The surviving flows tend to have shorter routes. This is an indication of a high level of unfairness, which mainly stems from the fact that flows being relayed by a node are unable to compete with the traffic generated at that node. A node needs to first access the medium to transmit its packet and then, provided successful reception, its packet may be queued at a forwarding node. It is thus easier for a node to fill up its queue with its own packets, which leads to severe unfairness. This issue is particularly problematic because forwarded packets dropped due to a queue overflow waste considerable bandwidth. It should be noted that unfairness as discussed here is different from MAC-layer unfairness, whereby a node with long data packets could monopolize the medium. Here, unfairness occurs at a higher layer when packets from different nodes compete for a place in the send queue of a relaying node.

Flow starvation also occurs in the multi-interface grids, even though with lower severity. The problem is more severe in Grid 2 because the high throughput that can be achieved along a path, can have a similar effect as the case where a node's own packets gain an unfair share of its packet transmission buffer. 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. Thus unfairness may not only occur at a node generating and relaying traffic, but also happen one or more links ahead on the path from that node. This phenomenon is also present in Grid 3, but is less severe because many paths have comparatively lower bandwidth.

In other experiments not presented here, we increased the offered load each source (starting from 100 Kbps) and studied the effect on the aggregate throughput. Increasing the load results in higher throughput in the multi-channel grids, but also results in increasing unfairness. The increase in throughput is due to gains achieved by a few short flows, which when aggregated, offset the loss experienced by other flows.

To summarize, in this section we have shown that besides the problems mentioned earlier for the chain topology, other issues such as fairness plague community wireless networks that follow an ad-hoc networking approach. Without traffic discipline, it will be hard to realize full benefits of such multi-interface networks.

D. Random Topology

We now turn to the performance of a multi-interface system in a random topology. In contrast to the regular setup studied earlier, here, we randomly place the nodes and also, constrain the topology by a minimum distance between any two nodes.

We further assume that once the nodes are turned on, they operate on a common channel for a short period of time in order to execute a protocol that allows them to select channels. We devised a simple one-pass channel-assignment protocol, which can be further optimized. It uses a combination of broadcast and collection of neighbourhood channel allocation information to select channel that are least used locally. Further refinement that takes traffic load into consideration, such as in [25], is also possible. An example of a random topology with 20 nodes and the result of running the channel allocation algorithm is depicted in Fig. 10. We have kept a minimum distance of 150 m between the nodes.

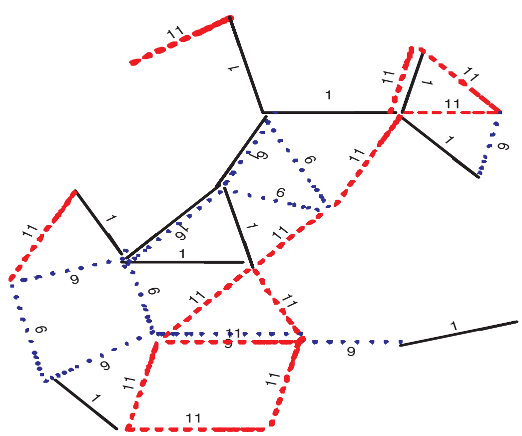

Fig. 10. Sample random networks with 20 access points

In the experiments described here we use an area of about 1100x1100 m with a network of 20 nodes. Three of the 20 nodes are selected as external gateways. Nine nodes receive data from one of the gateways, while the remaining eight send data to the gateways. Each source generates a load of 500 Kbps. We compare two configurations: one where all nodes can choose two channels and another where only a limited number of nodes can choose two channels.

Fig. 11 depict results of experiments with 10 random topologies with the constraint described above. Nodes in this set of experiments do not use RTS/CTS. We observe that there is a significant throughput gain using multiple interfaces in comparison to the one-channel network, even using a very simple channel selection scheme. In some cases, the improvement in throughput can be less than three folds despite using three independent channels. These results also show that having two non-interfering interfaces on all nodes does not necessarily perform better than the case where we have fewer dual-interface nodes. The choice of the right mix is dependent on the topology and traffic conditions. Finally, in urban community wireless networks nodes are closer to each other, which would further affect throughput gains.

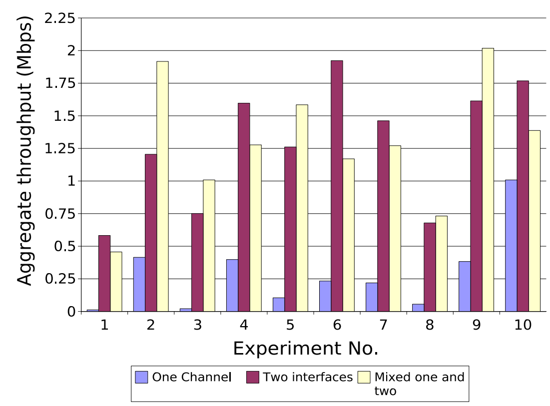

Fig. 11. Single-channel vs. Multi-channel random mesh

V. Conclusions and Future Work

In this paper we studied performance of dual-interface multi-channel wireless networks, which use commodity 802.11 products. We have shown that the performance achieved by such networks can bring about several-fold improvement in capacity. However, we pointed out several problems that restrict throughput gains. In many cases we can still achieve significant performance gains by using fewer dual-interface nodes. We use a simple channel allocation algorithm to appraise performance of a multi-interface system in a random topology. Our work exposes problems that must be addressed in order for such community wireless networks to become more widespread.

Various aspects of our work require further research. In the immediate future, our group is implementing the proposed system to evaluate its behaviour under real-world traffic conditions. We are interested in devising better channel allocation, fairness, and routing algorithms. We also plan to introduce of some dynamism in the channel allocation algorithm to accommodate change in activity and change in topology.

Acknowledgements

The authors would like to thank NSERC and the Software Telecommunications Group at the University of Waterloo for partially supporting this work.

References

[1] P. Gupta, P.and Kumar, "The capacity of wireless networks," in *IEEE Transactions on Information Theory*, vol. 46-2, March 2000, pp. 388–404.

[2] J. Li, C. Blake, D. S. J. D. Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Mobile Computing and Networking*, 2001, pp. 61–69.

[3] C. Perkins and P. Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," in *ACM SIGCOMM Conference on Communications Architectures, Protocols and Applications*, 1994, pp. 234–244.

[4] J. Broch, D. A. Maltz, D. B. Johnson, Y.-C. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad-hoc network routing protocols," in *Mobile Computing and Networking*, 1998, pp. 85–97.

[5] W.-C. Hung, K. E. Law, and A. Leon-Garcia, "A dynamic multi-channel MAC for ad-hoc LAN," in *Proceedings of the 21st Biennial Symposium on Communications*, June 2002.

[6] 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.

[7] N. J. Samir, S. R. Das, and A. 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.

[8] J. So and N. H. Vaidya, "Multi-channel MAC protocol for ad-hoc wireless networks," Coordinated Science Laboratory, University of Illinois at Urbana-Chamapaign, Tech. Rep., January 2003.

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

[10] 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)*, 2000.

[11] J. Li, Z. Haas, M. Sheng, and Y. Chen, "Performance evaluation of modified IEEE 802.11 MAC for multi-channel," in *Advanced Information Networking and Applications (AINA) conference, Xidian University, Xian, China,*, March 2003.

[12] J. Jun and M. L. Sichitiu, "The nominal capacity of wireless mesh networks," in *IEEE Wireless Communications*, October 2003.

[13] MIT LCS's PDOS Group. The Grid Ad-hoc Networking Project. http://www.pdos.lcs.mit.edu/grid/.

[14] Mesh Networks. http://www.meshnetworks.com/.

[15] Strix Systems. http://www.strixsystems.com/.

[16] Tropos Networks. http://www.troposnetworks.com/.

[17] Nortel Networks. Wireless Mesh Network: Extending the reach of Wireless LAN, securely and cost-effectively. White paper. Available at http://www.nortelnetworks.com/corporate/cm/wi/collateral/nn_106481.11-1903.pdf.

[18] L. Krishnamurthy, S. Conner, M. Yarvis, J. Chhabra, and E. T. C. Ellison, C. Brabenac, "Meeting the demands of the digital home with high-speed multi-hop wireless networks," in *Intel Technology Journal*, November 2002.

[19] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for IEEE 802.11 wireless networks," in *Proceedings of BroadNets*, 2004.

[20] A. Raniwala and T. Chiueh, "Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network," in *IEEE Infocom 2005*, 2005, to appear. Available at http://www.ecsl.cs.sunysb.edu/tr/hyacinth-infocom.pdf.

[21] 4G Systems http://www.4g-systems.biz.

[22] Locust World. http://www.locustworld.com.

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

[24] 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.

[25] A. Raniwala and T. Chiueh, "Evaluation of a wireless enterprise backbone network architecture," in *Proceedings of Hot Interconnects*, August 2004, to appear.