Effect of 802.11 Adaptive Exponential Backoffs on the Fluidity of Downlink Flows in Mesh Networks

Adel Aziz
EPFL
Lausanne, Switzerland
adel.aziz@epfl.ch

Roger Karrer
Deutsche Telekom Laboratories
Berlin, Germany
roger.karrer@telekom.de

Patrick Thiran
EPFL
Lausanne, Switzerland
patrick.thiran@epfl.ch

ABSTRACT

Efficient multi-hop traffic management is a need for successful Wireless Mesh Networks (WMNs) deployment. Using an analogy with fluid mechanism, we classify a link as laminar if the packets flow smoothly from the Wireless Access Point (WAP) to the different nodes of the network, and as turbulent otherwise. We identify a particular but frequent collision scenario, which sets the flow to be turbulent, resulting in a strongly reduced downlink end-to-end throughput. We show that the exponential back-off mechanism in a 802.11 WMN is responsible for this problem, which suggests in modification of the current exponential backoff policy of 802.11 for WMNs. We support these findings both with simulations and real measurements on a testbed infrastructure.

1. INTRODUCTION

The deployment of wireless mesh networks that cover large areas such as entire cities is rapidly increasing. This deployment is astonishing because business cases are far from certain and because our knowledge about building and operating mesh networks efficiently is still in its infancy. In particular the backhaul of a mesh network where data is forwarded over multiple hops from and to a wired mesh node and which therefore provides the key cost savings for mesh networks frequently shows dismal single-digit throughputs.

The culprit has been identified in many previous studies: the random access mechanisms of the 802.11 MAC are not efficient in backhaul networks. While the random access provides a fair access for randomly distributed nodes in a given area, it is far from efficient for the particular requirements of a wireless backhaul. The backhaul should forward flows in a 'laminar' way, i.e. packets should smoothly be passed from one node to the next one, in the same way traffic lights should sequentially show green lights, rather than creating a bumpy 'turbulent' traffic pattern due to unsynchronized traffic lights. We argue that such a laminar flow behavior improves the overall network throughput and provides better per-flow end-to-end behavior, such as lower delays and lower jitter.

Toward this objective, this paper makes three contributions. First, we provide evidence that turbulent behavior occurs in backhaul networks with 802.11 MACs. With a simple example, we show that the queues of some nodes rapidly build up whereas other nodes have empty queues and can access the medium as soon as a packet is in the queue. Second, we propose a solution for the above problem: (i) replace the exponential backoff policy of 802.11 by a fixed contention window and (ii) increase the retry limit of retransmitting packets. We show with analytical and simulation results that the total throughput of a mesh backhaul can be increased by 82% in a linear topology. Third, we experimentally evaluate our proposition in the Magnets indoor testbed. The measurements confirm that current 802.11 MACs create turbulent flow patterns, but that our modifications lead to a laminar behavior and that end-to-end throughput and total capacity increase. This evaluation also emphasizes that the benefits can be achieved by simple modifications of 802.11 parameters but without fundamentally changing the 802.11 protocol.

This paper is organized as follows. Section 2 provides background on the problem statement of multi-hop data forwarding, the failure of 802.11, our concept and related work. Section 3 verifies our concept with simulations, and Section 4 presents our experimental evaluation in an indoor mesh testbed. We then conclude in Section 5.

2. 802.11 IN MULTI-HOP BACKHAUL NETWORKS

This section provides background on the problem of multi-hop flow behavior and its causes.

2.1 Problem statement

Wireless mesh networks consist of two parts: an access part that provides connectivity to the user, and a backhaul network that transports data over multiple wireless hops called transit access points (TAPs) from and to a Wired Access Point (WAP) that is equipped with a fixed network line. The logical topology of the backhaul is typically arranged as a k-ary tree, with the WAP as the root and the access points that connect to the users as leafs. For simplicity reasons, we initially consider only linear topologies ($k <= 2$), as depicted in Figure 1. Due to its primordial role of connecting the backhaul to external networks such as the Internet, the WAP

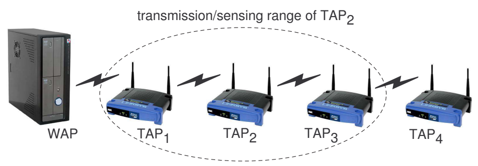

(a) 1-ary topology: De-centralized WAP

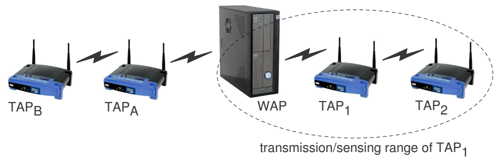

(b) 2-ary topology: Centralized WAP

Figure 1: Illustration of the mesh networks studied.

is the node through which all the traffic flows. Therefore, it is likely to be the bottleneck of the network. Moreover, we focus on downstream traffic, i.e. traffic from the WAP to the users because applications such as Web, multimedia streaming applications or P2P systems typically have larger downstream demands than upstream.

2.2 Objective

The objective of the backhaul network is to forward the packets as efficiently as possible. We define efficiency by two metrics. First, the achieved throughput should be maximized to match the total capacity of the network as closely as possible. Second, the end-to-end performance of each flow should be maximized. In particular, delays should be low and have low variations (for TCP as well as VoIP), and packet loss should be minimized.

We argue that these objectives are best achieved when the flows through the backhaul are *laminar*.

DEFINITION 1 (LAMINAR FLOW). *Laminar flows are characterized by a smooth propagation of packets through the network, where every packet only spends a negligible time in any TAP's buffer. They satisfy the following condition on the buffers B_i:*

$$Prob(B_i\ full) \approx 0 \qquad \forall i \neq WAP \tag{1}$$

The opposite of laminar flows are *turbulent* flows:

DEFINITION 2 (TURBULENT FLOW). *Turbulent flows are characterized by packets spending a significant amount of time in the buffer of TAPs.*

$$Prob(B_i\ full) \gg 0 \qquad for\ at\ least\ one\ i \neq WAP \tag{2}$$

When flows traverse multiple hops, this queuing delay creates perturbation in the flow propagation.

In the scenario described in the next subsection, we will see that it is the first TAP that creates turbulent flows, so that (2) is verified for $i = 1$.

To motivate our argument why laminar flows are desirable, consider the analogy of vehicular traffic. Along a road, traffic passes smoothly through a if the traffic lights are shifted in sequence. Under ideal conditions, a car can cruise at constant speed. Cars only have to wait at the first traffic light. Along the road, no car ever has to wait at a traffic light. Nor do cars have to break and therefore no collisions occur.

Reverting back this behavior to a backhaul network, laminar flows have a constant delay through the mesh network and therefore improve the stability of TCP-based flows as well as the quality of delay-sensitive applications such as VoIP and multimedia streaming. Moreover, having no packets waiting at the TAPs incurs no collisions that might reduce the overall network throughput.

2.3 Failure of 802.11

Unfortunately, the current 802.11 standard [1] that has been designed for fair resource sharing for a single communication range is far from achieving a laminar flow behavior over a multi-hop backhaul network. To understand this statement, it is necessary to understand the basic mechanisms in 802.11. To achieve a fair usage, a node that wants to transmit data sends the medium using RTS/CTS. If the physical layer does not detect activity on the link and the Network Allocator Vector (NAV) counter is null, the medium is considered idle and the node starts transmitting. In the other cases, the channel is considered busy and the node starts to backoff.

The backoff mechanism consists in a counter that is initially uniformly selected in the interval [0; *cw*], where the contention window *cw* has a value between CW_{min} (= 31 for 802.11b and = 15 for 802.11a/g) and CW_{max} (= 1023). The exact *cw* value is obtained by an exponential increase mechanism, i.e. *cw* is initialized at CW_{min} and it is doubled as long as the packet experiences a collision till reaching the CW_{max} limit. Finally, *cw* is reset to CW_{min} after a successful transmission of the packet. The backoff counter consists of slots of 20μs and is decremented as long as the channel is sensed idle and remains frozen if it is not the case. Eventually, when the counter reaches zero, the node sends the message over the medium following the RTS-CTS mechanism.

We now illustrate that these mechanisms lead to turbulent behavior in multi-hop backhaul networks. Figure 2(a) depicts the transmissions as a function of the time, whereas Figure 2(b) shows the corresponding queues and the values of CW_{min} for the topology depicted in Figure 1. We assume that the WAP has always traffic to send, so that its buffer is full (which we denote by ∞), and that TAP_1 has already 4 packets buffered. The build-up of the queues that lead to a turbulent behavior can be separated into 4 phases:

1. **Phase 1:** Packets are sent from TAP_1 to TAP_2 and TAP_3. At the end of this phase, each buffer contains at least one packet.

2. **Phase 2:** TAP_3 transmits a packet to TAP_4. TAP_1 is out of the sensing range of TAP_3: it is therefore

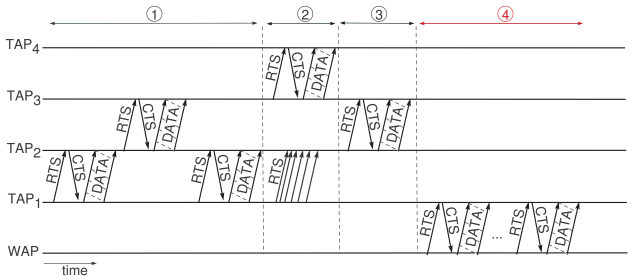

(a) Link activity. ACK messages are voluntarily omitted for readability purpose.

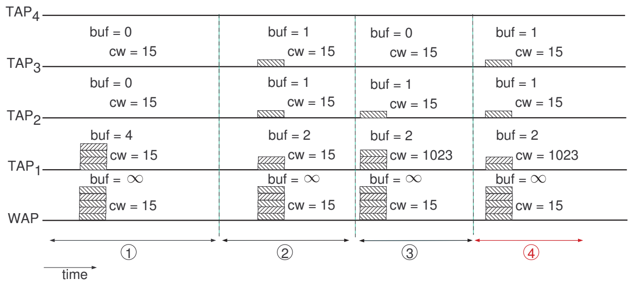

(b) Buffer size and *cw* evolution at the beginning of each phase.

Figure 2: Illustration of the perturbation creation due to the exponential backoff of MAC 802.11.

unaware of this transmission, and sends unsuccessful RTS. These RTS make WAP set its NAV properly, and increase the contention window of TAP_1 up to its maximal value of *cw*=$CW_{max} = 1023$.

3. **Phase 3:** TAP_2 transmits a packet to TAP_3. As the WAP is unaware of this transmission, its back-off counter is not frozen and will eventually reach 0. On the other hand, the NAV of TAP_1 is set by the RTS, which prevents it to decrement its contention window. Therefore, the contention window of TAP_1 remains at a high value (around CW_{max}).

4. **Phase 4:** The transmission of TAP_2 terminates. TAP_1 and WAP still have packets to send and compete for the channel. However their competition is not fair, because the contention window of TAP_1 is much larger than that of the WAP (1023 compared to 31 for 802.11b (or 15 for 802.11a) in our example, a ratio factor of 32 (or even 64)!). This unfair advantage will make WAP win the channel many times in a row. As a result, the buffer of TAP_1 builds up. This increase leads to the perturbation in the fluidity of the data flow.

2.4 Proposed Solution

To solve the buffer building-up issue, the consequences of the physical limitation should be reduced by preventing an unfair competition for the medium between TAP_1 and WAP due to *cw*. We argue that a possible solution to reach this goal is achievable with 2 modifications within 802.11:

- The exponential backoff mechanism is disabled and replaced by a fixed value for *cw* to ensure that unfair competition among the WAP/TAPs does not occur independently of the communication taking place previously.

- The 'Short Retry Limit' value which sets the maximum number of attempted transmissions before dropping a packet should be increased. When the exponential backoff mechanism is disabled, the time needed to reach the "Maximal Retry Limit" decreases. An increase in the retry limit avoids that packets are dropped too early. Packets that have left the WAP should not be dropped by any of the TAPs.

These two modifications require just changes in the parameter values of 802.11 and are therefore easy to implement.

2.5 Related work

Multi-hop wireless networks impose an interesting set of challenges in general [2] and in particular in experimental indoor and outdoor settings [3].

Our work focuses on mesh nodes with a single WiFi card because most mesh networks today are built with single cards. Our work therefore contrast solutions for multi-channel or multi-antenna systems ([4], [5], [6]).

Our work aims at understanding and addressing challenges for multi-hop networks at the MAC layer. Our approach therefore differs from related work aimed at MAC layers for single-hop communication, e.g. [7] and [8].

In [9], the authors also focus on MAC layer performance for multi-hop mesh networks. However, their approach is based on buffer queue management, while our solution targets MAC layer parameter.

Recent work ([10], [11]) also discuss the hidden node situation. In [11], the authors focus on the routing instability problem and propose source rate limiting as a solution. Complementary solutions to solve the inter-flow unfairness are analyzed in [10] through simulation. Our work differs from both these approaches by focusing on the intra-flow behavior and presenting simulation as well as experimental results to support our analysis of the impact of MAC 802.11 backoff policy.

Finally, the methodology of applying flow models from fluid physics has been successfully used, e.g for vehicular traffic [12]. We are exploiting and combining models from both areas now to model multi-hop traffic.

3. THROUGHPUT ANALYSIS

This section verifies the above claims and findings with simulations and assess the impact on the throughput.

3.1 Setup

We set up a linear (1-ary and 2-ary) topology that matches the topology shown in Figure 1. We simulate a topology of n nodes per branch, where n is varied from 4 to 20 TAPs, not including the WAP. The distance among the TAPs is chosen such that the sensing and the transmission range include the direct neighbors, but not the neighbors that are 2 hops away. The link capacity is set to $1Mb/s$ and we use packet sizes of 1500 Bytes.

Given these values, the theoretical single-link throughput can be calculated as

$$\frac{PAYLOAD}{DATA + RTS + CTS + ACK + t_{BACKOFF} * bw} * bw \tag{3}$$

With $PAYLOAD = 1500$ Bytes, $DATA = 1572$ Bytes, $RTS = 44$ Bytes, $CTS = ACK = 38$ Bytes, $t_{BACKOFF} * bw = 40$ Bytes and bandwidth $bw = 1/8 * 10^6$ B/s, we get a theoretical single-link throughput of 108.26 kB/s.

For multi-hop topologies, theoretical maximal throughput can be computed considering the maximal spatial reuse of a k-ary topology. Assuming the standard 2-hop collision model, i.e. 2 links can be only active simultaneously if they are separated by 2 other intermediate links, the throughput for a 1-ary topology (respectively, 2-ary topology) is easily computed to be one third (respectively, one half) of the capacity [7]. Therefore, the upper-bound on the throughput performance is $36.09kB/s$ for a 1-ary topology and $54.13kB/s$ for a 2-ary topology.

3.2 Simulation Results

Figure 3 shows the impact of the proposed modifications to the 802.11 parameter values as a function of the number of nodes for 1-ary topologies (Figure 3(a)) and 2-ary topologies (Figure 3(b)). The 4 lines denote the throughput derived from our analytical analysis, with exponential backoff (standard), with fixed contention window *cw* and with both fixed *cw* and significantly increased retry limit (1000). Such an extreme increase is motivated by the intuition that once a packet used bandwidth to leave the WAP, this resource is wasted if the packet is dropped further in the network.

First, we note that the standard 802.11 with exponential backoff achieves a dismal 44% of the theoretical throughput for $n = 20$ nodes. Moreover, significant throughput degradations are already visible for multi-hop networks of size 4. Second, with a fixed *cw* , the throughput achieves 66% of the theoretically achievable throughput for $n = 20$ nodes, and the throughput remains as high as 86% for a 4-hop topology. For fixed *cw* and increased retry limit, the throughput achieves 79% of the theoretical maximum even for network sizes of $n = 20$ nodes.

For 2-ary topologies, the results in relative terms are comparable to those of a 1-ary topology. In particular, our proposed solution achieves 87% of the theoretical limit for $n = 20$ nodes, 75% with fixed *cw* , and 70% for standard 802.11.

These results confirm that the modifications of the 802.11 parameter values have a significant impact on the effective throughput of a multi-hop wireless network.

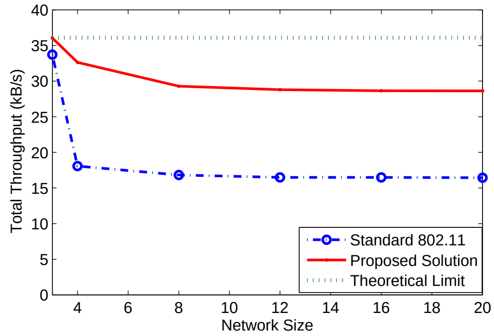

(a) 1-ary topology.

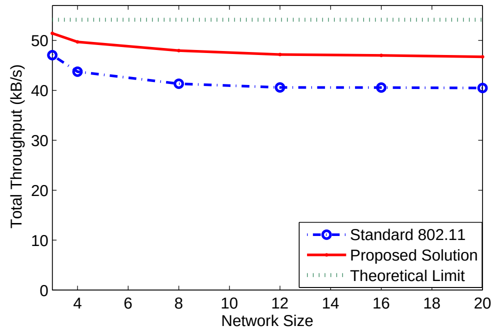

(b) 2-ary topology

Figure 3: Performance gain achievable by removing the exponential backoff policy and increasing the *short retry limit*.

4. EXPERIMENTAL EVALUATION

This section verifies the proposed modifications for 802.11 with measurements in a wireless testbed. We emphasize here also that the modifications were readily ”‘implemented”’ because we did not change the MAC layer protocols, only the parameter values.

4.1 Testbed setup

We perform our measurements in the indoor mesh testbed of the Magnets project [13]. We deployed two 5-hop topologies as depicted in Figures 4 and 5. Deploying the mesh nodes in this fashion allows us to overcome the hardware limitations that prevent us from adjusting the sensing range. The topologies allow us to closely match the interference model of our linear scenario in Figure 1. In particular, in the first topology, we deployed the WAP and the 4 TAPs on one floor, as depicted in Figure 1(a). Here, TAPs 1 and 3 are not entirely closed out as they still sense each other. In the second topology, we deployed the nodes on 4 adjacent floors. The construction of the building allows for a good visibility of nodes that are on neighboring floors, but prevents sensing when the nodes are 2 floors apart.

During the deployment and measurements on the testbed,

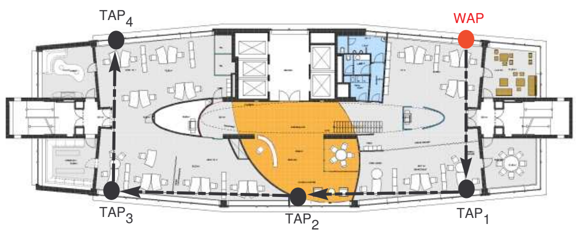

Figure 4: Topology 1: all nodes are on the same floor. The dotted arrows represent the directional flow through the wireless links.

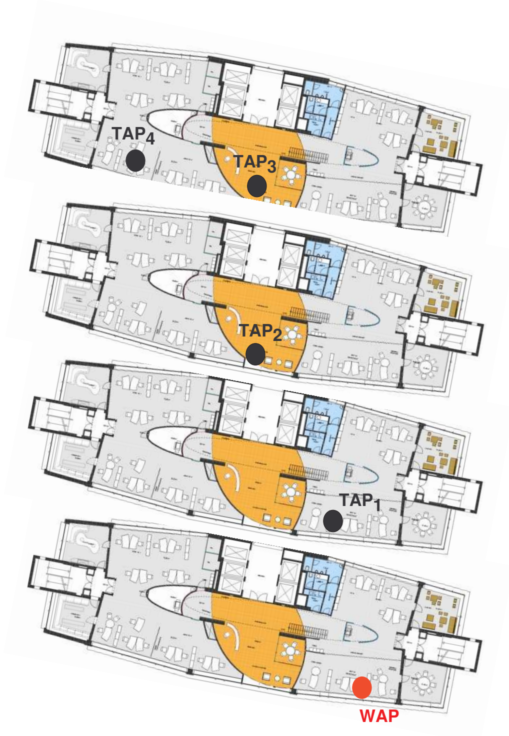

Figure 5: Topology 2: nodes are on adjacent floors.

we made similar observations as in [3] concerning the significant performance variability to millimeter changes of the position or direction of the antenna. Such variations do not impact our results as we maintained the location strictly unchanged during the simulation rounds.

The nodes consists of Routerboards 532 that can hold up to 6 WiFi cards, however, we only use 1 card per board. The WiFi card is an off-the-shelf Atheros-based 802.11a/b/g card. We use the 802.11a mode to avoid interference from other networks and fix the channel to 5.32 GHz. For the same reason, we run the experiments at night. The cards are connected to 3dB indoor omni-directional antennas. The boards run the Kamikaze version of OpenWRT 2.6.21.5 with the MadWifi driver. At the network layer, we use fixed routing to exclude routing messages and potential problems from route changes.

As traffic source and sink, we use 2 Linux-based PCs. On these PCs we run *iperf* [14]. The sender is connected via an Ethernet connection to the WAP, the receiver is also connected via a fixed line to TAP 4. An experiment consists of multiple runs with different values for CW_{min}. For each run, UDP traffic is generated at a rate of 10 Mb/s. This rate is far above the network capacity and therefore ensures that the WAP always has packets in its buffer to achieve the conditions described in Figure 2. Each run lasts for 150 seconds. In our evaluation, we ignore the first 50 seconds to avoid initial fluctuations such that we are sure to operate in a stationary regime.

For each run, we log the achieved throughput and average it for each second. The 100 obtained values are then use to compute an average over 100 seconds together with confidence intervals obtained using the normality assumption.

The results present the comparison of standard 802.11 with our proposed solution for different value of CW_{min}. By standard 802.11, we consider keeping the exponential backoff with $CW_{max} = 1023$ and all the parameter of 802.11 constant while only varying CW_{min}. On the other hand, as defined in Section 2.4, our proposed solution consist in fixing the contention window at *cw*= CW_{min} and increasing the retry limit to 1000 to match our simulation model.

4.2 Measurement Results

Figure 6 shows the multi-hop throughput obtained in our testbed, as a function of the value of CW_{min}. Note that the x-axis is logarithmically scaled because the values are typically powers of 2. First, considering the lines in Figure 6(a), we note that the throughput rapidly degrades as a function of the value of CW_{min} after some initial increase. The initial increase can be explained by the reduction of the collision probability due to the *cw* increase. Second, comparing 802.11 against our proposed solution, we note a difference of roughly 0.5 Mb/sec, or between 10% and 60% in relative terms.

For the second topology, the throughputs shown in Figure 6(b) shows three significant differences compared to the results in Figure 6(a). First, the throughput is significantly lower. This low throughput can be attributed to the larger distances and in particular the ceilings in the buildings that damp the signal. Therefore, the achieved rates are more than 50% lower than those of the previous experiment. Second, we do not see the initial increase in the throughput for low values of CW_{min}. This findings indicates that an optimal CW_{min} value is a topology dependent parameter. Finally, the difference between standard 802.11 and our proposed solution is more exposed. Our solution outperforms standard 802.11 by more than 1 Mb/sec, a net improvement of more than 100%!

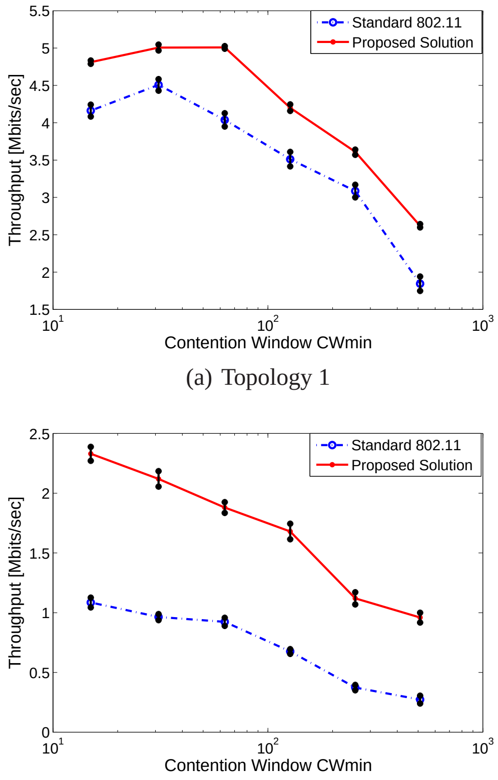

(b) Topology 2

Figure 6: 5-hop throughput as a function of *cw* values with confidence intervals.

5. CONCLUSIONS

This paper presents novel insights into the behavior of MAC layer protocols on the performance of a multi-hop wireless backhaul networks. The detailed understanding of the flow behavior over multiple hops is crucial for end-to-end flow properties and the use of the network capacity. The understanding that the backoff mechanism leads to turbulent flow behavior and thus the above drawbacks is vital for the design and deployment of wireless mesh networks.

Our results are consistent in model, simulations and the experimental evaluation in our testbed. This conclusion is particularly important because the effect of contention is local, i.e. affecting the communication of neighboring TAPs only. However, we show that this local event affects in fact the resource usage of the entire network as well as the end-to-end performance.

The concept of laminar and turbulent flows is a promising approach towards understanding and modeling MAC layer behavior, but it has the potential to be suited for higher layer behavior, such as routing or end-to-end congestion control. In future work, we will continue our study on flow behavior in general, as well as the impact of interacting flow behavior, such as TCP over multi-hop mesh networks.

6. REFERENCES

[1] *IEEE802.11, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std., August 1999

[2] R. Karrer, A. Sabharwal, and E. Knightly, *Enabling large-scale wireless broadband: the case for TAPs*. In *SIGCOMM Computer Communication Review*, 34(1), January 2004.

[3] K. Papagiannaki, M. Yarvis, and W. S. Conner, *Experimental Characterization of Home Wireless Networks and Design Implications*, In *INFOCOM'06*, Barcelona, April 2006.

[4] S. Roy, A. K. Das, R. Vijayakumar, H. M. K. Alazemi, H. Ma, and E. Alotaibi, *Capacity scaling with multiple radios and multiple channels in wireless mesh networks*. In *Proc. of The FirstIEEE Workshop on Wireless Mesh Networks (WiMesh)*, Sept. 2005

[5] J. Tang, G. Xue and W. Zhang, *Maximum Throughput and Fair Bandwidth Allocation in Multi-Channel Wireless Mesh Networks*. In Proc. IEEE Infocom, Barcelona, April 2006

[6] A. Raniwala and T. Chiueh, *Architecture and algorithms for an IEEE 802.11-based multi-channel wireless mesh network*. In *Proc. of IEEE Infocom*, 2005

[7] M. Durvy and P. Thiran, *A packing approach to compare slotted and non-slotted medium access control.*, In *INFOCOM'06* Barcelona, April 2006.

[8] V. A. Siris and G. Stamakis, *Optimal CWmin Selection for Achieving Proportional Fairness in Multi-Rate 802.11e WLANs: Test-bed Implementation and Evaluation*, In *WiNTECH'06*, Los Angeles, September 2006.

[9] N. S. Nandiraju, D. S. Nandiraju, D. Cavalcanti, and D. P. Agrawal, *A Novel Queue Management Mechanism for Improving Performance of Multihop Flows in IEEE 802.11s based Mesh Networks*. In *Proc. of IPCCC-2006*, Phoenix, Arizona, April 10-12, 2006.

[10] L. Li and P. A. S. Ward, *Structural Unfairness in 802.11-basedWireless Mesh Networks*, In *CNSR 2007*, New Brunswick, Canada, May 2007.

[11] P. C. Ng and C. C. Liew, *Throughput Analysis of 802.11 Multi-hop Ad hoc Networks*, In *IEEE/ACM Transactions on Networking*, June 2007.

[12] K. Nagel and M. Schreckenberg, *A cellular automaton model for freeway traffic*. In *J. Phys. I France 2*, pp. 2221-2229, December 1992.

[13] R. Karrer, I. Matyasovszki, A. Botta and A. Pescape, *MagNets: experiences from deploying a joint research-operational next-generation wireless access network testbed*, In *TridentCom'07*, Orlando, FL, May 2007.

[14] A. Tirumala, F. Qin, J. Dugan, J. Ferguson, and K. Gibbs, *Iperf: The TCP/UDP bandwidth measurement tool*, available at *http://dast.nlanr.net/Projects/Iperf/*