

Evaluation of the Throughput, Delay, Drop and Fairness Characteristics of Ethernet LAN

Project Report by

Vasily Tarasov and Yury Puzis (Computer Science, SUNY at Stony Brook)

Introduction

In this report we demonstrate throughput, delay, drop probability and fairness characteristics of Ethernet LANs. We used ns2 network simulator to get presented data. The network we simulated was typical Ethernet network with one server and multiple clients. For each experiment result we gave appropriate analysis.

Experiment A: Throughput vs Network Load

Network load in the Ethernet LAN fundamentally depends on the number of the nodes connected to the LAN and on the bit rate produced by each node. In order to take this fact into account we perform two experiments.

In the first one we fix the number of nodes to 3 (two of them transmit the information, and the remaining one receives). Then we vary the bit rate of each of two transmitting nodes from 0.1 to 10.0Mbps with a step 0.1Mbps. It means that the load on the network changes from 0.2Mbps to 20.0Mbps with a step 0.2Mbps.

In the second experiment we fix the bit rate of each node to 0.1Mbps, but change the number of nodes from 3 to 201 (again one node acts as a receiver). In this case again the network load ranges from 0.2Mbps to 20.0Mbps.

Figure A1: Throughput vs Load (varying bit rate) for different frame sizes




Figure A2: Throughput vs Load (varying the nodes number) for different frame sizes

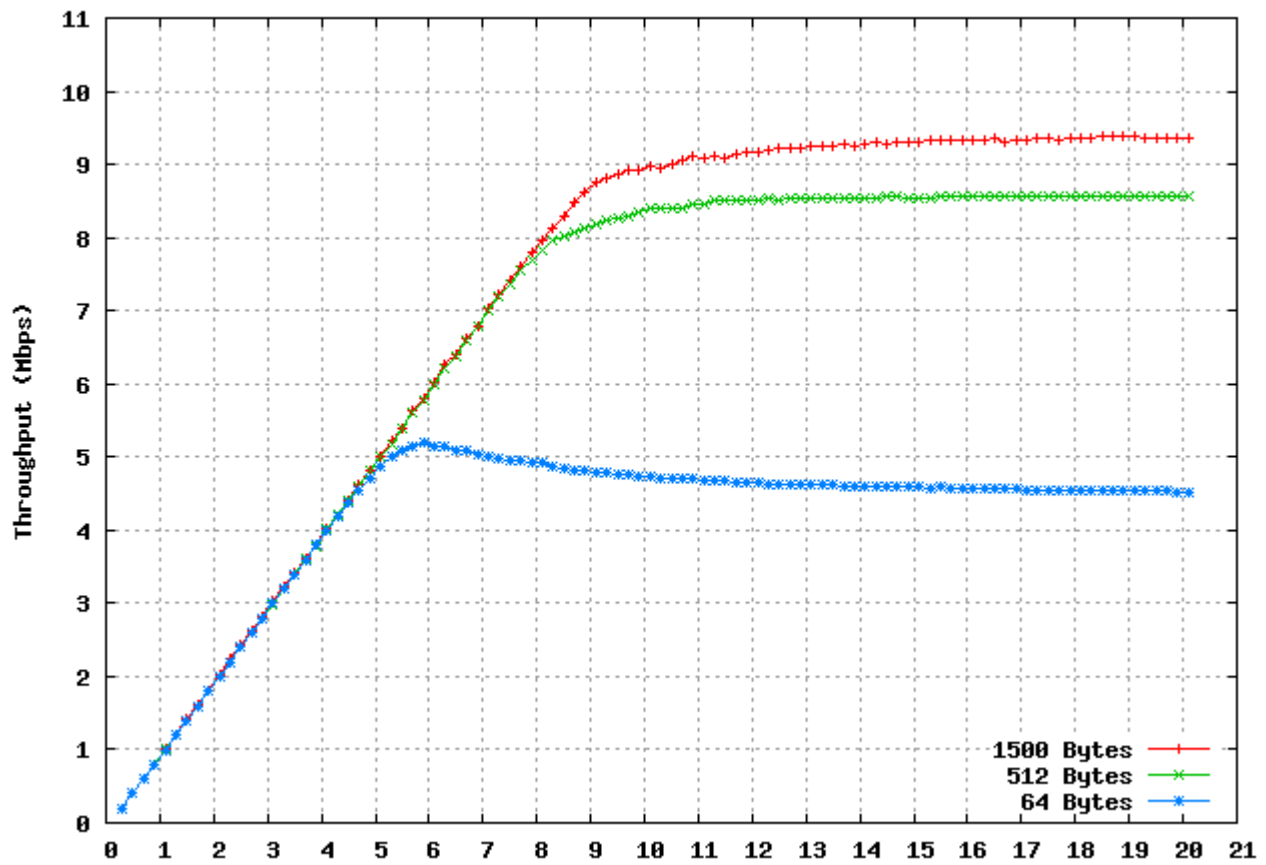



Figure A3: Throughput vs Load for two types of traffic increase




Each experiment was running for 10 seconds of simulation time, which is enough for stable results. The TX queue length was set to 1000 packets, because this is a default value for most of the cards on the Linux kernel, so it is reasonable to use it. We performed each experiment for the frames of size 64, 512 and 1500 bytes.

Figure A1 demonstrates the results of the first experiment for the different packet sizes. The throughput for 512 and 1500 byte frames graphs reaches maximum possible throughput of 10 Mbps when offered load reaches the same value. The reason is that the maximum network capacity is also 10 Mbps. The graph for 64 byte frames flattens out earlier, at 7 Mbps of offered load and throughput, because smaller frame size results in larger number of frames which consequently leads to a larger drop rate.

Figure A2 demonstrates the results of the second experiment for the different packet sizes. As evident, the result is nearly identical to the one on Figure A1.

Figure A3 compares the results of the first and second experiments for the packet size 512 bytes on one graph. The graph for varying bit rate maxes out slightly later, and provides throughput slightly closer to maximum possible value. There is however no substantial difference between the two scenarios.

Experiment B: Delay vs Network Load



In this experiment we measure the effect of change in network load on packet delay. To change network load we adopt approach demonstrated in Figure A2: 201 hosts of which 200 offer load of 0.1Mbps to a total of 20Mbps network load. Each experiment is executed for 10ms. Like in the

Figure B2: Pure Network Delay vs Load for Different Frame Sizes

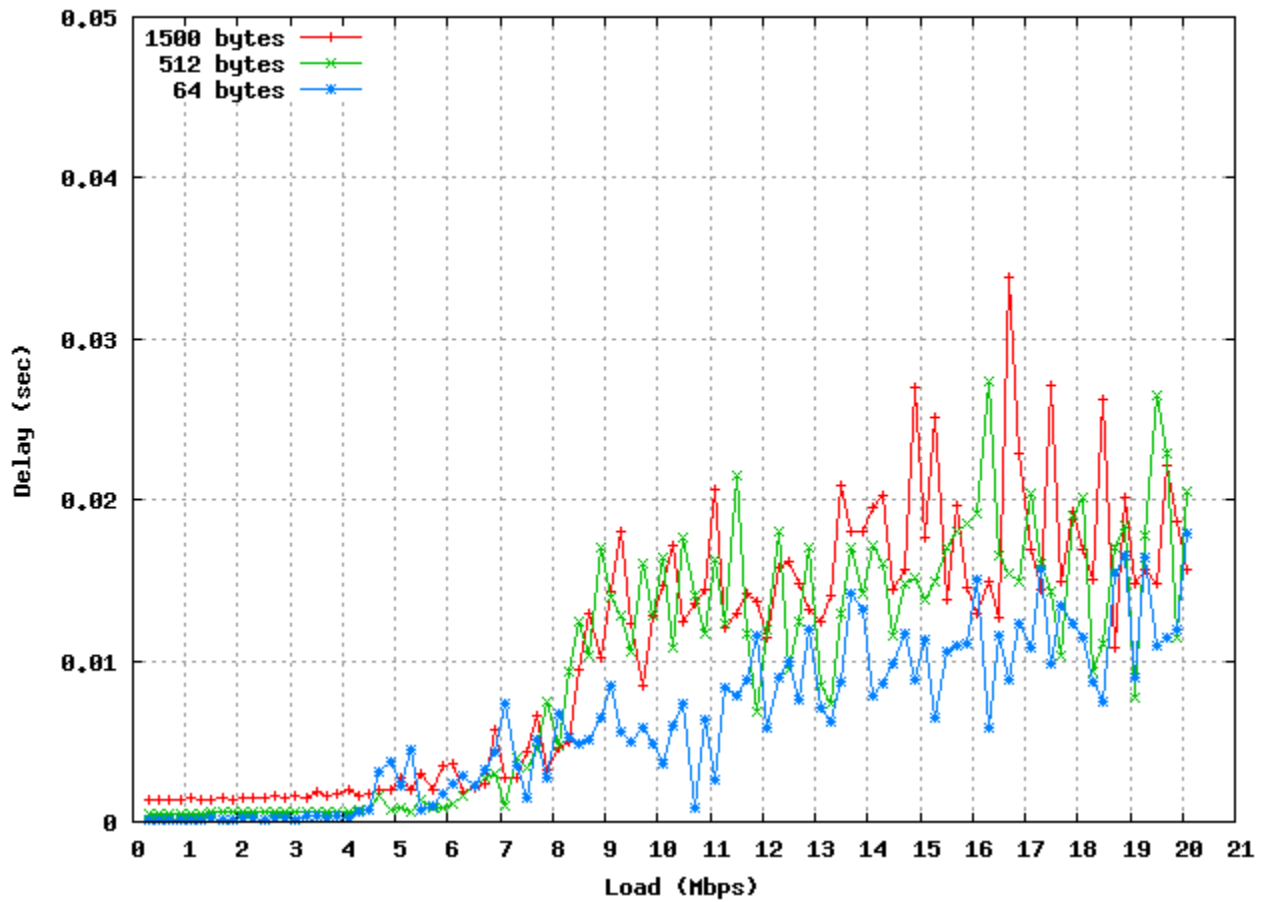
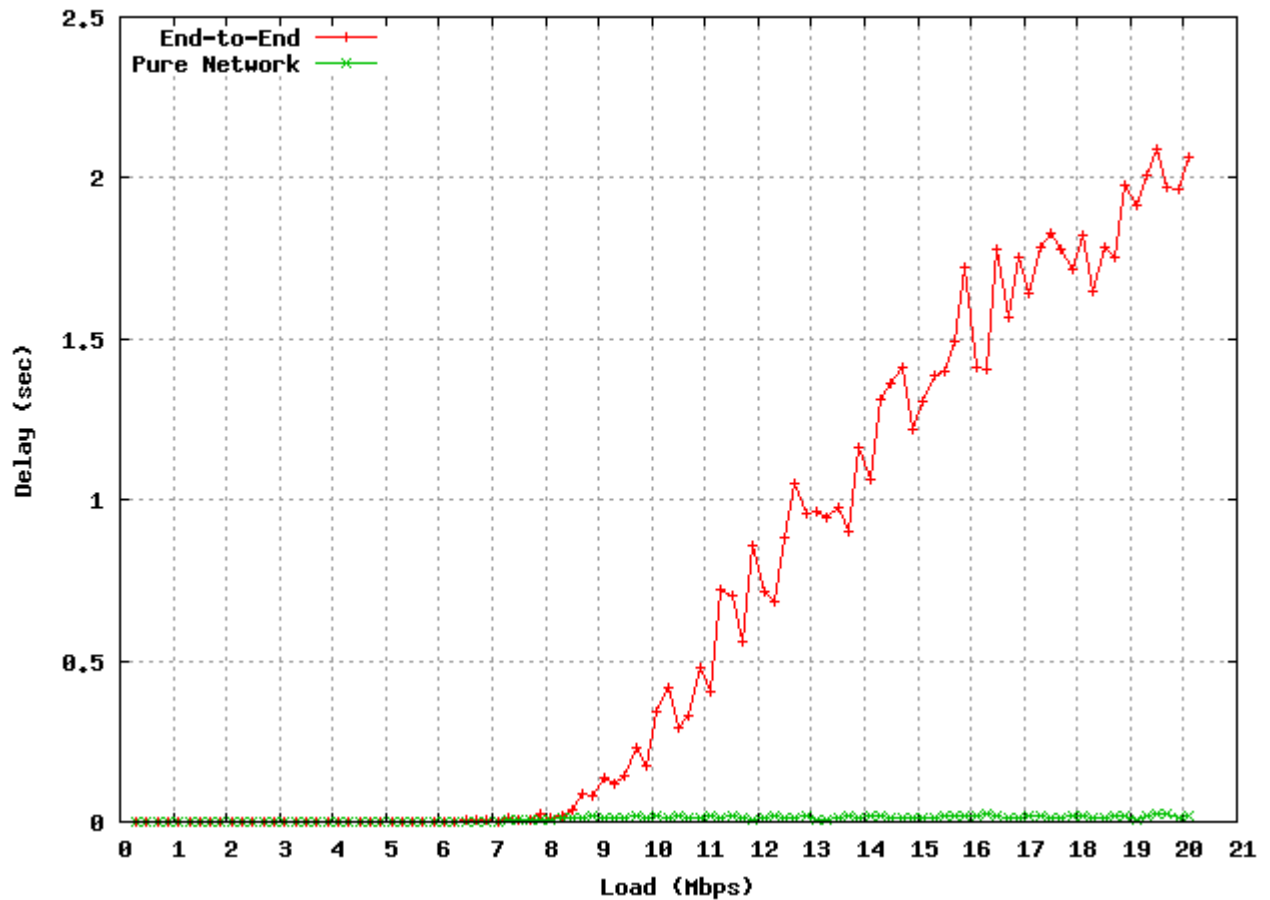


Figure B3: Comparison of Pure Network Delay and End-to-End Delay



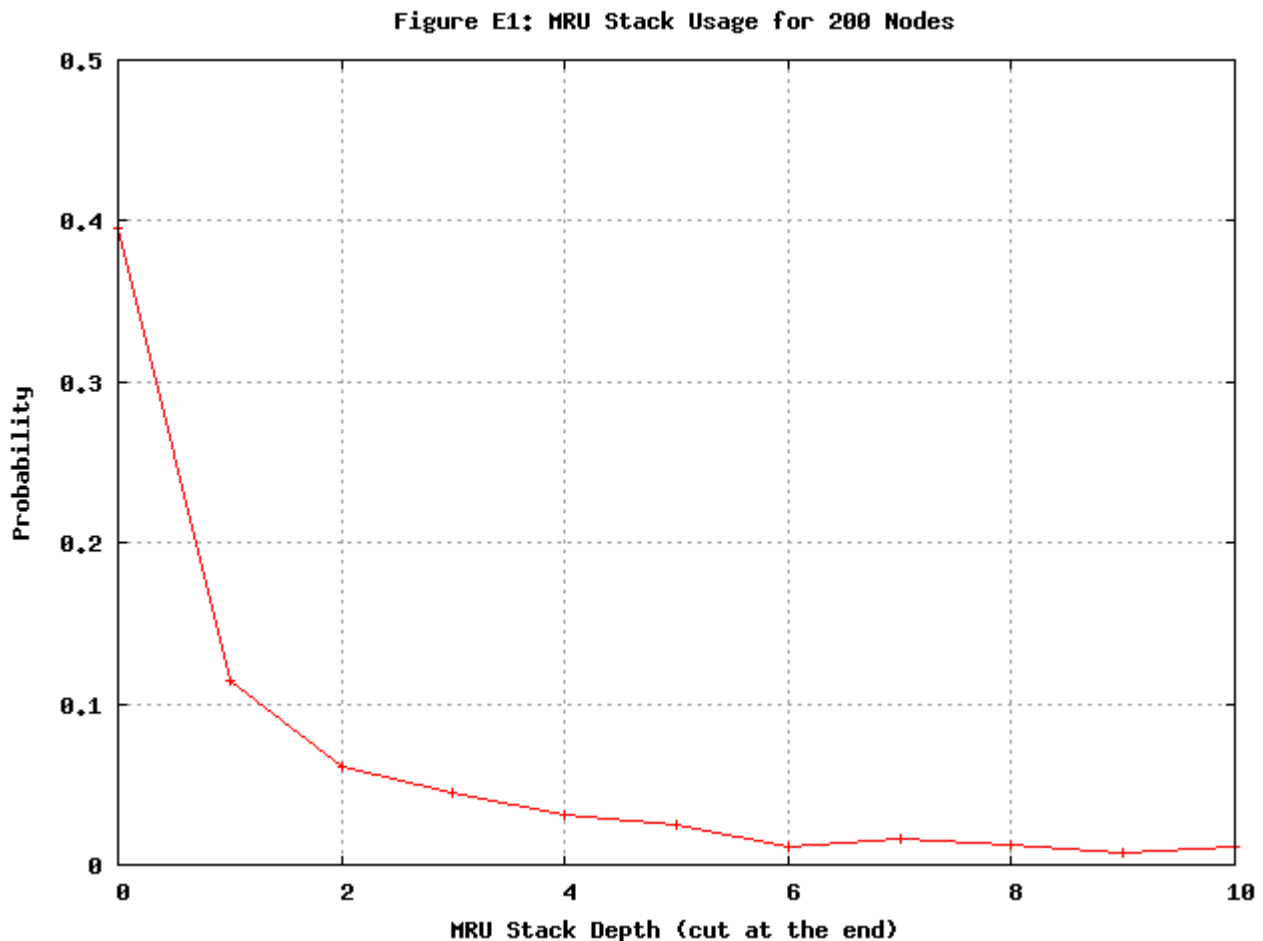
previous experiments the queue is set to 1000.

Figure B1 demonstrates result of experiment for end-to-end delay: delay from initial transmission of packet from source until packet is received at destination. The largest delay, 4.5 sec, is exhibited by graph for 64 byte frame size. We can speculate that the reason is in the fact that small frame size results in larger number of frames, which in turn results in more dequeue attempts and more collisions and backoffs.

Figure B2 demonstrates the above point by measuring pure network delay (which excludes delay at the buffer). Notice that the scale of Figure B2 is much smaller than that of Figure B1. This means that buffer delay is the major factor in causing packet delay, while network delay is the minor factor. This is explicitly shown on the Figure B3 where both delays can be compared side to side.

Experiment E(Problem 2): Ethernet Unfairness


In order to demonstrate short term unfairness of Ethernet we take the approach of Molle¹. As usual in our experiments we consider 200 nodes sending packets to a single destination node at 0.1Mbps. On Figure E1 you can see the dependency of probability that the identity of the sender of a randomly chosen packet will be the Kth most recently seen source address on the network from Most Recently Used (MRU) stack depth. In case of a completely fair network the graph should be a horizontal line.



However, you can notice that this is not so: the graph demonstrates the effect of channel capture by few

¹Molle, M. L., "New Binary Logarithmic Arbitration Method for Ethernet, "Technical Report CSRI-298, Computer Systems Research Institute University of Toronto , April 1994 (Revised July 1994).

Figure E2: MRU Stack Usage for 2 Nodes



“lucky” hosts at the expense of others. In this case the probability that the same node will send the next packet is 0.4.

Figure E2 demonstrates the same effect for two sender nodes. One of the nodes captures the channel and starves the other node. Here the probability that same node will send the next packet is almost 1. This result is logical since in the network with two nodes one of the nodes is likely to monopolize the channel after its initial success to transmit a packet.

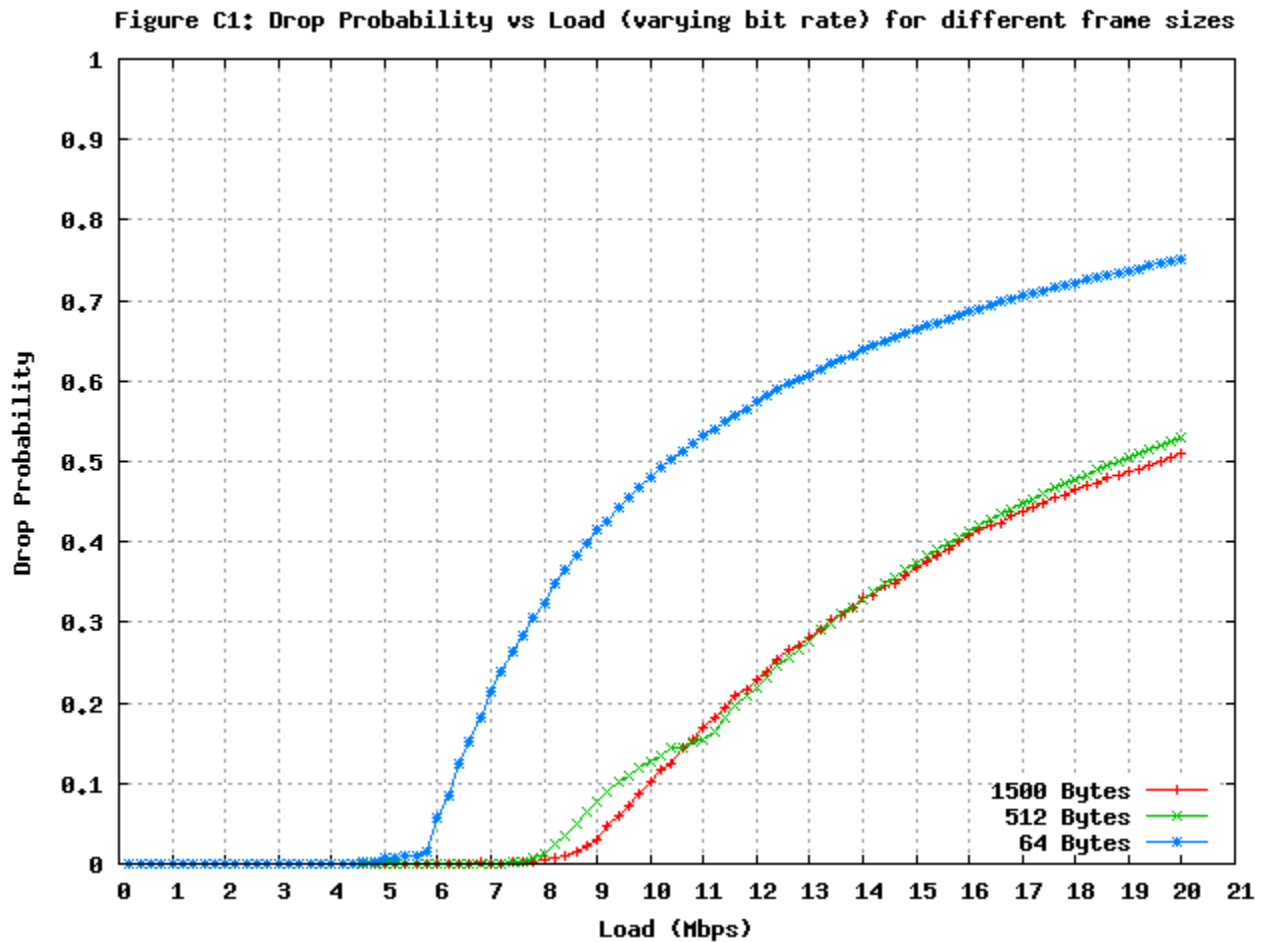
Experiment C: Drop Probability vs. Network Load

In this experiment we measure dynamics of drop probability as a function of aggregated offered load. The load on the network is increased by increasing number of hosts, with each host offering a fixed load of 0.1Mbps. Number of hosts is increased from 3 to 201 in increments of 1. One host acts as a sink, and the rest as packet sources. As a consequence the aggregated offered load grows from 0.1Mbps to 20Mbps. The experiment is done with buffer size 1000. This is done in order to ensure that buffer size is not contributing to drop rate of packets. Each experiment is run for 10ms.

Figure C1 demonstrates results for frames of sizes 64, 512 and 1500 bytes. It is evident from the figure that the drop probability stays near zero until offered load grows to 6Mbps for 64 byte frames, and until 8Mbps for 512 bytes and 1500 bytes. After that the drop probability exhibits logarithmic growth. The graphs for 512 and 1500 byte frames stay nearly identical, and substantially below graph for 64 byte frames.

The point of growth for 512 and 1500 byte frame graphs is very near the maximum capacity of the network, 10Mbps. When the network saturates the drop probability rises with the offered load as collisions become inevitable. Earlier breaking point in 64 byte frames graph demonstrates the fact that small frame size results in many more frames being sent. As a consequence buffer can overflow and

collisions when dequeuing from the buffer become much more frequent.



Experiment D: Drop Probability vs. Buffer Size

In this experiment we measure dynamics of drop probability as a function of buffer size. The tests are done for 512 byte frames. The network load is kept constant at 10Mbps. Network consists of two source hosts offering 5Mbps load each, and one sink host. We chose to use few hosts that offer large load, as opposed to many hosts that offer small load for two reasons. First, this allows us to control total size of buffers with finer granularity. Increasing buffer size by 1 for two hosts increases total size of buffers by $1*2=2$, while increasing buffer size by 1 for 200 hosts increases total size of buffers by 200. Second, fewer hosts result in fewer packet collisions which in turn results in more regular dequeue of packets from the buffer. Each experiment is run for 10ms.

Figure D1 demonstrates that increasing buffer size results in decreased drop probability. Intuitively, the larger is the buffer, the higher bursts in offered load it can accommodate without dropping packets. As soon as buffer's limit is reached, probability of packet drop increases dramatically.

Conclusion

The results presented in this report demonstrate most important characteristics of Ethernet networks (throughput, latency, drop probability) and their dependency from the load on the network. Additionally we were able to prove unfairness of the Ethernet during short periods of time.

Figure D1: Drop Probability vs Buffer Size

