The objective of the lab was to simulate an Ethernet LAN, to confirm the answer to a question from the midterm exam. The question was as follows:

1. a. Suppose you want to network a cluster of database servers with several client workstations using standard 10BaseT Ethernet, with 10BaseT hubs. If each of the workstations sends an average of one 1000 byte request to one of the servers each second, and the server responds to each request with a reply that averages 2750 bytes, estimate an upper limit on the number of workstations you could attach to the network before it starts to be overloaded by the traffic. Assume that the database servers can keep up with the load, and you have as many ports on the hubs as you want. (Hint: remember the difference between the theoretical throughput of a 10BaseT network and the practical throughput).

   b. Now suppose you get some extra money and can buy Ethernet switches instead of hubs. Would this change your estimated maximum number of workstations? Explain why or why not. If the number would change, would it go up or down?

The simulation was done using OpNet ITGuru. The first section describes the configuration of the network and the simulation parameters for each scenario. The second section presents the actual simulation results and analyzes them. The third section describes additional scenarios that were simulated to explore the behavior of a related network using coax Ethernet and an Ethernet switch. under different changes to the data center server and network configurations, including an analysis of the simulation results. The final section presents conclusions.

Configuration

The first simulation scenario, **SingleHub**, consisted of a single 10BaseT Ethernet LAN built using a single *ethernet128_hub*, *10BaseT* links, and attached *ethernet_station* workstations. For this first scenario, no server was present, but 100 workstations were added using the *Topology … Rapid Configuration* tool, selecting a star topology. Each of the workstations was configured identically to generate traffic using the following parameters: each was assigned an *On State Time* (time in which it was allowed to send packets) of **exponential(100)**, an *Off State Time* (dead time between packet transmissions) of **exponential(0)**, and a *start time* of **constant(5)**. The *packet size* was set to **constant(1250)**. The *interarrival time* (time distribution for generation of packets to send) was configured to be **exponential(0.3333)**. The latter two numbers were arrived at after some experimentation; an initial attempt was made to set the interarrival time to **constant(1)** and the packet size to **constant(3750)**, and the start time of **uniform(5, 6)**. The intent was to make each workstation send 3750 bytes once every second, but to stagger the start time of each workstation. However, the use of a packet size greater than 1500 bytes caused no bits to be received by any station on the network (the *Traffic received (bits/sec)* was always zero), so I split the data up into three separate packets and sent them every one-third second. It may have been possible to enable fragmentation of the packets instead; I did not experiment with that.

The diagram of this network is not very interesting, so I did capture the topology. For the scenario, the global statistics for *Ethernet Delay (sec.)*, *Traffic Sent (bits/sec)*, and *Traffic Received (bits/sec)* were collected.

Running the Simulation

The simulation was executed for fifteen seconds. The global attribute *Eth Hub Optimization* was enabled to reduce the simulation run time. After the simulation was completed, a graph was generated comparing the traffic sent vs. traffic received for the network; Fig. 1 shows the resulting graph.

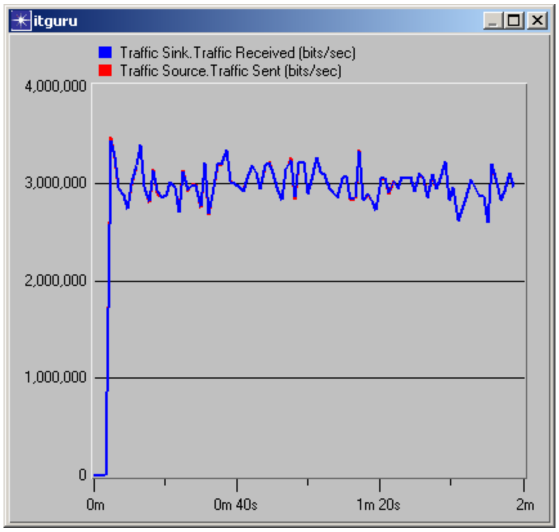

Figure 1 - Traffic Sent and Received, bits/sec

As expected, the network throughput (traffic received) quickly reached a steady state at around 3 Mbps (100 stations * 3 packets/sec * 1250 bytes/packet * 8 bits/byte = 3 Mbps). To confirm, a second scenario was created with identical parameters, but with only 75 workstations. Fig. 2 shows the corresponding graph.

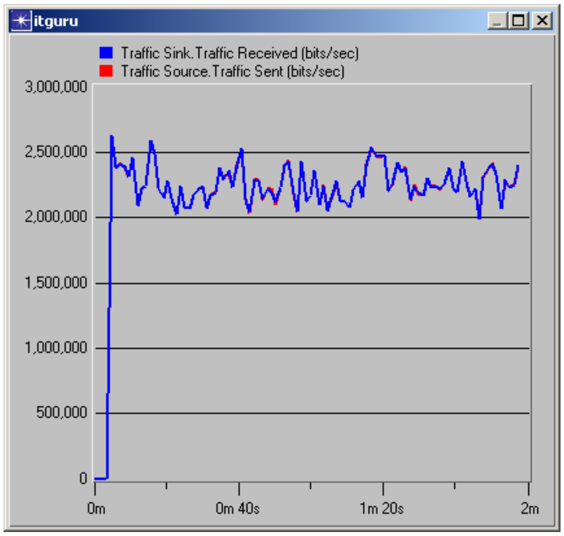

Figure 2 - Traffic Sent and Received for 75 Workstations

The network throughput was as expected, 75 * 3 * 1250 * 8 = 2.25 Mbps. A third scenario was added with 125 workstations. Fig. 3 shows the result.

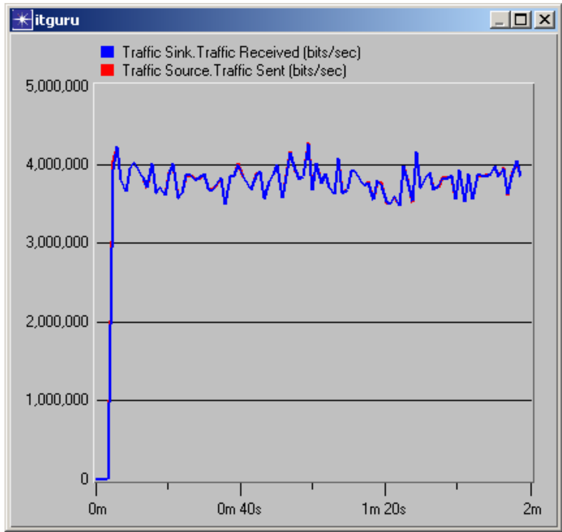

Figure 3 - Traffic Sent and Received, 125 Workstations

Traffic continued to increase beyond the expected 3 Mbps cap, reaching the expected theoretical maximum value of 3.75 Mbps. In order to investigate whether the network would handle more load, instead of increasing the number of workstations, the interarrival time was decreased to 0.1667 sec, to double the offered traffic. Fig. 4 shows the results.

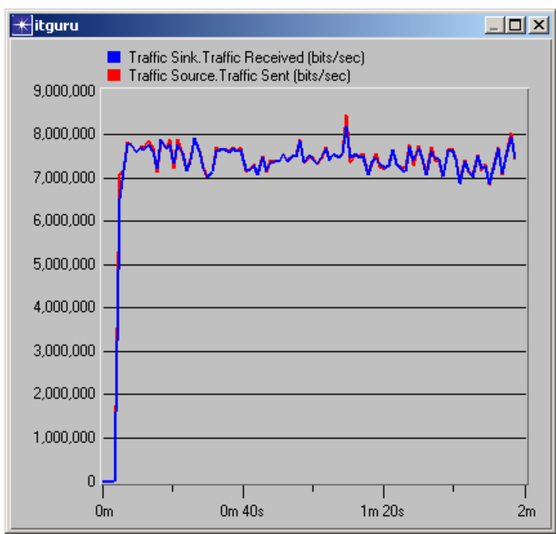

Figure 4 - 125 Workstations, Double Offered Traffic

This result was completely unexpected; doubling the offered traffic created double the throughput, rather than peaking at a significantly lower rate.

Coax Ethernet Scenario

In order to check whether the OpNet model for the hub was behaving significantly differently than expected, I simulated a coax Ethernet with the same number of attached workstations and same traffic pattern. The network used the *ethcoax* model, *ethcoax_station* workstations, *eth_coax* links, *eth_tap* taps. The simulation took considerably longer, so it was only run for 15 sec. Fig. 5 shows the result.

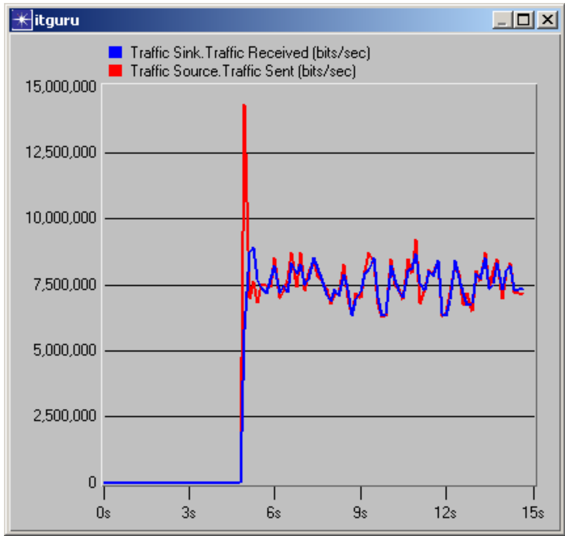

Figure 5 - Coax Ethernet Throughput, 125 Workstations

Again, this was definitely not the expected result. Nothing has been found to explain why the network throughput continues to increase beyond the expected practical maximum of around 3 Mbps.

Switched Ethernet Scenario

After the unexplainable results obtained using the 10BaseT hub, I went on to replace it with a 10BaseT switch (*ethernet128_switch*), in the same network with 125 workstations and the interarrival time of 0.1667 sec. The global setting *Switch Sim Efficiency* was enabled. Fig. 6 shows the network throughput.

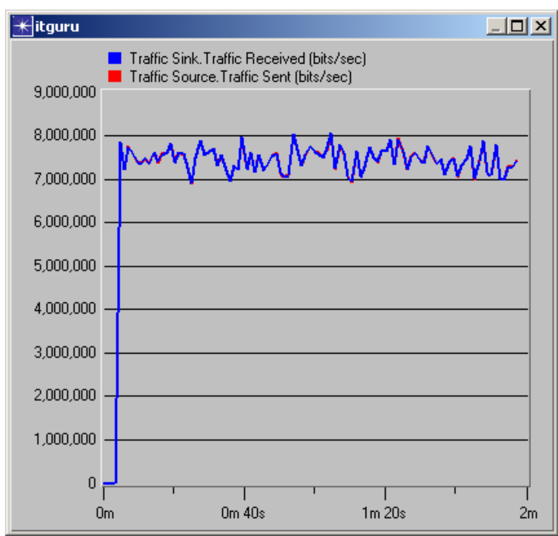

Figure 6 - Switched Network Throughput

To see how much additional traffic the network would support, the interarrival time was halved (to 0.08333 sec); Fig. 7 shows the resulting throughput.

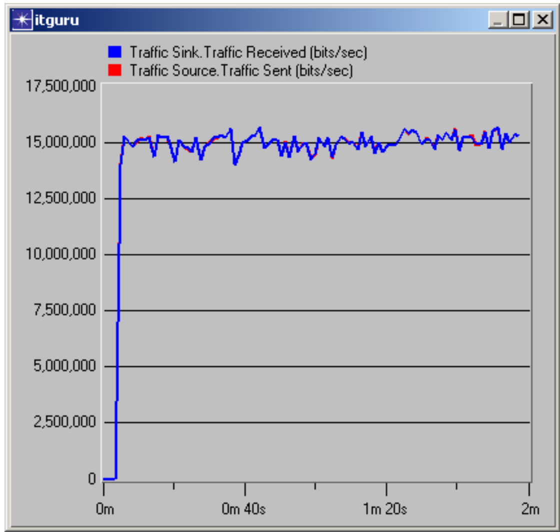

Figure 7 - Throughput with Double Offered Traffic

The throughput of the network doubled with the doubling of offered traffic, and exceeded the 10Mbps limit of 10BaseT Ethernet. This is explained by the fact that the switch is buffering packets and is capable of delivering packets to multiple unicast destinations simultaneously. An additional simulation with the inter-arrival time decreased to 0.005 sec was run; it took considerably longer to simulate, so the total time simulated was decreased to 15 sec. The resulting throughput is shown in Fig. 8.

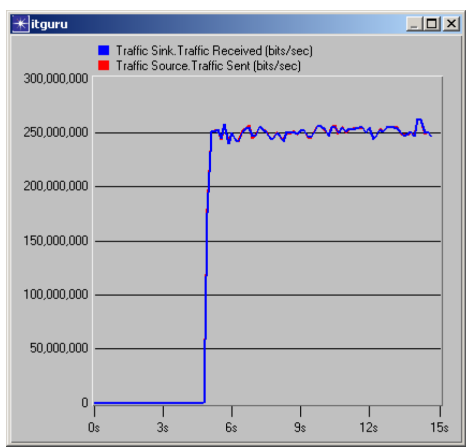

Figure 8 - Throughput with 5 msec Interarrival Time

The network continues to handle all offered traffic (125 stations * 200 packets/sec * 1250 bytes/packet * 8 bits/byte = 250 Mbps).

Conclusions

I cannot explain why the networks using the 10BaseT hub and the coax Ethernet did not reach a peak value significantly less that 5 Mbps, as expected. These results were completely unexpected and cast some doubt on the effectiveness of the OpNet simulation. The fact that the switched network was able to handle much more load was anticipated, but it is not clear whether these results are reasonable either.