Effects of Filler Traffic In IP Networks

Adam Feldman

April 5, 2001

Master's Project

Abstract

On the Internet, there is a well-documented requirement that much more bandwidth be available than is used on average. This extra bandwidth is required to handle "network spikes", those times when network traffic peaks, using much more bandwidth than is normally needed. Sometimes these spikes can be quite large, requiring a network be provisioned to provide two to twenty times the bandwidth than it uses on average in order for the network to maintain its adequate speed of operation and prevent degradation for its users. Thus, for a network to achieve acceptable performance, a good portion of its bandwidth must sit idle most of the time. Due to the oft high prices of bandwidth, this is either not feasible or leads to quite a waste of money. If this extra bandwidth could be used for other activities while it is not needed to handle a spike of traffic, then the network would be much more cost effective. The goal of this project is to determine a method of allowing the extra bandwidth to fill its idle time by carrying low priority traffic. When the network spike occurs, normal traffic would take priority over this "filler" traffic, allowing the network to use this bandwidth as if it had been idle. Since the filler traffic is composed of non-time sensitive data, the interruption the spike causes is expected and not disruptive. This way the network is equipped to handle all spikes of traffic without slowdown, while at the same time, the excess bandwidth is being used to accomplish work.

The problem in executing this idea is deciding the nature and specifics of the filler traffic in such a way that prevents it from interfering in any way with the normal, or pre-existing, traffic, yet ensures that it will eventually reach its destination. To see the effects different types of traffic have on the pre-existing traffic, the results of a network simulation, using traffic data from a Harvard trace and completed with NS-2, were studied. The simulation was executed many times for the same input data, to see the effects of changing different parameters, such as bandwidth and latency of the central link of the simulated network. The results of this experiment are packet dynamic, breaking down into three items, which were charted over the experiments, for a given changing parameter: average packet delay, percent of dropped packets, and the amount of bandwidth used. Each of these result categories were plotted separately based on which type of traffic (pre-existing and filler) and which direction, to show how the filler traffic in one direction affects the pre-existing traffic in both directions.

The resulting output charts indicate that the idea of implementing filler traffic is feasible. They also confirm that the parameters can have a noticeable impact on the effect the filler traffic has on the pre-existing traffic. If the central link has too low of a bandwidth, then the pre-existing traffic is slowed. On the other hand, the usefulness of the filler traffic is also dependent on the parameters used. For example, if the latency is very large, then the delay of the filler packets will increase.

Introduction

With today's heavy usage of the Internet, bandwidth is at a premium. Despite the high cost of bandwidth, it is necessary for a provider to make available much more bandwidth than is regularly used. This is due to the burstiness of the traffic. Although most of the time a given network will only use a certain amount of bandwidth, occasionally extra bandwidth will be needed to account for temporary increases in traffic (bursts), otherwise, a burst can cause a slowdown across the network [1]. In order to prevent this potential slowdown, a network must have much more bandwidth than it normally uses, as much as 20 times what is normally used [2], as shown in Figure 1.

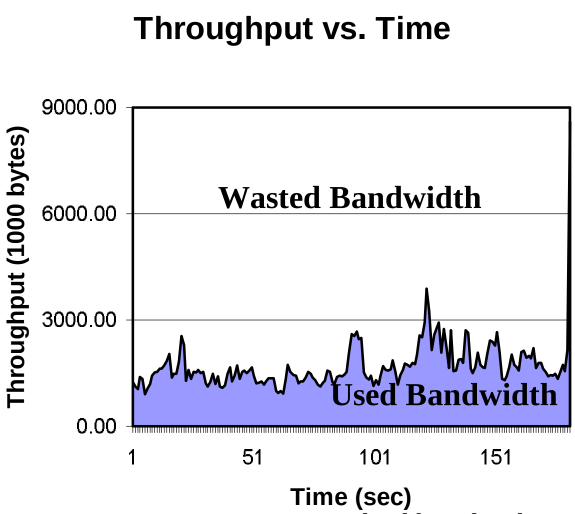

Figure 2. Extra Bandwidth Utilized

Bandwidth rental accounts for about 55% of an ISP operating expenses [3]. However, because these bursts only occur occasionally, this extra bandwidth is wasted most of the time. If only 5% to 50% of its bandwidth is typically used [4], then the ISP is throwing away a significant portion of its budget. However, if the extra bandwidth could be used, carrying "filler traffic", it would no longer all be wasted. ISPs would be able to sell this excess bandwidth for specific uses during the times it is not needed to handle a burst of traffic. Thus the bandwidth would be used to save money, making high levels of service cost effective.

In order to successfully implement filler traffic, there must be a way of giving the pre-existing traffic a higher priority than the filler traffic. Otherwise, the transmission of filler traffic will not be halted when the bandwidth is needed by pre-existing traffic. This can be done using strict prioritization, which works by assigning specific priorities to each packet. In this method, it is easy to determine which packets should be sent immediately, and which are forced wait until there are no other packets waiting. Without this prioritization method, all packets would have the same priority, and would be sent in the same order they are received. In practice, there is a buffer where packets “wait” until they can be transmitted. Normally, packets are placed in the buffer and sent using a first in, first out scheme. However, by using strict prioritization, packets with high priority are able to put in the front of the buffer, ahead of the packets with a lower priority. This is useful when certain packets need to be sent with preferential treatment. Streaming multimedia and other real-time applications have a need to be sent with bounded delay and packet loss. Strict prioritization would allow the pairing of these packets of high importance with packets which do not require as high of a level of service. In this way, more bandwidth can be utilized on a regular basis.

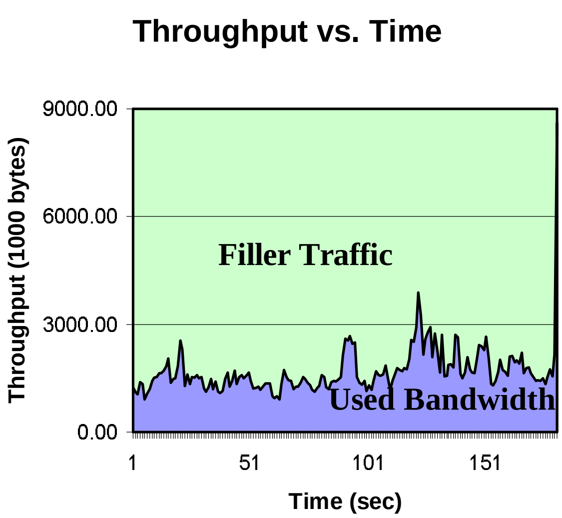

Figure 2. Extra Bandwidth Utilized

In practice, this filler traffic system would work as follows. During normal network operation, filler traffic uses excess bandwidth to transfer low priority data. When the amount of pre-existing network traffic spikes (or bursts) to high levels, the filler traffic transmission is temporarily halted by the router. This allows the pre-existing traffic to be transmitted unhindered. Upon the continuation of normal traffic levels, filler traffic can resume transmitting. This would allow bandwidth which was once wasted to be utilized, as shown in Figure 2. The problem at hand is to determine how filler traffic affects the transmission of pre-existing traffic.

Application

Because the excess bandwidth is needed at specific, unpredictable times, it must remain available to carry normal network traffic. Prioritization allows pre-existing traffic access to the extra bandwidth during those periodic times when it is needed to prevent a network slowdown. Additionally, filler traffic should not be time sensitive, as pre-existing traffic can greatly delay its delivery.

In spite of these requirements, there are many possible uses of this extra bandwidth. For example, two that are currently being researched are FastStart [5] and Prepushing [6]. Additionally, distributed computing applications, such as SETI@Home, and continuous sensor data acquisition, as in weather station readouts and Global Positioning Systems, can take advantage of acting as filler traffic to collate data. Finally, further utilizations include data backup and cataloging, such as web spidering, across the network.

Important Factors to Study

In order for filler traffic to be considered successful in fulfilling its objectives, there are two factors which must be considered. Filler traffic must be unobtrusive while maintaining a reasonable performance. Thus, it is important to study not only how it performs, but also its unobtrusiveness from the point of view of the pre-existing traffic.

Filler Traffic Unobtrusiveness

This is a measure of how the filler traffic affects the pre-existing traffic. Ideally, the pre-existing traffic (or, more specifically, the users who are transmitting the pre-existing traffic) should not even be able to tell that the filler traffic exists. As per the problem specifications, the filler traffic should only be using the bandwidth which is sitting idle – and when that bandwidth needs to be used by pre-existing traffic, filler traffic transmission is instantly put on hold. However, this exactness is not possible. For example, packet-switched networks, such as IP networks, do not allow for packet transmission to be halted mid-packet. Therefore, it is certain that the filler traffic will have some affect on the pre-existing traffic. The task at hand is to determine how great that affect is. In order to do this, the variables which must be examined are average utilized bandwidth, average delay, and percent of dropped packets, all of the pre-existing traffic. If the utilized bandwidth goes down, or the delay or dropped packets goes up, then the filler traffic is negatively affecting the pre-existing traffic, and it must be determined how great this affect is, and how to minimize it.

Filler Traffic Performance

The performance of the filler traffic is a study of how much useful work is being completed by the filler traffic – i.e. how much data is successfully being sent. It is great if the filler traffic is totally unobtrusive, remaining invisible to the pre-existing traffic, however, if this results in few or no filler packets being successfully sent, then its usefulness is very limited. To measure the performance of filler traffic, the same three variables as above must be considered – but with respect to the filler traffic. Simultaneously, the filler traffic delay and percent of dropped packets must be minimized, while its utilized bandwidth is maximized.

Experiment Information

Network Simulation

All experiments to study the effects of filler traffic on an IP network were conducted using Network Simulator, version 2 (NS-2). [7] NS-2 uses data gathered from a real network containing information about the source, destination, size, type, and time of each packet. TCPDump was used to gather the packet information between Harvard and the Internet in 30 minute segments. Three of these traces made up the pre-existing traffic for these experiments. They were gathered on March 13th, 1997 starting at 8:39, 12:39, and 16:39, respectively [8]. The traces represent all of the packets traveling between Harvard's local network and the rest of the Internet. This traffic flowed across Harvard's 10Mbps Ethernet link to the Internet, and was divided into two flows, incoming and outgoing.

Type of Filler Traffic

The type of filler traffic studied in these experiments is transmitted using FTP. Thus, the filler traffic was transmitted as an infinite file using TCP congestion control. Selective Acknowledgement TCP (SACK) was used – it was determined to be a good choice because it is a commonly used, modern version, being supported on almost 40% of clients as of March 2000, and its usage is growing [9]. Sack allows the receiver to return information to the sender about which packets were received, even if they were received out of order. Additionally, delayed acknowledgement was implemented, so that acks could be delayed and aggregated at the receiving side. After a bit of testing, the advertised window was set to the very large value of 20,000, to simulate infinity, so that it would not be a limiting factor in the experiments.

Network Topology

The network studied had a dumbbell topology, as shown in Figure 3. Nodes 4 and 5 were routers which managed the packet flow across the central link. Nodes 2 and 3 were connected to the routers, one each, via high speed links, and represented the computers in Harvard's network and the rest of the Internet, respectively. These were the nodes that the pre-existing traffic traveled between. Finally, node 0 represented the filler source, while node 1 represented the filler destination. These two nodes were also connected to the routers with high speed links. Because the end nodes are connected to the middle nodes using high speed links (much higher speed than the central link), ensuring that they are not the bottleneck of the network, only the information relating to packets traveling across the central link will be studied by these experiments.

Additionally, a control experiment was conducted with no filler traffic to get a basis of comparison for the rest of the experiments. The results of this experiment indicated that the network was set up correctly, and that the analysis software functioned accurately, as nothing surprising was found. It also acted as a comparison baseline for the rest of the experiments. Finally, experiments were executed which switched the filler source with the filler destination, thus reversing the flow of the filler traffic. Due to the symmetric nature of the network, this resulted in unremarkable changes to the results.

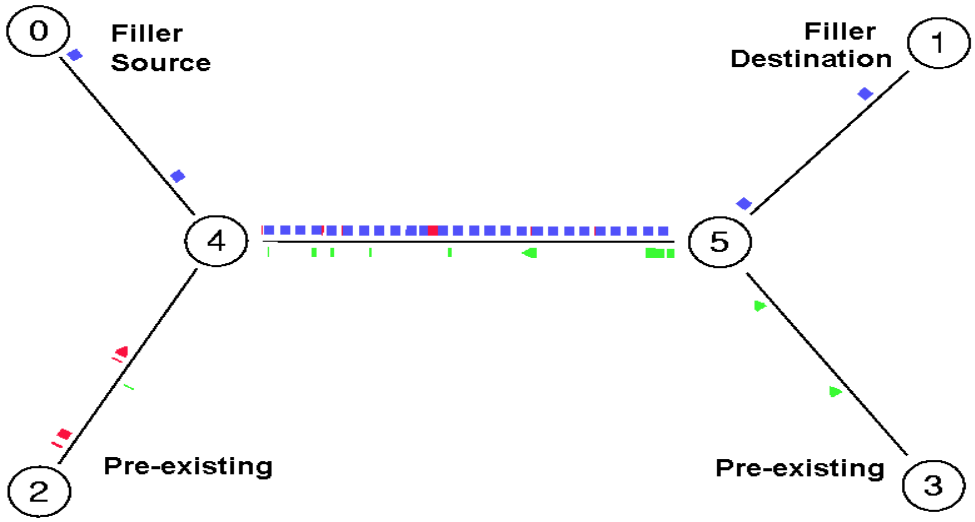

Figure 3: Dumbbell Network Topology

Central Link Parameters

Each experiment was designed to study how varying a specific parameter changed the effect that filler traffic had on the pre-existing traffic. Initially, a base value for each parameter was selected, and a single run of NS-2 was executed using these base values. Then, one by one, each parameter was changed to one of several alternate values while the rest of the parameters remained set to their base values. The parameters (and their values):

Bandwidth – This is the amount of bandwidth on the central link. The central link is designed to be a duplex link, meaning that the value of the bandwidth is the amount of available bandwidth in each direction. The base value for bandwidth is 10 Mbps, which was selected because it is the actual link of the Harvard traces. The varied values were initially 3, 4, 6, 8, and 15 Mbps.

Latency – The latency is the amount of time it would take for a zero size packet to travel across the central link. The base value of latency is 20ms. Experiments tested additional values of latency of 3ms, 10ms, 30ms, 40ms, and 250ms.

Filler Buffer – This is the size of the buffer which is used to store packets waiting to be sent across the central link. If this buffer is full, and another packet is ready to be sent, it will be dropped. Therefore it is important that the filler buffer is not too small. However, if the filler buffer is too large, the average delay will increase, because packets must "wait in line" for a very long time. Initially, the base value for filler buffer was 16KB, while the other tested values were 4KB, 8KB, and 32KB. Preliminary studies produced the predicted results of increasing delay and decreasing dropped packets as buffer size increases. However, having a filler buffer of 16KB caused a side effect while trying to study latency. As latency increased, utilized bandwidth went down because many packets were dropped as the filler buffer became full. Therefore, the final experiments were all conducted with a filler buffer equal to 1.3 times the product of the bandwidth and latency of the given experiment. This product is called the *bandwidth delay product* (BDP), and is a measure of the total number of bytes of traffic that can be on the link at any time [10]. This dynamic filler buffer, which increases as latency increases, can support any latency.

Packet Dynamics

Upon the completion of these experiments, statistical data relating to the packet dynamics was gathered from the ns-2 output. This data was gathered using software I wrote designed to focus on Average Delay, Percent of Packets Dropped, Average Throughput, and Ack Compression (the latter is not used). The Appendix contains documentation relating to the usage and output of this software. Each measure is calculated and studied by flow – the data about packets from 0 to 1 is separate from the data for packets from 1 to 0, etc.

Average Delay – This is the average amount of time it takes a packet to be sent, from the time it is dequeued from the buffer on one node of the central link, until it is received at the destination node at the other end of the central link. There are two types of packets, regular data packets and acknowledgement (ack) packets. Ack packets are packets which are sent from a destination to a source to let the source know that the destination has received data packets. The average delay for each of these is measured and calculated separately, then are averaged together. This is done to facilitate the calculation of ack compression (see below). Each time a packet of a given type (regular or ack) is encountered, its delay is added to a total delay variable, while a counter variable is incremented. At the end, the average delay is calculated by dividing the total delay by the number of packets encountered.

Percent of Packets Dropped – The percent of packets dropped is merely a measure of the number of packets dropped divided by the total number of packets, expressed as a percent. It indicates what percent of the total number of packets were enqueued but never dequeued or received because the filler buffer was full, preventing the packet from being saved.

Average Utilized Bandwidth – This is the average number of bytes per second that were sent along each flow. Normally throughput varies from utilized bandwidth because utilized bandwidth is a measure of the total number of packets sent, whereas throughput does not include duplicate packets which have been resent due to being dropped. However, in these experiments, utilized bandwidth is the same as throughput because the only dropped packets are dropped before they are sent across the central link. Utilized bandwidth is calculated by dividing the total number of received bytes (total bytes – dropped bytes) by the total time of the experiment (time last packet is received – time first packet is sent). This results in bytes/sec, which is converted to KB/sec for convenience.

The final two pieces of data gathered from the ns-2 output are distributions of data. First, is the distribution of the delays - for each 0.001 second increment, the number of packets which had a delay in that increment is recorded. Finally, the time series of the total number of bytes transmitted in every 0.1 second increment of the experiment (18,000 such increments for a 30 minute experiment) is recorded.

Results

Filler Traffic Unobtrusiveness

This section relates to how “invisible” the filler traffic was to the pre-existing traffic. The goal is for the filler traffic to have no effect on the pre-existing traffic. To determine the effect, if any, the packet dynamics of the pre-existing packets are examined. First, the average packet delay, the throughput, and the percent of dropped packets of the pre-existing traffic will be studied as bandwidth changes, then as latency changes. Since the usefulness of the filler traffic is not related to its unobtrusiveness, this section only deals with the statistics for pre-existing traffic. Additionally, normal filler packets only flowed in one direction. The ack packets flowing in the other direction were observed to have virtually no affect on the pre-existing traffic, and were thus not included. For ease of readability and comparison, all of the following charts are patterned the same. “Pre-existing” or “Filler” in the charts’ titles indicates which flow is shown on the chart, while “With Filler” and “Without Filler” in the legends represents experiments which were conducted with and without filler. On all charts, the solid line(s) pertain to trace 1, while dashed is for trace 2 and dotted for trace 3. Finally, all experiments which included filler traffic are represented by dark colored lines, while the corresponding filler-free experiments are a lighter shade of the similar color.

Average Delay vs. Bandwidth - Figure 4 shows the average delay of the pre-existing packets for each trace as bandwidth varies, both in experiments with filler traffic and with no filler traffic. The experiments with no filler traffic are included as a comparison, to help determine the exact effect that filler traffic has on the pre-existing traffic. Consider trace 2, as it has the largest initial gap between filler and non-filler experiments, albeit slightly. At the lowest bandwidth of 3 Mbps, which is only a little more than double the average throughput of the pre-existing traffic, there is only approximately a 4% increase in average delay. The difference drops steadily, until it is less than 2% with a bandwidth of 10 Mbps. In general, as bandwidth increases, the gap between filler and non-filler experiments decreases.

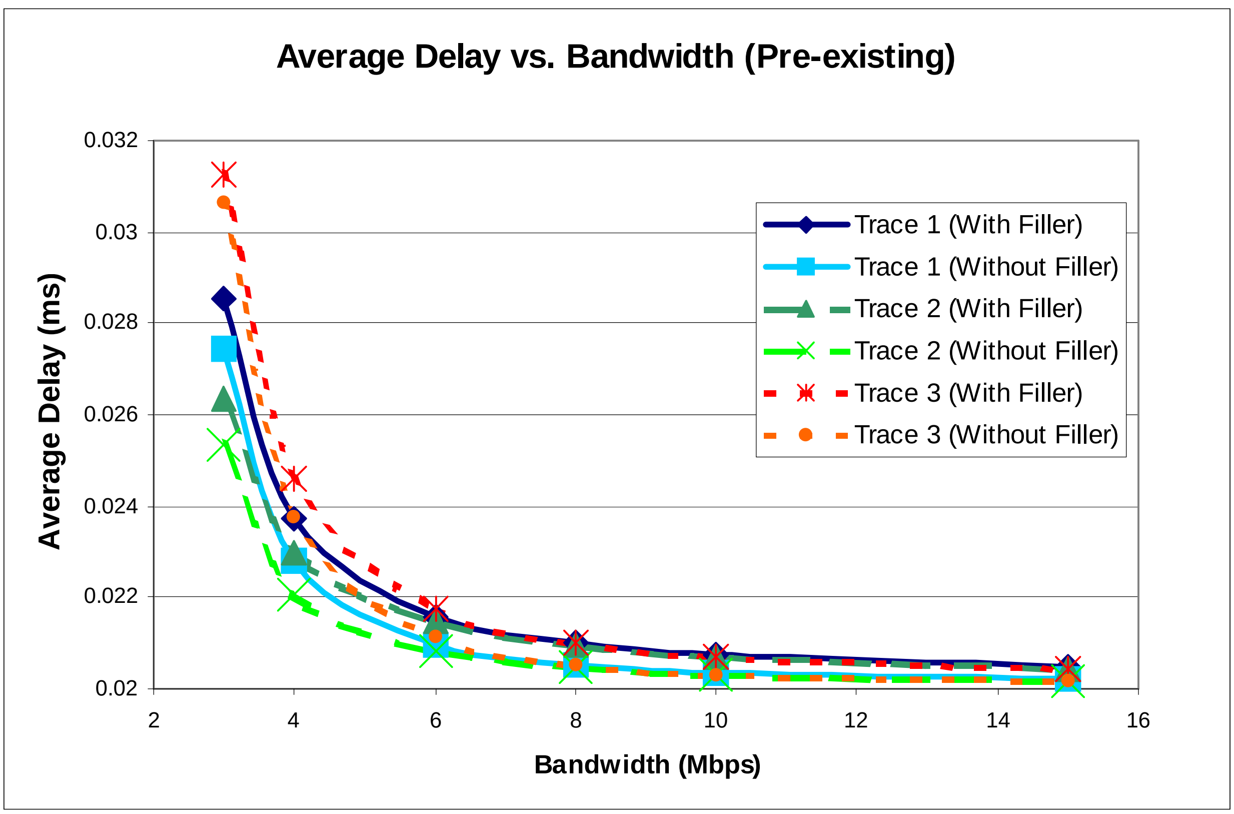

Figure 4. Average Delay vs. Bandwidth (Pre-existing)

Dropped Packets vs. Bandwidth - As seen in Figure 5, there are no dropped packets once bandwidth increases to between 4 Mbps and 6 Mbps. However, until that point, there are dropped packets in both filler and non-filler experiments. The data points that have the biggest difference between the filler and non-filler experiments are for trace 1, at 3 Mbps. It should be noted that this is an almost starved network. Here there is a 3% difference, meaning that the filler traffic causes only 3% more packets to be dropped in the pre-existing traffic. This increase is very small, amounting to 3 out of 25,000 packets, while without filler traffic, 1 packet in 250 are already being dropped.

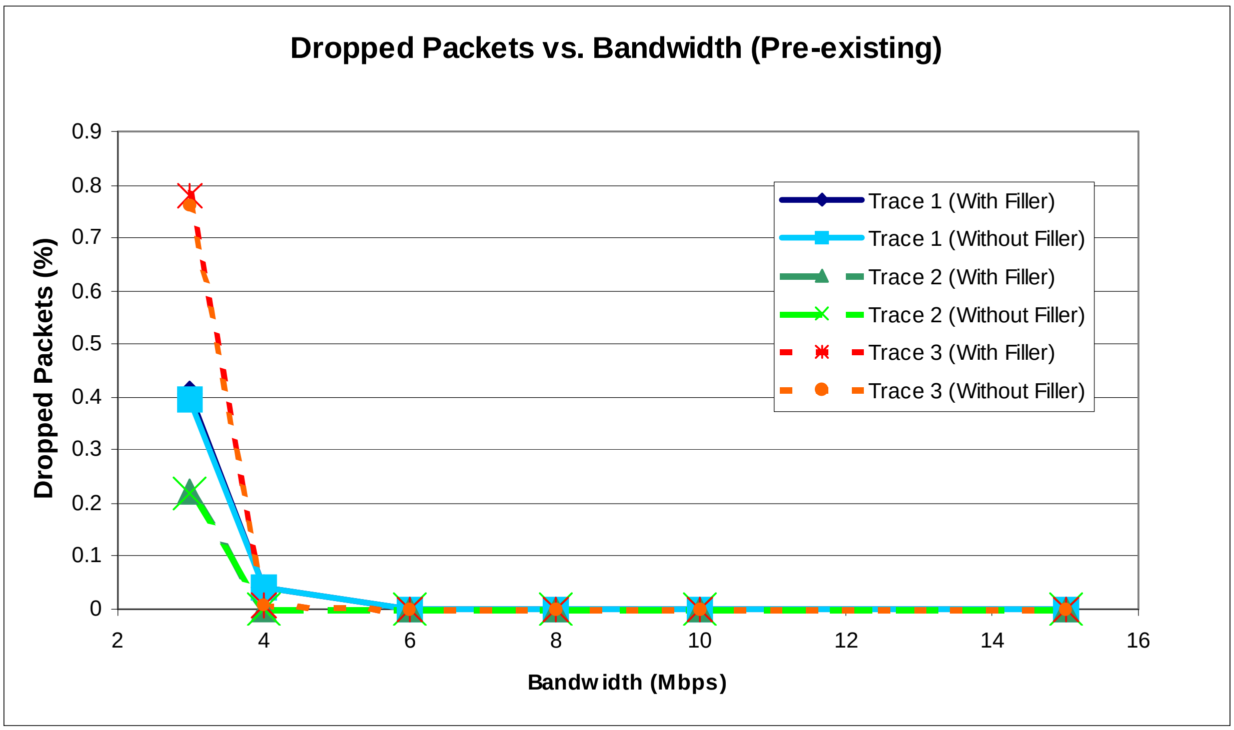

Figure 5. Dropped Packets vs. Bandwidth (Pre-existing)

Throughput vs. Bandwidth - There were no experiments conducted which involved a central link bandwidth of less than the average throughput of the pre-existing data. Because of this, the throughput of the pre-existing data, which is at a higher priority than the filler traffic, should not decrease due to the presence of filler traffic. Close examination of Figure 6 shows that this is what occurred. There was no reduction in throughput due to the existence of filler traffic.

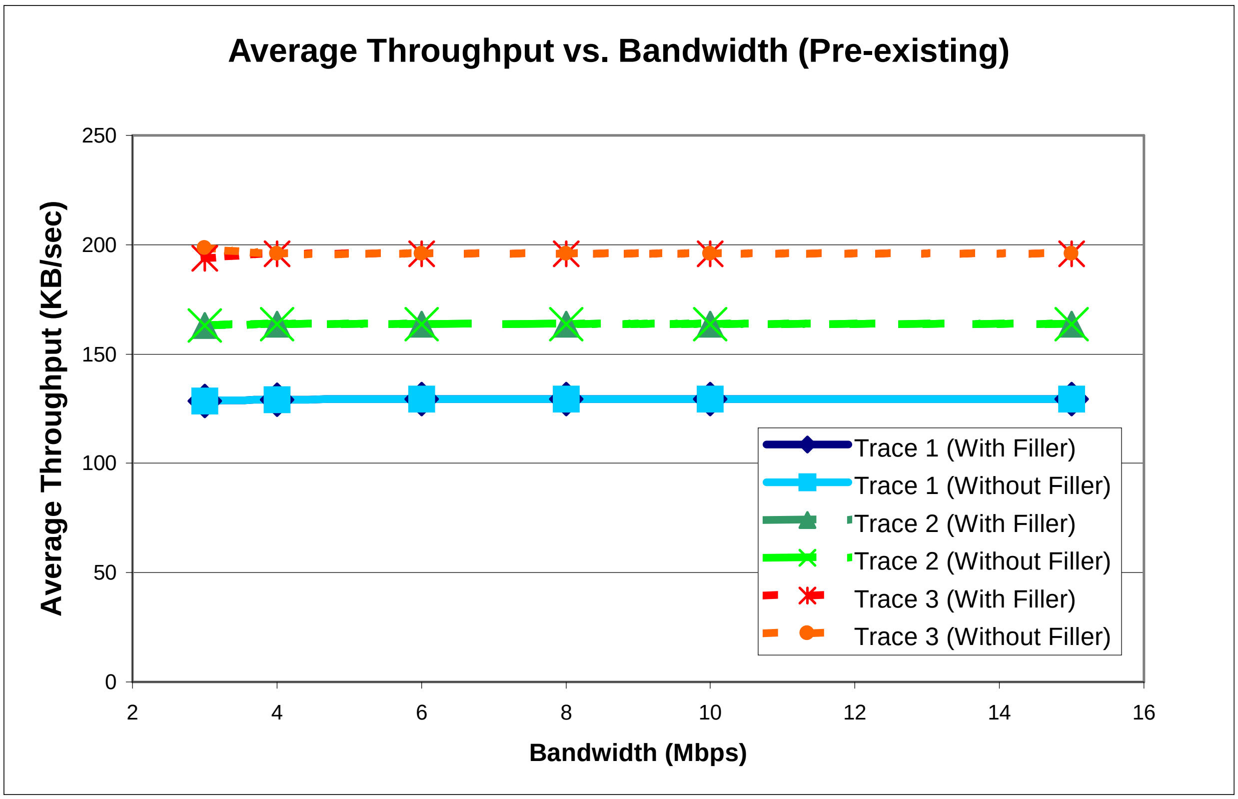

Figure 6. Average Throughput vs. Bandwidth (Pre-existing)

Average Delay vs. Latency - Since latency is a measure of how long it takes a packet to cross the central link, it is to be expected that the average delay increase linearly with time. Again, the experiments with no filler traffic are included in Figure 5 to demonstrate that the average delay of pre-existing packets barely increases due to filler traffic. For example, in trace 1 (of Figure 7), the average delay for packets with filler present, after link latency is subtracted out, is 0.3491ms, while the same value when filler traffic is present is in the range of 0.745ms. This appears to be a large increase (more than double), but when the link latency is figured in, this increase in delay is very small, less than 1% for a latency as low as 40ms (see Appendix for information related to low BDP networks). Again, filler traffic has only a very small affect on pre-existing traffic, regardless of latency.

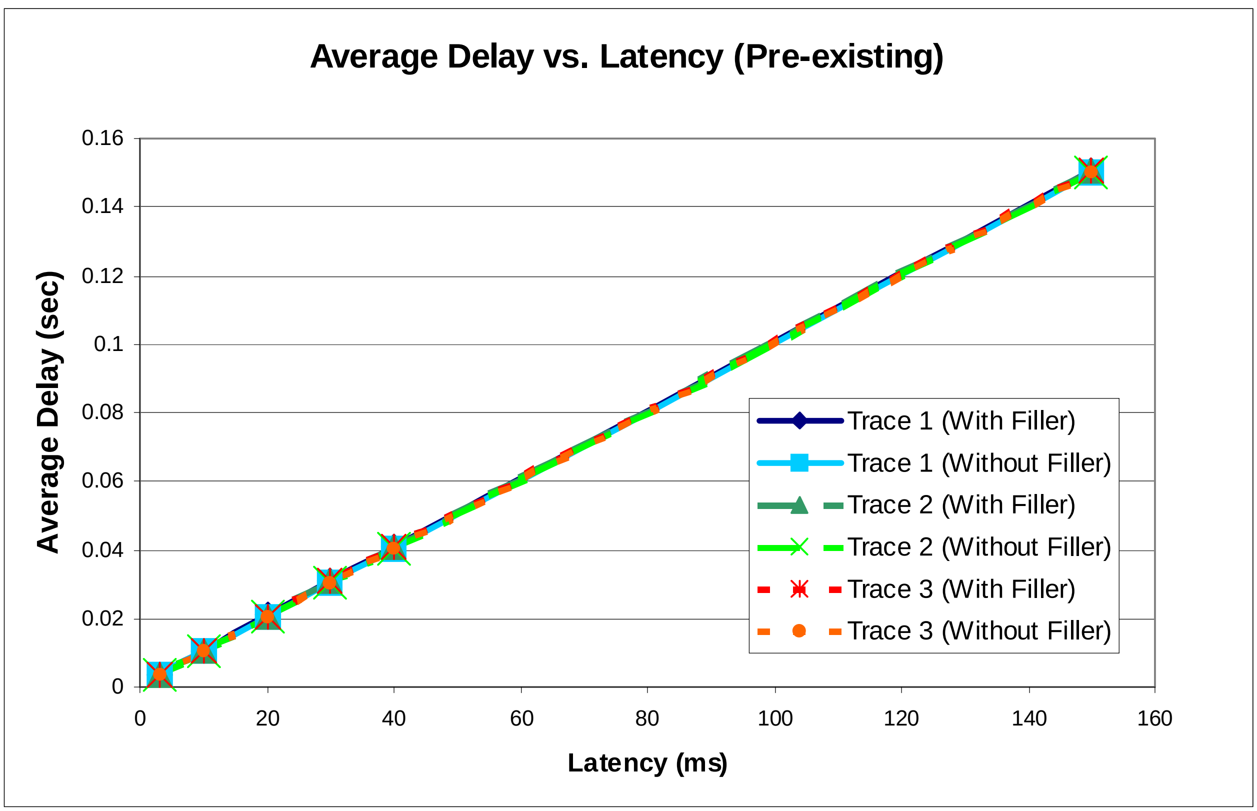

Figure 7. Average Delay vs. Latency (Pre-existing)