Wayne(Wenguang) Wang wew036@cs.usask.ca

Assignment 2: Self-Similar Traffic and Its Impact on Network Performance

1. (1 mark) Plot a graph of the traffic profile with a time interval granularity of I = 0.1 seconds.

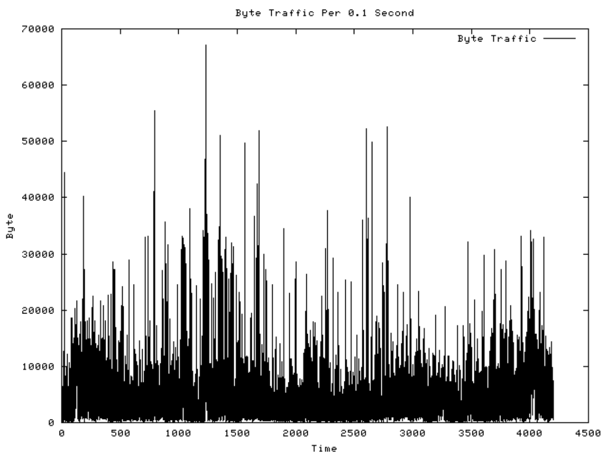

Time unit = 0.1 second

Figure 1 Byte traffic on time interval 0.1 second

2. (1 mark) Plot the autocorrelation function for the bytesper0.1sec.dat data set.

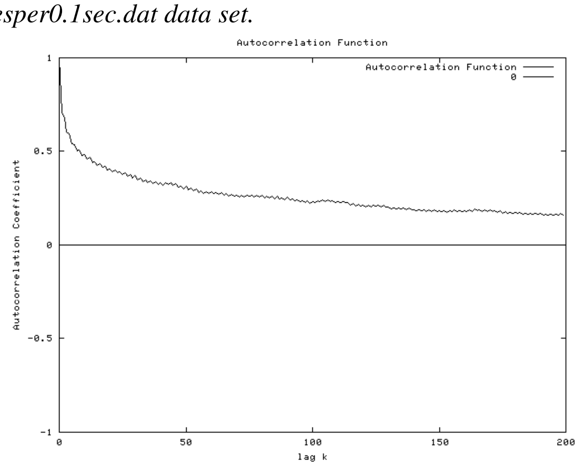

Figure 2 Autocorrelation Function

3. (1 mark) Plot a variance-time plot for the data set.

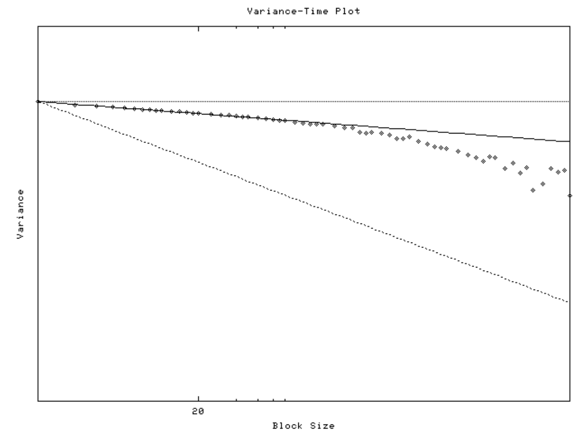

Figure 3 Variance-time plot

4. (1 mark) Plot an R/S pox diagram for the bytesper0.1sec.dat data set.

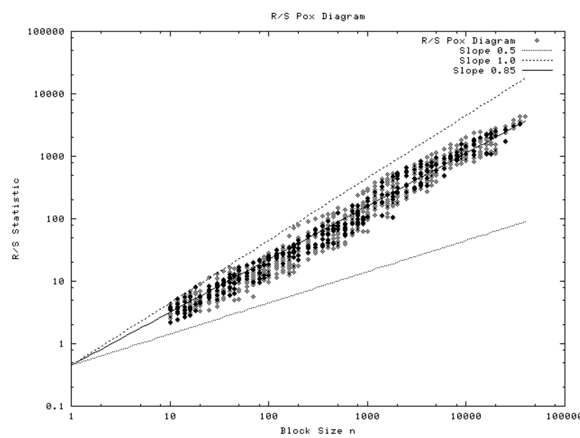

Figure 4 R/S Pox Diagram

5. (1 mark) Compute estimates of the Hurst parameter for the bytesper0.1sec.dat data set. Do this first with a crude estimate using the variance-time plot. Show this estimate on your variance-time graph. Then do this using a crude estimate from the R/S pox plot. Show this estimate on your R/S graph. Then estimate H using a periodogram technique.

In Figure 3 (variance-time plot), when the block size is big, the number of intervals is small, and the variance is not accurate. Therefore, the line which has slope β=0.2 only matches the former part of the graph. From Figure 3, H=1-β/2=**0.9**.

In Figure 4 (R/S Pox Diagram), the line with slope 0.85 matches the points well. Therefore, H is **0.85**.

The H is **0.9131** when using the periodogram technique.

6. (4 marks) Plot a graph that shows the cell loss ratio on the vertical axis and the link capacity C on the horizontal axis. Draw one line on the graph for each buffer size B used in the experiment, and clearly label which line is which.

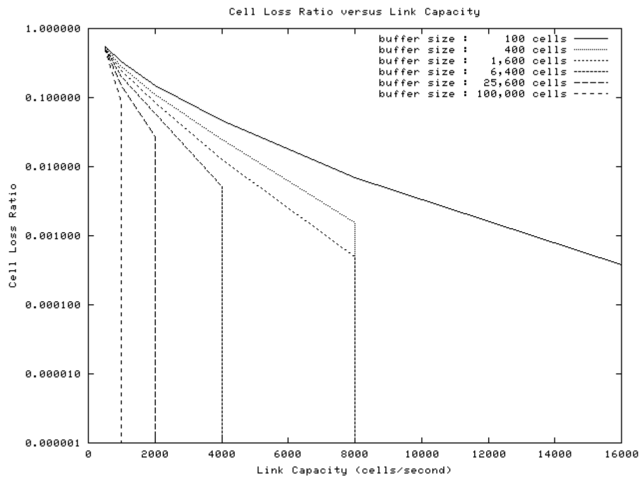

Figure 5 CLR on different link capacities and buffer sizes

7. (1 mark) a table of the CLR (cell loss ratio) results for different B and C.

Buffer Size (cells)	Link Capacity (cells/second)					
	500	1,000	2,000	4,000	8,000	16,000
100	0.548086	0.329600	0.151210	0.046813	0.006931	0.000383
400	0.526063	0.282095	0.112008	0.024717	0.001544	0
1,600	0.514457	0.242593	0.085219	0.012836	0.000495	0
6,400	0.504820	0.204064	0.059473	0.005056	0	0
25,600	0.499160	0.150254	0.027609	0	0	0
100,000	0.481605	0.087757	0	0	0	0

Table 1: the CLR (cell loss ratio) results for different B and C

8. (1 mark) a table of the PLR (packet loss ratio) results for different B and C.

Buffer Size (cells)	Link Capacity (cells/second)					
	500	1,000	2,000	4,000	8,000	16,000
100	0.298310	0.162153	0.069060	0.023677	0.004186	0.000414
400	0.283762	0.134092	0.046928	0.011789	0.000911	0
1,600	0.276906	0.112412	0.033701	0.005977	0.000292	0
6,400	0.271395	0.093849	0.023126	0.002388	0	0
25,600	0.268279	0.069125	0.010570	0	0	0
100,000	0.258203	0.039986	0	0	0	0

Table 2: the PLR (packet loss ratio) results for different B and C

9. (4 marks) a one-page (no more) discussion of the main results, observations, and conclusions from your simulation experiment.

The experiment is to determine the impact of buffer size (B) and output link capacity (C) on the cell loss ratio and the packet loss ratio. The selection of B and C is based on the characteristic of the TCP/IP traffic. The average traffic input is 44783 bytes/second, which is about 1000 cells per second. Therefore, link capacities that are smaller than, equal to, and bigger than this value are selected respectively. Furthermore, most of the B and C are selected in an exponential gap. From the experiment, I got the following observations:

1. Bigger buffer and faster link result in less data loss.
It is easy to observe that as the values of B or C become bigger, both the CLR and PLR become smaller. For the same B, the CLR and PLR decrease when C increases. Similarly, for the same C, CLR and PLR decrease when B increases. Unless B is smaller than the maximum packet size (1500 bytes or 32 cells) or C is smaller than the average input rate (1000 cells/s), the CLR and PLR can reach zero when B and/or C are big enough.

2. The PLR is smaller than the CLR.
If all lost cells are nearly evenly distributed across all packets, the PLR should be close to 100%. However, the experiment shows that the PLR is always smaller than the CLR when there are cell losses. This is because under the bursty input traffic, most lost cells come from a few packets during the burst time.

3. The link capacity C that is smaller than or equal to the average input traffic rate is unacceptable.
When C is less than the average input traffic rate (in Figure 5, when C is 500), the switch can never reach a stable state. When it is unstable, no matter how big the buffer is, the CLR and PLR are still quite high. Furthermore, when C is equal to the average input traffic rate (in Figure 5, it is 1000), if the traffic were CBR, the CLR and PLR should be 0. However, the input TCP/IP traffic is VBR with self-similarity, thus is very bursty. Therefore, even with a 100,000 cells huge buffer, which is impracticably big in an ATM switch, the CLR and PLR are still very high. Moreover, bigger buffers cause longer delays.
Therefore, the link capacities that are smaller than or equal to the average input traffic rate are both unacceptable.

4. When the link capacity increases linearly, both CLR and PLR decrease exponentially.
In Figure 5, all lines are very close to straight lines. Because the Y-axis is in log scale, this indicates that when increasing the link capacity linearly, the CLR decreases exponentially. The same trend also applies to the PLR (the graph is not shown in this report).

5. The relationship between the buffer size and the CLR/PLR.

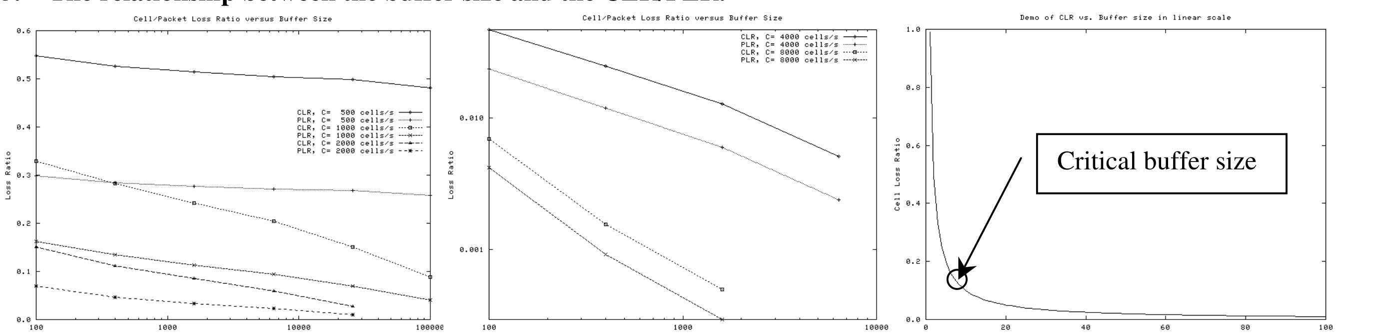

(a) Small Link Capacity (b) Large Link Capacity (c) Demonstration of the critical buffer size

Figure 6 Loss Ratio vs. Buffer Size

Figure 6(a) shows the relationship between the buffer size and the CLR/PLR when the link capacity C is relatively small (500–2000 cells/s). Figure 6(b) depicts the relation when C is relatively large (4000–8000 cells/s). All lines in Figure 6 (a) and (b) are very close to straight lines. However, in Figure 6(a), the Y-axis is in linear scale. This means that increasing buffer size exponentially can only decrease CLR and PLR linearly. This is because at that time, the bottleneck is the link capacity. Therefore, increasing buffer size is nearly useless. In Figure 6(b), the Y-axis is in logarithmic scale. Figure 6(b) shows that increasing buffer size exponentially can decrease CLR and PLR exponentially. This indicates that when the link capacity is large, the bottleneck moves to the buffer. Moreover, this relationship implies that there exists a critical buffer size B (see Figure 6(c)). A buffer which is smaller than B can cause very high cell loss, and a buffer which is bigger than B can only get a little benefit.

6. The self-similar input traffic cost m resources (buffer and link capacity).
Because the input traffic is self-similar, there are big traffic bursts across many time scales. Moreover, when the time scale is smaller, the burst rate is even higher. For example, the average rate of the input traffic in this experiment is about 1000 cells/s. In Figure 1 (the time interval is 0.1 second), the peak traffic rate is 67098 bytes per 0.1 second, or about 14000 cells/s, which is 14 times higher than the average. I also used 0.01 second time interval and got a 24000 cells/s peak traffic rate, which is 24 times higher than the average (not shown in the report). Figure 5 shows that when capacity C is equal to 8000 cells/s (where is 8 times higher than the average input traffic rate) and B equals to 1600 cells, the CLR is still not low enough (0.0005). This indicates that self-similar traffic needs more resources to achieve a satisfactory performance than traditional Poisson traffic.