>>> Solution for HW #2 for Capacity Planning (Fall 2001) <<<

The given RTT data for Meredith College (Internet1 path) and NCSU RTT (Internet2 path) from USF was used in this characterization [1]. Table 1 shows the characterization results for mean, variance, standard deviation, and CoV for the four data sets. Summary statistics were generated using `summary1.c`, from the tools page [2]. It can be seen that the Internet2 path has lower mean delay, lower variability, and less difference between "am" and "pm" that does the Internet1 path.

Table 1 – Summary statistics for Meredith and NCSU RTT data

Data set	Mean	Variance	Standard Deviation	CoV
ncsu_am.txt	34.17 ms	3.74 ms^2	1.93 ms	0.057
ncsu_pm.txt	34.47	19.63	4.43	0.129
meredith_am.txt	42.18	36.44	6.04	0.143
meredith_pm.txt	60.16	217.41	14.75	0.245

Histograms are shown in Figures 1 and 2 (histograms were created using `hist.c` from the tools page [2]). Figure 1 shows Meredith RTT histogram for morning (am series) and afternoon (pm series). Figure 2 shows the NCSU RTT histogram. The histograms show the same general trend as the summary statistics. The Internet1 path shows a large spread in RTT during the afternoon busy period. The Internet2, however, has very close to the same RTT between morning and afternoon, with only a very slight afternoon spread.

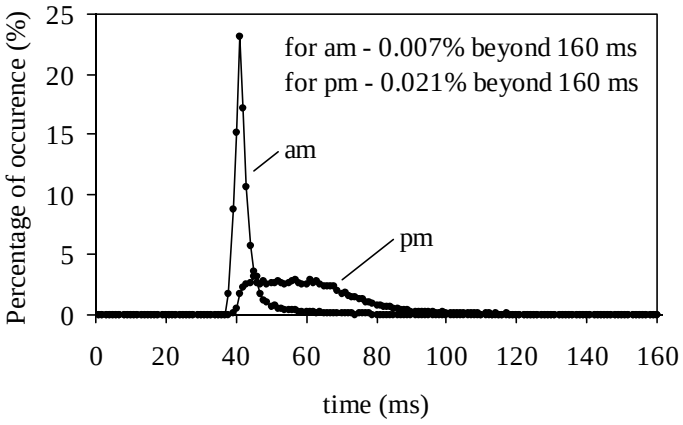

Figure 1 – Meredith RTT histogram

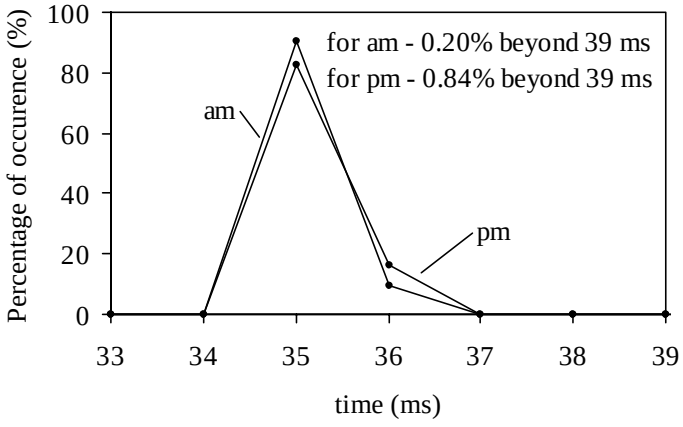

Figure 2 – NCSU RTT histogram

Autocorrelations are shown in Figures 3 and 4 (autocorrelations were created using `autoc.c` from the tools page [2]). Internet1 path RTT autocorrelation is high (close to 1) for small lags and decays as the lag increases. Internet2 path RTT autocorrelation is very low (less than 0.25 in all cases), but shows spikes every 5 lags. These spikes do not appear to be decreasing in magnitude as the lag increases.

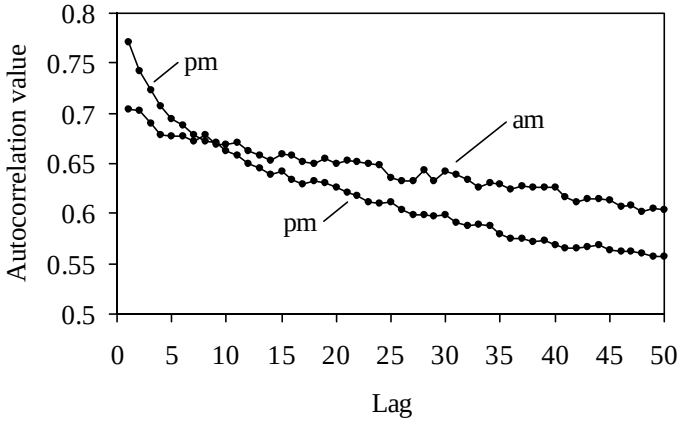

Figure 3 – Meredith autocorrelation

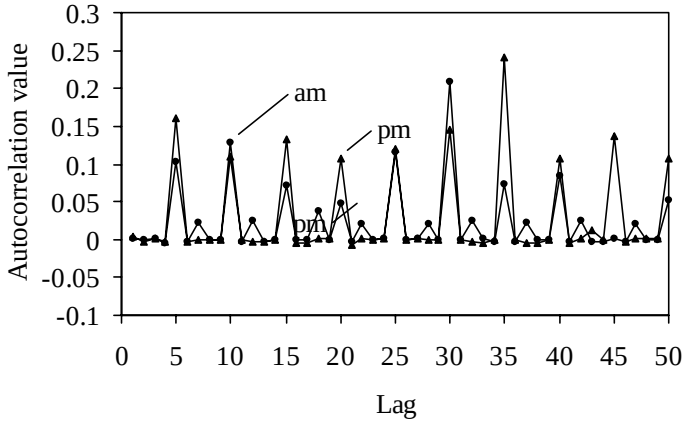

Figure 4 – NCSU autocorrelation

Some insights…

The Internet1 path had higher RTT that the Internet2 path for both am and pm. In addition, the Internet1 path exhibited overall greater variability and greater difference between am and pm. A possible cause for this may be that the Internet1 path is closer to its "knee" than is the Internet2 path. This would be the case if Internet1 utilization (i.e., offered load) were greater than that of the Internet2. This is a plausible explanation since the connectivity of the Internet is limited primarily to universities (i.e., there are less web sites to "surf to") and it is of greater bandwidth than the Internet1. This speculation could be confirmed by making utilization measurements of the two paths. An experiment that could be performed is to send increasingly higher rates of "ping" packets on the two paths and measure 1) the loss of pings and 2) the increase in RTT as a function of ping offered load. This experiment could show how close to the knee are the two Internet paths. The very different autocorrelation behaviors of the Internet1 and Internet2 are not so easy to explain. It is not clear what may be causing the periodic (?) autocorrelation spikes seen in the Internet2 measurements (Figure 4).

References

[1] Assignment #2 for Capacity Planning, Fall 2001. URL: *http://www.csee.usf.edu/~christen/class9/as9_2.txt*.

[2] Tools page for Kenneth J. Christensen, August 17, 2001. URL: *http://www.csee.usf.edu/~christen/tools/toolpage.html*.
