ThriftNet: A Reliable Networking Strategy

Donald F. Miller
Department of Computer Science
University of Mississippi
University, Mississippi 38677

Introduction

ThriftNet is a system of programs which employs a relatively simple protocol which effects the transfer of files between computers (Ferguson, et al., 1980). Since the notion of a simple protocol runs contrary to the prevailing thought on the design of computer networks (Maginnis, 1982), it is necessary to statistically establish the reliability of the protocol in accomplishing its objective of transferring files. Each time a file transfer process is completed using the ThriftNet system, statistical information is logged on the target operating system, consisting of the date and time, the user's name, the network node identification numbers for both the master and the target computer systems, the target system terminal line number and its baud rate, the number of byte count errors, the number of longitudinal redundancy check character errors, the number of successfully transferred 128 byte blocks, the total duration of the file transfer, and the effective data transmission rate between systems (i.e. the actual number of data bytes transferred per second).

Data Collection

During the period, June, 1981, to January, 1982, 4305 file transfers were logged into ThriftNet statistical files on five target operating systems: 1) a DECsystem-10 running the TOPS-10 operating system located on the campus of the University of Mississippi, 2) a PDP-11/34 running the Unix operating system also located on the University of Mississippi campus, 3) a DECsystem-10 running the TOPS-10 operating system located on the Western Michigan University campus, 4) a DECsystem-10 running the TOPS-10 operating system located at the University of Rochester School of Medicine and Dentistry, and 5) a PDP-11/44 running the RSX-11M operating system located on the campus of the Medical College of Wisconsin. The data collected at each of the target nodes has been combined so that the statistical analysis will focus on ThriftNet protocol robustness rather than operating system deficiencies. Even with this combination, three of the five target operating systems are TOPS-10 which may therefore influence some of the results. (We are currently working to expand the number and the identity of ThriftNet target operating systems to further minimize operating system dependent effects.)

ThriftNet Topology

Figure one shows a frequency matrix of file transfers between ThriftNet nodes. The five target nodes are listed across the top of the matrix while the eleven master nodes are indicated on the side. In addition to a node name, the node identification number is listed; in some cases a given node may be both a master and a target node. Each cell of the matrix contains the total frequency of file transfers between the two nodes, the percentage of file transfers relative to other target nodes, the percentage of file transfers relative to other master nodes, and finally, the percentage of the absolute total number of file transfers. As can be seen from the column totals at the bottom of the matrix, two systems account for 96.4% of all file transfers. The UM OCIS TOPS system on the University of Mississippi campus participates in 47.8% of all file transfers with this number being distributed over nine other ThriftNet nodes. 48.6% of all file transfers occurred on the UM CSCI UNIX system, however 99.7% of that 48.6% came from one other ThriftNet system. In spite of the combination of data from all fourteen

	TARGETID					
Count / Row % / Col % / Total %	UM OCIS TOPS 0.	UM CSCI UNIX 6.	WMU C.C. TOPS 15.	MCW PHYS IO RSX-1 16.	UR MCCF TOPS 20.	Row Total
MASTERID						
1. UM CSCI RT-11	101	2084	14	0	13	2212
	4.6	94.2	0.6	0.0	0.6	51.4
	4.9	99.7	18.4	0.0	33.3	
	2.3	48.4	0.3	0.0	0.3	
2. UM RIPS RTS	196	0	0	0	0	196
	100.0	0.0	0.0	0.0	0.0	4.6
	9.5	0.0	0.0	0.0	0.0	
	4.6	0.0	0.0	0.0	0.0	
3. UM RIPS OS8	126	0	1	0	0	127
	99.2	0.0	0.8	0.0	0.0	3.0
	6.1	0.0	1.3	0.0	0.0	
	2.9	0.0	0.0	0.0	0.0	
4. UM PHCL ADSS	106	0	0	0	0	106
	100.0	0.0	0.0	0.0	0.0	2.5
	5.2	0.0	0.0	0.0	0.0	
	2.5	0.0	0.0	0.0	0.0	
6. UM CSCI UNIX	650	0	1	0	0	651
	99.8	0.0	0.2	0.0	0.0	15.1
	31.6	0.0	1.3	0.0	0.0	
	15.1	0.0	0.0	0.0	0.0	
7. UM BYTECH	685	0	46	43	16	790
	86.7	0.0	5.8	5.4	2.0	18.4
	33.3	0.0	60.5	100.0	41.0	
	15.9	0.0	1.1	1.0	0.4	
8. UM PHYS RT-11	34	0	0	0	0	34
	100.0	0.0	0.0	0.0	0.0	0.8
	1.7	0.0	0.0	0.0	0.0	
	0.8	0.0	0.0	0.0	0.0	
9. WMU PSYC RT-11	0	0	10	0	10	20
	0.0	0.0	50.0	0.0	50.0	0.5
	0.0	0.0	13.2	0.0	25.6	
	0.0	0.0	0.2	0.0	0.2	
10. WMU PSYC OS8	0	0	4	0	0	4
	0.0	0.0	100.0	0.0	0.0	0.1
	0.0	0.0	5.3	0.0	0.0	
	0.0	0.0	0.1	0.0	0.0	
11. UM CHEM RSX-11M	146	7	0	0	0	153
	95.4	4.6	0.0	0.0	0.0	3.6
	7.1	0.3	0.0	0.0	0.0	
	3.4	0.2	0.0	0.0	0.0	
16. MCW PHYSIO RSX-1	12	0	0	0	0	12
	100.0	0.0	0.0	0.0	0.0	0.3
	0.6	0.0	0.0	0.0	0.0	
	0.3	0.0	0.0	0.0	0.0	
Column Total	2056	2091	76	43	39	4305
	47.8	48.6	1.8	1.0	0.9	100.0

Figure One: Frequency Matrix of File Transfers Between ThriftNet Nodes

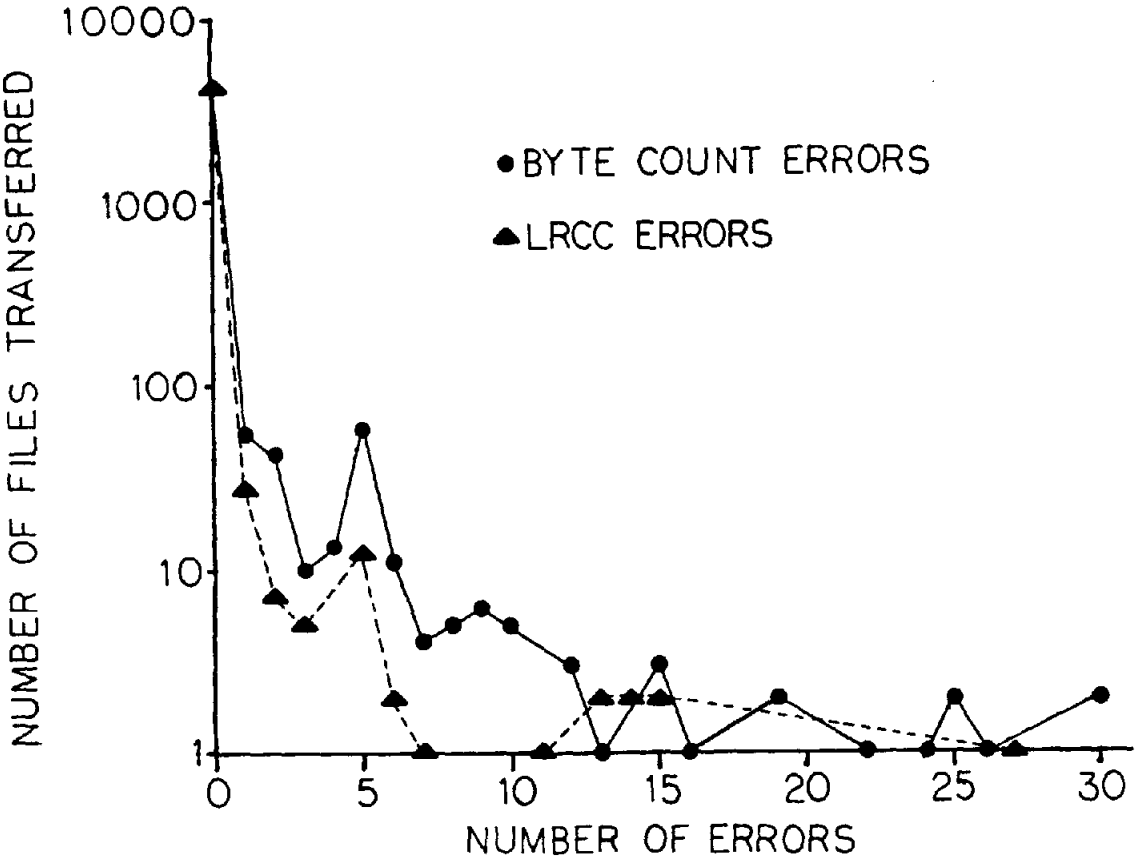

Figure Two: Semilogarithmic Plot of File Transfers vs. Byte Count and LRCC Errors

ThriftNet nodes, we have a disproportionately high percentage of transfers to or from two of the fourteen nodes. While this disproportion may tend to confound some of the data presented in the remainder of this paper, we still believe that the data will reveal a reliable protocol.

File Transfer Profile

ThriftNet file transfers tended to be evenly distributed across days of the week. The day having the lowest total number of file transfers was Sunday (11.5%) while the day having the highest total number of file transfers was Thursday (18.5%). The total number of file transfers were performed by 128 individual users. The number of files received and transmitted by target systems was approximately equal, with 49.2% of the files being received on target systems and 50.8% of the files transmitted from target systems. This finding is somewhat surprising in that ThriftNet was initially envisioned as a mechanism by which users of smaller systems could send files to larger timesharing systems. However, our user community found that file preparation on larger systems and subsequent transfer to smaller systems for execution an equally desirable alternative. 91.2% of all files were transferred in ASCII format while 8.8% of the files were transferred in binary integer format. 6.9% of all file transfers were cancelled by the user once the transfer was underway. Duration of file transfer was highly skewed, with the most frequent duration being seven seconds, the median transfer time being 35.5 seconds, and the mean transfer time being approximately five minutes.

File Transfer Error Analysis

Figure two shows a semilogarithmic plot of file transfers having various numbers of byte count and LRCC errors. As can be seen, the vast majority of the 4305

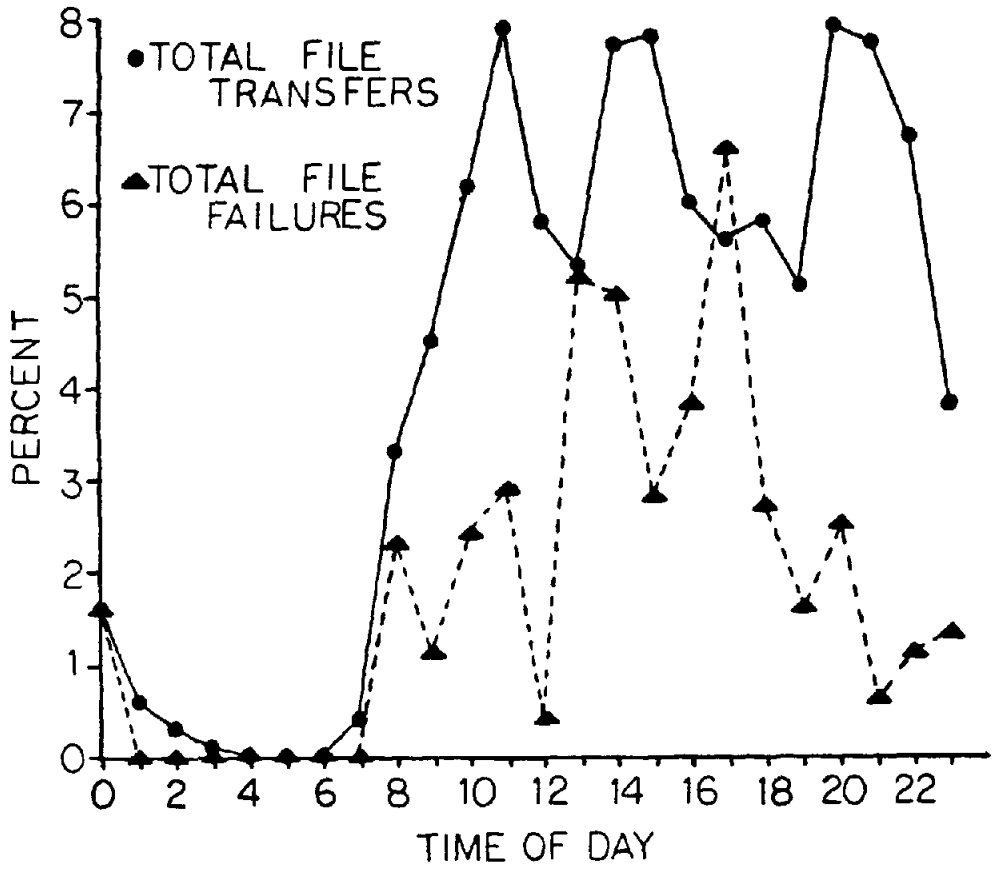

Figure Three: Percent File Transfers and Errors vs. Time of Day

file transfers had no errors. 94.8% had no byte count errors and 98.5% had no LRCC errors. For both byte count and LRCC errors, the curves are essentially negatively exponential with a small peak at five byte count or LRCC errors. This small peak is expected, since the protocol stops file transfers having five errors for a given ThriftNet block. Therefore, files having an error frequency greater than five had errors occurring "randomly" throughout the file transfer.

Figure three plots both percentage of total file transfers, and file transfer failure percentage over hours of the day. Percentage of total file transfers, symbolized by solid circles,can be seen to rise sharply beginning at 8 a.m. and stays elevated until 10 p.m. with valleys corresponding to the lunch and dinner hours. File transfer failure percentages, symbolized by solid triangles, also increase beginning at 8 a.m. and continue to increase until 5 p.m. except for valleys at noon and at 3 p.m. File transfer failure percentages sharply decrease after 5 p.m. and remain low throughout the evening. These data suggest that file transfer failure percentages are related to operating system load, and that a high number of ThriftNet file transfers in itself does not influence the success of the file transfer (i.e. does not significantly contribute to operating system load).

After user-aborted file transfers were excluded, 4010 file transfers remained. Of these, 4.49% had one or more byte count errors and 1.40% had one or more LRCC errors. The total file failure rate was three percent which consisted of 2.42% explained failures (i.e. failures due to five LRCC and/or byte count errors) and 0.58% unexplained failures (i.e. failures due to protocol deficiencies). 0.33% of the 217,633 attempted ThriftNet block transfers failed. Of the 4010 file transfers, 98 were able to recover from byte count and/or LRCC errors. However, 105 file transfers were unable to recover

from errors, thus creating an overall error recovery rate of 48.28%.

Summary

ThriftNet currently consists of fourteen active nodes, two of which are responsible for the bulk of the file transfers. The overall file transfer success rate was 97%. The overall error recovery rate from byte count and LRCC errors was 48.28%. Based on this information, we believe that we have established the basis for a reliable, yet simple networking system. Future ThriftNet research will concentrate on improving this reliability while not sacrificing network simplicity. To quote Thurber (1981), "Don't go to heroics to make lower levels [of your protocol] have an ultra reliable design: fault tolerance can be spread throughout the [system] design and you should seriously trade off the cost of losing an occasional (infrequent) datagram [block]."

References

Ferguson P.A., Miller D.F. and Maginnis P.T., "ThriftNet: An economical software package for interprocessor communications," Behavior Research Methods and Instrumentation, vol. 13, no. 2, pp. 251-254, 1981.

Maginnis P.T., "ThriftNet: A simple networking strategy," Proceedings of the ACM 20th Annual Southeast Regional Conference, In this issue.

Thurber K.J., "Conclusion," In: Distributed Systems-Architecture and Implementation, Eds: B.W. Lampson, M. Paul, and H.J. Siegert, Springer-Verlag, 1981.