Self-Similarity in Computer Networks

Juvyrose V. Santos
Kenwood High School

August 11,2000

Army Research Laboratory (ARL)
Mentors: Dr. John H. Brand II, Ann E. M. Brodeen

I. Abstract

This report documents the process of researching and outlining the foundation for an experiment on a military, low-bandwidth radio network. Background research was done on the existing literature on the self-similarity of civilian networks, focusing not only on the mechanisms and mathematics that prove the existence of self-similarity (the Hurst parameter) but also on the physical causes for that self-similarity (e.g. queue size, user "think time," protocols, etc.). Implications of the new information on network self-similarity imply a possibility of new processes in analyzing the data. Foundations were laid for a future experiment to investigate the effects of a self-similar model on a network consisting of SINCGARS military radios.

II. Introduction:

The purpose of the research was to lay the foundation for a future experiment. In 1994, ARL performed a number of radio network experiments. This included a BRL Memorandum report BRL-MR-3978 that investigated the effects of message length, message arrival rate, and frequency hopping on a small combat radio net. This resulted in the conclusion that the TACFIRE protocol could not accommodate the most severe of the data rates expected of the AFATDS. The traffic distribution was based on the Poisson model.

Recent work has found, however, that network traffic is not based on a standard Poisson distribution but is instead self-similar in nature, with long-range dependence. This throws an uncertainty over the original SINCGARS/TACFIRE experiment, as the testbed was essentially running a Poisson distribution and was a queued network. The seminal work that suggested a self-similar model might affect the possible network was Leland's (et al.) "On the Self-Similar nature of Ethernet Traffic."

This implies that certain conclusions in the original experiment may be affected with the adoption of a new model. Self-similarity, as shown in the original Bellcore traces, had affected the network management of other computer networks. There is no existing research on the effect of a self-similar traffic model on a low-bandwidth military radio network. Network management may be affected in such a network. Self-similar modeling provides more accurate results. The importance of the research lies not only in the field of network management but also in the purpose of the radios. SINCGARS is meant to be used in combat. Correct modeling and analysis must be used in order to provide proper data in order to avoid confusion on the front.

III. Background

Initial training for the research came in the form of a "crash course" in computers and computer networks. Among the resources consulted were Tanenbaum's Computer Networks [3], MCSE Certification books for Dummies, and practical experience on a Linux network. Also consulted were resources on statistics and SINCGARS radios. Following this process, I reviewed the Memorandum Report on the original AFATDS experiment (BRL-MR-3978) [1]. The summary of that report is as follows, taken from the technical note:

This experiment was designed to investigate the effect of message length, message arrival rate, and frequency hopping on the "throughput and delay of a small combat radio net." The combat net used the Tactical Fire Direction System (T ACFIRE) protocol, the Single Channel Ground and Airborne Radio System (SINCGARS), and Combat Net Radio (CNR). The experiment's purpose was to determine "the degree to which networks… satisfied the throughput and delay requirements of the Advanced Field Artillery Tactical Data System (AFATDS)."

The experiment was configured with four Sun workstation nodes, each containing a message driver. "Communications loading and a data collection program to log the sending and receipt of messages and acknowledgements as well as information of queues" were included in the message driver. These nodes were then connected to Magnavox Tactical Communications Modems (TCM) which, in turn, were connected to a SINCGARS. The SINCGARS were preloaded with "the F200 hopset." The radios were placed no more than three feet apart in order to minimize error, set to low power, and used resistor loads instead of antennas in order to avoid interference.

The following is verbatim from BRL Memorandum Report BRL-MR-3978:

> Four values within the range of the reported potential message rates were selected to emulate the rate of user generated messages as well as the user's system response to incoming messages. For this experiment, the arrival rate, λ, represented the total number of messages generated by all the nodes during the hour and queued for transmission on the net, *not the number of messages actually transmitted during the hour.* A scenario generator was written to create "messages" of character strings of a specified length and rate over a one hour period. In this baseline experiment message priorities were not considered. [1]

The TACFIRE protocol requires an acknowledgement of the receipt of the message from the remote receiver. If this is not received, the message is sent again. In this experiment, a limit of three retransmissions (after the original message transmission) was imposed for each message. If acknowledgement of the message had still not been received after four transmissions, the message was discarded.

The numbers of messages generated for transmission by the nodes (X_1, X_2, X_3, X_4) were "assumed to be mutually independent Poisson distributed variables, X_i, with parameter λ_i." It was stated that the "arrival rate of messages for transmission to a network is the sum of the message arrival rates of each node on the network." This same theory allowed the equal distribution of the "total expected loading among the four nodes during the experiment."

A network monitor contained a graphical display that illustrated the messages and acknowledgements transmitted between nodes, dynamically changing queue sizes, and network utilization for 15 second intervals during the experiment. "A message was assumed to enter network service when it reached the modem."

Two experiments were performed: one utilizing single channels, another for frequency hopping. Two factors in each of the experiments included the message arrival rate at 100, 250, 350, and 500 per node per one hour test cell and message lengths at 48, 144, 256, and 352 characters. Ultimately, there were 16 various test combinations. "These 16 test combinations were divided into blocks of size four, and the four blocks were run over a four day period." This was done because it was decided that the "shortest reasonable time to test anyone of the sixteen test combinations was one hour," and 16 hours could not be reasonably run in one day. The experiment was replicated three times.

Several statistical methods were employed in order to analyze the data garnered. These included the analysis of variance (ANOVA) and the Smirnov Nonparametric Test. The tests investigated the significance of difference among multiple sample means and tested the difference between single channel and frequency hopping data, respectively. It was found that both message length and arrival rate (how many messages were queued for each node) had "statistically significant effects on throughput, " with message length being the more significant of the two factors. There were no significant differences between the three replications of the experiment.

It was found that the increase in message length led to an increase in throughput and an increase in arrival rate, attributed to the increase in volume of data being sent across the network. The only meaningful change was found between the 400 and 1000 message arrival rates, after comparison of the change in mean throughput for the "different levels of message arrival rate for each level of message length." The small changes of mean throughput for arrival rates of 1000 and above indicate a lack of significance of a relationship between message lengths and arrival rate. "This is probably because with an arrival rate of 1000 and above, the queues were generally building, so network throughput reached its maximum." The network did not reach its maximum with arrival rate of 400 because the queues were generally empty. The increased messages reflected increased queue size, not throughput. The relationships outlined above can be seen on Figure 9, on page 18 of the report. It shows that an increase in message length led to greater queue delay. *"Loading the system with more messages only increased the queue sizes, not the throughput.* "

Further delays occurred and can be explained through the nature of the system.

> "In this system, when a node has a message to send, it senses the net to determine if it is idle. If it is, the node waits a fixed amount of time, called the Net Access Delay (NAD), and senses the net again. If the net is still idle, the message is transmitted. Different NADS are assigned to each node to minimize the possibility of message collisions. The NADS should be as small as possible to minimize the delay; but the differences between NADS should be large enough to prevent collisions. In the experiment, NADS of 1,2, 3, and 4 seconds were assigned to Nodes A, B, C, and D, respectively." [1]

TACFIRE protocol NADS are assigned according to priority, with a primary and secondary delay. The secondary NAD comes into effect after the original transmission has been sent in order to avoid monopoly of the net. "In conditions of heavy traffic, however, the two nodes with the smallest NADs will alternate transmitting, locking out the other nodes. The Magnovox TCMs did not implement two NADs per node. As a result, Node A monopolized the net when the traffic was heavy."

It was found that at 1000 messages per hour and beyond, queues are saturated. At 1000 msg/hour, "Nodes A and B again experienced only slight delays; however, the delays increased drastically as a function of message length at Nodes C and D." While Nodes A and B transmitted their queued messages, Node D had only sent 238 of its 250 messages at the end of the hour.

Further message and acknowledgement failures were also found. "The figures show that, even in the pristine environment of the experiment some tries and acknowledgements did fail." Graphical analysis of the data shows that short message tries failed more often than longer tries. "This indicates the errors observed were not the result of the bit error rate but were probably protocol related." A 1 second difference existed over all message length in round trip delay.

An excursion test was conducted in order to "eliminate confounding the effects of these possible causes of message failure." Failures resulting from erroneous detection of net idle times were removed. "One arrival rate, 1000 msg/hr, was selected and both FH and SC modes were tested, using the same four levels of message length in the full-scale experiment." The excursion test helped determine that the cause for the delay was the "erroneous detection of net idle times."

Conclusions garnered from the experimentation included the statement that

> "the results have shown even best case conditions to be worse than assumed in previous modeling efforts. The probability of messages failing was around 6% for SC transmissions. The effects of frequency hopping were substantial, doubling the number of fai1ed message transmissions. Average throughput never exceeded 648 bps. In FH mode, it never exceeded 566 bps. If the Hamming code is considered overhead, the maximum throughput dropped to 378 bps for SC and 330 bps for FH. Utilization never exceeded 81 %. Network delays were always more than twice the message transmission time… *This study leads us to conclude that the tested TACFIRE protocol at 1200 bps over a combat net radio channel cannot accommodate the most sever of the expected data rates of AFATDS."* [1]

This experiment can be considered the precursor of the current planned study. Once acquainted with the initial experiment, a literature search was embarked upon in order to lay the foundation for the significance of a self-similar distribution as opposed to a Poisson-distribution model.

Other facts were needed before I could reasonably parse the IEEE papers; namely, I had to know terms such as TCP/IP, Ethernet, OSI stack, UDP, etc. Therefore, some time was spent reading resources on networking. Important network information resources included Tanenbaum's *Computer Networks* [3]. I learned to acquire a more or less detailed knowledge of lnternet protocols, stacks, and bandwidth. It was necessary to know the differences between an Ethernet network, an ATM network, and one that depended on circuit switching (most dial-up connections).

To understand the hardware basis of the study, I was able to "dissect" an old Intel 8088 motherboard in order to gain a better understanding for the mechanics of computing. As a result, I gained knowledge on the construction and workings of a personal desktop computer, albeit an "antique."

Certain mathematical terms also had to be understood as well. This included the understanding of the Hurst parameter (and/or exponent), which was used to determine the degree of self-similarity of a given time series. That is, if the distribution had a Hurst parameter of between 0.5 and 1, it was self-similar. An elementary understanding of probability was also useful, as was the understanding of the Random Walk problem. A rudimentary understanding of other statistical processes were also acquired, as will be outlined later.

IV. Involvement

The crux of my involvement included the necessary literature search and review, as well as the learning of pertinent background facts on networking, protocols, and statistics. Once given the purpose of the study, I was able to review each resource garnered from the IEEE database in order to search for relevant information on the problem at hand. After sifting through numerous abstracts, I extracted the most relevant from the group and began to build a base on which the future experiment could rest.

I was able to gather enough resources to lead to suppositions on the self-similarity of a network, specifically, throughout all the layers in the OSI stack. This will be discussed later.

Further contribution to the project includes the co-authorship of an ARL Technical Note detailing the most relevant results of the literature search, summarizing the previous experiments on the subjects, and outlining several implications of the self- similarity within the different layers of the OSI (Open Systems Interface) stack.

Although the actual experiment was not performed during the duration of my apprenticeship, my involvement helped consolidate a body of knowledge on which the future experiment could be designed and analyzed. The Technical Note was not meant to be a conclusive paper-rather, it was meant to be a conceptual one. The research was also not meant to be conclusive but to provide a given background for the future experiments.

V. Conclusion and Future Work

The results of the literature searches formed the basis for the planned experiment. The literature described anew relationship, that of self -similarity, within the traced traffic of various networks. The visual "proof' of the self-similarity may be found in Leland's paper [2]. Self-similarity is then defined as the continued uniformity of data over widely differing time scales. In the Bellcore instance, it was a time scale that differed from 100 seconds to .01 seconds. Mathematical and statistical properties of the self-similarity within the data were also discussed, as were possible traffic models that were better suited to the traffic than the previously used Poisson and Markov Modulated models.

Fig. 1 The Bellcore self-similar trace and the traditional model [2]

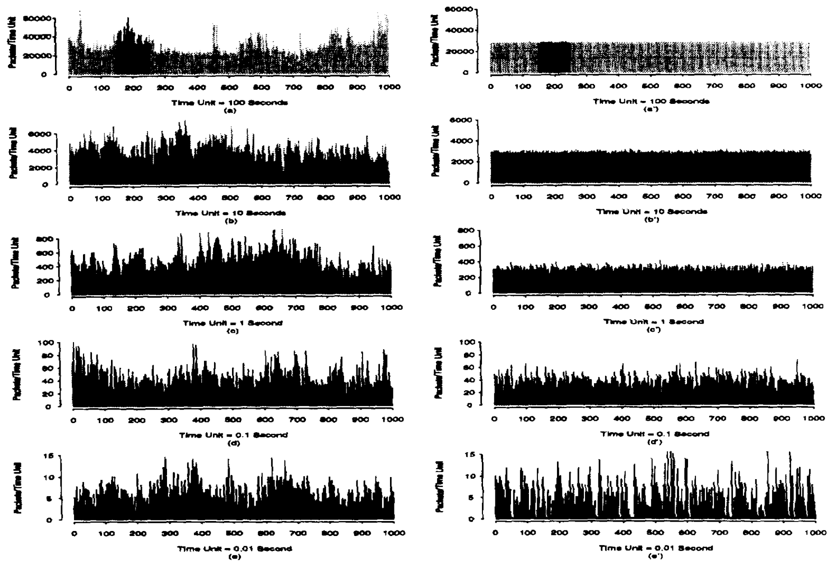

Fig. 4. Pictorial "proof" of self-similarity: Ethernet traffic (packets per time unit) on five different time scales (a)–(e). For comparison, synthetic traffic from an appropriately chosen compound Poisson model on the same five different time scales (a')–(e').

Other products of the literature search often included statements on the causes and the effects of the self-similarity in network traffic. Crovella's paper states how the self- similarity of traffic on the World Wide Web may be caused by the "underlying distributions" of document sizes, caching, user preference, and "think time." Other papers, such as Park's and Peha's papers, imply that the self-similarity is caused also by the protocol used by the hosts to transmit data. Park's paper attests, for example, that self-similarity is found only in TCP (Transfer Control Protocol) suites, and not in UDP (User Datagram Protocol) [5]. It attests that TCP exhibits self-similarity and long-range dependence on differing time scales because of its use of reliable and flow control mechanisms. UDP, on the other hand, shows no long-range dependence because it requires no headers and does not guarantee the arrival of data. Unlike TCP, UDP sends out the packets and does not wait for their acknowledgement. Therefore, there are no retransmissions and no control.

More papers focused on constructing models based on the new self-similarity paradigm. It was found that the 1/0 requests of a disk drive were also self-similar, implying further that the self-similarity was injected into every layer of the OSI stack [7]. Other papers described the mathematics used in the analysis of the new models and/or the gathered traffic.

These mathematical processes are listed below:

Four methods exist to determine the self-similarity of a given data set.

1. Variance time-plot: The variance time plot "relies on the slowly decaying variance of a self-similar series. The variance of $X^{(m)}$ is plotted against *m* on a log-log plot; a straight line with slope $(-\beta)$ greater than -1 is indicative of self-similarity, and the parameter *H* is given by H=1-β/2." [4]
2. R/S plot: The R/S method "uses the fact that four a self-similar dataset, the *rescaled range or R/S* statistic grows according to a power law with exponent *H* as a function of the number of points included *(n).* Thus, the plot of R/S against n on a log-log plot has a slope which is an estimate of H." [4]

3. Periodogram: The periodogram method "uses the slope of the power spectrum of the series as frequency approaches 0. On a log-log plot, the periodogram slope is a straight line with slope β-1=1-2H close to the origin." [4]
4. Whittle estimator: Unlike the preceding three methods, the Whittle estimator provides a confidence interval but is not as useful in exposing faulty assumptions. The drawback for the Whittle estimator is that the "form of the underlying stochastic process must be supplied." The methods cited for supplying this are there fractional Gaussian noise (fGN) with parameter 1/2<H<l, and fractional ARIMA (p, d, q) with 0<d<1/2... These two models differ in their assumptions about the short-range dependences in the datasets; FGN assumes no short range dependence, while fractional ARIMA can assume a fixed degree of short-range dependence." [4]
5. LLCD plots: Among other methods, the log-log complementary distribution plots (LLCD) are used to "assess the presence of heavy tails" in data. An approximately linear relationship over a significant range of at least three orders of magnitude is needed in the tail. [4]
6. Hurst parameter: The Hurst parameter *H* expresses the speed of decat of a time series' autocorrelation function. "As H---> 1, the degree of long-range dependence increases. A test for long-range dependence in a time series can be reduced to the question of determining whether His significantly different from 1/2." [5]

It was implied through the literature that the OSI stack exhibited self-similarity not only in one level but in all its layers. It was conjectured that self-similarity was "injected" into each layer as a function of the individual activity going on at each level. It was implied that each layer had to be considered when discussing the self-similarity (and its management implications) of the network as a whole. This provided a marked departure from the aforementioned literature in that it considered the analysis of self-similarity not only in terms of individual processes but in the interaction of many different, disparate factors. It was decided that the independence of each stack was yet to be analyzed and investigated. Thus, investigation of the separate layers of the OSI stack had to be carried out (fig. 2).

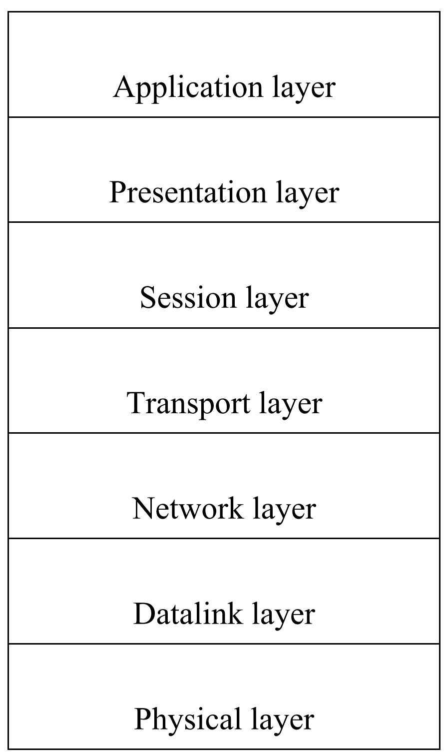

Figure 2. The OSI stack [3]

The OSI stack's application layer complies with the statement from the literature described above that heavy-tailed files lead to long-range dependence. The presentation and session layers are affected by the I/O self-similarity requests of the storage devices [7]. The transport layer is affected by the reliability and flow control mechanisms placed upon it by the TCP protocol (as well as other protocols that require acknowledgements). Because of the constant injection of self -similarity into the stack, as well as the recorded effects of the differences protocol may give, it has been conjectured that even a Poisson distribution may appear to be self-similar given the effect of the different factors that may cause self-similarity.

These implications form the basis of the concepts to be published in the Technical Note. Questions on the applicability of self-similarity on the low-bandwidth networks still need to be addressed. They engender the idea that anew experiment, both theoretical and experimental, may be carried out by ARL in next year's time frame in order to evaluate these same concepts.

Acknowledgements:

I would like to thank and acknowledge my mentors, Dr. John Brand and Ms. Ann E. M. Brodeen for their dedication in teaching me new things and their willingness to let me work along with one of their endeavors. I would also like to acknowledge everyone at ARL that helped me and "put up with me," including Dr. Malcolm S. Taylor, Mr. George Hartwig, Mr. Fred Brundick, Mr. Howell Caton, and Mr. John Dumer. ☺

References:

[1] Kaste, Virginia A. T. "The Information Distribution System: The Fact Exchange Protocol, A Tactical Communications Protocol." BRL-MR-3856 U. S. Army Ballistic Research Laboratory, Aberdeen Proving Ground, MD, 1990.

[2] Leland, Will E., Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. "On the Self-Similar Nature of Ethernet Traffic" (Extended Version). *IEEE/ACM Transactions on Networking,* Vol. 2, No. I, 1994.

[3] Tanenbaum, Andrew S. *Computer Networks.* Third edition. Upper Saddle River, NJ: Prentice Hall PTR, 1996.

[4] Crovella, Mark E. and Azer Bestavros. "Self-Similarity in World Wide Web Traffic; Evidence and Possible Causes." *IEEE Transactions on Networking,* vol. 5, no.6, December 1997.

[5] Park, Kihong, Gitae Kim, and Mark E. Crovella. "On the relationship between file sizes, transport protocols, and self-similar network traffic." 1996 *International Conference on Network Protocols,* pps. 171-180.

[6] Peha, Jon M. "Protocols can make traffic appear self-similar." *Proceedings of the 1997IEE/ACM/SCS Communication Networks and Distributed Systems Modeling And Simulation Conference,* 1997.

[7] Gomez, Maria E. and Vincente Santoja. "Self-similarity in 1/0 workload: Analysis and *Modeling." Workload Characterization: Methodology and Case Studies, pps. 97-104, 1999.*

Bibliography:

[1] Kaste, Virginia A. T. "The Information Distribution System: The Fact Exchange Protocol, A Tactical Communications Protocol." BRL-MR-3856 U. S. Army Ballistic Research Laboratory, Aberdeen Proving Ground, MD, 1990.

[2] Leland, Will E., Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson. "On the Self-Similar Nature of Ethernet Traffic" (Extended Version). *IEEE/ACM Transactions on Networking,* Vol. 2, No. I, 1994.

[3] Tanenbaum, Andrew S. *Computer Networks.* Third edition. Upper Saddle River, NJ: Prentice Hall PTR, 1996.

[4] Crovella, Mark E. and Azer Bestavros. "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes." *IEEE Transactions on Networking,* vol. 5, no.6, December 1997.

[5] Park, Kihong, Gitae Kim, and Mark E. Crovella. "On the relationship between tile sizes, transport protocols, and self-similar network traffic." 1996 *International Conference on Network Protocols,* pps. 171-180.

[6] Peha, Jon M. "Protocols can make traffic appear self-similar." *Proceedings of the* 1997 *IEE/ACM/SCS Communication Networks and Distributed Systems Modeling And Simulation Conference,* 1997.

[7] Gomez, Maria E. and Vincente Santoja. "Self-similarity in 1/0 workload: Analysis and *Modeling." Workload Characterization: Methodology and Case Studies,* pps. 97-104, 1999.

[8] Reif, F. *Fundamentals of Statistical and Thermal Physics.* New York: McGraw-Hill Book Company, 1965.

[9] Reichard, Kevin and Eric F. Johnson. *Teach Yourself... Unix.* Second Edition. New York: MIS Press, 1992.

[10] Tsybakov, B. and N. D. Georganas. "Self-similar processes in Communications Networks." *IEEE Transactions on Information Theory:* pps. 1713-1725, Sept. 1998.

[11] Anonymous. *Maximum Security.* Second Edition. Indianapolis: SAMS Publishing, 1998.

[12] Gleick, James. *Chaos: Making A New Science.* New York: Viking Penguin, 1987.

[13] Qureshi, Aziz *a.* "Modeling Long-Range Dependence in Multimedia Wireless Networks." *Proceedings of IEEE Tencon* '97. *IEEE Region 10 Annual Conference. Speech and Image Technologies for Computing and Telecommunications:* pps. 637-640, vol. 2. 1997.

[14] Feldman, A. and A. C. Gilbert, W. Willinger, and T. *a.* Kurtz. "The Changing Nature of Network Traffic: Scaling Phenomena."

[15] Kaste, Virginia A., Ann E. M. Brodeen, and Barbara Broome. "An experiment to examine protocol performance over combat net radios." BRL-MR-3978 U. S. Army Ballistic Research Laboratory, Aberdeen Proving Ground, MD, 1992.

[16] Grossglauser, Matthias and Jean-Chrysostome Bolot. "On the Relevance of Long-Range Dependence in Network Traffic." *IEEE/ACM Transactions on Networking: pps. 629-640, Vol. 7, no.5. October 1999.*

[17] Andrade Neto, Ernesto L., Antonio M. Alberti, Dalton S. Arantes, and Leonardo S. Mendes." A Realistic Model for Self-Similar Ethernet LAN Traffic in SimATM--An ATM network simulator; Design and Performance Implications.

[18] Shu, Yantai, Fei Xue, Zhigang Jin, and Yan O. "The impact of self-similar traffic on network delay ." *IEEE Canadian Conference on Electrical and Computer Engineering,* 1998. Vol.1, pps. 349-352. 1998.

[19] Tuan, Tsunyi and Kihong Park. "Multiple Time Scale Congestion Control for Self- Similar Network Traffic." *Network Systems Lab, Department of Computer Sciences, Purdue University,* May 1999.

[20] Borella, M. S., S. Uludag, *a.* B. Brewster, and I. Sidhu. "Self-Similarity of lnternet Packet Delay." 1997 *IEEE International Conference on Communications. ICC '97 Montreal, Towards the Knowledge Millennium.* Vol. 1, pps. 513-517, June 1997.

[21] Taylor, Angus E. *Calculus with Analytic Geometry.* Englewood Cliffs, New Jersey: Prentice Hall, 1959.

[22] Brandon, Cameron. *MCSE TCP/IP for Dummies.* Foster City, CA: IDG Books Worldwide, 1998.