



Self-Similarity in Tactical Network Traffic and Tactical Radio Experimentation

**by Ann E. M. Brodeen, John Brand, George W. Hartwig, Jr.,
Frederick S. Brundick, and María C. López**

ARL-TR-3181

April 2004

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Aberdeen Proving Ground, MD 21005-5067

ARL-TR-3181

April 2004

Self-Similarity in Tactical Network Traffic and Tactical Radio Experimentation

**Ann E. M. Brodeen, John Brand, George W. Hartwig, Jr.,
Frederick S. Brundick, and María C. López
Computational and Information Sciences Directorate, ARL**

| REPORT DOCUMENTATION PAGE | | | <i>Form Approved</i> OMB No. 0704-0188 | |
|--|------------------------------------|-------------------------------------|--|--|
| Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS. | | | | |
| 1. REPORT DATE (DD-MM-YYYY) April 2004 | | 2. REPORT TYPE Final | | 3. DATES COVERED (From - To) October 2001–September 2002 |
| 4. TITLE AND SUBTITLE Self-Similarity in Tactical Network Traffic and Tactical Radio Experimentation | | | 5a. CONTRACT NUMBER | |
| | | | 5b. GRANT NUMBER | |
| | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) Ann E. M. Brodeen, John Brand, George W. Hartwig, Jr., Frederick S. Brundick, and María C. López | | | 5d. PROJECT NUMBER AH48 | |
| | | | 5e. TASK NUMBER | |
| | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Research Laboratory ATTN: AMSRD-ARL-CI-CT Aberdeen Proving Ground, MD 21005-5067 | | | 8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-3181 | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | | 10. SPONSOR/MONITOR'S ACRONYM(S) | |
| | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited. | | | | |
| 13. SUPPLEMENTARY NOTES | | | | |
| 14. ABSTRACT Network traffic has been reported to exhibit self-similar properties over time. This behavior has been traced to a number of factors, including file-size distributions and characteristics of current network protocols. In this report, we present results from examining tactical military networks to determine if they also exhibit self-similar behavior. Tactical networks differ from typical commercial networks in that the effective throughput is orders of magnitude lower, protocols are particular to the tactical function they support, and the data transmitted is typically more homogeneous than that found on commercial networks. Tactical network utilization is often decidedly different than that found on commercial networks. The data used in this study were taken from two sources: (1) synthetic traffic streams generated using a fractional Gaussian noise generator and (2) data from field exercises. | | | | |
| 15. SUBJECT TERMS self-similarity, Hurst parameter, tactical communications | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT UL | 18. NUMBER OF PAGES 26 |
| a. REPORT UNCLASSIFIED | b. ABSTRACT UNCLASSIFIED | c. THIS PAGE UNCLASSIFIED | | |
| | | | 19b. TELEPHONE NUMBER (Include area code) 410-278-8947 | |

Contents

| | |
|--|-----------|
| List of Figures | iv |
| List of Tables | iv |
| 1. Introduction | 1 |
| 2. Analytical Methods | 1 |
| 2.1 Statistical Methods | 1 |
| 2.2 Synthetic Traffic Generator | 2 |
| 3. Data Analysis | 4 |
| 3.1 Synthetic Traffic | 4 |
| 3.2 Field Test Data | 7 |
| 4. Discussion | 8 |
| 5. Implications of Multifractal Behavior | 10 |
| 5.1 Implications for Current Work | 10 |
| 5.2 Implications for Intrusion Detection | 10 |
| 5.3 Implications for Analysis | 10 |
| 6. Summary | 11 |
| 7. References | 12 |
| Appendix A. Variance-Time Plots of Synthetic Traffic | 14 |
| Appendix B. Description of Fire Support Team Force Development Testing and Experimentation II (FIST FDT&E II) | 16 |
| Distribution List | 19 |

List of Figures

| | |
|--|----|
| Figure 1. Variance-time plot of synthetic traffic streams with differing H | 5 |
| Figure 2. Variance-time plot of synthetic traffic streams with differing H merged. | 6 |
| Figure 3. Variance-time plot showing effects of increased mean and variance on merged traffic streams. | 7 |
| Figure 4. FIST FDT&E II data with reference line, slope = -1 ($H = 0.5$). | 8 |
| Figure A-1. Variance-time plot of synthetic traffic streams with $H = 0.5$ (uncorrelated). | 14 |
| Figure A-2. Variance-time plot of synthetic traffic streams with $H = 0.7$ (moderately correlated). | 14 |
| Figure A-3. Variance-time plot of synthetic traffic streams with $H = 0.9$ (highly correlated). | 15 |
| Figure B-1. Layout of the FIST FDT&E II network. | 17 |

List of Tables

| | |
|--|----|
| Table B-1. Digital message traffic collected. | 17 |
| Table B-2. Good messages by type. | 18 |

1. Introduction

Network traffic has been discovered to be self-similar over time under numerous circumstances in commercial networks (Leland et al., 1994; Willinger et al., 1995). These are high-capacity networks primarily operating under the Transmission Control Protocol/Internet Protocol (TCP/IP) or asynchronous transfer mode packet systems (Park et al., 1996; Park and Willinger, 2000). Some tactical radio traffic data exist from the Fire Support Team Force Development Testing and Experimentation II (FIST FDT&E II). In addition, the U.S. Army Research Laboratory (ARL) has performed analyses of tactical communications protocols in the laboratory using controlled radio networks driven by Poisson-generated traffic. These channels are of extremely low throughput by commercial standards, by factors of 5000 or more, and the tactical communications protocols are different from those based on the TCP/IP suite.

This study examines the properties of self-similar traffic using the FIST FDT&E II data and synthetic traffic streams generated using Paxson's fast Fourier transform of fractional Gaussian noise (FFT FGN) method (Paxson, 1995). The field data and the generated traffic data are analyzed using several tools, which are discussed in the sections that follow.

2. Analytical Methods

2.1 Statistical Methods

Several tests have been described in the literature (Crovella and Bestavros, 1997; Grossglauser and Bolot, 1999; Leland et al., 1994; Park et al., 1996; Park and Willinger, 2000) and employed to analyze self-similar traffic. To date, only the variance-time plot test has been applied. This, in graphical form, is a method for determining whether the traffic stream is multifractal. Tests that analyze a traffic stream and determine a numerical value based on the assumption of monofractal composition may be in error if the stream is multifractal.

The sensitivity of the variance-time method is investigated by generating synthetic traffic with multiple self-similar components. The statistical tool used to analyze the traffic is a software program written to determine the variance of sample means vs. the sample size (the "variance-time plot"). For events that are independent over time, this tool is based on the relation

$$\text{Var}(\bar{X}) = \sigma^2 n^{-1}. \quad (1)$$

For events that are correlated over time and also display scale invariance,

$$\text{Var}(\bar{X}) = \sigma^2 n^{-\beta} \text{ (Beran, 1994)}. \quad (2)$$

If the variance of the sample means is plotted against the sample size on a log-log graph, the slope $-\beta$ of the resultant line is -1 for independent events. The factor β is related to the Hurst parameter (H), a heuristic measure of the degree of self-similarity, by the expression

$$\beta = 2(1 - H). \quad (3)$$

H typically varies between 0.5 and 1.0. A value of $H = 0.5$ corresponds to a sequence of events that is not correlated over time; $H > 0.5$ indicates increasing degrees of self-similarity. One attractive feature of this graphical technique is that it may reveal multifractal behavior. As the degree of self-similarity increases, the traffic stream becomes bursty over all time scales. Under certain conditions, this allows streams of data with more than one component to reveal the existence of multiple degrees of self-similarity in the traffic.

The question arises over just what the variance-time plot reveals. The literature on self-similar traffic treats three types of self-similarity: (1) strict self-similarity, where the traffic exhibits a set scale invariance that defines fractal behavior, (2) second-order self-similarity, and (3) asymptotic second-order self-similarity. The latter two involve the behavior of the autocorrelation function, a second-order function. Autocorrelation implies long-range order or dependence. The problem is that there are self-similar processes that are not long-range dependent and long-range dependent processes that are not self-similar. However, for asymptotic second-order self-similar processes and if $H > 0.5$, self-similarity implies long-range dependence and vice-versa (Park and Willinger, 2000). In this case, the synthetic traffic is generated from a self-similar process, fractional Gaussian noise, rather than being merely second-order self-similar.

The tool as written chose the initial aggregation size randomly, within limits. If the randomly chosen aggregation size is $>1/6$ of the total number of sampled traffic points or $<1/10$ of the total number of points, another aggregation size is drawn.* The start point is also randomly chosen. This usually leads to discarding “orphan” points at the beginning and end of the traffic stream, but the sample is chosen large enough that it is felt this could be disregarded. It is hoped, in this way, to compensate for any systematic variation in the traffic. The aggregation size is then decreased by a chosen factor, and the calculation is repeated. It might be noted that this leads to smaller and smaller numbers of orphan points. When the aggregation size reaches a chosen factor, usually <8 points, the calculation is terminated.

The results obtained using this tool are discussed in sections 3.1–3.2 and appendix A.

2.2 Synthetic Traffic Generator

Schuler’s program (Schuler, 1995), based on the work of Paxson (1995), was used to generate streams of synthetic traffic with controlled properties. The program option of generating a real

* These values were chosen arbitrarily to reduce scatter arising from the availability of only a few samples or a few points within the samples. Other values would work as well.

traffic trace was used. Traffic streams with differing means, variances, and degrees of self-similarity, as reflected in H , were generated. The traffic streams could be of any size but, typically, were over a million (2^{20}) increments. Larger traffic streams could easily be generated because the program is very fast, but a million time increments seemed adequate.

There are at least two ways of imposing two or more self-similar processes on a traffic stream. One is to generate a single traffic stream by successively imposing two or more processes that, individually, impose self-similarity. Another is to merge two or more streams of traffic, each with a degree of self-similarity. The latter was selected in this case, since all of the tools to do so were on hand. Traffic streams with multiple fractal properties were generated by merging two streams with differing values of H . These were used to estimate the ability of the statistical tool to distinguish between traffic streams of differing properties.

The merging of two streams may represent the case of a network with streams from two independent sources of differing self-similarity characteristics. A traffic stream generated by a series of processes, each of which, acting alone, would generate a stream with one set of self-similar characteristics, might represent an open system interconnect (OSI) stack with a heavy-tailed file-size distribution and a TCP that also imposes self-similarity on the traffic stream. The latter may be done in future research, since it is a very plausible representation of tactical network traffic. For simplicity, the analysis has so far been restricted to the simpler case—the merger of two independent traffic streams. No attempt has been made to impose carrier sense multiple access on the merging process; this amounts to an assumption of a link running at a sufficiently low fraction of the possible throughput such that collisions are unlikely. This is not a realistic representation of a tactical network but is a simple way to begin the analysis.

It is highly desirable to model the imposition of multiple self-similar processes on a single traffic stream; the Optimized Network Engineering Tool (OPNET)* provides excellent software for this, since it requires explicitly modeling the OSI stack. It is also highly desirable to model CSMA in heavily loaded networks and to have multiple traffic streams of differing characteristics. Again, OPNET may be a very useful tool for this.

The output from the traffic generator was analyzed using the variance-time plot discussed previously. The output was reasonably well-behaved, but there are several cautions in the use of the Paxson model. The first is that it is an approximation. According to Paxson, the output trace produced by his methodology is only approximately self-similar, even though the output trace passes four tests for self-similarity:

1. the variance of the sample means vs. sample size (the “variance-time plot”),
2. Beran’s goodness-of-fit test,

* OPNET is a registered trademark of OPNET Technologies, Inc., Bethesda, MD.

3. Whittle's estimator, and
4. normality or near-normality of the marginal distribution of the sample.

Paxson refers to this as the “quacks like a duck” phenomenon: if it walks like a duck and quacks like a duck, one may as well call it a duck. In that case, if the data act self-similar, then for engineering purposes they may be treated as such (Paxson, 1995).

The output trace from the traffic generator produced for this effort was tested using the variance-time plot. The data were well behaved; the estimated H of the resulting traffic stream was approximately that of the input value of H . Visual inspection of the goodness of fit should be validated by more rigorous statistical methods.

In addition, the traffic generator lists intermediate and final values of the traffic stream mean and variance and cautions that they may vary from the input values. The final output values for the mean and variance were found to replicate the input values exactly. In addition, when the traffic streams were analyzed by a separate program to calculate the mean and variance, the mean and variance of the generated output trace were found to be in very good agreement with the input values. The other traffic parameters also appeared to be well behaved. This is discussed below. Although the synthetic traffic was well behaved, future work should include new methods of generating the synthetic traffic to check this assumption.

3. Data Analysis

The ARL has performed a number of statistically controlled tactical radio network experiments. The baseline experiment and its results are fully described elsewhere (Kaste et al., 1992). This report focuses on exploring the self-similar properties of tactical communications data collected during a 1984 field experiment and synthetically generated laboratory data.

3.1 Synthetic Traffic

To gain some insight into the phenomenon of self-similarity and to estimate the utility of the variance-time tool, synthetic traffic was generated as described previously. Several observations are in order.

The tool should respond to some degree to the input mean value of the traffic streams because the variance is a second-order phenomenon. This can be seen in appendix A, figures A-1 through A-3. For example, figure A-1 shows uncorrelated traffic with means differing by 20:1. There is a substantial difference. Figure A-2 illustrates moderately correlated traffic with means varying by 2:1. There is little impact in that case. Streams with fixed values of H and variance but differing values of the overall (input) mean were plotted together. The plots were similar, increasing only slowly as the sample mean increased, for values of $H = 0.5, 0.7, \text{ and } 0.9$.

In addition, plots were made of streams with the same values of the input mean and variance but differing values of H (figures 1 and 2). Since this test is sensitive to the long-range correlation of traffic, it is not surprising that traffic with higher values of H should be differentiated at longer time scales, provided the long-term memory is strong enough. Likewise, since self-similarity accompanies burstiness, it is not surprising that traffic with a higher degree of self-similarity should be differentiated from traffic with little or no self-similarity at lower time scales as well. The implication is that, if the time scale can be sufficiently extended, traffic streams with components of differing self-similarity should be discernable with this graphical method.

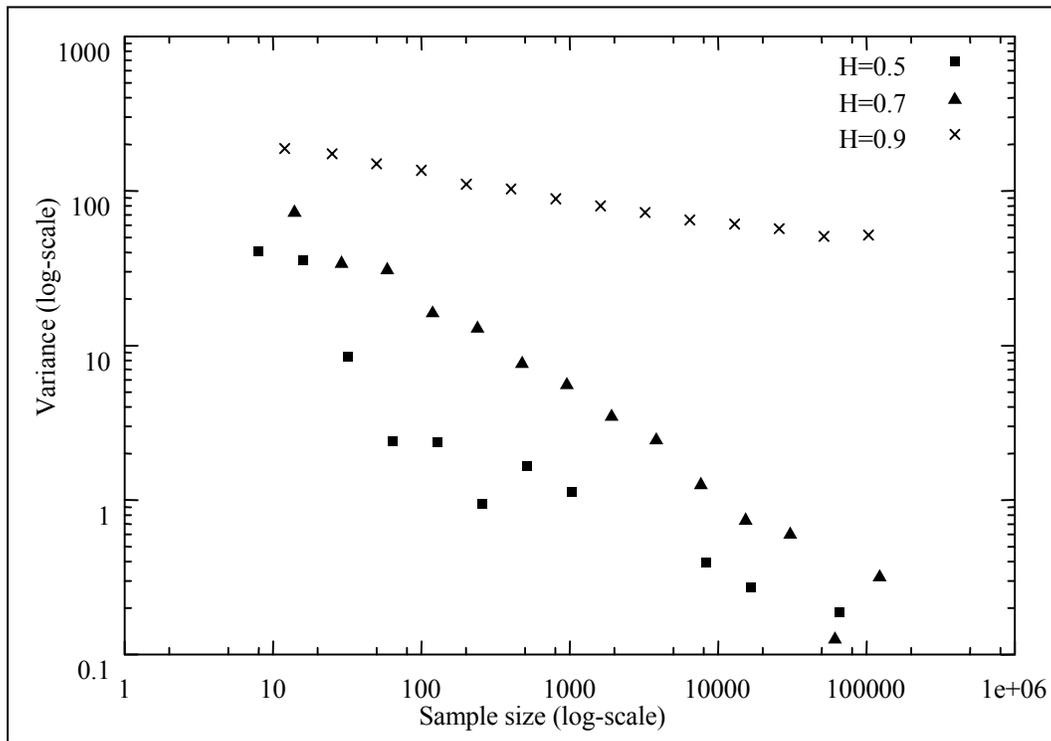


Figure 1. Variance-time plot of synthetic traffic streams with differing H .

Traffic streams of differing degrees of self-similarity were generated as discussed using Paxson's scheme for approximating self-similar traffic, which uses an FFT FGN. Traffic streams with $H = 0.5$ (uncorrelated), $H = 0.7$ (moderately correlated), and $H = 0.9$ (highly correlated) were generated and are plotted in figure 1. Initially, the traces were 1,048,576 arbitrary time steps (2^{20}) with traffic expressed in terms of arbitrary units. These traffic streams had input mean values of 1000 and variances of 300. That is, from the standpoint of an analysis, ignoring self-similarity, the traffic streams are nominally of identical statistical characteristics. As can be seen in figure 1, the variance of the sample means vs. sample size of the synthetic traffic indicates that the characteristics of the traffic streams are radically different.

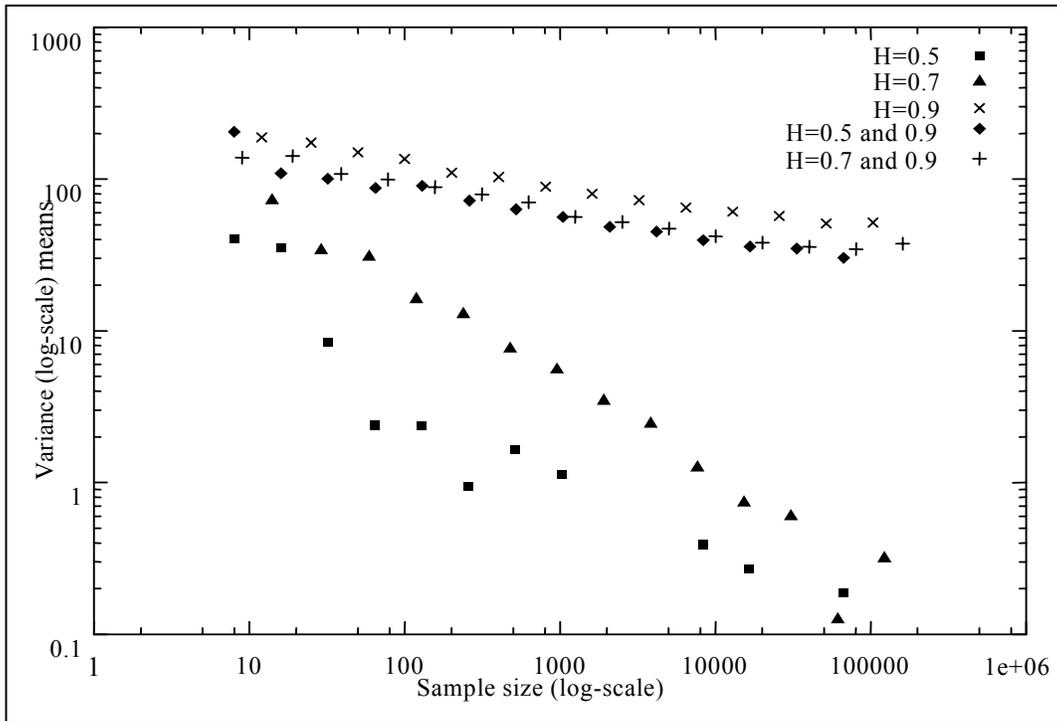


Figure 2. Variance-time plot of synthetic traffic streams with differing H merged.

Figure 2 indicates that when streams with similar input means and input variances but differing H s are summed pairwise, the trace of the summed traffic is essentially indistinguishable by inspection from the trace of the more self-similar traffic stream. That is, the contribution to the variance of the sample means (a measure of burstiness) from the more self-similar traffic, under these conditions, masks all other contributions from all other components of the traffic. It is also clear that a traffic stream with $H = 0.5$, but a substantially greater mean and variance, might well show up in the analysis of the summed stream. This is illustrated in figure 3. Indeed, analysis of the field test data indicates such a phenomenon might have taken place in two of the exercises.

In this variance-time plot, three synthetic traffic streams are shown. One stream ($H = 0.5$, input mean = 100,000, and input variance = 10,000) represents the background traffic. Another stream ($H = 0.9$, input mean = 500, and input variance = 300) represents a highly variable but highly correlated traffic stream, perhaps, of a network scan. The merged stream has a mean of 100,500 and a variance essentially the same as the larger stream. The merged stream shows a pronounced “knee” in the curve, since the second-order characteristics of the traffic stream with $H = 0.9$ dominate the curve for larger sample sizes and the characteristics of the uncorrelated stream dominate the curve for smaller sample sizes. Note this assumes the long-term memory of the $H = 0.9$ stream is quite strong. For background traffic with $H = 0.5$ but lower variance, the characteristics of the merged stream are indistinguishable by inspection from the highly self-similar stream, but these curves are not shown. Many permutations of the traffic characteristics

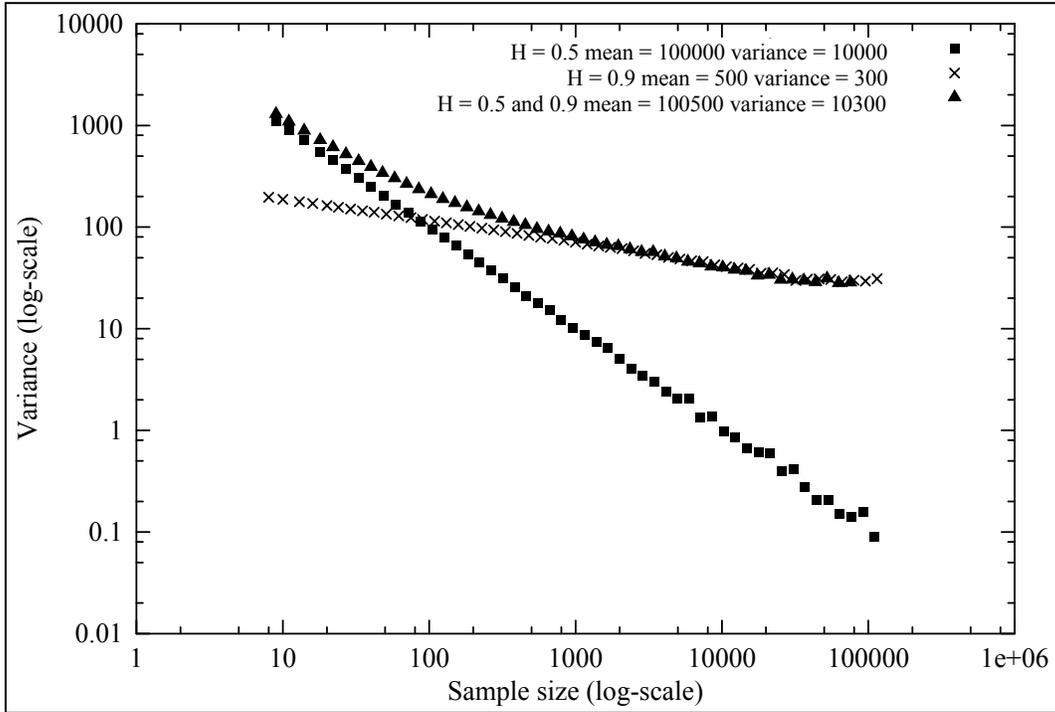


Figure 3. Variance-time plot showing effects of increased mean and variance on merged traffic streams.

are possible, including those where the smaller stream would be undetectable by graphical methods but possibly detectable by more subtle statistical analysis and those where the smaller stream would be simply undetectable.

3.2 Field Test Data

The three traffic streams from the field exercise are plotted in figure 4. From the variance-time plot, the traffic clearly exhibits characteristics of self-similarity. Approximate fits to the data are shown; all three traffic streams are quite different from the reference line, with the slope indicating $H = 0.5$. Interestingly, the traffic exhibits a considerable degree of self-similarity, with the appearance of $H \sim 0.9$. Even more interesting are the traffic streams, corresponding to the first week of the field exercise (FEX 1), which appear to consist of two components: one with $H \sim 0.5$ and the other with $H \sim 0.9$. The stream produced by the third week of testing (FEX 3) also appears to exhibit two components. The other, from the second week (FEX 2), does not appear to do so. Since the traffic was generated under similar conditions and constraints, it is not clear why the streams would differ or what phenomena produced the differences.

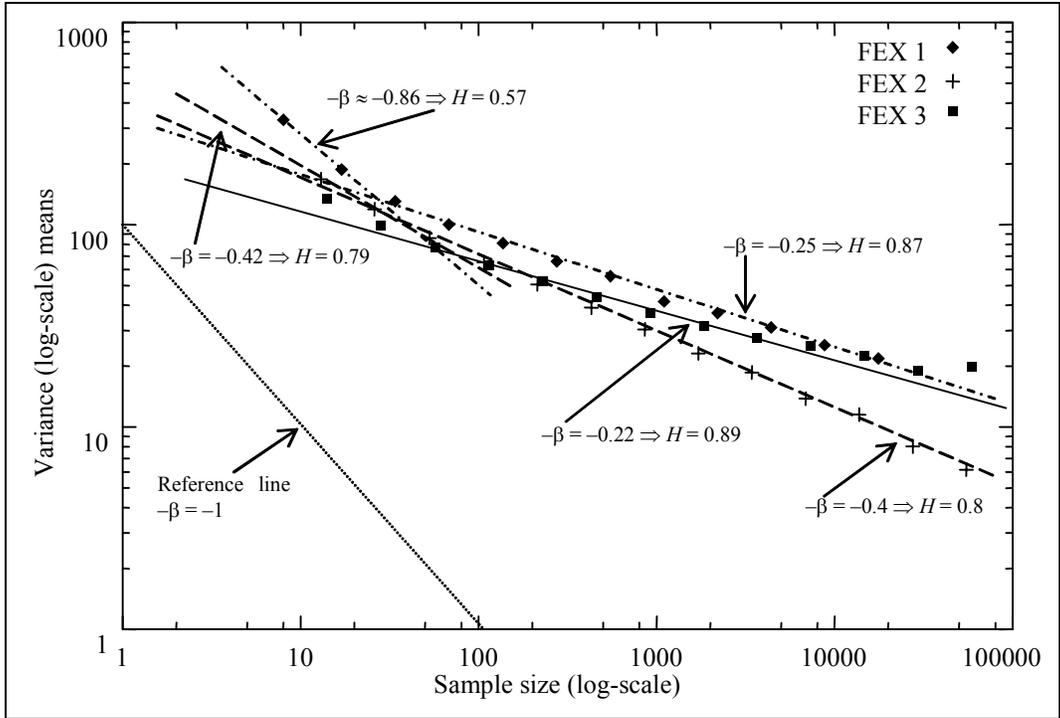


Figure 4. FIST FDT&E II data with reference line, slope = -1 ($H = 0.5$).

4. Discussion

There are a number of reasons to suspect that the field test data would demonstrate some degree of self-similarity. Some of these phenomena might also reasonably be expected to apply to the results of a laboratory test. Possible phenomena of interest in the field test data are the reliable transport mechanism and file size distribution.

Traffic using the TCP/IP suite has been reported as self-similar in contrast to User Datagram Protocol traffic, which shows little long-range dependence (Park et al., 1996). This is reasonable. TCP/IP, in imposing a reliability and control mechanism, produces a series of messages to ensure delivery. Thus, a message triggers predictable traffic into the future. It has been stated that retransmission mechanisms may make even a Poisson distribution traffic model appear self-similar (Tsybakov and Georganas, 1998).

The Tactical Fire Direction System (TACFIRE) protocol acts to a degree as a reliable transport protocol. The TACFIRE protocol requires responses from the message recipients similar to the SYN and ACK processes of the TCP/IP. Several possible factors that may impose self-similarity have been enumerated in the literature, among them being the effects of queued and buffered

traffic (addition of flow and reliability mechanisms), user “think time,” and the implications from the ON-OFF model, as described in (Grossglauser and Bolot, 1999; Willinger et al., 1997).

Parenthetically, it should be noted that the series of ACK messages typical of the TCP/IP are also found in the X.25 protocol, the staple underlying protocol of the Mobile Subscriber Equipment (MSE), the military cellular radio system. The implication is that the packet network based on MSE may also exhibit self-similar behavior, with its own characteristic time scale. This was not a factor in the ARL experiments but may well be a factor in actual tactical traffic. The field test data described in Appendix B do not involve an MSE and, therefore, do not involve an X.25 link either.

Any self-similarity actually generated by these phenomena in the laboratory test data may be hard to find because of low message rates that, in spite of being low compared to commercial applications, still loaded the net sufficiently to preclude a greater range of traffic rates. The laboratory test data were collected across a range of message sizes, but these messages were independent and not part of a train of packets composing a single large file. This may be a factor in tactical networks.

A file-size distribution weighted with a tail of large files (heavy-tailed file-size distributions) has been demonstrated to impose self-similarity on network traffic (Leland et al., 1994; Peha, 1997). It is likely that file-size distributions differ from place to place in the tactical net, at least at the present time. At the small-unit level, each combat vehicle has an appliqué computer, and updates to the situation maps occur frequently. These map updates are vector-based and, hence, shorter than bitmaps. Other messages such as calls for fire are preformatted and short; file sizes in the field test data included numerous preformatted messages of 20 and 52 bytes.

At the headquarters level, the situation is more complex. In a tactical headquarters today, there is a host of preformatted messages such as Maneuver Control System (MCS) data. These are individually short, but in a sense, MCS updates could be considered a very long message (file), since a particular update will surely be followed by another later on that is part of the same sequence. This also implies a contribution to autocorrelation over long time periods. Other command-related traffic could involve pictures, maps, and planning documents. The traffic at different levels of a tactical network might well be dominated by different file-size distributions. If self-similarity occurs at the two levels, there might well be different time scales or levels of self-similarity for the traffic streams, hence, different fractal properties. If traffic from two levels shares the same trunk, the net traffic would be a superposition of traffic with two or more time scales (multifractal), but relative characteristics of the two streams might mask all but a time scale associated with dominant traffic. Recent efforts to deliver imagery from satellites and air reconnaissance to units at the lowest level indicate that substantial file sizes may well extend to all areas of the tactical net, further complicating the issue.

5. Implications of Multifractal Behavior

5.1 Implications for Current Work

There is ongoing work to investigate multifractal properties of network traffic. This work is, however, on commercial traffic, not tactical. An unusually tacit assumption made by the investigators is that the traffic does not change its fundamental properties (i.e., stationarity) during the time intervals analyzed and the time the data are grouped. The emerging trends outlined in this report indicate that tactical network traffic may, indeed, be multifractal. The combat cycle may well produce traffic that is emphatically not stationary in its properties.

An example is defense, during a pause in the operational tempo, shifting to the assault; the times would be the most demanding and crucial to combat success. Network designs or specifications generated as a result of the application of the results of research on commercial traffic may well lead to massive inadequacies of network capability when the tactical networks in their turn violate stationarity.

5.2 Implications for Intrusion Detection

The possibility of differentiating between two or more streams of traffic of differing levels of self-similarity may make it possible to detect the presence of a new, alien traffic stream in a network. An example might be an attack or reconnaissance scan of a network. If the scan were to interrogate many different hosts in a particular block of IP addresses, switching back and forth from one host to another to avoid raising suspicions by dwelling on a particular host and not presenting a regular pattern by interrogating a set pattern of ports, the actual traffic to all of the targets might be distinguishable as such over the network. It is necessary to measure the self-similarity and other characteristics of traffic produced by several attack tools and reconnaissance tools (scanners) and to measure the properties of a network as time and circumstances change. A possible example of such a scenario is illustrated in figure 4. In this variance-time plot, several traffic streams are shown. In particular, a typical traffic stream with $H = 0.5$ merged with a traffic stream with $H = 0.9$ should present a clear indication of a change to the flow composed of two distinct streams, even though the overall (input) mean of the traffic stream with $H = 0.5$ is 200 times greater than the other. The usual network traffic should be characterized by $H > 0.5$, but the message flow from a network scan might be even larger.

5.3 Implications for Analysis

A method for discerning traffic characteristics of multifractal traffic must be obtained. Wavelet-based methods have been used by other researchers and seem to hold promise for this work as well (Park and Willinger, 2000).

6. Summary

Field test data gathered during the FIST FDT&E II indicated that a low bandwidth tactical radio network generated traffic exhibiting the characteristics of self-similarity. This was not unexpected. Furthermore, in two of the traffic streams, there appeared to be evidence of two components of self-similarity corresponding to near-zero correlation and highly correlated. The reasons were not clear.

7. References

- Beran, J. Statistics for Long-Memory Processes. *Monographs on Statistics and Applied Probability 61*; Chapman & Hall/CRC: New York, 1994.
- Crovella, M. E.; Bestavros, A. Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. *IEEE/ACM Trans. Networking* **1997**, 5 (1), pp 835–846.
- Grossglauser, M.; Bolot, J.-C. On the Relevance of Long-Range Dependence in Network Traffic. *IEEE/ACM Trans. Networking* **1999**, 7 (5), pp 629–640.
- Kaste, V. A. T.; Brodeen, A. E. M.; Broome, B. D. *An Experiment to Examine Protocol Performance Over Combat Net Radios*; BRL-MR-3978; U.S. Army Ballistic Research Laboratory: Aberdeen Proving Ground, MD, June 1992.
- Leland, W.; Taqqu, M. S.; Willinger, W.; Wilson, D. V. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Trans. Networking* **1994**, 2 (1), pp 1–15.
- Park, K.; Willinger, W. *Self-Similar Network Traffic and Performance Evaluation*; Wiley & Sons: New York, 2000.
- Park, K.; Kim, G.; Crovella, M. E. On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic. *Proceedings of the Fourth International Conference on Network Protocols*, Columbus, OH, 1996; pp 171–180.
- Paxson, V. *Fast Approximation of Self-Similar Network Traffic*; LBL-36750; Lawrence Berkeley Laboratory: University of California (Berkeley), April 1995.
- Peha, J. M. Protocols Can Make Traffic Appear Self-Similar (Expansion of research published as Retransmission Mechanisms and Self-Similar Traffic Models). *Proceedings of the 1997 IEEE/ACM/SCS Communication Networks and Distributed Systems Modeling and Simulation Conference*, Phoenix, AZ, 1997; pp 47–52.
- Schuler, C. Research Institute for Open Communication Systems, GMD FOKUS.
http://ita.ee.lbl.gov/html/contrib/fft_fgn_c.htm (accessed 18 April 2002), Hardenbergplatz 2, D-10623 Berlin, Germany, August 1995.

Tsybakov, B.; Georganas, N. D. Self-Similar Processes in Communications Networks. *IEEE Trans. Inform. Theory* **1998**, *44* (5), 1713–1725.

Willinger, W.; Taqqu, M. S.; Sherman, R.; Wilson, D. V. Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level. *IEEE/ACM Trans. Networking* **1997**, *5*, 71–86.

Willinger, W.; Taqqu, M. S.; Leland, W. E.; Wilson, D. V. Self-Similarity in High-Speed Packet Traffic: Analysis and Modeling of Ethernet Traffic Measurements. *Stat. Sci.* **1995**, *10* (4), 67–85.

Appendix A. Variance-Time Plots of Synthetic Traffic

Note that the plots in figure A-1 fall in approximately the same region over a considerable range of input means (500, 600, 1000, and 10,000 units) and a fixed input variance of 300 units.

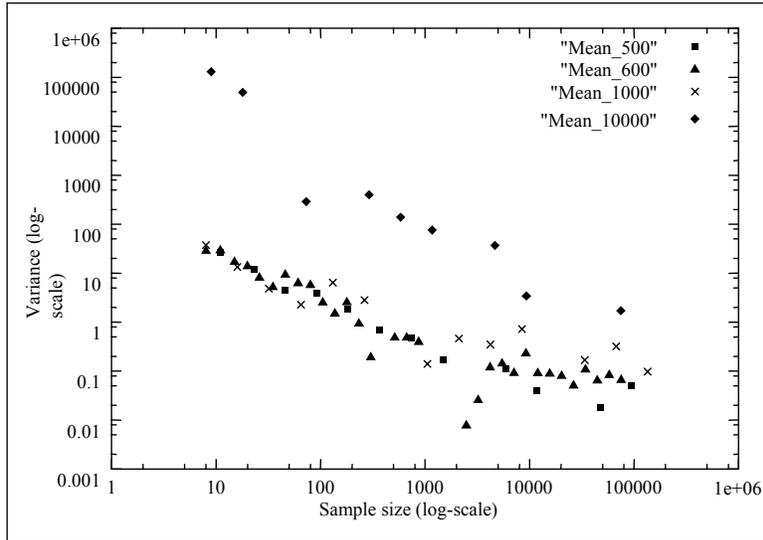


Figure A-1. Variance-time plot of synthetic traffic streams with $H = 0.5$ (uncorrelated).

Figure A-2 indicates that for moderately correlated traffic ($H = 0.7$), with a fixed input variance of 300 units and input means ranging from 500 to 1000 units, the plots fall approximately on the same line.

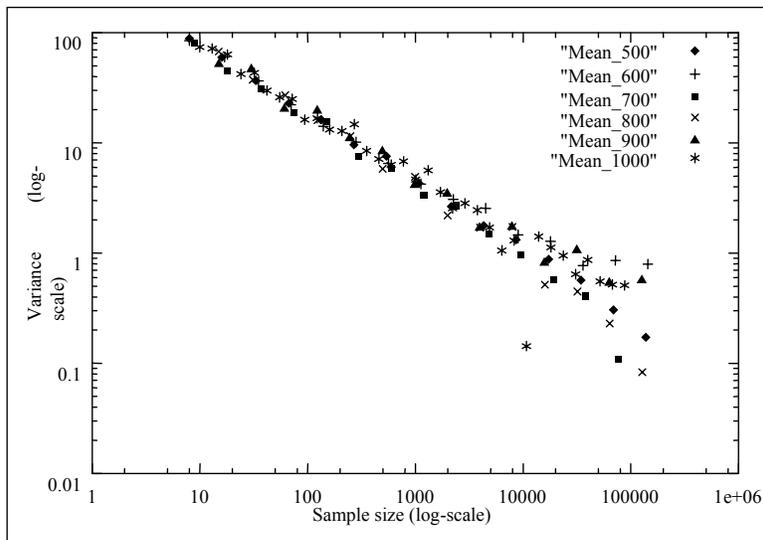


Figure A-2. Variance-time plot of synthetic traffic streams with $H = 0.7$ (moderately correlated).

The behavior of the plot in figure A-3 is similar to the behavior of the plot in figure A-1. The input variance in this plot has also been fixed to 300 units. Note the occasional outlying points. These are probably artifacts of the small sample sizes. Scatter due to large sample size may also occur as the number of samples becomes small. Truncation of the analysis program is set to a cell size of ≤ 8 .

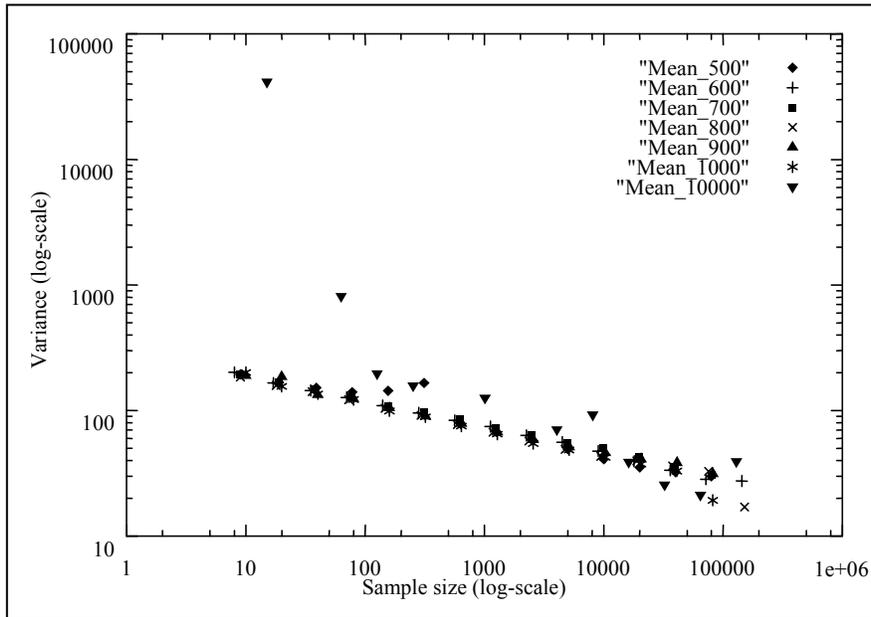


Figure A-3. Variance-time plot of synthetic traffic streams with $H = 0.9$ (highly correlated).

Appendix B. Description of Fire Support Team Force Development Testing and Experimentation II (FIST FDT&E II)

The field test data were gathered during the FIST FDT&E II. This test was conducted and extensively documented by the Human Engineering Laboratory and the U.S. Army Ballistic Research Laboratory.^{1,2} The test involved a mechanized infantry task force composed of two mechanized infantry companies and one armor company. The task force was one-third of the Field Artillery assets of a maneuver brigade. The experiment was run in three iterations at Ft. Riley, KS. Two iterations were Scenario-Oriented Recurring Evaluation System Europe V-based scenario driven field exercises (FEX), while the third was a free-play force-on-force exercise. The networks tested were devoted to digital traffic generated by digital message devices, using preformatted messages, over analog fixed frequency voice radios (AN/PRC 77). The layout of the network is shown in figure B-1.

Data were collected manually and by use of traffic recorded from a central point. The taped data were digitized and inserted into a message collection and recording system. The recorded traffic was frequency shift keyed, and the traffic was converted to RS 232 ASCII format by use of “bit box,” modems, which allow Tactical Fire Direction System (TACFIRE) hardware to communicate with commercial computers. Messages were sorted by fire mission into three categories: (1) messages grouped by fire mission number, (2) unknown messages, and (3) messages known but not part of a fire mission. The data were sent to the Field Artillery School at Ft. Sill, OK, for in-depth analysis and conversion into a comprehensive database. The data analyzed here are raw traffic, not sorted or otherwise reduced by purpose.²

Data were collected for three networks: (1) the Company Fire Control Net (CFC Net), (2) the Mortar Fire Direction Net (MFD Net), and (3) the Fire Direction Net (FD Net). It is not clear which data (FEX 1, 2, or 3) for a data stream used here came from which net, though they were apparently interleaved by time. A total of 344 hr of data was collected, 22 hr of which were from noncontrolled parts of the test and not digitized. The other 322 hr of data were categorized, as shown in table B-1.

The breakdown into message type is extremely interesting, as shown in table B-2. The variable format (VF) message traffic over the CFC3 and MFD nets is surprising because these nets contained digital devices capable of producing only fixed format messages. Kaste et al.¹ speculated that players with devices capable of generating VF messages were probably on the wrong communications nets. This is reasonable. Unlike a voice net, at that time, there was no

¹ Kaste, V. A. T.; Brodeen, D. C.; Winner, W. A. *Description of the Digital Data Collected From FIST FDT&E II*; BRL-TR-2676; U.S. Army Ballistic Research Laboratory: Aberdeen Proving Ground, MD, September 1985.

² Grynovicki, J. O.; Smith, J. H. *Experimental Design and Analysis for the FIST Force Development Testing and Experimentation II*; BRL-MR-3474; U.S. Army Ballistic Research Laboratory: Aberdeen Proving Ground, MD, October 1985.

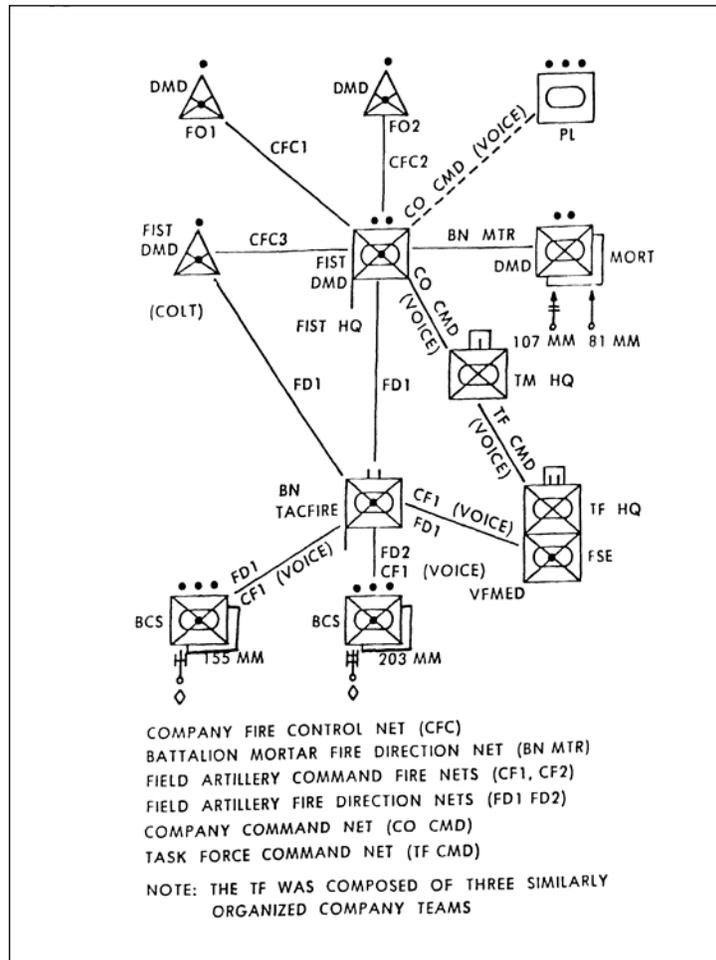


Figure B-1. Layout of the FIST FDT&E II network.

Table B-1. Digital message traffic collected.¹

| Nets | Good Messages | Unintelligible Messages | Percent Unintelligible | Total Messages |
|--------|---------------|-------------------------|------------------------|----------------|
| CFC1 | 10,103 | 746 | 7 | 10,849 |
| CFC2 | 9268 | 1341 | 13 | 10,609 |
| CFC3 | 11,839 | 1398 | 11 | 13,237 |
| FD1 | 97,880 | 2631 | 3 | 100,511 |
| FD2 | 24,552 | 1430 | 6 | 25,982 |
| MFD | 5780 | 3651 | 39 | 9431 |
| Totals | 159,422 | 11,197 | 7 | 170,619 |

Table B-2. Good messages by type.¹

| Nets | ACKs | Fixed Format | Variable Format | Total Messages |
|-------------|-------------|---------------------|------------------------|-----------------------|
| CFC1 | 4569 | 5534 | 0 | 10,103 |
| CFC2 | 3578 | 5690 | 0 | 9268 |
| CFC3 | 5503 | 4069 | 2267 | 11,839 |
| FD1 | 44,512 | 33,411 | 19,957 | 97,880 |
| FD2 | 10,688 | 118 | 13,746 | 24,552 |
| MFD | 2381 | 3233 | 76 | 5780 |
| Totals | 71,231 | 52,145 | 36,046 | 159,422 |

authentication or error-checking procedure for digital networks as there is in the voice radio procedure (e.g., Bravo Tango seven-four, this is Golf Sierra six; you do not belong on this net, over). Also, all but four of the messages on the CFC3 and MFD nets occurred during the free-play FEX 3; the other four occurred during the noncontrolled portion of FEX 2. That could account for a contribution of self-similarity in FEX 3; unfortunately, there is only a hint of such behavior (see figure 4).

VF messages can be lengthy (100–1500 characters) as opposed to fixed format messages (48 characters). The message length does not include synchronization characters or the preamble. It may be that the message traffic generated by VF messages introduces a degree of autocorrelation: if the message is being transmitted, there is a good chance it is still being transmitted some time from now. It is, however, more likely that the self-similarity is largely an artifact of the reliable transport function of the TACFIRE protocol.

NO. OF
COPIES ORGANIZATION

1
(PDF
Only) DEFENSE TECHNICAL
INFORMATION CENTER
DTIC OCA
8725 JOHN J KINGMAN RD
STE 0944
FT BELVOIR VA 22060-6218

1 COMMANDING GENERAL
US ARMY MATERIEL CMD
AMCRDA TF
5001 EISENHOWER AVE
ALEXANDRIA VA 22333-0001

1 INST FOR ADVNCD TCHNLGY
THE UNIV OF TEXAS
AT AUSTIN
3925 W BRAKER LN STE 400
AUSTIN TX 78759-5316

1 US MILITARY ACADEMY
MATH SCI CTR EXCELLENCE
MADN MATH
THAYER HALL
WEST POINT NY 10996-1786

1 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CS IS R
2800 POWDER MILL RD
ADELPHI MD 20783-1197

3 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CI OK TL
2800 POWDER MILL RD
ADELPHI MD 20783-1197

3 DIRECTOR
US ARMY RESEARCH LAB
AMSRD ARL CS IS T
2800 POWDER MILL RD
ADELPHI MD 20783-1197

NO. OF
COPIES ORGANIZATION

ABERDEEN PROVING GROUND

1 DIR USARL
AMSRD ARL CI OK TP (BLDG 4600)

NO. OF
COPIES ORGANIZATION

ABERDEEN PROVING GROUND

25 DIR USARL
AMSRD ARL CI CT
M C LOPEZ

