

# Diagnostics for Causes of Packet Loss in a High Performance Data Transfer System

Phillip M. Dickens<sup>1,2</sup>, Jay. W. Larson<sup>2</sup>, and David M. Nicol<sup>3</sup>

<sup>1</sup>Department of Computer Science, Illinois Institute of Technology

<sup>2</sup>Mathematics and Computer Science Division, Argonne National Laboratory

<sup>3</sup>Department of Electrical and Computer Engineering and Coordinated Science Laboratory, University of Illinois Urbana-Champaign

## ABSTRACT

*As computational Grids become an increasingly dominant force in the high-performance computing arena, the problem of efficiently transferring very large data sets, across geographically distributed computing resources, becomes increasingly difficult and important. Current approaches view the problem largely, if not exclusively, as a network-level problem. Thus all packet loss is interpreted and treated as a network congestion event, limiting the ability to detect or react to changes in the end-to-end system. We believe that a new approach to this problem is worth pursuing, and we are investigating techniques that can differentiate between data loss caused by contention in the network and loss caused by contention for shared CPU resources at the communication endpoints. The approach is to collect and analyze what we term packet-loss signatures that describe the patterns of packet-loss in the current transmission window. We analyze these signatures using Fourier analysis and symbolic dynamics, and present a simple set of experiments demonstrating the effectiveness of this approach. Our longer-term goal is to exploit such information in next-generation congestion control mechanisms.*

## 1 Introduction

Computational Grids create large-scale distributed systems by connecting geographically distributed computational/data-storage resources via high-performance networks. Such systems, which can harness and bring to bear tremendous computational resources on a single large-scale problem, are quickly becoming a dominant force in the high-performance computing arena. An important area of research in Grid computing is the development of high-performance communication mechanisms that can take full

advantage of the underlying bandwidth when system conditions permit, can back off in response to growing contention in the network, and can accurately distinguish between the two situations.

Current approaches view the data transfer problem largely, if not exclusively, as a network-level problem. That is, the information available to and leveraged by the transfer control mechanisms is generally limited to the current loss rate and some measure of historical loss rates. Arguably, this limited view of the problem space, coupled with the limited use of external feedback, makes it very difficult to recognize and adapt to changes in the state of the end-to-end system.

Our research is addressing the issue of identifying the root cause(s) of data loss as observed by a high-performance data transfer system during its execution. The goal is to develop *system-aware* control mechanisms that can use such information to formulate responses tailored to the particular set of conditions responsible for the loss. The approach is to analyze what we term *packet-loss signatures* (or *bitmaps*) that show the distribution (or pattern) of those packets that successfully traversed the end-to-end transmission path and those that were dropped. These signatures are collected by the receiver and delivered to the sender upon request. Thus the packet-loss signatures are essentially large selective-acknowledgment packets, and are so named based on our belief (supported by experimental studies) that different classes of error mechanisms have different “signatures.”

We are exploring the application of complexity theory and Fourier analysis to the problem of learning the underlying structure (or lack thereof) of these signatures, and studying the relationship between such underlying structure and the system conditions responsible

for its generation. It is our view, supported by experimental studies provided below, that such information can provide significant insight into whether observed packet loss is, in fact, due to contention for network resources or whether it is due to conditions outside of the network domain. This leads to the notion that not all data loss is created equal (i.e., not equally harmful), and opens the door to control mechanisms that can more accurately determine the maximum rate at which the data transfer may proceed without negatively affecting the overall system.

The testbed for this research is FOBS (Fast Object-Based data transfer System), a high-performance data transfer system for computational Grids [7, 8]. FOBS is a UDP-based transfer system that provides reliability through a selective-acknowledgment and retransmission mechanism. As noted above, it is precisely the information contained within the selective-acknowledgment packets that is collected and analyzed by our classification mechanism.

Three important factors, whose combination is unique among high-performance data transfer mechanisms for computational Grids, makes FOBS an excellent testbed for this research. First, FOBS is an application-level protocol. Thus the congestion control algorithms can collect, synthesize, and leverage information from a higher-level view than is possible when operating at the kernel level. Second, the complexity measures can be obtained as a function of a *constant* sending rate. Thus the values of the variables collected are (largely) unaffected by the behavior of the algorithm itself. Third, FOBS is structured as a feedback control system. Thus the external data (e.g., the complexity measures) can be (but is not currently) analyzed at each control point, and this data can be used to determine the duration of the next control interval and the rate at which data will be placed onto the network during this interval. We do not discuss further the design, implementation, or performance of FOBS here. The interested reader is directed to [7, 8] for detailed discussions of these issues.

There are two important contributions of this paper. First, and to the best of our knowledge, this is the first paper to outline mechanisms by which data loss can be classified as being either internal or external to the network (in a high-performance wired network). Second, the techniques outlined here are generally applicable to UDP-based data transfer systems and, in fact, can be employed by TCP when using the selective-acknowledgment

mechanism [15]. The goal of this work is to use, in a sophisticated feedback-control mechanism for large-scale data transfers, the information made available by the application of these techniques. However, this paper focuses on the ability to classify the root causes of data loss and leaves the inclusion of these techniques into the congestion-control mechanisms of FOBS as a focus of current research.

The rest of the paper is organized as follows. In Section 2, we describe the Fourier spectral and symbolic dynamics analysis techniques used in this paper and discuss their application to the packet-loss signatures. In Section 3, we describe a set of experiments designed to test the effectiveness of our analysis strategy, and the results of these experiments are presented in Section 4. In Section 5, we discuss related work, and in Section 6, state our conclusions and outline future work.

## 2 Diagnostic Methodology

The packet-loss signatures described in Section 2 can be analyzed as time series data. Our objective is the identification of useful diagnostics that may be used to characterize causes of packet loss. The long-term plan is to apply these diagnostics in a feedback-control system to improve the overall efficiency of the data transfer mechanism. A desirable attribute of a diagnostic is that it can describe the dynamical structure of the time series. One classic approach is *Fourier analysis*, in which time series data are reduced to a linear superposition of their constituent frequencies in the time domain. Another approach is the application of *symbolic dynamics* techniques, which have been developed by the nonlinear dynamics community, and are highly appropriate for time series of discrete data.

We believe that the approach of time series analysis of the packet loss signatures may be useful in distinguishing between network-based and host-based causes of packet loss due to the differing timescales over which such losses occur. Packet loss due primarily to network-based causes such as router contention or contention at the NIC is likely to show temporal structure over a wide variety of timescales reaching down to the spacing between packets. A platform-based cause such as CPU contention at the host upon which the data receiver is executing will more likely be associated with a narrower range of longer timescales such as the

size of the time slice in a time-sharing operating system.

## 2.1 Fourier Analysis

We may view a packet loss signature (or a sequence of bitmaps corresponding to a longer transmission sequence) as a time-series, and apply standard techniques such as the Fast Fourier Transform (FFT) [13] to look for underlying structure. Its strength lies in being able to reveal the whole power spectrum of the sequence, including low-frequency behavior corresponding to subsequences well beyond the practical limitations of the word length approach discussed in Section 3.2. The FFT has proven invaluable in a large number of contexts in science and engineering as a means of characterizing underlying system structure.

In this approach, the packet loss signature is a discrete time series of real numbers  $f_0, f_1, \dots, f_{N-1}$ , with each element valued at 1.0 or 0.0 for each 1 or 0 in the bitmap, respectively. The FFT transforms a sequence of  $N$  time samples into a different sequence of  $N$  frequency weights  $F_n$ :

$$F_k = \sum_{n=0}^{N-1} f_n \exp(2\pi i k n / N). \quad (1)$$

These transformed values are coefficients of an orthogonal function expansion approximation to the original sequence:

$$f_k = (1/N) \sum_{n=0}^{N-1} F_n \exp(-2\pi i k n / N). \quad (2)$$

Each coefficient  $F_n$  in the resulting FFT is a complex number and is the weight associated with the  $n^{\text{th}}$  harmonic of the basic frequency used in the approximation of the time series. The analyses we present in this paper use the modulus of  $F_n$ , which corresponds to the power present in the  $n^{\text{th}}$  harmonic.

## 2.2 Symbolic Dynamics

Another approach to the analysis of the packet-loss signatures is the use of techniques from symbolic dynamics [11]. In this view, the bitmap is a sequence of symbols drawn from a finite discrete set, in our case two symbols: 1 and 0. One diagnostic that quantifies the amount of structure in the sequence is *complexity*. There

are numerous ways to quantify complexity. In this discussion, we have chosen the approach of d'Alessandro and Politi [6] who define a hierarchical approach to measures of complexity that will be discussed in detail below. This novel approach has been applied with success to quantify the complexity and predictability of time series of hourly precipitation data [9].

The approach taken in [9] is to view the stream of 1s and 0s as a language and focus on subsequences (or *words*) of length  $\mathbf{n}$  in the limit of increasing values of  $\mathbf{n}$  (i.e., increasing word length). First-order complexity, denoted by  $C^1$ , is a measure of the richness of the language's vocabulary and represents the asymptotic growth rate of the number of *admissible words* of fixed length  $\mathbf{n}$  occurring within the string as  $\mathbf{n}$  becomes large. The number of admissible words of length  $\mathbf{n}$ , denoted by  $\mathbf{Na}(\mathbf{n})$ , is simply a count of the number of distinct words of length  $\mathbf{n}$  found in the given sequence. For example, the string **0010100** has  $\mathbf{Na}(\mathbf{1}) = 2$  (0,1),  $\mathbf{Na}(\mathbf{2}) = 3$  (00,01,10),  $\mathbf{Na}(\mathbf{3}) = 4$  (001, 010, 101, 100). The *first-order complexity* ( $C^1$ ) is defined as

$$C^1 = \lim_{n \rightarrow \infty} (\log_2 \mathbf{Na}(n)) / n. \quad (3)$$

The first-order complexity can be interpreted in terms of the types of behavior associated with the high and low values limits of the values it can take. A constant string will have  $C^1=0$  and a purely random string will have a value of  $C^1=1$ . A purely periodic string, or one comprising only a few periodic sequences will tend to have low values of  $C^1$ .

The *second-order complexity* ( $C^2$ ) complements  $C^1$  and represents the richness of the grammar of the language: that is, the number of rules constraining the combinations of symbols. For word length  $\mathbf{n}$ , this constraint is the number of *forbidden words*,  $\mathbf{N}_f(\mathbf{n})$ , or the number of distinct symbol combinations of length  $\mathbf{n}$  that do *not* appear in the string. For example, in the string considered above,  $\mathbf{N}_f(1) = 0$ , since all possible combinations of word length 1 (i.e., 0, 1) appear in the sequence.  $\mathbf{N}_f(2) = 1$ , since the word 11 does not appear in the string, and  $\mathbf{N}_f(3) = 4$  (000, 011, 110, 111). A more specific constraint is the number of *irreducible forbidden words* of length  $\mathbf{n}$ ,  $\mathbf{N}_{if}(\mathbf{n})$ , where the qualifier irreducible means that the forbidden word contains no shorter irreducible forbidden word. Thus  $\mathbf{N}_{if}(2) = 1$  (since 11 contains only admissible words), and  $\mathbf{N}_{if}(3) = 1$  (forbidden words 011, 110, and 111 all contain the

irreducible forbidden word 11, whereas 000 contains no irreducible forbidden words). Second-order complexity is defined as

$$C^2 = \lim_{n \rightarrow \infty} (\log_2 N_{\text{if}}(n)) / n. \quad (4)$$

Intuitively, this value represents the presence of underlying rules in the process generating the symbol string, where the rules prevent certain combinations of symbols from appearing in the string. Infinite strings generated by a completely random process should have no forbidden words and thus will have a second-order complexity value of 0. Strings generated by a process with rules governing the creation of the string should result in a second-order complexity measure greater than 0.

We have presented definitions of both  $C^1$  and  $C^2$  for the sake of completeness. In the analyses presented here, we have found  $C^1$  to be of greater immediate use than  $C^2$ .

### 3 Experimental Design

We performed two sets of experiments to evaluate the effectiveness of our approach. In one set, we compared packet-loss signatures generated when data loss was caused by contention for NIC resources and when the data loss was caused by contention for CPU cycles. The second set of experiments compared the packet-loss signatures generated when data loss was caused by contention at a router and when the data loss was caused by contention for CPU cycles at the host upon the receiver was executing.

All data transfers were between hosts at Argonne National Laboratory (ANL) and the National Center for Supercomputing Applications (NCSA). The host platform at ANL (Chiba City), was a Linux cluster with 256 dual Pentium III 500 MHz processors. The computational platform at NCSA (Titan) was an IA-64 Linux cluster consisting of 128 compute nodes each consisting of dual Intel 800 MHz Itanium 1 processors. The two sites are connected by the Illinois Wired/Wireless Infrastructure for Research and Education (I-WIRE) which operates at 10 Gbps. There was no discernable contention on the I-WIRE at the time these experiments were conducted.

In the first set of experiments the data receiver executed on a dedicated processor within Chiba City, and additional compute-bound processes were spawned on this same

processor to create CPU contention. As the number of additional processes increased, the amount of time the data receiver was switched out similarly increased. Since the data receiver was not available to take packets off of the network during the times it was switched-out, there was a direct relationship between CPU load and the resulting packet loss rate. We were interested in analyzing the structure of the bitmaps as a function of both the root cause of data loss (i.e., contention for CPU or NIC resources) and the loss-rate. We therefore varied the number of additional processes to obtain loss rates of (approximately) 1%, 5%, 10%, and 15%.<sup>1</sup>

To investigate loss patterns caused by contention for NIC resources, a second (background) data transfer was initiated. The data sender of the background transfer executed on a different node within Chiba City, and the receiver executed on the second processor within the same node as the primary data receiver. Since both processors of a given node share the same NIC, we were able to generate contention at the NIC without causing contention for CPU cycles with the two receivers. Initially, the combined sending rate was set to the maximum speed of the NIC (100 Mbps for the Chiba City compute nodes), and contention for NIC resources was increased by increasing the sending rate of the background transfer. The packet loss experienced by both data transfers was a function of the combined sending rate, and this rate was also set to result in loss-rates of (approximately) 1%, 5%, 10%, and 15%.<sup>2</sup>

In the second set of experiments, we used nine parallel UDP data streams (each sending at 100 Mbps) between Titan and Chiba City. We then created a 10<sup>th</sup> data stream between a HP N4000 node at the Center for Advanced Computational Research (CACR) and a BM IntelliStation Z Pro 6894 workstation within NCSA. The 10<sup>th</sup> stream shared a router with the 9 parallel UDP streams creating contention for that router's resources. The sending rate of the 10<sup>th</sup> stream was varied between 50 and 100 Mbps. When data was sent at 100 Mbps packets

---

<sup>1</sup> The particular process-scheduling algorithm also has a direct impact on the relationship between the number of additional processes and the time the data receiver is switched out (e.g., the data receiver may be a higher priority process). In these experiments, however, the relationship was strong enough to allow us to generate loss rates very close to the target rates.

<sup>2</sup> The results presented represent the loss rates and packet-loss signatures of the primary data receiver.

were dropped at the router. When data was sent at 50 Mbps, the router was able to process all ten streams. The result of interest was a time-series of the packet-loss signatures of a randomly picked stream at the receiving host (NCSA). For comparison, we also conducted experiments resulting in a time-series of packet-loss signatures when data was lost due to CPU load.

## 4 Experimental Results

Fourier analyses of packet loss signatures for the CPU and NIC contention cases are presented in Figures 1 and 2, respectively. In both cases, the approximate packet loss rate was 5%. In both cases, the mode coefficients for the first 35000 of 70000 frequencies are plotted. The characteristics of these two graphs are common across loss rates and different experiments in the set analyzed.

It is evident that these graphs are visually very different. We should add that they are both different from the graphs one obtains from experiments where errors are due to contention in network routers (which are much more random, and have much less structure). The FFT of the packet loss signature from the CPU contention experiment is striking in its narrow spectrum dominated by low frequencies. The underlying bitmap has a few largish blocks (approximately 500 bits) of zeros separated by roughly equal sequences of ones. The strong periodicity of this pattern is reflected in the very high coefficient at the lower end of the FFT graph, with very little power elsewhere. The NIC contention experiments produce bitmaps showing numerous strong modes spread across the spectrum as evidenced by the numerous spikes in Figure 2. Of course, a sequence of 0s and 1s appears to the FFT as a series of square wave pulses of varying duration, which create discontinuities in the signal being analyzed. A good deal of the spikiness in the graph is due to the FFT approximating a sequence of largely periodic appearances of 0 in a sea of 1s with a main sine wave and its harmonics. The key thing to observe here is the strong pattern of spikes and their distribution across the frequency spectrum. The clear conclusion from these results is that the FFT can distinguish between two major sources of packet loss.

Complexity-based analyses of these experiments are shown in Figures 3 - 8. Each figure shows the mean first-order complexity measure (calculated from the five 350,000 bit

strings), and the 95% confidence intervals for the mean, for word sizes  $n = 4$  to  $n = 16$ . Figures 3 and 4 represent  $C^1$  for bitmaps from experiments in which packet-loss is due to pure CPU contention and pure NIC contention, respectively. Figures 5 - 8 compare these complexity measures for CPU versus NIC contention at each of the four loss rates tested.

As can be seen from Figure 3, the values of  $C^1$  decay very quickly with increasing word size for all loss rates tested. This result is consistent with the hypothesis that loss due to the data receiver being switched-out will result in packet-loss signatures that are basically periodic. We note that the values of  $C^1$  tend to converge as the loss rate (and word length) increases, approaching a value of around 0.32 for all loss rates at  $n = 16$ . This is important in that it provides evidence that the complexity of the packet-loss signatures, when such signatures are caused by contention for CPU cycles, is largely independent of the loss-rate.

As Figure 4 shows, the complexity values of the packet-loss signatures resulting from contention for NIC resources are significantly higher than those observed in Figure 3. Also, and in contrast to Figure 3, the values of  $C^1$  continue to increase with increasing loss rate. This result is consistent with the hypothesis that packet-loss signatures generated by contention for NIC resources are significantly more complex than those generated when the loss is caused by preemption of the data receiver. This also shows that the "signal" associated with NIC contention becomes increasingly strong with increasing loss-rate.

The distinction in packet-loss signatures is even more evident in Figures 5-8, which compare the values of  $C^1$  as a function of the loss-rate. As can be seen, the complexity of the signatures is clearly distinguishable even at very low loss rates, and becomes more pronounced as the loss rate increases. Thus some experimentation may be required to learn the complexity values resulting from the two types of loss when the loss rate is very low. However, these results suggest that the loss rate does not have to become large before such a distinction becomes obvious.

Figures 9-10 show the values of  $C^1$  resulting from contention for router resources with those caused by contention for CPU resources. The disparity in values of  $C^1$  depicted in Figure 9 is quite striking. The complexity measures associated with contention for resources at the router are cyclic, and track quite well the gradual

filling and draining of the router buffers. In the case of the router, the loss rate fluctuated between 0% and 2.5%. The loss rates associated with CPU contention fluctuated between approximately 1% and 3% with little fluctuation in the values of  $C^1$ . This is another clear demonstration of how different causes of packet loss can produce significant differences in the underlying structure of the packet-loss signatures.

Figure 10 depicts the relationship between the loss rate and the corresponding values of  $C^1$  as a function of the circumstances under which the data was lost. As can be seen, the values of  $C^1$  are significantly higher (across all loss rates) when the data was lost due to contention in the network as opposed to contention at the receiving host. This again demonstrates that the structure of packet-loss signatures is quite sensitive to the root cause of the data loss even when the loss rate is quite low.

## 5 Related Work

It has been well established that TCP was not designed for and does not perform well in the high-bandwidth, high-delay network environments typical of computational Grids (see, for example, [12, 4, 10]). This paper advocates a different approach based on two components: gathering information about the current system dynamics using complexity analysis and information theory, and using such information in highly intelligent and adaptable application-level control mechanisms. Another approach is the development of mechanisms to improve the performance of the TCP protocol itself in this network environment (see, for example, [2, 16, 15]). There is also significant interest in the exploration of user-level techniques that can circumvent some of the problems inherent within TCP [12, 1, 17, 19].

Research aimed at modifying TCP for hybrid wired/wireless networks provides strong support to our belief that identifying the root cause of packet loss can significantly improve performance. The problem statement is similar in the sense that unmodified TCP cannot distinguish packet loss caused by network contention from packet loss caused by errors in the wireless network link. Thus TCP's aggressive congestion control mechanisms are triggered even though the loss is not a true indicator of network congestion, and the unnecessary reductions in bandwidth utilization results in significantly degraded performance

[3, 14].

## 6 Conclusions and Future Research

In this paper, we have described a strategy of analyzing packet-loss signatures from a high-speed data transfer mechanism, and how this strategy enables classification of the dominant cause of packet loss in the current transmission window. These techniques are based on both traditional Fourier analysis and first-order complexity measures (introduced for the first time in this paper) of packet-loss signatures to determine the root cause of packet loss. We outlined a series of simple experiments to test the efficacy of these techniques, demonstrating they are easily capable of distinguishing packet loss caused purely by CPU contention or NIC contention. We have also been successful in detecting network contention using first-order complexity analysis. In actual Grid settings packet loss will likely be caused by a combination of factors, and the resulting signals from the complexity measures will be harder to discern. However, the results presented here provide strong evidence that contention for network resources at the communication endpoints can be detected even when such contention is quite low. This useful information can be used by the control mechanism to identify conditions under which it may be very damaging to increase the sending rate. What is not clear, and is the focus of current research, is whether there can be significant contention within the wide area network(s) connecting the communication endpoints that can (or likely to) produce complexity measures that appear to represent periodic behavior. This issue is strongly related to the queuing discipline used by the routers in the end-to-end path, and is being investigated through experimental and simulation studies.

Given the success and shortcomings of the analysis techniques reported here, we are planning to dig more deeply into the symbolic dynamics of the packet loss signatures. An immediate area of interest is investigation of the usefulness of the second-order complexity  $C^2$  in our classification scheme. Beyond hierarchical measures of symbol stream complexity, there exist even more powerful techniques within the realm of *information theory* [18]. In particular, levels of *entropy convergence* [5] provide a set of quantities that promise to complement and expand the complexity measures  $C^1$  and  $C^2$  introduced above, offering the potential to detect

hidden structure, quantify periodicity, and form the foundation for prediction schemes.

The ability to classify the temporal dynamics of packet loss behavior (as expressed by the packet-loss signatures) offers two significant advantages. First, such classification allows the control mechanisms to apply corrective actions based on the particular cause of packet loss. For example, the control mechanisms may be able to migrate the data receiver, rather than drastically reducing the sending rate, when the root cause of packet loss is determined to be contention for CPU (rather than network) resources. Second, if the underlying dynamics has structure, it may be possible to construct simple predictors that allow the data transmitter to shape its behavior in such a way as to increase the probability that a sent packet is received successfully. These are enticing possibilities, and the exploration, evaluation, and integration of these techniques to the problem of large-scale data transfers represents a focus of future research activities.

A longer-term issue of interest is the eventual use of these schemes operationally in a feedback-control system for data transfer mechanisms. There will no doubt be trade-offs between comprehensive data collection and analysis and overall speed. Efficient implementations of the analysis schemes will be required. The FFT is already renowned for its speed, being an  $O(N\log_2 N)$  algorithm. The symbolic techniques we outlined here are currently implemented in a simple but inefficient fashion. Even with no optimizations, however, the value of  $C^1$  can be computed in a matter of milliseconds for word lengths  $n \leq 10$ . Symbolic techniques are popular in the chaos community because they lend themselves well to efficient implementations, and a faster algorithm can be built with relative ease. We note without elaboration that recent work in information theory as discussed in [5] may be quite helpful in defining a word length stopping criterion for the calculation of  $C^1$ .

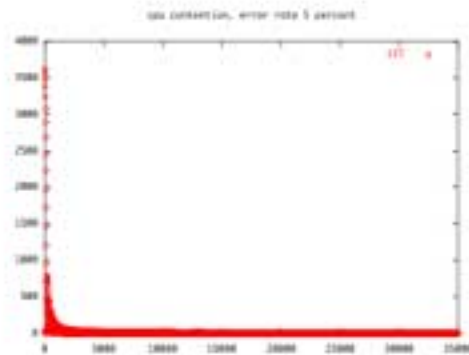
## References:

- [1] Allcock, W., Bester, J., Breshahan, J., Chervenak, A., Foster, I., Kesselman, C., Meder, S., Nefedova, V., Quesnel, D., and Tuecke, S. Secure, Efficient Data Transport and Replica Management for High-Performance Data\_Intensive Computing. In *Proceedings of IEEE Mass Storage Conference*, 2001.
- [2] *Automatic TCP Window Tuning and Applications*. National Laboratory for Advanced Networking Research Web Page [http://dast.nlanr.net/Projects/Autobuf\\_v1.0/autotcp.html](http://dast.nlanr.net/Projects/Autobuf_v1.0/autotcp.html)
- [3] Balakrishnan, S., Seshan, S., Amir, E., and Katz, R. Improving TCP/IP performance over wireless networks. In *Proceedings of ACM MOBICON*, November 1995.
- [4] Boyd, E.L., Brett, G., Hobby, R., Jun, J., Shih, C., Vedantham, R., and Zekauska, M. *E2E piPEline: End-to-End Performance Initiative Performance Environment System Architecture*. July, 2002. <http://e2epi.internet2.edu/e2epipe11.shtml>
- [5] Crutchfield, J., and Feldman, D. Regularities Unseen, Randomness Observed: Levels of Entropy Convergence, Santa Fe INstitute Working Paper 01-02-012, 2001.
- [6] D'Alessandro, G., and Politi, A. Hierarchical Approach to Complexity with Applications to Dynamical Systems. *Physical Review Letters*, 64 (14). 1609-1612. April, 1990
- [7] Dickens, P. FOBS: A Lightweight Communication Protocol for Grid Computing. In *Proceedings of Europar 2003*, 2003.
- [8] Dickens, P., and Gropp, B. An Evaluation of Object-Based Data Transfers Across High Performance High Delay Networks. In *Proceedings of the 11th Conference on High Performance Distributed Computing*, Edinburgh, Scotland, 2002.
- [9] Elsner, J., and Tsonis, A. Complexity and Predictability of Hourly Precipitation. *Journal of the Atmospheric Sciences*, 50 ((3)). 400-405. 1993
- [10] Feng, W., and Tinnakornsrisuphap, P. The Failure of TCP in High-Performance Computational Grids. In *Proceedings of Proceedings of Super Computing 2000 (SC2000)*.
- [11] Hao, B.-l. *Elementary Symbolic Dynamics and Chaos in Dissipative Systems*. World Scientific, 1989.
- [12] *Home Page of The Web100 Project* <http://www.web100.org/>
- [13] Korner, T.W. *Fourier Analysis*. Cambridge University Press, New York, 1988.
- [14] Liu, J., Matta, I., and Crovella, M. End-to-End Inference of Loss Nature in a Hybrid Wired/Wireless Environment. In *Proceedings of Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'03)*, Sophia-Antipolis, France, 2003.
- [15] Mathis, M., Mahdavi, J., Floyd, S., and Romanow, A. TCP Selective


- Acknowledgement Options, RFC 2018.  
 RFC 2018.
- [16] *Modifying TCP's Congestion Control for High Speeds* May, 2002  
<http://www.aciri.org/floyd/>
- [17] Ostermann, S., Allman, M., and Kruse., H. An Application-Level solution to TCP's Satellite Inefficiencies. In *Proceedings of Workshop on Satellite-based Information Services (WOSBIS)*, November, 1996.

- [18] Shannon, C. A mathematical theory of communication. *Bell System Technical Journal*, 27. 379-423.
- [19] Sivakumar, H., Mazzucco, M., Zhang, Q., and Grossman, R. Simple Available Bandwidth Utilization Library for High Speed Wide Area Networks. *Submitted to Journal of SuperComputing*.


**Figures:**



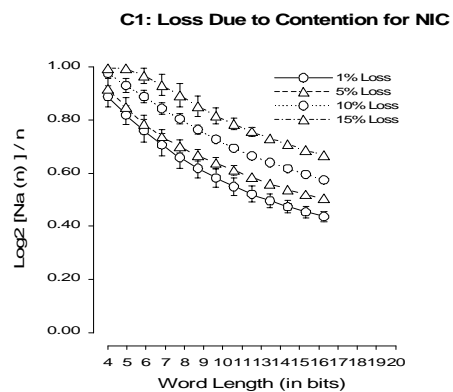
**Figure 1. Fourier power spectrum for a packet-loss signature from a CPU contention experiment with a packet loss rate of 5%.**



**Figure 2. Fourier power spectrum for a packet-loss signature from a NIC contention experiment with a packet loss rate of 5%.**



**Figure 3. This figure shows the complexity values obtained when the data loss was caused by contention for CPU resources. These measures were obtained for loss rates of 1%, 5%, 10%, and 15%. The complexity values were calculated for word sizes n = 4 to n = 16.**



**Figure 4. This figure shows the complexity values obtained when the data loss was caused by contention for NIC resources. These measures were obtained for loss rates of 1%, 5%, 10%, and 15%. The complexity values were calculated for word sizes n = 4 to n = 16.**






Figure 5. Complexity measures when loss is generated by contention for CPU resources versus contention for NIC resources when the loss rate is 1%.

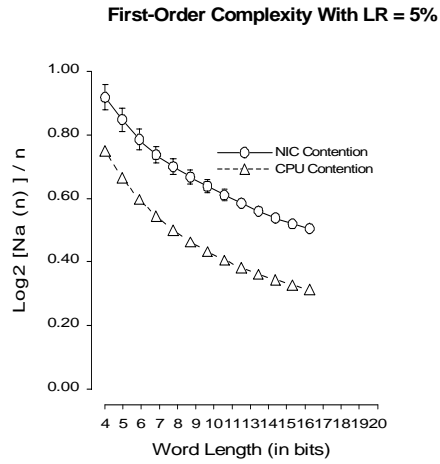


Figure 6. Complexity measures when loss is generated by contention for CPU resources versus contention for NIC resources when the loss rate is 5%.




Figure 7. Complexity measures when loss is generated by contention for CPU resources versus contention for NIC resources when the loss rate is 10%.

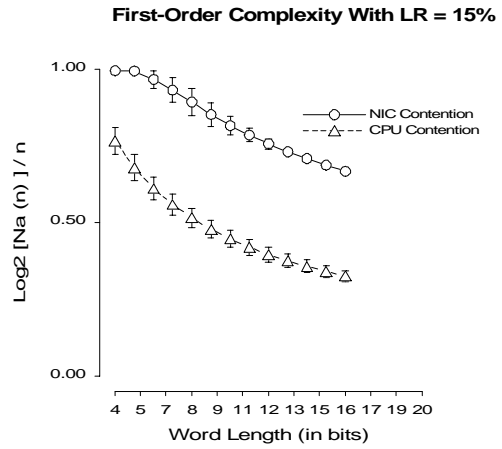
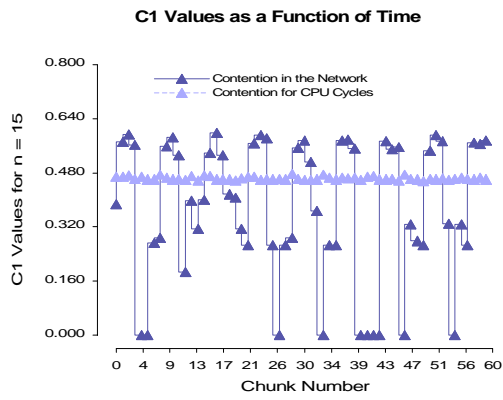



Figure 8. Complexity measures when loss is generated by contention for CPU resources versus contention for NIC resources when the loss rate is 15%.



**Figure 9.** C1 values obtained when data loss was caused by contention for network resources and when it was caused by contention for CPU resources.



**Figure 10.** C1 values shown as a function of loss rate.