Time-Series Investigation of Anomalous CRC Error Patterns in Fibre Channel Arbitrated Loops

Kenny C. Gross, Wendy Lu, and Dawei Huang
RAS Computer Analysis Laboratory
Sun Microsystems, Inc.

Abstract

A study was undertaken to investigate the origin of anomalous cyclic redundancy code (CRC) error measurements in an instrumented Fibre Channel (FC)-based Storage Area Network (SAN). Machine learning algorithms are being developed for fault detection and identification in FC loops. Macros that implement univariate and bivariate spectral decomposition techniques were employed to analyze time-series data recorded during a sequence of experiments conducted in a high-bandwidth laboratory SAN configuration. Digitized thermocouple readings from within the storage configuration were recorded in parallel with the CRC counts and were archived for off-line analysis. Results from spectral and cross-correlation analyses led to the identification of a nonlinear, temperature-dependent phenomenon that produces serial correlation in the stochastic CRC counts, impacting false-alarm probabilities for proactive fault monitoring of enterprise computing networks. A major finding from this investigation is that the efficiency of machine learning techniques, in terms of sensitivity, robustness, and false alarm avoidance, can be enhanced by use of methods that accommodate serial correlation in CRC time series.

Keywords: real-time fault monitoring; online defect identification; process anomaly detection

1. Introduction: Fibre Channel Storage Area Networks

In recent years, several technical developments have converged to drive enterprise needs for extremely fast data links. High performance computers have become the focus of much attention in the data communications industry. Performance improvements have spawned increasingly data-intensive and high-speed networking applications. These emerging applications include professional movie editing, scientific visualization, video-on-demand, large scale voice-over-IP, and medical imaging. Fibre Channel (FC) is the general name of an integrated set of standards developed by the American National Standards Institute (ANSI). The purpose of fibre channel storage area networks (SANs) is to provide practical and inexpensive means of quickly transferring data between supercomputers, mainframes, workstations, desktop computers, storage devices, displays, and other peripherals.

In fibre channel SANs, the data flows between hardware entities called N_ports, generally part of the termination card that contains hardware and software support for the fibre channel protocol. Each node must have at least one N_port, and usually has two; either may serve as a transmitter, receiver, or both. Fibre channel allows for an active intelligent interconnection scheme, called a Fabric, to connect devices. The Fabric is most often a switch, or a series of switches, that takes the responsibility for routing. The topology of the fibre channel network is transparent to the attached devices. If traditional networking components, such as switches and hubs, are added to the storage infrastructure, then multiple servers can be connected with multiple storage devices. Fibre channel's flexibility allows for the introduction of a much more modular building block architecture type that can be fine-tuned toward specific applications. These features facilitate design and future scalability of FC networks, but bring challenges in terms of machine-learning paradigms for proactive health monitoring.

2. Machine Learning Challenges

Although modularity and flexibility are virtues for FC-based network architectures, they present challenges for health monitoring agents designed to assure continuous availability and predictable performance for end users. These challenges have prompted the adoption of machine-learning best practices with the following key objectives:

- Identify hard failures and generate alerts as soon as they occur
- Minimize downtime for network users
- Maximize accuracy of root cause diagnostics
- Allow an acceptable level of soft errors before triggering a diagnosis/repair action (soft error rate discrimination, SERD)

- Minimize false alarms, which contribute to customer dissatisfaction as well as warranty costs for the manufacturer

Error detection in FC-based networks is performed by monitoring a 32-bit Cyclic Redundancy Check (CRC) parameter. Changes in CRC count frequencies can signify the incipience or onset of physical disturbances in network hardware, lasers, optical interconnects, or degrading component or cable connectors. For the experiments reported here, an experimental Fibre Channel Arbitrated Loop (FCAL) network was set up in a laboratory setting. High volumes of traffic were driven through the network in the form of random I/O to disk arrays. An optical attenuator with a controllable attenuation factor accelerated CRC error counts so that machine-learning algorithms for network health monitoring could be developed and evaluated. Evaluation criteria for network health monitoring algorithms include the missed- and false-alarm probabilities, and time-to-annunciation for alarms signifying subtle degradation in the FC network.

During early operation with the FCAL network, anomalous patterns were observed in the CRC error-count time series. Univariate and bivariate spectral decomposition techniques [1,2] were exploited to explore and characterize the anomalous CRC patterns, and to test hypotheses advanced to explain their origin. A bivariate spectral analysis method called the normalized cross power spectral density (NCPSD) method [3] was instrumental in proving that the anomalous patterns in CRC time series were due to fluctuations in ambient temperature from datacenter HVAC cycling. The macros developed as part of this investigation are quite general and can be applied with only minor modifications to explore and characterize any experimental time series that are suspected to be contaminated with periodic components.

3. CRC Counts from FCAL Loops

The health of FC devices directly impacts system availability and customer satisfaction. In the telecommunications industry, equipment manufacturers and service providers have been successfully applying machine-learning best practices for proactive health monitoring of networks to ensure quality-of-service metrics, and to provide failure annunciation in a variety of ways. One such method that has proven effective is to monitor bit error rates (BER) to accurately measure the quality of the optical transmission. In practice, current FCALs rely on CRC counts, reported through the link error status block register in every FC device, for assurance that the BER is acceptably low. Since the CRC count is equivalent to a bit error count at low BER, machine–learning methods may be applied to patterns of CRC counts as a function of I/O intensity for continuous health-check surveillance of FC-based optical networks.

A standard measure of the quality of transmission throughout FC networks is the optical signal-to-noise ratio (SNR). With the optical SNR, the power of each channel is measured relative to the broad background noise that is largely attributable to amplified spontaneous emission (ASE).

A sequence of experiments was performed using an optical attenuator that allowed us to systematically vary the SNR and collect BER counts. During the attenuation experiments, high volumes of random network traffic were generated by performing random I/O to an array of 44 9-Gbyte disk drives. The network-attached storage schematic used for these experiments is shown in Figure 1.

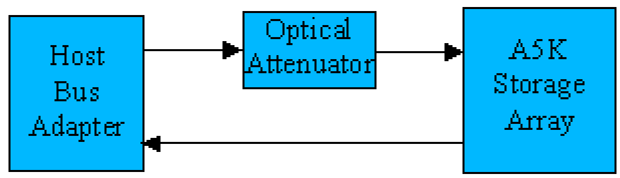

Figure 1. Component schematic for FCAL monitoring demonstration

By treating the optical-attenuation coefficient as an experimental control variable, we can systematically decrease the SNR at the receiver, thereby increasing the system BER in a sequence of parametric sensitivity experiments. Typical experimental results, collected over a several-week testing sequence using controlled attenuation levels of 0 dB to 1.2 dB, are plotted in Figure 2.

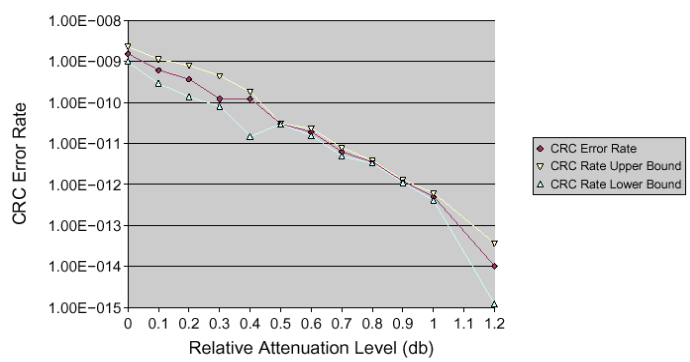

Since the network traffic was stationary (in the stochastic sense), CRC error counts were expected on the basis of theoretical considerations to constitute a white-noise time series process. Typical results for one of the experiments conducted with an attenuation level of 0.6 dB are plotted in Figure 3. The CRC time series plotted in Fig. 3 have an appearance that would suggest contamination by serial correlation. To quantitatively substantiate or refute this hypothesis, the Kolmogorov-Smirnov test was employed. The Kolmogorov-Smirnov whiteness test evaluates the hypothesis that the series of observations is random in time (i.e. a white noise process), vs. a hypothesis that the time series is contaminated by some type of serial correlation.

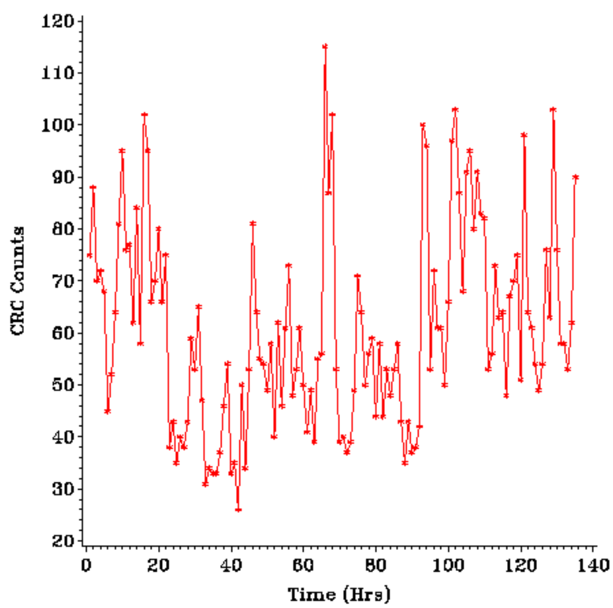

Figure 3. Time series plot of fibre channel CRC counts

For the data collected in the 0.6 dB attenuation experiment, the Kolmogorov-Smirnov statistic computed from the digitized counts is 0.220. The 99% critical value for the K-S test using the data plotted in Fig. 2 is 0.12. Since 0.220 is (significantly) greater than 0.12, the whiteness hypothesis can be rejected with at least a 99% confidence factor.

One hypothesis advanced to explain the anomalous behavior of the CRC counts is that changes in ambient temperature in the datacenter could be producing small changes in the signal-to-noise ratio for the FCAL optical interconnects. In view of the strong, non-linear dependence between CRC counts and SNR exhibited in the experimental data plotted in Fig. 2, it is quite plausible that subtle changes in temperature could be causing measurable changes in CRC rates.

To explore this hypothesis, digitized temperature signals were extracted from the array of network-attached disk drives. Although these temperature signals were measured from inside the drive array, the drives were cooled with ambient air, and would therefore be expected to reflect subtle changes in ambient air temperature. Temperatures from all drives were averaged at each sampling interval to produce one mean temperature reading for the array. Results of the digitized temperature time series are plotted in Figure 4.

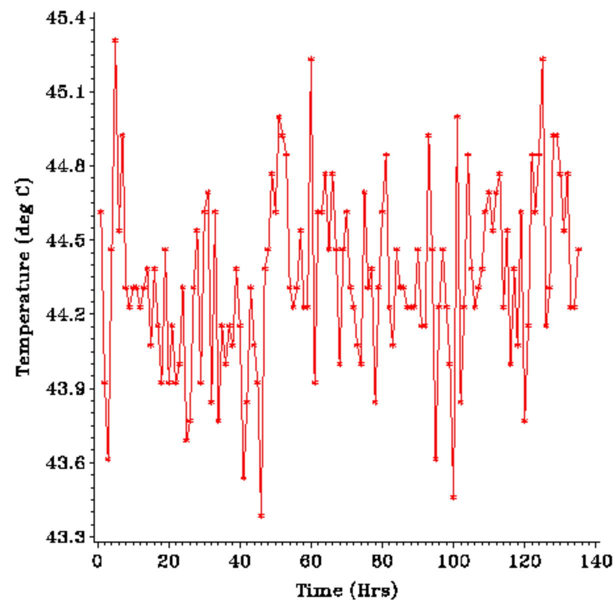

Figure 4. Drive array internal temperature signals

4. Univariate Spectral Decomposition Analyses

Univariate spectral-decomposition analysis was performed on the raw data recorded during steady-state experiments that were conducted at discrete attenuation settings. Output from the spectral calculations is provided in the form of periodograms and spectrograms, which plot the power spectral density (PSD) for the time series as a function of period, P, or frequency, w, respectively. A physical interpretation of the PSD function is that f(w)dw represents the contribution to variance of components with frequencies in the range (w,w+dw). When the spectrum is plotted, the total area under the curve is equal to the variance of the signal being analyzed. A peak in the spectrum indicates an important contribution to variance at frequencies in the appropriate region.

Figure 5 plots the PSD computed from the temperature time series observations. Note that the PSD of a white noise process is also a white noise process [3]. The PSD plotted in Fig. 5 is clearly non-white. The spectrogram reveals the presence of a strong periodicity at approximately 3 hrs, with a harmonic at approximately 6 hrs. The strong periodicities at approximately 3 hrs and 6 hrs were not expected at the outset of our experiments, and have subsequently been attributed to a subtle variation in datacenter ambient temperature from HVAC cycling. The fortuitous periodicity in datacenter ambient temperature provides a valuable exploratory diagnostic tool that has enabled us to exploit the power and sensitivity of bivariate spectral methods to investigate the physical origin of the high degree of serial correlation in FCAL CRC count rates.

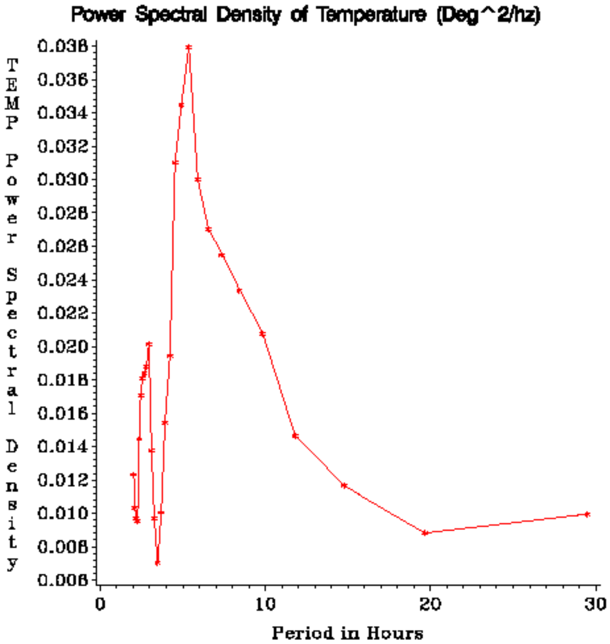

Figure 5. Spectral decomposition of temperature

Application of univariate spectral decomposition to the CRC count-rate time series demonstrates quite similar periodicities, as shown in Figure 6, and provides further corroboration of the non-whiteness findings from the Kolmogorov-Smirnov analysis.

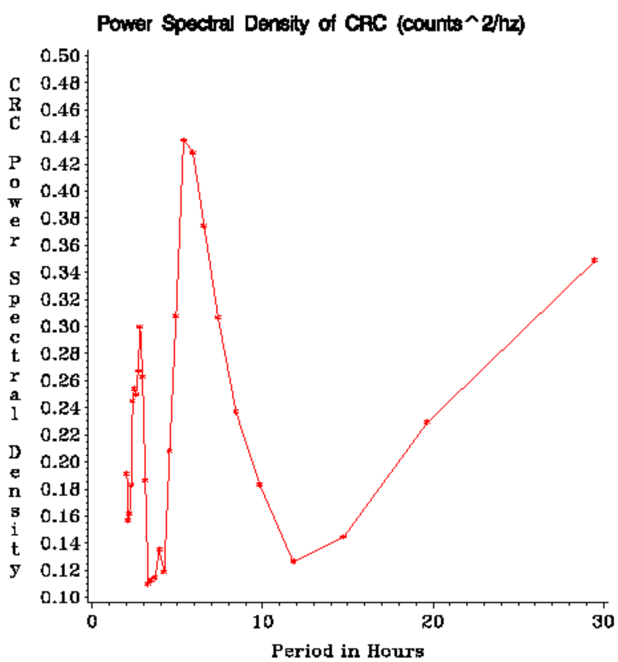

Figure 6. Spectral decomposition of CRC count rates

To quantitatively evaluate whether there is indeed a physical dependence of CRC counts on temperature, bivariate spectral-decomposition analyses were performed using the temperature and CRC signals as inputs. In bivariate analysis [4,5], the objective is to compute a cross power spectral density function, whose plot is similar in appearance and interpretation to the univariate PSD function used above. The normalized cross power spectral density (NCPSD) function computed from two time series essentially shows whether frequency components in one series are associated with large or small amplitudes at the same frequency in the other series.

Standard time-domain cross correlation between temperature and CRC counts was insufficient to establish a relationship between those variables. The correlation coefficients produced during the various attenuation experiments typically ranged from -0.1 to 0.1, due to the poor signal-to-random-noise performance for both of those stationary time series. Nevertheless, the NCPSD method has the capability to extract even very small common physical variations from processes contaminated by large random components. Results of the NCPSD analysis, plotted in Figure 7, show a strong concordance between temperature and CRC count rates with harmonic periodicities at approximately 3 and 6 hours.

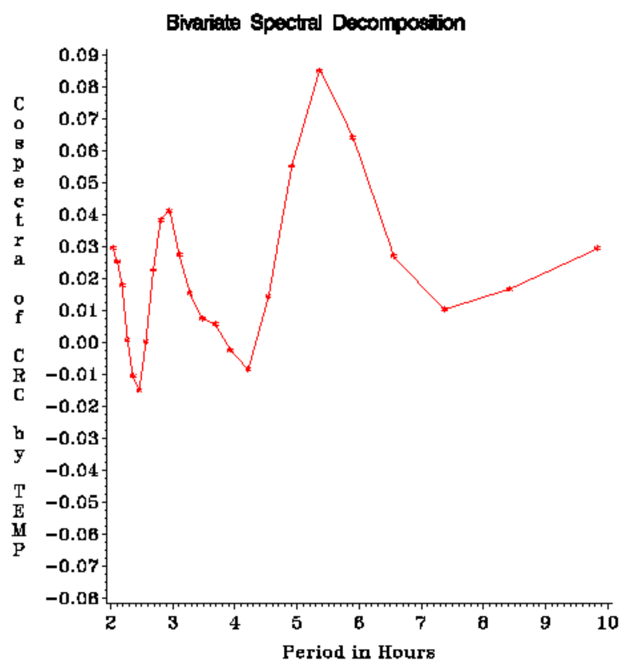

Figure 7. NCPSD analysis of temperature and CRC count rates

5. Summary and Conclusions

A detailed investigation of time-series data from a FCAL interconnected optical network-attached storage configuration was conducted using univariate and bivariate spectral decomposition to explore the physical origin of anomalous CRC response patterns. CRC count rates exhibit significant serial correlation, as demonstrated by Kolmogorov-Smirnov whiteness tests. The serial correlation components are a non-linear function of ambient temperature, as demonstrated by measurements of CRC rates that were measured in systematic parametric investigations as a function of signal-to-noise ratio. The presence of these non-linear, serially-correlated components can adversely impact the sensitivity and false-alarm avoidance of machine learning algorithms that are being separately designed for real-time health monitoring of FCAL networks [4]. For this reason, a major recommendation emerging from the research findings reported here is that real-time monitoring methods be employed that can accommodate serially correlated components superimposed on stationary time series [5], or that can learn relationships among multiple, correlated signals [6,7]. In a follow-on project launched at the conclusion of the study reported herein, Sun Microsystems is evaluating a variety of real-time machine-learning techniques that can autonomously learn the relationships between CRC counts for FCAL networks and ambient (ex-cabinet as well as localized, internal) temperatures for the purpose of enhancing the sensitivity, time-to-annunciation, and false-alarm avoidance for network health monitoring.

References

[1] K. C. Gross and L. K. Polley. "Response of EBR-II's DN Monitoring Systems During Sinusoidal Reactivity-Oscillation Experiments," ***Trans. Am. Nucl. Soc., 52***, 643 (June 1986).

[2] K. C. Gross and R. V. Strain. "Sinusoidal Source-Perturbation Experiments with a Breached Fuel Element in EBR-II," ***Nucl. Tech. (98)*** pp. 113-123 (Apr 1992).

[3] K. C. Gross and L. K. Polley. "Investigation of Non-recoil Fission-Product Release Phenomena Using Multifrequency Source-Perturbation Experiments in EBR-II," ***Annals of Nucl. Energy*** (Aug 1986).

[4] K. C. Gross and W. Lu. "Early Detection of Signal and Processing Anomalies in Enterprise Computing Systems," ***2002 IEEE International Conference on Machine Learning and Applications (ICMLA)***, Las Vegas, Nevada, June 24-27, 2002.

[5] K. C. Gross and K. K. Hoyer. "Spectrum-Transformed Sequential Testing Method for Signal Validation Applications," ***8th Power Plant Dynamics, Control & Testing Symp.*** Vol. I, pp. 36.01-36.12, Knoxville, TN, May 1992.

[6] K. C. Gross, R. M. Singer, S. W. Wegerich, J. P. Herzog, R. VanAlstine, and F. Bockhorst. "Application of a Model-based Fault Detection System to Nuclear Plant Signals," ***Proc. 9th Int'l. Conf. On Intelligent Systems.*** pp. 66-70, Seoul, Korea, July 6-10, 1997.

[7] K. Cassidy, K. C. Gross, and A. Malekpour. "Advanced Pattern Recognition for Detection of Complex Software Aging Phenomena in Online Transaction Processing Servers," ***Int'l Performance and Dependability Symposium***, Washington, DC, June 23-26, 2002.