Spectral Decomposition and Reconstruction of Telemetry Signals from Enterprise Computing Systems

Kenneth C. Gross and Eugenio Schuster

***Abstract*— The System Dynamics Characterization and Controls Laboratory of Sun Microsystems has developed novel continuous system telemetry technology for proactive fault monitoring and "electronic prognostics" capabilities for all types of executing computing platforms. Central to the goal of designing and evaluating prognostic and control systems for enterprise computing platforms is first the ability to simulate telemetric signals and computer system degradation events with real-time models in an interactive computer environment. This paper describes the development and implementation of a Telemetric Parameter Simulation System (TPSS) that employs Fourier-based decomposition and reconstruction methodology with the capability to efficiently decompose any signals into their deterministic and stochastic components, then reconstruct new, simulated signals that possess exactly the same statistical noise idiosyncrasies as the original telemetry variables. The TPSS has become an indispensable tool in Sun's ongoing development of innovative diagnostic techniques for telemetry operability surveillance and other electronic prognostic applications in enterprise in Sun's present and future computer platforms.**

I. Introduction

A System Dynamics Characterization and Control (SDCC) laboratory has been established as part of Sun Microsystems' Physical Sciences Research Center in San Diego, Ca. As part of this effort, new advanced simulation capabilities are being built upon an innovative distributed-processing computing system that is dedicated to at least four areas of SDCC support for enhancing the reliability, availability, serviceability (i.e. RAS), and quality-of-service (QoS) for present and future computing platforms: development of real-time server telemetry signal simulation for design, RAS, and QoS analysis; development and testing of server feedback-and-control algorithms; validation of control schemes prepared by universities and collaborating vendor partners for testing with the new server prototype platforms; and development and testing of advanced pattern recognition techniques to annunciate the incipience or onset of degradation of sensors and server components, field-replaceable-units (FRUs), and integrated systems. Central to the goal of designing and evaluating control and information systems for prototype server platforms is first the ability to simulate them with real-time models in an interactive, computer laboratory environment. In support of the ongoing development of the SDCC Laboratory, an ancillary effort has been undertaken to develop a Telemetry-Parameter Simulation System (TPSS). The TPSS has been designed to meet two key functional requirements:

(1) To be able to analyze any server telemetry signal, decompose that signal into its deterministic and stochastic components, then reconstruct a new, simulated signal that possesses exactly the same statistical noise idiosyncrasies (e.g. non- gaussian skewness, kurtosis, autocorrelation structure) as the actual server telemetry signal.
(2) To be able to filter out the principal serially-correlated, deterministic components from xerver telemetry variables so that the remaining stochastic signal (called the residual function) can be analyzed with signal validation tools that are designed for signals drawn from independent random distributions.

Functional requirement (1) addresses a capability that is lacking in current, state-of-the art system telemetry simulation systems: the ability to simulate and explore the effects of sensor degradation events. As is, conventional server modeling codes for modeling thermal and power dynamics in prototype server architectures are capable of providing high-fidelity prototype simulations that can enable users to demonstrate server behavior during a wide range of transient execution modes that might include load and memory dynamics, fan-speed changes, and Dynamic Reconfiguration (DR) events for CPUs and/or system boards. These modeling systems cannot, however, enable users to explore questions regarding expected autonomic or human-operator responses to sensor-degradation events. It is the objective of the investigation undertaken here to develop a systematic methodology for high-fidelity simulation of telemetric signals from dynamically executing computer servers. It is important at the outset to note that there are over 1000 physical sensors in today's high end servers. These are hardware transducers that measure such physical variables as temperature, voltage, current, and vibration levels throughout the system. Unfortunately, throughout the computing industry it is often the case that the physical sensors have a shorter mean-time-between-failure (MTBF) than the server assets that the sensors are supposed to protect. Sensor degradation events can either cause premature outages for servers (if the sensors drift out of calibration and trip a threshold); or, worse, the sensors can fail inside their appropriate operating range. This failure mode can result in more catastrophic failures of systems because the protection capability of the sensor is diminished or lost.

K. Gross is with Sun Microsystems Inc., RAS Computer Analysis Laboratory `kenny.gross@sun.com`

E. Schuster is with the Department of Mechanical Engineering and Mechanics, Lehigh University `schuster@lehigh.edu`

The TPSS, when fully integrated with conventional system modeling codes, will ensure that every parameter has exactly the same statistical structure, both deterministically and stochastically, as its corresponding "real" physical variable in the server. Such an integrated system will then provide the capability to simulate any postulated degradation modes, subtle or catastrophic, for sensors in any part of the server. This capability takes on an elevated importance in applications where the focus shifts from simulation to actual feedback-and-control applications, where a control mechanism acting on an erroneous input signal due to a degraded sensor could give rise to improper and possibly catastrophic control actions.

Functional requirement (2) above makes the TPSS a useful adjunct to expert systems that have been developed separately at Sun for signal validation and sensor operability surveillance applications. These expert systems employ a sensitive pattern-recognition technique, the sequential probability ratio test (SPRT), for early annunciation of subtle anomalies in noisy process variables. The standard SPRT test [1], [2] was designed to accommodate independently distributed, gaussian noise. Many physical telemetry variables from servers, however, are contaminated with noise that is either serially correlated or nongaussian (or both). These statistical anomalies, when present in stochastic processes monitored by a SPRT, can give rise to higher false-alarm and missed-alarm probabilities than specified in the design of the SPRT. To obviate these problems, we have demonstrate as part of the project reported here that we can use the TPSS to spectrally filter the predominant serial-correlated components from digitized telemetry variables, then apply the SPRT to the stochastic residual function. The reduction in empirical false alarm probabilities for the SPRT results from an improvement in the whiteness of the residual function, as demonstrated previously for similar applications to signals from commercial nuclear reactors [3].

The development of the Fourier-based spectral decomposition algorithms employed in the first stage of the TPSS calculations is presented in detail in Section II of this report. The four principal modules comprising the TPSS software package developed during this investigation are presented in Section III. In Section IV, we present actual results of applying the TPSS to a variety of measured telemetry signals from executing enterprise-class servers. For each signal used in the TPSS demonstration, a variety of frequency-domain and time-domain analyses are applied to the original and synthesized signals. These analyses have been selected to enable us to make a realistic quantitative assessment of the effectiveness and utility of the TPSS from the standpoints of functional requirements (1) and (2) above. The analyses employed include: the Fisher Kappa test for whiteness, the D'Agostino Pearson moment test for normality, the Kolmogorov-Smirnov test for normality, and a run-of-signs test, which is a non-parametric test to assess serial correlation [4].

II. Methods Used

A continuous system telemetry harness (CSTH) has been developed and patented by Sun Microsystems for a new approach to high-sensitivity proactive fault monitoring of computer servers and all of the active electronic components comprising those servers [5]. The CSTH collects time series signals relating to the health of dynamically executing components, subsystems and integrated systems. These time series provide quantitative metrics associated with physical variables (distributed temperatures, voltages, and currents throughout the system), performance variables (parameters having to do with throughput, transaction latencies, queue lengths, load on the cpu and memories, IO traffic, bus saturation metrics, FIFO overflow statistics, etc.), and various quality-of-service (QOS) metrics. The CSTH signals are continuously archived to an offline circular file (i.e. the "Black Box Flight Recorder"), and are also processed in real time using advanced pattern recognition for proactive anomaly detection. The CSTH coupled with advanced pattern recognition techniques provides sensitive early detection of a variety of mechanisms that are known to cause downtime in enterprise datacenters, including: Environmental issues (thermal anomalies, air-flow restrictions, failed fans); Software aging phenomena (memory leaks, resource contention); Degraded/failed sensors; Degradation of power supplies; and "inferential sensing" capability (wherein a failed sensor is replaced with a highly accurate analytical estimate). This new CSTH innovation coupled with advanced pattern recognition is helping to increase component reliability margins and system availability goals while reducing (through improved root cause analysis) costly sources of "no trouble found" events that have become a warranty-cost issue in computing (and other) industries.

CSTH time-series signals, X_t, contain a serially correlated component, Y_t, and some random contribution, e_t, so that $X_t = Y_t + e_t$. A standard Fourier series describes X_t as

$$X_t = \frac{a_0}{2} + \sum_{m-1}^{N/2} \left(a_m \cos(\omega_m t) + b_m \sin(\omega_m t) \right) \tag{1}$$

where $a_0/2$ is the mean value of the series, a_m and b_m are the Fourier coefficients corresponding to the Fourier frequency ω_m, and N is the total number of observations. Then Y_t represents the Fourier modes with the largest amplitude oscillations, and e_t is a discrete function of random residuals. The following numerical approximation to the Frourier transform is useful in determining the Fourier coefficients a_m and b_m [6]. Let X_j be the value of X_t at t_j, the j^{th} time increment. Then assuming 2π periodicity and letting $\omega_m = 2\pi m/N$, the approximation to the Fourier transform yields:

$$a_m \quad = \quad \frac{2}{N} \sum_{j=0}^{N-1} x_j \cos(\omega_m j) \tag{2}$$

$$b_m = \frac{2}{N}\sum_{j=0}^{N-1} x_j \sin(\omega_m j) \quad (3)$$

for $0 < m < N/2$. Furthermore, the power spectral density (psd), of the signal is given by

$$l_m = N\frac{a_m^2 + b_m^2}{2} \quad (4)$$

To keep the signal bandwidth as narrow as possible without distorting the psd, no spectral windows or smoothing shall be used in the subsequent development. The Fourier modes corresponding to the eight highest l_m provide the amplitudes and frequencies contained in Y_t. The highest eight l_m modes were found to give an accurate reconstruction of Y_t while reducing most of the serial correlation for the telemetry variables under investigation here (see Section III). By no means was an exhaustive study performed to identify the optimum number of harmonics to use in the reconstruction. However, increasing stepwise the number of harmonics used from four to eight was found to significantly reduce the serial correlation. Increasing the number of harmonics used from eight to ten introduces nonphysical high frequency oscillations in Y_t due to in error computing the psd of a noisy signal.

The reconstruction of Y_t uses the general form of (1), where the coefficients and frequencies are those with the eight highest psd values. This reconstruction generates a curve with the same mean and essentially the same deterministic behavior as X_t. Evaluating Y_t and the residual function $e_t = X_t - Y_t$ using various statistical tests to determine whiteness and normality demonstrates the effectiveness of the method. The test employed herein to quantify whiteness is the Fisher Kappa test [7]. The D'Agostino Pearson K^2 test [8] and the Kolmogorov-Smirnov test [9] are used to quantitatively evaluate deviation from normality.

The Fisher Kappa white noise test examines the periodogram of the signals by attempting to reveal periodicities in the data. The periodogram of a white noise process should contain no outstanding peaks or dips. Therefore, the Fisher Kappa test compares the largest psd value with the mean psd value to determine if the time series under study can be considered a white noise process. To accept the null hypothesis that the data is random, the kappa statistic must be lower than its corresponding critical value.

The D'Agostino Pearson K^2 test for normality looks at the third and fourth moments, called the skew and kurtosis respectively, of the signals. The skew determines the degree to which the data lean asymmetrically to one side of the bell curve. The kurtosis reveals whether the bell is too narrow or wide to be considered a true gaussian process. The sample skew and kurtosis are calculated and compared to the estimated skew and kurtosis of gaussian data with the same mean and variance. For a gaussian process, the skew should be zero and the kurtosis should be three. The resulting K^2 test statistic is chi-square with two degrees of freedom. Values of K^2 exceeding the critical value call for a rejection of the null hypothesis that the data are gaussian. Besides determining normality, this test also reveals characteristics of nonnormality from the skew and kurtosis.

The Kolmogorov-Smirnov tesf for normality is similar to the Fisher Kappa test for whiteness except the evaluation takes place here in the time domain as opposed to the spectral domain. The signals are compared to the corresponding value for a normal process, and the difference function evaluated. The test statistic must be lower than the critical value to accept the null hypothesis that the data are normal.

The run-of-signs test [4], used on the raw signal and residual function, is a simple nonparametric tests which checks for autocorrelation. It is based on the hypothesis that positive autocorrelation results in long sequences, or runs, of residuals of the same sign. A run is a sequence of residuals, all of the same sign, with the two residuals immediately surrounding the run having the opposite sign. The run-of-signs test determines the total number of runs in the residual function and compares that to the number of positive residuals (N_1), the number of negative residuals (N_2), and the total number of residuals. An uncorrelated residual function will have the value N_1 close to N_2, and the number of runs will be approximately $(2N_1N_2/(N_1 + N_2)) + 1$.

Finally, adding a random component e_t to the Fourier reconstruction Y_t completes the reconstruction of the signals. Ideally, we would like for the original signal X_t and the reconstructed signal X_t' to have the same mean and variance, or identical first and second moments. To accomplish this, the variance of X_t, Y_t, and the residual function e_t are compared. Since Y_t is a deterministic function comprised of sines and cosines, and e_t is a randomly generated function, the two are independent. This assures additivity of the variances: $\text{Var}(Y_t)+\text{Var}(e_t) = \text{Var}(Y_t+e_t)$. Therefore, generating a Gaussian random function e_t with $\text{Var}(e_t) = \text{Var}(X_t) - \text{Var}(Y_t)$ and mean 0, then adding it to Y_t creates a signal X_t which agrees with the original signal through the first and second moments.

III. Matlab Programs for Decomposition and Reconstruction of Server Telemetry Signals

Matlab version 13 was used in writing all of the code for TPSS. The Matlab Digital Signal Processing (DSP) toolbox was used for frequency domain decomposition, and performing the Fischer Kappa white noise test. The Matlab Statistics toolbox was invoked to perform the Kolmogorov-Smirnov test for normality. The Fisher-Kappa test was adapted from [7] and coded in Matlab language.

The TPSS code has been modularized using the Matlab programming language to maximize flexibility and to facilitate the extensibility, maintenance, and configuration control for the completed system. The Matlab routine FOURREC performs that spectral decomposition of the raw input signals and creates the composite N-harmonic Fourier curve. A second routine called ADDRAN was designed to efficiently

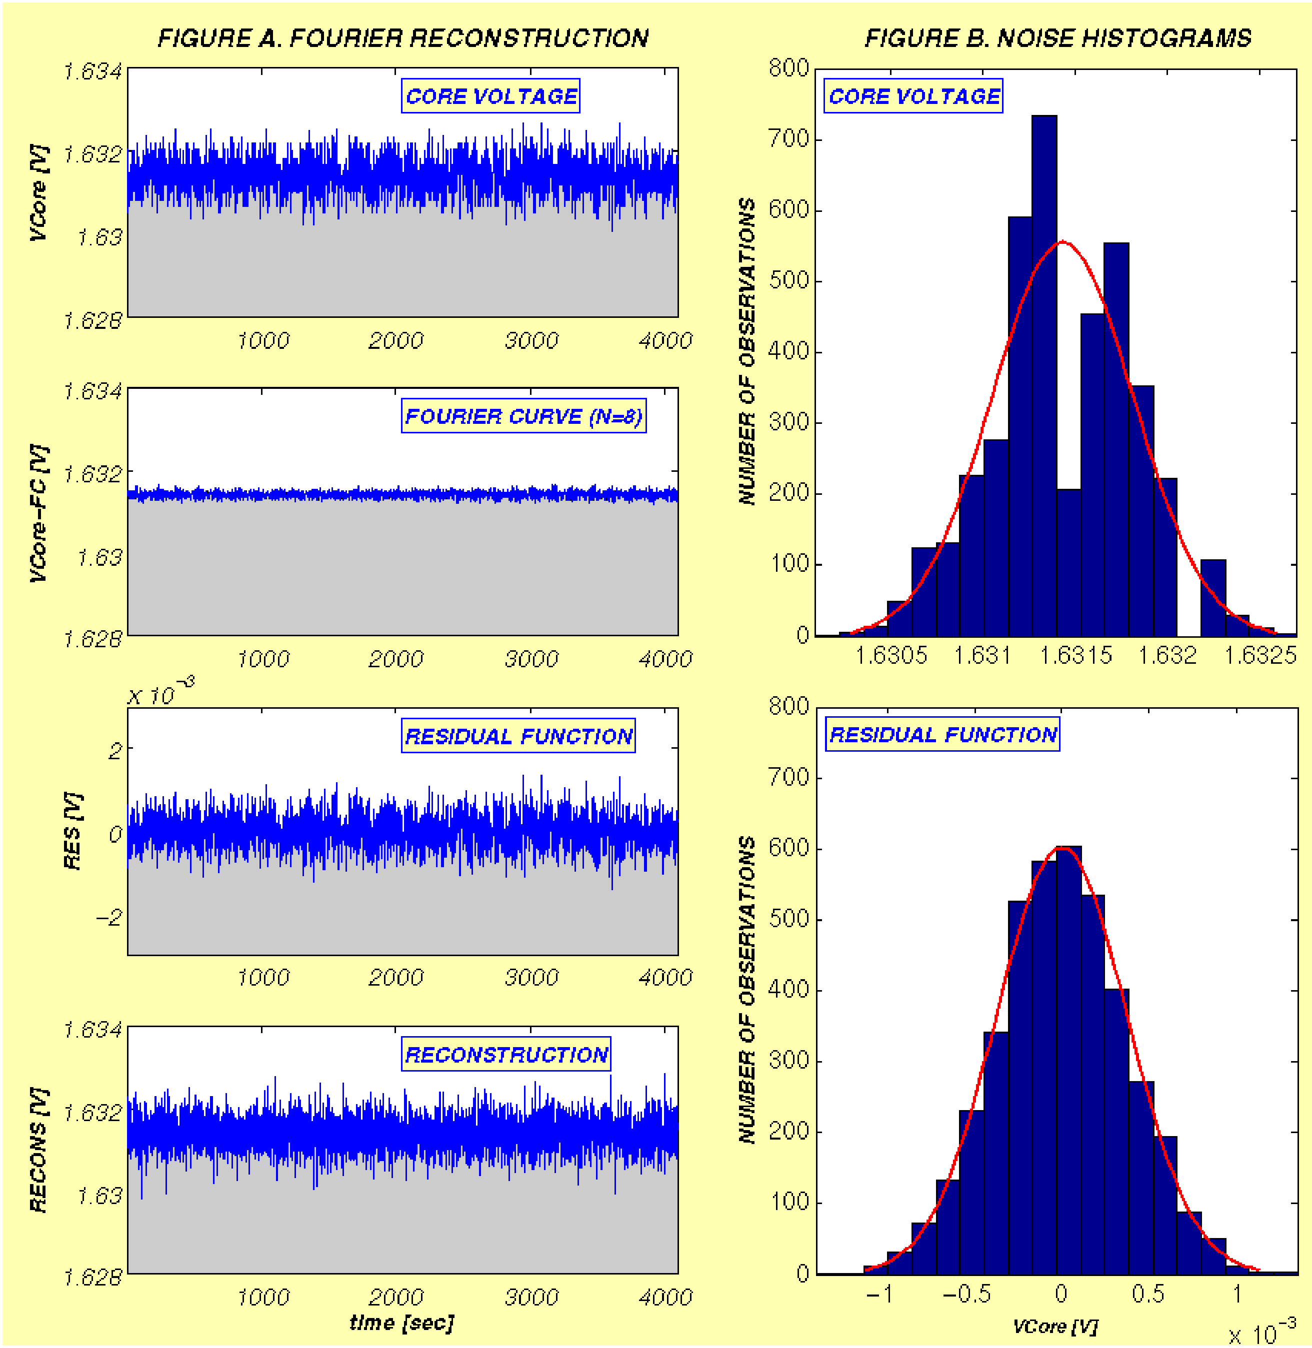

Fig. 1. Core voltage without load

superimpose pseudorandom noise components and generate univariate statistics for the raw and reconstructed signals. The customized Matlab routine NRMTST performs the tesf for normality on the raw signals as well as the residual data. Finally, the RUNS program conducts a run-of-signs test [4] designed to quantify the degree of autocorrelation in the residual data.

IV. Results for Specific Server Telemetry Variables

Various digitized telemetry signals from executing servers, such as various voltages, currents, and temperatures for CPUs, memory modules, and power supplies, are extracted from Sun's Continuous System Telemetry Harness (CSTH) to demonstrate the power and utility of the TPSS core methodology. All statistical tests have been performed for both the raw data and the residual functions. The 95 per cent confidence limit is used for all critical values.

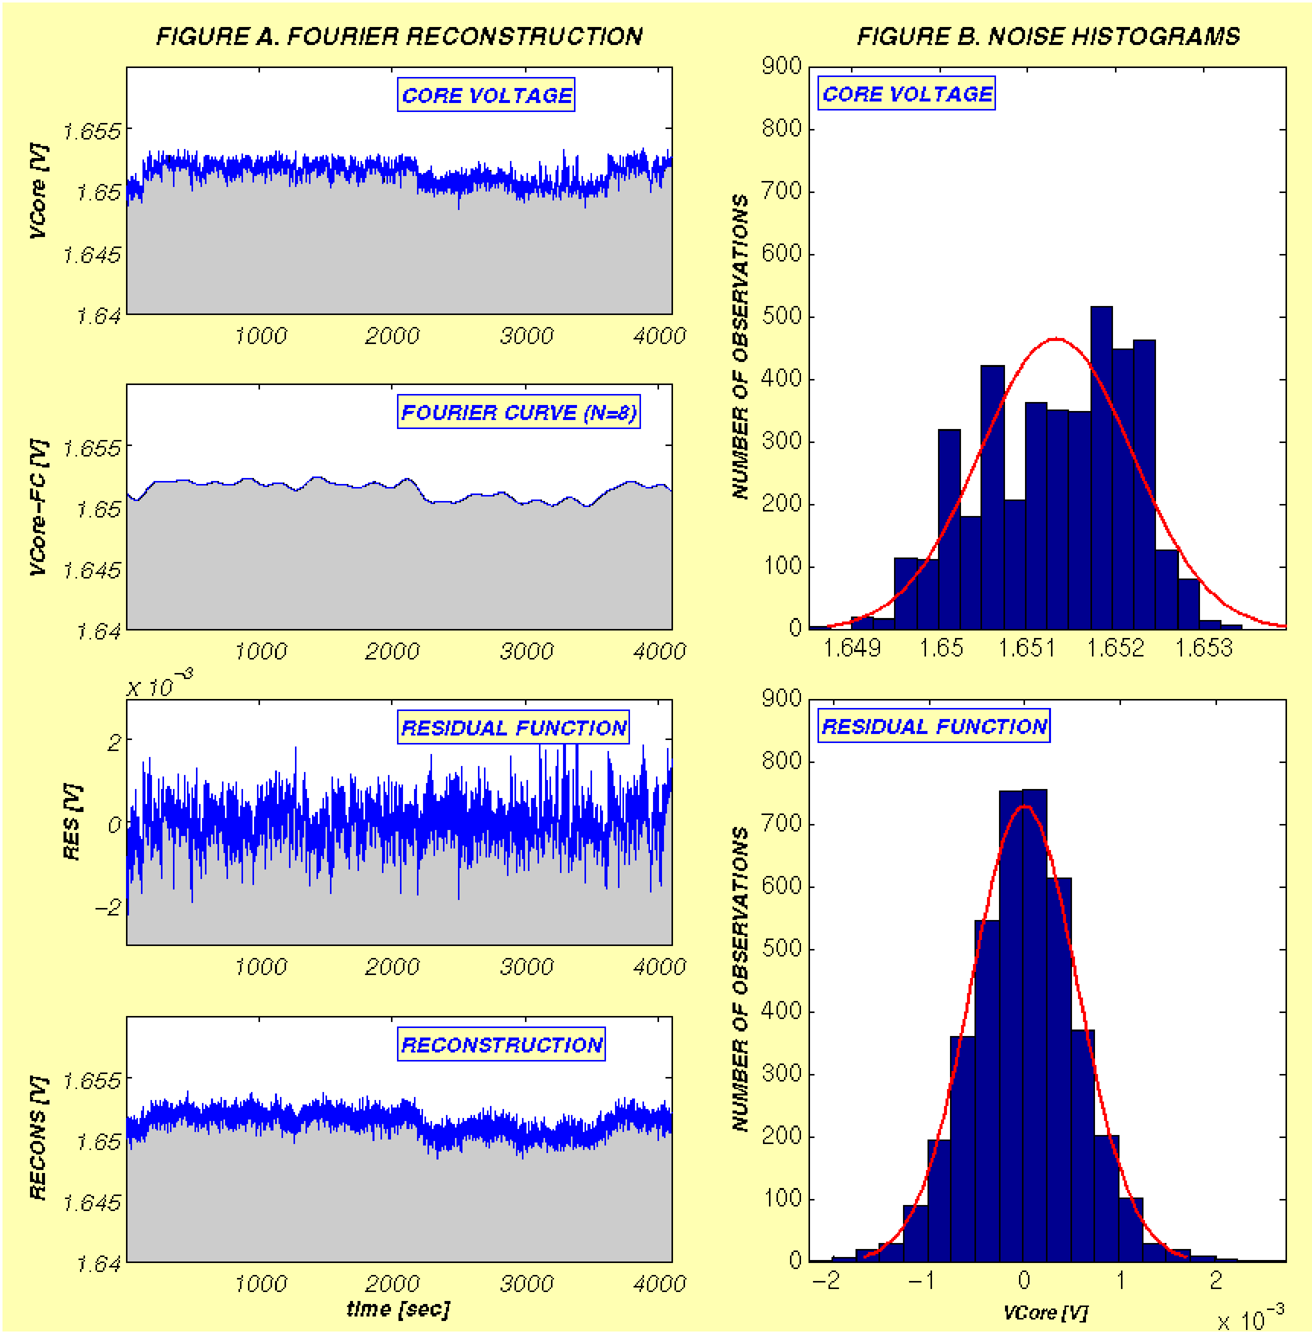

Fig. 2. Core voltage with load

Results at the four stages of the reconstruction are plotted. Periodograms and histograms of the raw data and the residual functions are also plotted. In all cases, signals have been measured from a Sun StarCat F15K server comprising 18 system boards, 72 CPU modules, and 0.5 TeraBytes of main memory. The telemetry data have been digitized at a 30-msec sampling rate using 16,384 observations for each data set.

Results for a core voltage signal are shown in Figure 1. The core voltage is the voltage supplied to each CPU processor on a system board. There are 4 CPUs modules per system board. Figure 1 depicts just one of 72 core voltage signals analyzed. The core voltage signal shown was collected during a time period when the system was idle, i.e. running no user loads. The top subplot shows 4000 observations of the raw voltage signal. The second subplot shows the Fourier composit curve as generated with 8 fourier modes. (N=8 was found to adequately capture the dynamical components of the signals analyzed in this study). The third time-series subplot shows the residual function, obtained by subtracting the Fourier composite

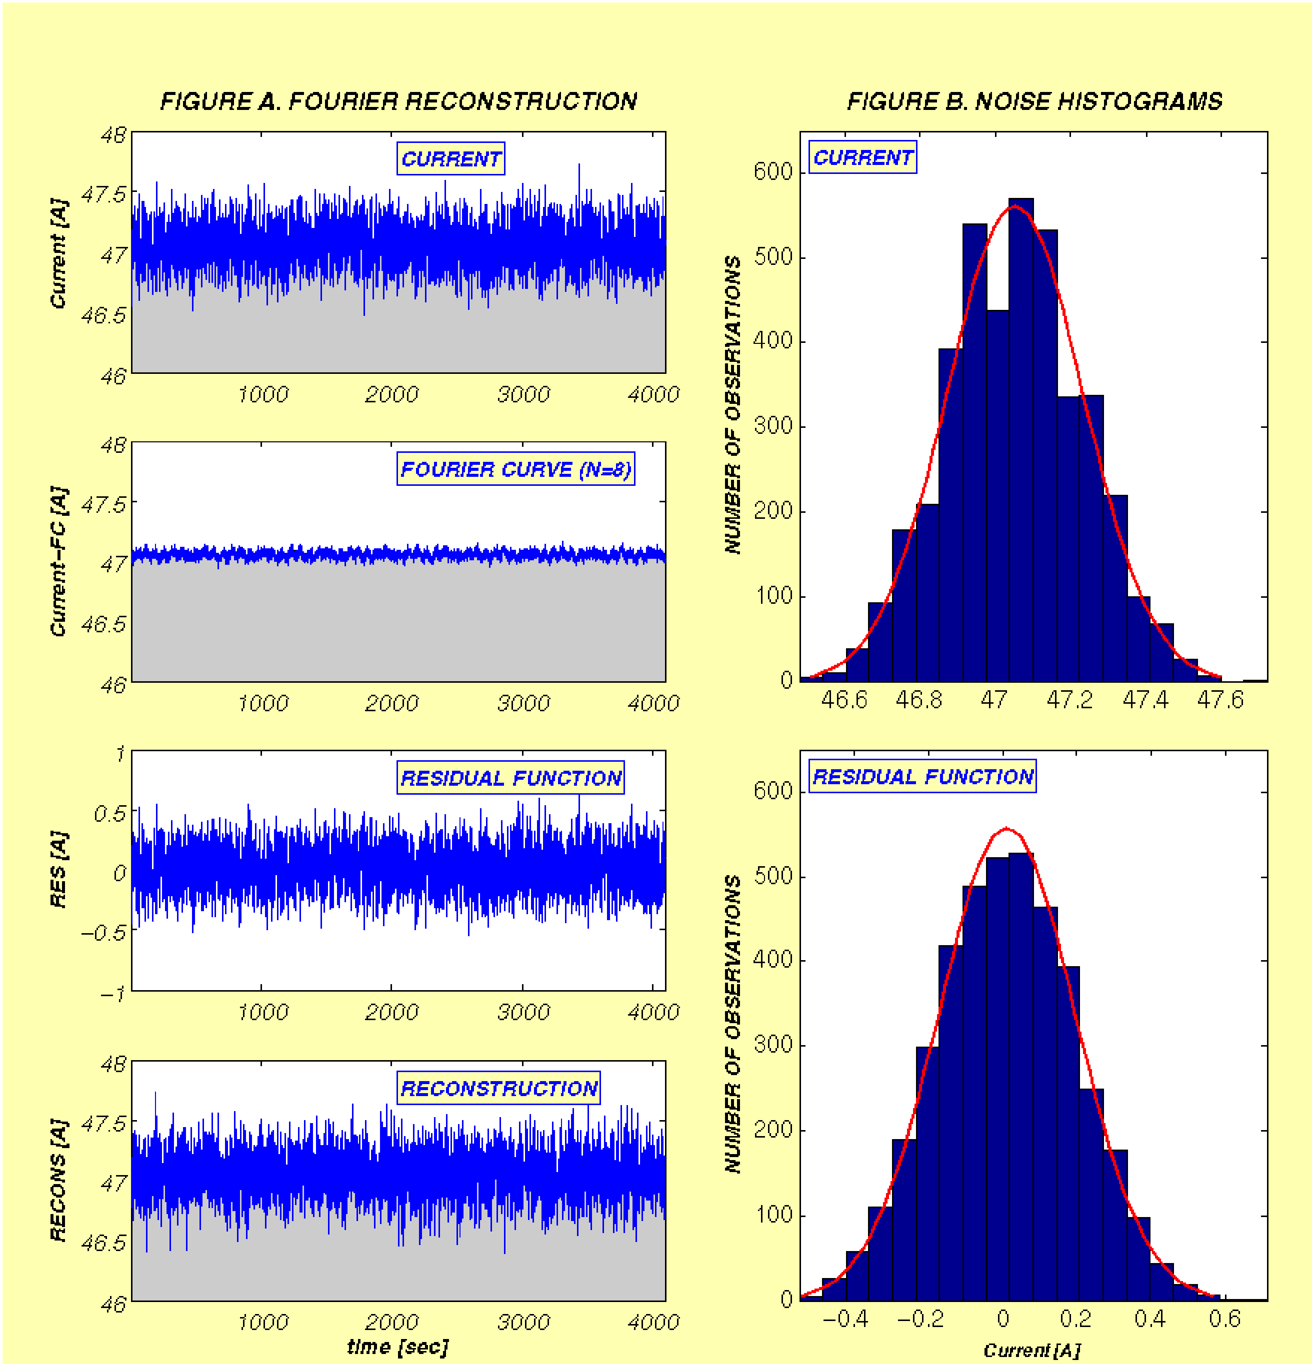

Fig. 3. Current

curve from the raw time series. The final reconstructed signal, obtained by adding pseudorandom noise to the Fourier curve, is shown in the 4th subplot.

The primary value of the technique introduced in this paper lies in the fact that the reconstructed signal shown in the 4th subplot possesses a statistically indistinguishable structure to that for the original raw signal. This means that after one learns the Fourier parameters and the moments of the residual function from a relatively brief experiment, it then becomes possible to synthesize millions of hours of data that can then be used in validation studies of expert system algorithms that are being designed by Sun for proactive fault monitoring applications.

The right side of Figure 1 shows histograms of the raw signal and residual data for the core voltage signal. The first histogram subplot is computed from the original raw signal and is superimposed against a pure Gaussian curve with the same mean and variance. The second subplot shows the histogram of the core voltage signal against a Gaussian curve with the same mean and variance. The improvement in normality (closeness to the Gaussian histogram) is quite apparent in a visual comparison of the top and bottom subplots in the right side of Figure 1. Similar results are shown for a core voltage signal measured during operation with a typical user load profile in Figure 2. Note that even though the raw signal contains a greater degree of dynamical behavior (because of the time-varying load on the system), the TPSS technique nevertheless synthesizes the raw signal with high fidelity. Finally, Figure 3 illustrates application of TPSS to one of very many current signals from the enterprise server used for these investigations. In addition to voltage and current signals, the technique outlined herein has been applied to temperatures, vibration levels, and a wide range of performance metrics that the CSTH extracts from the Solaris operating system.

It should be noted here that selective spectral filtering by the TPSS, which we have designed to reduce the consequences of serial correlation in our sequential testing schemes, does not in itself guarantee that the degree of non-normality in the data will also be reduced. Fortuitiously, for over 90 per cent of the signals we have investigated as part of this investigation, the reduction in serial correlation is accompanied by a reduction in the absolute value of the skewness for the residual function. Moreover, in cases where there is not a reduction in skewness, it can generally be observed that the skewness is very small to begin with. Finally, it has been shown in a separate investigation [3] that nonnormality is much less of a problem, in terms of affecting SPRT misidentification probabilities, than is nonwhiteness.

V. Conclusions

A telemetry parameter simulation system (TPSS) has been built and is now being used as an adjunct to the System Dynamics Characterization and Controls Laboratory at Sun Microsystems. The TPSS has the capability to process any stochastic telemetry variable from enterprise-class computer servers and then (1) generate a synthesized signal that has exactly the same statistical and correlation structure as the actual computer telemetry signal; and (2) accomplish spectral filtering of telemetry variables so that SPRT-based modules may be applied for Electronic Prognostic applications, even when the underlying physical processes internal to the server are contaminated with serially-correlated components. Capability (1) enables the TPSS to be used for high-fidelity telemetry signal simulation, providing rich functionality that has not heretofore existed in the computer-server signal simulation arena. Two of the TPSS's most important early uses will be in novel server feedback-and-control strategies for dynamical provisioning of cooling and/or load; and for system-upset investigations to explore a wide range of scenarios that can arise involving signal faults generated from degraded sensors.

Capability (2) of the TPSS, frequency-domain filtering, makes the TPSS an indespensible tool in Sun's ongoing development of innovative expert systems for signal validation, sensor operability validation, and real-time predictive failure annunciation (i.e. Electronic Prognostics) applications that require high-reliability, high-sensitivity parameter surveillance. Spectral filtering of process variables contaminated with serial correlation will assure that SPRT modules achieve pre-specified false-alarm and missed-alarm probabilities, thereby rendering the overall expert system amenable to formal quantitative reliability analysis methodology. This is an indespensible requriement if Electronic Prognostic systems are to be deployed in business-critical, mission-critical, and especially life-critical applications.

References

[1] A. Wald and J. Wolfowitz, "Optimum Character of the Sequential Probability Ratio Test," *Annals of Mathematical Statistics*, 1948.

[2] K. S. Fu, *Sequential Methods in Pattern Recognition and Machine Learning*, Academic Press, New York, 1968.

[3] K. E. Humenik and K. C. Gross, "Using Fourier Series Methods to Reduce Correlation of Nuclear Power Reactor Data," *Nuc. Sci. & Eng.* (112), 127-135, Nov 1992.

[4] Dickey, D. A. and Brocklebank, J. C., "Checking for Autocorrelation in Regression Residuals," *Proc. SAS Users Group International 11th Conf.*, pp. 959-965, Atlanta, GA, Feb 1986.

[5] K. Whisnant, K. C. Gross and N. Lingurovska, "Proactive Fault Monitoring in Enterprise Servers," *Proc. 2005 IEEE Intn'l Multiconference in Computer Science and Computer Engineering*, Las Vegas, NV (June 2005).

[6] K. C. Gross and K. K. Hoyer, "Reactor Parameter Signal Simulator," *Proc. 4th Intnl. Conf. on Simulation Methods in Nucl. Eng.*, Vol. I, Session 3C, Montreal, Canada, June 2-4, 1993.

[7] W. A. Fuller, "Introduction to Statistical Time Series," *John Wiley & Sons*, Inc., New York, 1976.

[8] R. B. D'Agostino, M. Stephens, "Goodness of Fit Techniques," *Marcel-Bekker, Inc.*, Ch. 6, 1986.

[9] B. Lindgren, "Statistical Theory, 3rd Ed.," *MacMillan Publishing Co.*, New York (1968).