

Dynamic System Characterization of Enterprise Servers via Nonparametric Identification

Eugenio Schuster and Kenneth C. Gross

Abstract—Dynamic characterization and fault detection are carried out in enterprise servers using nonparametric identification techniques based on sinusoidal excitation. The introduction of subtle sinusoidal perturbations in computer load variables or physical variables allows us to obtain a dynamic input-output characterization in the frequency domain. The input-output relationship is described in terms of coupling coefficients between a wide variety of physical and performance variables at different selected frequencies. This innovative approach in the field of computer science, based on a well-known system identification technique, has been demonstrated in empirical studies to provide valuable dynamic system characterization information that can be indispensable to datacenter operations personnel for the functions of performance management, capacity planning, quality-of-service (QoS) assurance, dynamic resource provisioning, and root cause analyses.

I. INTRODUCTION

An e-Commerce transaction server is a complex system with hundreds of resource, performance, and throughput parameters, making the study of relationships between and among variables quite difficult using traditional “static analysis” approaches. Currently, system performance for servers is characterized by testing their operation under maximum load, random load, and using performance benchmarks that mimic typical user loads. These conventional approaches are not able to fully characterize transfer function relationships among performance variables and establish cause/effect relationship between them.

Dynamic characterization of complex systems such as enterprise compute servers and web servers can be achieved by introducing perturbations in one or more “input” variables, and measuring the time-dependent responses in one or more “response” variables. These variables can be physical variables (e.g., from distributed temperature, voltage, and current transducers that are already built into the servers), system resource variables, or quality-of-service (QoS) variables derived from system performance parameters. We quantify this relationship between input variables and response variables with a “dynamic coupling coefficient,” which may be a function of load, or, more generally, may be a multivariate function of very many input variables. In a dynamically executing system such as a web server, distributed synthetic transaction generators can be employed for real-time continuous monitoring of system transaction latencies. These “canary tests” provide QoS performance

metrics on a 24*7 basis as a dynamical function of system load. Specifically, in order to measure the impact of some performance parameter X on another performance parameter Y, the synthetic transactions introduce an (ideally small—to preserve linearity) perturbation in X, from which the resulting effect on parameter Y, if any, can be measured. As an example, one might compress a 10 Mbyte file and attempt to discern the temperature effect on one or more ASIC modules on a system board. Using time domain techniques, such a measurement would very likely be impossible on a large, multi-user, multi-cpu server, because of the extremely small effect one is seeking to discern and the poor signal-to-noise ratio. If such experiments were to be conducted during times of high user activities, the perturbation in X would have to be quite large to infer accurate coupling coefficients. Such maneuvers would likely cause system overload events, and would certainly interfere with the normal day-to-day operation of the system one is seeking to characterize.

The well-known sinusoidal excitation technique for estimation of transfer functions [1] allows us to translate this input-output effect to the frequency domain. The advantage of working with this technique is that we concentrate our effort to a few number of frequency points (the frequencies of the sinusoidal excitations) where the correlation or coupling between variables is clearly seen. The technique has been already adapted by one of the authors for the dynamic system characterization of chaotic, nonlinearly interacting physical variables in nuclear power plants [2], [3], and has been used in many other industrial applications. The use of multifrequency sinusoidal excitation and Fourier-based techniques for identification and fault detection is already pointed out in the literature more than three decades ago [4], [5]. However, although there has been considerable success with applying system identification techniques to computing systems over the last few years [6], these techniques are exclusively based on a time-domain formulation. The sinusoidal excitation technique, used now for dynamical system characterization of large, multi-processor servers, is an elegant and powerful exploratory analysis tool to characterize complex system behavior, particularly the relationships among various dynamic system parameters.

The paper is organized as follows. Section 2 introduces briefly the mathematical underpinnings of the sinusoidal excitation method as a nonparametric identification technique. Section 3 explains how the coupling coefficients are computed taking into account the limitations imposed by the *discrete Fourier transform*. Several experimental results are presented in Section 4. Section 5 summarizes the paper.

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

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

II. MATHEMATICAL BACKGROUND

It is often the case that the mechanism of signal generation is so complex that it is very difficult, if not impossible, to represent a signal as deterministic. In these cases, modeling the signal as an outcome of a random variable is quite useful. It is common to characterize this random variable by simple statistical characteristics such as the mean, variance, skewness, kurtosis or autocorrelation function. For a stationary random process, where the statistical properties are invariant to a shift of time origin, we use the expectation operator E to define the mean, variance, autocorrelation and autocovariance of a sequence $u[n]$ (where n is an integer number) as

$$m_u = E\{u[n]\} \quad (1)$$

$$\sigma_u^2 = E\{(u[n] - m_u)^2\} \quad (2)$$

$$R_{uu}[\tau] = E\{u[n]u[n - \tau]\}. \quad (3)$$

$$C_{uu}[\tau] = E\{(u[n] - m_u)(u[n - \tau] - m_u)\}, \quad (4)$$

and the cross-correlation and cross-covariance as

$$R_{yu}[\tau] = E\{y[n]u[n - \tau]\}. \quad (5)$$

$$C_{yu}[\tau] = E\{(y[n] - m_y)(u[n - \tau] - m_u)\}. \quad (6)$$

While stochastic signals are not absolutely summable or square summable and consequently do not have *Fourier transforms*, many of the properties of such signals can be summarized in terms of the autocorrelation or autocovariance sequence, for which the *Fourier transform* often exists. We define the *power spectrum density* (PSD) as the *Fourier transform* of the auto-covariance sequence,

$$\begin{aligned} \Phi_{uu}(e^{j\omega}) &= \sum_{n=-\infty}^{\infty} C_{uu}[n]e^{-j\omega n}, \\ C_{uu}[n] &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_{uu}(e^{j\omega})e^{j\omega n}d\omega, \end{aligned} \quad (7)$$

and the *cross spectrum density* (CSD) as the *Fourier transform* of the cross-covariance sequence,

$$\begin{aligned} \Phi_{yu}(e^{j\omega}) &= \sum_{n=-\infty}^{\infty} C_{yu}[n]e^{-j\omega n}, \\ C_{yu}[n] &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_{yu}(e^{j\omega})e^{j\omega n}d\omega. \end{aligned} \quad (8)$$

By definition of the auto-covariance and the *inverse Fourier transform* we can note that

$$\sigma_u^2 = C_{uu}[0] = \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_{uu}(e^{j\omega})d\omega. \quad (9)$$

For a linear time-invariant (LTI) system with impulse response $h[n]$, the output sequence $y[n]$ is related to the input sequence $u[n]$ through the convolution sum,

$$y[n] = h[n] * u[n] + v[n] = \sum_{k=-\infty}^{\infty} h[k]u[n - k] + v[n]. \quad (10)$$

The noise sequence $v[n]$, assumed to be uncorrelated with the input sequence $u[n]$, may represent not only real measurement noise but also other phenomena such as uncontrollable inputs or disturbances. We assume for convenience and without any loss of generality, that $m_u = 0$ and $m_v = 0$. Then, we have $m_y = \sum_{k=-\infty}^{\infty} h[k]m_u = 0$.

The cross-correlation (cross-covariance) between the input $u[n]$ and output $y[n]$ is given by

$$\begin{aligned} R_{uy} &= E\{u[n]y[n - \tau]\} \\ &= E\left\{u[n] \left[\sum_{k=-\infty}^{\infty} h[k]u[n - \tau - k] + v[n - \tau] \right]\right\} \\ &= \sum_{k=-\infty}^{\infty} h[k]E\{u[n]u[n - \tau - k]\} + E\{u[n]v[n - \tau]\} \\ &= \sum_{k=-\infty}^{\infty} h[k]R_{uu}[\tau + k] + R_{uv}[\tau] \\ &= h[\tau] * R_{uu}[\tau] \end{aligned} \quad (11)$$

where we have used in the last step the fact that $R_{uu}[-\tau] = R_{uu}[\tau]$ and $R_{uv}[\tau] = 0$ (input and noise sequences are uncorrelated). The frequency response of the LTI system (10) is defined as the *Fourier transform* of the impulse response $h[n]$ and denoted by $H(e^{j\omega})$. Recalling the definitions of the PSD (7) and CSD (8), and taking into account that the *Fourier transform* of the convolution of $h[\tau]$ with $R_{uu}(\tau)$ is the product of their *Fourier transforms*, we apply *Fourier transform* to the last equation and obtain the relationship

$$\Phi_{uy}(e^{j\omega}) = H(e^{j\omega})\Phi_{uu}(e^{j\omega}). \quad (12)$$

For applications of interest to computer monitoring and characterization, a sequence is generally a representation of a sampled signal of finite duration (N samples). For this investigation we therefore seek to estimate the autocovariance and cross-covariance based on finite-length sequences. The estimators for the autocovariance and cross-covariance are respectively defined as

$$\hat{C}_{uu}(\tau) = \frac{1}{N} \sum_{n=0}^{N-1} (u[n] - \hat{m}_u)(u[n - \tau] - \hat{m}_u), \quad (13)$$

$$\hat{C}_{yu}(\tau) = \frac{1}{N} \sum_{n=0}^{N-1} (y[n] - \hat{m}_y)(u[n - \tau] - \hat{m}_u). \quad (14)$$

Both estimators are asymptotically unbiased ($\lim_{N \rightarrow \infty} E\{\hat{C}_{uu}(\tau)\} = C_{uu}(\tau)$, $\lim_{N \rightarrow \infty} E\{\hat{C}_{yu}(\tau)\} = C_{yu}(\tau)$) and in addition it can be showed that

$$E\{\hat{C}(\tau)\} = \frac{N - |\tau|}{N} C(\tau). \quad (15)$$

Assuming, without loss of generality, the case where the sample means $\hat{m}_u = \hat{m}_y = 0$, and taking into account the definition of the estimators for the autocovariance and cross-covariance, it is straightforward to show that

$$\hat{\Phi}_{uu}(\omega) = \sum_{n=-\infty}^{\infty} \hat{C}_{uu}[n]e^{-j\omega n} = \frac{1}{N} |U(\omega)|^2, \quad (16)$$

$$\hat{\Phi}_{yu}(\omega) = \sum_{n=-\infty}^{\infty} \hat{C}_{yu}[n]e^{-j\omega n} = \frac{1}{N} Y(\omega)U^*(\omega), \quad (17)$$

where $U(\omega)$ and $Y(\omega)$ are the *discrete Fourier transforms* of $u[n]$ and $y[n]$ respectively and $P_{uu}(\omega) \equiv \frac{1}{N} |U(\omega)|^2$ is

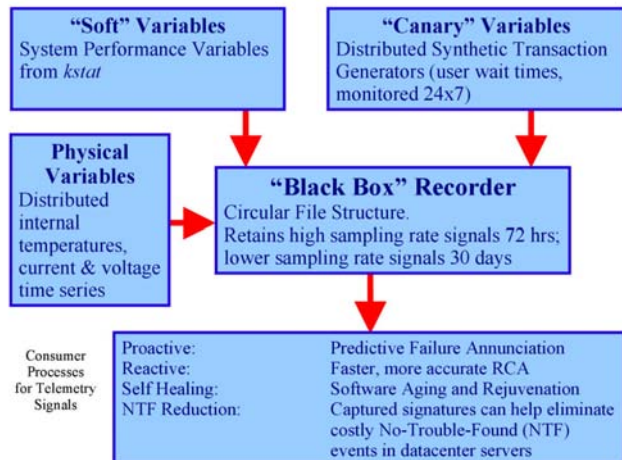


Fig. 1. Scheme of the system telemetry harness.

defined as the *Periodogram* of the sequence $u[n]$. Assuming that the sequence $u[n]$ is the sampled version (with sampling frequency f_s) of a continuous stationary random signal $s(t)$ whose PSD $\Phi_{ss}(\Omega)$ is bandlimited by the antialiasing lowpass filter $(-2\pi\frac{f_s}{2} < \Omega < 2\pi\frac{f_s}{2})$, its PSD $\Phi_{uu}(\omega)$ is proportional to $\Phi_{ss}(\Omega)$ over the bandwidth of the antialiasing filter, i.e.,

$$\Phi_{uu}(\omega) = \frac{1}{T_s} \Phi_{ss}\left(\frac{\omega}{T_s}\right), \quad |\omega| < \pi,$$

then

$$\Phi_{uu}(f) = \Phi_{ss}\left(\frac{\omega}{T_s}\right) = \frac{\Phi_{uu}(\omega)}{f_s}, \quad |f| < \frac{f_s}{2}.$$

Using the estimators (16) and (17), (12) suggests that the *frequency response* of the system at some specific frequency ω_o can be estimated as:

$$\hat{H}(e^{j\omega_o}) = \frac{\hat{\Phi}_{uy}(e^{j\omega_o})}{\hat{\Phi}_{uu}(e^{j\omega_o})}. \quad (18)$$

III. ANALYTICAL APPROACH

A. System Telemetry Harness

An advanced system telemetry harness has been developed as part of this effort to collect, preprocess, analyze, and archive hundreds of system performance, throughput, quality-of-service (QoS), and physical variables. This comprehensive system telemetry harness produces a plethora of performance variables to monitor. Their number is restricted to about 150 by eliminating those that are either very poorly correlated, those that are redundant (e.g., variables that may report the same metric, but with different units), or those with a high degree of colinearity. Detailed cross-correlation and coherence analysis was performed at the outset of the investigation. Since sampling for various performance metrics is done at disparate (and sometimes time-varying) sampling rates, an analytical resampling algorithm




Fig. 2. (a) Input-output relationship. The system S has inputs u , disturbances d and outputs y . (b) Transaction activity profile.

is invoked to align all the parameters and bring them to a uniform, synchronized sampling rate. A schematic overview of the system telemetry harness used for this investigation is illustrated in Figure 1.

B. Spectral Decomposition of Digitized Telemetry Signals

We introduce here a multifrequency sinusoidal excitation approach to perform dynamic system characterization of complex, multi-CPU enterprise computer servers. We are interested in determining the degree of coupling between any pairwise combination of performance and/or physical variables in our system. If such relationships can be established with sufficiently good accuracy, the approach can be used to proactively detect a wide range of system anomalies that have heretofore been obscured by conventional time-domain stress testing methods.

During system operation with normal user workloads, very subtle sinusoidal perturbations are introduced in what we call the input variables (typical user transaction requests). The magnitude of the sinusoidal perturbation must be small enough to preserve the linearity assumption around the equilibrium point and the validity of the theory. The effect of these perturbations in the output variables (quality-of-service metrics) is quantified in terms of dynamic coupling coefficients. Figure 2-a illustrates the procedure and Figure 2-b shows an example transaction activity profile used as sinusoidal perturbation.

The coupling coefficient is defined as the transfer function between the input and output variables at the excitation frequency $0 < \omega_o < \pi$. According to (18), we can estimate the coupling coefficient $C(\omega_o)$ as

$$C(\omega_o) = \hat{H}(e^{j\omega_o}) = \frac{\hat{\Phi}_{uy}(e^{j\omega_o})}{\hat{\Phi}_{uu}(e^{j\omega_o})}. \quad (19)$$

Based on the knowledge of $\hat{\Phi}_{uy}(e^{j\omega_o})$ and $\hat{\Phi}_{uu}(e^{j\omega_o})$ the computation of the coupling coefficient is straightforward, but subject to errors that depend upon (1) the quality of

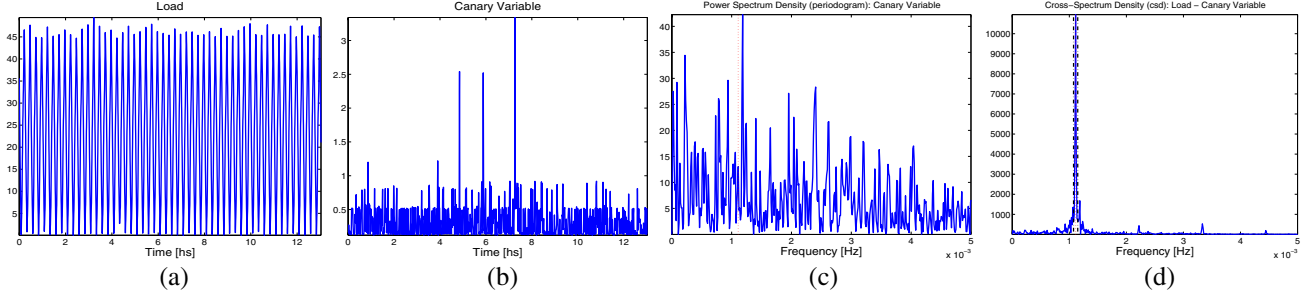


Fig. 3. PSD vs. CSD analysis.

the telemetry system that is measuring input and response variables, and (2) the sampling rate attainable for the target server. It is important to keep in mind that in the real world what is computed is the *discrete Fourier transform* and not the *Fourier transform*. This means, based on (16) and (17), that we know the PSD and CSD only at discrete points over the interval $-\pi \leq \omega < \pi$. It may happen that none of these discrete frequency points coincide exactly with the target excitation frequency ω_o . The value of the PSD at ω_o depends on the way we compute the DFT (number of points or sampling frequencies). To make our analysis independent of the DFT parameters, we consider (9). When the spectrum is plotted, the total area under the curve is equal to the variance of the time series and is independent of the shape of the curve. A physical interpretation of the PSD function is that $f(\omega)d\omega$ represents the contribution to variance of components within the frequency range $(\omega, \omega + d\omega)$. A peak in the spectrum indicates an important contribution to variance at frequencies in the appropriate region. This suggests that instead of computing the coupling coefficient as a ratio of PSDs, we can compute it as a ratio of power averages around the frequency ω_o ,

$$C(\omega_o) = \frac{\int_{\omega_o - \Delta\omega}^{\omega_o + \Delta\omega} \hat{\Phi}_{uy}(e^{j\omega}) d\omega}{\int_{\omega_o - \Delta\omega}^{\omega_o + \Delta\omega} \hat{\Phi}_{uu}(e^{j\omega}) d\omega}, \quad (20)$$

where $\Delta\omega$ is a multiple of the frequency resolution $2\pi/N$ of the DFT.

As an example of this technique, a sinusoidal perturbation with period of 15 minutes is generated in one of the input variables. This sinusoid has a very small amplitude in comparison with normal variations in user load patterns (typically $< 1\%$ of nominal variations). One or several synthetic client transactions (called the “canary” variables) are launched to execute typical user transactions (example: file compression, table lookup, inversion of a small matrix, sort a linear list). The response times for these canary tests are recorded to produce continuous time series that reflect QoS from the end-user perspective. In parallel with the canary tests, a large suite of system performance, throughput, and transaction latency variables as well as physical variables are recorded on a 24*7 basis using a continuous system telemetry harness, which has been separately developed by Sun Microsystems for high end UNIX® servers. Figures 3-a,b show respectively the transaction profile imposed as

perturbation load and the response time of the synthetic variable as functions of time. Figure 3-c plots the PSD of the response time of the synthetic client (canary variable) as a function of frequency. The stochastic noise (chaotic user load) associated with the canary signal is so high that the test period of 15 minutes is not discernable in the univariate PSD. This illustrates that characterization of typical web-server performance metrics is not amenable to univariate spectral decomposition calculations via conventional Fourier analysis because the signal-to-noise ratio is too small to discern the sinusoidal perturbation in PSD of the response variables. To overcome this limitation of conventional Fourier analysis methods, we employ the CSD, a bi-variate diagnostic technique that is highly sensitive, even to weakly coupled parameters with very poor signal-to-noise ratios, dramatically and selectively amplifying the input sinusoid harmonics in response variables such that the period of the sinusoidal perturbation in the control variable is readily apparent with excellent peak resolution and low noise and “side lobe” contamination [7]. In the CSD subplot in Figure 3-d, a well-defined peak corresponding to the period of the sinusoid is readily observable, implying a common periodicity and, hence, a cause-and-effect relationship between the sinusoidal perturbation in the load and the synthetic client’s response time.

IV. EXPERIMENTAL RESULTS

A. Coupling Coefficients: Load \rightarrow Physical & Computer Variables

One subject of our investigation is the relationship of different variables with respect to the dynamic load on the server. Figure 4 (left) shows the coupling coefficients at some specific frequency between both physical and performance variables and the load. The magnitudes of these coupling coefficients are a direct measurement of the correlation of the variables with respect to the load. By ranking the coupling coefficients from higher to lower, we can learn which variables are more sensible to load variations.

Understanding that this ranking can vary with the frequency, it is quite straightforward to extend the approach outlined above to span multiple, simultaneous excitation sources. As an example of this multifrequency approach, Figure 4 (right) shows the coupling coefficient at three different frequencies. In addition, information coming from

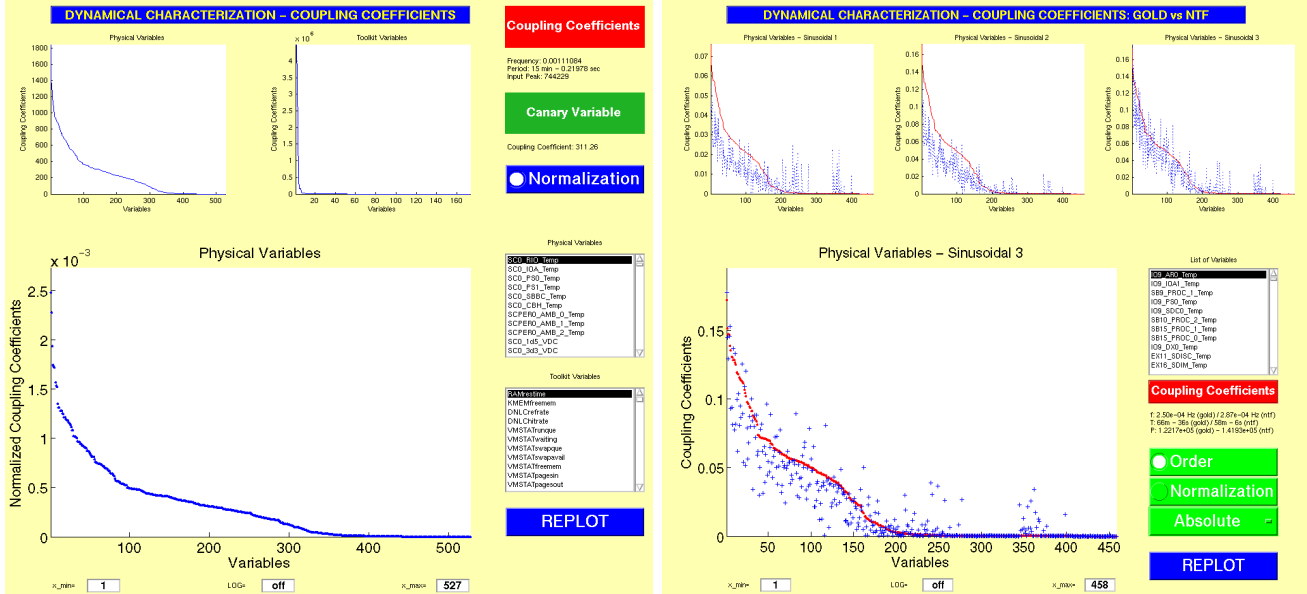


Fig. 4. Left: Coupling coefficients between load and physical and computer variables. Right: Coupling coefficients between load and physical variables. Comparison between reference (·) and test (+) boards.

two different servers has been superimposed in the figure. This technique is useful for detection of very subtle faults in a “test” machine by comparison with a “reference” machine. The test machine may be a server at a customer’s datacenter, and the reference machine may be an identically configured new server in the computing vendor’s laboratory. By running the same sinusoidal-load profile script in both the reference and the test machines, and taking special care of reproducing the same equilibrium point in both machines, we may expect to have similar coupling coefficients. Identifying differences in the coupling coefficients between the reference and the test machine provides a particularly sensitive method to expose subtle faults that have heretofore been difficult to root cause and have contributed to higher warranty and serviceability costs for computing server vendors. The coupling coefficients can be displayed in absolute values, as a ratio between the reference and the test machines, or as a difference between them in order to facilitate the detection of outliers.

B. Coupling Coefficients: Temperature \rightarrow System Voltages

We are also interested in studying the cross spectral correlation between physical variables, for instance, the coupling coefficients between temperature and voltage in several board CPU processors. In this experiment, sinusoidal temperature oscillations are introduced by use of a controller script that systematically varies the speeds of individual fans in an array of fans that are used to cool the server. This experiment is useful for “teasing out” sources of very subtle faults that show up as fluctuations in voltage, but that are accelerated by thermal variations. Anomalous behavior in the value of a coupling coefficient relating temperatures to voltage or temperatures to current may be an indication of a subtle incipient anomaly in system board components, and may enable service personnel to

identify and remove boards at elevated risk before their subsequent deterioration may adversely impact a business-critical datacenter.

C. Sensor Operability Validation

This multifrequency sinusoidal impulsion technique also proves to be very useful for the purposes of signal validation and sensor operability validation in complex servers [8]. Modern high-end computer servers may have over 1000 physical transducers measuring distributed temperatures, voltages, and currents throughout the system. It is very frequently the case that these inexpensive physical sensors have shorter Mean Time Between Failure (MTBF) values than the assets the sensors are supposed to protect. A common practice in the computing industry is to place threshold limits on the sensor signals so that, for example, if a temperature goes too high or a voltage too low, a warning message is generated or other, automated actions are taken to protect the system. If any of the sensors should fail during the life of the server, then the degree of protection for the system is degraded, as is the ability to proactively detect anomalous conditions by real time monitoring agents. The coupling coefficient method developed here provides a novel means to obtain “instant” sensor operability validation for all sensors in the server at any time desired. Figure 5 (left) shows the values of the coupling coefficients between load and the cpu voltages for both the reference and test machines. Variable 7 in the figure shows a very small value for the test machine as compared with the reference machine indicating that the sensor associated with that variable is not functioning properly. Root cause analysis for the system board containing Sensor 7 revealed that the sensor was in a “stuck-at” condition, a failure mode that is extremely difficult to catch with conventional threshold limits.

Data for the above example was generated by applying

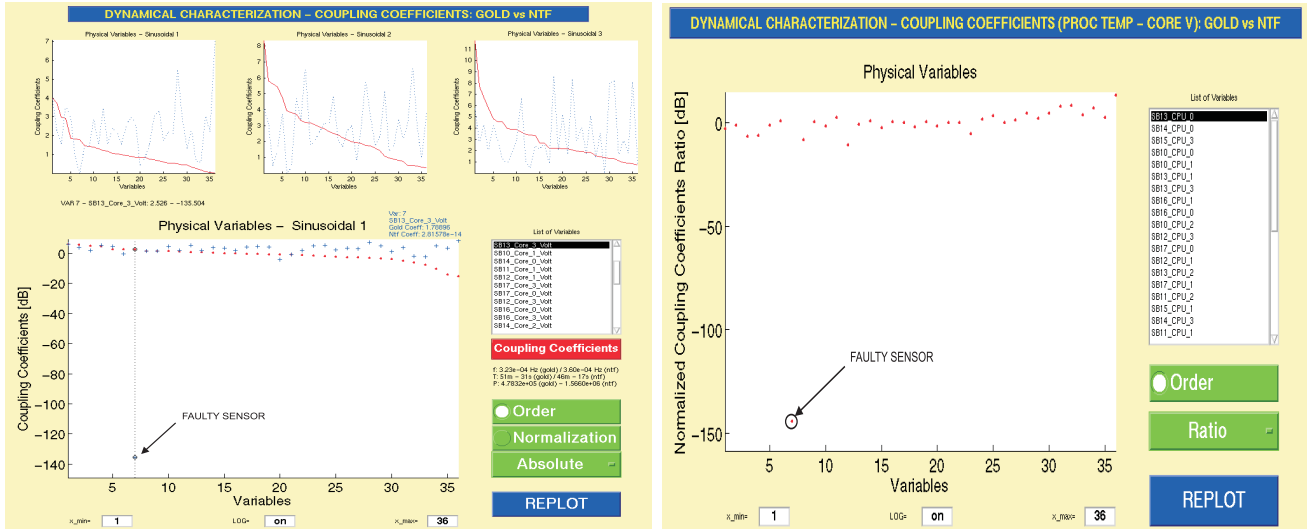


Fig. 5. Left: Coupling coefficients between load and system voltages. The outlier is evidence of a faulty sensor. Comparison between reference (-) and test (+) boards. Right: Coupling coefficients between temperatures and system voltages. Comparison between reference and test boards. An outlier in the ratio means a difference between the two boards and evidence of a potential fault.

the multifrequency sinusoidal impulsion method through system load. As an additional illustration of the diverse applicability of the techniques introduced herein, we illustrate in Figure 5 (right) the dynamic coupling coefficients between temperatures and the voltages for various system boards in the test server. In this case the information is presented as a ratio between the dynamic coupling coefficients of the reference machine and the test machine. The outlier in the figure is an indication of a potential fault in the associated CPU processor on one system board in the test machine. In this case a root cause analysis revealed that the sensor itself was functioning properly but there was a fault in the A/D converter used to digitize the sensor's output signal.

V. SUMMARY

Most present methods for qualification testing of enterprise computing systems involve putting a maximum expected load on one or multiple input variables, and seeing if the system hangs or crashes. While this type of qualification testing is necessary, we have found that dynamical response testing can provide an additional wealth of information that is useful for designing robust systems that deliver optimal QoS performance over a large range of system performance. Nonparametric identification by sinusoidal excitation described in this paper can be used to evaluate with a very high accuracy, dynamic coupling coefficients, transfer functions, and phase relationships among a wide range of physical, throughput and performance metrics. Experiments documented herein with large, multi-processor web servers have demonstrated that complex systems comprising chaotic performance dynamics and high stochastic content are not amenable to univariate spectral decomposition calculations via conventional Fourier analysis because the Signal/Noise ratio is too small. The cross-spectrum technique can accurately assess cross correlation and coherence relationships

among multiple, dynamic system parameters, even those characterized with extremely poor signal to noise ratios. In addition to establishing whether there is any causal association between input impulse variables (i.e., any types of typical user transactions) and measured response variables (temperatures, voltages, currents, as well as QoS metrics), the techniques introduced here also provide: dynamical coupling coefficients between input variables and response variables; and phase shifts between “cause” and “effect” variables throughout the dynamically executing system. The coupling coefficient method provides a novel means to obtain “instant” sensor operability validation for all sensors in the server at any time desired.

REFERENCES

- [1] L. Ljung, *System Identification: Theory for the user*, 2nd edition, Prentice-Hall, 1999.
- [2] K. C. Gross and L. K. Polley, “Investigation of Nonrecoil Fission-Product Release Phenomena Using Multifrequency Source-Perturbation Experiments in EBR-II,” *Annals of Nuclear Energy*, Vol. 18 (3), 1986, pp. 419-441.
- [3] K. C. Gross and R. V. Strain, “Sinusoidal Source Perturbation Experiments With a Breached Fuel Subassembly in the Experimental Breeder Reactor II,” *Nuclear Technology Journal*, Vol. 98, 1992.
- [4] Cusset, B. F. and Mellichamp, D. A., “On-line identification of process dynamics. A multi-frequency response method,” *Industrial and Engineering Chemistry, Process Design and Development*, v 14, n 4, Oct. 1975, p 359-68.
- [5] Rees, D., “Automatic testing of dynamic systems using multi-frequency signals and the discrete Fourier transform,” *Proceedings of the New Developments in Automatic Testing Conference*, 30 Nov.-2 Dec. 1977, Brighton, UK.
- [6] J. L. Hellerstein, “Challenges in Control Engineering of Computing Systems,” *Proceedings of the 2004 American Control Conference*, Boston, MA, June 30 - July 2, 2004.
- [7] G. M. Jenkins and D. G. Watts, *Spectral Analysis and its Applications*, Emerson-Adams Press, 2000.
- [8] K. C. Gross and E. Schuster, “Method and Apparatus for Validating Sensor Operability in a Computer System,” U.S. Patent Application filed July 2004 (SUN Microsystems, Inc. case number SUN04-0605-SPL).