N94-71106

Supply Current Diagnosis in VLSI

J. F. Frenzel
Electrical Engineering Department
University of Idaho
Moscow, ID 83843

P. N. Marinos
Electrical Engineering Department
Duke University
Durham, NC 27706

Abstract - **This paper presents a technique based upon the power supply current signature (cd) which allows for the testing of mixed-signal systems, in situ. Through experiments with a microprocessor, the cd is shown to contain important information concerning the operational status of the system which may be easily extracted using approaches based on statistical signal detection theory. The fault-detection performance of these techniques is compared to *that achieved through auto-regressive modeling of the cd.***

1 Introduction

The growth of mixed-signal technology has created a great need for new methods of system testing and fault prognostication. The main objective of this research was to develop a unified test methodology, applicable to digital as well as analog systems, that would reduce fault modeling requirements, eliminate completely the need of any partitioning of hybrid systems into their respective analog and digital subsystems for purposes of testing, and simplify the test generation process. To satisfy such a broad objective, it became necessary to search, both theoretically and experimentally, for system observables carrying information about the functional status of the system, and methods for extracting such information in a manner useful for purposes of fault detection and system prognosis.

1.1 Review of Supply Current Analysis

As early as 1975 it was postulated that monitoring of the supply current could provide certain advantages in the testing of digital integrated circuits [20], [21]. Yet, supply current testing lay essentially dormant until the explosion of CMOS technology led researchers to reexamine the benefits afforded by current testing. Levi was one of the first to comment upon the characteristics of CMOS technology which make it particularly amenable to what is referred as "I_{DD} Testing" [8]. This initial treatise was continued by Malaiya and Su, culminating in procedures for applying I_{DD} testing and estimating the effects of increased integration on measurement resolution [9], [10].

Recently, several researchers have examined I_{DD} testing as a method of quantifying reliability. Hawkins et al. have reported on numerous experiments where I_{DD} measurements have forecast potential reliability problems in devices which had previously passed conventional test procedures [17], [12]. This application has prompted research dedicated to improving the accuracy of measuring I_{DD} [7], [4]. Maly et al. have proposed a built-in current sensor which provides a pass/fail flag when the current exceeds a predetermined threshold. Combined with a switching mechanism, it provides a means of removing the faulty device from operation once excessive current flow is detected [18], [16], [15].

1.2 Power Supply Current Signature Analysis

All of the research on supply current testing, to-date, has been focused on comparisons of the quiescent current to a simple threshold for purposes of fault detection. No effort has been made to examine the AC characteristics of the supply current waveform for indications of potential failures. While Dorey et al. acknowledge the potential information to be gained from a study of switching currents, they dismiss this area due to the complexities of waveform acquisition and analysis [13]. Only recently, Hashizume et al. utilized an autoregressive (AR) model of the supply current waveform for detecting faults in combinatorial logic through pattern recognition [14]. By analyzing the *entire* cd as a continuous-time signal, it is possible to develop a test methodology, applicable to both analog and digital technology, capable of fault prognostication. However, estimating the coefficients needed in the AR model of the cd is computationally burdensome.

In this paper we will develop and evaluate an efficient method for extracting information from the cd using statistical signal detection theory. Section 2 describes the simulation of microprocessor functional faults which were used to evaluate the test technique. A model for the cd is presented in Section 3, and based upon limited assumptions, a method for detecting an unknown fault component, referred to as "the likelihood ratio test", is introduced. The performance of this technique against the simulated microprocessor failures is examined and a method for system prognosis presented. In Section 4 we compare the fault detection performance of AR modeling to that achieved by the likelihood ratio test. Finally, Section 5 summarizes the results of this research and presents recommended usage of the test technique.

2 Simulation of Microprocessor Functional Faults

In this section we describe the simulation of functional failures using the Intel 8086 microprocessor. The Intel 8086 was running at 2.45 MHz on an SDK-86 development board. The power to the processor was isolated from the board supply and the current drawn by the processor sampled at a 102.4 MHz rate using an AC current probe and a digitizing oscilloscope.

Three classes of functional faults were investigated: *data storage and transfer faults*, *register decoding faults*, and *instruction decoding faults* [6]. Simulated failures were introduced by modifying either initial register contents or the instructions stored in program

memory. It should be stressed that the instructions were modified such that the number of one's in the instruction code or operand(s), referred to as the *weight*, was constant for any given byte. This is important, as the amount of current drawn by the input buffers during the instruction fetch cycle contributes greatly to the overall current. If the weight of a particular byte was altered in order to simulate a fault, then the change in current drawn by the input drivers when that byte was fetched from memory might obscure any additional current variation. Consequently, due to strict adherence to this principle, it is possible to attribute deviations in the current signature to the presence of the simulated fault, rather than the modeling technique.

2.1 Data Storage and Transfer Faults

Data storage (or transfer) faults were modeled by altering the contents of a register used by a move instruction. The reference case initialized the register DX with the operand AAAA. For the simulation of fault-1 DX was initialized with the operand 5555 which has the same weight as the operand used in the reference case. Under fault-2, DX was initialized with the value ABAA which has a weight one greater than that of AAAA. The deviations of the resulting signatures from the reference signature are shown in Figure 1.

From these experiments we might project that it will be difficult to detect data faults that do not change the weight of either the operands or the result of a particular instruction, as the differences observed under data fault-1 appear to be random. A possible explanation is that the supply current signature is actually the sum of many individual currents; if one transistor draws less current due to a change in logic state, yet an equivalent transistor in another bit position draws more current due to the opposite state change, then the two effects will cancel, leaving no discernible differences in the resulting signature.

2.2 Register Decoding Faults

Register decoding faults were modeled by altering the register field of individual instructions, thus enabling the selection of incorrect source and/or destination registers. In each case, the total weight of the register field was kept constant and all registers were initialized to the same value to to prevent the introduction of any artificial effects.

Fault-1 involved modifying the instruction MOV BX,DX, encoded as 8B DA, to 8B D3 which caused the execution of MOV DX,BX. This modeled the occurrence of a register decoding fault which caused the selection of incorrect source and destination registers. Figure 2 shows the difference between the reference signature and that observed under the simulated fault.

As a second example, fault-2, we chose to model the selection of an incorrect source register. This was done by modifying the register field of the MOV instruction to D9, which caused the execution of MOV BX,CX. Because the register addresses of CX and DX have the same weight, it was expected that this fault would cause very little change in the cd. However, as shown in Figure 2, the difference between the reference signature and that obtained under the simulated fault shows a large peak at the point which corresponds

to the activation of the source register. From this we conclude that although CX and DX appear equivalent, since they are both general purpose registers and their register addresses are of equal weight, there are electrical variations which cause a large difference in the amount of current drawn during their use.

2.3 Instruction Decoding Faults

Instruction decoding faults were introduced by modifying the opcode fields of individual instructions, taking care to preserve the weight of the opcode. Two instruction decoding faults were modeled; fault-1 was introduced by changing the opcode field of the MOV instruction from 8B to 8E, which resulted in the execution of the instruction MOV DS,DX. The second fault was injected by changing the opcode field to 39 which executed the instruction CMP DX,BX. In both cases the weight of the opcode field remained consistent with that used to produce the reference signature. The differences between the reference signature and the signatures recorded under fault-1 and fault-2 are shown in Figure 3. These differences are significantly greater than the differences observed under data faults or register decoding faults. Evidently, the execution of an incorrect instruction severely impacts the current signature, allowing for simple detection of the fault. This is an important advantage of Power Supply Current Signature analysis, as control faults are typically the most difficult to detect. Control faults may give rise to the execution of spurious instructions which could contaminate register contents. Detection of such spurious instructions involves the time-consuming propagation of the processor state to observable outputs.

3 The Likelihood Ratio Test

It has been demonstrated, using the results of SPICE simulations and circuits comprised of 7400-series devices, that examination of the cd may be used for purposes of fault detection [23]. A method of analysis, referred to as "transition matching", was developed and demonstrated to be extremely effective. However, when applied to more complex devices, transition matching proved to be insufficient at completely utilizing the information contained within the cd. This experience exemplified the danger associated with optimizing an analysis technique for a particular system; each system will exhibit its own characteristics and fault responses, depending upon the technology (CMOS, ECL, hybrid, ...) and level of integration (module, board, ...). If we wish to obtain a methodology which may be applied successfully across all boundaries, then we must develop such a technique without placing any restrictions upon the form of the cd or the fault effects. In this section we will present a cd model that limits these assumptions and develop a method of analysis based upon statistical signal detection.

3.1 cd Model Development

We have chosen to model the observed cd as the sum of three signal components as shown in Equation 1. The observable supply current drawn by a device under test, $z(t)$, is equal to

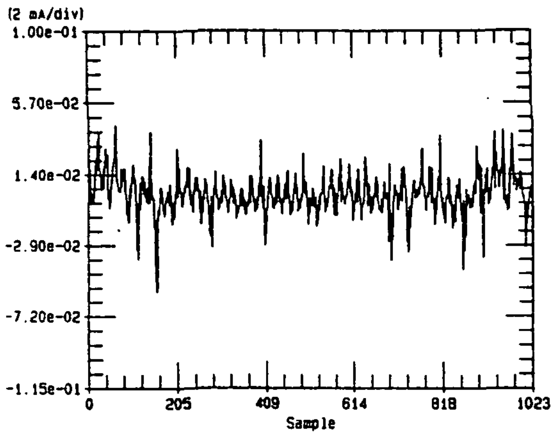

Fault-1

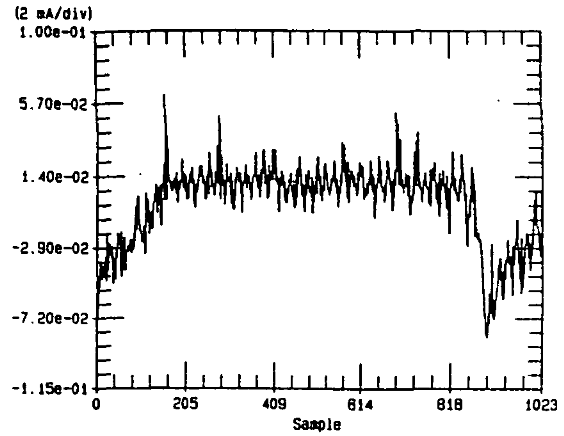

Fault-2

Figure 1: Data Fault Differences

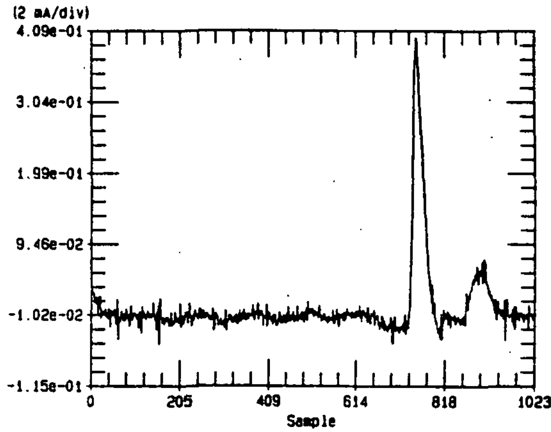

Fault-1

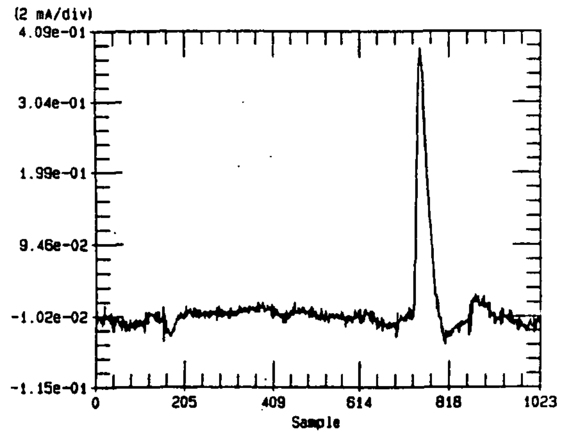

Fault-2

Figure 2: Register Fault Differences

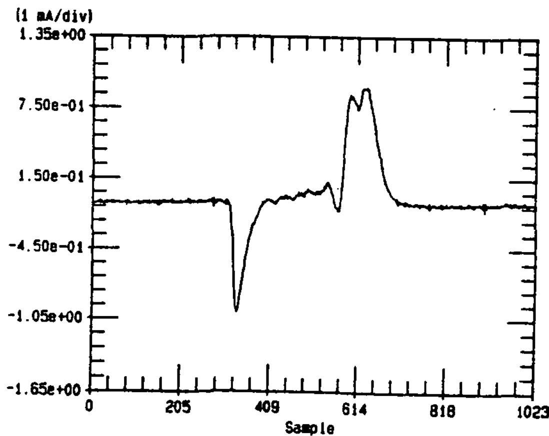

Fault-1

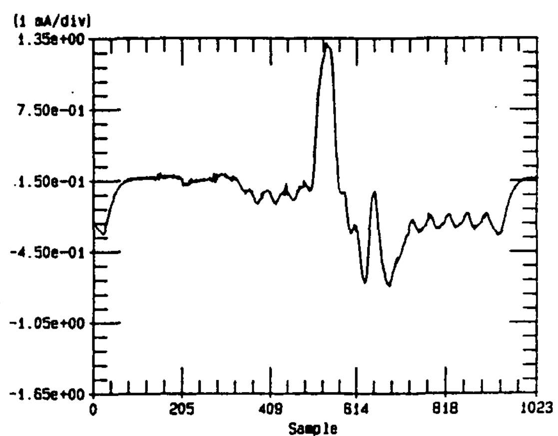

Fault-2

Figure 3: Instruction Fault Differences

the sum of the current drawn by a fault-free device, $w(t)$, any additional current (positive or negative) which is drawn as a consequence of faults, $F(t)$, and random noise, $n(t)$, caused by factors such as thermal effects and sampling error. For the case of a fault-free device $F(t)$ will equal zero; conversely, when multiple faults are present $F(t)$ will be a composite signal, referred to as the "fault component", representing the cumulative effect of the individual faults. The observed current signature may be given as

$$z(t) = w(t) + n(t) + F(t) \tag{1}$$

If we are able to form an estimate of $w(t)$, either through simulation or repeated observations of a "golden device", then this estimate, $\hat{w}(t)$, may be subtracted from the observed cd, leaving

$$z(t) - \hat{w}(t) = n(t) - e(t) + F(t) \tag{2}$$

where the term $e(t)$ represents the error in the estimate, and may be ignored provided enough trials are made to form an accurate estimate of the supply current drawn by a fault-free device. Consequently, the procedure of detecting a fault reduces to that of estimating $w(t)$ and determining whether $F(t)$ is equal to zero. This is equivalent to the classical signal estimation and detection problem involving non-random, unknown signals in noise.

3.2 Maximum Likelihood Estimator and Detector

With a "golden device", Equation 1 will reduce to $z(t) = w(t) + n(t)$, and the problem of estimating $z(t)$ is that of estimating an unknown signal in noise. If the signal is deterministic and the noise has a mean of zero with a Gaussian distribution, then an appropriate estimator is the maximum likelihood estimator (MLE), which is given as

$$\hat{w}_{MLE} = \frac{1}{N}\sum_{i=1}^{N} z_i t \tag{3}$$

Subtracting this estimate from the observed cd and discarding the error term, leaves $z(t) - \hat{w}_{MLE}(t) = n(t) + F(t)$, and the problem of fault detection becomes equivalent to detecting an unknown signal, $F(t)$, in noise.

Again, if the signal is deterministic and the noise is Gaussian with zero mean, an appropriate detector for $F(t)$ is the likelihood ratio test detector [19], given by

$$\lambda_t = [z(n) - \hat{w}(n)]' R_n^{-1} [z(n) - \hat{w}(n)] \tag{4}$$

with $z(n)$ representing the sampled current during application of the test patterns to an untested device, $\hat{w}(n)$ is the estimate of the current drawn by a fault-free device, and R_n^{-1} is the inverse of the noise covariance matrix. Under multiple observations of the cd, the test statistic becomes

$$\lambda_t = \sum_{i=1}^{N} [z_i(n) - \hat{w}(n)]' R_n^{-1} \left(\frac{1}{N}\sum_{j=1}^{N} [z_j(n) - \hat{w}(n)]\right) \tag{5}$$

We can now state a formal description of the test procedure which we define as "the likelihood ratio test".

1. Form the inverse noise covariance matrix, R_n^{-1}, using the acquired noise statistics.
2. Estimate the current drawn by a fault-free device, employing either simulation or a "golden device" and Equation 3.
3. Apply the appropriate test patterns to the device under test and form the test statistic λ_t according to either Equation 4 or Equation 5, depending upon whether multiple observations are available.
4. Compare λ_t to an established threshold; if the threshold is exceeded then the unit under test is classified as "faulty", otherwise it is accepted as having passed the test.

We have introduced a method for analyzing the cd, based upon knowledge derived from classical signal detection theory. We chose to model a faulty device as exhibiting a cd with an additional, but unknown, fault component. This allows greater flexibility in the types of faults which may be detected, as well as the systems to which this procedure may be applied, as no restrictions or assumptions have been made of the *form* of either the cd or the fault component. We will now evaluate the fault detection performance of the likelihood ratio test against the microprocessor faults described in Section 2.

3.3 Performance Evaluation of the Likelihood Ratio Test

It is possible to quantify the separation of the density function obtained under a simulated fault, hypothesis H_1, from the density function obtained from a fault-free system, hypothesis H_0, using a detectability index given as

$$d^2 = \frac{(\mu_{H_1} - \mu_{H_0})^2}{\sqrt{\sigma^2_{H_1}\sigma^2_{H_0}}} \tag{6}$$

where u and σ represent the parameters of the test statistic density functions. A second method of assessing fault detection capability is to use the empirical probabilities of fault detection (P_D) and false alarm (P_F), where false alarm implies the classification of a fault-free system as faulty. These probabilities are calculated by tallying the instances of correct and incorrect system classification under H_1 and H_0, respectively. Both of these methods will be used to compare the performance of the likelihood ratio test against the simulated faults under different processing environments.

As was noted in Section 2, the data faults had the smallest impact upon the supply current, while the instruction decoding faults had the largest. This observation would indicate that the more circuitry affected by the fault during execution of the test program, the greater the alteration of the supply current signature, and is supported by the indices of detectability shown in Table 1. It is important to note that it was predicted that fault-1 from the class of data storage and transfer faults would prove difficult to detect using this method, and based upon the results in Table 1, this appears to be correct. However, for all remaining faults the test algorithm yielded *complete* fault detection with no false alarms.

Table 1: Algorithm Performance versus Faults

Signature	d	μ	σ
Data Fault-1	0.07	1848	576
Data Fault-2	12.27	6115	338
Register Fault-1	80.37	1.8×10^5	2.5×10^3
Register Fault-2	57.24	1.9×10^5	5.5×10^3
Instruction Fault-1	1204.40	1.6×10^6	5.4×10^3
Instruction Fault-2	424.24	3.9×10^5	2.6×10^3

Table 2: Detectability under Subsampling

Detectability Index d for Data Fault-2					
Number of Points					
31	62	125	250	500	1000
6.71	8.56	10.58	11.41	11.98	12.27

3.3.1 Data Reduction

The utility of any test method is dependent upon the amount of resources required for implementation. Consequently, an effort was made to evaluate the effects of subsampling upon the fault detection performance.

Subsampling was accomplished by discarding evenly spaced samples from the original data. The effect of subsampling upon the detectability index is shown in Table 2 for data fault-2. As might be expected, the fewer the number of points used to make a decision, the lower the detectability index. However, if we examine the case which yielded the lowest detectability index, based upon 31 data points, we find that there was still a significant amount of separation under the two hypotheses. Assuming equal *a priori* probabilities, if one operates the detector at the threshold which yielded the minimum probability of error, then the probability of detection was 0.98 and the probability of false alarm was zero, based upon histograms of 50 reference signatures and 50 signatures under data fault-2.

3.4 Use of the Likelihood Ratio Test in System Prognosis

The previous sections chronicled the effectiveness of cd analysis for the purpose of detecting system faults. However, the increasing use of electrical systems in "critical mission" applications has created an urgent need for test methods capable of exposing *potential* faults *prior* to actual system failure. Supply current analysis in general, and the likelihood ratio test, in particular, possess several unique attributes which provide for the capability of system prognosis. This section explores the relationship between system failures, their effect upon the supply current, and the potential for system prognosis.

A fault may be defined as the alteration, in electrical behavior, of a circuit component

or signal path. "Hard" failures, such as open and short circuits, may be caused by metal migration, poor bonding, and insulator breakdown, resulting in an alteration of circuit connectivity. "Soft" failures, such as a change in component value or switching speed, may not *immediately* lead to an operational failure, yet over time may deteriorate into a hard fault which does affect the functionality of the system. However, both types of faults cause a change in the electrical current drawn by the affected subnetwork *as a function of time*. Depending upon the amount of circuitry identified with the fault, this deviation may be reflected in the power supply current signature. It is these two attributes, sensitivity to system changes and immediate observation of system behavior, that allow for system prognosis under cd analysis. Because cd analysis removes the requirement for propagating system status to observable outputs, fault prognostication may be accomplished *prior* to experiencing functional failures.

The likelihood ratio test is particularly appropriate for prognostication as it allows for the statistical comparison of present behavior, captured in the test statistic, to past behavior, represented by the distribution of the test statistic under fault-free conditions ($p(\lambda|\mathrm{H}_0)$. For purposes of system prognosis, it is possible to quantify the deviation of the present behavior from historical observations by calculating the integral

$$\int_{\lambda_t}^{\infty} p(\lambda|\mathrm{H}_0)\, d\lambda \tag{7}$$

Generally $p(\lambda|\mathrm{H}_0)$ will be represented by a histogram; thus, the integral may be calculated by tallying the number of observations for which the calculated test statistic exceeded the present test statistic and normalizing. If the result is greater than one half, then the agreement between the present cd and the reference signature is better than was normally observed. However, as this number approaches zero, indicating a strong deviation from previous system behavior, the probability of falsely classifying the system as faulty approaches zero, and the system should be taken off-line for extensive testing and examination, even if no malfunctions have been detected.

4 Autoregressive Modeling of the Supply Current Signature

It has been suggested that fault detection in digital devices might be realized through autoregressive modeling of the supply current waveform [14]. In this section we will review the theory behind autoregressive (AR) modeling and apply the technique against the simulated faults described in Section 2. Finally, we will compare the effectiveness of AR modeling against the performance of the likelihood ratio test as detailed in Section 3.3.

4.1 The Theory of Autoregressive Modeling

Autoregressive (AR) modeling is an area of time series analysis in which the time series in question is assumed to be the output of a linear system according to the following equation

$$s_n = -\sum_{k=1}^{p} a_k s_{n-k} + G u_n \tag{8}$$

where G and a_k, $1 \leq k \leq p$, are the parameters of the system, u_n is the present input, and s_n is the present output. This approach has proven useful in exposing the underlying structure of many complicated systems, ranging from the human vocal tract to wind turbulence [3]. In this particular case we intend to explore the use of the AR coefficients, a_k, as a means of compressing the information contained in the supply current signature.

Often the input signal, u_n, is unknown and it is necessary to estimate the present output as a linearly weighted summation of the past outputs

$$\tilde{s}_n = -\sum_{k=1}^{p} a_k s_{n-k} \tag{9}$$

The error of the estimate, $\tilde{s}_n$, is given by

$$e_n = s_n - \tilde{s}_n = s_n + \sum_{k=1}^{p} a_k s_{n-k} \tag{10}$$

and is typically referred to as the residual.

4.1.1 Determining the Model Order

The first step in AR modeling is the determination of the appropriate model order. Two methods are commonly used to arrive at a selection: computation of the residual variance, and analysis of the partial autocorrelation function (PACF) [1]. The former is a straightforward procedure; AR models of increasing order are successively applied against the time series under study until the variance of the resulting residuals reaches a satisfactory threshold.

An alternative method is based upon study of the partial autocorrelation function. The PACF is a plot of the correlation between observations at increasing lags, with the effects of the intervening observations removed. It has been shown, that for an AR process of order p, the PACF will cut off after lag p, where cut off implies that the function truncates abruptly with the remaining values less than twice the standard error of the coefficient estimate [2]. As a result, it is possible, through evaluation of the estimated PACF, to determine the appropriate model orders to select for experimentation.

4.1.2 Reducing Nonstationarity through Differencing

While Equation 8 is effective at modeling a wide class of times series, there are many signals which exhibit some degree of nonstationarity, indicated by a slowly decaying ACF [2].

In order to effectively model these waveforms it is necessary to first reduce the effect of the nonstationarity. This may be accomplished through suitable first-order differencing. A time series which is nonstationary in the mean may be transformed into a stationary process through the application of a single difference operator, whereas a series which is nonstationary in both the mean and the slope will require that the operation be performed twice [2].

4.1.3 Refining the System Model

Rare is the case where the scientist is presented with data which calls for a specific model order. More often, time series analysis is an iterative process, involving many attempts at improving the model performance through the selection of difference operators and model order. An initial model is formed based upon the information presented in the ACF and PACF. After this, it is necessary to analyze the residuals, using the ACF, for any remaining process structure which has not been included in the model. The model is then updated to reflect this additional information and the process repeated until the residuals resemble those of a random process.

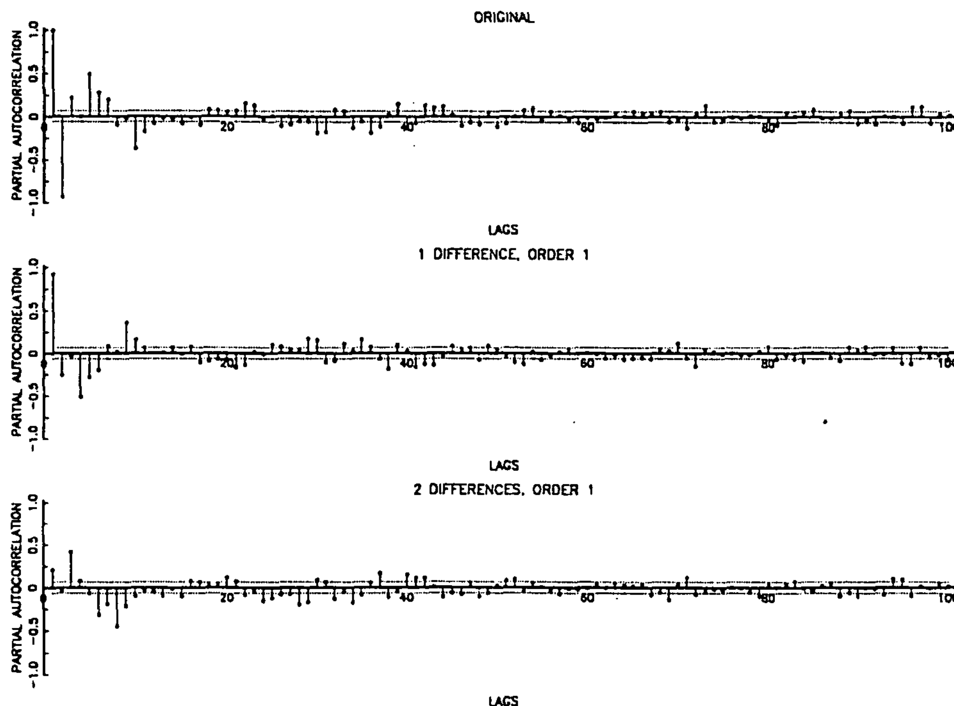

Figure 4: PACF of Original and Differenced Data Reference Signatures

4.2 Application of AR Modeling to Microprocessor Faults

We began our investigation with the cd observed during application of the data storage and transfer test program. Figure 4 shows the PACF for the original time series, as well as the first and second order differenced time series. Bearing in mind that, for an AR process

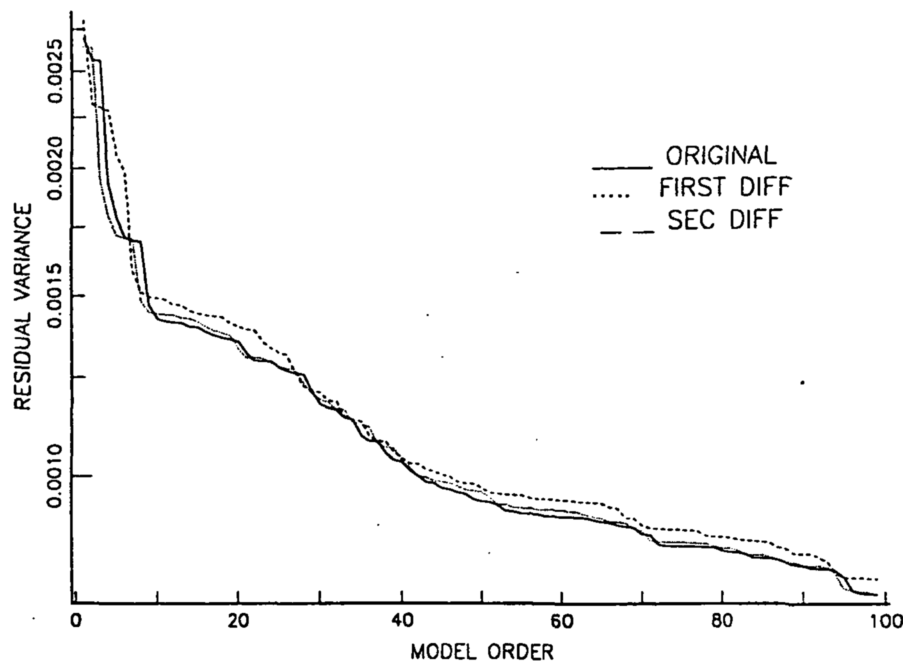

Figure 5: Residual Variance versus Model Order

of order k, the PACF will cut off after k lags, it appears that there is no clear choice of a particular model order for any of the time series. However, we can deduce that the model order must include at least ten terms. This deduction is supported by the information in Figure 5, which is a logarithmic plot of the residual variance versus model order for each of the time series. Based on these observations we selected two AR models for exploration with these time series, one of order 12 and one of order 100.

Figure 6 shows the ACF of the residuals obtained when modeling each of the time series using 12 terms. There are a significant number of coefficients which are greater than twice the standard error of the estimate, the most prominent of which occurs at lag 40. This is supported by Figure 5, where the slope of the plot seems to change slightly in that vicinity. There is also a large coefficient at lag 94. Figure 7 shows the ACF of the residuals obtained using 100 terms and we see that there are no coefficients greater than the margin of error for lags less that 100. From this we would conclude that there is no structure remaining in the process which needs to be incorporated into the model.

4.2.1 Performance Evaluation of AR Modeling

We now turn to the appraisal of autoregressive modeling using the performance metrics introduced in Section 3.3. As the test statistic, we will use the Euclidean distance between a vector containing the AR coefficients of the reference cd and a vector formed from the coefficients corresponding to the signature of the device under test (DUT). This distance

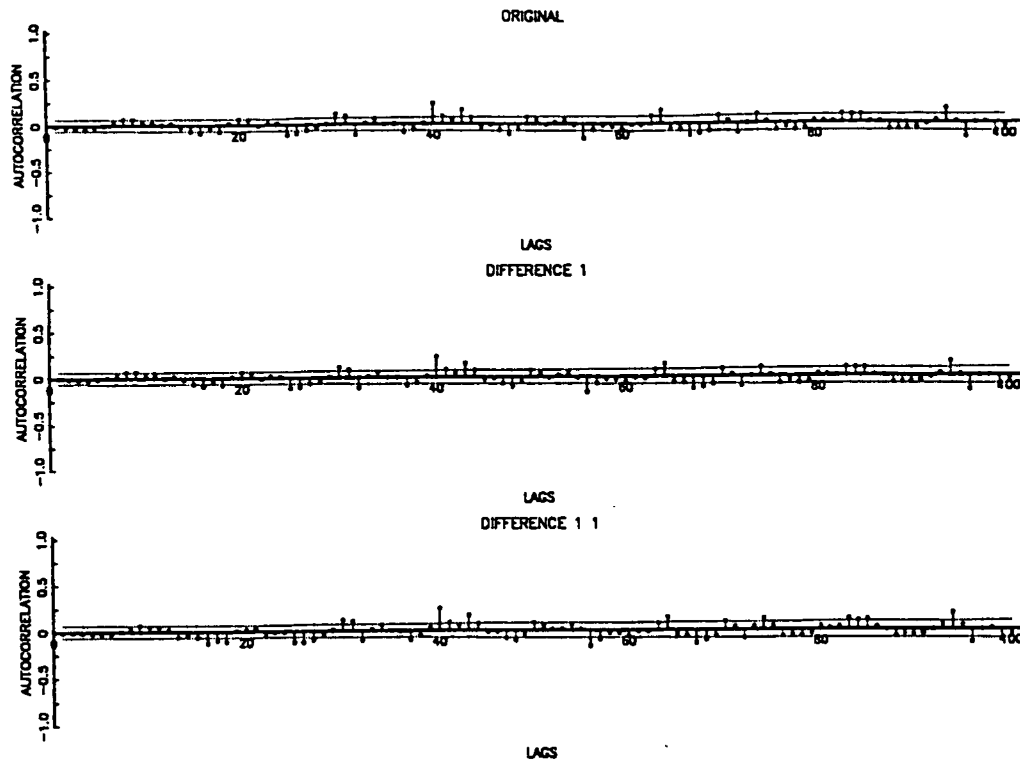

Figure 6: ACF of the Residuals for Model Order of 12

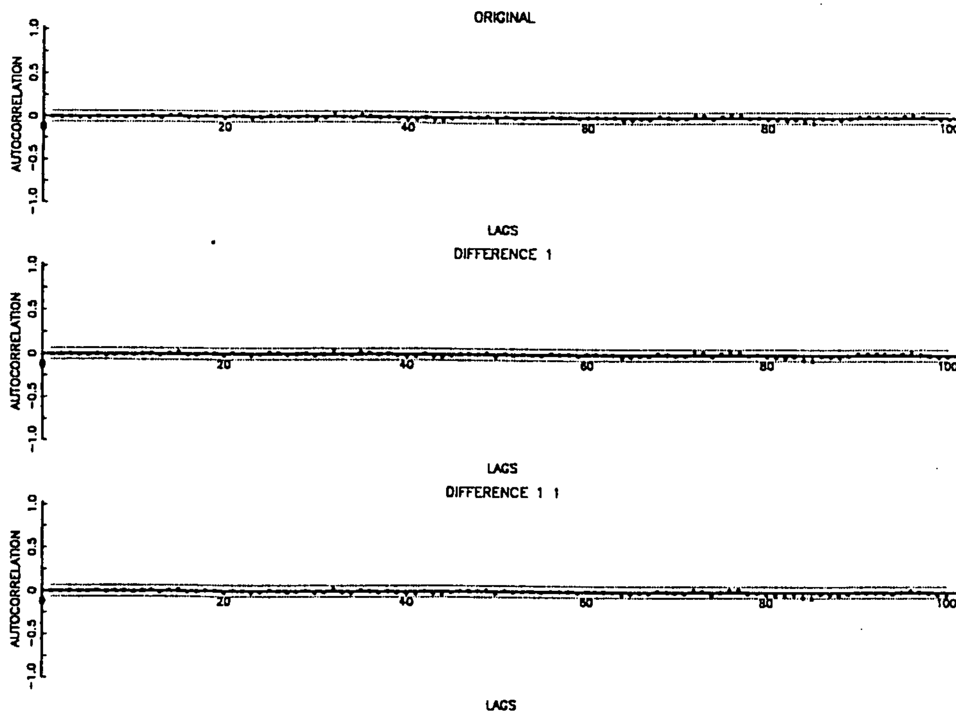

Figure 7: ACF of the Residuals for Model Order of 100

Table 3: Performance of AR modeling against Data Fault-2

Data Fault-2					
Time Series	Order	d	P_D	P_F	MPE
original	12	1.22	0.94	0.46	0.26
	100	1.44	0.72	0.18	0.23
difference (1)	12	1.14	0.86	0.38	0.26
	100	1.24	0.86	0.36	0.25
difference (1 1)	12	0.58	0.86	0.48	0.31
	100	0.44	0.84	0.44	0.30

D_E, in contrast to the λ_t used earlier, is given by

$$D_E = \sum_{k=1}^{M}(a_{Rk} - a_{Tk})^2, \tag{11}$$

where a_R and a_T correspond to the AR coefficients of the reference signature and the signature from the DUT, respectively.

Table 3 lists the results obtained with each of the time series, using AR models with 12 and 100 terms, against data fault-2. From this we can draw several conclusions: first, the application of the difference operator consistently resulted in poorer performance; second, based upon the minimum probability of error (MPE), the model with 100 terms yielded superior results against the model with 12 terms, although this effect diminished with each application of the difference operator.

Compared to the perfect fault detection demonstrated by the likelihood ratio test in Section 3.3, autoregressive modeling would appear to be a poor candidate for modeling the supply current signature in devices of this complexity. We elected to perform the comparison using data storage fault-2, which affected the cd to a lesser degree than either register or instruction decoding faults. AR modeling is normally used to characterize the spectral density of a process in a general sense, and is insensitive to minor variations. Consequently, for the sake of completeness we applied AR modeling against fault-1 of the register decoding and instruction decoding fault classes.

Table 4 shows the results of applying an AR model with 100 terms against the original supply current signatures obtained under fault-1 of the register and instruction fault classes and compares the performance to that achieved against data fault-2. Contrary to expectation, the performance of AR modeling was poorer against the register and instruction faults than the data fault, even though their effect upon the cd is much greater. One explanation for this phenomenon is that the *magnitude* of the variations is not the dominant factor in AR modeling, as it was in the likelihood ratio test, but rather the *shape* of the supply current deviations. In autoregressive modeling, the transfer function must be represented as an all-pole model, as the signal output is based only upon its previous values. This imposes restrictions upon the types of waveforms which may be accurately modeled. Although instruction fault-1 caused a supply current variation that is roughly

Table 4: Performance of AR modeling against Each Fault Class

Model Order 100				
Time Series	d	P_D	P_F	MPE
Data Fault-2	1.44	0.72	0.18	0.23
Register Fault-1	0.57	0.50	0.26	0.38
Instruction Fault-1	0.25	0.94	0.76	0.41

eight times greater than that experienced under data fault-2, its effect upon the AR coefficients was less, rendering it more difficult to detect. A contributing factor may be the effect of the noise upon the AR coefficients. It was shown in Section 3.3 that the fault detection performance could be greatly enhanced by incorporating the noise covariance matrix into the test statistic. A disadvantage of AR modeling is that there is no method for noise compensation.

4.2.2 Use of the Residual Variance

It has been reported that the residual variance may be used, in conjunction with the AR coefficients, to improve the performance of AR modeling [14], [22]. To evaluate the effectiveness of this technique when applied to cd analysis, we have repeated certain experiments using the residual variance and the AR coefficients to form the comparison vector. It was found that while the use of the residual variance consistently increased the distance between the reference signature and those of simulated faults, the contribution was minor. The maximum increase was on the order of 10^{-3}, with most of the values being on the order of 10^{-8}. Consequently, we conclude that use of the residual variance contributes very little to AR modeling in this particular application.

4.3 Summary

Based upon the experiments reported in this section, one must conclude that autoregressive modeling is not an effective technique for characterizing the information contained within the cd of a device as complex as a microprocessor. The effect of faults upon the signature is too small to be reflected in the AR coefficients to an extent that would exceed the normal deviations due to noise. However, for systems in which failures cause a drastic alteration of the cd spectrum, it is possible that AR modeling might prove useful in the area of fault diagnosis, as the system observables would be captured in a vector, rather than a single test statistic, allowing for the use of a fault dictionary.

5 Conclusions

In this paper we have presented a method of testing, referred to as *Power Supply Current Signature* (cd) *Analysis*, and demonstrated its potential for purposes of fault detection

using examples of failures in a general purpose microprocessor. A model for the experimental cd was introduced and used to develop a method of signature analysis based upon statistical signal estimation and detection, referred to as the *likelihood ratio test*. The performance of this technique was shown to be excellent at detecting all decoding faults and most data faults, and a methodology for system prognosis using the likelihood ratio was introduced. Finally, performance comparisons between the likelihood ratio test and autoregressive modeling of the cd were presented.

There are two applications for which cd analysis may be an attractive alternative to conventional testing, the first being the production testing of cost-sensitive products. Once a mature manufacturing process has been installed, cd analysis could be used in place of expensive and time-intensive testers. This would apply to both modules and boards, for cd analysis provides for the testing of mounted modules *in situ*, eliminating the need for board partitioning and module isolation. The second application involves the field-test of critical systems. Because cd analysis does not require any external observation points, systems may be tested on-line, with the application which is assigned to that board or subsystem serving as the test patterns during normal operation. Periodically, signatures could be captured and compared, using the likelihood ratio test, to those observed previously. This procedure would allow for on-line monitoring and prognostication, or failure prediction, of critical systems. In such an application, once environmental effects such as temperature fluctuation had been eliminated, any detectable cd perturbations could be directly attributed to a change in the system behavior.

Areas of future research include methods for improving the accuracy of the supply current measurement and refinement of the cd model. The concept of a built-in current sensor as proposed by Maly et al. would provide several advantages. Specifically, conversion of supply current to a voltage for off-chip measurement should provide greater immunity to system noise, thus increasing the signal to noise ratio of the cd. Furthermore, it allows for the partitioning of a VLSI module into smaller sections, providing greater distinguishability than would otherwise be possible. This concept of partitioning would involve designing for testability for cd analysis, and could be applied with similar expectations at the board level. Finally, an on-chip current sensor would provide for the implementation of the likelihood ratio test as a built-in test function.

In this paper we chose to model the cd as an unknown but nonrandom signal which could be estimated through observation of a "golden device". By subtracting this estimate from the cd of the DUT, the problem of detecting a fault reduced to detection of the unknown and nonrandom fault component. The advantage afforded by such a decision was widespread applicability across all levels of integration. However, one could chose to model the fault component as a random signal, with several uncertain parameters, such as phase and amplitude. Another possibility would be to model the cd as one of M possible signals. The task of fault detection would then be to determine which of the M possible signatures the cd of the DUT best resembled. However, this limits the number of detectable faults, or perhaps fault classes, to $(M-1)$. A more practical solution might involve modeling minor variations in the cd as uncertainties in the noise statistics.

In closing, we have presented the development of a statistical approach to fault detection

and system prognosis which has demonstrated potential at detecting faults in complex digital devices. Furthermore, no restrictions have been placed upon the nature of the system or the possible faults, other than the requirement of access to the supply current for observations. As a result of these precautions, it is hoped that this technique will prove useful in the testing of hybrid, mixed-signal systems.

References

[1] Gwilym M. Jenkins and Donald G. Watts, *Spectral Analysis and its Applications*, Holden-Day, 1968.

[2] Douglas C. Montgomery and Lynwood A. Johnson, *Forecasting and Time Series Analysis*, McGraw-Hill, 1976.

[3] John Makhoul, "Linear Prediction: A Tutorial Review," *Proceedings of the IEEE*, Volume 63, Number 4, April, 1975.

[4] Charles Crapuchettes, "Testing CMOS I_{DD} on Large Devices," *Proceedings of the International Test Conference*, pp 310–315, 1987.

[5] J.A. Starzyk and J.W. Bandler, "Nodal Approach to Multiple-fault Location in Analog Circuits," *IEEE Transactions on Circuits and Systems*, pp 1136–1139, 1982.

[6] James F. Frenzel and Peter N. Marinos, "Functional Testing of Microprocessors in a User Environment," *Proceedings of the Fourteenth International Conference on Fault-tolerant Computing*, pp 219–224, 1984.

[7] Mike Keating and Dennis Meyer, "A New Approach to Dynamic I_{DD} Testing," *Proceedings of the International Test Conference*, pp 316–321, 1987.

[8] Mark W. Levi, "CMOS is Most Testable," *Proceedings of the International Test Conference*, pp 217–220, 1981.

[9] Yashwant K. Malaiya and Stephen Y. H. Su, "A New Fault Model and Testing Technique for CMOS Devices," *Proceedings of the International Test Conference*, pp 25–34, 1982.

[10] Yashwant K. Malaiya, "Testing Stuck-On Faults in CMOS Integrated Circuits," *Proceedings of the International Conference on Computer-Aided Design*, pp 248–250, 1984.

[11] Luther K. Horning and others, "Measurements of Quiescent Power Supply Current for CMOS ICs in Production Testing," *Proceedings of the International Test Conference*, pp 300–309, 1987.

[12] J. M. Soden and C. F. Hawkins, "Test Considerations for Gate Oxide Shorts in CMOS ICs," *IEEE Design and Test*, August, pp 56–64, 1986.

[13] A. P. Dorey and others, "Reliability Testing by Precise Electrical Measurements," *Proceedings of the International Test Conference*, pp 369–373, 1988.

[14] Masaki Hashizume and others, "Fault Detection of Combinational Circuits Based on Supply Current," *Proceedings of the International Test Conference*, pp 374–380, 1988.

[15] L. Richard Carley and Wojciech Maly, "A Circuit Breaker for Redundant IC Systems," *Proceedings of the Custom Integrated Circuits Conference*, pp 27.6.1–27.6.6, 1988.

[16] Wojciech Maly and Phil Nigh, "Built-In Current Testing — Feasibility Study," *Proceedings of the International Conference on Computer-Aided Design*, pp 340–343, 1988.

[17] Luther K. Horning and others, "Measurements of Quiescent Power Supply Current for CMOS ICs in Production Testing," *Proceedings of the International Test Conference*, pp 300–309, 1987.

[18] Derek B. I. Feltham and others, "Current Sensing for Built-In Testing of CMOS Circuits," *Proceedings of the International Conference on Computer Design*, pp 454–457, 1988.

[19] James F. Frenzel, "Power Supply Current Signature Analysis: A Tool and a Methodology for use in Fault-Detection and System Prognosis," PhD thesis, Duke University, 1989.

[20] George F. Nelson and William F. Boggs, "Parametric Tests Meet the Challenge of High-Density ICs," *Electronics*, pp 108-111, December, 1975.

[21] A. J. Melia, "Supply-Current Analysis (SCAN) as a Screen for Bipolar Integrated Circuits," *Electronics Letters*, 1978, volume 14, number 14, pp 434–436.

[22] Masaki Hashizume and others, "Estimating the Level of Anesthesia by EEG Analysis," *Systems and Computers in Japan*, 1985, volume 16, number 1.

[23] James F. Frenzel and Peter N. Marinos, "Power Supply Current Signature (cd) Analysis: A New Approach to System Testing," *Proceedings of the International Test Conference*, pp 125–125, 1987.