

Predicting Device Performance From Pass/Fail Transient Signal Analysis Data

Jim Plusquellic, Amy Germida, Jonathan Hudson, Ernesto Staroswiecki, Chintan Patel
Department of CSEE, University of Maryland, Baltimore County
plusquel, germida, jhudso1, estaro1, cpatel2@cs.umbc.edu

Abstract

Transient Signal Analysis (TSA) is a Go/No-Go device testing method that is based on the analysis of voltage transients at multiple test points. In this paper, a technique based on an extension to TSA is presented that is able to predict critical path delay using data from non-critical (predictor) path tests. A characterization phase is performed a priori in which both predictor path and critical path delays are measured from a set of defect-free devices. The characterization data is used to define the relationship between the power supply transient signal data and the actual delays. Once established, prediction is performed during production test by simply re-analyzing the data from the predictor path Go/No-Go TSA tests, and therefore, no speed bin testing is required. Simulations on an 8-bit multiplier are used to demonstrate a linear relationship between a range of supply rail Fourier Phase harmonics and delay under various process models. The accuracy of the prediction is evaluated statistically against the measured delays from an additional set of critical path simulations.

1.0 Introduction

Transient Signal Analysis (TSA) is a parametric approach to testing digital integrated circuits [1][2]. In TSA, defect detection is accomplished by analyzing the transient signals measured at multiple test points on a device. The test points are located at or near the core logic supply pins on the device. By cross correlating the multiple test point measurements, a characterization profile of the device is obtained. The profile provides tolerance to process variation effects, and enhances the capacity of the method to distinguish between defect-free and defective devices.

The power supply transients capture performance attributes of the device as well as pass/fail information. In this work, the performance information is analyzed for the purposes of predicting critical path delays. The basic idea is to reuse the data collected for the predictor path Go/No-Go tests as a means of determining the maximum operational frequency of the device. This can reduce the cost and complexity of speed binning, a procedure that is rou-

tinely carried out on devices such as microprocessors.

In order to show that such information exists in the power supply transient waveforms, a set of simulations were carried out on an 8-bit combinational multiplier. In these experiments, the delays to the outputs of the multiplier were measured under 16 non-critical path test sequences. The experiments were repeated on circuit models that incorporated common variations in process parameters. The transient signals measured on ten power supply test points were analyzed in the frequency domain under each of these test sequences.

It is shown that the area under a narrow frequency band in a “calibrated” version of Fourier Phase is linearly related to the propagation delay through the device. In other words, the phase analysis of the power supply transients allows propagation delay to be estimated to the primary outputs under any test sequence. Correlation coefficients are reported to show the strength of the linear relationship between phase-area and delay in each of the experiments.

The linear property between phase-area and delay makes it possible to easily transform phase-area/delay graphs derived under predictor path tests to phase-area/delay graphs for critical paths. The latter is used to predict critical path delay using the phase-area measurements from predictor path tests. The statistical means and standard deviations of the error in the prediction of critical path delays under the various process models are reported.

The method described in this paper was developed from numerical techniques. For instance, the frequency band was determined by repeatedly applying the analysis procedure to data from different regions of the frequency spectrum. It is recognized that a supporting theoretical basis is valuable and will be addressed in a future work.

The rest of this paper is organized as follows. Section 2.0 outlines some related work. Section 3.0 describes the experimental method. Section 4.0 describes the device and simulation experiments. Section 5.0 presents the experimental results and Section 6.0 gives a summary and conclusions and discusses areas for future investigation.

2.0 Background

Huisman analyzed the variance in propagation delay along scan paths and within ring oscillators and found that




Figure 1. Time domain difference waveforms from three POs (223, 329 and 370) of seven ISCAS85 hardware devices illustrating the high degree of correlation that exists between outputs across devices.

the correlations between structurally identical paths can be low, particularly for shorter paths [3]. He extrapolates these findings to conclude that the use of an easily measurable predictor path is poor at estimating the delay of some critical path. However, he found that the correlation between a predictor path and performance is much better since the effect of uncorrelated delay variations in the performance are reduced or averaged out. He concludes that a predictor path is effective at predicting performance if it has similar “building blocks” and is sufficiently long, so that the effects of intra-chip variations can be ignored.

Eisele, *et. al.* found that the effects of uncorrelated gate delay variations (local variations) are most significant in circuits with a large number of critical paths with low logic depth [4]. They also found that the sensitivity increases for smaller device dimensions and reduced supply voltages. For data path elements in pipelined systems, which incorporate many critical paths of low logic depth, local uncorrelated variations increase the standard deviation defining the delay distribution function along these paths. This, in turn, increases the clock period required to achieve a specific yield. They conclude that testing a single critical path is not sufficient to characterize the performance of the entire device.

In previous experiments, a high degree of correlation in delay was found along paths driving the outputs of a set of small hardware devices implementing the ISCAS85 c432 benchmark circuit [5][6]. Figure 1 shows the time domain waveforms from three primary outputs (POs) on the c432 for eight devices. The topmost waveform in each of the three plots is the Reference device waveform used to compute the seven point-wise difference waveforms shown below it. The length of the sensitized path driving these POs increases from left to right in the figure, with right-most plot showing the delays along the critical path. The areas of the seven test device difference waveforms computed in the region surrounding the rising and falling edge of the Reference waveform were used as input to correlation analysis. The areas are shaded and the regions are

labeled as front edge and back edge in the left plot of Figure 1. The smallest correlation coefficient computed from the 15 scatterplots of edge/output area pairings, e.g. front edge of 223 (x-axis) vs. back edge of 329 (y-axis), was 0.95. Perfect correlation is 1.00.

The high correlation between non-critical paths (223 and 329) and the critical path (370) in these experiments suggests that uncorrelated variation does not play a significant role in these devices. This may be attributable to the long critical path and/or the older 2.0 μ m fabrication technology. As suggested by Eisele, *et. al.*, uncorrelated delay variations are less significant under these conditions. Consequently, this result may not continue to hold true for larger devices in newer technologies. However, the method presented in this work is based on power supply transients, and not on delay measurements. The global nature of the supply rail, in providing a common connection for transistors driving many disjoint logic paths, may “average out” local delay variations and reduce their impact on prediction accuracy.

In this work, we use tests that propagate signals along paths of varying lengths (from a 5 to 50 gates) as a means of predicting delay along a single critical path. Our results support the argument that critical path delay is more accurately predicted as the predictor path length increases. However, we recognize that the positive results obtained from the shorter path experiments may not hold true in hardware experiments on larger devices. Intra-chip variations are not modeled in the simulations, and therefore, cannot affect the accuracy of the predictions. The objective of this preliminary work is to establish a procedure that accurately estimates critical path delay under the premise that the long term objective of its application is in predicting performance, and not critical path delay.

3.0 Experimental Method

In previous works, TSA was shown to identify defective devices by cross-correlating the waveforms measured simultaneously at each of the core logic supply pads. TSA




Figure 2. Time and Frequency Domain Signature Waveforms.

exploits the device coupling mechanisms that allow the underlying parametric properties of the device to be conveyed to the supply rails. The most significant parametric property reflected in the supply transients is the behavior of the dynamic current. Dynamic current is responsible for charging and discharging the load capacitances of logic nodes and, consequently, has a significant influence on a related parametric property, propagation delay. Other parametric properties such as noise and steady-state current affect the supply rail transients but play a less significant role.

The theoretical basis for this work is founded on the notion that propagation delay is “encoded” in the supply rail transient signals and a clever “decoding” procedure can extract this information easily. The procedure presented in this work analyzes the Fourier phase behavior of the supply rail transients over a narrow frequency band. In these experiments, the frequency band is close to the maximum operational frequency of the devices. This suggests that the frequency band and the technology generation are related. More experiments are needed to determine if this relationship continues to hold true in other technology generations.

In order to realize the mapping between the transient signal information and critical path delay, tests that exercise both non-critical and critical paths are conducted on a sample of defect-free devices. The non-critical or predictor path tests are selected from the Go/No-Go test patterns. The propagation delays and supply rail transients are measured under each of these tests. As shown in the following sections, this data is used to define the mapping function that is used during production test to predict performance. It is important to realize that the mapping function is established *a priori* using both types of measurements, but only the supply rail transient signal measurements are required during production test. If TSA is used as part of the test suite for Go/No-Go production test, then the data from defect-free devices can be further analyzed to obtain performance information.




Figure 3. Phasor Diagram for Phase at 20MHz.

3.1 Signature Waveforms

Since a major advantage of the proposed prediction method is to reuse TSA data, the signal post-processing performed in defect-oriented TSA is reviewed here. TSA captures the change in transient behavior between devices in Signature Waveforms (SWs). SWs are created by subtracting the waveforms measured from the same test points on two devices. An example is shown in Figure 2. The V_{DD} time domain waveforms from a reference (Ref) and Test device are shown along the top on the left side of the figure. The Time Domain SW that results from the point-wise difference operation is shown along the top right, shaded to a zero baseline. The area under the curve is referred to as a Signature Waveform Area (SWA).

Although the Time Domain SWs capture the changes in the parametric behavior of the test device, it is the frequency domain SWs, more specifically, the Fourier Phase SWs, that have proven to be most effective in identifying performance differences. The bottom left side of Figure 2 shows the Fourier Magnitude and Phase “waveforms” that result from the application of a discrete Fourier transform to the Time Domain waveforms displayed along the top. The Magnitude SW is created in a manner similar to the Time Domain SW. The Phase SW is the focus of this work.

The Phase SW is also created using a point-wise subtraction but it is “calibrated” so that the phase angles in the difference waveform are made relative to the reference waveform. Figure 3 shows an example in which the reference device 20MHz phase harmonic is shifted Δ degrees forward with respect to some fundamental (not shown). An absolute point-wise difference operation applied to the Reference and a Test device phase waveforms produces a Phase SW with values covering the range of ± 360 degrees. However, the absolute difference waveform is less effective at predicting performance than the calibrated version. The calibration operation computes the relative phase shifts with respect to the phase of the reference. Figure 3 shows an example in which α and β represent the data values of the 20MHz harmonic in the Phase SWs of Test devices B and C, respectively.




Figure 4. Phase SWs at test point V_{DD1} from six transistor beta simulations of Experiment #5.

Similar to the Time Domain and Fourier Magnitude SWs, the test device Phase SWs capture phase variation relative to a reference device. It will be shown that this phase variation or shift is proportional to the difference in delay along the logic paths in the devices. Moreover, the relationship holds for a range of harmonics. By computing the area over a region of the Phase SWs, the contribution of more than one frequency component can be taken into account. In order to preserve the magnitude of the shifts across all frequency components in the region, the areas above and below the baseline are considered positive in the sum. In other words, the direction of the phase shift is ignored. This policy correctly preserves the magnitudes of opposing phase shifts in adjacent frequency components, such as those shown in Figure 4 at 16 and 20MHz.

3.2 Phase and Performance

In previous defect-oriented work, it was evident in the scatterplots that device performance was related to the area under the time and frequency SWs. Figures 4 and 5 show an illustrative example using Phase SWs. In Figure 4, six Phase SWs are shown from seven simulations (one nominal reference simulation and six test simulations). The signals were generated on a supply rail test point, V_{DD1} . Each of the six test simulations uses a circuit model in which the transistor betas are varied globally by +5% (#1), +10% (#2), +25% (#3), -5% (#4), -10% (#5) and -25% (#6) of the nominal reference value. These simulation models represent simple forms of process variation. The areas under the Phase SWs (SWAs) over the frequency range of 12MHz to 24MHz were computed. Figure 5 shows the scatterplot of the Phase SWAs for V_{DD1} and the Phase SWAs from a second test point, V_{DD3} (not shown). The data points are defined by the six corresponding pairs of V_{DD1} Phase SWAs (x-axis) and V_{DD2} SWAs (y-axis).

It is evident in the scatterplot that global process varia-

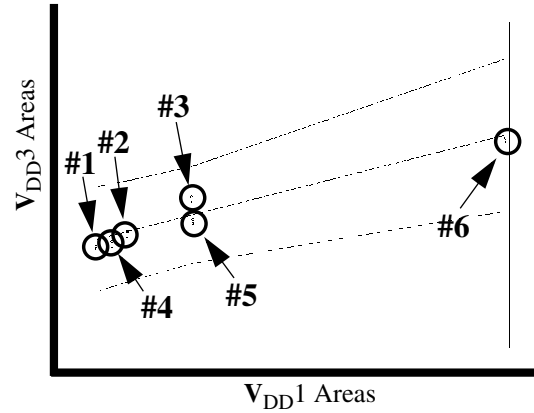


Figure 5. Scatterplot of Phase SWAs measured at two test points, V_{DD1} and V_{DD3} , from six transistor beta simulations of Experiment #5.

tions produce data points that track linearly across devices. Three sigma prediction limits are shown to delineate the (defect-free) region of interest in this work. Outlying data points (none shown) are generated by defective devices and are not considered here. However, for defect-free devices, the position of the data points within these limits is related to the performance of the device.

The scales are left unspecified in Figure 5 but it should be clear that the reference device data point (not plotted) would occur at the origin. It is also evident that the distance from the origin of each of the data points is related to the percentage change in beta, e.g. a 10% increase in beta generates a data point further removed from the origin than a 5% increase in beta. However, it is not possible to determine whether the test device is faster or slower than the reference in the scatterplot since both sets of data points (decreases in beta and increases in beta) overlap. This is an artifact of the procedure described earlier, in which all areas above and below the baseline are considered positive.

3.3 Phase Product Waveforms

The patterns in the waveforms of Figure 4 suggest a method for dealing with this problem. The top three “inc” SWs are from simulations in which the transistor betas are increased relative to the nominal value. These represent devices with superior performance over the nominal. In contrast, the bottom three “dec” SWs represent devices that are slower than the nominal. As pointed out earlier, the areas under the curve between 12 and 24 MHz increase monotonically with corresponding increases and decreases in beta. However, a second important characteristic is evident in these waveforms. The behavior or shape of the “inc” waveforms are “mirrored” across the x-axis in the “dec” waveforms. This is expected given that phase and performance are related through a continuous monotonic func-




Figure 6. Phase product waveforms using the SWs from Figure 4 and the 5%-inc Phase SW as the multiplicand.

tion, with zeros indicating a zero phase shift or match to the reference value.

Under these conditions, a simple and robust correlation procedure can be used to distinguish between devices that are slower or faster than the reference. This procedure also preserves the linear relationship between phase area and the magnitude of the performance difference. Figure 6 shows a set of Phase product waveforms computed using the Phase SWs shown in Figure 4. The products are computed using the 5%-inc SW (5% increase in beta) as the multiplicand and each of the remaining SWs as multipliers. The device used as the multiplicand will be referred to as the **calibration** device. The choice of the calibration device is arbitrary, but our analysis suggests that better results are obtained if its performance is centered between the reference and either of the two extremes delimiting the performance variation of devices in the process. In this example, a device superior in performance to the reference device is chosen. Therefore, other devices with superior performance are expected to generate positive values at each harmonic while those slower than the reference are expected to generate negative values.

Figure 6 shows that this is indeed the case. The region over which the product is computed is again delimited by vertical lines in the figure. However, in contrast with the previous defect-oriented procedure, the area below the baseline subtracts from the sum. As shown in this example, the best result is obtained when the phase product values at each harmonic are all positive or all negative. The sign of the sum is used to determine whether the test device is faster or slower than the reference and its magnitude is used to determine the relative speedup or slowdown. The multiplication provides a means of separating the data points in the scatterplot of Figure 5 into two spaces: those that are positively correlated with the calibration device

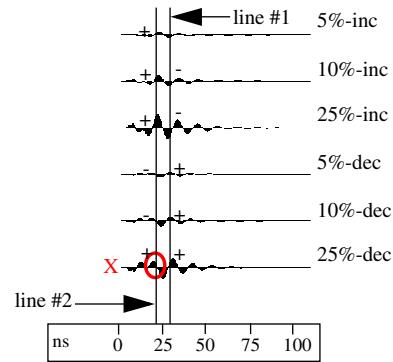


Figure 7. Time Domain SWs from six simulations at test point V_{DD1} . +/- indicates the sign of waveform at the intersection of the vertical lines #1 and #2.

and those that are negatively correlated with it.

The multiplication adds robustness to the method that is not evident from the example shown in Figure 6. In this example, all frequency components in the range analyzed are polarized in a direction consistent with the relative performance of the device. This suggests a simpler method that uses a single frequency as a means of making the decision. For example, the components along the line intersected by the 20MHz harmonic can be used to correctly predict the performance polarity with respect to the reference. However, it is possible that external factors, such as measurement noise or a significant increase in uncorrelated (local) process variation, may produce an erroneous value(s) in the SW, yielding a positive value for an otherwise slower device or vice versa. As an illustrative example, let's assume that the frequency range of interest originally extended through to the 28MHz frequency component. As shown in Figure 4, this harmonic, if used alone, would falsely predict the three "dec" waveforms as faster than the reference. However, if a range of frequencies is used instead, e.g. the 12 to 28MHz harmonic range, the positive value at 28MHz in the "dec" SWs only partially offsets the negative values of the other harmonics. Although this adds error in the prediction of the magnitude of the performance difference, it permits a correct binary decision to be made over the relative speedup or slowdown of these devices.

Similar conclusions can be made concerning an analysis method that uses the Time Domain SWs as shown in Figure 7. Although intersecting line #1 correctly predicts the behavior, intersecting line #2 falsely predicts the bottommost device as faster than the reference. Other examples indicate that choosing the line of intersection is difficult at best. Alternative techniques are under consideration, such as using the minimum and maximum values of the transient waveforms in the Time Domain SWs, but it is expected that the Phase product offers the most robust solu-




Figure 8. Phase product SWAs (y-axis) vs. relative delay (x-axis) for V_{DD6} in Experiment #6.

tion.

3.4 Linearity of Phase and Delay

In this section, we show that the Phase product SWAs vary linearly with path delays, and proceed to develop a method for the prediction of critical path delays. Given that the Phase SWs (and consequently Phase product SWAs) capture variation relative to a reference device, the analysis will also treat path delays as relative to the reference. The path delay of the reference device defines the origin in this scheme with faster devices generating negative values and slower devices generating positive values. The term “relative delay” will be used by itself in the following discussion with the qualifier “with respect to the reference” implied.

Figure 8 plots the Phase product SWAs against the relative delay to a primary output in the multiplier (described in Section 4.0.) Ten curves are shown in the graph; one for each of the power supply test points monitored in the simulations. In this example, each of the curves is composed of forty-one data points, one from each of forty-one process models in which one or more circuit or transistor parameters were varied over the range of $\pm 25\%$ of the nominal values. As noted in the figure, the variations in the process models produce a relative delay range of 15ns (-6ns to $+9\text{ns}$). The absolute path delay along the longest sensitized path is approximately 22ns in the reference device. Therefore, the modeled process parameter variations cause significant differences in the simulated device performance.

The curves in this graph illustrate a strong linear relationship between the Phase product SWAs and relative delay across the process models. Although the nearly ideal linear behavior portrayed in Figure 8 did not hold for all cases, the trend far outweighed the exceptions. Section 5.0 summarizes this result for the other experiments and supply test points through correlation analysis.

In contrast, the Time Domain analysis yields a non-linear relationship between area and delay, as shown in Fig-




Figure 9. Experiment #5 Time SWs from six simulations at test point V_{DD6} .




Figure 10. Time Domain SWAs (y-axis) vs. relative delay (x-axis) for Experiment #5.

ures 9 and 10. Figure 9 shows the Time Domain SWs from the six beta simulation experiments described above. Similar to the phase results, the amount of variation in the time domain transient is proportional to the change in beta. However, the corresponding phase-area/delay graph of Figure 10 shows the relationship is non-linear. In fact, the function appears to be non-invertible, which adds difficulty to the prediction procedure described in the next section.

3.5 Method for Predicting Critical Path Delay

The linear relationship between Phase product SWAs and relative delays suggests a straightforward means of predicting critical path delay. A set of defect-free devices is tested under a predictor path Go/No-Go test pattern and the transient signals on a supply rail test point are measured. The delays along the longest predictor path sensitized under the test sequence are measured for these devices. The reference device and calibration device are selected from the set based on the measured path delays. The performance of the reference device should be close to the mean performance of the set. The performance of the calibration device should be chosen to split the difference between the performance of the reference and the performance of either the fastest or

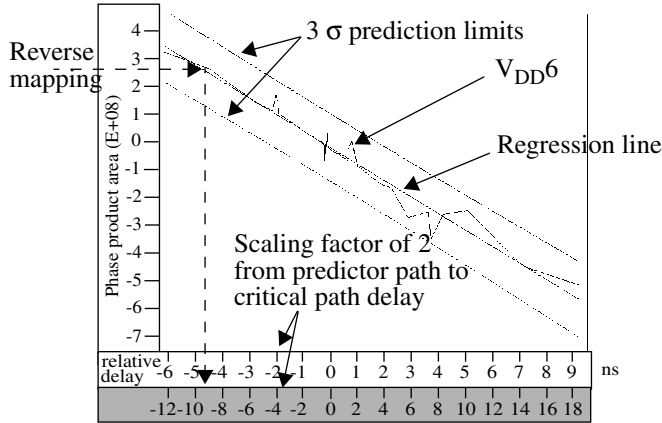


Figure 11. Predictor path Phase product SWAs (y-axis) vs. predictor path (upper x-axis) and critical path (lower x-axis) relative delay from Experiment #6.

slower device in the set. The Phase SWs are computed using a reference device and the Phase product SWAs are computed using a calibration device (as described previously). The relative delays are computed from the path delays and a graph similar to the one shown in Figure 11 is derived. A linear regression line is derived through the data points.

A critical path test is then applied to at least two devices (the reference and calibration device) and the relative delay(s) computed. An x-axis scaling factor is computed as the ratio of the calibration device's relative delays under the critical path and predictor path test sequences. The scaling factor is used to rescale the x-axis as shown along the bottom of Figure 11 for a scaling factor of 2. A more robust method, which incorporates all devices, uses the critical path regression line computed from the graph of the predictor path Phase product SWAs and corresponding critical path relative delays. In this case, it is necessary to measure the critical path delays of all devices in the set.

The rescaled predictor path graph is used to predict critical path delay for defect-free devices in production test. The Phase product SWAs are computed from the Phase SWs of the test devices and used in a reverse mapping operation on the rescaled predictor graph to derive the critical path delay. Figure 11 shows an example in which a Phase product SWA computed under the predictor test pattern yields a predicted value of -9.5ns for critical path relative delay. The critical path delay is obtained by simply adding this value to the reference device critical path delay.

4.0 Experiment Setup

The simulation experiments were conducted on a full-custom design of an 8-bit 2's complement multiplier. A block diagram of this device is shown in Figure 12. The power supplies for the core logic are labeled as V_{DD1}




Figure 12. Block-level diagram of the multiplier.

through V_{DD10} and are joined internally through a series connection of 2Ω resistors (shown on the right in the figure) to simulate the supply grid configuration of a larger device. (Core logic test pads are contact opens in the passivation layer to metal below.) The transients were measured from the Metal2 core logic test pads (shown on the left in the figure). This is representative of conducting TSA at wafer probe. The input test sequences were run at 4 MHz for a duration of 250ns.

4.1 Experiment Description

Accurate circuit models were generated using the SPACE extraction tool [7]. The lot averaged circuit parameters reported by MOSIS for hardware devices fabricated at 2μ were used to derive the technology file used by SPACE. 16 experiments are reported on in this work, each dedicated to exciting a different predictor path through the device.

A simulation run was made using a nominal defect-free simulation model for each of the 16 predictor path experiments and one for the critical path experiment. In addition, 81 simulation models were extracted from the layout and used to analyze the effects of process drift. In these models, one or more of the transistor and/or circuit parameters reported by MOSIS were varied over the range -25% and +25% of the nominal value.

Table 1 summarizes the simulation experiment models and runs. For example, in Experiments 1-3, 9-11 (column 3), models were extracted and simulated in which the parameters shown were changed individually for 30 models, and in groups of nine for 10 models by plus and minus 5%, 10% and 25% of the nominal values. Two additional simulations (Min and Max) were conducted for each of these experiments in which the nine process parameters were all set to -25% and +25% respectively to approximate the worst case process model. The simulations identified in column 2 were designated as the reference device and were

	Reference Model	Exp. Group 1	Exp. Group 2	Exp. Group 3	Critical Path Exps.
Experiments	1-17	1-3, 9-11	4-6, 12-14	7-8, 15-16	17
# of models	1	42	48	27	81
Total # of sims	17	252	288	108	81
Transistor/Circuit parameters varied by +/- 5%, 10% and 25%	None	Beta, V_t , poly Ω , metal2 contact Ω , metal cap. over p-/n-well	Beta, V_t , p-/n-diff Ω , metal1 contact Ω , poly cap. to substrate, metal1 to metal2 cap.	All 9 parameters listed to the left.	All combinations of models listed to the left.
# of circuit parameters varied per model.	None	1 (30 models). 9 (10 models). 9 at Min/Max (2 models).	1 (36 models). 9 (10 models). 9 at Min/Max (2 models).	9 (25 models). 9 at Min/Max (2 models).	1 (54 models). 9 (25 models). 9 at Min/Max (2 models).

Table 1: Simulation Experiments and Models.

used to create the Phase SWs as described in Section 3.1.

An additional 81 critical path simulations were carried out under each of the process models, as shown in the right-most column of Table 1. The critical path delays from these simulations were used to compute the prediction error as given in the next section.

5.0 Experimental Results

Two sets of results are presented in this section. As indicated in Table 1, 664 simulations were carried out under 16 predictor path test sequences using 81 process models. In each of these experiments, the transient signals were monitored on ten supply rail test points. In addition, the path delays along the longest sensitized path under each test sequence were measured. In all, 160 phase-area/delay curves were analyzed for linearity using correlation coefficients. The results are shown in Table 2. The Experiments are listed along the rows and the test points are listed across the columns. The entries greater than 90% are highlighted to illustrate the high degree of correlation that was found among many of the curves. Of the entries that are less than 90%, there are seven less than 70%. The correlation coefficients for these curves indicate poor correlation, and cannot be used for prediction. Many of the lower values appear in the right-most column. It is apparent that the circular topology of the V_{DD10} supply rail produces uncorrelated phase variations under certain test sequences.

Space limitations prevent us from showing the individual critical path predictions, which were computed using the procedure described in Section 3.5. Instead, a summary of the error in the prediction is given. A critical path regression line was derived for each of the 160 test cases as described in Section 3.5. The longest path driven by the test sequence was always chosen as the predictor path in each experiment. The error in the estimates are given only for

supply test point V_{DD6} . The correlation results given in Table 2 indicate that phase-area/delay graphs for V_{DD6} are highly correlated (as are V_{DD5} and V_{DD7}). Therefore, good results in the prediction of critical path delay are expected independent of the test sequence.

The results are shown in Table 3 which lists the experiments row-wise. The second column gives the mean magnitude of the error in delay misprediction among the entire set of process model simulations (up to 48 as indicated in Table 1). This statistic is computed using the absolute value of the difference between the predicted and measured relative delays. It is a more meaningful statistic than the true mean since prediction error can be negative or positive under each process simulation, biasing the result toward zero. The third column gives the standard deviation of the error distribution. In this case, the true mean is used in the expression. This allows the worst case error in misprediction to be estimated as approximately 3 times the standard deviation. The fourth column gives the individual worst case error in misprediction as a percentage of the critical path relative delay value (column six). The fifth column gives the range of the predictor path relative delays (the “predicted from” range) while the sixth column gives the range of critical path relative delays (the “predicted into” range). These were deemed important because some of the predictor paths are extremely short with respect to the critical path, e.g. Experiment #9. As discussed in Section 2.0, prediction error is expected to increase as the length of the predictor path decreases.

The shaded cells in the fifth column identify the experiments with “predict from” ranges greater than 9ns. These cases are expected to have the lowest misprediction statistics. The shaded cells in the second and third columns identify the experiments with mean errors and standard

Exp/Vdd	V _{DD1}	V _{DD2}	V _{DD3}	V _{DD4}	V _{DD5}	V _{DD6}	V _{DD7}	V _{DD8}	V _{DD9}	V _{DD10}
#1	.916	.904	.919	.955	.998	.999	.996	.998	.985	.997
#2	.972	.926	.926	.960	.993	.985	.908	.999	.999	.998
#3	.973	.891	.964	.993	.971	.978	.994	.997	.997	.995
#4	.962	.972	.996	.997	.996	.993	.996	.998	.996	.990
#5	.959	.162	.892	.995	.997	.999	.999	.950	.951	.995
#6	.932	.950	.796	.965	.986	.976	.998	.999	.999	.999
#7	.969	.890	.898	.688	.950	.962	.999	.999	.996	.989
#8	.868	.975	.915	.983	.955	.949	.942	.954	.989	.996
#9	.986	.955	.956	.943	.903	.868	.933	.876	.806	.183
#10	.990	.991	.974	.953	.939	.904	.845	.758	.699	.050
#11	.947	.978	.991	.986	.990	.992	.992	.983	.964	.074
#12	.900	.960	.978	.990	.989	.987	.986	.988	.987	.070
#13	.923	.923	.958	.982	.988	.984	.968	.988	.987	.773
#14	.944	.968	.965	.986	.995	.988	.995	.980	.981	.907
#15	.949	.912	.870	.958	.974	.995	.995	.984	.987	.839
#16	.825	.956	.922	.931	.937	.959	.973	.987	.978	.972

Table 2: Correlation Coefficients of Phase product SWAs vs. predictor path relative delay.

Exp	Mean Magnitude of Error in seconds	Standard Deviation of Error	Worst Case Prediction %	Predictor path relative delay range	Measured critical path relative delay range
#1	2.95e-10	4.25e-10	3.8%	6.69e-09	2.59e-08
#2	4.79e-10	8.60e-10	9.4%	2.75e-08	2.59e-08
#3	5.09e-10	1.01e-09	14.7%	1.31e-08	2.59e-08
#4	3.90e-10	5.98e-10	5.5%	1.17e-08	3.42e-08
#5	1.79e-10	2.34e-10	1.8%	3.51e-08	3.42e-08
#6	8.54e-10	1.34e-09	12.2%	1.52e-08	3.42e-08
#7	1.74e-09	2.30e-09	19.7%	1.11e-08	3.42e-08
#8	2.43e-09	2.92e-09	18.5%	7.77e-09	3.42e-08
#9	1.53e-09	2.92e-09	37.8%	1.79e-09	2.59e-08
#10	1.32e-09	2.34e-09	31.6%	6.31e-09	2.59e-08
#11	3.60e-10	5.58e-10	7.9%	9.24e-09	2.59e-08
#12	4.93e-10	8.63e-10	10.4%	1.56e-08	3.42e-08
#13	5.71e-10	9.44e-10	8.1%	1.86e-08	3.42e-08
#14	5.32e-10	8.37e-10	7.9%	1.89e-08	3.42e-08
#15	6.52e-10	8.08e-10	4.7%	1.16e-08	3.42e-08
#16	1.62e-09	2.11e-09	14.7%	7.76e-09	3.42e-08

Table 3: Error analysis of critical path delay from predictions using predictor path delays and V_{DD6}.

deviations less than 1.4ns. For ten of the eleven shaded cells in column 5, columns 2 and 3 are shaded, which asserts the higher prediction accuracy for longer predictor paths. Column 4 shows that the worst case prediction error is less than 15% for these experiments. It is also clear that the shorter predictor path experiments (#7 through #10 and

#16) yield significantly larger prediction errors. Therefore, it is important to select Go/No-Go test sequences that sensitize long predictor paths. Although prediction error was large (>30%) for two of the experiments, it is reasonably bound to less than 20% in other experiments and less than 15% for experiments with sufficiently long predictor paths.

6.0 Summary and Conclusions

A method for predicting critical path delay is proposed that is based on a simple extension of a Go/No-Go parametric testing method called Transient Signal Analysis (TSA). In the defect-oriented TSA method, the transient signals from multiple power supply test points are cross correlated to increase both the tolerance of the method to process variations and its sensitivity to defects. In the proposed performance-oriented extension to TSA, the frequency domain data from previous Go/No-Go tests is re-analyzed for performance information. The novelty of the method is the discovery that a linear relationship exists between a narrow band of supply transient Fourier Phase harmonics and the delay along logic paths of the device. A characterization process defines the mapping function between phase-areas and critical path delay using measurements from a set of defect-free devices. The mapping function is used in subsequent Go/No-Go testing as a means of predicting performance using only predictor path Go/No-Go tests. This eliminates the need to speed bin the defect-free devices.

Although it is well known that phase harmonics shift linearly with logic path delay, this property does not apply directly to power supply transient signals. The RLC coupling network of the supply rail and testing environment distort the frequency components that characterize the propagation of logic signals through the device. The procedure proposed in this work makes use of reference and calibration device behavior as a mechanism to “normalize” the phase harmonics carrying the performance information. The procedure restores the linearity and allows the prediction of delay from phase area.

The procedure is demonstrated using simulation data from an 8-bit combinational multiplier circuit. Eighty-one process models of the circuit were simulated to examine the effects of process variation on device performance. The results of correlation analysis show that the relationship between phase product area and delay is linear. Correlation coefficients greater than 80% were found in 151 of 160 experimental cases and many were greater than 95%. The error in the predicted critical path delay was subsequently analyzed using standard statistical methods. Worst case prediction error was found to be less than 15% for predictor paths lengths between 35 and 100% of the critical path length.

There are several issues that require further investigation. The relationship between the supply rail topology and the accuracy of the prediction is a complex issue that is

best resolved through more experiments. In this work, the analysis of the supply test point signals closest to the center of the design yielded the best results. A theoretical framework may aid in determining the most effective measurement locations. Similar issues exist over the structure of the underlying logic and logic signal path orientations. The multiplier circuit’s regular structure and signal path orientation may have contributed to the small prediction errors in many of the experiments. More complex designs need to be investigated to determine the suitability of the proposed procedure to accurately predict performance. Hardware experiments are currently underway to verify these results.

Acknowledgments

The authors wish to thank Fidel Muradali at the Imaging Electronics Division of Agilent Technologies for his continued support and discussions concerning this topic. We also wish to acknowledge Intel Corporation for contributing computer equipment for this research.

References

- [1] Amy Germida, Zheng Yan, J. F. Plusquellic and Fidel Muradali. “Defect Detection using Power Supply Transient Signal Analysis,” *ITC*, p. 67-76, September 1999. (<http://www.csee.umbc.edu/~plusquel>).
- [2] James F. Plusquellic, Donald M. Chiarulli, and Steven P. Levitan. “An Automated Technique to Identify Defective CMOS Devices based on Linear Regression Analysis of Transient Signal Data,” *IEEE International Workshop on IDDQ Testing*, November, 1998.
- [3] Huisman, L. M. “Correlations Between Path Delays and the Accuracy of Performance Prediction,” *International Test Conference*, p. 801-808, October 1998 .
- [4] Eisele, M. J. Berthold, D. Schmitt-Landsiedel and R. Mahnkopf, “The Impact of Intra-Die Device Parameter Variations on Path Delays and on the Design for Yield of Low Voltage Digital Circuits,” *Proceedings of the 1996 international symposium on Low power electronics and design*, 1996, Pages 237 - 242.
- [5] F. Brglez and H. Fujiwara. A neutral netlist of 10 combinational benchmark circuits and a target translator in FORTRAN. *Special Session on ATPG and Fault Simulation, Int. Symposium on Circuits and Systems*, pages 663-698, June 1985.
- [6] J. F. Plusquellic. “Digital Integrated Circuit Testing Using Transient Signal Analysis,” Ph.D. Dissertation, Department of Computer Science, University of Pittsburgh, August, 1997.
- [7] Arjan van Genderen, Nick van der Meijs, Frederik Beeftink, Peter Elias, Ulrich Geigenmuller and Theo Smedes. SPACE, Layout to Circuit Extraction software module of the Nelsis IC Design System. Delft University of Technology, 1996. (space@cas.et.tudelft.nl).