

First-Order Incremental Block-Based Statistical Timing Analysis

C. Visweswariah[†], K. Ravindran[‡], K. Kalafala^{*}, S. G. Walker[†], S. Narayan^{*}

[†]IBM Research

[‡]Department of EECS

^{*}IBM Microelectronics

T. J. Watson Research Center
Yorktown Heights, NY

University of California
Berkeley, CA

East Fishkill, NY
& Burlington, VT

ABSTRACT

Variability in digital integrated circuits makes timing verification an extremely challenging task. In this paper, a canonical first order delay model is proposed that takes into account both correlated and independent randomness. A novel linear-time block-based statistical timing algorithm is employed to propagate timing quantities like arrival times and required arrival times through the timing graph in this canonical form. At the end of the statistical timing, the sensitivities of all timing quantities to each of the sources of variation are available. Excessive sensitivities can then be targeted by manual or automatic optimization methods to improve the robustness of the design. This paper also reports the first incremental statistical timer in the literature which is suitable for use in the inner loop of physical synthesis or other optimization programs. The third novel contribution of this paper is the computation of local and global criticality probabilities. For a very small cost in CPU time, the probability of each edge or node of the timing graph being critical is computed. Numerical results are presented on industrial ASIC chips with over two million logic gates.

Categories and Subject Descriptors

B.8.2 [Hardware]: Performance and reliability—*Performance Analysis and Design Aids*

General Terms

Algorithms, performance, verification

Keywords

Statistical timing, incremental, variability

1. INTRODUCTION AND BACKGROUND

The timing characteristics of gates and wires that make up a digital integrated circuit show many types of variability. There can be variability due to manufacturing, due to environmental factors such as V_{dd} and temperature, and due to device fatigue phenomena such as electromigration, hot electron effects and NBTI (Negative Bias Temperature Instability). The variability makes it extremely difficult to verify the timing of a design before committing it to manufacturing. Nominally sub-critical paths or timing points may

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2004, June 7–11, 2004, San Diego, California, USA.

Copyright 2004 ACM 1-58113-828-8/04/0006 ...\$5.00.

become critical in some regions of the space of variations due to excessive sensitivity to one or more sources of variation. The goal of robust design, to first order, is to minimize such sensitivities.

Traditional static timing methodology is corner-based or case-based, e.g., best-case, worst-case and nominal. Unfortunately, such a methodology may require an exponential number of timing runs as the number of independent and significant sources of variation increase. Further, as described in [1], the analysis may be both pessimistic and risky *at the same time*. At corners that are timed, worst-case assumptions are made which are pessimistic, whereas, since it is intractable to analyze all possible corners, the missing corners may lead to failures detected after the manufacturing of the chip. Statistical timing analysis is a solution to these problems.

Statistical timing algorithms fall into two broad classes. The first is *path-based algorithms* wherein a selected set of paths is submitted to the statistical timer for detailed analysis. This set of methods can be thought of as “depth-first” traversal of the timing graph. Path-based statistical timing is accurate and has the ability to realistically capture correlations, but suffers from other weaknesses. First, it is not clear how to select paths for the detailed analysis since one of the paths that is omitted may be critical in some part of the process space. Second, path-based statistical timing often does not provide the diagnostics necessary to improve the robustness of the design. Third, path-based timing does not lend itself to incremental processing whereby the calling program makes a change to the circuit and the timer answers the timing query incrementally and efficiently [2]. Finally, path-based algorithms are good at taking into account *global* correlations, but do not handle independent randomness in individual delays. Doping effects and gate oxide imperfections are usually modeled as uncorrelated random phenomena. In fact, few if any statistical timing attempts in the literature include support for *both* correlated and independent randomness.

The statistical timer described in this paper belongs to the second class of statistical timers, namely *block-based statistical timers*. This set of methods traverses the timing graph in a leveled “breadth-first” manner. In [3], probability distributions are assumed to be trains of discrete impulses which are propagated through the timing graph. However, correlations both due to global dependencies on the sources of variation and due to path-sharing are ignored, as is the case with [4]. In this same general framework, [5] describes how correlations due to reconvergent fanout can be taken into account, but not dependence on global sources of variation. In [6], an approximate block-based statistical timing analysis algorithm is described to reduce pessimism in worst-case static timing analysis. The concept of parameterized delay models is proposed. Recently, [7, 8] focus on handling spatial correlations due to intra-die variability. Unfortunately, all these efforts suffer from some weaknesses. First, they do not provide diagnostics that can be used by a human designer or synthesis program to make the circuit more robust. Second, they are not immediately amenable to incremental processing. Third, they do not provide for a general enough timing model to accommodate correlation due to dependence on

common global sources of variation, independent randomness and correlation due to path sharing or reconvergent fanout. This paper describes a statistical timing algorithm that possesses the following strengths.

1. A canonical first-order delay model is employed for all timing quantities. The model allows for both global correlations and independent randomness. Thus timing results such as arrival times and slacks are also available in this canonical form, thereby providing first-order sensitivities to each of the sources of variation. These diagnostics can be used to locate excessive sensitivity to sources of variation and to target robust circuit designs by reducing these sensitivities.
2. The statistical timing algorithm is approximate, but has linear complexity in the size of the circuit and the number of global sources of variation. The speed of the algorithm and its block-based nature allow the tool to time very large circuits and *incrementally* respond to timing queries after changes to a circuit are made. To the best of the authors' knowledge, this is the first incremental statistical timer in the literature or industry.
3. The algorithm computes, with a very small CPU overhead, local and global criticality probabilities which are useful diagnostics in improving the performance and robustness of a design.

2. CANONICAL DELAY MODEL

All gate and wire delays, arrival times, required arrival times, slacks and slews (rise/fall times) are expressed in a standard or canonical first-order form below:

$$a_0 + \sum_{i=1}^n a_i \Delta X_i + a_{n+1} \Delta R_a, \quad (1)$$

where a_0 is the mean or nominal value, $\Delta X_i, i = 1, 2, \dots, n$ represent the variation of n global sources of variation $X_i, i = 1, 2, \dots, n$ from their nominal values, $a_i, i = 1, 2, \dots, n$ are the sensitivities to each of the global sources of variation, ΔR_a is the variation of an independent random variable R_a from its mean value and a_{n+1} is the sensitivity of the timing quantity to R_a . By scaling the sensitivity coefficients, we can assume that X_i and R_a are unit normal or Gaussian distributions $N(0, 1)$. Not all timing quantities depend on all global sources of variation; in fact [7, 8] suggest methods of modeling ACLV (Across-Chip Linewidth Variation) by having delays of gates and wires in physically different regions of the chip depend on different sets of random variables. In chips with voltage islands, the delay of an individual gate will depend only on the variability of the power supply of the island in which it is physically located.

3. THE CONCEPT OF TIGHTNESS PROBABILITY

Given any two random variables X and Y , the *tightness probability* T_X of X is the probability that it is larger than (or dominates) Y . Given n random variables, the tightness probability of each is the probability that it is larger than all the others. Tightness probability is called *binding probability* in [9, 10]. The tightness probability of Y , T_Y is $(1 - T_X)$. Below we show how to compute the max of two timing quantities in canonical form and how to determine their tightness probabilities. Given two timing quantities

$$A = a_0 + \sum_{i=1}^n a_i \Delta X_i + a_{n+1} \Delta R_a, \text{ and} \quad (2)$$

$$B = b_0 + \sum_{i=1}^n b_i \Delta X_i + b_{n+1} \Delta R_b, \quad (3)$$

their 2×2 covariance matrix can be written as $cov(A, B) =$

$$\begin{bmatrix} a_1 & b_1 \\ a_2 & b_2 \\ \vdots & \vdots \\ a_n & b_n \\ a_{n+1} & 0 \\ 0 & b_{n+1} \end{bmatrix} [V] \begin{bmatrix} a_1 & b_1 \\ a_2 & b_2 \\ \vdots & \vdots \\ a_n & b_n \\ a_{n+1} & 0 \\ 0 & b_{n+1} \end{bmatrix}, \quad (4)$$

where V is the covariance matrix of the sources of variation. Assuming that the X_i are independent random variables for the purposes of illustration, V is the unity matrix, and thus $cov(A, B) =$

$$\begin{bmatrix} \sum_{i=1}^{n+1} a_i^2 & \sum_{i=1}^n a_i b_i \\ \sum_{i=1}^n a_i b_i & \sum_{i=1}^{n+1} b_i^2 \end{bmatrix} = \begin{bmatrix} \sigma_A^2 & \rho \sigma_A \sigma_B \\ \rho \sigma_A \sigma_B & \sigma_B^2 \end{bmatrix}. \quad (5)$$

By comparing terms in (5) above, σ_A , σ_B and the correlation coefficient ρ can be computed in linear time. Now we seek to determine the distribution of $\max(A, B)$ and the tightness probabilities of A and B . We appeal to [11, 12] for analytic expressions to solve this problem. Define

$$\phi(x) \equiv \frac{1}{\sqrt{2\pi}} \exp(-\frac{x^2}{2}) \quad (6)$$

$$\Phi(y) \equiv \int_{-\infty}^y \phi(x) dx \quad (7)$$

$$\theta \equiv (\sigma_A^2 + \sigma_B^2 - 2\rho \sigma_A \sigma_B)^{1/2}. \quad (8)$$

Then, the probability that A is larger than B is

$$\begin{aligned} T_A &= \int_{-\infty}^{\infty} \frac{1}{\sigma_A} \phi\left(\frac{x-a_0}{\sigma_A}\right) \Phi\left(\frac{\left(\frac{x-b_0}{\sigma_B}\right) - \rho \left(\frac{x-a_0}{\sigma_A}\right)}{\sqrt{1-\rho^2}}\right) dx \\ &= \Phi\left(\frac{a_0 - b_0}{\theta}\right). \end{aligned} \quad (9)$$

The mean and variance of $\max(A, B)$ can also be analytically expressed as

$$\begin{aligned} E[\max(A, B)] &= a_0 T_A + b_0 (1 - T_A) + \theta \Phi\left(\frac{a_0 - b_0}{\theta}\right), \\ var[\max(A, B)] &= (\sigma_A^2 + a_0^2) T_A + (\sigma_B^2 + b_0^2) (1 - T_A) + \\ &\quad (a_0 + b_0) \theta \Phi\left(\frac{a_0 - b_0}{\theta}\right) - \{E[\max(A, B)]\}^2. \end{aligned} \quad (10)$$

Thus, the tightness probabilities, expected value and variance of $\max(A, B)$ can be computed analytically and efficiently. Similar formulas can be developed for $\min(A, B)$. The CPU time of this operation increases only linearly with the number of sources of variation.

Tightness probabilities have an interpretation in the space of the sources of variation. If one random variable has a 0.3 tightness probability, then in 30% of the weighted volume of the process space it is larger than the other variable, and in the other 70%, the other variable is larger. The weighting factor is the joint probability density function (JPDF) of the underlying sources of variation.

4. BLOCK-BASED STATISTICAL TIMING: THE KEY IDEA

To apply these ideas to static timing, we need probabilistic equivalents of the “max,” “min,” “add” and “subtract” operations. Addition and subtraction of two quantities in canonical form is easy, so we focus here on the max and min operations. Consider two edges of a timing graph that suggest two arrival times A and B at a node on which they are incident. Using the formulas of the previous section, we seek to express $C = \max(A, B)$ back into canonical form for further correlated propagation through the timing graph. The concept of tightness probability helps us in this difficult step.

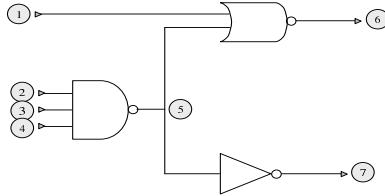


Figure 1: Sample circuit.

From (9) and (10), we know the mean and variance of C . In traditional static timing, C would take the value of the larger of A and B , and for all downstream purposes, the characteristics of the dominant edge that determined the arrival time C are preserved, and the other edge is ignored. This is like having a tightness probability of 100% and 0%. In the probabilistic domain, the characteristics of C are determined from A and B in the proportion of their tightness probabilities. Thus if the probabilities were 0.75 and 0.25, the sensitivities of A and B would be linearly combined in a 3 : 1 ratio to obtain the sensitivities of C . Mathematically,

$$c_i = T_A a_i + (1 - T_A) b_i, i = 1, 2, \dots, n, \quad (11)$$

where T_A is the tightness probability of A .

The mean of the distribution of $\max(A, B)$ is preserved when converting it to canonical form. The only remaining quantity to be computed is the independently random part of the result. This is done by matching the variance of the canonical form to the variance computed analytically from (10). Thus the first two moments of the real distribution are always matched in the canonical form. Interestingly, the coefficients computed in this manner preserve the correct correlation to the global sources of variation as suggested by [11] and are similar to the coefficients computed in [7].

The max of two Gaussians is not a Gaussian, but we re-express it in the canonical Gaussian form and incur an accuracy penalty for doing so. However, this step allows us to keep alive and propagate correlations due to dependence on the global sources of variation, which is absolutely key to performing timing in a realistic fashion. Monte Carlo results will be shown in the results section to assess the accuracy of this method.

When more than two edges of the graph converge at a node, the max or min operation is conducted one pair at a time, just as with deterministic quantities. The tightness probabilities are treated as conditional probabilities and post-processed to compute the final tightness probability of each arc incident on the node whose arrival time is being computed. As more equally critical signals are max'ed, accuracy degrades slightly since the asymmetry in the resulting probability distribution increases, making it harder to approximate in canonical form.

Slews (rise/fall times) are propagated in much the same manner. If the policy is to propagate the worst slew, then a separate tightness probability is computed for the slews and applied to represent the bigger slew in canonical form. If the policy is to propagate the latest arriving slew, then the same arrival tightness probabilities are applied to combine the incoming slews to obtain the output slew.

In this manner, by replacing the “plus,” “minus,” “max” and “min” operations with probabilistic equivalents, and by re-expressing the result in a canonical form after each operation, regular static timing can be carried out by a standard forward and backward propagation through the timing graph [13]. Early and late mode, separate rise and fall delays, sequential circuits and timing tests are therefore easily accommodated just as in traditional timing analysis.

5. CRITICALITY COMPUTATION

The methods presented in the previous section enable statistical timing analysis, during which the concept of tightness probability is

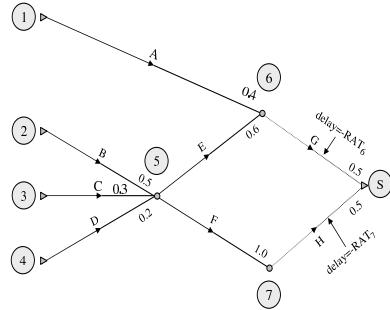


Figure 2: Timing graph of the sample circuit.

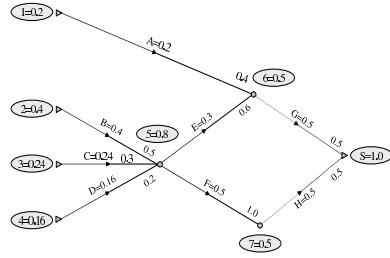


Figure 3: Backward traversal of the timing graph.

leveraged to propagate arrival and required arrival times in a parametric canonical form. In this section, the use of tightness probabilities in computing criticality probabilities [14] is presented. One of the important outcomes of deterministic timing is the ability to find the most critical path. In the statistical domain, the concept of the most critical path is probabilistic. The criticality probability of a path is the probability that the path is critical; the criticality probability of an edge is the probability that the edge lies along a critical path; and the criticality probability of a node is the probability that a critical path passes through that node. Computing these probabilities will obviously have important benefits in enumerating critical paths, enabling robust optimization and generating test vectors for at-speed test.

5.1 Forward propagation

The ideas behind criticality computations are described by means of an example. Consider the combinational circuit of Fig. 1. In this example, separate rising and falling delays and slew effects are ignored for simplicity, but the ideas can be extended in a straightforward manner. Likewise, sequential circuits pose no special problem. The example assumes late-mode timing, but early-mode follows the same reasoning.

The timing graph of the circuit is shown in Fig. 2. During the forward propagation phase of timing analysis, each edge of the timing graph is annotated with an *arrival tightness probability* (ATP), which is the probability that the edge determines the arrival time of its sink node. The ATPs in this example have been chosen arbitrarily, and are shown at the tail of each edge of the timing graph. Once the primary outputs are reached, a virtual output edge is added from each primary output to a sink node, shown as edges G and H in Fig. 2. Each such edge is considered to have a delay equal to the negative of the asserted required arrival time at the corresponding primary output. In the presence of timing tests (such as setup, hold or clock pulse width tests), a virtual edge is added to the sink node whose delay is the negative of the computed required arrival time. Then the standard forward propagation procedure is continued to compute the “arrival time” of the sink of the graph, and the ATPs of the virtual output edges. In this case, for illustration purposes, the ATP of each of the virtual output edges is 0.5.

Property 1: The sum of the ATPs of all edges incident on any

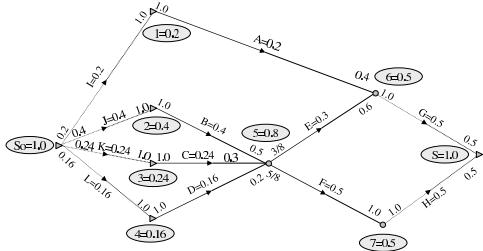


Figure 4: Source node of the timing graph.

node of the timing graph is 1.0.

Property 2: The probability of a path being critical is the product of the ATPs of all edges along the path. For path 2B5E6GS to be critical, for example, edge B has to determine the arrival time of node 5 (probability=0.5), edge E has to determine the arrival time of node 6 (probability=0.6) and edge G has to determine the arrival time of node S (probability=0.5), for a total probability of 0.15, assuming independence between these events.

Property 3: The sum of the criticality of all paths in a timing graph is 1.0.

5.2 Backward propagation

Fig. 3 shows the criticality calculations during the backward propagation phase of timing analysis. During the backward propagation, we will compute the *global criticality* of each edge and each node of the timing graph, and the *required arrival tightness probability* (RATP) of each edge of the timing graph, which is the probability that the edge determines the required arrival time of its source node.

Property 4: The sink node has a node criticality probability of 1.0. This property is obvious since all paths must pass through the sink node. The sum of the ATPs of the virtual output edges is therefore also 1.0.

Starting with the sink node S, the backward propagation first considers edges G and H. They each have a 0.5 edge criticality since they each determine the arrival time of S with 0.5 probability. The criticality of nodes 6 and 7 are likewise 0.5 each.

Property 5: The criticality of an edge is the product of its ATP and the criticality probability of its sink node. Clearly, an edge is globally critical only if its sink node is critical and it determines the arrival time of that sink node.

Property 6: The criticality of a node in the timing graph is the sum of the criticality of all edges leaving that node. Using the above two properties, the criticalities of edges and nodes are easily computed during a leveled backward traversal of the timing graph, and are shown in Fig. 3. The criticality computations can piggy-back on top of the usual required arrival time calculations. Note that the criticality of edge A, for example, is the product of the criticality of node 6 (0.5) and the ATP of edge A (0.4). The criticality of node 5, for example, is the sum of the edge criticalities of edges E and F.

Property 7: The sum of the node criticalities of all the primary outputs is 1.0. For general sequential circuits, this property would apply to all slack-determining end-points (primary output and timing test points).

Property 8: The criticality of any node in the timing graph is the sum of the path criticalities of all paths in its fanout cone. For example, node 5 has two paths in its fanout cone, path 5E6GS with a path criticality of 0.3 and path 5F7HS with a path criticality of 0.5, totaling to a node criticality of 0.8 for node 5.

As the backward propagation progresses, *required arrival tightness probabilities* (RATPs) are computed and annotated on to the timing graph. These probabilities are shown close to the source node of each edge in Fig. 4.

Property 9: The sum of the RATPs of all edges originating at

any node of the timing graph is 1.0. At a node such as 5 where there are multiple fanout edges, the RATPs will be in the proportion of the edge criticalities of the downstream edges. When the primary inputs are reached during backward traversal, a new node of the timing graph called the source node is postulated, with virtual input edges from the source node to each of the primary inputs, shown as edges I, J, K and L in Fig. 4. Each virtual input edge is considered to have a delay equal to the arrival time of the corresponding primary input, and the required arrival time of the source node is computed. During this computation, the RATPs of the virtual edges are also determined.

Property 10: The ATPs of each of the virtual input edges is 1.0.

Property 11: The criticality of the source node is 1.0. This property is obvious since every path passes through the source node.

Property 12: The sum of the node criticalities of all the primary inputs is 1.0.

Property 13: The sum of the edge criticalities of the virtual input edges is 1.0 as is the sum of their RATPs.

Property 14: The criticality of any path is the product of the RATP of all edges of the path. Thus the criticality of path SoJ2B5E6GS is $0.4 \times 1.0 \times 3/8 \times 1.0 = 0.15$.

Property 15: The product of the ATPs along any path of the graph is equal to the product of the RATPs.

Property 16: The criticality of an edge is the sum of the criticality of all paths through that edge.

Property 17: The sum of the edge criticalities of any cutset of the timing graph that separates the source from the sink node is 1.0. In other words, any cut through the graph that leaves the source node on one side and the sink node on the other will cut edges whose criticality probabilities sum to 1.0. This must be the case since every critical path will have to pass through exactly one edge of the cutset.

It is important to note that the edge and node criticalities can be computed on a global basis, or on a per-end-point basis, where an end point is a slack-determining node of the graph (a primary output or either end of a timing test segment). The application will dictate which type of computation is more efficient and suitable.

5.3 Path enumeration

Enumeration of paths in order of criticality probability is useful in a number of different contexts, such as producing reports, providing diagnostics to the user or a synthesis program, listing paths for test purposes, listing paths for CPPR (Common Path Pessimism Removal) purposes [15], and enumerating paths for analysis by a path-based statistical timer [10]. One straightforward manner of enumerating paths is by means of a breadth-first visiting of the nodes of an augmented graph as shown in Fig. 4, while following the unvisited node with the highest criticality probability at each juncture. A running total of the criticality probability of the listed paths is maintained, and the path enumeration stops when the set of critical paths has been covered with a certain confidence.

During the path enumeration, the following properties are useful.

Property 18: The ATP of an edge is an upper bound on the criticality of any path that passes through that edge.

Property 19: The RATP of an edge is an upper bound on the criticality of any path that passes through that edge.

Property 20: The criticality probability of an edge is an upper bound on the criticality of any path that passes through that edge.

Property 21: The criticality probability of a node is an upper bound on the criticality of any path that passes through that node.

6. INCREMENTAL STATISTICAL TIMING

Optimization or physical synthesis programs often call an incremental timer millions of times in their inner loop. To suit this purpose, a statistical timer needs to incrementally and efficiently answer timing queries after one or more changes to the circuit has

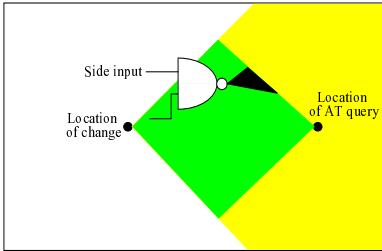


Figure 5: Incremental timing analysis.

been made.

Consider the situation shown in Fig. 5. Assume a single change has been made to the circuit at the location shown. The change could be the addition of a buffer, the resizing of a gate, the removal of a latch, and so on. Assume that the calling program queries the timer for the arrival time at the “Location of AT query” point. Clearly, only the arrival times in the yellow cone of logic change (on black-and-white hardcopies, the lightest grey region). Further, only arrival time changes in the fanin cone of the query point can have an effect on the query. The intersection of these two cones of logic is shown in green (or the darker grey region). Thus by purely topological reasoning, the portion of the circuit that must be re-timed to answer this query is limited. This kind of limiting is called *level-limiting* and is accomplished by storing AT, RAT and AT-RAT levels for each gate [2]. The leveling and limiting procedures are identical for the statistical timing situation, and the implementation can easily ride on top of an existing deterministic incremental capability.

In addition to level-limiting, the amount of re-computation can be further reduced by *dominance-limiting*. Consider the NAND gate shown in Fig. 5. One input of the NAND gate is from the “changed” cone of logic and the other from an unchanged region. If the arrival time at the output of the NAND gate is unchanged because it was determined both before and after the change by the side input, then the fanout cone of the NAND gate (shown in dark black in Fig. 5) can potentially be skipped in answering the query. This type of limiting is called *dominance-limiting*. In our statistical timer, the notion of “change” is treated probabilistically by examining the tightness probabilities. If the ATP of the side input is sufficiently close to 1.0 both before and after the change, then the arrival time of the output of the NAND gate need not be recomputed, and its fanout cone can potentially be skipped until some other input of that fanout cone is known to have materially changed. Similar concepts are applicable during backward propagation of required arrival times.

Of course, there are several complications that must be faced in a real application such as slew propagation, latches, multiple clock phases and phase changes, and the dynamic adaptation of data structures to such changes. These details are omitted due to lack of space.

7. IMPLEMENTATION

The above ideas have been implemented in a prototype called EinsStat. EinsStat is implemented on top of the static timing analysis program EinsTimer in C++ with TCL scripting under Nutshell. Multiple clock phases, phase renaming, rule tests (such as setup and hold tests), automatic tests (such as clock gating, clock pulse width and clock inactive tests), arbitrary timing assertions and timing adjusts anywhere in the timing graph, and clock overrides are supported as in EinsTimer. The timer works permanently in incremental mode [16], even if a complete timing report is requested.

Each timing assertion, gate delay, wire delay and timing test guard time must be modeled in canonical form, i.e., with a mean part, a dependence on global sources of variation and an independent random portion. The EinsStat implementation allows each

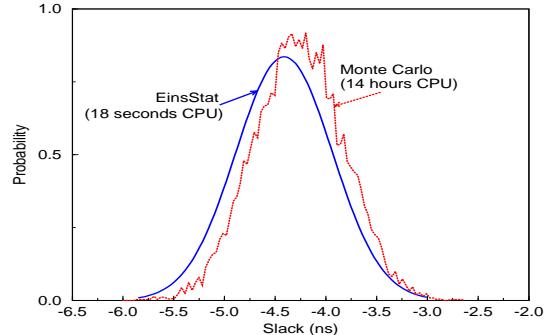


Figure 6: EinsStat vs. Monte Carlo analysis on chip “A.”

gate and each wire to have its own customized variability model, provided the model can be expressed in the canonical form. For testing purposes, however, three global sources of variation were implemented. The first is gate vs. wire delays. Each of these sets of delays can have an independent and correlated variability, and a mistrack coefficient. In the case of gate vs. wire delays, mistrack implies that when gates get faster, wires get slower, and vice versa, and in general expresses correlations between the two sets of delays. The second supported global source of variation is rise vs. fall delays of gates (to model N/P mistrack due to manufacturing variations or fatigue effects). Again, each of these can have a random and correlated part and a mistrack coefficient. The third supported source of variation is meant similarly to study mistrack between normal V_t and low V_t gates. In the benchmark results presented in the next section, sensitivities to these three global sources of variation were provided in a blanket fashion as a percentage of the nominal delay. EinsStat supports TCL commands to express a situation in which, for example, “all normal V_t gates have a 1% independent randomness and a 4% correlated variability, all low V_t gates have a 2% independent randomness and 5% correlated variability, and the two sets of variations mistrack with respect to each other.”

8. NUMERICAL RESULTS

A set of industrial ASIC designs was timed with 3 global sources of variation as well as independent randomness built into every edge of the timing graph. The benchmark results are shown in Table 1, in which the chips are code named A, B, etc., to preserve confidentiality. “Propagate segments” is the number of edges in the timing graph with unique source-sink pairs of nodes. “Load” is the CPU time to load the netlist, timing rules and assertions. “EinsTimer” is the CPU time of the deterministic base timer, while the “EinsStat” column shows the CPU time taken when the statistical timer runs alongside (and in addition to) the deterministic timer. All CPU times were measured on an IBM Risc/System 6000 model 43P-S85 on a single processor. All timing runs included forward propagation of early and late arrival times, and reverse propagation of early and late required arrival times. Similarly, the memory consumption to load each design, assertions and delay models (Base), run deterministic timing (EinsTimer) and statistical timing alongside (and in addition to) deterministic timing (EinsStat) are shown in subsequent columns of Table 1. The CPU and memory overhead of statistical timing are very reasonable, considering the wealth of additional data being generated. In the small test case A, memory consumption was dominated by the delay models, so the overhead due to statistical timing was dwarfed. In test case E, the larger overhead was due to nodes in the timing graph having extremely high incidence due to SoC timing macromodels.

The statistical experiments were performed both with and without criticality computations, and the CPU time and memory overhead were observed to be nearly identical (within 1%), lending credence to the efficiency of the criticality computations.

The primary goal of EinsStat is to produce timing results in a

Table 1: CPU and memory results.

Name	Gates	Clock domains	Propagate segments	CPU time (secs.)			Memory (MB)		
				Load	EinsTimer	EinsTimer + EinsStat	Base	EinsTimer	EinsTimer + EinsStat
A	3,042	2	17,579	5.1	2.8	3.8	111	53	60
B	183,186	79	959,709	140.5	121.3	187.6	423	177	723
C	1,085,034	182	5,799,545	5131.5	809.9	1233.1	3200	600	4300
D	1,213,361	18	6,969,860	783.5	1079.3	1485.7	2990	1160	4380
E	2,095,176	51	13,460,759	1494.9	1316.9	2724.3	4590	3320	11330

parameterized form, and therefore to give the designer information regarding the robustness of the design. However, EinsStat produces these timing results as random variables, and the correctness of the mean and spread of these random variables can be verified by Monte Carlo analysis. To render the analysis tractable, EinsStat makes a number of assumptions that prevent it from obtaining the exact result. Inaccuracy creeps in every time the probability distribution resulting from a max or min operation is re-expressed in canonical form. Specifically, the max or min of two Gaussians is not Gaussian, but EinsStat forces it back into a Gaussian form. The extent of these inaccuracies is revealed by Monte Carlo analysis.

Test chip “A” (3,042 logic gates) was used to demonstrate the importance of global correlations, and to compare EinsStat results with Monte Carlo results. The critical path in this chip is a long combinational path passing through about 60 stages of logic, with a nominal delay of 23.06 ns including wire delay. With 5% *correlated* variability on every gate and wire delay, the longest path delay is 23.01 ns with a σ of 0.9 ns. With 5% *independent* variability on every gate and wire delay, the longest path delay is 23.62 ns with a σ of 0.13 ns. Clearly, with more independent randomness, there is more cancellation of variability along a long path, yielding a tighter distribution but with a more pessimistic mean. The correlated case produces a more optimistic mean path delay, but with a much bigger spread. EinsStat allows the modeling of these extreme situations and anything in-between.

Test chip “A” was analyzed both by EinsStat and by Monte Carlo analysis with 10,000 samples. Of the 47,048 unique slacks to choose from, the comparison is shown on one representative slack, that of the nominally critical end-point. Fig. 6 shows the comparison between EinsStat and Monte Carlo. The mean value, spread and tails are predicted with reasonable accuracy. The Monte Carlo analysis required 14 hours of CPU time, while EinsStat required 18 seconds on the same computer.

A repowering experiment on chip “A” was used to evaluate incremental operation of EinsStat. For each of 493 gates with negative slack, the gate power level (size) was modified, and EinsStat was queried for the new slack on each pin of the modified gate. Incremental EinsStat was 6 times faster than non-incremental EinsStat with identical results. For large designs and for different types of changes and queries, we expect the run time improvement obtained by incremental processing to be quite dramatic.

9. FUTURE WORK AND CONCLUSIONS

This paper presents a novel statistical timing algorithm which propagates first-order sensitivities to global sources of variation through a timing graph. Each edge of the timing graph is modeled by a canonical delay model that permits global dependence as well as independent randomness. The timing results are presented in a parametric form, which can help a designer or optimization program target robustness in the design. A novel theoretical framework for computing local and global criticality probabilities is presented, thus providing detailed timing diagnostics at a very small cost in run time.

The following avenues of future work suggest themselves. The assumption of linear dependence of delay on each source of variation is valid only for small variations from nominal behavior. Extending the theory to handle general nonlinear models and asym-

metric distributions would be a big step forward. Second, the impact of variability of input slews and output loads on the delay of timing graph edges can be chain-ruled into the canonical delay model as suggested in [6]. Finally, the criticality computations in this paper assume independence between the criticality probabilities of any two paths, an assumption that is valid to first order, but not quite correct. Extending the theory to remove dependence on this assumption is a challenging task.

10. ACKNOWLEDGMENTS

The authors would like to thank the following for useful discussions and contributions: J. D. Hayes, P. A. Habitz, J. Narasimhan, D. K. Beece, M. R. Guthaus, D. S. Kung, D. J. Hathaway, V. B. Rao, A. J. Suess and J. P. Soreff.

11. REFERENCES

- [1] C. Visweswarah, “Death, taxes and failing chips,” *Proc. 2003 Design Automation Conference*, pp. 343–347, June 2003, Anaheim, CA.
- [2] R. P. Abato, A. D. Drumm, D. J. Hathaway, and L. P. P. van Ginneken, “Incremental timing analysis,” *U. S. Patent 5,508,937*, April 1993.
- [3] J.-J. Liou, K.-T. Cheng, S. Kundu, and A. Krstic, “Fast statistical timing analysis by probabilistic event propagation,” *Proc. 2001 Design Automation Conference*, pp. 661–666, June 2001, Las Vegas, NV.
- [4] M. R. C. M. Berkelaar, “Statistical delay calculation: a linear time method,” *Proc. TAU (ACM/IEEE workshop on timing issues in the specification and synthesis of digital systems)*, December 1997.
- [5] A. B. Agarwal, D. Blaauw, V. Zolotov, and S. Vrudhula, “Computation and refinement of statistical bounds on circuit delay,” *Proc. 2003 Design Automation Conference*, June 2003, Anaheim, CA.
- [6] L. Scheffer, “Explicit computation of performance as a function of process variation,” *Proc. TAU (ACM/IEEE workshop on timing issues in the specification and synthesis of digital systems)*, pp. 1–8, December 2002, Monterey, CA.
- [7] H. Chang and S. S. Sapatnekar, “Statistical timing analysis considering spatial correlations using a single PERT-like traversal,” *IEEE International Conference on Computer-Aided Design*, pp. 621–625, November 2003, San Jose, CA.
- [8] A. Agarwal, D. Blaauw, and V. Zolotov, “Statistical timing analysis for intra-die process variations with spatial correlations,” *IEEE International Conference on Computer-Aided Design*, pp. 900–907, November 2003, San Jose, CA.
- [9] J. Jess, “Dfm in synthesis,” research report, IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598, December 2001.
- [10] J. A. G. Jess, K. Kalafala, S. R. Naidu, R. H. J. M. Otten, and C. Visweswarah, “Statistical timing for parametric yield prediction of digital integrated circuits,” *Proc. 2003 Design Automation Conference*, pp. 932–937, June 2003, Anaheim, CA.
- [11] C. E. Clark, “The greatest of a finite set of random variables,” *Operations Research*, pp. 145–162, March–April 1961.
- [12] M. Cain, “The moment-generating function of the minimum of bivariate normal random variables,” *The American Statistician*, vol. 48, pp. 124–125, May 1994.
- [13] C. Visweswarah, “System and method for statistical timing analysis of digital circuits,” *Docket YOR9-2003-401*, August 2003. Filed with the U. S. Patent office.
- [14] C. Visweswarah, “System and method for probabilistic criticality prediction of digital circuits,” *Docket YOR9-2003-402*, August 2003. Filed with the U. S. Patent office.
- [15] D. J. Hathaway, J. P. Alvarez, and K. P. Belkhale, “Network timing analysis method which eliminates timing variations between signals traversing a common circuit path,” *U. S. Patent 5,636,372*, June 1997.
- [16] C. Visweswarah, “System and method for incremental statistical timing analysis of digital circuits,” *Docket YOR9-2003-403*, August 2003. Filed with the U. S. Patent office.