

# Improved *A Priori* Interconnect Predictions and Technology Extrapolation in the GTX System

Yu Cao, *Student Member, IEEE*, Chenming Hu, *Fellow, IEEE*, Xuejue Huang, *Student Member, IEEE*, Andrew B. Kahng, Igor L. Markov, Michael Oliver, Dirk Stroobandt, *Member, IEEE*, and Dennis Sylvester, *Member, IEEE*

**Abstract**—*A priori* interconnect prediction and technology extrapolation are closely intertwined. Interconnect predictions are at the core of technology extrapolation models of achievable system power, area density, and speed. Technology extrapolation, in turn, informs *a priori* interconnect prediction via models of interconnect technology and interconnect optimizations. In this paper, we address the linkage between *a priori* interconnect prediction and technology extrapolation in two ways. First, we describe how rapid changes in technology, as well as rapid evolution of prediction methods, require a dynamic and flexible framework for technology extrapolation. We then develop a new tool, the GSRC technology extrapolation system (GTX), which allows capture of such knowledge and rapid development of new studies. Second, we identify several “nontraditional” facets of interconnect prediction and quantify their impact on key technology extrapolations. In particular, we explore the effects of interconnect design optimizations such as shield insertion, repeater sizing and repeater staggering, as well as modeling choices for RLC interconnects.

**Index Terms**—*A priori* interconnect prediction, crosstalk noise, inductance, interconnect delay, system performance models, technology extrapolation, VLSI.

## I. INTRODUCTION

TECHNOLOGY extrapolation—the calibration and prediction of achievable design in future technology generations—drives the evolution of VLSI system architectures, design methodologies, and design tools. To fully explore future possibilities, technology extrapolation systems must contain models for *a priori* performance prediction, i.e., without exact knowledge about the final system. Therefore, *a priori* interconnect predictions (of wire length distribution, interconnect embedding and interconnect performance) are at the core of technology extrapolation models of achievable system power, area density and speed. Technology extrapolation, in turn, informs *a priori* interconnect prediction via models of

interconnect technology and interconnect optimizations. In this paper, we address the linkage between *a priori* interconnect prediction and technology extrapolation.

### A. *A Priori* Interconnect Prediction

Currently, *a priori* interconnect prediction techniques are focused on estimating average wire lengths, or wire length distributions, of a design in advance of the actual placement and routing. These estimates can be useful in floorplanning and design exploration, for example. *A priori* predictions must be very fast and reasonably accurate to drive floorplan changes, circuit optimizations, and other aspects of the design process.

A primary tenet of interconnect estimation is *Rent’s Rule* [1], [2] which predicts a power law relationship between the number of gates in a module of a partitioned circuit and the number of terminals needed for communication between the module and the remaining of the circuit. Wire length estimates based on Rent’s rule and a hierarchical placement model were introduced by Donath [3] in 1979, and several improvements were published by Stroobandt [4], [5] almost 20 years later. Around the same time, Davis [6] independently found a similar result for a flat placement model and this model has been frequently used by other authors. Since the introduction of the system-level interconnect prediction workshop (SLIP) in 1999 [7], the progress in the field of interconnect prediction has been tremendous [2], [8], [9]. Still, the field of *a priori* interconnect prediction has not matured enough to make its results easily applicable. It is difficult to keep up with the rapid evolutions in this field. Also, current techniques do not provide the accuracy necessary to apply them with confidence. In order to improve accuracy, more detailed knowledge about the entire system design is needed. The latter problem can be solved by coupling *a priori* prediction techniques with technology extrapolation systems. If such a technology extrapolation system can be made flexible enough to be easily adaptable to new *a priori* prediction results, much progress could be made in both *a priori* interconnect prediction and technology extrapolation at the same time. This is the goal of this paper.

### B. Technology Extrapolation

Leading edge VLSI system design aggressively exploits new process technologies, circuit techniques, design methodologies and design tools. It is thus difficult to predict the envelope of *achievable design*—e.g., with respect to performance, power, area, manufacturing cost, etc.—for a given behavior or function, in a given (future) process technology. On the other hand, such

Manuscript received August 1, 2001; revised January 16, 2002.

Y. Cao, C. Hu, and X. Huang are with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720 USA (e-mail: ycao@eecs.berkeley.edu; hu@eecs.berkeley.edu; xuejue@eecs.berkeley.edu).

A. B. Kahng is with the Computer Science Department, University of California, San Diego, CA 92093 USA (e-mail: abk@ucsd.edu).

I. L. Markov and D. Sylvester are with the Department of EECS, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: imarkov@eecs.umich.edu; dennis@eecs.umich.edu).

M. Oliver is with the Computer Science Department, University of California, Los Angeles, CA 90095 USA (e-mail: oliver@cs.ucla.edu).

D. Stroobandt is with the RUG-ELIS Department, Ghent University, Ghent B-9000, Belgium (e-mail: dstr@elis.rug.ac.be).

Digital Object Identifier 10.1109/TVLSI.2002.808479

*technology extrapolation* activity directly influences the evolution of future VLSI system architectures, design methodologies and design tools, as well as broader investment strategy in the semiconductor and electronics sectors.

Highly influential technology extrapolation systems, developed 5–10 years ago, are due to Bakoglu and Meindl (SUSPENS) [10], Sai-Halasz [11] and Hewlett-Packard Laboratories (AIM) [12]. More recent second-generation systems include GENESYS [13], RIPE [14]–[16] and BACPAC [17], along with Roadmap-related efforts [18], [19]. Typically, each system provides a plausible performance prediction model and estimates of die size and power dissipation, based on a small set of descriptors spanning device/interconnect technology through system architecture. The most critical aspect of performance prediction is the idealized *critical path* in the system of interest. For example, a model on-chip critical path might be described as “12 fanout-4 gates driving average-length local interconnects, plus an optimally buffered corner-to-corner 2  $\mu\text{m}$ -wide global wire.” A key component of the model is the *a priori* interconnection length model, as can be seen from the example earlier. The validity and accuracy of the *a priori* interconnect prediction models used are critical to the conclusions drawn from the extrapolation system. Therefore, it is important in technology interpolation systems to see and to understand the underlying models and to be able to easily adjust them.

In Section II-A, we observe that: i) current technology extrapolation systems are often incomparable; ii) they are “hard-coded” (hence, it is difficult to assess their quality and to explore changes through modeling choices); and iii) their development has entailed a near-total duplication of effort. These observations limit their application on system level technology extrapolation, including interconnect prediction and motivate efforts toward an entirely new level of technology extrapolation capability. Our GTX system has been developed with the goals of *flexibility*, *quality* and *prevention of redundant effort* in mind. The GTX system addresses these goals by providing an open, portable *framework* for specification and comparison of alternative modeling choices. A fundamental design decision in GTX is to separate model specifications from the derivation engine. GTX adopts a paradigm wherein *parameters* and *rules* allow users to flexibly capture an essentially unbounded space of attributes and relationships that are germane to VLSI technology and design. GTX, as well as documentation for numerous studies, is downloadable at <http://vlsicad.ucsd.edu/GTX/>.

### C. Contributions of This Paper

The main contributions of this paper are as follows:

- adding awareness of circuit optimization degrees of freedom into the predictions of interconnect performance used in technology extrapolation;
- adding more detailed interconnect modeling into interconnect predictions;
- identifying sensitivities in extrapolation models, showing where the most stringent accuracy requirements are with respect to how interconnect properties and optimizations are modeled;
- describing GTX—a new, portable, and general framework for reusable technology extrapolation effort.

With respect to interconnect prediction, we conduct a number of studies, within the GTX framework, on the *optimized global interconnect* portion of on-chip critical paths. Specifically, we assess the impact on critical path models of several potentially important, yet previously unmodeled, *optimization degrees of freedom* and *design constraints*, including

- adding extracted inductance estimates (and analytic *RLC* line delay estimates) to the interconnect model;
- modern repeater optimizations, such as detailed repeater size and interconnect width optimizations [20];
- *engineering considerations*, e.g., repeater area/size bounds, deliberate backing off of optimal values to the “knee of the curve,” and limiting the number of allowed wire widths;
- *switch factor* based bounds on delay uncertainty due to crosstalk from neighboring wires;
- using *real-world design technology* in the global interconnect models, e.g., repeater staggering and shielding techniques.

Our paper attempts to dispel some of the “vagueness” of current performance predictions that arises from the gaps noted above. We do not make any value judgments with respect to existing models; rather, we simply build a comprehensive modeling environment with GTX that allows us to identify the issues that *must* be considered by current and future performance predictions.

The remainder of our paper is organized as follows. Section II reviews relevant previous work in VLSI technology extrapolation and reveals large sensitivities to particular input parameters and to modeling choices. The observations lead to general goals for technology extrapolation systems. The remainder of Section II describes the architecture and implementation of our GTX technology extrapolation system. In Section III, we link *a priori* interconnect models to technology extrapolation and assess the impact of several interconnect design optimizations as well as modeling choices for *RLC* interconnects. We study: i) inductance on critical paths in terms of shielding, driver sizing and slew rates (and their impact on coupling noise); ii) the cost-performance tradeoffs inherent in signal shielding; and iii) a comprehensive study on wire sizing and repeater optimization. This analysis attempts to give a realistic depiction of what an optimal repeater topology should look like in terms of repeater sizing, wire widths, pitch allocation, etc.

## II. GTX: A NEW TECHNOLOGY EXTRAPOLATION FRAMEWORK

### A. VLSI Technology Extrapolation

A number of previous systems attempt to forecast and estimate the performance of microprocessors. Four systems—SUSPENS, GENESYS, RIPE, and BACPAC—are especially noteworthy.

**SUSPENS** [10] is the forerunner for most technology extrapolation systems. It predicts the clock frequency, chip area and power dissipation. SUSPENS ignores on-chip cache and memory structure, as well as details of multilayer interconnect structure and clock distribution. SUSPENS is also oblivious to such DSM effects as scaling and noise.

**GENESYS** [13] offers both a GUI for MS Windows (95, 98, NT) and a command-line interface (it has no Web interface). The GENESYS output file is divided into four main sections: device/material, circuit, interconnect and system. System-level outputs include throughput, maximum clock frequency, CPI and delay times for random logic and interconnect.

**RIPE** [14] explores the effect of interconnect design and technology tradeoffs on IC performance. Default input data are extracted from the NTRS roadmap. Memory and the multilayer interconnect structure are taken into account. Recent additions include estimations of linewidth variability, yield, signal integrity and electromigration [15], as well as *RLC* interconnect models [16]. The user cannot add new parameters and rules.

**BACPAC** [17] is based on a system-level performance model that consists of smaller-scale analytical models. The innovations of BACPAC include attention to power dissipation, on-chip memory, process variation and other effects. BACPAC is applicable to both ASICs and microprocessors. It attempts to enhance the accessibility of technology extrapolation via a Web-based interface. However, the derivation flow is mostly fixed and users cannot add new parameters and rules.

With respect to the previous systems for technology extrapolation, we make the following observations.

- 1) Different systems may predict the same “parameter” (e.g., microprocessor clock frequency), yet be incomparable due to differing sets of inputs and assumptions, as well as lack of documentation and visibility into internal calculations.
- 2) Each system typically offers exactly one “inference chain” for any given output of interest (e.g., cycle time). Furthermore, this inference chain can involve a large spectrum of modeling choices. The *quality* of such modeling choices cannot be assessed since the system is “hard-coded” and no exploration of modeling sensitivity or robustness is possible.
- 3) The hard-coded nature of previous systems also means that they are inflexible: the user cannot define studies of other system parameters and interaction with the system is limited.
- 4) Finally, development of previous systems has entailed near-total duplication of effort since each system attempts to bind the same envelope of achievable design in gathering, interpreting and systematizing data and models. Redundant efforts are made even though no single entity—EDA vendor, system house, or academic group—can achieve “best-possible modeling” of all aspects of technology and design.

The second point above leads to a simple and motivating experiment that reveals how existing technology extrapolations have large *sensitivities* to particular input parameters and to modeling choices. We incorporate the *cycle-time models* from SUSPENS [10] (with extensions of Takahashi *et al.* [21]), BACPAC [17] and Fisher *et al.* [19] within GTX and reproduce published results with each model. As will be explained in Section II-C, our implementations are tuned to ensure maximal interchangeability of GTX rules for each model, allowing extensive evaluation of various model sensitivities.

TABLE I  
LOGIC STAGE DELAY  $t_l$ , GLOBAL DELAY  $t_g$  AND OVERALL CLOCK FREQUENCY ( $f_c$ ) FOR INTERCONNECT MODELS

model	$t_l$ (ps)	$t_g$ (ps)	$f_c$ (MHz)
BACPAC	893	115	745
Fisher	1162	204	659
SUSPENS	665	—	1505

Our experiments address two basic types of sensitivity: *parameter sensitivity* and *model (or rule) sensitivity*. The former describes the influence of changes in the primary input parameters to the model, while the latter describes the influence of changes in the estimation model itself. (We do not aim to make value judgments about or compare the models; rather, our goal is to show the value of being able to try variant estimation methods.) We perform the following experiments:

- 1) For the same primary inputs, compare the results for different models (model sensitivity).
- 2) For each model, vary the input parameters by  $\pm 10\%$  and note the difference in the resulting clock frequency (parameter sensitivity).
- 3) For each rule out of one rule chain (model), replace one rule by a rule from another model that computes the same parameter and record the change in clock frequency (model sensitivity).

For all experiments and models, we use a common primary input (PI) parameter base derived for 0.25  $\mu\text{m}$  technology and mainly following the default parameter values of BACPAC (additional PI’s for other models are tuned to these parameter values). Despite the common parameter base, our initial model sensitivity assessment of the SUSPENS, BACPAC, and Fisher models shows very different values for respective predictions of logic stage delay ( $t_l$ ), global delay ( $t_g$ ) and overall clock frequency ( $f_c$ ) (see Table I).<sup>1</sup> A more detailed type of model sensitivity analysis tests the sensitivity of a given model to “hybridization” with other models. In other words, we take the rule chain for a single model and replace exactly one rule by an equivalent rule (or set of rules) from another model. Details of such experiments are in [22].

Parameter sensitivity studies that vary *single* PI parameter values in each model’s evaluation, changing each PI value by  $\pm 10\%$ , are also detailed in [22]. More extensive studies *simultaneously* change more than one parameter value, again by  $\pm 10\%$ . Since this produces three values for each parameter and since there are between 15 (SUSPENS) and 46 (BACPAC) primary inputs, it is not possible to consider all possibilities and we therefore only vary smaller parameter subsets (up to seven parameters at the same time). Fig. 1 plots the relative occurrence of clock frequency values in small intervals that result from the sweeping. If we say that a more “robust” (to changes of its input

<sup>1</sup>SUSPENS does not have a model for global delay on chip. We believe this was compensated by taking into account more stages but we chose to use a number of stages equal to that of the other models to maintain interchangeability. While the very high 1.5 GHz frequency predicted by SUSPENS is largely due to the lack of a global interconnect model, the logic stage delay is still significantly different from the other models.

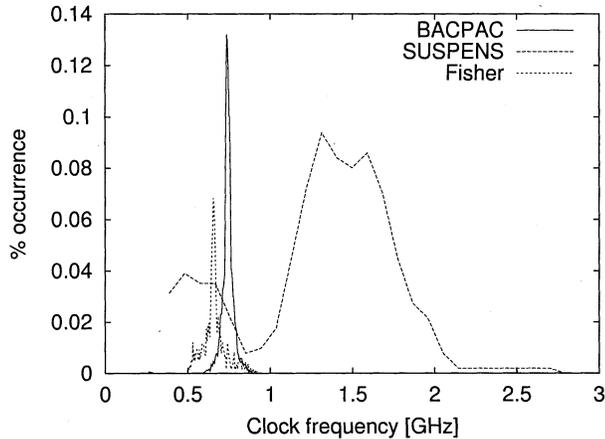


Fig. 1. Parameter sensitivity: BACPAC, Fisher, and SUSPENS.

parameters) model is one with a narrower and higher peak, then BACPAC would seem to be the most robust and SUSPENS the least robust.<sup>2</sup>

### B. Goals of a Technology Extrapolation Framework

The above observations motivate three key goals as we seek a new level of technology extrapolation capability.

1) *Flexibility*: To experimentally determine model sensitivity and robustness, users must have the ability to: i) (interactively) edit available inference chains and collections of “rules;” ii) define new parameters and rules; and iii) request specific types of studies, such as parameter optimization or trade studies. Support for interaction (GUI, session management, etc.) is an implicit requirement.

2) *Quality*: We seek adoptability in the sense of having an easy learning curve and providing much “value” in the form of high-quality embedded data, embedded models and user interface. We aim for a system that can be continuously improved to have “best-possible models” across the entire scope of technology extrapolation. Since no single group can achieve this alone, we require an open-source mechanism that is conducive to distributed ownership and maintenance.

3) *Prevention of Redundant Effort*: To avoid redundant effort, we seek a “permanent repository of first choice” for rules and data (calibration points) related to technology extrapolation. Beyond the open distribution mechanism aforementioned, adoptability (by academics open to collaboration, or by companies with proprietary data and firewalls) and maintainability become key concerns. A lower bound for adoptability is a platform-independent implementation that subsumes the functionality of all previous “hard-coded” systems. This recognizes the proprietary nature of user data and offers usability behind firewalls, with frequent releases to update the state of model/data collection.

<sup>2</sup>In general, we find BACPAC to be much less sensitive to either hybridization with other models, or variation of input parameter values. This does not necessarily imply that BACPAC is a better model (e.g., if a model predicts a clock frequency of 700 MHz independent of any input parameter value, then this is “robust” but not practical or correct). Note also that sweeping over more than seven parameters at once will widen the peaks shown in the plot.

### C. The GTX System

GTX establishes a clear separation between knowledge and implementation (Fig. 2). Knowledge is represented independently from its implementation in a serializable public-domain format. It contains data (parameters), the models (rules) that can operate on them and studies (rule chains), a collection of rules to obtain a particular result. The implementation then consists only of a derivation engine and a graphical user interface (GUI). The engine can load modules of parameters, rules and a rule chain and automatically operate on them. The result of the operation is new data. Known studies are supplied in prepackaged rule chains; additional modules can be written and shared by users.

1) *Parameters, Rules, and Rule Chains*: As aforementioned, the values of interest are encapsulated in parameters and potential inferences between them in rules. Each rule accepts as inputs a fixed collection of parameters and its evaluation computes a single output parameter. The collection of available rules and parameters is naturally viewed as a bipartite digraph in which an edge extends from a rule to a parameter if the parameter is the output of the rule, or from a parameter to a rule if the parameter is an input to the rule.

Two or more rules may compute the same output (i.e., alternative models of the same value) and the above digraph may contain cycles. However, any particular calculations must avoid such irregularities to prevent value conflicts and infinite loops. This is supported through the notion of a rule chain—an acyclic subgraph of the graph of available rules and parameters such that no two rules compute the same output.

2) *Parameters*: Parameters are the common base on which rules of different types operate. The main attributes of a parameter are its name, data type and its units. In order to obtain the goal of high reuse-ability of rules and parameters, the parameter names have to be carefully chosen so that they are easy to understand. Also, we must ensure that no physical attribute receives two different names in GTX and that no GTX parameter name is used for two different physical attributes. Therefore, we have devised strict rules for the parameter names [22]. The grammar for parameters is specified at our website [23]. Following is a very simple example representing the chip edge length.

```
#parameter dl_chip
#type double
#units {m}
#default
  1e-2
#description
  chip edge length
#endparameter
```

3) *Namespaces*: To provide a better indication of the source of rules and parameters, we have implemented *namespaces* in which a rule or parameter may be found. For example, rules derived from BACPAC are in the BACPAC namespace. “Meaningful” parameters should be in the *global namespace*; parameters that are used only for the purpose of specific rules should be in the namespace of those rules. A rule can “see” parameters in its own namespace or in the global namespace. The :: operator is used to indicate the namespace, similar to C++.

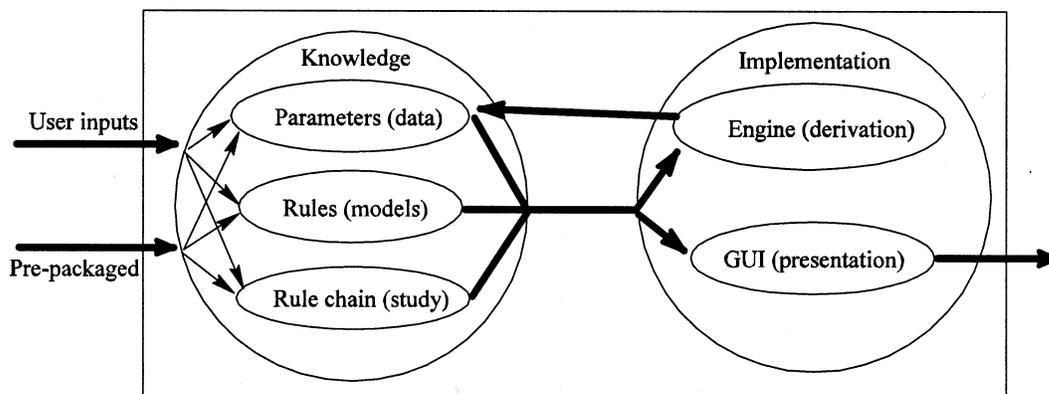


Fig. 2. Schematic view of GTX framework.

#### 4) Rules: GTX supports the following types of rules:

**ASCII Rules** provide a closed-form expression language that allows calculating the output from the input using common mathematical functions or operations, interpolation or table lookup and if-then-else. There is no program flow and, therefore, no iteration per se, but vector operations are provided that allow common computations such as sums. In a new addition to the “ASCII rules” language, it is now possible for rules to call other rules as functions. This means that one or more of the input parameters of the called rule may be substituted by an expression in the calling rule. The specific motivation for this extension is the ability to take, e.g., a rule that computes a parameter based on calendar year and use it in another rule, but shifted by some number of years.

**External Executable Rules** cause the engine to invoke a specified executable file (e.g., a PERL script), passing the input values on the command line or through a file. The external executable saves its output into a temporary file to be read by the engine. External executable rules allow the inclusion of executables for which source code is not available or for computations that cannot be expressed in ASCII rules.

**Code Rules** are hard-coded into the engine itself and require recompilation of the engine code. Therefore, they are appropriate only when execution speed is an issue.

These types provide a reasonable expressive power and facilitate easy updates to GTX with new models. The following is an example of an ASCII rule computing the chip edge length from the chip area. The #output and #inputs sections declare the types and units of output and input parameters. The formula in the #body section specifies the evaluation of the rule.

```

#namespace BACPAC
#rule dl_chip
#description
rule from BACPAC for the chip edge length
#output
double {m} dl_chip; // chip edge length
#inputs
double {m2} dA_chip; // chip area
#body
sqrt(dA_chip)
#reference
BACPAC
#endrule
  
```

5) **Rule Chains:** The GTX user indicates to the engine which of the currently available rules should be evaluated, by providing a simple list of those rules. The order in which rules are executed forms the rule chain and is decided by the engine based on the relations between the rule inputs and outputs. If we had a rule “BACPAC::dA\_chip” that computes the chip area, e.g., as a function of number and size of the gates, then the chip edge length could be computed by executing the following rule chain

```

BACPAC::dA_chip
BACPAC::dl_chip
  
```

6) **Engine Structure and Operation:** For each parameter, the engine maintains zero, one or more values. Values can be set by default, loaded from files, entered by the user or computed. Multiple values can be computed by sweeping, i.e., evaluating rules over multiple combinations of input parameters. When instructed to evaluate a rule chain, the engine clears values that can be computed by rules of the chain. For each combination of values of primary inputs of the chain, the engine evaluates rules in topological order and adds their output values to respective collections of values, unless some constraints fail. A faster algorithm is possible to produce all derivable sets of values, but with our simple algorithm the inputs of any particular value can be recovered (e.g., for minimization along a rule chain).

7) **Graphical User Interface:** The GUI is implemented with the cross-platform toolkit wxWindows; we have run it successfully on Windows 95/98, Windows NT, Solaris and Linux. At any given time, the user may view: i) current parameters; ii) current rules; iii) current rule chain; iv) values of parameters in the current chain; or v) the graph of rules and parameters. When a particular parameter or rule is selected, its details are shown and can be edited. The chain view shows all rules in the chain and helps the user to add new rules to the chain. The values view shows both inputs to and outputs of the current chain. The inputs may be edited. This view permits invoking the chain and observing the output, sweeping over multiple input values, observing the trace of such a sweep (including optimization) and plotting (Fig. 3). The graph view facilitates understanding what parameters are inputs (or outputs) to what rules, or optionally what parameters are used to compute what other parameters. In addition to the five views, the GUI handles extensive file I/O and interactive addition of new parameters and rules.

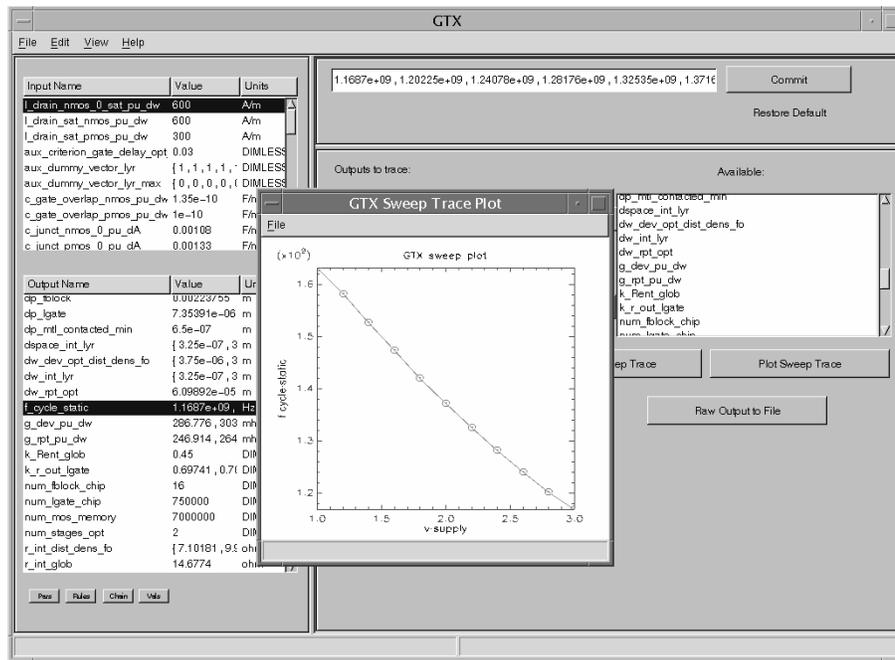


Fig. 3. Screen shot of GTX GUI.

### III. NEW DIMENSIONS OF A *PRIORI* INTERCONNECT PREDICTION

The very nature of technology extrapolation requires the use of *a priori* prediction techniques based on simple models with a limited number of parameters, preferably represented by simple ASCII rules. One of the key aspects of interconnect-centric design is the interconnection length. The early technology extrapolation systems used the simple *a priori* interconnection length estimation technique of Donath [3] to estimate average wire lengths. More recent systems included Davis' interconnection length model [6] and progress in *a priori* interconnect prediction has accelerated with the introduction of the System-Level Interconnect Prediction Workshop (SLIP) in 1999 [7]. For a technology extrapolation system to remain up-to-date, it is necessary to adapt to such evolutions easily. The flexibility of GTX allows such adaptations (simply by rewriting a few ASCII rules without having to recompile) and ensures that the effects of the latest improvements in wire length estimation can be investigated. In particular, we have found that GTX is useful as a development tool for adding new rules that model a very particular part of the design behavior, as an emulation tool for existing estimator tools, as a comparison tool between different estimation methods and as an evaluation tool for those methods.

To better explain the capabilities of the GTX system, the remainder of this paper is dedicated to the combination of *a priori* interconnect prediction models and models of the process and/or the design technology. Both types of models are easily implemented in GTX rules which allows interesting studies of the impact (quantitatively and qualitatively) of potential design or process improvements. In the next two subsections, examples of such studies— 1) with respect to inductance and *RLC* interconnect analysis and 2) with respect to interconnect design optimizations—are presented. As can be seen from the results, these represent two important new dimensions in achieving ac-

curate *a priori* interconnect performance predictions. In general, these studies are aimed at achieving reusable, transparent, well-engineered prediction models for optimized interconnects and on-chip critical paths. The default technology used in our studies (exceptions will be noted) is a  $0.18\ \mu\text{m}$  CMOS process with a supply voltage of  $1.8\ \text{V}$ .  $V_{th}$  is  $0.3\ \text{V}$  and the  $I_{dsat}$  values for NMOS/PMOS are  $700/350\ \mu\text{A}/\mu\text{m}$ . The critical global interconnect we assume is a  $1.5\ \text{cm}$  top-level copper line with thickness of  $1.3\ \mu\text{m}$  and  $\epsilon_r = 4.0$ .

#### A. Inductance and *RLC* Interconnect Analysis

The effect of inductance on the wire delay is well demonstrated in [10]. Interconnects in deep-submicron designs operating at high frequencies, whose inductive impedance cannot be neglected, must be modeled using *RLC* segment models. When the ratio of inductive impedance to resistance exceeds a certain threshold in an interconnect line, a nonmonotone voltage response (i.e., oscillation before settling to a steady-state value) results. This makes threshold delay calculation much more difficult than in the *RC* line case. In such regimes, Elmore and other *RC* line models cannot accurately estimate signal delay.

Inaccuracies in delay estimation are not only harmful to technology projections, but can also damage performance-driven routing methods that try to optimize interconnect segment length, width, spacing and repeater/buffer sizing, etc. based on analytic delay formulas. Our study quantifies the impact of using analytic threshold delay formulas derived from *RLC* line models as opposed to *RC* line models.

1) *RLC Delay Modeling*: Inductance has a larger impact on inductive noise peak and indirectly affects the capacitive coupling noise peak because the slew times at all the nodes of the wire are faster when the line is modeled as *RLC*. Inductance is calculated based on expressions from [24], [25] and the partial inductance concept [26]. We focus on analytical *RLC* interconnect delay models because their continuous, closed-form

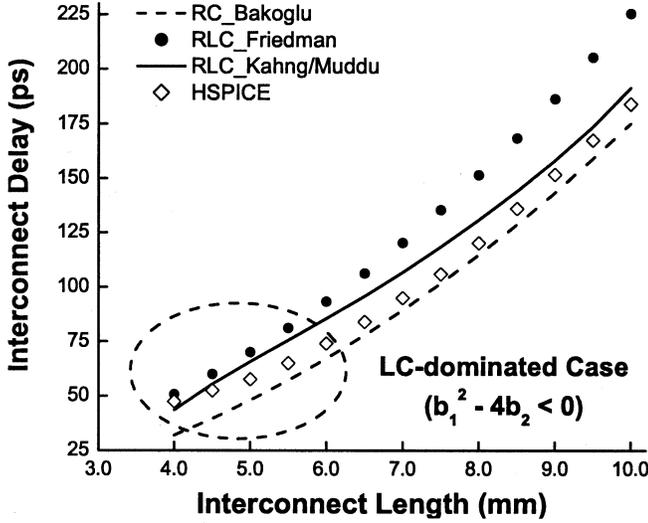


Fig. 4. Comparison of RC/RLC delay models.

nature is well suited to modern iterative-improvement interconnect design methodologies and global optimization techniques. Gate delay is computed separately using a Thevenin model with voltage source and source resistance corresponding to the driver and the load is modeled with a capacitance.

The two-pole delay model we use in this study was originally presented in [27] and is briefly described here. The transfer function for the two-pole model is given by

$$H(s) = \frac{1}{1 + b_1s + b_2s^2}. \quad (1)$$

The coefficients in this transfer function are given by

$$\begin{aligned} b_1 &= R_S C + R_S C_L + \frac{RC}{2} + RC_L \\ b_2 &= R_S \frac{C^2}{6} + \frac{R_S RC C_L}{2} + \frac{RC^2}{24} \\ &\quad + \frac{R^2 C C_L}{6} + \frac{LC}{2} + LC_L \end{aligned} \quad (2)$$

where  $R_S$  is source resistance,  $C_L$  is load capacitance at the end of the line, and  $R, C, L$  are the total electrical characteristics of the line. When the input at the source is modeled as a step input the output response is computed separately for the underdamped and overdamped cases.

We have implemented three different interconnect models and compared them with SPICE results. Fig. 4 shows the results with varying line lengths. Linewidth is fixed at  $1 \mu\text{m}$ . The driver size and receiver size are fixed at  $(W_p, W_n) = (54, 18) \mu\text{m}$ . For long wire lengths or for narrow line widths, the line tends to be more resistive ( $RC$  dominant) and Bakoglu's  $RC$  model produces results closely matching with SPICE. However, when delay is more  $LC$  dominated (i.e., large inductance value), the  $RC$  model underestimates delay by more than 10%. Friedman's model [28] matches well with SPICE for  $LC$ -dominated cases but overestimates delay by up to 30% in  $RC$ -dominated cases. Finally, the two-pole model in [27] described above matches SPICE for both  $RC$  and  $RLC$  cases within 10% error. (Given its acceptable accuracy, we use the two-pole model for subsequent studies below.) Note that with increasing line length,

TABLE II  
UNDERSHOOT VOLTAGE NORMALIZED TO  $V_{dd}$  WITH VARYING DRIVE STRENGTHS AND RETURN PATH DISTANCES; WIDTH =  $2 \mu\text{m}$

Repeater Size ( $W_n/L_n$ )	Return path distance ( $\mu\text{m}$ )			
	25	50	100	150
200	0.0004	0.008	0.021	0.029
300	0.0098	0.035	0.061	0.074
400	0.0262	0.062	0.093	0.108
500	0.0420	0.083	0.116	0.132

the 2-pole model changes from the complex pole case (overdamped or  $LC$ -dominated) to the real pole case (underdamped or  $RC$ -dominated). The condition to determine the case is from  $b_1^2 - 4b_2 > 0$  (real poles) or  $\leq 0$  (complex or double poles).

We also study the reduction of threshold delay by controlling overshoot/undershoot of the voltage response. Typically, circuit design guidelines will define the amount of overshoot and undershoot allowed in a response. These can be translated into a condition between the first and second moments of the interconnect transfer function, which are in turn functions of driver and interconnect parameters. As shown in Table II, undershoot conditions in  $0.18 \mu\text{m}$  technology can be easily avoided with proper repeater sizing and by providing reasonable signal return paths.

**Shielding Topologies:** Shielding is an important technique that designers can leverage to maximize interconnect performance at the cost of increased routing area [29]. By inserting ground and  $V_{dd}$  shield wires, current return paths can be clearly defined and loop inductance can be reduced compared to cases without explicit shielding. The extreme case of shielding is described in [30] where every signal wire has a ground and  $V_{dd}$  wire as its two nearest neighbors. In this paper, we seek to minimize the cost of a design while achieving good performance. The width of the shield wires ( $W_{\text{shield}}$ ) and signal wires ( $W_{\text{sig}}$ ), the spacing between signal wires ( $S_{\text{sig}}$ ) and the spacing from signal to neighboring shield wires ( $S_{\text{shield}}$ ) are all parameters in this paper. We examine the following three scenarios:

- No shielding (NS)—all current returns through a regular power grid. Wiring pitch is equal to  $(W_{\text{sig}} + S_{\text{sig}})$ .
- Single shielding (1S)—each signal wire has one shield wire as a nearest neighbor, while the other neighbor is another signal wire. If signal wires are denoted by S and shield (ground) wires by G, the order is G-S-S-G-S-S-G-S-S-G. Wiring pitch is  $(2S_{\text{shield}} + S_{\text{sig}} + 2W_{\text{sig}} + W_{\text{shield}})/2$ .
- Double shielding (2S)—signal and shield wires alternate. This case is identical to the dense wiring fabric in [30]. Wiring pitch is  $(W_{\text{sig}} + W_{\text{shield}} + 2S_{\text{shield}})$ .

The cost function is defined as the product of wiring pitch, repeater sizing factor and the number of repeaters inserted in the path. We attempt to minimize this cost function based on the following constraints.

- 1) Maximum delay is set at 1 ns and calculated according to each of the three delay models we have implemented.
- 2) Peak noise is fixed at 20% of  $V_{dd}$  and calculated based on the exponential model in [31].
- 3) Delay uncertainty is constrained and defined to be the difference between the  $RC$  (2-pole) and  $RLC$  delays.
- 4) The maximum allowable input slew time is 0.5 ns.

TABLE III  
COST FUNCTION COMPARISON FOR VARYING SWITCH FACTORS, DELAY MODELS, AND SHIELDING SCENARIOS

Model	Shielding	SF = 1	SF = 2	SF = 3
RC, 1 pole	NS	3.45	5.75	8.75
	1S	5.55	7.4	9.25
	2S	7.65	7.65	7.65
RC, 2 pole	NS	3.45	6.25	9.0
	1S	5.55	7.4	9.25
	2S	7.65	7.65	7.65
RLC	NS	2.85	4.6	6.75
	1S	5.1	6	7.4
	2S	7.05	7.05	7.05

Using these constraints, we can examine the impact of shielding topology on circuit performance via coupling capacitance (constraints 2,4) and inductance (constraints 1–3). Recall that switch factors account for the capacitive Miller effect—the impact of neighboring wires switching in the same (opposite) directions can be modeled by lumping their coupling capacitances to ground and multiplying by some switch factor. Switch factors are 1, 2, and 3 in this study [32].

We sweep repeater size, number of repeaters,  $W_{sig}$  and  $S_{sig}$  to find the minimal layout cost while meeting the above constraints. We also set  $W_{shield}$  to  $2W_{sig}$  and  $S_{shield}$  equal to  $S_{sig}$  to reduce the total number of variables. Results are presented in Table III, which shows the achievable cost (in arbitrary units) with varying switch factors and delay models. The 2S case can yield the minimal cost when a high switching factor is used. This is true in both RC models—these two models show very similar results from the optimization runs. The RLC model gives the overall best-cost results. Also, the slew time constraint can be more easily met if inductive effects are accounted for. The third constraint described above turns out to be a limiting factor for many input combinations—we find that RLC delay uncertainty is within bounds for smaller repeater sizes and for the 1S and 2S cases where inductance is small due to nearby current return paths.

### B. Design Optimization Studies

In this subsection, we introduce a number of techniques to optimize the use of repeaters in critical paths. Models are developed and used to account for many effects that are currently dealt with in an *ad hoc* manner.

1) *Wire Sizing*: We next turn to the impact of wire sizing on important design metrics such as delay, noise and cost. We begin with an expression for optimal wire width as a function of line length,  $l$ , from [20]

$$W_{opt}(l) = \sqrt{\frac{R_{sh}(C_f l + 2C_L)}{2R_D C_a}}. \quad (3)$$

Here,  $C_a$  and  $C_f$  denote the area and fringing capacitances per unit length,  $R_{sh}$  is the sheet resistance,  $R_D$  is the driver resistance, and  $C_L$  is the load capacitance at the end of the line.<sup>3</sup> We first examine the impact of line spacing on optimal wire width by changing spacing from 0.5 to 2  $\mu\text{m}$ —Fig. 5 plots the optimal

<sup>3</sup>Fringing capacitance is taken as the difference between the total line capacitance and the parallel-plate capacitance from [36]

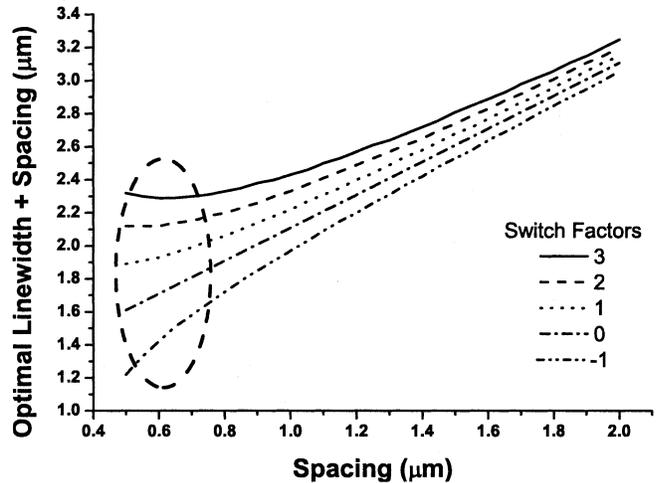


Fig. 5. Translation from optimal linewidth to optimal pitch demonstrates inflection points for certain switch factors.

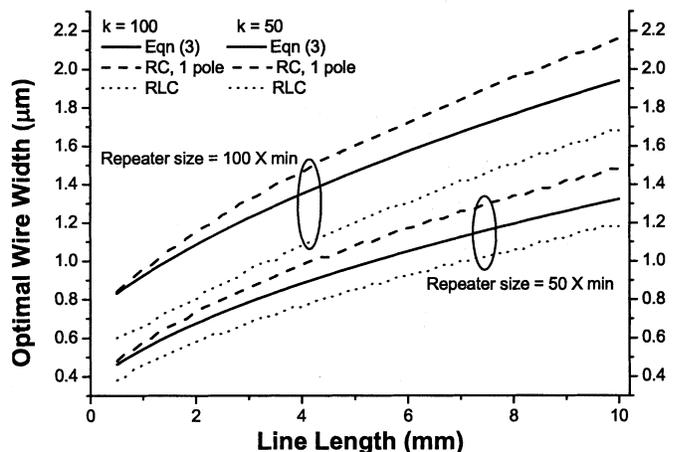


Fig. 6. Optimal wire width expression (3) exhibits 30% error with respect to RLC model, less error compared to RC.

linewidth + spacing for a 1.5 mm line, versus spacing alone on the  $x$ -axis. This plot shows an inflection point for switch factors 3 and 2, which corresponds to the optimal *pitch*, not just the optimal linewidth.

Nominal and optimistic switch factors may have such an inflection point, but they do not fall in the design space of the process technology. Fig. 5 uses (3) to calculate optimal line width.

We compare line widths obtained using (3) to the optimal line widths as found by sweeping the linewidth in GTX, for a range of driver and interconnect topologies. In addition, we incorporate inductance into the delay expressions and again perform exhaustive sweeping to find optimal line widths based on minimizing RLC as well as RC stage delay. As shown in Fig. 6, our results demonstrate that (3) matches the GTX results within 10% and often less than 5% error. However, the presence of inductance causes the optimal line width to shrink substantially and (3) therefore overestimates  $W_{opt}$  for RLC lines. Also, increasing repeater size leads to a rise in  $W_{opt}$  for all models studied—expression (3) shows slightly more error for larger drivers.

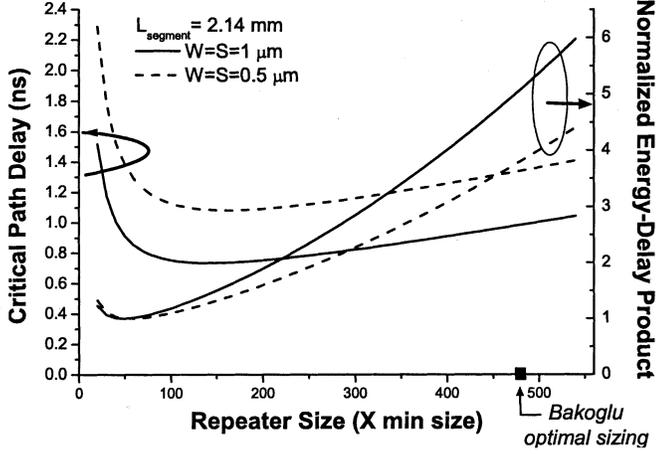


Fig. 7. This plot clearly demonstrates the severe oversizing resulting from simple expressions such as (4).

2) *Repeater Sizing*: The most commonly cited optimal buffer sizing expression is that of Bakoglu [10]

$$S = \sqrt{\frac{R_D C_{\text{int}}}{R_{\text{int}} C_{\text{in}}}}. \quad (4)$$

$R_D$  reflects the minimum-sized driver resistance,  $C_{\text{in}}$  is the input gate capacitance of a minimum-sized inverter, and  $R_{\text{int}}$  and  $C_{\text{int}}$  are, respectively, the line resistance and capacitance per unit length. Although this expression can give accurate results in some cases when optimizing for delay only, the delay versus device size relationship lends itself to further optimization due to its insensitivity near the optimal point. Results obtained from (4) are often unrealistically large—typical standard cell libraries may include inverters or buffers up to 54–96X the minimum size ( $W_n = L_{\text{drawn}}$ ) whereas (4) can give results in the range of 400–700X minimum. To compensate for this, an expression was derived in [17] to optimize a weighted delay-area product rather than purely delay—it gave results on the order of 50%–60% smaller than (4). Even with this modification, however, so-called optimal repeater sizes seem impractical in the face of power and area constraints.

Here and in the remainder of the subsection, we present a more experimental approach to finding optimal repeater size. For various wire geometries, noise conditions, area and placement constraints and delay models, we develop a complete picture of the optimal repeater topology solution. We begin with a simple sweep of the repeater size for a single stage of a chain and examine both delay and energy-delay product versus repeater size in Fig. 7.

As Fig. 7 shows, the optimal buffer sizing as calculated from (4) is 480 times the minimum-sized inverter. From pure delay analysis, GTX optimization results indicate that the ideal buffer size for our standard critical path is  $\sim 140$ – $150$  times the minimum size. When optimizing the energy-delay product, that value drops all the way to 50–60 times minimum. Any range of weighting functions can be easily incorporated into the rule chains—for instance,  $(\text{energy-delay})^2$  or  $(\text{energy-delay})^3$ . Results from such functions are not included here, but will push the optimal size toward the delay-only size of 140–150

times minimum. It is also important to note from Fig. 7 that the path delay function around the delay-optimal repeater size is very flat: a buffer which is 43% smaller than optimal yields only a 6.8% delay penalty. Since the energy-delay optimal size is found in the steep part of the delay curve, a truly ideal choice would more closely reflect the knee of the delay curve. In the case of Fig. 7, our choice of “optimal repeater size” is in the range of 80–100 times the minimum inverter size.

3) *Repeater Placement Uncertainty*: The placement of repeaters in a deep submicron design is nontrivial—many thousands of repeaters must be used to meet timing and noise objectives and this number will increase with process scaling. As a result, the area consumed by these buffers is substantial and may no longer be ignored during the floorplanning design phase. Particularly in a hierarchical design methodology, such as that proposed in [33], it may not be possible to place repeaters at any given location either inside a predesigned block or at the top-level of the hierarchy. A potential solution to this problem involves the formation of repeater block regions located around the chip at the floorplanning stage that provide specified areas for repeaters to be placed [34]. However, with such an approach the feasible distances between repeaters are discrete, not continuous.

Here, we study the impact on critical path delay of this inability to place repeaters at arbitrary locations. As before, we examine a top-level metal 1.5 cm route in the default technology. We define an uncertainty parameter  $\epsilon$ , which can range from 0 (no uncertainty) to 1 (maximum uncertainty). We express the location uncertainty as  $(1 \pm \epsilon)L_{\text{seg}}$  where  $L_{\text{seg}}$  is the nominal distance between repeaters when there are no placement restrictions. Given these bounds on segment length between consecutive buffers, we examine the worst-case scenario when half of the segments in the critical path have length  $(1 - \epsilon)L_{\text{seg}}$  and the other half are of length  $(1 + \epsilon)L_{\text{seg}}$  while total path length is fixed. Given uniform buffer sizing, half of these segments will be overdriven while the other half are underdriven.

While sweeping  $\epsilon$ , we vary the switch factor and plot the path delay and peak noise normalized to the  $\epsilon = 0$  case. Results shown in Fig. 8 indicate that the impact of repeater placement uncertainty is small for total path delay but large for peak noise. This can be understood by realizing that the path delay effectively averages out the resulting fast and slow stages while peak noise is a function of the segment length  $(1 \pm \epsilon)L_{\text{seg}}$  and not the total path length. Since the peak noise results are normalized to the  $\epsilon = 0$  case, the switching factor does not play a major role. With a conservative  $\epsilon$  of 0.3, the worst-case peak noise increases by approximately 30%.

4) *Staggered Repeaters*: The use of staggered repeaters for global buses was first described in [35]. The layout structure is shown in Fig. 9. This approach uses offset buffers in a bus-like structure to minimize the impact of coupling capacitance on delay and crosstalk noise. If repeaters are offset so that each gate is placed in the middle of its neighboring gates’ interconnect loads, the effective switching factor is limited to one. This is because potential worst-case simultaneous switching on adjacent wires can be present for only half the victim line’s length and in such conditions the other half of the victim line will consequently experience best case neighboring switching activity.

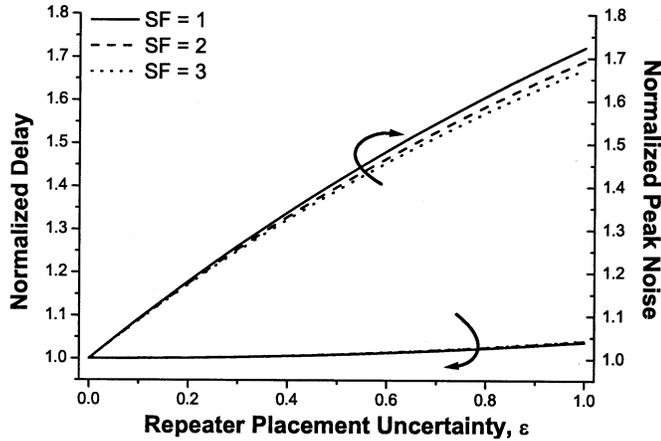


Fig. 8. Repeater placement uncertainty has a large impact on noise but little on delay.

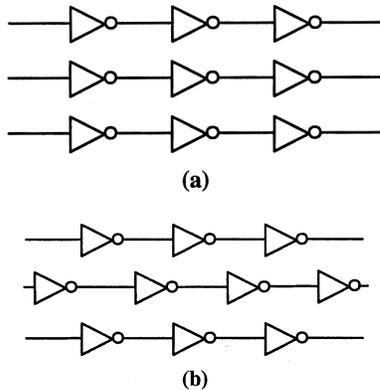


Fig. 9. Reduction of worst-case Miller coupling by staggered repeaters. In usual layouts (a), inverters on left and right neighbors are at phase = 0 with respect to the inverters on middle line. Staggering (b) places inverters on left and right neighbors at phase = 0.5.

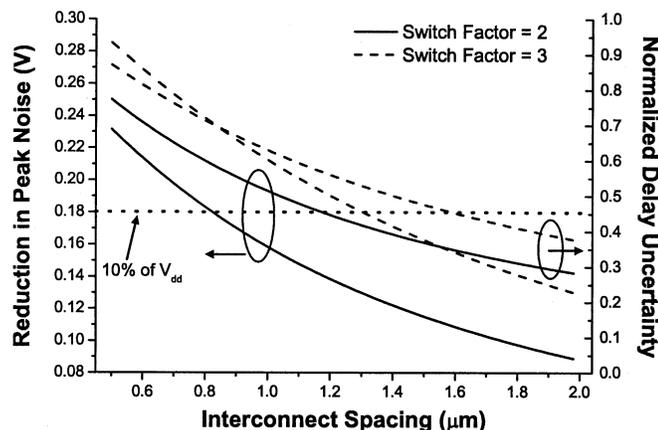


Fig. 10. Staggered repeater topology substantially reduces peak noise and delay uncertainty compared to traditional case.

In our analysis, we examine the potential reduction in delay uncertainty, as well as in peak crosstalk noise, due to staggered repeaters. Fig. 10 shows that the noise reductions can easily be greater than 10% of  $V_{dd}$  for realistic spacing and switch factors. The delay uncertainty when using nonstaggered repeaters can exceed 50% of the nominal delay—but staggering almost

completely eliminates this uncertainty that stems from capacitive coupling.

#### IV. CONCLUSION

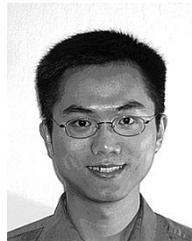
*A priori* interconnect prediction is a key enabler for technology extrapolation engines. In addition, the rapid evolution of technology and prediction methods leads to a need for dynamic and flexible prediction frameworks. This paper outlines one such framework that aims at building upon *a priori* interconnect prediction and comprehensive interconnect models to quantify the available performance envelope for future IC design as well as to investigate the impact of various design optimization strategies in these designs. We have described the architecture and implementation of GTX, the MARCO GSRC Technology Extrapolation system. GTX has the potential to change how we extrapolate the impact of new process and design technology: it can provide a “living roadmap” that incorporates—and serves as a repository for—essentially unlimited forms of domain knowledge. The structure of GTX (with the flexibility of the ASCII rules) is especially useful for including *a priori* interconnect prediction techniques and for keeping the model database up-to-date in this rapidly evolving field.

The combination of *a priori* interconnect prediction techniques with a highly flexible technology extrapolation system such as GTX (together with the addition of models for process and design technology), enables a quick and easy evaluation of modeling choices through a new set of studies of interconnect properties. Awareness of the latest interconnect optimizations and design degrees of freedom represent “new dimensions” in *a priori* interconnect prediction. In this paper, we have used GTX to examine the topics of *RLC* delay modeling, wire shielding, optimal repeater and wire sizing, repeater staggering and repeater placement uncertainty effects. We demonstrated that when including inductance, errors in estimates of optimal line delay could increase up to 30%, implying that an *RLC*-based model could be necessary. A closed-form wire sizing expression was evaluated and found to yield good results compared to a 1-pole *RC* delay model, but more substantial error compared with an *RLC* model. We also found that conventional models for optimal repeater sizing are insufficient—our examples show significant overestimation of repeater size up to 500%. A more effective sizing criterion would weight energy and delay so that the size closely approximates the knee of Fig. 4. We have also modeled the impact of repeater staggering (a layout technique which limits delay uncertainty and peak noise due to capacitive Miller effect).

In summary, we have identified the importance of including a wide range of circuit optimization techniques into predictions of interconnect performance in technology extrapolation. GTX strives to bring about this awareness by allowing for more detailed interconnect modeling and optimization degrees of freedom in its analysis. Using GTX, we demonstrated a wide range of modeling and parameter sensitivities in existing extrapolation models—this isolates the most critical modeling needs in extrapolation engines. Bringing these contributions together, we see that while technology extrapolation engines such as GTX rely heavily on *a priori* interconnect predictions, these estimation methods also require substantial augmentation to develop a truly useful design prediction framework.

## REFERENCES

- [1] B. S. Landman and R. L. Russo, "On a pin versus block relationship for partitions of logic graphs," *IEEE Trans. Comput.*, C-20, pp. 1469–1479, 1971.
- [2] P. Christie and D. Stroobandt, "The interpretation and application of Rent's rule," *IEEE Trans. VLSI Syst., Special Issue on SLIP* 8, vol. 6, pp. 639–648, Dec. 2000.
- [3] W. E. Donath, "Placement and average interconnection lengths of computer logic," *IEEE Trans. Circuits Syst.*, vol. 26, pp. 272–277, 1979.
- [4] D. Stroobandt and J. V. Campenhout, "Estimating interconnection lengths in three-dimensional computer systems," *IEICE Trans. Inform. Syst., Special Issue on Synthesis and Verification of Hardware Design*, vol. E80-D, no. 10, pp. 1024–1031, 1997.
- [5] —, "Accurate interconnection length estimations for predictions early in the design cycle," *VLSI Design, Special Issue on Physical Design in Deep Submicron*, vol. 10, no. 1, pp. 1–20, 1999.
- [6] J. A. Davis, V. K. De, and J. D. Meindl, "A stochastic wire-length distribution for gigascale integration—Part I: Derivation and validation," *IEEE Trans. Electron Devices*, vol. 45, pp. 580–589, 1998.
- [7] D. Stroobandt and A. B. Kahng, *Workshop Notes of the 1st Int. Workshop on System-Level Interconnect Prediction*, 1999, <http://www.elis.rug.ac.be/~dstr/SLIP.html>.
- [8] D. Stroobandt, "Recent advances in system-level interconnect prediction," *IEEE CAS Soc. Newslett.*, vol. 11, pp. 1;4–20;48, 2000.
- [9] —, *A Priori Wire Length Estimates for Digital Design*. Boston, MA: Kluwer Academic, 2001.
- [10] H. B. Bakoglu, *Circuits, Interconnections and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.
- [11] G. A. Sai-Halasz, "Performance trends in high-performance processors," *Proc. IEEE*, pp. 20–36, Jan. 1995.
- [12] P. Raje, "A framework for insight into the impact of interconnect on 0.35  $\mu\text{m}$  VLSI performance," *Hewlett-Packard J.*, pp. 1–8, 1995.
- [13] J. C. Eble, V. K. De, D. S. Wills, and J. D. Meindl, "A generic system simulator (GENESYS) for ASIC technology and architecture beyond 2001," in *Proc. ASIC, 1996*, pp. 193–196.
- [14] B. M. Geuskens and K. Rose, *Modeling Microprocessor Performance*. Boston, MA: Kluwer Academic, 1998.
- [15] R. Mangaser, C. Mark, and K. Rose, "Interconnect constraints on BEOL manufacturing," in *Proc. Adv. Semicond. Manufact. Conf.*, 1999, pp. 304–308.
- [16] C. Mark and K. Rose, "Interconnect strategies for deep submicron CMOS manufacture," in *Proc. Adv. Semicond. Manufact. Conf.*, 2000, pp. 413–418.
- [17] D. Sylvester and K. Keutzer, "System-Level performance modeling with BACPAC—Berkeley advanced chip performance calculator," in *Proc. SLIP*, 1999, <http://www.eecs.umich.edu/~dennis/bacpac/>, pp. 109–114.
- [18] ITRS Semiconductors, "International Technology Roadmap for Semiconductors," <http://www.itrs.net/>, 1999.
- [19] P. D. Fisher and R. Nesbitt, "The test of time: Clock-cycle estimation and test challenges for future microprocessors," *IEEE Circuits Devices Mag.*, vol. 14, pp. 37–44, 1998.
- [20] J. Cong and D. Z. Pan, "Interconnect estimation and planning for deep submicron designs," in *Proc. Design Automation Conf.*, 1999, pp. 507–510.
- [21] S. Takahashi, M. Edahiro, and Y. Hayashi, "A new lsi performance prediction model for interconnection analysis of future LSIs," in *Proc. ASP-Design Automation Conf.*, 1998, pp. 51–56.
- [22] A. E. Caldwell *et al.*, "GTX: The MARCO GSRC technology extrapolation system," in *Proc. Design Automation Conf.*, 2000, pp. 693–698.
- [23] [Online]. Available: <http://vlsicad.cs.ucla.edu/GSRC/GTX/>
- [24] L. He, N. Chang, S. Lin, and O. S. Nakagawa, "An efficient inductance modeling for on-chip interconnects," in *Proc. Custom Integrat. Circuits Conf.*, 1999, pp. 457–460.
- [25] X. Qi *et al.*, "On-Chip inductance modeling and RLC extraction of VLSI interconnects for circuit simulation," in *Proc. Custom Integrat. Circuits Conf.*, 2000.
- [26] A. E. Ruehli, "Inductance calculations in a complex integrated circuit environment," *IBM J. Res. Dev.*, pp. 470–480, Sept. 1972.
- [27] A. B. Kahng and S. Muddu, "An analytical delay model for RLC interconnects," *IEEE Trans. Computer-Aided Design*, vol. 16, no. 12, pp. 1507–1514, 1997.
- [28] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Equivalent elmore delay for RLC trees," *IEEE Trans. Computer-Aided Design*, vol. 19, no. 1, pp. 83–97, 2000.
- [29] Y. Massoud, S. Majors, T. Bustami, and J. White, "Layout techniques for minimizing on-chip interconnect self-inductance," in *Proc. Design Automat. Conf.*, 1998, pp. 566–571.
- [30] S. P. Khatri, A. Mehrotra, R. K. Brayton, A. Sangiovanni-Vincentelli, and R. H. J. M. Otten, "A novel VLSI layout fabric for deep submicron applications," in *Proc. Design Automat. Conf.*, 1999, pp. 491–496.
- [31] K. L. Shepard *et al.*, "Design methodology for the S/390 parallel enterprise server G4 microprocessors," *IBM J. Res. Dev.*, pp. 515–554, July–Sept. 1997.
- [32] A. B. Kahng, S. Muddu, and E. Sarto, "On switch factor based analysis of coupled RC interconnects," in *Proc. Design Automat. Conf.*, 2000, pp. 79–84.
- [33] D. Sylvester and K. Keutzer, "Getting to the bottom of deep submicron," in *Proc. Int. Conf. Computer-Aided Design*, 1998, pp. 203–211.
- [34] J. Cong, T. Kong, and D. Z. Pan, "Buffer block planning for interconnect-driven floorplanning," in *Proc. Int. Conf. Computer-Aided Design*, 1999, pp. 358–363.
- [35] A. B. Kahng, S. Muddu, and E. Sarto, "Tuning strategies for global interconnects in high-performance deep submicron IC's," *VLSI Design 10(1)*, pp. 21–34, 1999.
- [36] T. Sakurai, "Closed-form expressions for interconnect delay, crosstalk and coupling in VLSI's," *IEEE Trans. Electron Dev.*, pp. 118–124, Jan. 1993.



**Yu Cao** (S'98) received the B.S. degree in physics from Peking University in 1996 and the M.A. degree in biophysics from the University of California (UC), Berkeley, in 1999. He is currently a Ph.D. candidate in electrical engineering at UC Berkeley.

In summer 2000, he was with Hewlett-Packard Labs, Palo Alto, CA, as a research intern and during summer 2001, he was with IBM Microelectronics division, East Fishkill, NY, working on RF spiral inductor modeling. From 2001 to 2002, he consulted Celestry Design Technologies, Inc., San Jose, CA, on RF design tool development. His research interests include high-speed interconnect analysis, nano-CMOS design solutions and RF design automation. He holds one pending U.S. patent.

Mr. Cao received the 2000 Beatrice Winner Award at ISSCC.



**Chenming Hu** (S'71–M'76–SM'83–F'90) received the B.S. degree from the National Taiwan University and the M.S. and Ph.D. degrees from the University of California (UC), Berkeley.

He is the TSMC Distinguished Professor of Electrical Engineering and Computer Sciences at UC Berkeley. His research areas include microelectronic devices and technology and device modeling for circuit simulation. He has authored or coauthored five books and more than 700 research papers. He leads the development of the MOSFET model BSIM, the industry standard model for IC simulation.

Dr. Hu is a Member of the National Academy of Engineering, a Fellow of the Institute of Physics, and an honorary professor of the Chinese Academy of Science, Beijing, and Chiao Tung University, Taiwan. He received the 1997 Jack A. Morton Award for contributions to MOSFET reliability. He has received UC Berkeley's highest honor for teaching—the Distinguished Teaching Award and the DARPA Most Significant Technological Accomplishment Award for codeveloping the FinFET transistor structure.



**Xuejue Huang** (S'97) received the B.E. degree from Huazhong University of Science and Technology, Wuhan, China, in 1994 and the M.S. degree in Electrical Engineering and Computer Sciences from the University of California (UC), Berkeley, in 1999. Currently, she is pursuing the Ph.D. degree at UC Berkeley.

From 1994 to 1996, she was with China Integrated Circuit Design Center. During summer 2000, she worked at Hewlett-Packard Labs, Palo Alto, CA, as a research intern and in summer 2001 she worked at T. J. Watson Research Center, IBM, on interconnect modeling and optimization for multi-GHz clock network design. Her current research interests include high-speed IC design and modeling.



**Andrew B. Kahng** received the A.B. degree in applied mathematics/physics from Harvard College and the M.S. and Ph.D. degrees in computer science from the University of California at San Diego (UCSD).

He was with the University of California, Los Angeles, Computer Science Faculty for 11 years. Since 2001, he has been a Professor with the UCSD CSE and ECE Departments. He currently leads the "C.A.D." theme in the MARCO Gigascale Silicon Research Center (<http://vlsicad.ucsd.edu/GSRC/>) and is Chair of the Design International Technology

Working Group for the International Technology Roadmap for Semiconductors (ITRS). His research is mainly in VLSI CAD, centering on physical layout and performance analysis. Other research interests include combinatorial and graph algorithms and large-scale heuristic global optimization.

Dr. Kahng has received the NSF Young Investigator and three Best Paper Awards. He was the founding General Chair of the ACM/IEEE International Symposium on Physical Design and since 1997 has defined the physical design roadmap for the ITRS.



**Igor L. Markov** received the Master's and Doctorate degrees from the University of California, Los Angeles, where he received the Best Ph.D. Student Award for 2000 from the Computer Science Department.

Currently, he is an Assistant Professor of electrical engineering and computer science at the University of Michigan. He is working in the area of combinatorial optimization with applications in the design of integrated circuits and quantum computing. His results on mathematical and combinatorial optimization

have led to the development of state-of-the-art CAD tools now used at Intel, IBM, Cadence, Synplicity, and Monterey. He has coauthored more than 30 publications.

Professor Markov is a member of Technical Program Committees of three ACM/IEEE symposia and two ACM/IEEE workshops.



**Michael Oliver** received the B.S. degree from the California Institute of Technology, Pasadena, and the M.A. degree, in mathematics from the University of California, Los Angeles (UCLA) in 1984 and 1987, respectively.

He is currently developing the GTX software and pursuing the Ph.D. degree in mathematics at UCLA, in the theory of Borel equivalence relations.



**Dirk Stroobandt** (M'98) was born in Gent, Belgium, in 1972. He graduated in 1994 and received the Ph.D. degree in electrotechnical engineering in 1998 from Ghent University.

From 1994 to 1998, he was a Research Assistant and is currently a Postdoctoral Fellow with the Fund for Scientific Research—Flanders (Belgium) (F.W.O.). He is affiliated with the Department of Electronics and Information Systems (ELIS) of Ghent University. He visited the group of A. B. Kahng at the Computer Science Department of UCLA as a Postdoctoral Researcher from July 1999 to June 2000. His research is oriented toward *a priori* estimations of interconnection lengths in electronic systems and its applications to CAD, computer architecture evaluation (especially optoelectronic three-dimensional architectures), and design optimization.

Dr. Stroobandt is the inaugural winner of the ACM Outstanding Doctoral Thesis Award in Design Automation and coorganized (in 1999) and chaired (in 2000) the International Workshop on System-Level Interconnect Prediction (SLIP). He is Guest Editor of two Special Issues of the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS on System-Level Interconnect Prediction.



**Dennis Sylvester** (S'95–M'00) received the B.S. degree in electrical engineering *summa cum laude* from the University of Michigan, Ann Arbor, in 1995. He received the M.S. and Ph.D. degrees in electrical engineering from the University of California (UC), Berkeley, in 1997 and 1999, respectively.

Currently, he is an Assistant Professor with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor. Prior to joining the University of Michigan, he was a Senior R&D Engineer with the Advanced Technology Group, Synopsys, Mountain View, CA. He was also with Hewlett-Packard Laboratories in Palo Alto, CA, from 1996 to 1998. He has published numerous papers in his field of research, which includes interconnect characterization and modeling, low-power circuit design techniques, noise-aware timing analysis, and statistical process variation.

Dr. Sylvester received a 2002 NSF Career Award and the 2000 David J. Sakrison Memorial Prize for the Best Dissertation in the UC Berkeley EECS Department. He also received the 2000 Beatrice Winner Award at ISSCC, two Outstanding Research Awards from the SRC, and a Best Student Paper Award at the 1997 International Semiconductor Device Research Symposium. He is on the Technical Program Committee of numerous design automation conferences, and is the General Chair for the 2003 System-level Interconnect Prediction (SLIP) Workshop. In addition, he is part of the International Technology Roadmap for Semiconductors (ITRS), U.S. Design Technology Working Group and made significant modeling contributions to the Design and System Drivers chapters of the 2001 ITRS. He is a Member of ACM, American Society of Engineering Education, and Eta Kappa Nu.