

Estimating the efficiency of collaborative problem-solving, with applications to chip design

M. Y. L. Wisniewski
E. Yashchin
R. L. Franch
D. P. Conrady
G. Fiorenza
I. C. Noyan

We present a statistical framework to address questions that arise in general problems involving collaboration of several contributors. One instance of this problem occurs in the complex process of designing ultralarge-scale-integration (ULSI) semiconductor chips. In these processes, computer-aided design tools are treated as "black boxes." In most cases, the automated design tools operate on designs and successfully complete a specified task to create designs that satisfy specified design criteria. In other cases involving complex designs, however, the tools are unable to create designs that satisfy the specified criteria. In both situations, the performance of the tools can be enhanced with systematic external intervention that is implemented with some supplemental algorithm. This algorithm can either be fully automated or be implemented by hand, relying on formally describable human expertise. In this intervention, the supplemental algorithm and the automated design tool take turns to move the design from one configuration to another until either the task is complete or further improvements are not possible or necessary. In such a setting, a number of questions arise about how to measure the effectiveness of the external intervention. One question, for example, is whether the external intervention consistently assists the progress of the automated program. This situation is an instance of a general problem that we address in this paper. As an example, we apply the statistical framework to the problem of routing a functional unit of the IBM POWER4 microprocessor.

1. Introduction

Achievement of many complex tasks requires cooperation between a number of contributors, each of whom performs a given subtask, possibly several times. The ability of a contributor to perform a subtask well cannot be assessed exclusively on the results of this subtask, since each contributor operates independently and has a potential to impede the collaborative effort. For example, a single contributor can achieve results that appear independently advantageous; however, the actions may interfere with

those of other contributors and impede progress toward attaining the goal.

In this paper, we consider two types of contributors: 1) mandatory contributors, whose actions are required to achieve the goal and whose participation is thus mandated *a priori*; and 2) supplemental contributors, whose respective subtasks are performed only to improve the results of the mandatory contributors. We assume that criteria exist to compare results of these tasks, whether or not the goal is achieved. In such a formalism, a few questions

to address include the following: a) Which contributors are essential to achieve the goal? b) Is the addition of a given contributor likely to increase the design quality? c) Is the goal achievable? In this paper, we formulate criteria and present an approach to address such questions. For simplicity, we do not present a mathematical framework of this general problem. Instead, we consider in detail a specific problem related to chip design of ultralarge-scale-integrated (ULSI) circuits, such as systems-on-chips (SOCs).

In ULSI design, many steps in the design process are implemented with computer-aided design tools such as automated routers. In most cases, these tools are able to operate on the design data without external intervention and successfully complete a specified task to create a design that satisfies a set of objective criteria. For example, the task of an automated routing program is to connect an assortment of blocks with wires. An automated router successfully completes this task by creating a wire layout with correctly routed connections that satisfy the required project cycle time and that contain zero violations such as intersecting wires. In other cases, such as for complex designs, the tools are not able to complete the task successfully, and external intervention is required. This intervention can be implemented with a supplemental algorithm that is either fully automated or relies on formally describable human expertise. In these cases, the collaborative effort involves a mandatory wire router and a supplemental algorithm that operate in turns, and progress is monitored after each trial. Here the term *trial* refers to a pair of *moves*; the first move is taken by the supplemental algorithm, and the second move is taken by the router. The sequence of trials is continued until either the goal is achieved or for as long as one acts in accordance with one of several policies addressed later in the text; for example, one policy is to continue the trials until forecasts indicate that further improvements are neither probable nor significant. Once the goal is achieved (that is, the wire route configuration contains no violations and meets the cycle time requirement), the quality of the design layout can be assessed in terms of physical characteristics such as total wire length or yield. At this point, one may consider conducting additional trials based on cost/benefit analysis and time constraints.

For the external intervention to work in a cooperative manner with the tool, the supplemental algorithm must identify those portions of the problem that pose difficulty for the tool and then attack those portions in the next move. The external intervention can present additional constraints such as wire routes and can also relax other constraints with wires that implement difficult-to-route connections. The interactions of the external intervention and the tool are helpful if the tool is able to make additional progress, and the results tend to improve with each trial. Therefore, in the context of this example,

we formulate a method for external intervention and introduce a statistical framework to provide quantitative information on the effectiveness of the intervention. Specifically, we address the following questions: a) How does one select the next move? b) Is the external intervention proceeding along a useful path without adversely affecting physical properties of other wire segments that are routed with the automated routing tool? c) What is the highest-quality design that one can ultimately obtain? d) Given the cost of each step, when should one stop the trials?

2. Background

The ability to route ULSI designs is nontrivial, since these designs contain large numbers of signals. An understanding of signal routing is also important to achieve optimal performance [1–5], power dissipation [6], and manufacturability [7–10] of increasingly complex circuits. The standard objective criteria for a successfully routed chip design are that the connections for all signals must be routed with zero violations and satisfy the project cycle time requirement. For some designs, the automated routing tool is unable to route all signals without introducing violations in some signal routes. For other designs, the tool is able to create a wire layout with zero violations, yet there is a strong possibility that the layout can be improved further to enhance overall chip performance, yield, and reliability. Improvements can manifest themselves, for example, as changes in design physical characteristics, such as shorter overall wire length or lower overall number of vias, where a via is defined as an interlayer connection that enables the circuit to span two or more layers on the chip. For both types of designs, one can take advantage of external intervention that amounts to pre-routing part of the circuitry to facilitate the subsequent work of the router. The pre-routed wires are implemented with a separate algorithm that inserts custom interconnections in the design layout.

Effective custom interconnections enable one to develop a replicable approach that can be ported from one design to another. This approach enables the creation of designs that consistently satisfy specified criteria and that have a low level of interconnect complexity. The phrase *interconnect complexity* describes the following physical characteristics of interconnections: 1) signal interconnection length; 2) extra signal length compared with a benchmark estimate such as the signal Steiner length;¹ 3) number of nonredundant vias in a signal route; 4) signal density; and

¹ The Steiner route length of a signal in a circuit design is the complete specification of all of the lengths of straight horizontal and vertical route segments that are estimated to connect the points (pins) on a signal. The Steiner problem is NP-complete; this statement means that the time required to determine an exact route solution for a signal that connects an arbitrary number of pins increases exponentially as the number of pins increases. The Steiner algorithm discussed in this paper is available in the IBM Hierarchical Design Planner.

5) via density. Benefits of custom interconnections include a) the ability to implement violation-free route solutions for signals that are incorrectly wired by an automated router; b) the ability to decrease interconnect complexity by reducing signal density and via density; c) the ability to remove timing violations that result from routes implemented by the automated router [11]; and d) the ability to equalize wire delays for selected groups of signals with special geometric arrangements.

3. Method for external intervention

In this section, we describe a method for external intervention in ULSI design. Details of the method are described in [11]. In this discussion, we focus on the analysis of signal net lengths in a design layout; however, a similar analysis for vias or other physical properties can be obtained from the discussion.

1. Identify signals with excess route length

The first step is to identify signals with excess route length; these signals are potential targets for custom interconnections. In this step, the automated routing tool routes the signals in a design, and the net length L and Steiner length estimate L_s are measured for each signal. We introduce a measure of signal route complexity called the normalized excess Steiner length $NESL$,

$$NESL = \frac{L - L_s}{L_s}, \quad (1)$$

where $NESL$ is a ratio that measures how closely the actual signal length L approaches L_s . In the case where the automated router creates a long wire route for this signal, with L much greater than L_s , the large value of $NESL$ indicates increased route complexity for this signal. In this framework, signal routes are ordered in decreasing value of $NESL$.

2. Identify congested regions

The next step is to identify congested regions in a design layout. In this step, two-dimensional maps of the signal density and the via density are obtained from the routed design layout [11]. We use the term *congested region* to refer to a region that exhibits high values of both signal density and via density compared with values in surrounding regions.

3. Identify signals to target for custom interconnections

The next step is to compare the locations of congested regions with locations of signals with large values of $NESL$. In this step, those signals in congested regions are targeted for custom routes. In our experience, regions that are both congested and contain signals with large values of $NESL$ contain signals that should

be targeted for external intervention with custom interconnections [11].

4. Insert custom interconnections

The next step is to insert custom interconnections for targeted signals identified in step 3. In this step, the remaining signals are then routed with an automated routing tool, and the routes are checked for electrical shorts and geometry violations. Steps 1 to 4 are iterated until the wire layout contains no violations and satisfies the project cycle time requirement. It is also possible to iterate these procedures until $NESL$ is reduced below a design-specified limit for all signals [11].

5. Quantify complexity of the interconnections

At this step in the process, a working design layout exists; in general, this design layout is not unique. In this example, since we consider net lengths of signal routes, the wire layouts containing shorter signal routes are better layouts, with all other factors assumed equivalent. It is possible therefore to assign a figure of merit to a design layout according to the following procedure [11].

First, the total interconnect length L_T and total interconnect Steiner length L_{ST} for all signals are obtained by taking the sum of the contributions for all signals. We introduce a measure of complexity of the design interconnections called the total excess Steiner length $TESL$,

$$TESL = L_T - L_{ST}, \quad (2)$$

where $TESL$ is the length by which L_T exceeds L_{ST} .

We also introduce a second measure of interconnect complexity called the average normalized excess Steiner length $\langle NESL \rangle$,

$$\langle NESL \rangle = \frac{\sum_{i=1}^{N_{\text{group}}} NESL(i)}{N_{\text{group}}}, \quad (3)$$

where the average is taken over a group of signals N_{group} , and where $NESL(i)$ represents the $NESL$ of signal i . The group of signals can consist of the entire set of N signals in the design or a subset of signals $N_{\text{group}} \leq N$ in a congested region.

4. Statistical model of effectiveness

In the discussion above, we have described an algorithm to insert custom interconnections and have provided figures of merit to evaluate the quality of design routes. The quantities provided by these expressions are not absolute numbers and require a statistical treatment to assess their validity. In this section, we present the salient points of a statistical model of effectiveness of external intervention. Complete details are in [12].




Figure 1

Region of influence in a design: (a) trial $(i - 1)$; (b) trial i . The region of influence R_i is shaded blue. The remainder $\bar{R}_i$ is unshaded. (1') and (1) show custom interconnections inserted in both trials $(i - 1)$ and i with total length $L_c^{(i-1)}$; (2') shows targeted routes in R_i with length $\Delta L_t^{(i-1)}$; (2) shows custom routes in R_i with length $\Delta L_t^{(i)}$; (3') shows routes that partially pass through R_i in trial $(i - 1)$ with length $L_o^{(i-1)}$; (3) shows routes that partially pass through R_i in trial i with total length $L_o^{(i)}$; (4') and (4) show remaining routes that do not pass through R_i with total length $L_r^{(i-1)}(\bar{R}_i)$ and $L_r^{(i)}(\bar{R}_i)$, respectively, in trials $(i - 1)$ and i .

Let R_i denote the area, or region of influence, that physically encloses new custom interconnections for ΔN_c^i signals in trial i , as shown in **Figure 1**. $\bar{R}_i$ denotes the complement of R_i and is the design area that does not contain new custom interconnections. In this figure, R_i is shaded blue; $\bar{R}_i$ is unshaded; (1) indicates an example of custom interconnections inserted in previous trials; (2) indicates an example of a new custom interconnection in trial i ; (3) indicates an example of an automated signal route that partially passes through R_i ; (4) indicates an automated signal route that does not pass through R_i . The variable N_c^i denotes the cumulative number of signals routed with custom routes in all trials, including trial i .

The total route length $L^{(i)}$ of all signal routes after completion of trial i is composed of four separate components: total length $\Delta L_c^{(i)}(R_i)$ of new custom interconnections contained in R_i in trial i ; total length $L_c^{(i-1)}$ of all custom routes in previous trials 0 to $(i - 1)$, where $L_c^{(i-1)} = \sum_{j=1}^{i-1} \Delta L_c^{(j)}(R_j)$; total length $L_o^{(i)}(R_i)$ of route segments routed by an automated router where at least part of the route passes through R_i ; and total length $L_r^{(i)}(\bar{R}_i)$ of route segments routed by an automated router where no part of the route passes through R_i . The total route length $L^{(i)}$ in trial i and $L^{(i-1)}$ in trial $(i - 1)$ are described by the expressions

$$L^{(i)} = L_c^{(i-1)} + \Delta L_c^{(i)}(R_i) + L_o^{(i)}(R_i) + L_r^{(i)}(\bar{R}_i), \quad (4)$$

and

$$L^{(i-1)} = L_c^{(i-1)} + \Delta L_t^{(i-1)}(R_i) + L_o^{(i-1)}(R_i) + L_r^{(i-1)}(\bar{R}_i), \quad (5)$$

respectively, where $\Delta L_t^{(i-1)}(R_i)$ is the total route length in trial $(i - 1)$ of signals targeted for custom routes in trial i .

Figure 1(a) shows examples of the four types of design routes in trial $(i - 1)$: custom interconnections (1') inserted with length $L_c^{(i-1)}$, targeted routes (2') in R_i with length $\Delta L_t^{(i-1)}$, other routes (3') that partially pass through R_i with length $L_o^{(i-1)}$, and remaining routes (4') that do not pass through R_i with length $L_r^{(i-1)}(\bar{R}_i)$. In Figure 1(b), the lengths of signals represented by those shown in (1)–(4) are described by the quantities $L_c^{(i-1)}$, $\Delta L_c^{(i)}(R_i)$, $L_o^{(i)}(R_i)$, and $L_r^{(i)}(\bar{R}_i)$, respectively. Note that custom interconnection lengths shown in (1) and (1') are equal to $L_c^{(i-1)}$.

In the following discussion, the simplified notation $\Delta L_c^{(i)}$, $\Delta L_t^{(i-1)}$, $L_o^{(i-1)}$, and $L_r^{(i-1)}$ is employed instead of $\Delta L_c^{(i)}(R_i)$, $\Delta L_t^{(i-1)}(R_i)$, $L_o^{(i-1)}(R_i)$, and $L_r^{(i-1)}(\bar{R}_i)$, respectively, and the full notation is provided only where confusion is possible. For trial $(i - 1)$, the variable $\tilde{L}^{(i-1)} = L^{(i-1)} - L_c^{(i-1)} = \Delta L_t^{(i-1)} + L_o^{(i-1)} + L_r^{(i-1)}$ represents the automated signal length that may be affected in trial i by the addition of custom interconnections with length $\Delta L_c^{(i)}$.

To evaluate the effectiveness of external intervention with custom interconnections, custom routes with length $\Delta L_c^{(i)}$ are first added to a design layout. The automated route program then completes the remaining routes. At this point, *variables that characterize the effectiveness of the external intervention in trial i are calculated*. The effectiveness $\epsilon_L^{(i)}$ of the external intervention on the total route length $L^{(i)}$ in trial i compared with the length in the preceding trial is represented by the expression

$$\epsilon_L^{(i)} = 1 - \frac{L^{(i)} - L_c^{(i-1)}}{\tilde{L}^{(i-1)}} = 1 - \frac{\Delta L_c^{(i)} + [L_o^{(i)} + L_r^{(i)}]}{[\Delta L_t^{(i-1)} + L_o^{(i-1)} + L_r^{(i-1)}]}, \quad (6)$$

where the numerator in both fractions represents the sum of $\Delta L_c^{(i)}$ and the total automated route length, and where the denominator in both fractions represents the total automated route length in trial $(i - 1)$. For $\epsilon_L^{(i)}$ to be positive, the numerator of the fractions in Equation (6) must be smaller than the denominator; in this case, the total automated wire length of the previous trial is replaced in the next trial with a lower combined wire length of custom routes and new automated routes.

The total effectiveness in Equation (6) can be represented as a weighted average of the following variables: effectiveness $\epsilon_{L_c}^{(i)}$ of the addition of custom interconnections in R_i , effectiveness $\epsilon_{L_o}^{(i)}$ of segments

created by an automated router where at least part of the signal route passes through R_i , and effectiveness $\epsilon_{L_r}^{(i)}$ of remaining route segments that do not pass through R_i , and are given by the expressions

$$\epsilon_{L_c}^{(i)} = 1 - \frac{\Delta L_c^{(i)}}{\Delta L_t^{(i-1)}}, \quad (7)$$

$$\epsilon_{L_o}^{(i)} = 1 - \frac{L_o^{(i)}}{L_o^{(i-1)}}, \quad (8)$$

and

$$\epsilon_{L_r}^{(i)} = 1 - \frac{L_r^{(i)}}{L_r^{(i-1)}}, \quad (9)$$

where each weight is the portion of the total length allocated to each type of route (custom, other, or rest), as described in [12].

We postulate that these three basic effectiveness variables are random variables distributed according to the normal distribution with (unknown) means μ_{L_c} , μ_{L_o} , and μ_{L_r} , respectively. These means are basic parameters that characterize the inherent effect of external intervention; the means are assumed to be invariant from one trial to another. The (unknown) variances of the effectiveness variables in Equations (7)–(9) are assumed to be inversely proportional to the domains of impact $\Delta L_t^{(i-1)}$, $L_o^{(i-1)}$, and $L_r^{(i-1)}$, respectively, with individual proportionality constants that are also assumed to be trial-invariant. In the following discussion, estimates of the means of the effectiveness variables and these proportionality constants are obtained from data in trials 1, 2, $\dots$, n . We assume that the correlation structure of the observed variables $(\epsilon_{L_c}^{(i)}, \epsilon_{L_o}^{(i)}, \epsilon_{L_r}^{(i)})$ is trial-invariant; this correlation structure is taken into account when performing statistical inference on these and derived parameters. We also assume that effectiveness variables observed in different trials are mutually independent.

To justify these assumptions, we note that external intervention in a given trial consists of a sequence of subtasks of roughly the same type and magnitude; accordingly, the domains of impact consist of subdomains of an appropriate unit size. The observed trial effectiveness is then an average of subdomain effectiveness variables, and under some relatively weak assumptions about the interdependence of variables related to subdomains, normality of the trial effectiveness variables is justified by the Central Limit Theory [13]. The assumption that the means of the effectiveness variables and the correlation structure remain invariant from trial to trial is justified by subtasks and related subdomains being similar from trial to trial; the proportionality constants represent variance of effectiveness observed for a subdomain of unit size. These assumptions may be oversimplified and should

be viewed as a first step toward understanding the properties of the collaborative effort that enable practical decision-making in an environment with a small amount of available data. With a better understanding of interactions between the router and supplemental algorithm, one may also wish to explore more elaborate models. There is also a possibility that this work can lead to further simplification; for example, one can conclude that μ_{L_r} can be assumed to be equal to zero, or some correlations can be assumed to be equal to zero, or both. In practice, the adequacy of the assumed models should be monitored with standard goodness-of-fit tests [13] to detect deviations from the model and to learn about their nature.

To describe the overall progress of external intervention in the design process, we introduce another figure of merit called the cumulative net length effectiveness, $\epsilon_L^{(i)(0)}$, defined by $\epsilon_L^{(i)(0)} = 1 - L^{(i)}/L^{(0)}$; this cumulative effectiveness is computed for trials $i = 1, 2, \dots, n$, where $L^{(0)}$ is the initial total net length. The additional superscript (0) in $\epsilon_L^{(i)(0)}$ indicates that the cumulative effectiveness is measured with respect to the initial wire layout. The distribution of the cumulative effectiveness depends entirely on the basic parameters introduced above.

Estimation of model parameters

The wire layout in each trial contains a different number of custom interconnections. From n independent trials, one can obtain unbiased estimates $\hat{\mu}_{L_c}$, $\hat{\mu}_{L_o}$, and $\hat{\mu}_{L_r}$ of the means of the effectiveness variables by calculating a weighted average of the appropriate effectiveness of individual trials. The standard errors and correlation structure of these estimators can be obtained from the data with standard statistical methods, as described in [12]. Furthermore, one can obtain standard errors for derived quantities, such as total effectiveness or projected effectiveness of future trials, as illustrated in the next section.

In practice, it is important to establish as soon as possible that the observed effects are not due to chance (that is, that the underlying means of the effectiveness variables are not zero or are even of the opposite sign). Evidence against the hypothesis of zero effect of external intervention can be measured with so-called *p-values*; these quantities are denoted as *p-value*(μ_{L_c}), *p-value*(μ_{L_o}), and *p-value*(μ_{L_r}) for custom routes, other routes, and the rest of the routes, respectively. For example, under the assumption that $\mu_{L_c} = 0$, *p-value*(μ_{L_c}) is the probability that the estimated effectiveness will exceed the value of $\hat{\mu}_{L_c}$ actually observed in n trials [14].

Confidence bounds for the effectiveness variables of external intervention with custom interconnections can also be constructed from the trial data. We postulate *a priori* that the effect of the use of custom interconnections is greater than zero; therefore, only the lower confidence




Figure 2

Number of route violations in the instruction fetch unit (IFU). Design-rule violations (circles) and electrical shorts (squares) are shown as a function of the number of custom interconnections N_c (lower abscissa) and as a fraction of custom interconnections N_c/N (upper abscissa).




Figure 3

Total interconnect length L_T (meters) of all N IFU signals as a function of N_c (lower abscissa) and as a fraction of custom interconnections N_c/N (upper abscissa). $L_{TS} = 5.43$ meters is the Steiner limit.

bound is computed. In the following discussion, the 95% lower confidence bound (LCB) for μ_{L_c} is denoted by $LCB_{0.95, \mu_{L_c}}$; the two-sided 95% confidence intervals for μ_{L_o} and μ_{L_r} are given by pairs of bounds $[LCB_{0.975, \mu_{L_o}},$

$UCB_{0.975, \mu_{L_o}}]$ and $[LCB_{0.975, \mu_{L_r}}, UCB_{0.975, \mu_{L_r}}]$, where UCB indicates the 95% upper confidence bound.

Statistical method to forecast further external intervention

In standard chip design practice, work on wire layout stops as soon as the signal connections in a wire layout contain no violations and satisfy the project cycle time requirement. Since the completion of this work typically occurs before the deadline to transfer the layout to manufacturing, an opportunity exists to improve the layout characteristics further. For example, one may have time to implement an additional trial. In this case, the question to address is this one: To what extent is this trial likely to improve the design quality? One can decide to proceed with an additional trial if the cumulative effectiveness variables of physical characteristics of interest are likely to improve; note that the proposed trial can improve some characteristics and degrade others. Tools to forecast the effects of additional external intervention can be most useful in this context.

We have developed a statistical framework for physical design to forecast the potential effect of custom interconnect design on physical properties of routes. Complete details are found in [15]. This framework provides designers with quantitative metrics of their efforts and with the ability to direct and plan additional efforts.

In this framework, we represent the estimated total effect for net lengths with the variable $\hat{\mu}_L^{(n+1)}$. The underlying mean effectiveness, $\mu_L^{(n+1)}$, is a function of the basic parameters; its estimate can be obtained in terms of $(\hat{\mu}_{L_c}, \hat{\mu}_{L_o}, \hat{\mu}_{L_r})$ from all previous trials 1, 2, $\dots$, n . Furthermore, one can forecast the total net length $L^{(n+1)}$ in the next trial $(n + 1)$ with the expression

$$L^{(n+1)} = L_c^{(n)} + \Delta L_t^{(n)}(R_{(n+1)})(1 - \hat{\mu}_{L_c}) + L_o^{(n)}(R_{(n+1)})(1 - \hat{\mu}_{L_o}) + L_r^{(n)}(R_{(n+1)})(1 - \hat{\mu}_{L_r}), \quad (10)$$

where $R_{(n+1)}$ is the region of influence projected to contain additional custom interconnections $\Delta N_c^{(n+1)}$ [15].

With this framework, one can forecast whether the proposed trial $(n + 1)$ should be attempted and whether this trial has a reasonable chance to improve physical properties of the design routes. In particular, the projected value of the total net length $L^{(n+1)}$ is compared with the value $L^{(n)}$ of trial n to ascertain whether intervention in trial $(n + 1)$ with proposed custom interconnections will further reduce the total cumulative length of the design interconnections.

In practice, a decision to proceed with an additional trial will not be made from predictions of net length alone; the predicted effect on other physical characteristics, such as the total number of vias, can be

obtained with the same approach. The projected effect on these characteristics also has to be considered and might point to an opposite conclusion. In this case, a decision to proceed must consider all projected changes in relevant physical characteristics of the layout. An example of this situation is discussed in the next section of this paper. In the following discussion, the p-value for net lengths in the projected trial ($n + 1$) is represented by the variable $p\text{-value}(\mu_L^{(n+1)})$. The 95% LCB for $\mu_L^{(n+1)}$ is represented by $LCB_{0.95, \mu_L^{(n+1)}}$.

5. Application to POWER4 instruction fetch unit

We now apply the statistical framework described in the preceding sections to the problem of routing a functional unit of the IBM POWER4 microprocessor [16–19].

External intervention

An automated routing tool² is not able to route three congested regions of the POWER4 instruction fetch unit (IFU) without introducing violations in some of the signal routes, as discussed in detail in [11], [12], and [15]. We refer to the congested regions as 1, 2, and 3 in the following discussion. **Figure 2** shows that the largest number of violations (electrical shorts and design-rule violations) occurs for $N_c = 0$, where the automated routing tool operates without external intervention. To assist the routing tool, custom interconnections are added for signals with large *NESL* in the congested regions according to the procedure described in Section 3.

Figure 3 shows the total length L_T and total Steiner length L_{TS} for IFU interconnections as a function of the number of custom interconnections N_c (lower abscissa) and the fraction of custom interconnections N_c/N (upper abscissa). This figure shows that L_T is reduced for signals targeted for custom interconnections as N_c is increased. For $N_c/N = 0.37$, the total interconnect length is reduced by approximately 25 mm (0.5%). **Figure 4** shows *TESL* as a function of N_c and N_c/N . For $N_c/N = 0.37$, the figure shows that *TESL* is reduced by 13%.

Figure 5 shows $\langle NESL \rangle$ for signals targeted for custom interconnections in three congested regions as a function of N_c (lower abscissa) and N_c/N (upper abscissa). The solid symbols in the figure indicate external intervention. Figure 5 shows that the external intervention is immediately effective in Region 3, where custom routes are introduced in trial 3, as shown by the red circle. In Region 2, the first external intervention occurs in trial 2, and the second external intervention occurs in trial 5, as shown by the green hexagons; the latter is more effective. As a result, $\langle NESL \rangle$ is reduced from 0.02 to 0.01 in




Figure 4

TESL (meters) for all N signals, signals targeted for custom interconnections, and signals with noncustom interconnections as a function of the number of custom interconnections N_c (lower abscissa) and as a fraction of custom interconnections N_c/N (upper abscissa).




Figure 5

Average normalized excess Steiner length $\langle NESL \rangle$ for IFU signals that are targeted for custom routes as a function of N_c (lower abscissa) and as a fraction of custom interconnections N_c/N (upper abscissa) in three congested regions. For Region 1, the solid black square indicates external intervention in trial 2. For Region 2, the solid green hexagons indicate external intervention in trials 1 and 5. For Region 3, the solid red circle indicates external intervention in trial 3.

Region 2 and from 0.055 to 0.02 in Region 3. In Region 1, the effect of the external intervention is not immediately

² For some routers, the automated routing tool might generate different route solutions from different random seeds. In our experiments, the Cadence Silicon Ensemble** router generates the same solution for each specified set of input cost factors.

Table 1 Effectiveness of external intervention for trials $i = 1, 2, \dots, 6$ in the instruction fetch unit. For net lengths L (upper half of the table): $\epsilon_{L_c}^{(i)}, \epsilon_{L_o}^{(i)}, \epsilon_{L_r}^{(i)}, \epsilon_L^{(i)}$ (%) with the corresponding values for $\Delta L_t^{(i-1)}, L_o^{(i-1)}, L_r^{(i-1)}, L^{(i-1)}$ (meters); For vias V (bottom half of the table): $\epsilon_{v_c}^{(i)}, \epsilon_{v_o}^{(i)}, \epsilon_{v_r}^{(i)}, \epsilon_v^{(i)}$ (%) with the corresponding values for $\Delta v_t^{(i-1)}, v_o^{(i-1)}, v_r^{(i-1)}, v^{(i-1)}$.

L	Custom interconnect		Other routes		Rest of routes		All routes	
i	$\epsilon_{L_c}^{(i)}$	$\Delta L_t^{(i-1)}$	$\epsilon_{L_o}^{(i)}$	$L_o^{(i-1)}$	$\epsilon_{L_r}^{(i)}$	$L_r^{(i-1)}$	$\epsilon_L^{(i)}$	$L^{(i-1)}$
1	9.2	1.05	-4.5	1.78	-0.57	2.68	0.027	5.51
2	7.0	0.163	-1.3	1.14	-0.064	4.20	-0.11	5.51
3	7.7	0.52	-0.19	0.100	0.013	3.78	0.92	5.51
4	-13.0	0.13	-1.2	0.24	0.099	3.53	-0.40	5.47
5	5.3	0.061	0.0043	0.50	-0.0012	3.20	0.086	5.49
6	-4.3	0.23	-1.4	0.83	-0.026	2.64	0.051	5.48

V	Custom interconnect		Other routes		Rest of routes		All routes	
i	$\epsilon_{v_c}^{(i)}$	$\Delta v_t^{(i-1)}$	$\epsilon_{v_o}^{(i)}$	$v_o^{(i-1)}$	$\epsilon_{v_r}^{(i)}$	$v_r^{(i-1)}$	$\epsilon_v^{(i)}$	$v^{(i-1)}$
1	69	7630	-1.7	26210	-2.1	32515	6.2	66355
2	59	2016	-7.3	16109	0.44	41730	0.34	62255
3	65	8327	-0.52	1162	0.80	49339	9.8	62052
4	25	1119	0.66	2290	0.088	46702	0.67	56279
5	78	2682	2.7	5048	-0.41	41206	4.2	55946
6	73	976	-16	5181	1.1	40129	0.71	53892

Table 2 Estimated effectiveness of external intervention in the instruction fetch unit. The mean effectiveness, standard error, p-value, lower confidence bound LCB, and upper confidence bound UCB are shown in % for custom interconnections, the other routes, and the rest of the routes. The confidence intervals have 95% coverage; for μ_{L_c} and μ_{v_c} , only the lower 95% bound is given.

Net lengths	Mean (%)	Standard error (%)	p-value	LCB (%)	UCB (%)
Custom	$\tilde{\mu}_{L_c} = 6.7$	$\hat{\sigma}(\hat{\mu}_{L_c}) = 2.3$	0.015	2.2	—
Other	$\tilde{\mu}_{L_o} = -2.4$	$\hat{\sigma}(\hat{\mu}_{L_o}) = 0.85$	0.039	-4.5	-0.18
Rest	$\tilde{\mu}_{L_r} = -0.074$	$\hat{\sigma}(\hat{\mu}_{L_r}) = 0.091$	0.46	-0.31	0.16

Vias	Mean (%)	Standard error (%)	p-value	LCB (%)	UCB (%)
Custom	$\tilde{\mu}_{v_c} = 65.4$	$\hat{\sigma}(\hat{\mu}_{v_c}) = 4.4$	0.000013	56.5	—
Other	$\tilde{\mu}_{v_o} = -4.1$	$\hat{\sigma}(\hat{\mu}_{v_o}) = 2.2$	0.12	-9.7	1.5
Rest	$\tilde{\mu}_{v_r} = 0.085$	$\hat{\sigma}(\hat{\mu}_{v_r}) = 0.43$	0.85	-1.0	1.2

visible, and $\langle NESL \rangle$ remains approximately constant at 0.02 as N_c increases, because the overall length of the custom interconnections is approximately equal to the overall net length of the automated routes in this region.

During the process of external intervention described in Section 3, the total interconnect length can increase, as indicated in Figure 3 for trial 2. In addition, the observed effectiveness for some characteristics can be negative,

as shown later for net lengths in Figure 6 for trial 2. These observations can occur as a result of statistical fluctuations; recall that the goal of the process is to achieve a wire layout with zero violations as opposed to minimal net length. Reduction in total interconnect length or number of vias or both is a useful by-product of the external intervention and can predict the degree of success toward achieving the goal, but is not a goal itself. Even if

the goal is achieved, and if a decision is made to continue the collaboration, the new focus is on other design properties such as yield and performance. A reduction in total net length or number of vias or both is primarily of interest because this reduction is associated with better design properties.

As discussed in the next section, the main benefit of the collaborative effort for this particular design layout is a reduction in the number of vias as opposed to reduction in the net length. In essence, the message is that the collaborative effort enables one to reduce the number of vias significantly without paying a penalty in terms of increased net length.

Statistical model

Custom interconnections are added in the IFU in six trials ($n = 6$). Complete details may be found in [12]. Associated with each trial is a different number of custom interconnections N_c and different region of influence R_i . **Tables 1 and 2** show that the estimated effect of custom interconnections is distinctly positive, with 6.7% effectiveness and 95% LCB of +2.2%. The effect is statistically significant, as the associated $p\text{-value}(\mu_{L_c}) = 0.051\%$. For the other routes that pass through R_i , the effectiveness is -2.4% with $p\text{-value}(\mu_{L_o}) = 0.039 < 0.05$; this result provides evidence that routes in R_i are negatively affected by the custom routing activity. The magnitude of this impact can be characterized by the confidence interval for μ_{L_o} of $[-4.5\%, -0.18\%]$. Although net lengths of other signals are increased, the longer routes of these signals do not negatively affect other design requirements, such as signal timing requirements. For the rest of the routes that do not pass through R_i , the estimated effect is $\tilde{\mu}_{L_r} = -0.074\%$, with standard error 0.091%. (In this discussion, $\tilde{\mu}$ is used instead of $\hat{\mu}$ to designate values of estimators that are actually measured from our data.) There is no evidence of negative impact caused by the custom routing activity since the $p\text{-value}(\mu_{L_r}) = 0.46 > 0.05$. The 95% confidence interval for μ_{L_r} is $[-0.31\%, 0.16\%]$.

Tables 1 and 2 indicate that external intervention with custom interconnections has a strong positive effect on the number of vias in R_i . For example, the number of vias is reduced by 65.4% in signal routes targeted for custom interconnections; this effect is statistically significant, as the associated $p\text{-value}(\mu_{v_c}) = 0.000013\%$; the 95% LCB is $56.5\% \gg 0$. For the other routes and the rest of the routes, there is no evidence of negative impact caused by the custom routing activity. For the other routes, the estimated effect is $\tilde{\mu}_{v_o} = -4.1\%$ with standard error 2.2%, resulting in $p\text{-value}(\mu_{v_o}) = 0.12 > 0.05$; the 95% confidence interval for μ_{v_o} is $[-9.7\%, 1.5\%]$. For the rest of the routes, the estimated effect is $\tilde{\mu}_{v_r} = 0.09\%$ with standard error




Figure 6

Measured cumulative effectiveness variables for (a) net lengths $\epsilon_L^{(i)(0)}$ and (b) vias $\epsilon_v^{(i)(0)}$ in the IFU for each trial $i = 1, 2, \dots, 6$. The solid dots show the projected values of cumulative effectiveness in trial $i = 7$ for net lengths (a) and vias (b).

0.43%, resulting in $p\text{-value}(\mu_{v_r}) = 0.85 > 0.05$; the 95% confidence interval for μ_{v_r} is $[-1.0\%, 1.2\%]$.

Figures 6(a) and 6(b) show values of the cumulative effectiveness variables for net lengths $\epsilon_L^{(i)(0)}$ and vias $\epsilon_v^{(i)(0)}$, respectively, for trials $i = 1, 2, \dots, 6$ as well as for projected trial 7. The figure shows that the values of the cumulative effectiveness variables for net lengths and vias tend to increase with each additional trial $i = 1, 2, \dots, 6$; they reach 0.41% and 19.3%, respectively, in trial $i = 6$. For projected trial $i = n + 1 = 7$, Figure 6 shows that the value for cumulative effectiveness is approximately constant for vias and is reduced for net lengths. The reduction for net lengths is predicted primarily because of the negative value of $\hat{\mu}_{L_o}$ and the relatively large value of the domain of impact $L_o^{(6)}$. In this situation, since there is a conflict between the projected results for net lengths and vias, the question that arises at trial 6 is whether to

Table 3 Projected effectiveness of external intervention in the instruction fetch unit. The projected effectiveness, standard error, p-values, and 95% lower confidence bounds (LCBs) are shown for net lengths and vias for the proposed trial $i = n + 1 = 7$.

<i>All nets</i>	<i>Mean (%)</i>	<i>Standard error (%)</i>	<i>p-value</i>	<i>LCB (%)</i>
Net lengths	$\hat{\mu}_L^{(7)} = -1.0$	$\hat{\sigma}(\hat{\mu}_L^{(7)}) = 0.24$	$p\text{-value}(\mu_L^{(7)}) = 0.996$	-1.5
Vias	$\hat{\mu}_v^{(7)} = 0.10$	$\hat{\sigma}(\hat{\mu}_v^{(7)}) = 0.75$	$p\text{-value}(\mu_v^{(7)}) = 0.45$	-1.4

proceed with trial 7. To answer this question, we can decide to act in accordance with one of several policies. For example, one policy is to proceed only if we see strong evidence for improvement in the projected trial; another possible policy is to proceed unless there is evidence that the result would be worse. In both cases, the decisions whether to proceed are primarily based on the p-values of the projected effectiveness variables.

Table 3 summarizes the projected results for proposed trial $i = 7$; the table shows a negative projected mean effectiveness for net lengths and a positive projected mean effectiveness for vias. The table also shows that there is strong evidence for negative impact on net lengths. In fact, the probability to measure $\hat{\mu}_L^{(7)} \leq -1.0$ under the assumption that there is no negative impact is smaller than $1 - p\text{-value}(\mu_L^{(7)}) = 1 - 0.996 = 0.004$. Therefore, if we act in accordance with the second policy, we do not proceed with trial 7. This conclusion is also reached under the first policy; in fact, for the observed data, one does not even need p-values to establish that there is no evidence for strong improvement in the projected trial. Therefore, we decided not to add the projected custom routes to the design layout, and declare the design layout complete at trial $n = 6$. Complete details are found in [15].

6. Discussion

The approach presented in this paper is an initial attempt to formalize interactions in a collaborative effort aimed at achieving a certain goal. It is quite likely that some of the assumptions made in developing this approach will require modification for more complex forms of interaction; such modifications can either be developed empirically or with detailed knowledge about the nature of the interaction. With statistical methods for monitoring model adequacy and simulation analysis, one should be able to arrive at models that are useful for a problem of interest and examine the effects of possible incorrect specification of the model or violation of the basic assumptions.

The problem of wire-route configurations that we choose to illustrate the proposed methodology appears to have long-term benefits. Because of increasing design complexity, it is generally difficult for routing programs to achieve a routable design layout within specified constraints such as geometry or real estate, in addition to multiple objectives imposed by current engineering

requirements. Furthermore, if a router fails to produce an acceptable design layout, it is typically impossible to determine whether this failure occurs because of shortcomings of its own algorithm or because no routable design layouts exist.

In principle, the development of forms of external intervention that predictably improve on the performance of a router is similar to the development of a new drug that enhances the ability of the body to heal itself. In the latter case, the collaborative effort involves a drug and the body's immune system. Indeed, our external interventions can be viewed as a form of drug; as with drugs, one cannot assume blindly that the intervention will succeed, and a rigorous trial phase is required. Implementation of the proposed methodology can lead to a library of proven external interventions that has the potential to become a differentiator in the field, providing a framework for knowledge accretion and preservation.

Another situation relevant to this approach is robust design of complex systems; these systems include, for example, queueing networks or inventory management systems. Design of these systems typically involves the use of some heuristic algorithm in conjunction with computerized experimentation and simulations. When considering a potential enhancement of the heuristic approach, one may treat this enhancement as a supplemental algorithm until its effectiveness is firmly established on the basis of a statistical analysis similar to that described in this paper. At some later point, such a supplemental algorithm is included as an integral part of the core heuristic approach, with another candidate supplemental algorithm taking its place. This form of collaboration between the basic approach and candidate supplemental algorithms can be used as a basis for continual improvement.

7. Conclusions

To illustrate a collaborative system in which several contributors strive to achieve a goal, we have considered a specific example of interaction between an automated routing program and a supplemental pre-routing algorithm. The statistical framework that we have introduced to assist computer-aided design tools in typical design processes should be applicable for other problems involving collaborative efforts. The relevant class of

problems consists of those problems in which contributors seek to achieve a goal by proceeding one step at a time without interfering with one another.

In high-technology industries, collaborative solutions to complex problems involve the participation of highly skilled personnel and sophisticated computer programs. This work is a preliminary attempt to establish a mathematical framework to formalize the interaction of these entities and to optimize these interactions.

Acknowledgments

We thank Brian Wallace of the Customer Support Group at Cadence Design Systems in Bellevue, WA, for support of the automated routing tool Silicon Ensemble used in this research. We thank Edward Hughes and Philip Honsinger of the IBM Microelectronics Division in East Fishkill, New York, for their explanation of the Steiner algorithm presented in the footnote. We thank the POWER4 IFU team members in Austin, Texas, East Fishkill, New York, Poughkeepsie, New York, and Yorktown Heights, New York; Ed Seymour at the IBM Server Group in Austin, Texas; and Evanthis Papadopoulou, Jim Harper, Jim Ryan, and Leon Stok at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, for discussions. We thank Ellen Eide, Pong-Fei Lu, Sampath Purushothaman, Zhigang (David) Pan, Anna Topol, and Robert Wisniewski at the IBM Thomas J. Watson Research Center and Brian Konigsburg at the IBM Server Group in Austin, Texas, for their comments on the text. We thank Izzy Bendrihem, Tom Bucelot, and Kelvin Lewis at the IBM Thomas J. Watson Research Center for providing a seamless computing environment. We thank Jim Leonard and Katie Pathe of IBM Watson Library Services for literature searches used in this work.

**Trademark or registered trademark of Cadence Design Systems, Inc.

References

1. H. B. Bakoglu and J. D. Meindl, "Optimal Interconnection Circuits for VLSI," *IEEE Trans. Electron Devices* **ED-32**, 903–909 (May 1985).
2. M. T. Bohr, "Interconnect Scaling—The Real Limiter to High Performance ULSI," *IEDM Tech. Digest*, pp. 241–244 (1995).
3. A. Deutsch, G. V. Kopcsay, C. W. Surovic, B. J. Rubin, L. M. Terman, R. P. Dunne, Jr., T. A. Gallo, and R. H. Dennard, "Modeling and Characterization of Long On-Chip Interconnections for High-Performance Microprocessors," *IBM J. Res. & Dev.* **39**, 547–567 (September 1995).
4. G. A. Sai-Halasz, "Performance Trends in High-End Processors," *Proc. IEEE* **83**, 20–36 (January 1995).
5. A. K. Stamper, "Interconnection Scaling 1 GHz and Beyond," *IBM MicroNews* **4**, 1–12 (1998).
6. M. Pedram, "Tutorial and Survey Paper: Power Minimization in IC Design: Principles and Applications," *ACM Trans. Design Automation Electron. Syst.* **1**, 3–56 (January 1996).
7. ITRS Semiconductor Industry Association, *The International Technology Roadmap for Semiconductors*, 2001.
8. R. Ho, K. Mai, H. Kapadia, and M. Horowitz, "Interconnect Scaling Implications for CAD," *IEEE International Conference on Computer-Aided Design, Digest of Technical Papers*, 1999, pp. 425–429.
9. R. W. Keyes, "The Wire-Limited Logic Chip," *IEEE J. Solid-State Circuits* **SC-17**, 1232–1233 (December 1982).
10. J. D. Meindl, "Opportunities for Gigascale Integration," *Solid State Technol.*, pp. 85–89 (December 1987).
11. M. Y. L. Wisniewski, E. Yashchin, R. L. Franch, D. Conrady, D. Maynard, G. Fiorenza, and I. C. Noyan, "The Physical Design of On-Chip Interconnections, Part I: Quantification of Interconnect Properties," *Research Report RC-22345*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, January 20, 2002.
12. M. Y. L. Wisniewski, E. Yashchin, R. L. Franch, D. Conrady, G. Fiorenza, and I. C. Noyan, "The Physical Design of On-Chip Interconnections, Part II: Statistical Model of Effectiveness," *Research Report RC-22344*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, January 20, 2002.
13. Ralph B. D'Agostino and Michael A. Stephens, *Goodness-of-Fit Techniques*, Marcel Dekker, New York 1996.
14. Bernard Rosner, *Fundamentals of Biostatistics*, Duxbury Press, Boston, 1986.
15. M. Y. L. Wisniewski, E. Yashchin, R. L. Franch, D. Conrady, G. Fiorenza, and I. C. Noyan, "The Physical Design of On-Chip Interconnections, Part III: Statistical Model of Forecasting," *Research Report RC-22346*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, January 20, 2002.
16. C. J. Anderson, J. Petrovick, J. M. Keaty, J. Warnock, G. Nussbaum, J. M. Tendler, C. Carter, S. Chu, J. Clabes, J. DiLullo, P. Dudley, P. Harvey, B. Krauter, J. LeBlanc, P.-F. Lu, B. McCredie, G. Plumb, P. J. Restle, S. Runyon, M. Scheuermann, S. Schmidt, J. Wagoner, R. Weiss, S. Weitzel, and B. Zoric, "Physical Design of a Fourth-Generation POWER GHz Microprocessor," *IEEE International Solid-State Circuits Conference, Digest of Technical Papers*, 2001, pp. 232–233.
17. K. Diefendorff, "Power4 Focuses on Memory Bandwidth," *Microprocessor Report* **13**, 1–8 (October 1999).
18. IBM Enterprise Server pSeries 680; see <http://www-1.ibm.com/servers/eserver/pseries/hardware/enterprise/>.
19. J. D. Warnock, J. M. Keaty, J. Petrovick, J. G. Clabes, C. J. Kircher, B. L. Krauter, P. J. Restle, B. A. Zoric, and C. J. Anderson, "The Circuit and Physical Design of the POWER4 Microprocessor," *IBM J. Res. & Dev.* **46**, 27–51 (January 2002).

Received March 14, 2002; accepted for publication May 21, 2002

Mary Y. L. Wisniewski *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (myl@us.ibm.com).* Dr. Wisniewski received an A.B. degree *summa cum laude* from Harvard-Radcliffe Colleges, an M.Phil. degree from Cambridge University, and a Ph.D. degree in 1997 from Cornell University, all in physics. In 1996, she joined IBM as a Research Staff Member at the Thomas J. Watson Research Center, Yorktown Heights, New York. She is a member of the VLSI Design Department and was the integrator of the instruction fetch unit of the IBM POWER4 microprocessor. She is currently working on issues in CMOS VLSI design. Dr. Wisniewski is a member of the Institute of Electrical and Electronics Engineers (IEEE), IEEE-Lasers and Electro-Optics Society (LEOS), American Physical Society, Materials Research Society, and Optical Society of America. She was recently elected to the IEEE-LEOS Board of Governors.

Emmanuel Yashchin *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (yashchi@us.ibm.com).* Dr. Yashchin received a Diploma in applied mathematics from Vilnius State University (U.S.S.R.) in 1974 and an M.Sc. degree in operations research and a D.Sc. degree in statistics from the Technion-Israel Institute of Technology in 1977 and 1981, respectively. In 1982 he was a Visiting Assistant Professor at Iowa State University. He joined IBM in 1983 as a Research Staff Member in the Statistics group of the Mathematical Sciences Department at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York. Dr. Yashchin's research interests include change point models, quality control, reliability, and spatial statistics. He is a member of the American Statistical Association, Institute of Mathematical Statistics, Royal Statistical Society, and American Society for Quality.

Robert L. Franch *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (franch@us.ibm.com).* Mr. Franch received the B.S.E.E. degree from the Polytechnic Institute of New York, Brooklyn, in 1980. That same year he joined IBM in East Fishkill, New York, where he worked on bipolar device reliability and stress testing of VLSI devices and interconnects. In 1984 he joined the IBM Research Division at the Thomas J. Watson Research Center as a member of the Test Systems group, where he worked on high-speed functional testing of n-MOS and CMOS chips developed at IBM Research. In 1993 he joined the VLSI Design group, where he has worked on circuit design for microprocessors. Mr. Franch has since been engaged in the design of CMOS circuit macros and has contributed to several generations of IBM S/390 microprocessors.

David P. Conrady *Sequence Design Inc., 112 Avenue Kleber, 75016 Paris, France (dconrady@sequencedesign.com).* Mr. Conrady received the B.A. degree from Rice University in computer science and French in 1985. That same year he joined IBM in East Fishkill, New York, where he developed automatic placement software for bipolar and CMOS VLSI designs. In 1993 he was the integrator of a chip in the RS/6000 microprocessor family in Austin, Texas. He was subsequently engaged in layout and timing optimization of several generations of IBM S/390 microprocessors at IBM in Boeblingen, Germany, and Poughkeepsie, New York. For the work reported in this paper, Mr. Conrady worked at IBM in

the Microelectronics Division in Hopewell Junction, New York. In 2001 he joined Sequence Design Inc., where he supports timing and signal-integrity optimization of high-speed nanometer VLSI designs in Europe.

Giovanni Fiorenza *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (gfiorenz@us.ibm.com).* Dr. Fiorenza received the B.S., M.S., and Ph.D. degrees in physics from Stevens Institute of Technology in 1979, 1984, and 1992, respectively. He joined the U.S. Air Force in 1979 and worked at the Air Force Weapons Laboratory, Kirtland AFB, Albuquerque, New Mexico, on laser damage in thin films until 1981. He joined IBM Component Vendor Assurance, Poughkeepsie, New York, in 1984 and worked on technology assessment, functional reliability, and soft-error evaluations of n-MOS and CMOS technologies until 1988, when he moved to IBM East Fishkill, New York, where he worked on the development of high-performance bipolar/BiCMOS technologies until 1992, and on device design and development of device simulation software until 1996. In 1996, he joined the IBM Thomas J. Watson Research Center in Yorktown Heights, New York, where he worked on circuit and physical design for the POWER4 microprocessor. He is currently working on circuit and physical design for high-performance microprocessors. Dr. Fiorenza holds one patent and is a member of the IEEE and Computer Society.

I. Cevdet Noyan *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (noyan@us.ibm.com).* Dr. Noyan is a Research Staff Member with the Thin Film Metallurgy group at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York. He has been with IBM since 1985 and has worked on computer structures at the box, package, and chip levels. He is currently working on thin-film wiring issues.