Calibration of Rent's Rule Models for Three-Dimensional Integrated Circuits

Shamik Das, *Member, IEEE*, Anantha P. Chandrakasan, *Fellow, IEEE*, and Rafael Reif, *Fellow, IEEE*

***Abstract*—In this paper, we determine the accuracy of Rahman's interconnect prediction model for three-dimensional (3-D) integrated circuits. Utilizing this model, we calculate the wiring requirement for a set of benchmark standard-cell circuits. We then obtain placed and routed wirelength figures for these circuits using 3-D standard-cell placement and global-routing tools we have developed. We find that the Rahman model predicts wirelengths accurately (to within 20% of placement and of routing, on average), and suggest some areas for minor improvement to the model.**

***Index Terms*—Rent's Rule, system-level interconnect prediction (SLIP), three-dimensional (3-D) integration.**

I. Introduction

TECHNOLOGY feature sizes continue to shrink to meet performance demands on integrated circuits. This, coupled with growing overall chip dimensions, leads to greater consumption of the available power and delay budgets by the interconnect structures on these chips [2]. As global and semiglobal wires become increasingly expensive and clock frequencies become higher and higher, designers seek new architectures and technologies that rely less on sending signals across the chip [3]. However, few scalable solutions have been proposed.

One such solution is three-dimensional (3-D) integration. In a 3-D integrated circuit, transistors may be fabricated on top of other transistors, resulting in multiple layers of active components. These transistors may then be wired to other transistors on the same device layer, to transistors on different device layers, or both, depending on the process technology. Several different approaches to fabricating 3-D circuits or 3-D-compatible transistors have been taken [4]–[7]. These vary in terms of the maximum number of device layers and the maximum density of interconnects between these layers. Some technologies, such as vertical multi-chip module packaging (MCM-V), only permit interconnects at the boundary of device layers, thus, making MCM-V circuits effectively like folded two-dimensional (2-D) circuits in most cases, in terms of their interconnect performance. Other technologies, such as epitaxial overgrowth, permit direct connections between adjacent transistors on different device layers, but are not easily scaled to more than just a couple such layers. The wafer-bonding approach shown in [7], where discrete wafers are "glued" together using a copper interconnect interface, permits multiple wafers and multiple third-dimension interconnects, overcoming the above limitations.

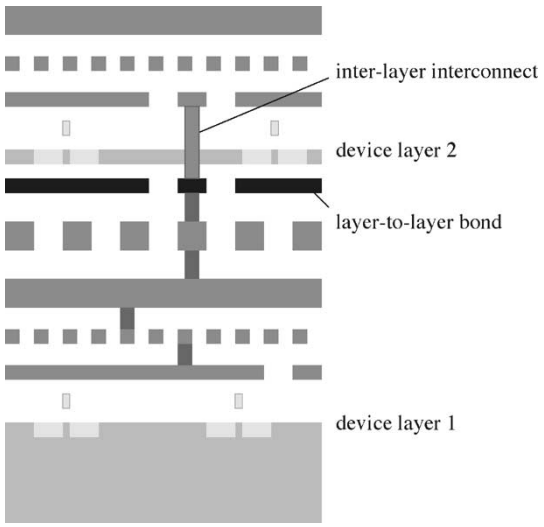

Fig. 1. Wafer-bonded structure with two device layers and copper interconnect interface. (Figure courtesy A. Fan.)

The particular characteristics of a chosen 3-D technology will assuredly impact the degree to which circuits may benefit from the technology. In the above wafer-bonding technique, for example, the bottom device layer is formed from a bulk silicon wafer 500–700 microns thick; subsequent device layers may be as thin as 1 or 2 microns if silicon-on-insulator (SOI) wafers are used. The thickness of metallization layers also adds to the total layer-to-layer thickness in this type of 3-D integrated circuit, as does the thickness of the copper interwafer interface. Wafers may be bonded face-to-face (i.e., such that the metallizations are adjacent) or face-to-back; note, however, that face-to-back bonding is required for three or more device layers. With face-to-back bonding, interconnections between wafers are formed using high-aspect-ratio vias through the device area of the upper wafer. A diagram of a face-to-back wafer-bonded structure is shown in Fig. 1.

In order to evaluate the potential of these 3-D integration technologies, Rahman *et al.* developed an interconnect-prediction model for 3-D integrated circuits [1]. This model (to which we will refer to as the Rahman model) is an extension of the Davis stochastic predictive model for conventional 2-D integrated circuits [8]. However, while the accuracy of the Davis model has been verified with numerous test cases, until now it has not been possible to verify the Rahman model due to a lack of placed instances of circuits in a 3-D integration technology.

Manuscript received August 30, 2002; revised July 3, 2003. This work was supported in part by the MARCO Focused Research Center on Interconnects, under the Massachusetts Institute of Technology, subcontract B-12-M00; in part by the Georgia Institute of Technology, Atlanta, GA, and in part by DARPA under Grant B-12-D00.

The authors are with the Microsystems Technology Laboratories, Massachusetts Institute of Technology, Cambridge, MA 02139-4307 USA.

Digital Object Identifier 10.1109/TVLSI.2004.825833

To this end, we have developed standard-cell placement and global-routing tools for 3-D integrated circuits by extending state-of-the-art 2-D placement and routing techniques to the 3-D case. Using these tools, we have obtained multiple-device-layer placements for a set of benchmark circuits. With these results from placement and routing, we may then determine the accuracy of the Rahman model.

II. Predictive Model and Assumptions

A. Overview of the Rahman–Davis Models

The Rahman model uses Rent's Rule [9] to predict a wirelength distribution for a given circuit. This distribution, $f_{3D}(l)$, gives the number of wires of length l within this circuit. Taking an appropriately-weighted sum of f_{3D} over all valid values of l (i.e., from the shortest possible length to the die-edge length) yields the total wirelength, interconnect power consumption, or interconnect delay of the circuit. This distribution function can thus be used to make important determinations about circuit performance based solely on *a priori* measures of circuit parameters.

Specifically, Rent's Rule stipulates that in a well-partitioned design, the number of logic gates N within a submodule of the design and the number of interconnection terminals T of this module obey the relationship

$$T = kN^p \tag{1}$$

where k, the Rent coefficient, is the average number of terminals per logic gate and p is the Rent exponent, which is a design-dependent fixed quantity that is around 2/3 for typical circuits. Using Rent's Rule, we may predict the number of interconnects of length l gate pitches from a given logic gate

$$I_{3D}(l) = \frac{\alpha k}{N_c}\left[(N_a+N_b)^p - N_b^p + (N_b+N_c)^p - (N_a+N_b+N_c)^p\right] \tag{2}$$

where $\alpha = (\text{f.o.})/(1+\text{f.o.})$ and f.o. is the average fan-out, $N_a = 1$ is the logic gate in question, N_c is the number of gates at Manhattan distance l gate pitches from this gate (scaled so as to avoid double-counting), and N_b is the number of gates between the gate in question and the N_c gates at distance l, as derived in [1], [8]. We assume the gates to be of uniform size. A diagram depicting the distribution N_a, N_b, N_c of logic gates is shown in Fig. 2.

Knowing the number of interconnects of length l originating at a gate, we now must compute the number $M_{3D}(l)$ of gate pairs separated by length l. This is done by counting gates in the grid in Fig. 2. Specifically, the number $M_{2D}(l)$ of gate pairs in a single device layer that are separated by Manhattan distance l gate pitches is given in [8] by

$$M_{2D}(l) = \begin{cases} \frac{1}{3}l^3 - l^2 l_{\max} + \frac{1}{2}l_{\max}^2 l, & 1 \le l < \frac{l_{\max}}{2} \\ \frac{1}{3}(l_{\max}-l)^3, & \frac{l_{\max}}{2} \le l < l_{\max} \end{cases} \tag{3}$$

where $l_{\max}$ is the maximum length of a 2-D wire in gate pitches. Given $M_{2D}(l)$, we may write

$$M_{3D}(l) = \sum_{i=0}^{N_z-1} \beta_i M_{2D}(l - it_z)u(l - it_z) \tag{4}$$

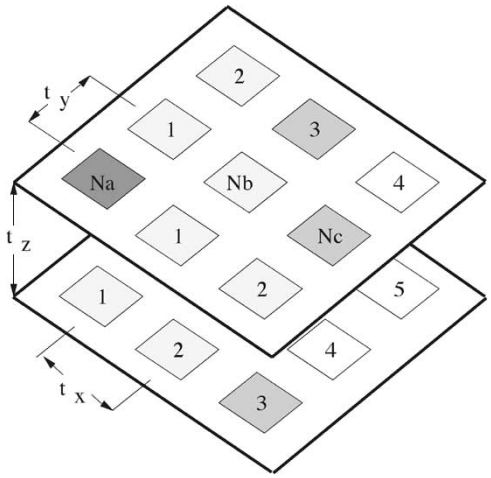

Fig. 2. Derivation of the 3-D wirelength distribution: $N_a = 1$ is the logic gate in question, N_c is the number of target logic gates at Manhattan distance l gate pitches, and N_b is the number of logic gates in between. t_x, t_y, and t_z are the gate width, height, and interlayer thickness in microns, respectively. (Figure courtesy A. Rahman.)

where u is the unit-step function, N_z is the number of device layers, t_z is the interlayer thickness (expressed here in units of gate pitches), and β_i are constants that depend on the number of device layers and the range of interlayer interconnects [1]. (We assume for the purposes of this paper that interlayer interconnects may range over all device layers, i.e., that there is no restriction on the number of device layers that a third-dimension interconnect may span.)

Having computed I_{3D} and M_{3D}, the wirelength distribution f_{3D} may be computed as in [1]

$$f_{3D}(l) = \Gamma M_{3D}(l) I_{3D}(l). \tag{5}$$

The quantity Γ is a normalization constant such that the sum of f_{3D} over all valid l matches the total number of interconnects in the circuit.

For a given circuit, then, we may compute N_a, N_b, and N_c and thereby obtain I_{3D}, M_{3D}, and f_{3D}. In order to do so, we require certain *a priori* values: the Rent exponent and coefficient, die size in gate pitches, gate fan-out, and number of device layers to be targeted. To convert these values from units of gate pitches to meters of wire, we require the gate height, average gate width, and interlayer thickness t_z in microns. We discuss how these values are obtained in the following sections.

B. Adaptations for 3-D Standard-Cell Placement and Routing

1) Dimensions: In standard-cell circuits, the gates and low-level modules are synthesized as rectangular cells of fixed height and variable width. The die area is specified *a priori* as a fixed number of rows with fixed height and width and fixed inter-row spacing (i.e., a *fixed-die* context). For 3-D circuits, we preserve the fixed-die context by scaling the number of rows and width of rows by the square root of the number of device layers, so as to maintain constant area for cell placement.

To determine the gate count and gate pitch, we use the size of the narrowest cell as the unit gate. The gate count is then equal to the total cell width divided by the width of the unit gate, and the horizontal and vertical gate pitch are given by the width of

the unit gate and the row-to-row pitch given by the fixed-die context respectively.

The layer-to-layer thickness t_z is given by the technology. In a wafer-bonded circuit, for example, t_z may be as low as a few microns. If a solder-bump interconnect interface is used, on the other hand, the thickness may not increase appreciably, but the capacitance of interwafer interconnects may increase by as much as an order of magnitude. For MCM-V packages, the average die-to-die interconnect length scales with the die size. Thus, one may wish to determine t_z as a function of electrical parameters (or other parameters) of the interlayer interconnect, rather than strictly as the distance between device layers.

2) Determination of Rent Parameters: Rent parameters for a given circuit traditionally are determined by recursive partitioning of the circuit netlist. However, it has been shown that a similar version of the Rent parameters can be derived from placement. Furthermore, Rent parameters from placement are believed to reflect more accurately the distribution of wires and the quality of placement. Indeed, wirelength estimation for 2-D circuit placements using placement-based Rent parameters is more accurate [10].

There are two generally-accepted ways to compute the Rent parameters. In both cases, gate counts and terminal counts are computed for submodules of the circuit at various levels of hierarchy [9]. One may then find simultaneously the Rent coefficient and Rent exponent that provide the best fit to the data, as in [10]. This typically produces a Rent coefficient that differs from the average number of terminals per gate. Alternatively, one may compute the Rent coefficient using its definition as the average number of terminals per gate, and then find the Rent exponent that best fits the data using this Rent coefficient. We will use the latter method, since the predictive model assumes this value for the Rent coefficient.

It is expected that the greater number of nearest neighbors available to transistors in 3-D integrated circuits leads to a shift in the wirelength distribution toward local wires and a reduced need for inter-partition interconnects at any fixed partition size. In other words, the Rent exponent derived from a 3-D placement of a given circuit is expected to be less than that derived from a 2-D placement; in particular, for a given circuit

$$p_{\text{partition}} < p_{n+1} < p_n < p_1 \tag{6}$$

for modest values of n. ($p_{\text{partition}}$ denotes the Rent exponent derived from partitioning, and p_i is the Rent exponent derived from a placement using i device layers.)

Since it is our goal to evaluate the Rahman model as an *a priori* estimation tool for circuits for 3-D integration, we will utilize the Rent parameters extracted from partitioning and from 2-D placement.

3) A Note on Hierarchy: It is observed that nearly all state-of-the-art placement tools for standard-cell circuits utilize some form of hierarchical or top-down algorithm. Most commonly, placement is divided into *global* and *detailed* stages. During global placement, the placement problem is partitioned into smaller problems that are manageable by a detailed placement engine. This detailed placement engine, working separately on each of these small problems, determines the eventual final location of all the cells.

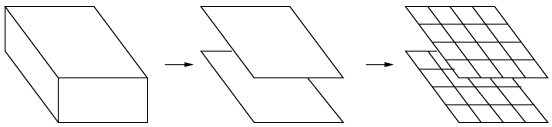

Fig. 3. Partitioning strategy where device-layer assignment is done first in order to minimize the number of interlayer vias.

The Davis model can be extended to reflect this hierarchical placement methodology [11]. In doing so, two separate Rent exponents are used, one for the partitioning process and one for the assignment of cell locations within a partition. This corresponds to the use of separate Rent exponents for global and detailed placement stages. For example, in [11], the situation where detailed placement locations are randomly assigned (corresponding to Rent exponent $p = 1$) is considered. Conversely, if detailed placement locations are assigned using a partitioning algorithm, it is expected that the two Rent exponents will be equal and that the hierarchical model will reduce to the Davis model. Since this is the case with our placement tool (as will be described in the following section), the use of the Davis and Rahman methodologies is suitable.

Using the dimensional values and Rent parameters obtained as described above, we are able to predict the wirelength distribution $f_{3\text{D}}(l)$ of a given circuit. In the next section, we discuss briefly the tools we have developed to determine the actual placed and routed wirelengths of the same circuits.

III. Overview of the 3-D Place-and-Route Tools

A. 3-D Standard-Cell Placement

We have developed a standard-cell placement tool for 3-D integrated circuits [12]. The placement is done at both global and detailed stages by recursive min-cut hypergraph bipartition based on [13], with fine tuning of the final cell locations done optimally as in [14]. That is, the circuit is represented as a hypergraph to be embedded into the die, where each cell corresponds to a node and each net corresponds to a hyperedge of the graph. This hypergraph is then partitioned repeatedly such that the number of nets cut by each partitioning step is minimized. Each partitioning step localizes the cells of the circuit to smaller and smaller subregions of the overall die area.

To extend this algorithm to three dimensions, we take the sequence of partitioning steps for 2-D placement and interleave additional steps that localize cells to specific device layers. The ordering of these steps with respect to the partitioning steps for 2-D placement determines the overall placement quality. If, for example, the partitioning into device layers is done first, this yields a placement that minimizes the number of interlayer vias (Fig. 3). Alternatively, partitioning of a die subregion into device layers can be done considering the aspect ratio of the subregion—that is, the partitioning is performed when the z axis

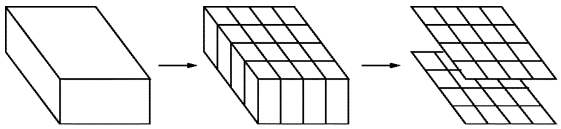

Fig. 4. Partitioning strategy where device-layer assignment is done by considering aspect ratio in order to minimize total wire length.

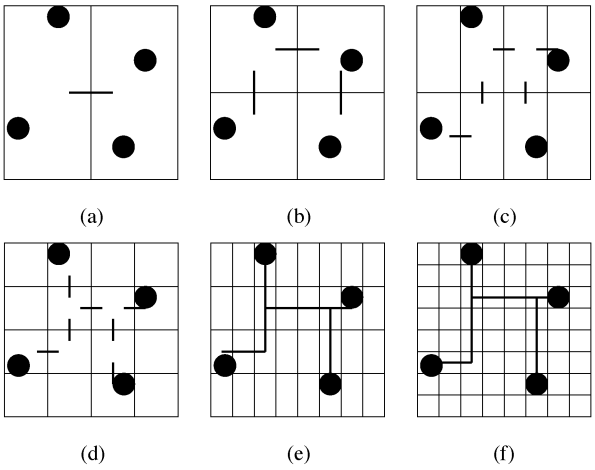

Fig. 5. Single-net example of the hierarchical routing procedure. Routing proceeds from stage (a)–(f) by recursive partitioning.

thickness of the subregion exceeds the height or width (Fig. 4). In this case, the resulting wire length is minimized. Wirelengths are determined using the half-perimeter metric, which in the 3-D case is the sum of the length, width, and height of the bounding box containing all terminals of a given net.

At the interface, the user may specify the number of device layers, as well as whether it is desired to minimize the overall wire length or the number of interlayer vias. The user must specify the size of the interlayer vias in the same units used to provide the die dimensions to the tool; this is akin to specifying the interlayer pitch t_z to the predictive model.

B. 3-D Global Routing

Our global router is an over-the-cell router based on the hierarchical approach in [15]. The routing substrate, consisting of the wiring surface above the placed cells as well as any area unoccupied by cells (for placement of interlayer vias), is recursively partitioned into routing subregions. Wires are either wholly contained within a subregion or are cut by the partition; the cut points become terminals for detailed routing by a switchbox router. For routing of interlayer wires where the wire must go through the upper device layer, area is allocated during placement, either between cell rows (i.e., in the fixed-die specification) or within rows (i.e., as 3-D feed-throughs in the extra space in the row). A sample routing for a single net is shown (from a 2-D perspective for clarity) in Fig. 5.

The results of global routing are computed as the sum over all routing regions of the half-perimeter wirelengths of the wires contained within each region. This measurement should more closely reflect the final aggregate wire length.

Our standard-cell tools are implemented in about 17 000 lines of C code. For partitioning during placement, the user may select between `hMetis` [16] and `PaToH` [17], two high-quality hypergraph partitioners. Circuit netlists and die specifications may be given in either GSRC or Cadence LEF/DEF formats.[1]

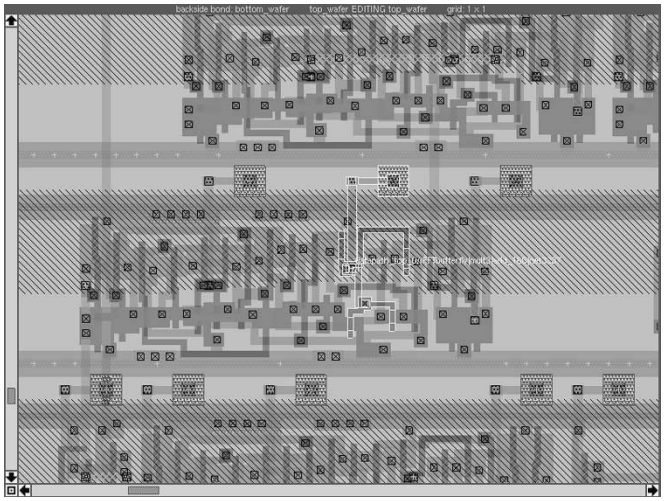

Fig. 6. Top layer of a two-device-layer placement of a benchmark circuit.

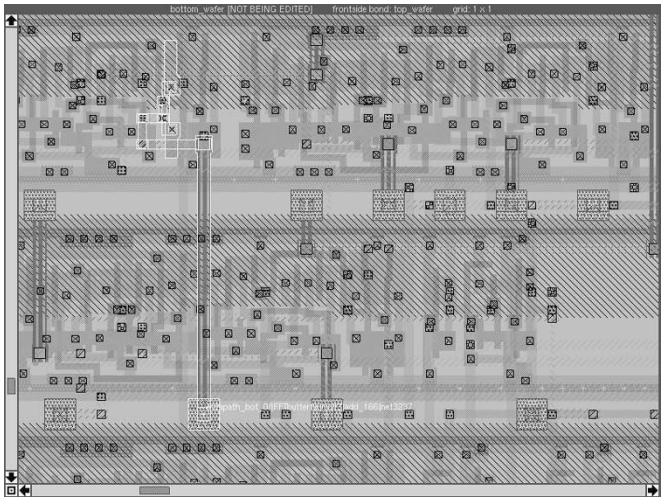

Fig. 7. Bottom layer of a two-device-layer placement of a benchmark circuit.

In Figs. 6 and 7, a two-device-layer placement of a benchmark circuit is shown.

IV. Performance of the Model

In order to evaluate the predictions of the Rahman model, we first calibrate our placement tool against leading-edge tools for 2-D placement. Table I provides wirelengths obtained for circuits from the IBM-PLACE 2.0 benchmark set, for which wirelengths obtained by placers Dragon [19], Capo MetaPlacer [20], and Cadence QPlace were obtained in [21]. We observe that our placer is competitive with state-of-the-art tools for 2-D placement.

We have used our tools to place and route circuits from the ISPD '98 benchmark suite [] using one through five device layers for each circuit. We choose this benchmark suite since it contains several large circuits; for purposes of comparison between the predictions of the Rahman model and placement

[1]Available: http://vlsicad.cs.ucla.edu/GSRC/bookshelf/Slots/Placement/plFormats.html

TABLE I
PERFORMANCE OF OUR PLACER AND OTHER STATE-OF-THE-ART PLACERS ON THE IBM-PLACE 2.0 CIRCUIT BENCHMARK SET. WIRE LENGTHS ARE IN METERS

	QPlace	Dragon	Capo	our placer
ibm01-easy	0.59	0.58	0.56	0.56
ibm01-hard	0.59	0.56	0.56	0.55
ibm02-easy	1.59	1.54	1.55	1.56
ibm02-hard	1.57	1.44	1.52	1.59
ibm07-easy	3.79	3.55	3.73	3.63
ibm07-hard	3.66	3.32	3.60	3.59
ibm08-easy	3.97	3.66	3.94	3.96
ibm08-hard	3.78	3.41	3.77	3.83
ibm09-easy	3.45	3.10	3.18	3.20
ibm09-hard	3.25	3.07	3.23	3.16
ibm10-easy	6.47	6.00	6.26	6.16
ibm10-hard	6.28	5.97	6.35	6.12
ibm11-easy	5.15	4.78	4.99	4.96
ibm11-hard	4.97	4.55	4.99	4.81
ibm12-easy	9.31	8.54	8.65	8.85
ibm12-hard	8.53	8.46	8.35	8.30
dev. from avg.	+3.2%	-3.5%	+0.4%	-0.1%

TABLE II
CELLS OF THE ISPD '98 BENCHMARK SUITE USED IN THIS STUDY

circuit	# cells	# nets	Rent k	Rent p (1)	Rent p (2)	Rent p (3)
ibm11	68119	67016	3.48	0.662	0.753	0.692
ibm12	69026	67739	4.26	0.685	0.755	0.715
ibm13	81018	83806	3.67	0.677	0.764	0.665
ibm14	145492	143202	3.51	0.689	0.787	0.719
ibm15	157861	161196	3.99	0.669	0.766	0.667
ibm16	181633	181188	4.16	0.675	0.765	0.705
ibm17	182359	180684	4.55	0.694	0.759	0.725
ibm18	210051	200565	3.89	0.671	0.741	0.707

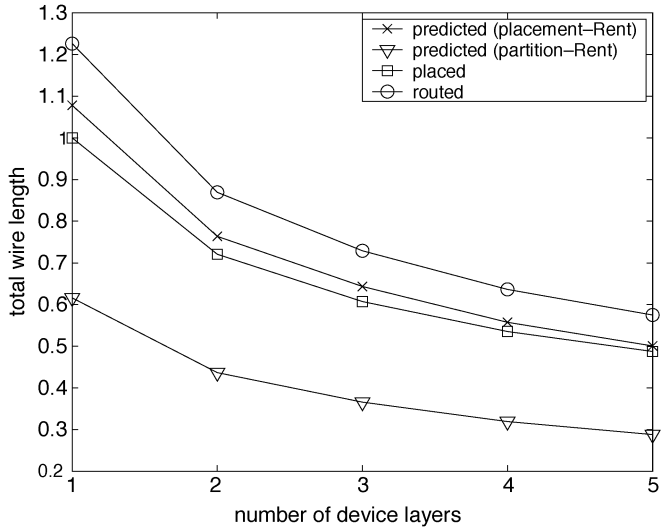

Fig. 8. Predicted versus placed and routed wirelengths of the average benchmark circuit. Wire length is given relative to the 2-D placed wire length. Interlayer pitch t_z is 1 micron.

outcomes, we use the eight largest circuits from this benchmark. Some of the figures of merit for these circuits are given in Table II. In particular, we provide the Rent exponent as calculated from partitioning and from 2-D placements using both methods described in Section II. The value (1) is the Rent exponent from partitioning. Value (2) is the one used in the model. Value (3) is computed using the method in [10].

In Figs. 8 and 9, we compare the total wire length of these circuits as predicted by the Rahman model versus as obtained by

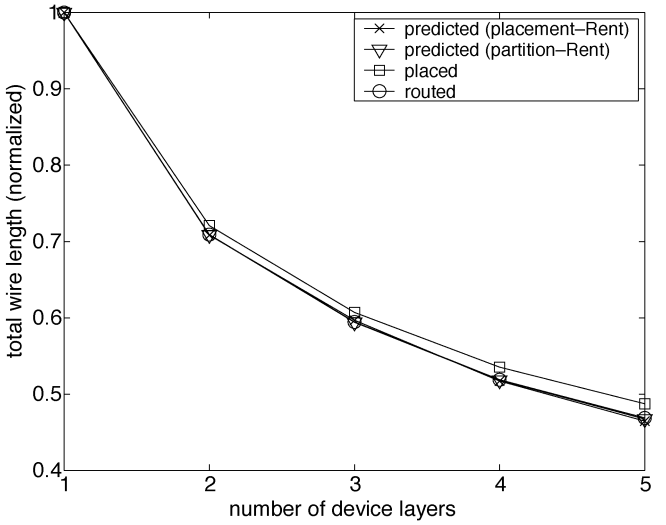

Fig. 9. Predicted versus placed and routed wirelengths of the average benchmark circuit. Wire length is normalized to exhibit the percentage reduction due to 3-D integration. Interlayer pitch t_z is 1 micron.

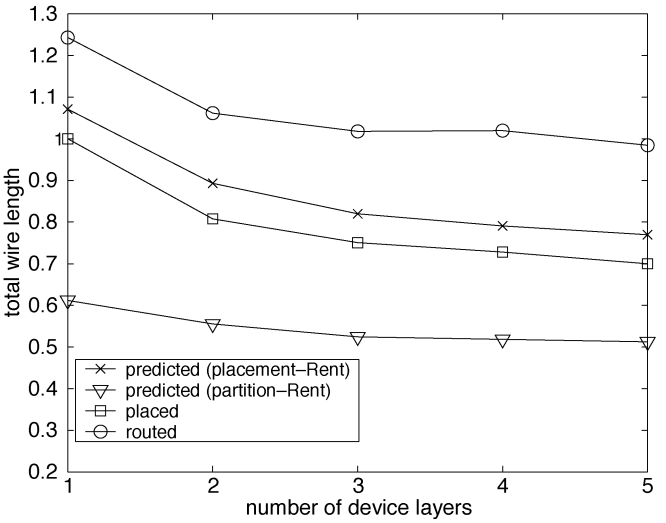

Fig. 10. Predicted versus placed and routed wirelengths of the average benchmark circuit. Wire length is given relative to the 2-D placed wire length. Interlayer pitch t_z is 250 microns.

placement and routing, assuming an interlayer pitch of 1 (i.e., assuming that an interlayer via is equivalent to 1 micron of metal wire). In Fig. 8, the wirelengths are normalized to the 2-D placement case and averaged over all eight circuits. Fig. 9 shows the same data where all wirelength curves are normalized to their 2-D cases, so that one may see how the model predicts the percent reduction in total circuit wire length as a function of number of device layers.

Similarly, Figs. 10 and 11 compare the prediction of the Rahman model to the placement and routing outcomes for the same circuits, but where an interlayer pitch of 250 is used (i.e., where an interlayer via is equivalent to 250 microns of metal wire).

Finally, we may verify some of the underlying assumptions of the model by comparing the predicted wire length distribution with the distribution obtained from placement. As an example, Fig. 12 gives the predicted wire length distribution (number of

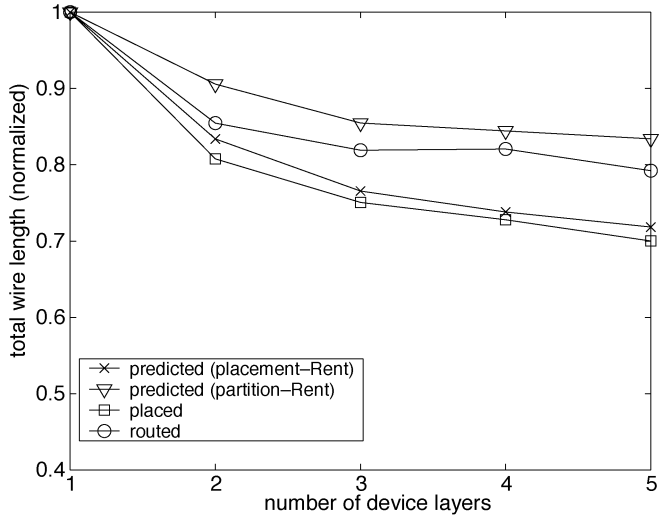

Fig. 11. Predicted versus placed and routed wirelengths of the average benchmark circuit. Wire length is normalized to exhibit the percentage reduction due to 3-D integration. Interlayer pitch t_z is 250 microns.

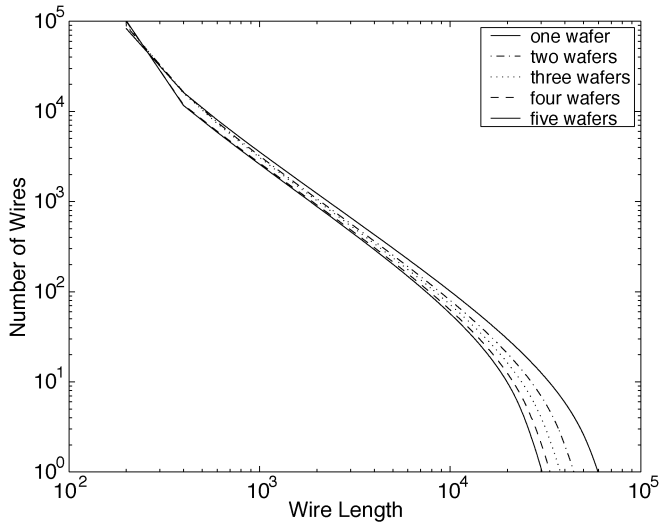

Fig. 12. Predicted wire length distribution for the ibm14 benchmark circuit with interlayer pitch t_z of 1 micron.

wires in the circuit as a function of wire length) for the ibm14 circuit. Fig. 13 provides the distribution from placement for the same circuit.

V. Discussion

It can be seen that the Rahman model is fairly accurate in predicting how 3-D integration affects the wirelengths of circuits. We see in Tables III and IV that the wire length predictions using the 2-D placement Rent exponent are within about 20% of the wirelengths obtained from placement and of those obtained from global routing. As expected, we find that the Rent exponents from 2-D placement more accurately reflect the character of the 3-D placements than do the Rent exponents from partitioning, because of the use of terminal propagation in both 2-D and 3-D placement algorithms. Both the model and the placement and routing data thus show that 3-D integration provides useful benefits for digital circuits: a reduction in total wire length of up to 28% using two wafers to 51% using five wafers is possible.

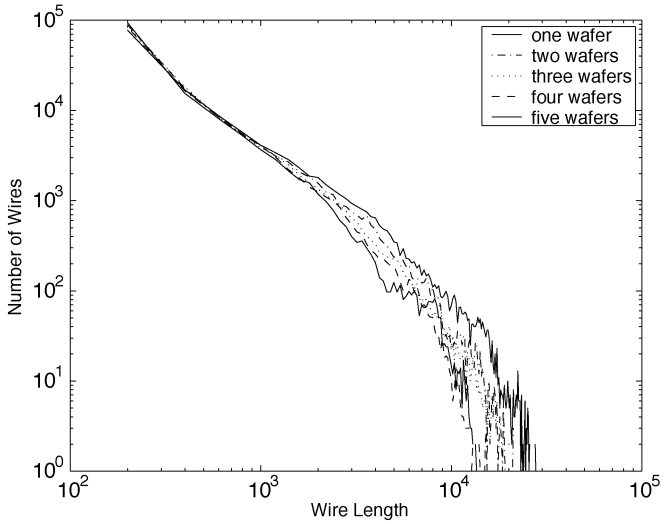

Fig. 13. Placed wire length distribution for the ibm14 benchmark circuit with interlayer pitch t_z of 1 micron.

TABLE III
Absolute Prediction Error Relative to Placed Wire Length as a Function of Number of Device Layers and Inter-Layer Thickness

	one	two	three	four	five
$t_z = 1$	21.2%	19.1%	21.3%	20.0%	18.6%
$t_z = 250$	21.2%	20.9%	20.1%	19.1%	20.1%

TABLE IV
Absolute Prediction Error Relative to Routed Wire Length as a Function of Number of Device Layers and Inter-Layer Thickness

	one	two	three	four	five
$t_z = 1$	17.4%	17.0%	17.8%	18.1%	17.9%
$t_z = 250$	18.2%	18.1%	20.1%	22.5%	21.8%

In Figs. 12 and 13, we observe that the model predicts the shape of the wirelength distribution fairly well. In addition, we see that increasing the number of device layers shifts the distribution leftward, yielding more local wires and fewer global and semiglobal wires. However, we also observe that the model tends to underestimate the number of medium-length wires while overestimating the number of global wires. It is believed that this discrepancy arises from the assumption within the model that the gates are laid out in a square array, while in actual placements, the aspect ratio may deviate from unity. Additionally, the use of a constant fan-out, independent of wire length, may affect the distribution.

Another area for improvement of the 3-D aspects of the model is in the model's prediction of percent reduction in wire length (Figs. 9 and 11), which appears to be slightly inaccurate. It is not clear, however, whether the error lies in strictly 2-D aspects of the model, or in the extension to 3-D, or both.

Therefore, we examine the percentage of interconnects that span multiple device layers. As we see in Fig. 14, placements with a high interlayer pitch use less of the available interlayer bandwidth. However, within the Rahman model, the division of interconnects into 2-D $(\beta_0 M_{2D}(l))$ and 3-D $(\sum_{i=1}^{N_z-1} \beta_i M_{2D}(l-it_z)u(l-it_z))$ components is less strongly dependent on the number of wafers, as can be seen in the figure. This error thus possibly may be explained by the computation of $M_{3D}(l)$.

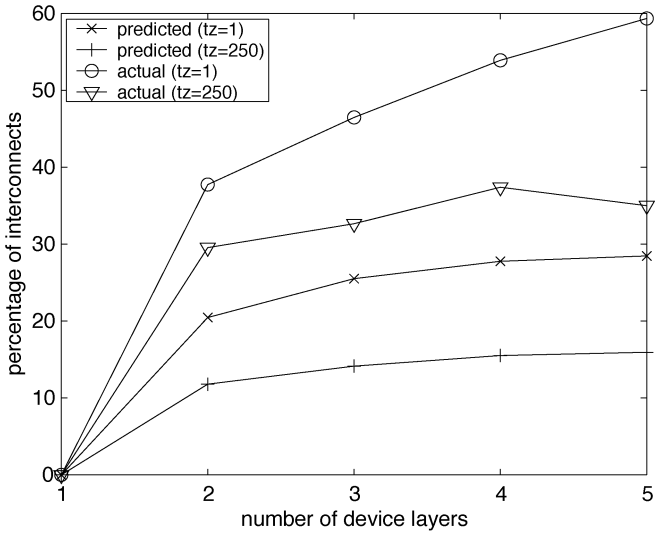

Fig. 14. Predicted percentage of interconnects that span multiple device layers, compared with placement and routing data for $t_z = 1$ and $t_z = 250$.

We conjecture that β_i should be a function of t_z. In practice, the coefficients β_i are set discretely, based on the range of interlayer interconnects. Specifically, for $0 < i \leq r_{\max}$, $\beta_i = 2(N_z - i)$, where a third-dimension wire is permitted to span at most $r_{\max}$ device layers; for all other $i > 0$, $\beta_i = 0$. (We have assumed that $r_{\max} = N_z - 1$, i.e., that an interlayer interconnect may range over all device layers.) β_i is thus independent of t_z. By contrast, our placement tool adjusts the point of partitioning into device layers (and thus the density of interlayer interconnects) based on t_z, in order to minimize wire length. It is not clear, however, how one may determine *a priori* the dependence of β_i on t_z.

Nevertheless, as shown in Figs. 8–11, the Rahman model is a valid extension of the Davis model to 3-D integrated circuits. It proves useful for determining system-level performance characteristics of circuits that are targeted for 3-D integration.

VI. Conclusion

We have developed standard-cell placement and global-routing tools for 3-D integration. These tools are competitive with state-of-the-art tools for conventional 2-D circuits; we use them to obtain placement and routing figures for circuits from the ISPD '98 benchmark suite. Using this data, we have found that A. Rahman's interconnect prediction model for 3-D integrated circuits produces wirelength results that are accurate to within about 20% of placement and of routing, on average. Both the model and the corroborating placement and routing data thus indicate that 3-D integration presents significant benefits to designers of digital circuits.

Acknowledgment

The authors would like to thank A. Rahman for helpful discussions and clarifications of the interconnect model used in this paper. The authors also thank the comments from the reviewers of this manuscript which were extensively useful and are appreciated.

References

[1] A. Rahman and R. Reif, "System-level performance evaluation of three-dimensional integrated circuits," *IEEE Trans. VLSI Syst. Special Issue System-Level Interconnect Prediction*, vol. 8, pp. 671–678, Dec. 2000.
[2] M. T. Bohr, "Interconnect scaling—The real limiter to high-performance ULSI," in *Proc. Int. Electron Devices Meeting (IEDM)*, 1995, pp. 241–244.
[3] (2003) International Technology Roadmap for Semiconductors. [Online] Avaliable: http://public.itrs.net/
[4] C. W. Eichelberger, "Three-dimensional multichip module system," U.S. Patent 5 111 278, May 1992.
[5] G. Roos, B. Hoefflinger, M. Schubert, and R. Zingg, "Manufacturability of 3D-epitaxial-lateral-overgrowth CMOS circuits with three stacked channels," *Microelectron. Eng.*, vol. 15, pp. 191–194, 1991.
[6] V. Subramanian, P. Dankoski, L. Degertekin, B. T. Khuri-Yakub, and K. C. Saraswat, "Controlled two-step solid-phase crystallization for high-performance polysilicon TFTs," *IEEE Electron Device Lett.*, vol. 18, pp. 378–381, Aug. 1997.
[7] A. Fan, A. Rahman, and R. Reif, "Copper wafer bonding," *Electrochem. Solid-State Lett.*, vol. 2, pp. 534–536, 1999.
[8] J. Davis, V. K. De, and J. D. Meindl, "A stochastic wire-length distribution for gigascale integration (GSI)—part I: derivation and validation," *IEEE Trans. Electron Devices*, vol. 45, pp. 580–589, Mar. 1998.
[9] B. S. Landman and R. L. Russo, "On a pin versus block relationship for partitions of logic graphs," *IEEE Trans. Computers*, vol. C-20, pp. 1469–1479, Dec. 1971.
[10] X. Yang, E. Bozorgzadeh, and M. Sarrafzadeh, "Wirelength estimation based on rent exponents of partitioning and placement," in *Proc. ACM/IEEE Int. Conf. SLIP*, 2001, pp. 25–32.
[11] P. Christie and D. Stroobandt, "The interpretation and application of Rent's Rule," *IEEE Trans. VLSI Syst.*, vol. 8, pp. 639–648, Dec. 2000.
[12] S. Das, A. Chandrakasan, and R. Reif, "Design tools for 3-D integrated circuits," in *Proc. ASP-DAC*, 2003, pp. 53–56.
[13] A. E. Dunlop and B. W. Kernighan, "A procedure for placement of standard cell VLSI circuits," *IEEE Trans. Computer-Aided Design*, vol. 34, pp. 92–98, Jan., 1985.
[14] A. B. Kahng, A. E. Caldwell, and I. L. Markov, "Optimal partitioners and end-case placers for standard-cell layout," in *Proc. ACM Int. Symp. Physical Design*, Apr. 1999, pp. 90–96.
[15] M. Burstein and R. Pelavin, "Hierarchical wire routing," *IEEE Trans. Computer-Aided Design*, vol. CAD-2, pp. 223–234, Apr. 1983.
[16] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar, "Multilevel hypergraph partitioning: applications in VLSI design," in *Proc. 34th Design Automation Conf.*, 1997, pp. 526–529.
[17] U. V. Catalyurek and C. Aykanat, "Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 10, pp. 673–693, July 1999.
[18] M. Wang, X. Yang, and M. Sarrafzadeh, "Dragon2000: fast standard-cell placement for large circuits," in *IEEE Int. Conf. Computer-Aided Design*, 2000, pp. 260–263.
[19] A. E. Caldwell, A. B. Kahng, and I. L. Markov, "Can recursive bisection alone produce routable placements?," in *Proc. 37th Design Automation Conf.*, 2000, pp. 477–482.
[20] X. Yang, B.-K. Choi, and M. Sarrafzadeh, "Routability driven white space allocation for fixed-die standard-cell placement," in *ACM Int. Symp. Physical Design*, Apr. 2002, pp. 42–47.
[21] C. J. Alpert, "The ISPD98 circuit benchmark suite," in *ACM Int. Symp. Physical Design*, Apr. 1998, pp. 80–85.

Shamik Das (M'03) received the S.B. degree in electrical engineering and mathematics and the M.Eng. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, in 2000. He is currently pursuing the Ph.D. degree in electrical engineering at the same university, where he is working on digital-circuit aspects of and computer-aided design tools for three–dimensional integrated circuits.

Mr. Das is a Member of Phi Beta Kappa.

Anantha P. Chandrakasan (S'87—M'95–SM'01–F'00) received the B.S., M.S., and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively.

Since September 1994, he has been with the Massachusetts Institute of Technology (MIT), Cambridge, where he is currently a Professor of electrical engineering and computer science. His research interests include the low-power digital integrated circuit design, wireless microsensors, ultrawideband radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Norwell, MA: Kluwer, 1995) and *Digital Integrated Circuits (2nd edition)* (Englewood Cliffs, NJ: Prentice-Hall, 2003). He is also a coeditor of *Low Power CMOS Design* (Piscataway, NJ: IEEE Press, 1998) and *Design of High-Performance Microprocessor Circuits* (Piscataway, NJ: IEEE Press, 2000).

Dr. Chandrakasan has received several awards including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, and the 1999 Design Automation Conference Design Contest Award. He served as a Technical Program Cochair for the 1997 International Symposium on Low-Power Electronics and Design (ISLPED), Very Large Scale Integration (VLSI) Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Subcommittee Chair for the International Solid-State Circuits Conference (ISSCC), from 1999 to 2001, the Program Vice Chair for ISSCC 2002, the Program Chair for ISSCC 2003, and the Technology Directions subcommittee Chair for ISSCC 2004. He was an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS from 1998 to 2001. He serves on the IEEE Solid-State Circuits Society (SSCS) AdCom. He is the Technology Directions Chair for ISSCC 2005.

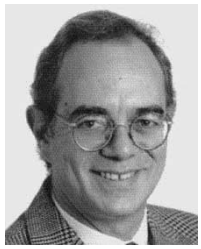

Rafael Reif (M'79–SM'90–F'93) received the Ingeniero Electrico degree from Universidad de Carabobo, Valencia, Venezuela, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1973, 1975, and 1979, respectively.

From 1973 to 1974, he was an Assistant Professor at Universidad Simon Bolivar, Caracas, Venezuela. In 1978, he became a Visiting Assistant Professor in the Department of Electrical Engineering, Stanford University. In 1980, he joined the Massachusetts Institute of Technology, Cambridge, where he is currently a Professor and the Associate Department Head for Electrical Engineering in the Department of Electrical Engineering and Computer Science. He was the Director of MIT's Microsystems Technology Laboratories (MTL) from 1990 to 1999. He is presently working on future microelectronics interconnect technologies, and on environmentally benign replacement chemistries for microelectronics fabrication. He held the Analog Devices Career Development Professorship in the Department of Electrical Engineering and Computer Science at MIT, and was awarded the IBM Faculty Fellowship of MIT's Center for Materials Science and Engineering from 1980 to 1982.

Dr. Reif received a United States Presidential Young Investigator Award in 1984. He is a recipient of the Semiconductor Research Corporation's 2000 Aristotle Award. He is a Member of Tau Beta Pi, the Electrochemical Society, and the American Physical Society.