

Global Interconnect Design in a Three-Dimensional System-on-a-Chip

James W. Joyner, Payman Zarkesh-Ha, *Member, IEEE*, and James D. Meindl

Abstract—A stochastic model for the global net-length distribution of a three-dimensional system-on-a-chip (3D-SoC) is derived. Using the results of this model, a global interconnect design window for a 3D-SoC is established by evaluating the constraints of: 1) wiring area; 2) clock wiring bandwidth; and 3) crosstalk noise. This window elucidates the optimum 3D-SoC global interconnect parameters for minimum pitch, minimum aspect ratio, and maximum clock frequency. In comparison to a two-dimensional system-on-a-chip (2D-SoC), the design window expands for a 3D-SoC to allow greater flexibility of interconnect parameters, thus increasing the guardbands to process variations. In addition, the limit on the maximum global clock frequency is revealed to increase as S^2 , where S is the number of strata. This increase in on-chip signaling rate, however, comes at the expense of I/O density, highlighting the need for new high-density-I/O packaging techniques to exploit the full potential of 3D-SoC.

Index Terms—Global wiring, integrated circuit interconnection, net length, system-level interconnect prediction, system-on-a-chip (SoC), three-dimensional (3-D) integration, wirelength distribution.

I. INTRODUCTION

WIRING requirements, particularly those of global interconnects, are fast becoming a bottleneck to the performance of future gigascale integrated (GSI) systems [1], [2]. The advent of three-dimensional (3-D) integration in which interstratal interconnects link multiple strata of transistors and wiring as illustrated in Fig. 1 is commonly regarded as a promising solution in satisfying these ever-growing wiring demands [3]. A stratum is defined as a single layer of transistors with its corresponding metal levels [3] while an interstratal interconnect is one that connects gates in different strata.

Previous work [4] has demonstrated that the use of 3-D integration for homogeneous systems of identical microcells reduces the length of the longest interconnects significantly while providing a lesser advantage for average interconnects. In addition, the density of interstratal interconnects is restricted by the alignment tolerance of wafer-bonding techniques, posing a strict limit on the level of vertical integration [4]. To exploit the tremendous opportunities of 3-D integration for long interconnects while steering clear of the interstratal interconnect density limitations, a system-on-a-chip (SoC) consisting of previously




Fig. 1. Cross-section of a three-dimensional integrated circuit showing the interstratal interconnects connecting two strata.




Fig. 2. An example layout of a 3D-SoC from an existing 2D-SoC, where each megacell is contained in a single stratum.

designed and optimized heterogeneous megacells connected by relatively few, long nets can be implemented in three dimensions. The resulting 3-D layout of heterogeneous megacells effects a reduction in the net lengths between megacells, which directly benefits the global clock frequency.

In Section II, a stochastic net-length distribution for a 3D-SoC is developed as an extension of previous work for a 2D-SoC [5]. Then, an integrated architecture for global interconnects in a 2D-SoC [6] is reviewed and extended to quantify the advantages of a 3D-SoC in Section III. In addition, the limitations of 3-D heterogeneous systems as well as key recommendations to maximize the potential of 3D-SoC are discussed. Finally, conclusions are presented in Section IV.

II. A STOCHASTIC GLOBAL NET-LENGTH DISTRIBUTION FOR 3D-SoC

A global net-length distribution model describes the net density function (n_{ndf}), the number of nets connecting heterogeneous megacells in an SoC per unit length, with respect to the minimum rectilinear Steiner tree (MRST) length of the net [5]. To determine the distribution stochastically, three models representing the netlist, placement, and routing information are required. In transitioning to a 3D-SoC, each megacell is confined to a single stratum such that the area and performance of each megacell remain constant as the number of strata increases. An example of shifting to a two-stratal layout is visualized in Fig. 2.

Manuscript received August 26, 2002; revised July 24, 2003. This work was supported in part by Semiconductor Research Corporation (SRC), by SRC-Center for Advanced Interconnect Systems Technologies (CAIST), and by Defense Advanced Research Projects Agency (DARPA).

J. W. Joyner and J. D. Meindl are with the Georgia Institute of Technology, Atlanta, GA 30332-0269 USA (e-mail: joyner@ece.gatech.edu).

P. Zarkesh-Ha is with the LSI Logic Corporation, Milpitas, CA 95035
Digital Object Identifier 10.1109/TVLSI.2004.825835

In a general case, the total area of the megacells in the upper stratum may not be equal to the total area of those in the lower stratum. Thus, total area may not be conserved in changing the number of strata utilized. In addition, through-wafer vias must be routed through the substrate to connect to other strata. If the megacells are independently designed, these through-wafer vias may be constrained to the area between megacells, creating another source of “white space.” To take these effects into account, a stratal area inefficiency factor $\eta_s(S) \geq 1$ is used.

The netlist information provides a distribution of the nets connecting megacells in an SoC according to the fanout of the nets. The fanout distribution is based on a relationship known as Rent’s Rule [7]. Using this relationship, the fanout distribution $N[f.o.]$ [8] is calculated as

$$N[f.o.] = \frac{k_{eq}n \left(f.o.^{p_{eq}-1} - (f.o. + 1)^{p_{eq}-1} \right)}{f.o. + 1} \quad (1)$$

where $f.o.$ is the fanout of the net, k_{eq} and p_{eq} are the equivalent Rent’s coefficient and exponent [9], respectively, and n is the total number of megacells. The netlist information is independent of the number of strata utilized since the netlist of connections between megacells is independent of the relative placement of the cells involved.

The placement information describes the average dimensions of the bounding area of a net connecting a group of megacells. For a 3D-SoC, a group of m megacells ($m \leq n$) is bounded, on average, by a rectangle with an edge of length

$$e'_{3D}[m] = \sqrt{\frac{\eta_s(S) A_{meg}}{S} [m\eta_p + n(1 - \eta_p)]} \quad (2)$$

where A_{meg} is the average megacell area and the placement efficiency (η_p) [5] is given by

$$\eta_p = \frac{1 - \frac{\text{Block_bounding_area}}{\text{Stratal_area}}}{1 - \frac{m}{n}} \quad (3)$$

The block-bounding area of m megacells is the minimum rectangular area encompassing all m megacells. Likewise, the net-bounding area of an m -terminal net is the minimum rectangular area covering all m terminals. The stratal area inefficiency factor $\eta_s(S)$, as described earlier, captures the effective increase in average megacell size as compared to that in a 2D-SoC ($\eta_s(1) = 1$). According to the relationship of block-bounding edge-length to net-bounding edge length derived in [5], the edge length of the net-bounding area can be written as

$$e_{3D}[m] = \frac{m-1}{m+1} e'_{3D}[m]. \quad (4)$$

The routing information provides a distribution of the length of a net for a given number of terminals and net-bounding area. A minimum rectilinear Steiner tree (MRST) model calculates the minimum total length of a tree connecting the terminals using manhattan geometry. The routing information is determined through a random walk method using [10] for net length. For a given number of terminals m , this simulation maps a net-bounding edge length $e_{3D}[m]$ to a probability density of net length. Since the layout of the MRST is independent of the actual dimensions of the net-bounding area, the net length

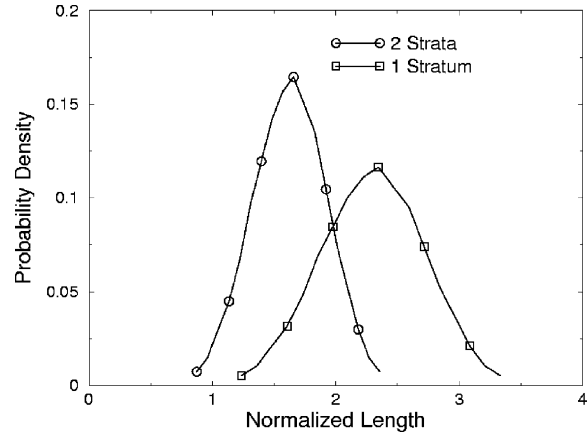


Fig. 3. The probability density of length for a five-terminal net in a 2D-SoC (1 stratum) and a 3D-SoC (2 strata). The length is normalized to the net-bounding edge length of the net in a 2D-SoC.

scales linearly with $e_{3D}[m]$. An m -terminal net-length probability density function $f_{m,3D}(l_{norm})$ with the length normalized to $e_{3D}[m]$ is determined through simulation. Assuming the interstratal segments of a 3-D net are negligible in length in comparison to the total net length, the density function scales as the net-bounding area with respect to the number of strata. The probability density function $f_{m,3D}(l_{norm})$ can then be expressed as

$$f_{m,3D}(l_{norm}) \approx \sqrt{S} \cdot f_m(\sqrt{S} \cdot l_{norm}) \quad (5)$$

where $f_m(l_{norm})$ is the corresponding probability density function for a global net in a 2D-SoC. Illustrating this relationship, Fig. 3 plots the probability density of length normalized to net-bounding edge length in a 2D-SoC for a five-terminal net in both a 2D-SoC and a 3D-SoC.

The complete global net-length distribution for 3D-SoC is derived from the three models of netlist (1), placement (2)–(4), and routing information (5). In summary, the netlist information provides the number of nets $N[f.o. = m - 1]$ for a given fanout $f.o.$, the placement information predicts the edge length $e_{3D}[m]$ of the net-bounding area for an m -terminal net that connects m megacells, and the routing information is the probability density function $f_{m,3D}(l_{norm})$ for an m -terminal net. Combining these three distinct models, the global net-length distribution for a 3D-SoC, the product of the number of nets and the length probability density function of these nets, is written as

$$n_{ndf,3D}(l) = \sum_{m=2}^n N[m-1] \cdot f_{m,3D}\left(\frac{l}{e_{3D}[m]}\right). \quad (6)$$

The global net-length distribution is compared with data from a RISC microprocessor described in [11], [5] for the two-dimensional (2-D) case. The microprocessor consists of 20 megacells ranging in size from 9500 to 380 000 gates with a total chip area of 16.6 mm \times 17.8 mm. The placement efficiency is assumed as $\eta_p = 0.80$ according to the typical value found from experimentation. The stratal area inefficiency is assumed as $\eta_s(S) = 1$ to determine the theoretical limits of 3D-SoC technology. This assumption is valid for a system with a large number of smaller megacells. If this value were to increase above unity, it would

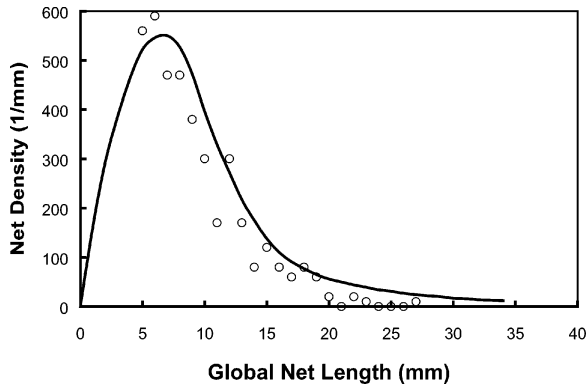


Fig. 4. Global net-length distribution for a 2D-SoC provides the number of nets per unit length versus net length.

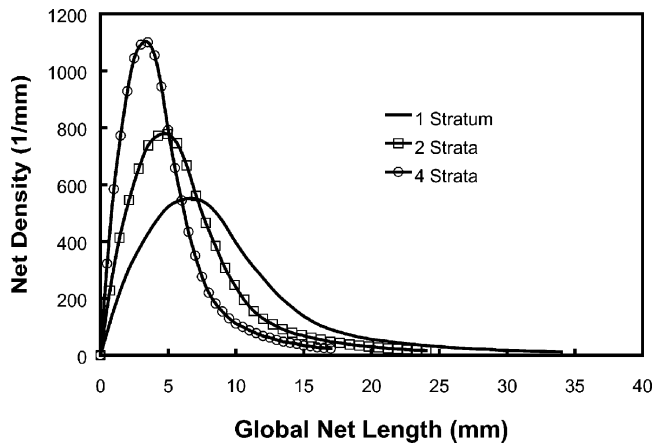


Fig. 5. The global net-length distributions for 1, 2, and 4 strata illustrate that the global net-length reduces as the square root of the number of strata.

have the impact of increasing the net length. The net-length distribution scales as the square root of the stratal area inefficiency. The 2-D global net density as a function of net length is compared to experimental data in Fig. 4 [5]. Further implications of the results for 2D-SoCs are discussed in [5].

Defining $n_{\text{ndf}}(l)$ as the net density function for a 2D-SoC, the effects of scaling the number of strata can be seen as

$$n_{\text{ndf},3\text{D}}(l) \approx \sqrt{S} \cdot n_{\text{ndf}}(\sqrt{S} \cdot l). \quad (7)$$

Using the same system parameters as before, a comparison of the global net density functions for a 2D-SoC and 3D-SoC of 2 and 4 strata is provided in Fig. 5. The resulting 3D-SoC distributions highlight the reduction in global net length by the square root of the numbers of strata.

III. A GLOBAL INTERCONNECT DESIGN WINDOW FOR 3D-SoC

In the design of a global interconnect architecture that consists of the signal, clock, and power-supply networks, an optimization of the interconnect cross-sectional dimensions and spacings can be achieved. Fig. 6 illustrates the global wiring with signal interconnect width w , interconnect height h , interconnect spacing or dielectric width s , and dielectric thickness t . In determining the values of w , h , s , and t for an optimal design, three constraints are considered: 1) wiring area; 2) clock wiring bandwidth; and 3) crosstalk noise [6].




Fig. 6. Cross-section of global interconnects.

The area required for global signal routing is expressed as

$$A_{\text{Signal}} = (w + s) l_{\text{tot}} \quad (8)$$

in which the total length of the global signal wiring l_{tot} is expressed as

$$l_{\text{tot}} = \int_{l=0}^{l_{\text{long}}} l \cdot n_{\text{ndf},3\text{D}}(l) dl. \quad (9)$$

The length of a net l is the variable of integration, l_{long} is the length of the longest net, and $n_{\text{ndf},3\text{D}}(l)$ is the 3D-SoC net-length distribution (7). In the design of an SoC, the area available for global signals is commonly limited by the SoC area (A_{SoC}) such that

$$\frac{A_{\text{Signal}}}{e_w} + A_{\text{Clock}} + A_{\text{Power}} \leq n_{\text{ml}} A_{\text{SoC}} \quad (10)$$

where A_{Clock} and A_{Power} are the wiring areas for the global clock and global power-supply networks, respectively, n_{ml} is the number of metal levels dedicated to global wiring, and e_w is the signal wiring efficiency [6]. Solving (10) for signal area

$$A_{\text{Signal}} \leq e_w (n_{\text{ml}} A_{\text{SoC}} - A_{\text{Clock}} - A_{\text{Power}}). \quad (11)$$

While A_{Clock} is typically negligible in comparison to the chip area [6], A_{Power} is disruptive to the area available for the routing of signal nets. Assuming an area-array placement of power-supply bonding pads and a uniform load distribution, the power-supply network area is estimated [6] as

$$A_{\text{Power}} = \left\{ 2A_{\text{SoC}} \left(\frac{P_{\text{tot}} \rho_w}{16\delta V_{\text{dd}}^2 h n_{\text{pg}}} \right) \times \left(2 - \frac{P_{\text{tot}} \rho_w}{16\delta V_{\text{dd}}^2 h n_{\text{pg}}} \right) \right\} \quad (12)$$

where P_{tot} is the total chip power, ρ_w is the metal resistivity, δ is the ratio of worst-case IR-drop to power supply voltage V_{dd} , and n_{pg} is the number of power and ground pads. Assuming the use of two metal levels for global interconnects and neglecting A_{Clock} , the bound from the wiring area constraint is determined by substituting (8) and (12) into (11)

$$w + s \leq 2e_w \frac{A_{\text{SoC}}}{l_{\text{tot}}} \left(1 - \frac{P_{\text{tot}} \rho_w}{16\delta V_{\text{dd}}^2 h n_{\text{pg}}} \right)^2. \quad (13)$$

The second constraint on global interconnects is imposed by the bandwidth required for global clock distribution. Assuming an RC-limited bandwidth [6], the global clock frequency f_c must satisfy

$$f_c \leq \frac{1}{2\pi r c \left(\frac{1}{2} l_{\text{cc}} \right)^2} \quad (14)$$

where r is the interconnect resistance per unit length, c is the interconnect capacitance per unit length, and l_{cc} is the length of

	180 nm	50 nm
V_{dd} (V)	1.8	0.6
f_c (GHz)	1.2	3
P_{tot} (W)	90	174
A_{SoC} (cm ²)	4.5	8.2
δ (%)	5	5
p_{noise} (%)	25	25
n_{pg}	1536	1536
l_{tot} (m)	32.1	72.3
e_w (%)	25	25

Fig. 7. Technology parameters for the 180- and 50-nm technology generations.

the corner-to-corner interconnect which scales roughly as the inverse square root of the number of strata. This constraint assumes the clock frequency cannot exceed the 3-dB bandwidth of a worst-case H -tree pathlength from the clock I/O to the corner of the chip. Expressing the resistance and capacitance in terms of material and interconnect design parameters using a parallel-plate model

$$f_c \leq \frac{1}{2\pi \left(\rho_w \frac{1}{wh} \right) \left[2\epsilon_r \epsilon_o \left(\frac{h}{s} + \frac{w}{t} \right) \right] \left(\frac{1}{2} l_{cc} \right)^2} \leq \frac{1}{4\pi \rho_w \epsilon_r \epsilon_o \left(\frac{1}{w_s} + \frac{1}{ht} \right) \left(\frac{1}{2} l_{cc} \right)^2} \quad (15)$$

where ϵ_r is the relative permittivity, and ϵ_o is the permittivity of free space. The parasitics of any through-wafer vias and their impact on the bandwidth of the clock network are assumed to be negligible. The bound from the bandwidth constraint is found by rearranging (15) as

$$ws \geq \left(\frac{1}{\pi \rho_w \epsilon_r \epsilon_o l_{cc}^2 f_c} - \frac{1}{ht} \right)^{-1}. \quad (16)$$

The impact of interconnect crosstalk noise provides the third constraint placed on the global interconnect optimization. Using a distributed RLC model, the simplified ratio of worst-case peak crosstalk noise (V_n) to V_{dd} [12] must satisfy

$$\frac{V_n}{V_{dd}} = \frac{\pi}{4} \frac{c_m}{c_m + c_{gnd}} \leq p_{noise} \quad (17)$$

where c_m and c_{gnd} are the mutual and ground capacitances per unit length, respectively, and p_{noise} is the maximum noise allowed as a percentage of V_{dd} . The bound from the noise constraint is determined as

$$ws \geq ht \left(\frac{\pi}{4p_{noise}} - 1 \right). \quad (18)$$

The three constraints of: 1) wiring area (13), 2) clock wiring bandwidth (16), and 3) crosstalk noise (18) establish three bounds in the selection of the global interconnect parameters w , h , s , and t [6]. Using technology parameters outlined by the ITRS for a 2D-SoC at the 180-nm technology generation given in Fig. 7 and assuming that s is equal to w , and t to h , these three constraints are plotted in Fig. 8 with interconnect width




Fig. 8. Global design space for global interconnects bounded by: 1) wiring area; 2) clock bandwidth; and 3) crosstalk noise constraints for the 180-nm technology node.

and height as the x and y axes, respectively.¹ The central triangular region is the design window for which all three constraints are met simultaneously. The three corners of the window correspond to designs for minimum pitch, minimum aspect ratio (h/w), and maximum global clock frequency as indicated in Fig. 8.

Since both the total length of all nets (9) and the corner-to-corner distance (14) scale inversely with the square root of the number of strata used in a 3D-SoC, the wiring area and clock wiring bandwidth constraints on the global interconnect design window can be relaxed through use of more strata. Using the scaled constraints, the global interconnect design windows for systems of 1, 2, and 4 strata at the 180- and 50-nm technology nodes are illustrated in Fig. 9. The window for a single stratum at the 50-nm node is significantly smaller than that at the 180-nm node, illustrating the increasing restrictions placed on global interconnect design with advancing technology [6]. As the global design window shrinks for future technology generations, manufacturing process variations may force the relaxation of projected system parameters at the 50-nm technology node. Transitioning to a 3D-SoC enables the design window to be greatly expanded with increasing strata, thus providing increased flexibility in global interconnect design. This expanded window also increases the width of the guardbands that separate the design points in the center of the window from the design constraints as compared to 2D-SoC. By increasing the width of the guardbands, larger deviations in interconnect dimensions due to process variations can be tolerated before one of the design constraints is violated, thus increasing design margins.

In addition, the maximum global clock frequency corner is impacted by increasing the number of strata used. Increasing the clock frequency shifts the bandwidth constraint curve toward this corner. The maximum achievable clock frequency can be found by simultaneously solving (13) and (18) for interconnect dimensions and then solving (15) for clock frequency using these dimensions. For a large number of power supply pads and large-wiring thicknesses, the percentage area required for power distribution is negligible in comparison to the signal wiring area. Letting the number of power supply pads approach

¹ITRS Semiconductor Industry Association, 1999




Fig. 9. Global design windows for systems of 1, 2, and 4 strata at the (a) 180-nm and (b) 50-nm technology nodes.

infinity and solving (13) and (18) for the interconnect width and height yields

$$\begin{aligned} w &= \frac{e_w A_{\text{SoC}}}{l_{\text{tot}}} \\ h &= \frac{w}{\sqrt{\frac{\pi}{4\rho_{\text{noise}}} - 1}}. \end{aligned} \quad (19)$$

Substituting these values into (15) and simplifying gives the maximum global clock frequency for which all three constraints can be simultaneously met as

$$f_{c,\text{max}} \approx \frac{4A_{\text{SoC}}^2 e_w^2 \rho_{\text{noise}}}{\pi^2 \rho_w \epsilon_r \epsilon_o l_{\text{tot}}^2 l_{\text{cc}}^2} \propto S^2. \quad (20)$$

The maximum global clock frequency, therefore, increases as S^2 through the simultaneous decreases in the signal net length and in the clock-network pathlength. For instance, the use of 2 strata increases the limit on the maximum global clock frequency by $4\times$. This increase in the clock frequency is a limit on performance enhancement through 3-D integration. As an achievable design, however, it is limited by both the manufacturing restrictions for the interconnect dimensions and the skin effect on bandwidth for large interconnects. For more aggressive designs using two strata, however, an increase in frequency close to that projected by this limit is feasible.

Two key obstacles that arise in the implementation of a 3D-SoC are related to its layout. The ability to place the megacells such that total chip area also remains constant, and thus the extent of vertical integration in a heterogeneous system, is limited by the size of the largest megacells. To overcome this obstacle, large megacells, memory in particular, should be divided among multiple strata as performed in [1]. Another implicit penalty incurred in transitioning to a 3D-SoC is the reduction of the surface area available for I/Os and heat removal. Effective techniques are needed to remove heat from the limited surface area of a 3D-SoC as discussed in [13]. In addition, the required I/O density increases as the number of strata. A limitation to the effectiveness of implementing a 3D-SoC is the necessity for high-bandwidth, high-density I/Os. In an effort to alleviate this limitation, exploration of packaging technologies specifically related to high-density I/Os is highly encouraged.

IV. CONCLUSION

A stochastic model for the global net-length distribution of a 3D-SoC is derived. In comparison to a 2D-SoC, the resulting distribution projects that the use of three-dimensional architectures potentially reduces net length as the square root of the number of strata. A global interconnect design window for a 3D-SoC is developed by evaluating the constraints of: 1) wiring area; 2) clock wiring bandwidth; and 3) crosstalk noise. The resulting window provides insight into optimizing the 3D-SoC global interconnect parameters for minimum pitch, minimum aspect ratio, and maximum frequency. In comparison to a 2D-SoC, the global interconnect design window is greatly expanded for a 3D-SoC, increasing the flexibility in interconnect design. The expanded design window has larger guardbands between the center design points and the edge of the window. This allows greater tolerance in deviations of the interconnect dimensions due to process variations before a design constraint is violated, thus improving design margins. In addition, the maximum global clock frequency is shown to increase as S^2 , where S is the number of strata. For example, a 3D-SoC with two strata has a maximum global clock frequency four times that of a 2D-SoC. This increase in on-chip bandwidth, however, derives from a tradeoff of off-chip I/O density, highlighting the need for high-density I/O packaging technologies.

ACKNOWLEDGMENT

The authors deeply appreciate the insightful comments of K. Bowman of Intel Corporation.

REFERENCES

- [1] K. W. Lee *et al.*, "Three-dimensional shared memory fabricated using wafer stacking technology," in *Proc. Int. Electron Devices Meeting (IEDM)*, 2000, pp. 165–168.
- [2] D. Sylvester and K. Keutzer, "A global wiring paradigm for deep submicron design," *IEEE Trans. Computer-Aided Design*, vol. 19, pp. 242–252, Feb. 2000.
- [3] D. A. Antoniadis, A. Wei, and A. Lochtefeld, "SOI devices and technology," in *Proc. 20th European Solid-State Device Research Conf. (ESSDERC)*, 1999, pp. 81–87.
- [4] J. W. Joyner, R. Venkatesan, P. Zarkesh-Ha, J. A. Davis, and J. D. Meindl, "Impact of three-dimensional architectures on interconnects in gigascale integration," *IEEE Trans. VLSI Syst.*, vol. 9, pp. 922–928, Dec. 2001.

- [5] P. Zarkesh-Ha, J. A. Davis, and J. D. Meindl, "Prediction of net length distribution for global interconnects in a heterogeneous system-on-a-chip," *IEEE Trans. VLSI Syst.*, vol. 8, pp. 649–659, Dec. 2000.
- [6] P. Zarkesh-Ha and J. D. Meindl, "An integrated architecture for global interconnects in a gigascale system-on-a-chip (GSoC)," in *Proc. Symp. Very Large Scale Integration Technology*, 2000, pp. 194–195.
- [7] B. S. Landman and R. L. Russo, "On a pin versus block relationship for partitions of logic graphs," *IEEE Trans. Computer*, vol. C-20, pp. 1469–1479, Dec. 1971.
- [8] P. Zarkesh-Ha, J. A. Davis, W. Loh, and J. D. Meindl, "Prediction of interconnect fan-out distribution using Rent's Rule," in *Proc. Int. Workshop System-Level Interconnect Prediction (SLIP)*, 2000, pp. 107–112.
- [9] —, "On a pin versus gate relationship for heterogeneous systems: Heterogeneous Rent's Rule," in *Proc. Custom Integrated Circuits Conf. (CICC)*, 1998, pp. 93–96.
- [10] A. B. Kahng and G. Robins, "A new class of iterative Steiner tree heuristics with good performance," *IEEE Trans. Computer-Aided Design*, vol. 11, pp. 893–902, July 1992.
- [11] N. Vasseghi *et al.*, "200-MHz superscalar RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1675–1686, Nov. 1996.
- [12] J. A. Davis and J. D. Meindl, "Compact distributed RLC interconnect models: Parts I and II," *IEEE Trans. Electron Device*, vol. 47, pp. 2068–2087, Nov. 2000.
- [13] S. Im and K. Banerjee, "Full chip thermal analysis of planar and vertically integrated high performance ICs," in *Proc. Int. Electron Devices Meeting (IEDM)*, 2000, pp. 737–740.



James W. Joyner received the B.E.E., and M.S. degrees in electrical and computer engineering, and the Ph.D. degree from the Georgia Institute of Technology, Atlanta, in 1999, 2001, and 2003, respectively. He is currently studying at the Candler School of Theology at Emory University in Atlanta, Georgia.

His research interests include the interaction of science, technology, and religion in society and the historical relationship of science and religion.



Payman Zarkesh-Ha (S'94–M'01) received the B.S. degree in electrical engineering from the University of Science and Technology, Tehran, Iran, in 1992, an M.S. degree in electrical engineering from Sharif University, Tehran, Iran, in 1994, and the Ph.D. degree in electrical engineering from the Georgia Institute of Technology, Atlanta, in 2001.

In March 2001, he joined LSI Logic Corporation as a Research Staff Member in the Interconnect Modeling Group in the Device Technology Division of LSI Logic Corporation, Milpitas, CA. He has worked on interconnect architecture design and optimization for the next ASIC generations. He has published more than 30 technical papers and a book chapter in this area and holds four U.S. patents.



James D. Meindl (M'56–SM'66–F'68–LF'97) received the B.A., M.A., and Ph.D. degrees in electrical engineering from Carnegie Institute of Technology, Carnegie Mellon University, Pittsburgh, PA.

He is the Director of the Joseph M. Pettit Microelectronics Research Center and the Joseph M. Pettit Chair Professor of Microelectronics at the Georgia Institute of Technology, Atlanta. He is also Director of the Interconnect Focus Center, a multiuniversity research effort managed jointly by the Microelectronics Advanced Research Corporation and the Defense Advanced Research Projects Agency for Department of Defense. His current research interests include the physical limits on gigascale integration and nanotechnology. Meindl is a Life-Fellow the American Association for the Advancement of Science, a Member of the American Academy of Arts and Sciences, and the National Academy of Engineering.

Dr. Meindl was recently awarded first place on the IEEE International Solid-State Circuits Conference 50-Year Anniversary Author Honor Roll. He received the Georgia Institute of Technology 2001 Class of 1934 Distinguished Professor Award, the IEEE Third Millennium Medal in 2000, the 1999 SIA University Research Award, and the 1997 Hamerschlag Distinguished Alumnus Award from Carnegie Mellon University.