Three-Dimensional Integrated Circuits: Performance, Design Methodology, and CAD Tools

Shamik Das, Anantha Chandrakasan, and Rafael Reif
Microsystems Technology Laboratories
Massachusetts Institute of Technology
60 Vassar St., Cambridge, MA 02139
{shamikd, anantha, reif}@mtl.mit.edu

Abstract

Three-dimensional integration technologies have been proposed in order to mitigate design challenges posed by deep-submicron interconnect. By providing multiple layers of active devices together with high-density local interconnects between these layers, 3-D technologies give digital-circuit designers greater freedom in meeting power and delay budgets that are increasingly interconnect-dominated. In this paper, we quantify the benefits 3-D integration can provide, using specific circuit benchmarks. We perform this analysis using a suite of circuit design tools we have developed for 3-D integration. We observe that on average, 28% to 51% reduction in total wire length is possible over two to five wafers respectively; similarly, 31% to 56% reduction in the length of the longest wire is achievable. We also characterize the impact of technology parameters on these reductions.

1. Introduction

Performance demands on integrated circuits will almost certainly scale in Moore's-Law fashion. The physics of integrated circuits, however, does not permit indefinite scaling in this manner. Devices may become smaller and smaller, up to a point. Yet at the same time, chip sizes become larger and larger. Global and semi-global wires therefore become longer and longer while the pitch of these wires becomes smaller. As a result, interconnects consume larger and larger portions of the power and delay budgets available to designers [2, 10].

Solutions to this problem range from architectural to technological in nature. *Interconnect-centric* architectures, for example, seek to avoid using global wires, relying instead on local wires and a network-style global communications framework. One ideally would have a technology that supports such architectural exploration while simultaneously ameliorating the interconnect performance of conventional architectures. Three-dimensional integration is one such technology. In a 3-D integrated circuit, transistors may be fabricated in multiple interconnected planes, as opposed to the single active plane provided in conventional integration. There are several approaches to fabricating such circuits [6, 14, 7].

Circuit designers, however, have only limited means by which to explore these new technologies. To date, only system-level predictive studies have been performed to assess in a general way the benefits provided by 3-D integration [12]. In this paper, we provide analyses of specific circuits that exhibit significant performance improvement when targeted for 3-D fabrication. We perform these analyses using a design flow we have developed for 3-D integration. The tools in this flow consist of a standard-cell placement and global routing tool and a 3-D IC layout editor.

In the following sections, we will examine the characteristics of potential 3-D integration technologies that impact circuit design and performance. We will describe our design flow for 3-D ICs. Finally, we will present placement and routing data, obtained using this flow for several circuits, exhibiting how circuit performance improves with the use of 3-D integration.

2. 3-D Integration Technologies

There are many approaches to fabricating three-dimensional integrated circuits or packages. These approaches may be characterized in terms of the maximum number of device layers and the maximum density of interconnects between these layers. We classify as packaging approaches those that involve stacking of pre-packaged ICs. At this end of the spectrum lies MCM-V, a vertically-stacked multi-chip module technology [6]. At the other end lie fully-integrated technologies, where multiple dice

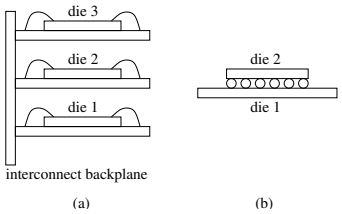

Figure 1. Basic structure of (a) a three-layer MCM-V package and (b) a flip-chip package (with dice bonded face-to-face using a solder-bump interconnect).

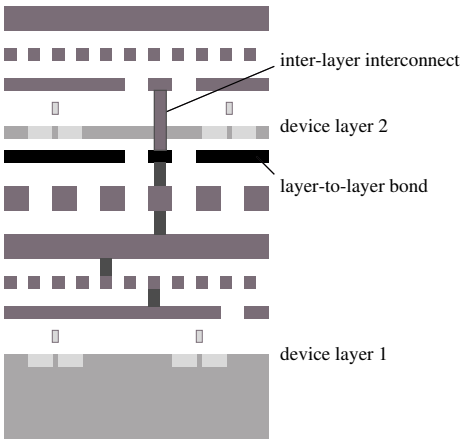

Figure 2. Wafer-bonded structure with two device layers and copper interconnect interface.

or wafers are stacked, or multiple device planes are fabricated on a single die, prior to packaging. Such technologies include solid-phase crystallization [14] and wafer bonding [7]. Along the spectrum exist several other alternatives, such as liquid-phase [13] and flip-chip technologies.

From a circuit-design perspective, the choice of technology involves some trade-off analysis. MCM-V, for example, is a proven, low-cost technology. However, the crucial drawback with MCM-V is that all inter-layer interconnect must pass through a backplane, so that within each layer, the inter-layer interconnect must be routed to the perimeter (Figure 1). As a result, MCM-V is generally used only where single-die integration is not feasible.

Conversely, fully-integrated technologies offer high-density interconnect between device layers. The density and performance of this interconnect varies with the technology. Flip-chip packaging, for example, uses solder bumps (Figure 1) with pitch on the order of 100 microns. On the other hand, crystallization and wafer-bonding techniques use vias to form the inter-layer interconnect, with pitch as low as a few microns.

A wafer-bonded structure is shown in Figure 2. In this technology, discrete wafers are processed and stacked. The bond is formed using a dielectric glue layer or (in the case of [7]) using two layers of copper metallization that are fused together under pressure. The bottom device layer is formed from a bulk silicon wafer 500 to 700 microns thick; subsequent device layers may be as thin as one or two 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 inter-wafer interface.

Wafers may be bonded face-to-face (i.e. such that the metallizations are adjacent, as in Figure 3) or face-to-back (Figure 2). Even in this implementation choice, there are trade-offs; while face-to-back bonding is required for three or more device layers, the inter-layer interconnects in face-to-back bonding are formed using vias through the device area of the upper layer, thus reducing the die area available for devices. Furthermore, there are metal-coupling and substrate-coupling differences between the two scenarios.

For the purposes of this study, we will consider only those technologies that offer dense inter-layer interconnect over the entire die surface. Assuming a high-density inter-layer interconnect, this means that every transistor has more "nearest neighbors" available to it. This in turn leads to a greater percentage of communication that can be performed using local wires. Indeed, stochastic wire-length models for 3-D integration [12] have been used to predict how the distribution of wire lengths in a circuit is shifted (Figure 4). These sorts of predictions strongly motivate the development of design tools with which we can analyze performance improvements for specific circuits.

3. 3-D Design Methodology and Tools

Our design tools are meant to replace existing tools in a digital design flow where such tools need to be aware of the three-dimensionality of the circuit. A typical design flow is given in Figure 5. In order to adapt this flow for a 3-D integration technology, we have developed a standard-cell placement tool, global routing tool, and a layout editor for 3-D ICs [4].

Our standard-cell place-and-route tool, PR3D, operates on circuit netlists in GSRC [9] or Cadence LEF/DEF [8] formats, which may be generated from Verilog or VHDL descriptions of the circuits. The user specifies the number of device layers to be targeted, as well as the wire-length cost of inter-layer vias relative to the cost of 2-D wire (for example, if least interconnect power is the target, the cost of inter-layer vias may be given as the capacitance of such vias relative to the capacitance of a unit length of 2-D wire). The user also must choose whether to minimize the total wire length or to minimize the number of inter-layer vias

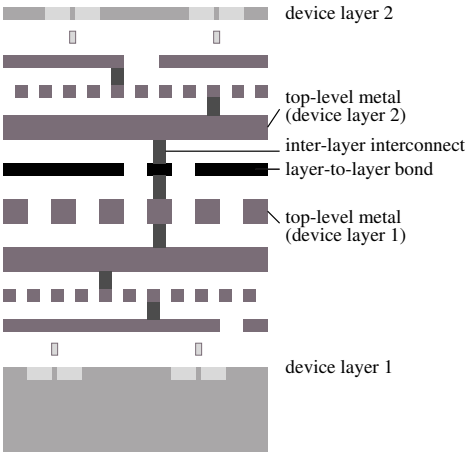

Figure 3. Wafer-bonded structure where the metallizations of the wafers are adjacent to the bonding interface (face-to-face configuration).

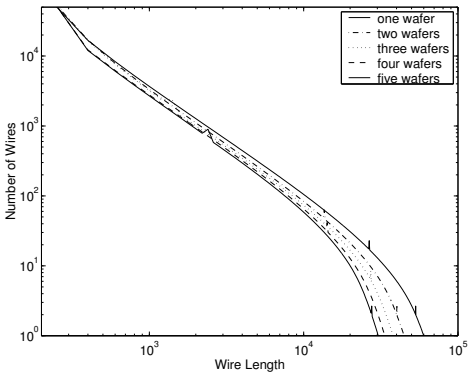

Figure 4. Predicted wire length distribution for a sample circuit as a function of the number of device layers used.

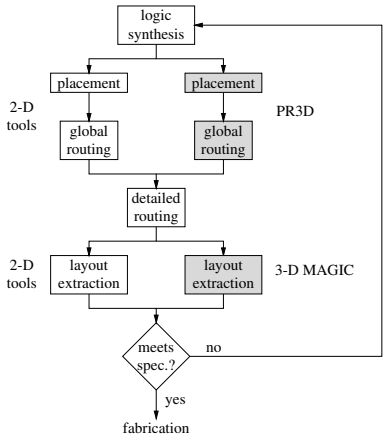

Figure 5. Design flow for 3-D integrated circuits. Highlighted areas indicate where our tools replace tools for conventional ICs.

(for technologies where the inter-layer wire pitch dominates over other factors such as capacitance). Placement is performed using a min-cut partitioning algorithm [5] that we have adapted for 3-D placement [4].

Once placement is complete, global routing is performed in order to localize the routes taken by nets in the circuit. 2-D global routing is commonly done by partitioning the die surface into a rectangular grid and allocating routes to the grid boxes according to a congestion criterion. In our 3-D global router, we additionally must allocate area for inter-layer vias: in the face-to-back configuration, such vias present obstacles not only for device placement, but also for routing in first-level metal. We perform the allocation of routes, including inter-layer vias, using a hierarchical algorithm based on [3].

When the locations of inter-layer vias have been set, the individual layers of the placement may be passed to existing tools for routing. Once routing is complete, however, the layers must be merged in order to extract and simulate the circuit. We have extended the Berkeley layout editor Magic [11] to manage multiple-layer designs. Our tool, called 3-D MAGIC, allows hand-design of 3-D integrated circuits as well as manipulation of designs generated by our place-and-route tool. 3-D MAGIC automates the display of inter-layer vias between the two relevant layers of the circuit, and maintains electrical-connectivity information for circuit nodes across all layers. 3-D MAGIC can also extract an entire 3-D circuit at once, using a user-configurable lumped-parameter model for the inter-layer interconnect. With this extraction, existing tools can be used to simulate circuit performance and close the design loop.

4. Circuit Performance Characterization

Using our tools, we placed and routed the ISPD '98 set of 18 digital benchmark circuits (ibm01 – ibm18) [1]. In our first analysis (Figure 6), we show the total wire length of the average of these circuits as a function of number of device layers, determined from placement; Figure 7 gives the results for routing. We observe that relative to a 2-D (conventional single-die) placement, 28% to 51% reduction in total wire length is possible; we expect that capacitance (and therefore energy consumption) follows the wire length. Furthermore, with our tool we can set the capacitance of an inter-layer via relative to the capacitance of an equal length of wiring from, say, metal 1. The four curves in the figure show quantitatively how as the inter-layer via capacitance increases, the benefit of 3-D integration decreases. We conclude that optimizing inter-layer interconnect is of key importance for 3-D integration technologies. For example, we compare placements of the ibm03 benchmark circuit using inter-layer vias versus using solder bumps for the inter-layer interconnect (Figure 8). (It is important to note that

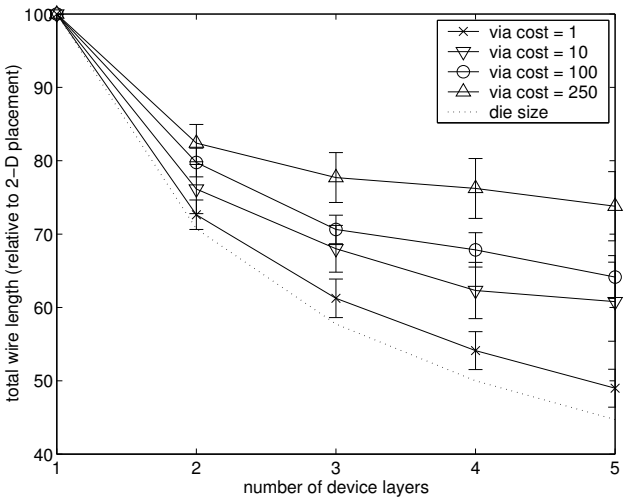

Figure 6. Total wire length (as a function of number of device layers) for various inter-layer via capacitances, obtained from placement. Total wire length is minimized by the placement tool. Via cost is the via capacitance expressed relative to the capacitance of one micron of metal wire.

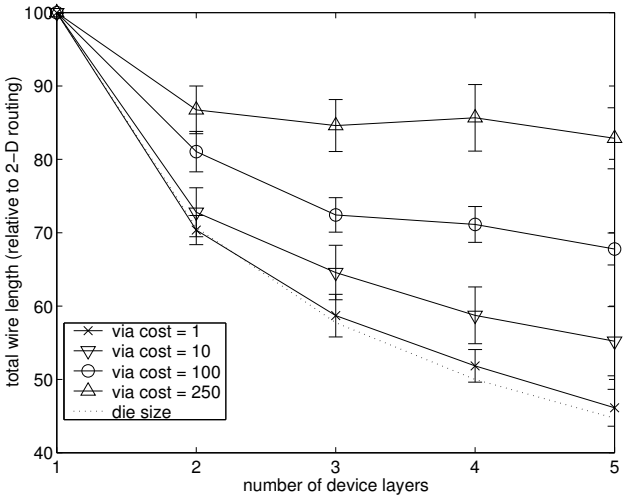

Figure 7. Total wire length (as a function of number of device layers) for various inter-layer via capacitances, obtained from routing. Total wire length is minimized by the routing tool. Via cost is the via capacitance expressed relative to the capacitance of one micron of metal wire.

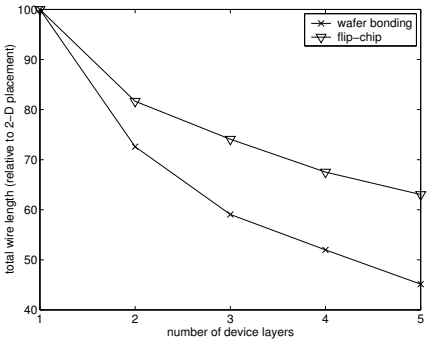

Figure 8. Total wire length (as a function of number of device layers) of the ibm03 benchmark, using vias vs. using solder bumps for the inter-layer interconnect.

for three or more device layers, through-wafer vias must be used in conjunction with the solder bumps. However, in the two-wafer case, we assume a face-to-face implementation, where such vias are not required.) It is clear that technologies such as wafer bonding offer significant performance improvement even when compared to existing die-stacking methods.

Our second analysis (Figures 9 and 10) shows the same experiment, but where we minimize the number of inter-layer vias (as opposed to minimizing the total wire length). We see that the benefit of 3-D integration is not so great here (up to 7% to 17% reduction in total wire length), but that the results are more immune to inter-layer via capacitance variation.

Figures 11 and 12 show the length of the longest wire from the same two analyses, where we observe up to 31% to 56% reduction in this length; this confirms the expected reduction in global signal delays. We see that the longest wire is not affected strongly by the inter-layer via capacitance. We can also see that this length tracks the side length of the die fairly well.

Additionally, for purposes of comparison with Figure 4, we obtain the placed wire length distribution for the same circuit, shown in Figure 13.

Finally, in Figures 14 and 15, we show layouts of a sample circuit that has been placed and routed in two device layers using our tools.

5. Conclusion

In response to the increasing demands of interconnect on circuit performance budgets, many forms of three-dimensional integration technologies have been proposed. As such technologies mature, circuit designers will need to be able to exploit them. To this end, we have presented a design flow for 3-D integrated circuits.

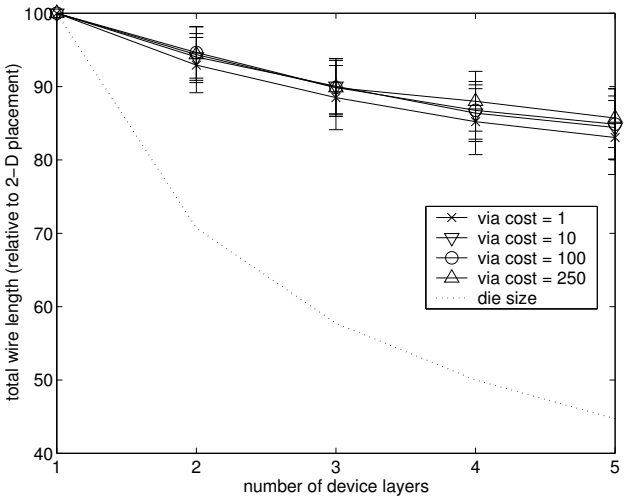

Figure 9. Total wire length (as a function of number of device layers) for various inter-layer via capacitances, obtained from placement. The number of inter-layer vias is minimized by the placement tool. Via cost is the via capacitance expressed relative to the capacitance of one micron of metal wire.

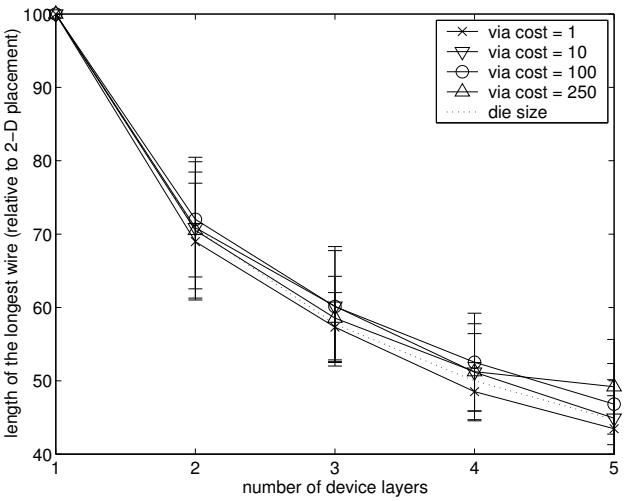

Figure 11. Length of the longest wire (as a function of number of device layers) for various inter-layer via capacitances. Total wire length is minimized by the placement tool. Via cost is the via capacitance expressed relative to the capacitance of one micron of metal wire.

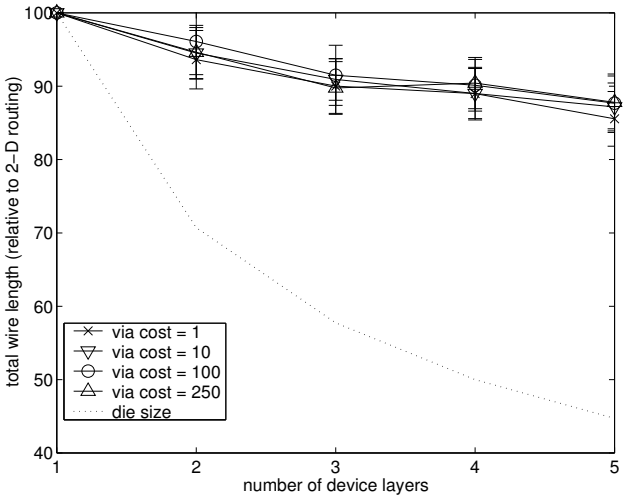

Figure 10. Total wire length (as a function of number of device layers) for various inter-layer via capacitances, obtained from routing. The number of inter-layer vias is minimized by the routing tool. Via cost is the via capacitance expressed relative to the capacitance of one micron of metal wire.

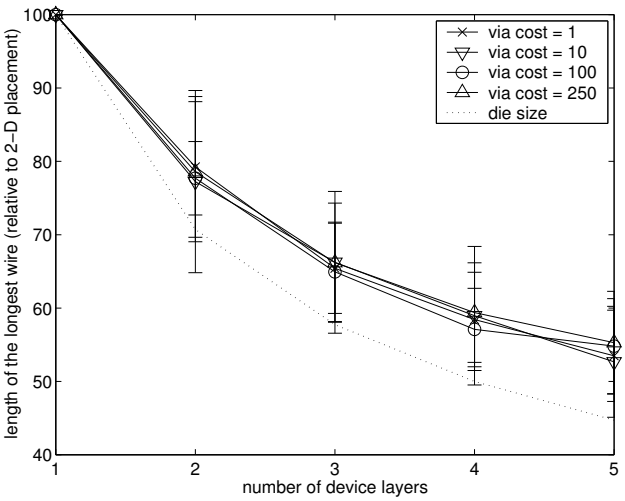

Figure 12. Length of the longest wire (as a function of number of device layers) for various inter-layer via capacitances. The number of inter-layer vias is minimized by the placement tool. Via cost is the via capacitance expressed relative to the capacitance of one micron of metal wire.

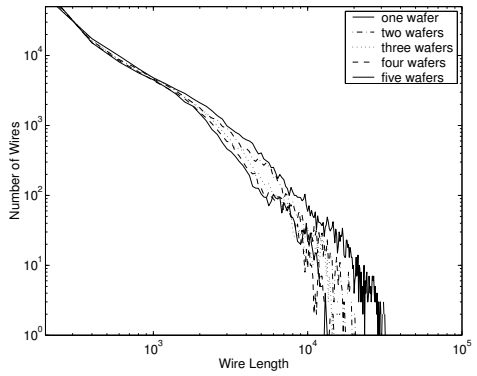

Figure 13. Placed wire length distribution for a sample circuit as a function of the number of device layers used.

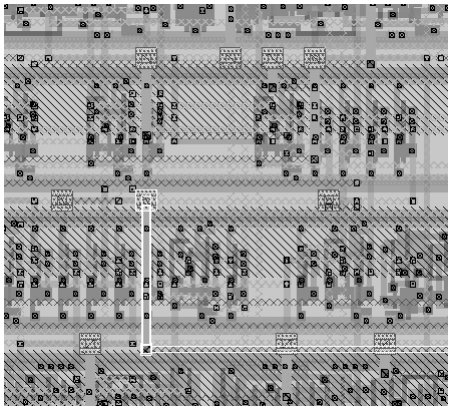

Figure 14. First device layer of a two-layer placement showing a selected net spanning both device layers.

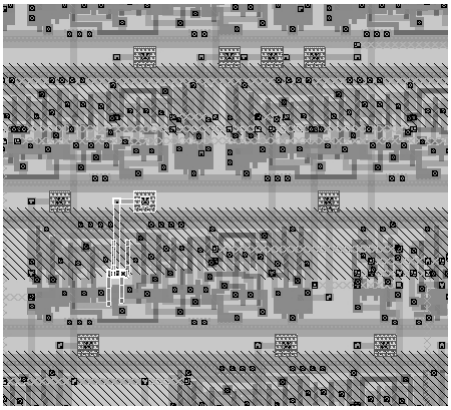

Figure 15. Second device layer of a two-layer placement showing a selected net spanning both device layers.

We examine 3-D technologies using these tools and a set of circuit benchmarks. We find that the most promising of these technologies offer significant benefits to circuit designers. For example, we see up to 28% to 51% total wirelength reduction and 31% to 56% reduction in length of the longest wire using two to five device layers. We also observe that 3-D technology parameters such as inter-layer via capacitance have significant impact on these savings. With the design tools we present, circuit designers will certainly be able to leverage the benefits 3-D integration provides.

6 Acknowledgments

This paper acknowledges support from the MARCO Interconnect Focus Center, funded at MIT through subcontract B-12-M00 from the Georgia Institute of Technology. This paper also acknowledges support from DARPA under grant B-12-D00.

References

[1] C. J. Alpert. The ISPD98 circuit benchmark suite. In *ACM International Symposium on Physical Design*, pages 80–85, April 1998.

[2] M. T. Bohr. Interconnect scaling – the real limiter to high-performance ULSI. In *Proc. IEDM*, pages 241–244, 1995.

[3] M. Burstein and R. Pelavin. Hierarchical wire routing. *IEEE Transactions on Computer-Aided Design*, CAD-2(4):223–234, 1983.

[4] S. Das, A. Chandrakasan, and R. Reif. Design tools for 3-D integrated circuits. In *Proc. ASP-DAC 2003*, to appear.

[5] A. E. Dunlop and B. W. Kernighan. A procedure for placement of standard cell VLSI circuits. In *IEEE Transactions on Computer-Aided Design*, pages 92–98, 1985.

[6] C. W. Eichelberger. Three-dimensional multichip module system. United States Patent 5,111,278, May 1992.

[7] A. Fan, A. Rahman, and R. Reif. Copper wafer bonding. *Electrochemical and Solid-State Letters*, 2:534–536, 1999.

[8] http://openeda.org/download_lefdef.html.

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

[10] International technology roadmap for semiconductors. http://public.itrs.net/.

[11] J. Ousterhout et al. Magic: A VLSI layout system. In *Proceedings of the 21st Design Automation Conference*, pages 152–159, 1984.

[12] A. Rahman and R. Reif. System-level performance evaluation of three-dimensional integrated circuits. *IEEE Trans. on VLSI Systems, Special Issue on System-Level Interconnect Prediction*, 8(6):671–678, 2000.

[13] B. A. Ridley, B. Nivi, and J. M. Jacobson. All-inorganic field-effect transistors fabricated by printing. *Science*, 286:746–749, 1999.

[14] 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 Letters*, 18:378–381, Aug. 1997.