CAD Implications of New Interconnect Technologies

Louis K. Scheffer
Cadence
555 River Oaks Parkway
San Jose, CA
lou@cadence.com

ABSTRACT

This paper looks at the CAD implications of possible new interconnect technologies. We consider three technologies in particular: three dimensional ICs, carbon nanotubes as a replacement for metal interconnects, and optical interconnections for longer range on-chip communication. Each of these requires new CAD support to be used effectively.

Categories and Subject Descriptors

B.7.1 [**Integrated Circuits**]: Types and Design Styles—*Advanced technologies, VLSI*; B.7.2 [**Integrated Circuits**]: Design Aids—*Placement and routing*

General Terms

Algorithms, Design, Performance

Keywords

3D interconnect, Nanotubes, On-chip optical

1. INTRODUCTION

Since interconnect is one of the main performance determinants of modern ICs, there are many proposals to improve it. As of early 2007, some of the main proposals are

- 3D Interconnect
- Carbon nanotubes to replace wires and/or vias
- Optical interconnect for long on-chip connections

Each of these requires CAD changes to be effective.

2. 3D CHIP CONSTRUCTION

Existing IC technology has a single surface for active devices, and a number of interconnect layers above it. If more layers of active devices can be added, there is the possibility of much shorter connections. This could save power, increase speed, and add to density. Chips built of many different technologies could be combined, each optimized for a function such as RF circuits or memories. On chip busses might be much shorter, leading to higher bandwidths between processors and memory. There have been many studies of these and other benefits that might be gained by such technology. See for example [1][2][3][4].

3D chips may potentially be constructed in many ways. Conventional chips can be stacked, or new techniques can possibly grow active layers on top of existing interconnect. In the terminology of 3D ICs, each layer is called a *tier*.

There are also many alternatives for CAD tools for 3D. One possibility is to do 3D CAD as an analog of the existing 2D CAD - this could include 3D placement[5], routing[5], congestion estimation[6], partitioning, FPGA tools[7], and so on. Many of the needed algorithms are relatively straightforward extensions of existing 2D algorithms - placement by many means (min-cut, analytic, force-directed, and simulated annealing) generalizes well to 3D [1]. Similarly, global routing in 3D should be quite similar to 2D since global route that includes layer assignment is already working with a 3D structure. Detailed routing is currently unclear since there is no consensus on the vertical cross-tier via structure, and hence no-one has built a detailed router for this case yet. However, no serious obstacles are apparent.

One important difference in 3D is that thermal issues become much more important. 2D chips are built on thin substrates, and the thermal path to the heatsink is therefore short. In 3D, the heat generated in the interior tiers must travel through the other tiers and interconnect to escape. Therefore 3D algorithms, particularly placement, must be sensitive to thermal concerns, and optimize them during design. Power distribution also becomes more difficult, since power can only be supplied to the outside of the stack, and power driven placement may be more important as well.

From an industrial perspective, (in the author's opinion) full 3D design seems likely only for FPGAs. For ASICs, only the partitioning and thermal analysis will be done in 3D. After this, conventional 2D tools will be used to design each tier. There are several practical reasons in favor of this approach:

- There is a large infrastructure for 2D design, and 2D design is likely to persist for quite a while. Unless 3D catches on in a big way (unlikely because of cost) it won't be economical to build 3D specific tools. One previous problem, that 2D tools were not set up for

[1]There is a gold mine of potential PhD theses here.

DAC 2007, June 4–8, 2007, San Diego, California, USA.

IOs over the whole surface of the chip, is now addressed fairly well since flip-chip shares this characteristic with 3D.

- Designers currently partition 2D designs anyway, just to get a handle on complexity, reduce interaction between groups, hit the sweet spot of DA tools, include IP, and so on. As long as the designer is partitioning, and cross-tier connections are expensive, then these make good partitioning boundaries, too.
- If the design is partitioned into tiers, most engineering changes (ECOs) will take place within a single tier. If the design is fully 3D, ECOs could very well be spread across many tiers, resulting in a much higher cost for mask changes.
- IP will mostly be available in 2D form. IP providers will not want to limit their market to those advanced customers building 3D chips, and so will provide 2D blocks.

Once a 3D IC is designed, it will need extraction, which takes the specified geometries and computes the corresponding resistors, capacitors and inductors for the interconnections. In 3D, resistance and capacitance extraction will be little changed, or at most a straightforward extrapolation of current 2D techniques. This is because resistance and capacitance are local effects, so within-tier and between-tier components can be extracted independently. Inductance calculation, however, will be considerably changed, since inductance depends strongly on the return path for the signal current. The return path may be quite far from the signal under consideration, and determining the return path is a difficult problem. Existing extractors use a number of heuristics (such as [8]), which have been tuned for existing 2D chip architectures. In 3D the problem will be more complex, since conductors on the adjacent tier may or may not contribute to the return path, depending on the connections. At the very least, new heuristics will be needed, and where there is significant current flow through the between-tier connections, a full 3D approach to inductance may be required.

3. NANOTUBES REPLACING WIRES

Carbon nanotubes are one of the few materials that are better electrical conductors than copper. They are potentially better in several ways - some nanotubes have higher conductivity than copper, all are much more resistant to electromigration, and some are more thermally conductive than copper. See [9][10][11] for discussion of the electrical properties of nanotubes. This naturally leads to the suggestion of replacing copper or aluminum conductors with conductors made of nanotubes [12][13][14]. However, the increased conductivity comes with a number of limitations. First, it is only along the length of the tube, not across it. This is not a serious problem for chip design, where the direction of current flow is well known, though it may well be a difficult manufacturing problem.

One of the nice points of nanotubes is that they do not suffer from electromigration, for all practical purposes. The atoms are tightly bound into a lattice, and the electron wind cannot move any of them. Also there are no grain boundaries. Carbon nano-tubes will melt before they suffer from electromigration.

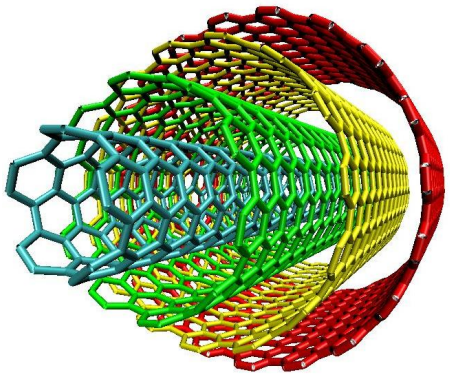

Figure 1: A multi-wall nanotube (MWNT) consists of concentric nested nanotubes, as shown in this diagam, courtesy of A.V. Krasheninnikov. Single-wall nanotube (SWNT) interconnects consist of parallel bundles of the smallest diameter nanotubes, like those in the center of this MWNT.

Electrical conductivity in nanotubes is more complex than conductivity in metals. The resistance of a nanotube is not a linear function of its length - instead the resistance is constant up to a certain length, then increases linearly thereafter. The minimum resistance, present even for very short nanotubes, is defined by the electron states available in its 1-D conductor. It appears half at each end, where the 1D conductor meets the 3D electron sea in the contact. In the nanotube itself, there is no additional resistance (electron transport is ballistic) for lengths up to (roughly) the mean free path. This mean free path is proportional to the nanotube diameter, and is about 1.6μ for the smallest diameter tubes. Above this length, the resistance increases linearly[15]. Furthermore, the temperature coefficient of a nanotube conductor can be either sign - higher temperatures reduce the mean free path, increasing resistance, but also can promote (in multi-wall tubes) more electrons into the conduction band, reducing resistance. Finally, nanotubes built with current techniques can be of one of three chiralities (basically where you end up after traversing around the cylinder following the atoms). Only one of these conducts - the other two are semiconductors, with very poor conduction compared to wires. The comparisons below all assume random construction. If techniques for manufacturing only the metallic type nanotubes can be perfected, the results below could improve by a factor of 3. Although we have no idea how to do this now, we can hope since nanotubes are constructed by catalysts, and biological catalysts often exhibit this type of specificity.

Given all this, how do nanotube conductors compare to copper? This depends on both the width and length of the wire. Copper conductors become worse at smaller widths, from scattering off the grain boundaries and rough edges, and from the higher percentage of low-conductivity cladding. Even so, for short lengths where the unavoidable quantum R dominates, a wire made of single-wall nanotubes is always a worse conductor than a copper wire of the same cross section. As the length increases, a bundle of single wall nanotubes becomes slightly better than an equivalent area of narrow copper, but still somewhat worse than bulk copper.

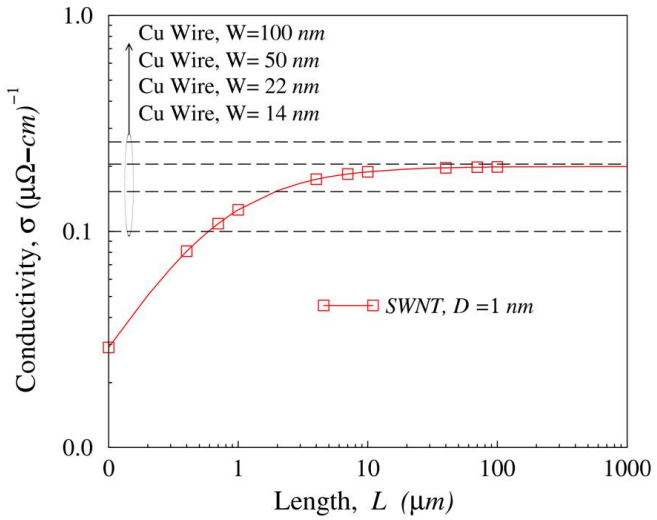

Figure 2: Long single wall carbon nanotubes are better conductors than narrow copper wires of the same area, but worse than wide wires or bulk copper. Very short wires are worse since contact resistance dominates. From [16].

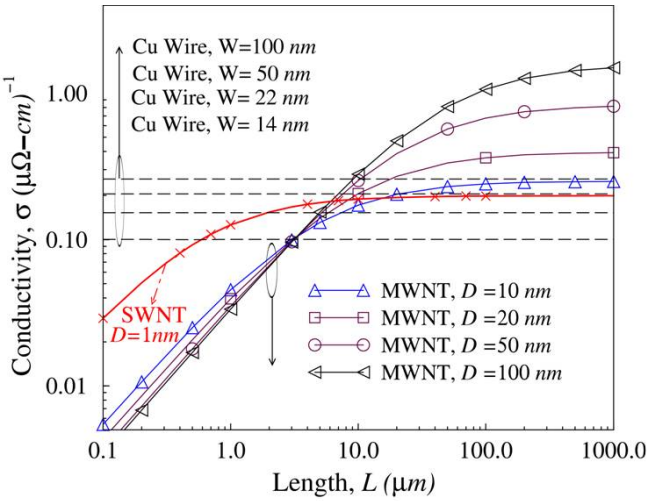

Figure 3: Large diameter carbon nanotubes are better conductors than bulk copper if they are long enough and big enough diameter. This is because the mean free path is directly proportional to the diameter of the tube. From [16].

This is shown in Fig. 2. Multi-wall carbon nanotubes are even worse at short lengths than single wall nanotubes. This is because there are fewer of them per unit area, so the unavoidable contact R hurts more. However, as the wire gets longer, the ballistic transport dominates, especially in large diameter tubes with their long mean free paths. A long enough length of a large diameter nanotube may be many times as conductive as copper. This is shown in Fig. 3.

Since conduction through nanotubes consists of relatively few electrons travelling far and fast, there is another term in the inductance, called *kinetic inductance*. This is due (loosely) to the kinetic energy of the electrons[17]. Unlike a conventional inductance, it does not depend on the distance to a return path, and decreases linearly with the number of parallel nanotubes. This effect occurs in normal wires too, but is negligable compared to the classical inductance from the induced magnetic field. In a single nanotube, the kinetic inductance is the dominant one, to the point where only a parallel bundle of nano-tubes makes any sense. Fortunately nanotubes are thin, and so generally will be used in parallel bundles, even at 22nm or 10nm technologies. In these cases the kinetic inductance can in general be neglected[16].

Finally, the surface of the nanotube is built differently than a metal surface, leading to slight differences in capacitance (a few percent).

With all of these effects, plus difficulties in manufacturing, it is not currently clear if nanotubes will be a practical alternative. See [18] for a pessimistic view, and [19] for an optimistic view.

There are at least three ways nano-tubes may be used in IC design. For short interconnects, a conventional copper interconnect with high aspect ratio can be replaced by a thin sheet of nanotubes, as shown in Fig 4. This reduces the capacitance of wires, particularly the lateral capacitance. This in turn implies better performance, less unwanted coupling, and lower power dissipation. Such thin conductors cannot be used with copper for two reasons - it is hard to fabricate such shapes because of CMP problems, and even if they could be built they would rapidly fail due to electromigration. Nanotubes may be capable of being grown as a thin layer (and hence not need CMP) and do not suffer from electromigration, for all practical purposes. The capacitance of a thin nanowire ribbon will be much lower than that of the copper wire, perhaps as little as 30% as much, depending on the aspect ratio of the copper. On the other hand, for such thin conductors the resistance of nanotubes will be higher than that of traditional thick copper, but short interconnects (up to 100 gate pitches or so) are not dominated by wire R, but by driver R, and gate and interconnect C. Therefore replacing short copper wires by a thin layer of nanotubes could provide both lower power dissipation and increased performance[20]. Since a large fraction of power is dissipated due to the capacitance of local nets (50% in the example of Intel in [21]), this could help the system power considerably even if restricted to local nets only. For longer lengths, the increased R becomes more important, and the thin nanowires become slower than conventional copper wires, as shown in Fig. 5.

The second way nano-tubes might be used is to replace longer wires with better electrical conductors. In this cause multi-walled nanotubes must be used, since only these nanotubes are better conductors than bulk copper[22]. This means they will be fairly thick, and have a similar capacitance to a copper wire, but lower resistance. They may be particularly useful as power supply wires, due to their lower resistance and lack of electromigration problems.

Finally, vertical bundles of nano-tubes may be used to replace vias. This has two advantages - they are essentially immune to electromigration, and they are excellent thermal conductors. This reduces the temperature gradients through the interconnect stack. [23][24][25]. This might go very well

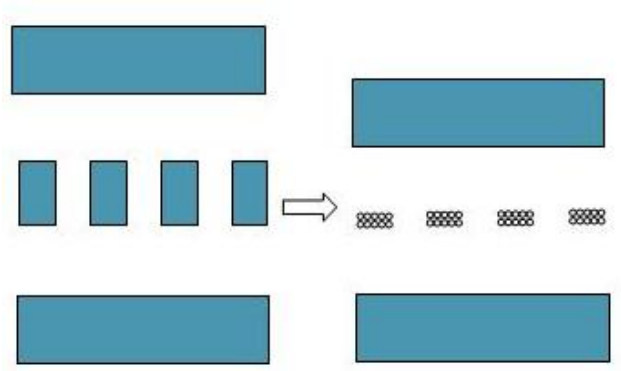

Figure 4: Proposed use of nanotubes to replace copper wires. From presentation associated with [16].

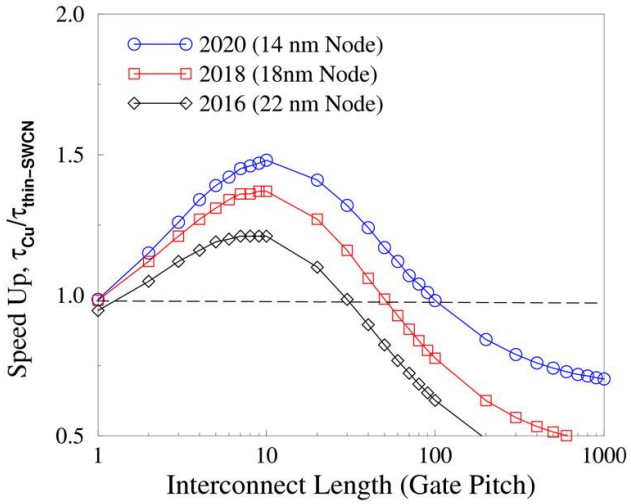

Figure 5: Range of lengths at which a thin nanowire ribbon is faster than a copper wire of conventional thickness. From [16].

with the 3D interconnections above, which can utilize higher vertical thermal conductivity.

3.1 CAD changes for nanotubes

The constant contact resistance of nanotubes means the calculation of the resistance of a connection must change and become length dependent. Also, since the length dependence of nanotube R depends on the mean free path, this will need to be specified. Since the mix of nanotubes used by each wiring layer may be different, it will need to be a parameter supplied by the foundry or user.

Nanotubes cannot turn tight corners, so at every corner there will be a block of metal, and 2 metal-nanotube contacts. This will affect the routers, probably resulting in one layer of purely vertical wires and one for purely horizontal wires. This resembles a proposal for *restricted design rules*[26], which proposes this same restriction to make lithography easier. Electromigration checking will also need to change - through the nanotubes themselves are virtually unaffected by electromigration, the normal metal used to make the contacts can still deteriorate. New models and new algorithms will be needed to check this.

Layer assignment may also be affected. If thin sheets of nanotubes replace the lower layers, they will have much higher resistance than the copper wires they might replace. We may need to resurrect old algorithms from the days of polysilicon that are very careful to route only short wires in these lower metal layers. Alternatively, timing driven routing might accomplish the same objectives, but it would need to be tested for use with such a wide range of resistances.

Finally, the impact and forms of process variation may differ for nanotube interconnect[27]. This must be taken into account in any statistical analysis.

4. ON-CHIP OPTICAL INTERCONNECT

Optical interconnect has promise for long distance connections within a chip. The main disadvantage is that it requires electrical-optical conversion at the sender, then optical-electrical conversion at the receiver. The advantages are that a much greater distance can be covered without repeaters, and that many signals can share a route with wavelength division multiplexing, and that optical is very promising for inter-chip and inter-board communication as well. Optical signals can also cross without interacting. Optical transmitters can be *direct*, where the transmitter generates the light directly, or *indirect*, where the transmitter modulates an external light source. See [28] or [29] for overviews of possible on-chip optical technology, and [30] for a comparison with copper interconnect.

The length of connection where optical makes sense is a strong function of optical technology. The electrical technology is by comparison quite static. For example, the capacitance of a wire is roughly 200 ff/mm, almost independent of scaling. The power P required to switch a wire is $P = CV^2f$, where C is the capacitance, V is the voltage swing, and f is the frequency. For a 1V, 1GHz signal, for example, this gives 0.18 mw/mm for the wire alone. Including the necessary repeaters might double this, giving 0.36 mw/mm. So if an optical transmitter/receiver pair takes 3.6mw, this starts to look good at 10mm. This break-even distance will be inversely proportional to the number of transitions, so the activity of a signal must be considered as well as its length.

Optical transmitters, receivers, and waveguides also take up a fair amount of real estate. Direct transmitters require a large buffer chain to generate the laser drive current, and indirect transmitters are larger yet - the smallest one reported as of early 2007 was about 40 microns long. Receivers are also large, also needing significant amplifying circuitry. Any specific figures for the sizes will need to be obtained from the current technology - this is an area of much active research. On the other hand, optical waveguide cannot be much narrower than a wavelength, or about 1μ (1000nm wide), and this is unlikely to improve. However, by using different wavelengths, many signals can travel in the same connection without interference. Some studies indicate this is needed if optical is to achieve the same bandwidth density as copper interconnections[31].

4.1 CAD changes for on-chip optical

Since the optical components needed to launch or receive

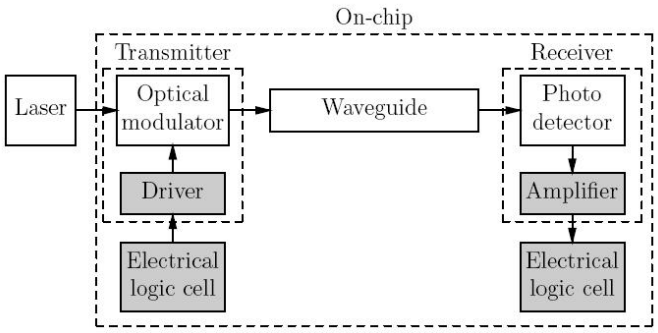

Figure 6: Overview of on-chip optical interconnect. An indirect optical source is shown, which modulates an external light source. A direct source would be similar, but with no external light input. From [28].

signals are much larger than the electrical equivalents, they will be needed to be included in the floorplan, which means the optical signals must be identified early. Also since optical signals only make sense past a certain length that depends on activity, floorplanners and other chip planning tools will need to be more aware of the switching activities of signals.

On the plus side, optical signals do not need repeaters, for the lengths we are worried about on-chip. This can considerably simplify floorplanning. Also they do not suffer from crosstalk. This can simplify extraction and noise analysis.

Turning an optical signal through an angle is more involved than doing so with an electrical signal. With a waveguide structure, only fairly gradual turns are possible, with a radius of at least a few microns. Other methods such as mirrors or diffraction gratings may allow sharper turns, but may be bigger or carry other restrictions.

Most interconnect layers will be constructed to carry optical or electrical signals, but not both. It is not yet clear whether separate vertical and horizontal optical layers are needed, or whether one layer suffices (since optical signals can cross each other under the right conditions). Clearly this will affect layer assignment and pin assignment, since a signal cannot casually be swapped between an optical and an electrical layer, nor a pin changed from optical to electrical or vice versa.

Many optical signals can be carried in one waveguide. This might be particularly helpful for busses, where many signals are routed to the same place. Of course this relies on the availability of optical transceivers running at different wavelengths. This is a subject of current research but by no means guaranteed.

Power analysis does not need to consider the impact of repeaters on optical signals. On the other hand the transmitters and receivers may consume significant static power, unlike their electrical equivalents.

One purely practical matter will be the development of new library formats, routing rules, and other technology files. Optical tranceivers will need to have delays, power requirements, and perhaps wavelengths, or wavelength ranges, specified. Routing rules will need to include size limits, crossings if the signals can be crossed in a given layer, how to turn signals, and so on.

5. PROCESS VARIATION/CORRELATION

All these new technologies will probably share at least one characteristic with existing methods - process variation. This means that the constructed interconnects will not all be exactly the same, and not exactly as desired. In addition, the variations from nominal will almost surely be correlated, with nearby components much more alike than components spaced far away on the chip.

This problem is not well dealt with today, even with the existing metal interconnects. The fab typically only states a range of variation for identical wires in identical contexts within a short distance of each other. Much more general models of correlation are needed, that can find the correlation between (for example) wires of different widths on the same layer, wires on different layers, and wires far apart on the chip. There have been many statistical timing papers on dealing with correlation, but without data from the fabs it is hard to get much traction on this issue.

Since these new interconnect technologies are not yet well specified, we can only hope that the solution that is finally adopted for the traditional metal interconnnect will apply to these new interconnects.

6. CONCLUSIONS

There are several new technologies that have the potential to change the way we deal with interconnect on chips. Three of the most likely are 3D chip construction, carbon nano-tube interconnect, and optical interconnect. Each of these requires CAD support before the advantages can be used. Fortunately this does not look terribly difficult. Nanotube interconnect is the easiest, where changes to electromigration checking and resistance calculations suffice. 3D is probably the next hardest, though the considerable research that has gone into 2D can help. The main new feature will be mandatory thermal analysis and improved power distribution analysis. Finally optical interconnect would stress CAD tools the most. Widespread adoption would require changes in floorplanning, global and detailed routers, pin assignment, congestion calculations, and many components of today's IC design flows.

7. ACKNOWLEDGMENTS

The 2006 SRC interconnect forum had a number of excellent talks that presented progress in these areas. Azad Naeemi kindly granted permission for many of the nanotube pictures.

8. REFERENCES

[1] W.R. Davis, J. Wilson, S. Mick, J. Xu, H. Hua, C. Mineo, A.M. Sule, M. Steer, and P.D. Franzon. Demystifying 3 D ICs: the pros and cons of going vertical. *IEEE Design & Test of Computers*, 22(6):498–510, 2005.

[2] SJ Abou-Samra, PA Aisa, A. Guyot, and B. Courtois. 3D CMOS SOI for high performance computing. *Low Power Electronics and Design, 1998. Proceedings. 1998 International Symposium on*, pages 54–58, 1998.

[3] MJ Alexander, JP Cohoon, JL Colflesh, J. Karro, and G. Robins. Three-dimensional field-programmable gate arrays. *ASIC Conference and Exhibit, 1995., Proceedings of the Eighth Annual IEEE International*, pages 253–256, 1995.

[4] C. Ababei, P. Maidee, and K. Bazargan. Exploring potential benefits of 3D FPGA integration. *Field-Programmable Logic and its Applications*, pages 874–880, 2004.

[5] C. Ababei, Y. Feng, B. Goplen, H. Mogal, T. Zhang, K. Bazargan, and S. Sapatnekar. Placement and routing in 3 D integrated circuits. *IEEE Design & Test of Computers*, 22(6):520–531, 2005.

[6] L. Cheng, WNN Hung, G. Yang, and X. Song. Congestion estimation for 3D routing. *VLSI, 2004. Proceedings. IEEE Computer society Annual Symposium on*, pages 239–240, 2004.

[7] J. Karro and JP Cohoon. A Spiffy tool for the simultaneous placement and global routing for three-dimensional field-programmable gate arrays. *VLSI, 1999. Proceedings. Ninth Great Lakes Symposium on*, pages 230–231, 1999.

[8] KL Shepard and Z. Tian. Return-limited inductances: a practical approach to on-chip inductance extraction. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 19(4):425–436, 2000.

[9] S. Li, Z. Yu, C. Rutherglen, and P.J. Burke. Electrical properties of 0.4 cm long single-walled carbon nanotubes. *Arxiv preprint cond-mat/0408332*, 2004.

[10] PJ Burke. An RF circuit model for carbon nanotubes. *IEEE-NANO 2002: Proceedings of the 2nd IEEE Conference on Nanotechnology*, pages 393–396, 2002.

[11] HS Gokturk. Electrical properties of ideal carbon nanotubes. *Nanotechnology, 2005. 5th IEEE Conference on*, pages 800–803, 2005.

[12] N. Srivastava and K. Banerjee. Performance analysis of carbon nanotube interconnects for VLSI applications. In *ICCAD '05: Proceedings of the 2005 IEEE/ACM International conference on Computer-aided design*, pages 383–390, Washington, DC, USA, 2005. IEEE Computer Society.

[13] Kaustav Banerjee and Navin Srivastava. Are carbon nanotubes the future of VLSI interconnections? In *DAC '06: Proceedings of the 43rd annual conference on Design automation*, pages 809–814, New York, NY, USA, 2006. ACM Press.

[14] Yehia Massoud and Arthur Nieuwoudt. Modeling and design challenges and solutions for carbon nanotube-based interconnect in future high performance integrated circuits. *J. Emerg. Technol. Comput. Syst.*, 2(3):155–196, 2006.

[15] Y. Massoud and A. Nieuwoudt. Accurate resistance modeling for carbon nanotube bundles in VLSI interconnect. In *Proceedings of the Sixth IEEE Conference on Nanotechnology*, pages 288– 291, June 2006.

[16] A. Naeemi and JD Meindl. Carbon Nanotube Interconnects. In *ISPD 07: Proceedings of the International Symposium on Physical Design*, pages 77–84, New York, NY, USA, 2007. ACM Press.

[17] A. Nieuwoudt and Y. Massoud. Understanding the Impact of Inductance in Carbon Nanotube Bundles for VLSI Interconnect Using Scalable Modeling Techniques. *IEEE Transactions on Nanotechnology*, 5:758–765, Nov 2006.

[18] A. Raychowdhury and K. Roy. A circuit model for carbon nanotube interconnects: comparative study with Cu interconnects for scaled technologies. *ICCAD 2004: IEEE/ACM International Conference on Computer Aided Design,*, pages 237–240, 2004.

[19] A. Naeemi, R. Sarvari, and JD Meindl. Performance comparison between carbon nanotube and copper interconnects for GSI. *Electron Devices Meeting, 2004. IEDM Technical Digest. IEEE International*, pages 699–702, 2004.

[20] A. Naeemi and JD Meindl. Monolayer metallic nanotube interconnects: promising candidates for short local interconnects. *Electron Device Letters, IEEE*, 26(8):544–546, 2005.

[21] Nir Magen, Avinoam Kolodny, Uri Weiser, and Nachum Shamir. Interconnect-power dissipation in a microprocessor. In *SLIP '04: Proc. 2004 international workshop on System Level Interconnect Prediction*, pages 7–13, New York, NY, USA, 2004. ACM Press.

[22] A. Naeemi and JD Meindl. Compact physical models for multiwall carbon-nanotube interconnects. *Electron Device Letters, IEEE*, 27(5):338–340, 2006.

[23] F. Kreupl, AP Graham, GS Duesberg, W. Steinhögl, M. Liebau, E. Unger, and W. Hönlein. Carbon nanotubes in interconnect applications. *Microelectronic Engineering*, 64(1-4):399–408, 2002.

[24] W. Hoenlein. New Prospects for Microelectronics: Carbon Nanotubes. *Jpn. J. Appl. Phys*, 41(6b):4370–4374, 2002.

[25] Y. AWANO. Carbon Nanotube Technologies for LSI via Interconnects. *IEICE Transactions on Electronics*, 89(11):1499, 2006.

[26] L. Liebmann and J.A. Carballo. Layout Methodology Impact of Resolution Enhancement Techniques. *Proc. Electronic Design Processes Workshop*, pages 55–61, 2003.

[27] A. Nieuwoudt and Y. Massoud. On the Impact of Process Variations for Carbon Nanotube Bundles for VLSI Interconnect. *IEEE Transactions on Electron Devices*, 54:446–455, March 2007.

[28] Guoqing Chen, Hui Chen, Mikhail Haurylau, Nicholas Nelson, Philippe M. Fauchet, Eby G. Friedman, and David Albonesi. Predictions of CMOS compatible on-chip optical interconnect. In *SLIP '05: Proceedings of the 2005 international workshop on System level interconnect prediction*, pages 13–20, New York, NY, USA, 2005. ACM Press.

[29] M. Haurylau, H. Chen, J. Zhang, G. Chen, NA Nelson, DH Albonesi, EG Friedman, and PM Fauchet. On-chip optical interconnect roadmap: challenges and critical directions. *Group IV Photonics, 2005. Second IEEE International Conference on*, pages 17–19, 2005.

[30] G. Chen, H. Chen, M. Haurylau, N.A. Nelson, D.H. Albonesi, P.M. Fauchet, and E.G. Friedman. On-Chip Copper-Based vs. Optical Interconnects: Delay Uncertainty, Latency, Power, and Bandwidth Density Comparative Predictions. In *Interconnect Technology Conference, 2006 International*, pages 39 – 41, June 2006.

[31] M.J. Kobrinsky, B.A. Block, J.F. Zheng, B.C. Barnett, E. Mohammed, M. Reshotko, F. Robertson, S. List, I. Young, and K. Cadien. On-chip optical interconnects. *Intel Technology Journal*, 8(2):129–141, 2004.