Multi-GHz Interconnect Effects in Microprocessors

Phillip J. Restle
914-945-2697
restle@us.ibm.com

Albert E. Ruehli
914-945-1592
ruehli@us.ibm.com

Steven G. Walker
914-945-1368
sgwalker@us.ibm.com

IBM T. J. Watson Research Center, P.O. Box 218 Yorktown Heights, NY 10598, USA

ABSTRACT
High frequency on-chip interconnect examples are accurately analyzed using full-wave PEEC (Partial Element Equivalent Circuit) analysis. All wire currents and voltages (or delays) are visualized using 3D animations to aid intuitive understanding of new, high frequency interconnect effects.

General Terms
Design

Keywords
Interconnect, Inductance, Clock Distribution, Full-wave Analysis, Extraction, Simulation, Circuit-Tuning.

1. INTRODUCTION
Large high-performance chip designs require the design and modeling of interconnects using thick copper metal layers with low resistance and capacitance. These fast metal layers are primarily used for power, ground, I/O wiring, and critical global signal nets such as the global clock distribution. Due to the decrease in resistance and capacitance, and the increase in signal frequencies, inductance effects have become significant for specific cases such as clock distributions and other highly optimized networks[1,2]. Technology developers, designers and CAD tool developers are still searching for ways to deal with these new effects. Unfortunately, accurate on-chip inductance extraction and simulation in the general case is much more difficult than resistance and capacitance. Even if ideal extraction and simulation tools existed, few engineers have experience designing with lossy transmission lines. This talk will attempt to demystify on-chip inductance through the discussion of several illustrative on-chip wiring examples analyzed using full-wave extraction and simulation methods. Several high-speed on-chip interconnect examples will be studied using a specialized full-wave PEEC (Partial Element Equivalent Circuit) analysis method tailored for on-chip applications [3]. In addition, a new visualization method will be employed to allow us to see and understand all the currents and voltages in these interconnect structures ranging from simple transmission lines to multi-GHz clock networks. The final example describes a new global clock network along with a new tuning method implemented for several server microprocessors.

ISPD'01, April 1-4, 2001, Sonoma, California, USA.

2. Technology
A few years ago many microprocessors incorporated on-chip power and ground planes. A primary motivation for these power and ground planes was the need to supply power to the center of the chip, using only I/O on the edge of the chip. With the more widespread use of ball-grid arrays for chip I/O, this is no longer necessary. Since reference planes on-chip are no longer necessary for power distribution, it is generally found more efficient to use some of each metal layer for VDD and VDD grids, with the wiring on adjacent levels being perpendicular. To allow both high-density local wiring and high performance global wires, microprocessors require both fine-pitch thin layers as well as much thicker low resistance Copper layers. The top levels are generally several times thicker than lower levels. The thicker levels provide improved on-chip power distribution as well as low-resistance and fast wiring for critical global signals.

3. Power and Ground Grid Design
In the past, the wire widths and periodicity used for the power grid were chosen primarily to minimize the resistive voltage drop between the power and ground pads and the devices. This optimization then resulted in a relatively coarse grid of wide wires on the top wiring levels, since this minimized the number of wire-to-wire spaces, and minimized resistance. While resistance is still an issue, it is now replaced in importance by other considerations. Both capacitive and inductive coupling have become relatively more important. One powerful method for reducing capacitive coupling on critical nets is to route longer, critical nets next to "quiet" VDD or VDD wires. This drives us to use finer VDD and VDD meshes to provide more "quiet" tracks adjacent to them. Inductive effects provide another powerful motivation to have VDD and GND wires very near critical nets and busses. In fact, it has become necessary in some cases to simultaneously design both the signal wires and nearby return-paths and shields (which are also part of the VDD/VDD grids). Thus the power grids are becoming "finer" and more customized, and the power grid design is more and more determined by capacitive and inductive coupling concerns.

4. PEEC Analysis
Effective wire capacitance extraction including reasonably accurate capacitive coupling terms can be calculated with acceptable accuracy for almost any purpose by considering only the nearest neighbor wires in every general direction. However, as exhibited in the examples, inductance effects can be significantly affected by wires as far as 100 microns away and how they are being terminated or driven. These numerous yet significant "far-mutuals" make

inductance extraction and simulation a more difficult problem than capacitance. In addition, when inductance is important, it is the transient interactions among the capacitive and inductive currents that determine the actual waveforms. Full-wave PEEC (Partial Element Equivalent Circuit) analysis is a powerful method for complete analysis of high-frequency on-chip interconnects [4]. The PEEC examples shown include the complete coupling matrix where every metal shape is coupled with capacitances and inductances to every other shape. While it is unfeasible to perform such a complete analysis on every net, full wave PEEC analysis is useful when a complete understanding of many important effects is needed for fundamental decision-making.

5. Visualization

To aid understanding of these effects, a new visualization method has been developed to help display all the currents and voltages in complex interconnect structures. These 3D animations are used to help develop physical intuition about issues such as inductive return-path design, inductive coupling, and power grid integrity. The visualization method used starts with a simple 2D layout view of the interconnects in the X-Y plane. This layout view is then distorted to show the voltages and currents everywhere at a single point in time. Figure 1 shows such a visualization of a structure containing a 1 cm long copper on-chip signal wire in a VDD and VDD grid. Note that the Y-axis is exaggerated since this layout is actually 70 times as long as it is wide. First, the Z-axis is used to represent voltage: Every point on each wire is displaced in the Z-axis a distance proportional to the voltage at that point on the wire. Next, the diameter of the wires is distorted so that the diameter of each wire is proportional to the current at each point in the wire. In figure 1, most wires shown are part of the GND grid, and have a Z coordinate of approximately 0 Volts. Two fairly straight wires near the edges of the structure are VDD wires, and thus have a Z coordinate of approximately 1.8 Volts. A single signal wire is being switched from GND to VDD at this snapshot in time. The near end has already been pulled up most of the way to 1.8 V, but the far end is still at GND. When available, color is used to show the direction of the current flow. With this visualization method in hand, we apply it to a few examples to help us understand basic on-chip transmission-line properties. First, basic effects such as reflection, overshoot, and return-path discontinuity will be illustrated. Finally, the evolving technology and design challenges for the special problem of global clock distribution will be discussed, culminating with strategies for large Multi-GHz chips.

6. Examples

In this sections, several on-chip copper interconnect structures are analyzed, visualized, and discussed.

6.1 High-Freq. Return Current Distribution

The first example consists of a signal wire, with 18 parallel GND wires, and 2 VDD wires, as well as perpendicular wires creating a power grid (see figure 1). The interconnect technology being modeled is IBM's first copper back-end process, CMOS7S. The whole structure is 1 cm long by 140 um wide. The input transition time is 40 ps. The largest current (shown by the largest diameter) is at the beginning of the signal wire (current direction is into page). Note, however, that at these high frequencies, the return current in the nearest GND wire is almost as large as the signal current. In the absence of inductance, the return current would spread out much more uniformly between all the other power wires to reduce resistance. However at these high frequencies, the return current mostly chooses the path closest to the signal wire to reduce the area of the current loop and thus the inductance.

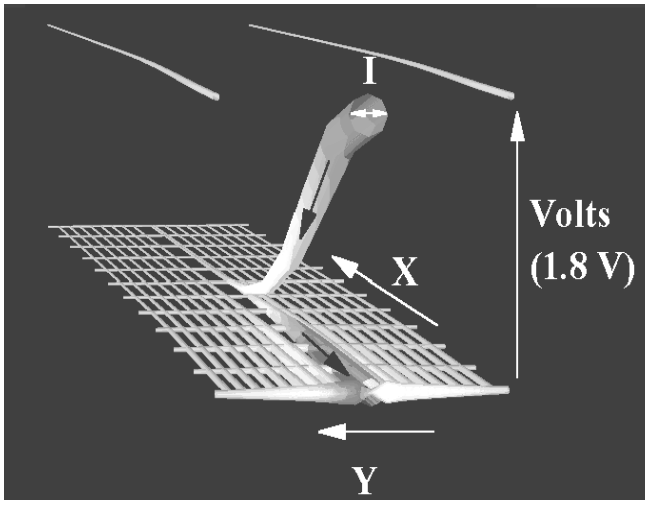

Figure 1. A 3-D visualization of currents and voltages. The Z-axis is voltage, and the wire diameter and color represent current.

6.2 Overshoot

In the figure 2, we see another new phenomenon seen when inductance is included: voltage overshoot and undershoot. In this case, the transmission line is under-terminated, so that the current "piles up" at the far end causing a temporary voltage overshoot above VDD. While this can be undesirable, resulting in problems with pass-gate circuits or gate-oxide reliability, this tendency to overshoot is also associated with a faster transition time at the far end of the transmission line. This faster transition time can increase the speed of the receiving circuit if the interconnect and receiver is designed carefully.

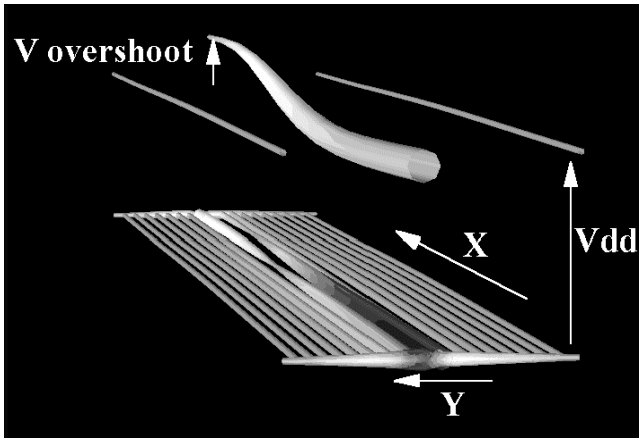

Figure 2. When wide wires are used to drive small loads, undesired voltage overshoot (>VDD) can occur due to inductance.

6.3 Wide Wire Effects

Most inductance effects are exacerbated when wide wires are used in an attempt to increase interconnect speed. The diagrams at the top of figure 3 illustrate that when wide wires are used, the average distance to the return path increases, resulting in larger inductance. At these frequencies, the current density within each wire is relatively constant, showing little skin effect. Figure 3 also shows that the use of excessively wide wires actually results in an increase in the delay, even assuming infinite driver strength, due to the increasing inductance effect.

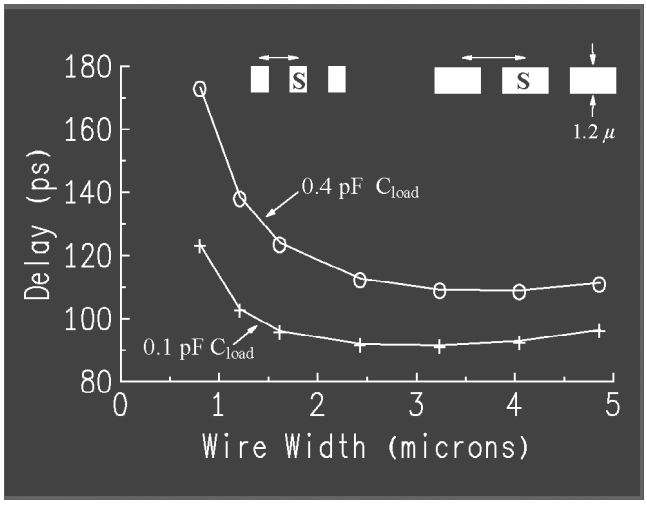

Figure 3. Increasing wire width can actually cause increased wire delay due to larger inductance from greater distances to coplanar return paths.

6.4 Discontinuous Power Grid

While the examples above can be analyzed with reasonable accuracy using a number of inductance extraction methods, some problems (such as the one presented in this section) inherently require full 3D analysis. In this example half of the GND wires nearest the signal wire were cut, representing a partial discontinuity in the power grid. Such discontinuities can occur due to mismatches between separate circuits having slightly different power grid structures. Figure 5 shows a structure very similar to figure 1, however half of the 20 GND wires parallel to the signal wire have been cut. This change results in no significant differences unless inductance is included in the analysis. When inductance is included a number of significant changes are noted. First, we see that the return current detours around the cut in the power grid, while remaining as close to the signal wire as possible. Especially interesting is the large disturbance (300 mV) seen at the ends of the cut power wires due to both capacitive and inductive coupling to these GND wires.

7. Global Clock Distributions

This section describes practical clock distribution networks presently used by industry.

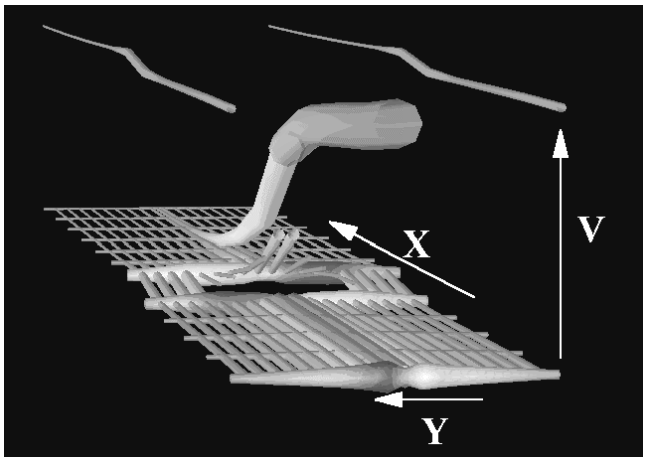

Figure 4. Current pattern due to power-grid discontinuity show the complexity of determining

7.1 Simple Clock Mesh

Perhaps the simplest global clock distribution network is a grid of wires, typically routed using two metal layers, one for the wires in the X-direction, the other for the Y direction. The intersection points are connected with vias to form a clock grid or mesh. This mesh can then be switched at the clock frequency using one or more lines of drivers. To switch at higher frequencies, the trend is to drive the mesh at two parallel edges, or all 4 edges. As clock frequencies increase, the interconnect delay between the drivers and the center of the mesh adds to the clock skew. Figure 5 shows such a mesh 5 mm on a side, operated at 2 GHz, showing the center of the mesh lagging behind the edges. At these high frequencies the center of the mesh switches significantly later in the cycle, and it becomes necessary to use a larger number of smaller meshes to scale this approach to higher frequencies. The first such mesh used a single line of drivers to drive the entire chip [5]. A recent chip uses 32 separate meshes each driven by two lines of drivers [6].

Another potential drawback of this approach is that the global clock waveform inevitably varies with the distance from the driver, and the waveforms may have plateaus near the drivers, and overshoots at the center of the mesh.

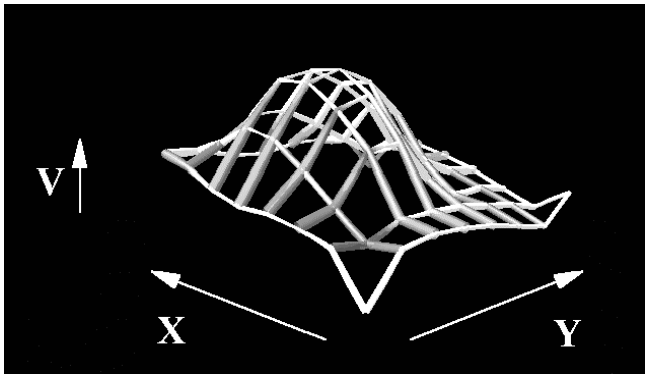

Figure 5. A 5 mm by 5 mm global clock mesh with drivers at all 4 edges shows significant skew when operated at 2 GHz.

7.2 Symmetrical H-trees

The most traditional clock distribution method is to use some form of tree networks. If the clock loads being driven were uniformly distributed, and all had the same value, it would be optimal to route symmetrical H-trees as shown in Figure 6. This would result in no skew, with identical waveforms for each load. Unfortunately, on a real chip the capacitive clock loads are far from uniform in capacitance or distribution, which make simple symmetric trees impractical.

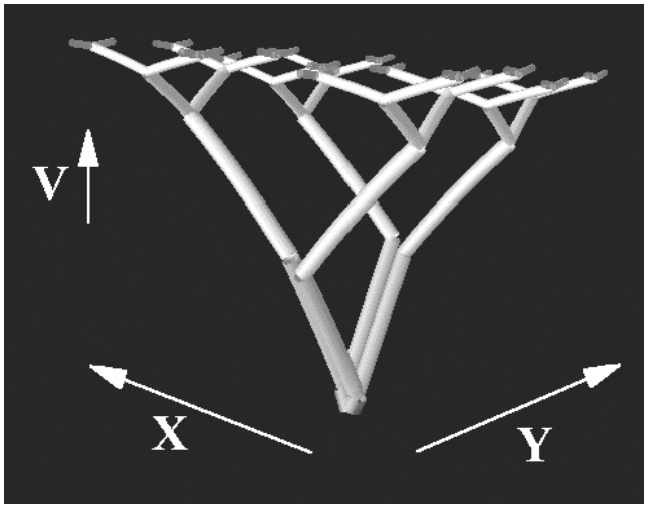

Figure 6. A symmetric tree driving a uniform load distribution.

7.3 Tunable Trees

There are a variety of algorithms for routing, buffering, and tuning clock trees to achieve low skew for asymmetric problems, or in fact produce any desired skew. One strategy is to route simple approximately symmetric trees before detailed clock loads are known. When the clock loads were determined, the tree is simulated to determine if the skew is acceptible. Figure 7 shows another type of visualization of a tree having non-uniform loads. In this visualization, the Z-axis represents delay rather than voltage. Figure 8 shows the result of a wire-width-tuning process used to reduce the skew without the need for re-routing. While this tuning process can result in low skew without re-routing, it is more susceptible to model errors or process variations than the simpler grid approach, which can at least guarantee low skew for nearby points on the grid. Note the one wire on the right that appears to have negative delay. Due to impedance mismatches and reflections it is possible for wires to have negative delay due to transmission line effects. The delays are defined using the VDD/2 crossing times.

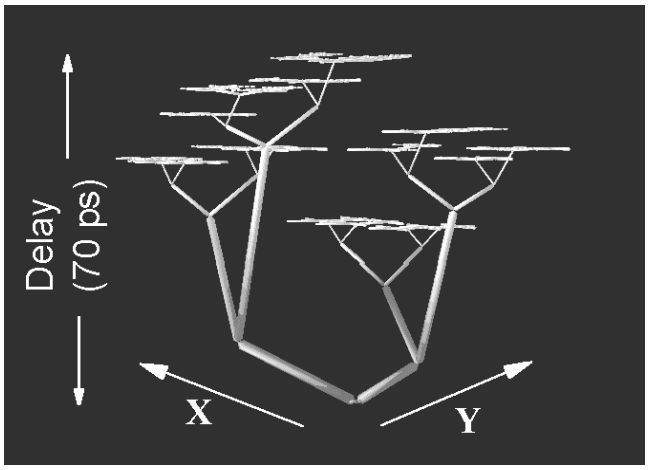

Figure 7. A tree driving a non-uniform load distribution where significant skew is observed between different branch end-points. Unlike previous figures, the Z axis represents delay rather than voltage, and the wire width shown represents the designed wire width, rather than current.

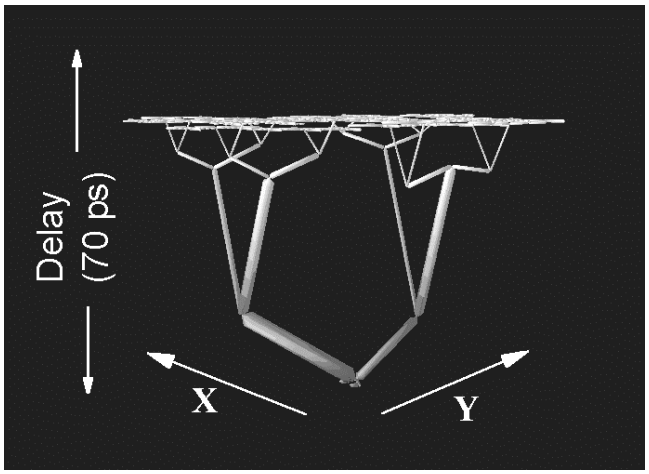

Figure 8. This shows the tree in figure 7 after wire-width tuning to achieve lower skew.

7.4 Tunable Grid-Tree

The final global clock distribution strategy is a combination of tunable trees all driving a grid or mesh that combines many of the advantages of both grid and tree strategies [7]. Figure 9 shows one quarter of the network for a Power4 GHz microprocessor [8]. There are 16 buffers in the network shown in figure 9, and these appear as vertical lines because they have delay, but do not transport the signal in the X-Y plane. The grid guarantees low skew between nearby points, while the tunable trees can provide low skew and similar waveforms to all locations on the grid. Tuning a large single interconnect network such as this is difficult because a single full simulation can take several hours of CPU time. To speed the tuning, a method was developed that involves approximating the actual grid-tree by a network with the grid loops cut, and using diffused loads instead of the actual loads. This approximation is then divided up into 64 smaller tuning jobs (one for each sector buffer) that

execute in parallel on 64 processors in less than 30 minutes. This method has been used on a number of microprocessors, yielding skews measured in hardware as low as 20 ps [7].

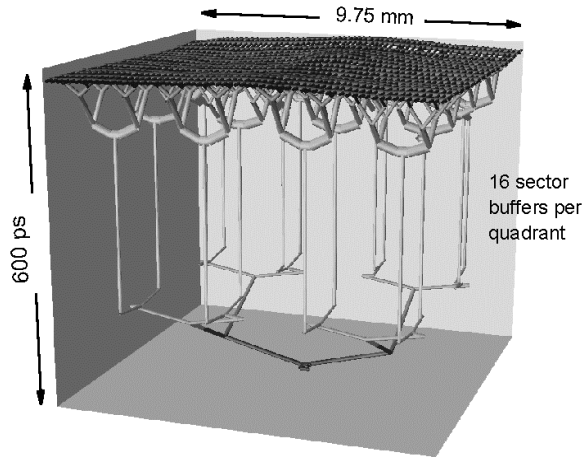

Figure 9. An optimized combination of tunable trees driving a single grid that has been used for several server micro-processors.

8. CONCLUSION

For many design purposes, only rough estimates of self and mutual inductance are sufficient to help guide routine design decisions. However, it is unwise to use questionable approximations when making decisions such as whether to include on-chip ground planes, and when developing optimal interconnect cross-sections and basic design guidelines. For these purposes, a complete and accurate interconnect extraction and simulation (such as those provided by full-wave PEEC), followed by clear understanding of the results, are needed to correctly guide technology and design choices. In addition, for critical networks such as global clock distribution, PEEC analysis can be applied directly to significant parts of the network to provide an accurate baseline for the faster approximations typically needed to design, tune, and simulate the full product network.

The number and variety of methods used for global clock distribution will continue to expand as we solve the challenging prospect of large Multi-GHz processors.

9. ACKNOWLEDGMENTS

We thank many IBM employees including Alina Deutsch, Peter Cook, Keith Jenkins, Barry Rubin, J. Scott Neely, James Warnock, Thomas Bucelot, David Boestler, Charles Alpert, and Michael Rosenfield of the IBM Research Division for many valuable interactions. The colleagues that turned these ideas into products for the IBM Enterprise Server Group include Timothy McNamara, David Webber, Peter Camporese, Kwok Eng, Howard Smith, and Dale Hoffman in Poughkeepsie N. Y., David Allen, Mike Rohn, Mike Quaranta of Rochester Minn., and Craig Carter, Byron Krauter, Brian Veraa, Steve Weitzel, John Petrovic, Dave Boerstler, Charles Alpert, Brad McCredie, and Carl Anderson of Austin Texas.

10. REFERENCES

[1] A. Deutsch, G. V. Kopcsay, P. Restle, *et al*, "When are Transmission-Line Effects Important for On Chip Interconnections?" *IEEE Trans. Microwave Theory Tech.* (USA) Vol. 45, No. 10, pt. 2, pp. 1836-46, Oct. 1997.

[2] P. J. Restle, K. A. Jenkins, A. Deutsch and P. W. Cook, "Measurement and Modeling of On-Chip Transmission-Line Effects in a 400 MHz Microprocessor," *IEEE Journal of Solid-State Circuits*, Vol. 33 No. 4, pp. 662-665, Apr. 1998.

[3] P. J. Restle, A. E. Ruehli, S. G. Walker, "Dealing with Inductance in High-Speed Chip Design," In *Proc. of the Design Automation Conf.*, pp 904-909, June 1999, New Orleans, LA.

[4] P. J. Restle, A. E. Ruehli, S. G. Walker, G. Papadopoulos, "Full-Wave PEEC Time-Domain Method for the Modeling of On-Chip Interconnects," *IEEE Trans. on Computer-Aided Design*, (in press).

[5] D. Dobberpuhl et al., "A 200 MHz 64-bit Dual Issue CMOS Microprocessor," *IEEE J. Solid-State Circuits*, Vol. 27, No. 11, 1992, pp. 1,555-1,567.

[6] S. Rusu, S. Tam, "Clock Generation and Distribution for a IA-64 Microprocessor," *IEEE ISSCC Tech. Dig.* pp. 176-7, Feb. 2000.

[7] P. J. Restle *et al*, "A Clock Distribution Network for Microprocessors," *Symposium on VLSI Circuits Digest of Technical Papers*, June 2000, pp. 184-187, Honolulu, HI

[8] C. Anderson *et al*, "Physical Design of a Fourth-Generation POWER GHz Microprocessor," *IEEE ISSCC Tech. Dig.* pp. 232-233, Feb. 2001.