

## Designing the Best Clock Distribution Network

Phillip J. Restle, Alina Deutsch

IBM T. J. Watson Research Center  
Yorktown Heights, NY

### Abstract

Clock distribution has become an increasingly challenging problem for VLSI designs, consuming an increasing fraction of resources such as wiring, power, and design time. Unwanted differences or uncertainties in clock network delays degrade performance or cause functional errors. Three dramatically different strategies being used in the VLSI industry to address these challenges are compared. Novel modeling and measurement techniques are used to investigate on-chip transmission-line effects that are important for high performance clock distribution networks.

### Introduction

The rapid increase in clock frequency and transistor count poses many problems for VLSI designers. One challenge receiving increased attention is the seemingly simple task of supplying one or more clock signals to all the latches and clocked dynamic gates on a chip.

Interconnects have become more important due to fundamental scaling effects that increase the fraction of the chip cycle time that is consumed by interconnect parasitic capacitances, delays, and coupling effects. Significant effort is devoted to modeling and reducing these delays through design optimization and technology improvements. Using these design and technology advances with careful micro-architecture and chip organization, it is still possible for experienced design teams to keep interconnect delays a small fraction of critical path delays. Clock distribution, however, is unique in that the total delay of the clock network is already dominated by interconnects. Buffer delays and capacitance have scaled to keep up with cycle time, but interconnect delay from the center of the chip to corners has not, while the number of clocked gates continues to grow. Thus clock distribution problems lead to unique modeling and design techniques as well as technology advances.

While the delay of a clock distribution network is relatively unimportant, any modeling error or uncertainty in the clock signal arrival times between key points in the clock distribution can cause performance or functional problems.

To reduce model, process, and noise induced clock distribution uncertainties, the total delay through the clock distribution is in general minimized. This leads to the use of long, wide wires placed on the lowest resistance wiring levels, that are driven with fast transition times, which in turn leads to significant transmission-line effects [1]. Uncontrolled transmission-line effects are a growing source of uncertainty and clock skew as will be discussed below.

### Network Topologies

Most high performance microprocessors distribute a single performance-critical clock signal to many locations on the chip, although local regions may be gated for power management. The different local clock phases needed for various latch circuits, arrays, or dynamic logic are then generated locally from this global clock signal. There is a wide variety of clock distribution network topologies now being used for global clock distribution.

Simplified electrical models very roughly inspired by three commercial microprocessor clock network topologies were studied for illustrative purposes: *grids* like the DEC 21264 [2], *trees* like the IBM S/390 [3,4], and length matched *serpentes* like the Intel P6 [5]. The goal of the comparison is to understand the advantages and disadvantages of these very different topologies, (without reproducing many important details) by creating simple simulation models of each topology. The effects of wiring technology will be simulated, considering Al and Cu wires, with and without on-chip dedicated reference layers. The implications of non-ideal real-world cases will be discussed, followed by design and measurements of a 400MHz product.

All three microprocessors use tree-like networks driving roughly 16 buffer or spine locations for the first, longest wires in the global clock distribution. This consensus occurs because perfectly symmetric H-trees driving identical loads result in zero nominal skew. The very different networks driven by these buffers will be the focus of this paper.

Each of the simplified topologies is designed to drive only one quadrant of a 17 x 17 mm chip containing 150 pF of uniformly distributed load in each quadrant. In figures showing physical wiring, all wire widths are drawn 10X wider to allow better visual comparison of wire widths.

### Grids

The DEC Alpha series of chips uses grid-based clock distributions driven by one or more lines of buffers. This robust topology guarantees very low skew in any local region, and can be routed early in the design. Fig. 1 shows a grid based network for one quadrant of the DEC 21264, driving 150 pF of gate load, using 350 pF of grid wire capacitance. The number of grid wires was chosen arbitrarily, then a wire width of 14  $\mu\text{m}$  approximately reproduced the published wire capacitance [2]. The DEC process includes reference planes above and below the two planes used for the clock grid wiring [2]. Fig. 2 shows simulated waveforms for this grid. Fig. 3 shows the expected increase in transmission-line effects if the Al reference-layers were instead used as standard orthogonal




Figure 1: Clock grid for a chip quadrant, crossed boxes show drivers, circles show locations of simulated waveforms.

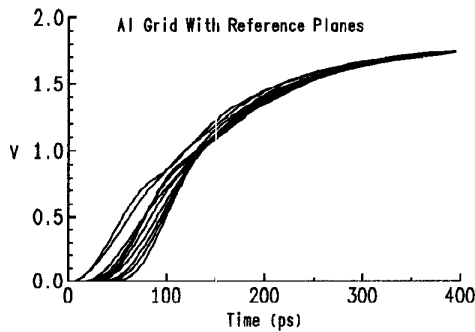


Figure 2: Simulations of clock grid of Fig. 1 using Al technology with reference planes.




Figure 3: Simulations of grid of Fig. 1, but without reference planes for low inductance return paths.

wiring levels, and the Vdd/Gnd return conductors were instead routed on both sides of each grid wire. Due to the increase in the average distance to the return current path when the reference planes are removed, the inductance increases, and the transmission-line effects such as plateaus and reflections become more pronounced, increasing the clock skew. These transmission-line effects can be reduced by using a finer grid having a larger number of narrower wires, but this increases total capacitance, wire delay, and skew. A modified grid was also simulated assuming a Cu wiring technology [6], that does not use dedicated reference planes (Fig. 4). For the Cu technology twice as many 4.5




Figure 4: Simulation of clock grid redesigned for copper grid with twice as many narrower wires, and reduced driver




Figure 5: Tree topology driving same loads as grid above.




Figure 6: Simulation of copper tree in Fig. 5.

$\mu\text{m}$  wide grid wires were used, with reduced driver sizes. The thinner, narrower, Cu wires results in reduced power ( $C_{\text{wire}}$  reduced from 350 pF to 228 pF) even with the finer grid, and exhibited no need for dedicated reference planes.

### Trees

Fig. 5 shows a symmetric H-tree designed to drive the same 150 pF loads as the grid in the previous section, with the maximum wire width chosen to be 14  $\mu\text{m}$  the same width used for the Al grid but (as shown) wire widths were optimized for minimum delay. Due to assumed idealized symmetry of the loads, the simulated skew is trivially zero. Fig. 6 shows significant transmission-line effects at various internal nodes within the trees, but only the smooth signals at the ends of the trees are relevant. Any overshoot can




Figure 7: Serpentine wires, driven by a line of clock drivers at the left edge. Only 32 of the 256 serpentine wires needed to drive the 256 loads of grid and tree examples are shown.

easily be controlled by reducing wire widths, and may be desirable as it is accompanied by faster transition times.

### Length-Matched Serpentine

Fig 7 shows another topology where each load is driven by a single point-to-point wire, and lengths are matched using a serpentine structure. To achieve delays and transition times similar to the trees, wire widths were chosen to be  $1.6 \mu\text{m}$ . This topology is relatively simple to design, and like symmetric trees, has trivially zero skew for identical loads, as long as coupling and uncontrolled variables are insignificant.

### Comparisons

Table 1 compares the three topologies assuming uniform load distribution and no environmental or process variations for the Cu wiring simulations (with no reference planes).

**Table 1: Uniform Load Distribution**

	$C_{\text{wire}}$	Delay	Skew
Grid	228 pF	21 ps	21 ps
Trees	15.5 pF	130 ps	0 ps
Serpentines	480 pF	130 ps	0 ps

For these symmetric cases, the tree topology provides low skew and much lower capacitance than the other topologies, although it requires placement of buffers at four locations internal to each chip quadrant. For real designs, a number of complications arise that further differentiate the topologies. First, actual capacitance of individual clock pins can vary from a few fF to a few pF at each pin. In addition, the loads are distributed non-uniformly over the chip, and across-chip process and power supply variations can be significant. One major advantage of the grid topology is that even very non-uniform load distributions affect the local skew very little. Thus, changes in clock loads, locations, or electrical models cause little change in clock timing, and rarely require re-tuning of the grid wires or drivers.

Although trees are potentially more efficient, wiring and tuning tree topologies to drive highly non-uniform loads with low skew can be much more difficult. Modeling errors or process variations can produce large skew even between nearby clock pins. An example is discussed below.

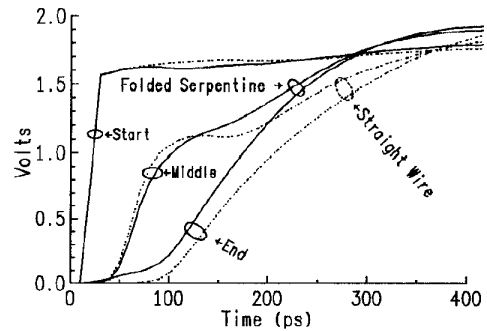


Figure 8: Simulated wave-forms at the start, middle, and end of straight vs. folded identical length wires of 1 cm length and  $5.3 \mu\text{m}$  width driving 2 pF. Grounded  $1.2 \mu\text{m}$  shield wires provide capacitance shielding but inadequate inductive shielding (with no reference planes).

The serpentine structure is simple to design for different load locations, however, the number of serpentine wires is large and changeable, contributing to wire congestion. Small loads require capacitance padding, while wide serpentine wires used to drive larger loads can lead to inductive self-coupling (Fig. 8). Since a full-chip distribution using this method requires a line of drivers on the left and right chip edges, there is potentially large skew between the left and right halves of the chip due to cross-chip variations [5].

### A 400 MHz Clock Tree Design

To take advantage of the efficiency, flexibility, low power, and potentially low skew of the tree structure, a proprietary low skew clock routing tool was developed to drive arbitrary load distributions with arbitrary wire widths while avoiding blockages. An optimization process also considers power, wiring tracks, and process variations. The tool has been used on several IBM microprocessor and ASIC designs. Measurements were made using an e-beam system [4], and backside photo-emission [7] (not shown).

Fig. 9 shows the importance of including transmission-line effects in the design of high-performance clock distributions. Although the topologies and loads were similar, a product chip designed with transmission-line effects included in the routing, extraction and tuning algorithms showed a 5 X reduction in skew compared to the test chip designed without these considerations [4].

Fig. 10 shows a unique representation of the product's clock-tree network. The clock routing tool matches delays, but for efficiency does not match lengths or loads, so subtle differences remain: Fig. 11 shows that clock pins at the ends of the longer trees exhibited more measured and simulated overshoot due to the faster signal speed. The magnitude of the overshoot is adequately modeled by the frequency independent inductance and resistance model used in the design-tools (Fig. 11), but a frequency-dependent model (not shown) including the extended wiring environment is needed to match the details of the measurement.

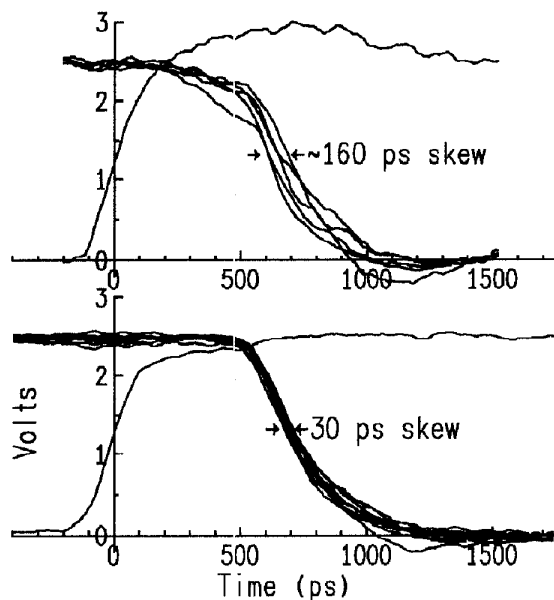


Figure 9 Hardware e-beam measurements. The top shows a test chip designed with RC wire models, the bottom shows waveforms from the 400MHz microprocessor clock network designed with transmission-line considerations. The rising signals are the output of the single central chip buffer, while the falling waveforms are 10 of 580 clock pins.




Figure 10: A 3-D (X,Y,Time) representation of the microprocessor tree network distributing the global clock to 580 local clock pins. The signal starts at the bottom, and the 9 vertical lines represent buffers. Gate loads are represented by cubes with volume proportional to capacitance.

### Conclusion

The challenges of clock distribution design continue to increase. A variety of alternatives including package wiring, and asynchronous logic can reduce the need for accurate on-chip clock distributions, however the increased cost and

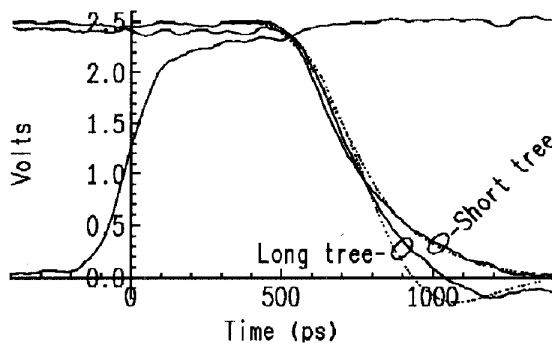


Figure 11: The measured (solid) and simulated (dashed) waveforms from the shortest and longest wiring trees shown in figs. 9 and 10.

complexity of these methods will limit their widespread use as long as strategies and technologies to design efficient synchronous clocking continue to succeed. The sophisticated interconnect modeling techniques (previously used mainly for chip-to-chip interconnects) and design tools required to design clock networks are also becoming important for other more numerous medium and long on-chip wiring. On chip reference planes simplify extraction, but are not presently needed to control on-chip transmission-line effects.

### Acknowledgements

We thank Peter Cook, Keith Jenkins, Barry Rubin, Albert Ruehli, Steve Walker, and J. Scott Neely of the IBM Research Division, Timothy McNamara, David Webber, and Peter Camporese of the IBM S/390 Division, Joaquin Fentanes, P.T. Patel, Keith Carrig and Dave Allen of the IBM Microelectronics Division, and Lawrence Pileggi of Carnegie Mellon University for profitable interactions.

### References

- [1] A. Deutsch *et al.*, "When are Transmission-Line Effects Important for On Chip Interconnections?," *IEEE Trans. Microw. Theory Tech.* (USA) Vol. 45, No. 10, pt. 2, pp. 1836-46, Oct. 1997.
- [2] B. Gieseke *et al.*, "A 600 MHz Superscalar RISC Microprocessor With Out-of-Order Execution," *ISSCC Tech. Dig.* pp. 176-7, Feb. 1997.
- [3] C. Webb *et al.*, "A 400MHz S/390 Microprocessor," *IEEE JSSC*, vol. 32, no. 11, pp. 1665-1675, 1997.
- [4] 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 JSSC*, Vol. 33 No. 4, pp. 662-665, Apr. 1998.
- [5] G. Geannopoulos, X. Dai, "An Adaptive Digital Deskewing Circuit for Clock Distribution Networks," *IEEE SSCC Tech. Dig.* p. 400, Feb. 1998.
- [6] D. Edelstein *et al.*, "Full Copper Wiring in a Sub-0.25  $\mu\text{m}$  CMOS ULSI Technology," *IEEE IEDM Tech. Dig.* pp. 773-6, Dec. 1997.
- [7] J.A. Kash and J.C. Tsang, "Dynamic Internal Testing of CMOS Circuits Using Hot Luminescence," *IEEE Electron Device Lett.*, vol. 18, pp. 330-332, 1997.