Evaluating the Potential of Future On-Chip Clock Distribution using Optical Interconnects

Nitya Ranganathan and Norman P. Jouppi
Advanced Architecture Program
HP Laboratories Palo Alto
HPL-2007-163
October 5, 2007*

clock distribution, optics, microprocessor

Designing clock distribution networks is a big challenge for future microprocessors due to increasing frequency, power, transistor counts and process variations. As technology scales, implementing conventional clock distribution networks that meet low power and skew requirements is becoming more difficult. On the other hand, optical interconnects are being proposed as an alternative to electrical interconnects due to their speed-of-light transmission, high bandwidth and low power dissipation. We look at the clock distribution trends in major microprocessor families, possible techniques for optical clocking, and the potential of moving from electrical to optical clock distribution. We found that clock distribution is largely a power amplification problem, and since electrical power amplification is more efficient that optical power amplification, wholly electrical clock distribution technologies would be preferred.

* Internal Accession Date Only

Approved for External Publication

Evaluating the Potential of Future On-Chip Clock Distribution using Optical Interconnects

Nitya Ranganathan and Norman P. Jouppi

August 16, 2006

Abstract

Designing clock distribution networks is a big challenge for future microprocessors due to increasing frequency, power, transistor counts and process variations. As technology scales, implementing conventional clock distribution networks that meet low power and skew requirements is becoming more difficult. On the other hand, optical interconnects are being proposed as an alternative to electrical interconnects due to their speed-of-light transmission, high bandwidth and low power dissipation. We look at the clock distribution trends in major microprocessor families, possible techniques for optical clocking, and the potential of moving from electrical to optical clock distribution. We found that clock distribution is largely a power amplification problem, and since electrical power amplification is more efficient that optical power amplification, wholly electrical clock distribution technologies would be preferred.

1 Introduction

The clock is used as the synchronization signal for various state elements in a microprocessor chip. It is the most widely distributed signal on the chip. Lower frequency system clocks are typically generated using Phase Locked Loops (PLLs) and multiplied to get higher frequency core clocks. The core and system clocks are distributed across the entire length and breadth of the chip to drive various latches, flip-flops and other state elements. The desired characteristics of a good clock distribution system are low jitter, low skew, low power, and low metal resources.

Jitter refers to the fact that the same state element can receive its clock signal with different delays in different cycles. Clock skew refers to the fact that different state elements receive the clock signal at different times. The allowed desired skew is a small fraction of the cycle time (typically less than 5%). Skew is caused by unbalanced trees, temperature variations, process variations, on die voltage variations, etc.

Maximum clock frequencies are currently going through a period of little change, even though the speed of the underlying transistors has gotten much faster. For example, the fastest shipping microprocessor clock speed has remained constant at around 3GHz over the last three years. This is expected, as by the year 2000 high-end microprocessors had become very heavily pipelined, which is not as power efficient as less deeply pipelined machines. It is expected that the number of pipestages will decrease

keeping the clock rate comparatively constant through at least 2010, at which point the maximum clock frequency should increase again[29]. This means that the impact of clock skew in proportion to device speed is becoming relatively less important.

Another design trend caused by power limitations is a trend to simpler cores and chip multiprocessors[29]. This has two beneficial effects. First, the use of simpler cores means that there are fewer state elements per core, reducing the clock load per core. Second, the use of chip multiprocessors means that clock skew is only important within individual cores and between the cores at their interfaces, simplifying the problem somewhat.

The amount of power spent in clock distribution can be divided into 2 parts: power spent in transmitting the clock signal from the PLL's through the global and intermediate clock wiring, and the power spent in driving the state elements (local distribution). The power due to transmission depends on the distribution technique employed. The power dissipated in driving the latches depends on several factors including transistor count, type of latches (latch capacitances) and transistor sizes.

Typical network topologies used for clock distribution are described below.

- Grid - Grid style networks have the lowest skew because of uniformity but they require lot of metal resources and dissipate a lot of power.
- Tree - H-trees, X-trees and binary trees are commonly used for distribution. Trees are attractive because of their low power and low metal usage. Balanced H-trees have very low skew, though they are very difficult to construct. Unbalanced or poorly balanced H-trees show very high skew.
- Spines - Spines are high power clock distribution lines running across the chip from which different local zones in the chip draw their own clocks. They dissipate lot of power but are useful in reducing skew if many spines are spread throughout the chip.
- Hybrid - Hybrid clock distribution systems combine one or more of the above techniques. They may use, for example, a few spines at different parts of the chip and local grids that are driven from the region in the spine closest to the grid. Similarly, hybrids of global H-trees and local grids are also possible. Some hybrid systems have different parts of the chip clocked using different distribution techniques. Most of the current processors use hybrid techniques as they simplify skew reduction and have relatively low power dissipation.

Current generation processors use one or more hybrid global/local networks for clock distribution. The local clock distribution typically consumes significantly more power than the global clock lines - for example the Intel Motecito spends 80% of its power in local distribution [33].

Optical interconnects are a promising future technology for speed-of-light communication with low power, low interference and high communication bandwidth. In the future, interconnects between chips, between cores on a chip and within components on a processor core could be made optical to achieve lower power and higher performance. In this document, we look at the potential of replacing on-chip electrical clock

interconnects with photonics technology. Clocking does not require higher bandwidth but low skew and power are very important in clock distribution.

In the following sections we discuss various distribution techniques that have been used in several major microprocessors in CMOS technology to show how the clock distribution has evolved to accommodate faster clock frequencies, increasing transistor counts and high core power. We then discuss how photonics can be used in clocking and whether there are any significant benefits to migrating from electrical to optical interconnects for clock transmission.

2 Conventional Clock Distribution Systems (Electrical)

The clock signal has to be routed to more area in the chip than any other interconnect. Also, it drives the maximum number of transistors (latches, flip-flops etc.). There have been wide variations in the observed area and power metrics for clock distribution depending on the technology, type of distribution and other factors. In general, the fraction of power used in clock transmission is lower than the total power required for clock distribution.

Clock distribution networks were studied by Restle and others [44] to estimate the best possible network to reduce both skew and power. They discuss the importance of reducing local skew within each local network. Clock skew variations across different local networks will affect performance if a path from a device in one network to a device in another is present. They also discuss the importance the clock gating in reducing local distribution power (which is significantly higher than the global network since lot of buffering and amplification is done in the local networks).

Various clock distribution architectures have also been studied by Yeh et al [53]. They study grid, tree and hybrid networks using several industry processor designs. They conclude that grid style distribution is more effective than trees to obtain small skew. Among grids, finer grids are more effective than coarser grids. They also estimate the extra power spent in grids compared to trees could be up to 30%. They evaluate hybrid architectures and find that a global tree driving several local grids (with the capability of gating each of those grids whenever possible) achieves significant power savings while not sacrificing skew.

A survey of clock distribution techniques seen in several major microprocessors starting from the 1990's is shown in Table 1. On-chip interconnects in processors were originally constructed using aluminum but have more recently been fabricated with copper, since copper has less resistance than aluminum. One thing to note is that the power consumed by the clock network in the early Alpha processors was as high as 40% of the total chip power, but in more modern processors the the power required by the clock has been reduced to as low as 25% of the total chip power.

Skew trends are shown in Figure 1. The outlier in the graph is the Alpha EV7 network clock, which did not have as tight skew requirements as the EV7 core clock. Although skew as a percentage of clock cycle time varies widely from implementation to implementation, the skew of the most recent crop of processors (UltraSPARCIV+, Cell, and the Itanium Montecito) remains well under 5% of the clock cycle time. This is true even though the clock cycle time measured in both picoseconds and logic depth

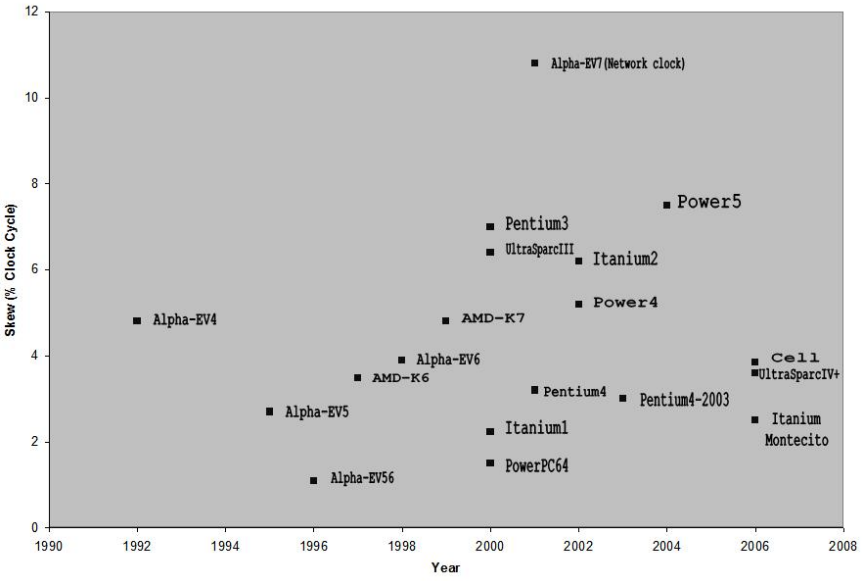

Figure 1: Clock skew as a percentage of cycle time for different processors by year.

per pipestage as measured in fan-out of 4 (FO4) has decreased dramatically over the past 16 years.

The DEC Alpha line of processors (EV4-EV7) used grid style networks which helped reduce the skew. The grids were driven by one or more driver panes which were high power regions. Earlier processors (EV4, EV5) had fewer panes and this resulted in thermal hotspots in the pane region. In the later processors more drivers were added to uniformly distribute the power and temperature.

Unlike the earlier Alpha processors, current processors do not dissipate a lot of power in the global network or the main clock driver. For example, in the Itanium 2 processor [1], 30% of the total chip power is used in the clock distribution: 4% in global and intermediate stages and 26% in local distribution and driving latches. Other current microprocessors require much more power in the local clock distribution compared to the global and intermediate clock transmission. One commonly employed technique to reduce the active power in the local distribution network is clock gating (conditional clocking).

A more detailed discussion of electrical clock distribution networks in high performance microprocessors, along with comparison of power and skew across processors appears in Appendix A.

Processor	Year	Trans. Count (M)	Feature size (mm)	Die size (mm^2)	Metal Layers	Volt (V)	Power (W)	Freq (MHz)	Clock Distribution	Clock Skew (ps)	Skew %Cycle	Clock Power	Clock Load (pF)	Jitter PLL (ps)	Ref
Alpha EV4	1992	1.68	750	233.5	3	3.3	30	200	1 driver pane driving grid	240	4.8	40%	3250		[13, 12]
Alpha EV5	1995	9.3	500	298.7	4	3.3	50	300	2 panes driving grid	90	2.7	40%	3750		[6, 3]
Alpha EV56	1996	9.6	350	208.8	4	3.3/2	25	433	2 panes driving grid	25	1.1	40%			[23]
Alpha EV6	1998	15.2	350	317.7	6	2.2	72	600	16 panes, X,H-trees, multiple grids	65	3.9	32%	2800	16	[18, 16, 21, 2, 50]
Alpha EV7	2001	152	180	396.7	7	1.5	125	1200	4 clocks, multiple trees and grids	90	10.8				[52, 28]
Pentium 3	2000	28	180		6			1000	3 spines, tree, grid	70	7.0				[22]
Pentium 4	2001	42	180		6			2000	3 spines, tree, grid	16	3.2			35	[32, 31]
Pentium 4	2003		90		7	1.2		3000	spines, tree, grid	10	3.0				[5]
Itanium 1	2000	25.4	180		6			800	global H-tree, regional grid, local	28	2.2				[47, 49]
Itanium 2	2002	221	180	421	6	1.5	130	1000	global H-tree, local tree	62	6.2	30%			[37, 38, 1]
Montecito	2006	1700	90	596	7	1.2	100	2500	H-tree,H-tree	10	2.5	25%			[17, 33, 39]
PPC64	2000	19	120		6	1.8		1000	binarytree,grid	14.9	1.5				[27]
Power4	2002	174	180	400	7	1.5	125	1300	tree, 1 grid	40	5.2			35	[45, 51]
Power5	2004	276	130	389	8	1.3		1500	tree	50	7.5				[9, 10, 46]
CELL	2006	234	90	221	8	1.2	<60	3200	3 grids	12	3.8			12.7	[42, 43]
Ultra SPARCIII	2000	23	150	244	7	1.5	60	800	grid, grid	80	6.4			62	[25]
Ultra SPARCIV+	2006	295	90	336	9	1.1	90	1800	grid, grid	20	3.6				[24]
AMD K6	1997	8.8	350	155.6	5	3.2		233	grid	150	3.5			67.2	[14, 15]
AMD K7	1999	22	250	184	6	2.5		600	btree,lines	96	4.8				[26, 19]

Table 1: Clock distribution in various microprocessors. Data for some processors has not been published.

3 Opportunities for Photonics in Clock Distribution

Optical Interconnects for clock distribution were first studied by Goodman et al [20]. They postulate that interconnect delays will be the limiting factor for performance in future MOS circuits and suggest moving to optical and electro-optic technologies. With near-infrared optical sources, modulators and detectors with media such as free space, optical fibers and integrated optical waveguides, they state the interconnect scaling problem could be solved.

They list five advantages in moving to photonics: freedom from capacitive loading effects which allows greater fan-in and fan-out, immunity to mutual interference effects, lack of planar constraints resulting in reduced cross-coupling for criss-crossing waveguides, reconfigurability of free space focused interconnects and possibility of direct injection of optical signals into electronic devices without the need for optical-to-electrical conversion.

We consider potential optical technologies for global clock distribution, and after opto-electronic conversion electrical signals would be distributed locally. Four different ways of optical clocking are described as listed below.

- Index-based with waveguides - light is carried from a single source generating the optical clock signal to the other parts of the chip using waveguides which are integrated on a suitable substrate.
- Index-based with fiber optics - light is transmitted similar to the above case except that fibers are used instead of waveguides in a separate core.
- Unfocused free space interconnect - optical clock signal broadcast to the entire chip by focusing light (through a lens or diffusers) perpendicular to the chip from above.
- Focused free space interconnect - an optical element like a hologram (which acts as a grating) sends the optical signals onto a multitude of detection sites simultaneously.

Miller et al [34, 35] discuss various opportunities for optics in on-chip interconnects and notes that one of the main reasons to move away from long electrical on-chip wires is to avoid growing transmission line and inductance effects. Power requirements of long point-to-point optical interconnects are much lower (with no need for repeaters for optical clock transmission), and optical signals, whether long or short, perform equally well. Other benefits include the ability to send signals across in the 3rd dimension and the possibility of integration with electronic devices. Debaes et al [11] and Bhatnagar [4] discuss the advantages of optical clocking and propose a receiver-less optical clocking scheme which reduces the latency of transmitting the clock signal to the local network. They conclude that such a latency-reducing scheme reduces skew and jitter but does not reduce the power consumed compared to an electrical clocking scheme.

Recent studies comparing electrical and optical clocking schemes have analyzed the power, skew and area usage for different technologies and estimated the potential

of photonics in clock distribution [8, 30]. Based on the 2001 ITRS roadmap, using analytical models, they estimate that most of the power dissipation is associated with local clock distribution. Since alternatives to electrical clock distribution have been proposed for replacing the global clock distribution only, they conclude that there are no significant power benefits to replacing electrical distribution with optical distribution. The paper also shows that low skew can be obtained with optical as well as non-scaled electrical interconnects (130nm) and concludes that skew benefits of optical transmission are also not significant if non-scaled electrical interconnects are used. They consider a global balanced and buffered H-Tree driving local grids as representative of current clock distribution techniques. For optical distribution, they replace the electrical H-tree by an H-tree structure built using waveguides with detectors at the end points to convert the light pulses to clock signal. The local grids and buffers driving the local grids remain the same in both implementations.

Mule et al [36] discuss the pros and cons of electrical and optical clock distribution systems. Among the different schemes used to transmit the optical clock onto different parts of the chip, they find the waveguide based approach most feasible since free-space approaches work in 3 dimensions (and hence are not compact) and this would complicate power distribution and cooling. To accomplish local clock routing with optical technology, the fanout should correspond to the total number of latches on the chip (which can be $>$ 100,000). Since optical signals do not use repeaters and rely on the original source strength to drive all the loads, it is extremely difficult to make the fanout more than a few tens or hundreds of loads at the most. Hence for the regions which require short wires like the local distribution regions, current systems can only use electrical routing.

A possible implementation of optical clocking is to replace the global transmission network with optical waveguides and use electrical local distribution. At the end of the global transmission network, when the global clock signal is converted to electrical signals using a detector, it has to be buffered and amplified and sent to the local network. However, given expected optical technologies, the number of fanouts on the global optical transmission must be relatively small (less than 100). This means that the electrical fanout driven by each optical receiver will be quite large (thousands of latches). In order to make the optical receiver fast and to support large optical fanouts, each optical receiver must be physically small and have a very small capacitance. However, taking a signal on a very small capacitance and buffering it to drive a very large capacitance is a classical *Logical Effort* problem [48]. This will require many stages of buffering, which introduces additional skew due to different process, thermal, and voltage conditions between different local network buffers. In contrast, a similar sized local distribution network driven from an electrical input will be connected to a global electrical transmission spine with very large capacitance due to wire parasitics. Thus the additional delay caused by connecting large electrical buffers to the global spine is modest. Hence a smaller number of buffer stages are needed to drive a local clock distribution network given an electrical input in comparison to an optical input. This could give local distribution buffers with electrical inputs lower skew than local distribution buffers with optical inputs.

Because of the extremely large fanouts of the clock, high capacitive loads, and extremely low effective load impedance, the clock distribution problem is primarily a

problem in power amplification. Electrical power amplification technologies are more power and skew efficient than expected optical technologies. For example, since the CMOS buffers act as non-linear amplifiers, their power added efficiency can easily be over 60%.

4 Other Clock Distribution Technologies

There have been several alternative proposals to improve clock distribution in terms of power and performance. Some of them are briefly described below.

- Wireless clocking with on-chip antennas - Clock distribution with on-chip antennas and RF waves has been explored and the construction of global clocking using wireless network is examined. RF technology does not bring in power savings if not used at the local distribution also [40, 8].
- High frequency clock network with standing waves - Standing waves (to reduce skew) and coupled oscillators (for phase averaging to reduce skew and jitter) are used to generate and transmit high frequency clocks at the global and intermediate levels [41].
- 3D clocking - 3-dimensional architectures are being explored as an alternative to conventional chips to reduce die size and hence long wire delays between various components on the chip. By routing wires across the 3rd dimension from one chip to the other bonded chip, clock wires become shorter and hence have reduced skew [34, 8].
- Resonant clocking - Conventional clock distribution is augmented with inductors to resonate the clock capacitance. Use of fewer buffers results in reduced power and jitter [7].
- Other approaches - Superconductivity to increase conductance of wires on the chip and improve latency, increasing the number of metal layers to shorten wire delays are among other approaches to improve clocking [34].

5 Conclusion

Expected future optical technologies appear capable of only replacing global and intermediate clock transmission. However, less power is dissipated in transmission for global and intermediate networks compared to local networks. Thus any use of optical technologies, even if they required no power, would only eliminate a minority of the clock distribution power. A further complication of such an approach would be that the local electrical clock distribution in a hybrid optical/electrical system would be required to have more stages of electrical buffers, introducing more skew. To the extent that clock distribution is a power amplification problem, and electrical power amplification is more efficient that optical power amplification, wholly electrical clock distribution technologies would be preferred.

References

[1] F. E. Anderson, J. S. Wells, and E. Z. Berta. The Core Clock System on the Next Generation Itanium Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2002.

[2] D. W. Bailey and B. J. Benschneider. Clocking Design and Analysis for a 600-MHz Alpha Microprocessor. *IEEE Journal of Solid-State Circuits*, 33(11), November 1998.

[3] B. J. Benschneider et al. A 300-MHz 64-b Quad-Issue CMOS RISC Microprocessor. *IEEE Journal of Solid-State Circuits*, 30(11), November 1995.

[4] A. Bhatnagar. Low Jitter Clocking of CMOS Electronics Using Mode-locked Lasers. Phd Dissertation, Department of Electrical Engineering, Stanford University, 2005.

[5] N. Bindal, T. Kelly, N. Velastegui, and K. L. Wong. Scalable Sub-10ps Skew Global Clock Distribution for a 90nm Multi-GHz IA Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2003.

[6] W. J. Bowhill et al. A 300MHz 64b Quad-Issue CMOS RISC microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 1995.

[7] S. C. Chan, K. L. Shepard, and P. J. Restle. Uniform-Phase Uniform-Amplitude Resonant-Load Global Clock Distributions. *IEEE Journal of Solid-State Circuits*, 40(1), January 2005.

[8] K.-N. Chen, M. J. Kobrinsky, B. C. Barnett, and R. Reif. Comparisons of Conventional, 3-D, Optical, and RF Interconnects for On-Chip Clock Distribution. *IEEE Transactions on Electron Devices*, 51(02), February 2004.

[9] J. Clabes et al. Design and Implementation of the POWER5 Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2004.

[10] J. Clabes et al. Design and Implementation of the POWER5 Microprocessor. In *Proceedings of the Design Automation Conference (DAC)*, 2004.

[11] C. Debaes et al. Receiver-Less Optical Clock Injection for Clock Distribution Networks. *IEEE Journal of Selected Topics in Quantum Electronics*, 9(2), March/April 2003.

[12] D. Dobberpuhl et al. A 200-MHz 64-bit Dual-issue CMOS Microprocessor. *Digital Technology Journal*, 4(4), 1992.

[13] D. Dobberpuhl et al. A 200MHz 64b Dual-Issue CMOS Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 1992.

[14] D. A. Draper et al. An x86 Microprocessor with Multimedia Extensions. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 1997.

[15] D. A. Draper et al. Circuit Techniques in a 266-MHz MMX-Enabled Processor. *IEEE Journal of Solid-State Circuits*, 32(11), November 1997.

[16] H. Fair and D. Bailey. Clocking Design and Analysis for a 6OOMHz Alpha Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 1998.

[17] T. Fischer, F. Anderson, B. Patella, and S. Naffziger. A 90nm Variable-Frequency Clock System for a Power-Managed Itanium-Family Processor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2005.

[18] B. A. Gieseke et al. A 600MHz Superscalar RlSC Microprocessor with Out-0f-Order Execution. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 1997.

[19] M. Golden et al. A Seventh-Generation x86 Microprocessor. *IEEE Journal of Solid-State Circuits*, 34(11), November 1999.

[20] J. W. Goodman, F. Leonberger, S.-Y. Kung, and R. A. Athale. Optical Interconnections for VLSI Systems. *Proceedings of the IEEE*, 72(7):850 – 866, July 1984.

[21] M. K. Gowan, L. L. Biro, and D. B. Jackson. Power Considerations in the Design of the Alpha 21264 Microprocessor. In *Proceedings of the 35th Design Automation Conference (DAC)*, 1998.

[22] P. K. Green. A GHz IA-32 Architecture Microprocessor Implemented on 0.18um Technology with Aluminum Interconnect. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2000.

[23] P. Gronowski et al. A 433-MHz 64-b Quad-Issue RISC Microprocessor. *IEEE Journal of Solid-State Circuits*, 31(11), November 1996.

[24] J. M. Hart et al. Implementation of a Fourth-Generation 1.8-GHz Dual-Core SPARC V9 Microprocessor. *IEEE Journal of Solid-State Circuits*, 41(1), January 2006.

[25] R. Heald et al. Implementation of a 3rd-Generation SPARC V9 64b Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2000.

[26] S. Hesley et al. A 7th-Generation x86 Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 1999.

[27] P. Hofstee. A 1GHz Single-Issue 64b PowerPC Processor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2000.

[28] A. K. Jain et al. A 1.2GHz Alpha Microprocessor with 44.8GB/s Chip Pin Bandwidth. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2001.

[29] N. P. Jouppi. The Future Evolution of High-Performance Microprocessors. ACM Micro-38 conference keynote, 2005. http://pcsostres.ac.upc.edu/micro38/01_keynote2.pdf.

[30] M. J. Kobrinsky et al. On-chip Optical Interconnects. *Intel Technology Journal*, 08(02), May 2004.

[31] N. A. Kurd, J. S. Barkatullah, R. O. Dizn, T. D. Fletcher, and P. D. Madland. A Multigigahertz Clocking Scheme for the Pentium 4 Microprocessor. *IEEE Journal of Solid-State Circuits*, 36(11), November 2001.

[32] N. A. Kurd, J. S. Barkatullah, R. O. Dizon, T. D. Fletcher, and P. D. Madland. Multi-GHz Clocking Scheme for Intel Pentium 4 Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2001.

[33] P. Mahoney, E. Fetzer, B. Doyle, and S. Naffziger. Clock Distribution on a Dual-Core, Multi-Threaded Itanium-Family Processor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2005.

[34] D. Miller. Rationale and Challenges for Optical Interconnects to Electronic Chips. *Proceedings of the IEEE*, 88, 2000.

[35] D. Miller, A. Bhatnagar, S. Palermo, A. Emami-Neyestanak, and M. Horowitz. Opportunities for Optics in Integrated Circuits Applications. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2005.

[36] A. V. Mule, E. N. Glytsis, T. K. Gaylord, and J. D. Meindl. Electrical and Optical Clock Distribution Networks for Gigascale Microprocessors. *IEEE Transactions on VLSI Systems*, 10(5), October 2002.

[37] S. D. Naffziger and G. Hammond. The Implementation of the Next-Generation 64b ItaniumTM Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2002.

[38] S. D. Naffziger and G. Hammond. The Implementation of the Next-Generation 64b ItaniumTM Microprocessor. Presentation at the International Solid State Circuits Conference (ISSCC), 2002. http://www.intel.com/design/itanium2/download/isscc_2002_1s.pdf.

[39] S. D. Naffziger, B. Stackhouse, and T. Grutkowski. The Implementation of a 2-core Multi-Threaded Itanium-Family Processor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2005.

[40] K. K. O et al. On-Chip Antennas in Silicon ICs and Their Application. *IEEE Transactions on Electron Devices*, 52(7), July 2005.

[41] F. O'Mahony, C. P. Yue, M. A. Horowitz, and S. S. Wong. A 10-GHz Global Clock Distribution Using Coupled Standing-Wave Oscillators. *IEEE Journal of Solid-State Circuits*, 38(11), November 2003.

[42] D. Pham et al. The Design and Implementation of a First-Generation CELL Processor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2005.

[43] D. C. Pham et al. Overview of the Architecture, Circuit Design, and Physical Implementation of a First-Generation Cell Processor. *IEEE Journal of Solid-State Circuits*, 41(1), January 2006.

[44] P. J. Restle et al. A Clock Distribution Network for Microprocessors. *IEEE Journal of Solid-State Circuits*, 36(05), May 2001.

[45] P. J. Restle et al. The Clock Distribution of the Power4 Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2002.

[46] P. J. Restle et al. Timing Uncertainty Measurements on the Power5 Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2004.

[47] S. Rusu and S. Tam. Clock Generation and Distribution for the First IA-64 Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2000.

[48] I. Sutherland, B. Sproul, and D. Harris. *Logical Effort: Designing Fast CMOS Circuits*. Morgan-Kaufmann, 1999.

[49] S. Tam, S. Rusu, U. N. Desai, R. Kim, J. Zhang, and I. Young. Clock Generation and Distribution for the First IA-64 Microprocessor. *IEEE Journal of Solid-State Circuits*, 35(11), November 2000.

[50] V. R. von Kaenel. A High-Speed, Low-Power Clock Generator for a Microprocessor Application. *IEEE Journal of Solid-State Circuits*, 33(11), November 1998.

[51] J. D. Warnock et al. The Circuit and Physical Design of the Power4 Microprocessor. *IBM Journal of Research and Development*, 46(1), January 2002.

[52] T. Xanthopoulos et al. The Design and Analysis of the Clock Distribution Network for a 1.2GHz Alpha Microprocessor. In *Proceedings of the International Solid State Circuits Conference (ISSCC)*, 2001.

[53] C. Yeh et al. Clock Distribution Architectures: A Comparative Study. In *7th International Symposium on Quality Electronic Design (ISQED)*, 2006.

A More on Clock Distribution

This appendix section gives more details on clock distribution trends, clocking characteristics and power and skew trends. We look at more detailed clock distribution techniques in 3 processors: DEC Alpha processors which were among the first of the high frequency processors implemented using custom design, Intel Pentium processors with high frequency and multiple GHz clocking and finally, Intel Itanium processors which have large die sizes with reasonably high frequencies.

DEC Alpha processors

One of the first 64-bit RISC processors, the Alpha processors employed one or more clock driver panes to drive the clock to multiple regions on the chip and local grid networks that supplied the clock to various latches. The driver panes were high power lines driving the local grids. The grid style networks dissipated lot of power but helped reduce the skew easily.

For EV4 [13, 12], level sensitive single phase clocking with a single driver pane for distribution was used. The 3rd metal layer was used for clock and power routing. The single driver pane was in the horizontal center of the chip extending vertically from top to bottom of the core. This pane was a hotspot in the chip which was significantly hotter (30C more) than the rest of the areas in the core. The single driver received the clock through a binary fanning tree with five levels of buffering. Grid style distribution was used to route the clock to the all areas of the chip. In EV5 [6, 3], upper metal layers 3 and 4 were used for clock distribution. It had 2 driver panes which led to more uniform distribution of the heat in the core. Two phase single wire clock distribution was used.

EV6 [18, 16, 21, 2, 50] had significantly more transistors than EV5 and dissipated more power. This led to the use of 16 driver panes with 4 independent phase-locked major clocks (one reference clock and 3 derived clocks) that were distributed over separate chip sections. Latches were clocked by secondary local grids. The core was divided into "Boxes" which represented different major sections of the core. Each Box clock driver drove a large grid over its region and was individually tuned to the capacitive characteristics of the Box. Local clocks were generated as needed from any clock including other local clocks. Different local clocks could have different numbers of local buffers and different buffer sizes. Time borrowing from temporally adjacent clock cycles was possible. Power dissipated from the clocking network was reduced from 40% (in the earlier processors) to 30% in EV6. Conditional logic gates were employed as clock buffers to realize power savings. Further, dynamic logic was used to reduce the switching capacitance of the clock network. These techniques saved about 10W of power compared to the previous generation processors.

In EV7 [52, 28], the core clock was similar to the EV6 clock but a separate NCLK powered the memory and network subsystem and was distributed along a grid surrounding the core. The NCLK region had a high skew (close to 11% of clock cycle). Two other major clocks were used to distribute clocks to the level-2 cache banks. X-trees and H-trees were used for distributing the global clocks.

Intel Pentium processors

The Pentium4 NetBurst architecture introduced in 2001 [32, 31] had multi-GHz clock generation and distribution. A total of 6 different core frequencies on the chip were used: 3 core frequencies for different sections of the core and 3 I/O frequencies for common, address and data I/O buses.

The global core clock is distributed through 3 spines to cover the large die for distribution. Each spine contains a binary distribution tree. This distribution tree terminates in 47 domain

buffers producing 47 independent clock domains. A four-stage hierarchical network of phase detectors is used to compare the rising edge clock timings of all domain clocks. Domain buffers can be disabled to power down large functional units to save power. To reduce clock jitter due to supply switching noise, the clock repeaters in the global distribution network use a RC-filtered power supply. Skew due to die variations is reduced using a static clock-deskewing scheme using delay adjustment. Interdomain skew is observed to be less than 20ps. The local clock drivers are used to drive the local clock load as well as produce the frequency appropriate for the particular logic block (one of the 3 different core frequencies 4GHz, 2GHz and 1GHz).

The Pentium4 of 2003 [5] had a 3-level clock distribution: Pre-Global Clock Network (PGCN, 27 inversion stages), Global Clock Grid (GCG, single inversion stage) and Local Clocking (4 inversion stages). Skew over multiple stages is reduced using shorting of inputs to receivers and also shorting of adjacent receivers. Reduced capacitance overhead (because of shorting) keeps the power consumption low (unlike typical grids). The local network supports clock gating for reducing power dissipation. The final grid stage and its driver dissipate 1.75W/GHz in addition to 0.75W/GHz in the PGCN.

Intel Itanium processors

The First Itanium 64-bit processor (2000) had a balanced H-tree for global clock to distribute the core and reference clocks to 8 deskew clusters each with 4 deskew buffers [47, 49]. The output of the deskewing logic is routed through a balanced tree to the distributed regional clock drivers (RCDs). There are 30 RCDs in the chip. The RCDs drive a uniform clock grid over a cluster. Local clock buffers tap these RCDs and provide the clocks to latches and other clocked elements. A separate reference clock is distributed along with the global clock to complete the deskew arch. Other features include intentional clock skew, clock borrowing and clock gating. The active distributed deskewing technique has programmable deskew units. Local skew control compensates for load mismatches, die variations and temperature and voltage gradients. Since there are several local grids, each grid spanning only a small portion of the chip area, skew minimization is easier.

In contrast to Itanium 1, the Itanium 2 processor (McKinley) use an H-tree networks for both global and local distributions [37, 38, 1]. Later releases of McKinley use a fused based deskewing technique.

The Montecito processor [17, 33, 39] has two Itanium cores and contains 1.7B transistors, the highest number of devices currently shipping on a microprocessor. It is a variable frequency clocking system and contains a single PLL that generates a multiple of the system clock frequency distributed to 14 digital frequency dividers (DFDs) for division to the proper zone frequency. Clock zones consist of major chip zones like the cores, bus logic etc. There are four major segments in the clock distribution. The first segment connects the PLL to the digital frequency divider (DFD). It is the only segment that has constant voltage and frequency. The second segment connects the DFD to the second level clock buffer (SLCB). The DFD output varies in frequency and it operates on a varying core supply voltage. The third segment connects the SLCB to the local clock buffers. A skew-matched RLC tree network technique is used. A typical SLCB drives 400 local clock buffers called clock vernier devices (CVDs) at 200 different locations. The fourth segment is the post-gater route in which gater loads are matched to the load of the latches and the routes that they drive. Worst case skew is observed to be 10ps.

Power Considerations

Clock distribution power breakdown for some processors is given below.

- The Alpha 21264 (1998) has an average power dissipation of 72W which includes 5.8W in the driver panes driving the global clock network and 10.2W in the gridded global network. The gridded clock distribution uses 14W, with local unconditional clocks using 7.6W and local conditional clocks using a maximum of 15.6W. More details are available in [2].
- The Itanium Montecito (2006) spends 25% of its power in clock distribution, with about 5% in its global and intermediate network and 20% in its local distribution driving the latches [33, 39].

Skew Considerations

Skew breakdown for some processors is given below.

- Alpha 21264 had substantially smaller local skew (box clock skew) compared to the global clock skew of 65ps.
- AMD K7 skew - 96ps for channel length variation and 32ps for RC simulated skew [26].
- Itanium 2 (McKinley) has overall skew around 62ps and local skew of 25ps [1].
- Dual-core Itanium (Montecito) has intermediate network skew of 6ps and overall skew of less than 10ps [33].