

Electrical and Optical Clock Distribution Networks for Gigascale Microprocessors

Anthony V. Mule', *Student Member, IEEE*, Elias N. Glytsis, *Senior Member, IEEE*, Thomas K. Gaylord, *Fellow, IEEE*, and James D. Meindl, *Life Fellow, IEEE*

Abstract—A summary of electrical and optical approaches to clock distribution within high-performance microprocessors is presented. System-level properties of intrachip electrical clock distribution networks corresponding to three microprocessor families are summarized. It is found that global clock interconnect performance and short-term jitter present the greatest challenges to the continued use of conventional clock distribution methodologies. An extrapolation of trends describing the percentage of clock period consumed by global skew and short-term jitter identifies the 32-nm technology generation of the 2002 International Technology Roadmap for Semiconductors (ITRS) as the first technology generation within which alternate methods of clock distribution may be warranted. Research efforts investigating interboard through intrachip optical clock distribution are also summarized. An optical distribution network compatible with high volume manufacturing in conjunction with a suitable means of providing optical-to-electrical signal conversion comprise the two fundamental challenges facing successful implementation of an optical clock distribution network. It is found that a global guided-wave distribution capable of efficient input and output coupling of optical power is required to meet the first challenge. The identification of a suitable means of optical-to-electrical conversion, however, remains an active topic of research.

Index Terms—High performance, high-speed interconnect, optical, optoelectronic integrated circuits, system level, VLSI.

I. INTRODUCTION

MANY unique approaches for implementing intrachip clock distribution in future high-performance gigascale microprocessors have been presented in the literature. These approaches include off-chip interconnection [1], package-level distribution [2], wireless microwave distribution [3], asynchronous distribution [4], distributed oscillator array networks [5], and optical distribution [6]. The desire to minimize the fraction of clock period consumed by skew and jitter exists equally in all cases irrespective of the implementation methodology. The fundamental differences among these approaches lie with:

- a) the respective method of implementing the global clock distribution network;
- b) the amount of skew and jitter generated by the distribution;
- c) area and power consumption associated with the driver and wiring distributions of the clock network;

Manuscript received April 13, 2001; revised November 22, 2001. This work was supported by the Semiconductor Research Corporation under Contract SJ-374.

The authors are with the Microelectronics Research Center, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: tvmule@ieee.org).

Digital Object Identifier 10.1109/TVLSI.2002.801604

- d) the novel advances in microelectronic technology required for realization of the proposed network (if any);
- e) the impact on clock network design methodologies.

Given that the success of an alternate approach depends on the relative gains offered in terms of performance, power dissipation, and design complexity, limitations associated with electrical clock distribution techniques must first be understood before the full potential of an alternate methodology can be assessed. Once these limitations are defined, the initial entry point of an alternate clock methodology into gigascale microprocessor technology can be identified.

In this paper, a summary of system-level properties describing the electrical clock distribution networks of high-performance microprocessors reported in the literature is presented. The continued practice of overcoming global clock interconnect performance limitations through the addition of global clock repeaters is questioned based upon the extrapolation of short-term jitter trends reported in the literature. The 32-nm technology generation of the 2002 ITRS is identified as the first technology generation within which an alternate global clock distribution methodology may be warranted. As optical clock distribution represents a possible alternative, a summary of interboard through intrachip optical clock distribution networks presented in the literature is provided to identify strengths and weaknesses associated with different distribution methods. Previous work on optical clock distribution has proceeded without a clear assessment of the potential for electrical clock distribution networks to extend performance into the GHz frequency regime. This paper presents for the first time a comprehensive summary of performance attributes associated with both electrical and optical clock distribution networks to underscore practical considerations for integration. By identifying the strengths and weaknesses of optical clock distribution networks, new research efforts in this field can be focused on the optical methodology or methodologies offering the greatest promise for successful integration.

Section II describes system-level properties that characterize any clock distribution network, regardless of implementation methodology. Section III provides a summary of key performance limitations challenging the extrapolation of electrical clock distribution methodologies to future technology generations. Section IV identifies the first technology generation within which an alternate global clock distribution methodology may be warranted. Section V presents an overview of general attributes associated with optical methods of clock distribution. Section VI summarizes key features of representative optical clock distribution networks achieving interboard through

intrachip clock distribution. Section VII presents a compilation of the advantages and disadvantages associated with various approaches to optical clock distribution. Section VIII presents concluding remarks.

II. CHARACTERIZATION OF CLOCK DISTRIBUTION NETWORKS

Over the course of the past three decades, microprocessor clock frequencies have increased from 108 kHz in 1971 [7] to over 3.0 GHz [8]. Two sources of timing uncertainty challenge the successful synchronous distribution of a clock signal. The first source of uncertainty, skew, represents static differences in the arrival time of the clock signal as measured between two or more distribution points. The second source of uncertainty, jitter, represents dynamic differences in the arrival time of the clock signal as measured from the same distribution point. Each source of timing uncertainty has a unique effect on critical path timing constraints. In the case of clock skew, a fixed deviation in the arrival time of the clock between successive storage elements can be either harmful or beneficial, depending upon the nature of combinational logic delay between the elements [9]. In addition, a distinction between *known* skew and *unexpected* skew is necessary. Known skew can be compensated for during design, resulting in negligible impact on performance. Unexpected skew, however, requires the addition of a timing margin to cycle time and degrades performance [10]. The effect of skew on the ability of a local logic path to meet cycle-time constraints, therefore, depends upon the nature of both delay and skew along that path. To allow flexibility in tailoring the effect of skew on local critical path performance, minimization of global skew is essential. In the case of jitter, two versions exist that present challenges to both intrachip and interchip synchronization. The first version, short-term (or cycle-to-cycle) jitter, affects intrachip synchronization and results from the modulation of clock buffer delay due to power-supply noise generated during random changes in digital switching activity across a die [11]–[13]. Short-term jitter must be added to the cycle-time budget of critical logic paths, reducing the maximum performance of the chip. The second version of jitter, long-term jitter, is represented by long-term shifts in the clock edge due to noise-induced shifts in the voltage-controlled oscillator (VCO) operating frequency of the on-chip phase-locked loop (PLL). Although long-term jitter challenges interchip input/output synchronization, intrachip critical path timing budgets are not affected since the same frequency is sent globally to all latches [13].

To understand limitations facing electrical clock distribution with respect to performance, it is first important to distinguish between open-loop distributions, closed-loop distributions without active compensation, and closed-loop distributions with active compensation. The basic premise behind open-loop synchronous timing is that no measurement or adjustment of clock phase is made to compensate for phase discrepancies between two clock signals stemming from the same source. Closed-loop distributions typically involve the use of on-chip PLL(s) to compensate for delay between the off-chip board-level reference clock and the on-chip clock proceeding

some portion of the global distribution network.¹ The use of a PLL alone, however, does not compensate for global skew. To reduce global skew, a closed-loop distribution network with active compensation can be used to measure and adjust phase differences between clock signals reaching different portions of the chip. In such distributions, active compensation (typically achieved through a delay-locked loop) can eliminate the majority of skew from the global clock signal, leaving jitter as the dominant limitation on performance [11]. The majority of electrical clock distribution networks reported in the literature have employed closed-loop synchronization both with and without active compensation to reduce global clock skew. All approaches to optical clock distribution investigated to date, however, can be categorized as open-loop distributions.

III. PERFORMANCE LIMITATIONS OF CONVENTIONAL ELECTRICAL CLOCK DISTRIBUTION NETWORKS

In discussing limitations associated with electrical clock distribution networks, the difference between the maximum frequency of operation and distribution must first be noted. The maximum frequency of *operation* for a microprocessor refers to the number of gate delays included along frequency-limiting critical logic paths. Only through a combined reduction in the delay of each gate through technology scaling as well as the total number of gate delays have the reported increases in clock frequency of high-performance microprocessors been achieved. Limits with respect to the maximum operating frequency of a microprocessor are therefore dictated by circuit and system limits with respect to the minimum number of gate delays along critical paths in conjunction with device and circuit limits on minimum allowable gate delay within a specific technology generation. By contrast, the maximum frequency of *distribution* is related to performance limitations imposed by global clock interconnection and the percentage of clock period consumed by skew and jitter. Hence, the goal of a clock distribution network is to enable the maximum operating frequency of a given chip architecture/technology iteration by minimizing the impact of global interconnection, skew, and jitter.

In determining the most restrictive limitations on the maximum clock frequency of distribution, the following factors must be considered: a) performance limitations imposed by global clock interconnection; b) skew associated with global and local critical logic paths; c) short-term jitter associated with global clock interconnect and distributed clock drivers; and d) microprocessor architecture. Absolute delay between global clock source and local latches is unimportant provided that transition times associated with individual edges of the clock signal do not degrade. As propagation delay along clock interconnect increases, however, clock-edge integrity degrades due to resistance-inductance-capacitance (RLC) parasitics. Minimization of propagation delay along global clock interconnect is therefore required to avoid clock-edge degradation. This minimization is achieved in practice through the addition of clock repeaters. An increase in the number of clock repeaters, however, increases global skew and short-term

¹An exception is seen in the 433 MHz Alpha 21164, for example, which uses a duty cycle equalization circuit.

TABLE I
SUMMARY OF HIGH-PERFORMANCE ELECTRICAL CLOCK DISTRIBUTION NETWORKS REPORTED IN THE LITERATURE

Microprocessor	Ref.	f_c (MHz)	t_x (1×10^6)	V_{dd} (V)	L (μm)	Die Area (mm^2)	Global Wiring Distributions ⁽⁴⁾	OL, CL, or CLAC	Global Skew (ps)	% Cycle Time (%)
Alpha 21064	[14]	200	1.7	3.3	0.75	234	Binary Tree, Grid	OL	293 ⁽¹⁾	5.86
Alpha 21164	[15]	300	9.3	3.3	0.50	299	Binary Tree, Grid	OL	80 ⁽¹⁾	2.40
Alpha 21164	[16]	433	9.6	2.0	0.35	209	Binary Tree, Grid	CL	90 ⁽¹⁾	3.90
Alpha 21264	[17]	600	15.2	2.2	0.35	318	X-Tree, H-Tree, RC Tree, Grid	CL	72 ⁽¹⁾	4.32
Alpha 21364	[18]	1200	152.0	1.5	0.18	397	X-Tree, H-Tree, Grid, Spine	CLAC	~ 90 ⁽¹⁾	10.8
Pentium II	[19]	300	7.5	2.8	0.35	203	NA	CL	120 ⁽³⁾	3.60
Pentium II	[20], [21]	450	7.5	1.4- 2.2	0.25	131	Binary Tree, Spine	CLAC	15 ⁽²⁾	0.68
Pentium III	[22]	650	9.5	1.4- 2.2	0.25	123	Binary Tree, Spine	CLAC	15 ⁽²⁾	0.98
IA-64	[23]	800	25.4	NA	0.18	NA	H-Tree, Binary Tree, Grid	CLAC	28 ⁽²⁾	2.24
Pentium III	[24]	1000	28.0	NA	0.18	NA	NA	NA	35-70 ⁽³⁾	3.50-7.00
Pentium IV	[25]	2000/4000 ⁽⁵⁾	42.0	NA	0.18	NA	Binary Tree	CLAC	16 ⁽²⁾	3.20
S/390 G4	[26]	411	7.8	2.5	NA	300	H-Tree	CL	30 ⁽²⁾	1.23
S/390 G5	[27,29]	609	25.0	1.9	0.25	215	H-Tree, Grid	CL	12 ⁽²⁾	0.73
S/390 G6	[28, 29]	760	25.0	1.9	0.20	215	H-Tree, Grid	CL	12 ⁽²⁾	0.91
PowerPC	[30]	1150	19.0	1.9	NA	NA	H-Tree, Grid	CL	15 ⁽¹⁾	1.72
Power4	[31] ⁽⁶⁾	1300	174.0	NA	NA	NA	H-Tree, Grid	CL	25 ⁽²⁾	3.25

(1) Simulated (2) Measured (3) Not Specified (Simulated or Measured) (4) Multiple instances of each topology may exist per chip

(5) Global/Local Frequencies

(6) Includes data from conference presentation

f_{clk} : Clock Frequency

t_x : Transistor Count

V_{dd} : Supply Voltage

L: Channel Length

OL: Open Loop

CL: Closed Loop

CLAC: Closed Loop with Active Compensation

NA: Not Available

jitter due to the impact of process variations and power supply noise, respectively, on the integrity of the global clock signal. Other common techniques for delay reduction along global clock interconnection offer limited relief. For example, delay reduction afforded through reverse scaling of global clock interconnection is limited to that associated with an optimal wire width for a particular interconnect technology, since delay begins to increase for wire widths larger than this value due to the increased parasitics between clock lines and non-return path interconnects. In addition, despite the approximately 40% reduction in resistivity when compared to aluminum, copper-based global clock interconnections still require an increasing number of repeaters with frequency due to losses induced by wire resistance [10].

Various techniques exist in the literature for eliminating skew from the global clock signal. Short-term jitter, however, remains uncompensated, as active compensation circuitry cannot make instantaneous corrections for delay variations in distributed clock buffers [13]. Global skew has been effectively managed through the use of various distribution geometries and forms of active compensation, as reflected in Table I. In this table, the clock distribution networks corresponding to the microprocessor architectures of the a) Alpha 21064, 21164, 21264, and 21364, b) Pentium, Pentium II, Pentium III, IA-64, and Pentium IV, and c) PowerPC, S/390 G4, G5, and G6, and Power4 are summarized with respect to technology, global clock distribution geometry, nature of distribution topology [open-loop (OL), closed loop without active compensation (CL), or closed loop with active compensation (CLAC)], worst case global skew, and percentage of clock period consumed by global skew based on information reported in the literature. Trends describing the consumption of clock period by global

skew versus clock frequency for individual microprocessors of Table I are illustrated in Fig. 1. From these data, it is evident that global skew has been minimized through intelligent choices of closed-loop active compensation techniques and/or global distribution geometry. The amount of short-term jitter is reported in conjunction with either the global phase-locked loop used to drive the distribution or the microprocessor itself. Table II summarizes measured values for short-term jitter corresponding to the PLL/global clock distribution network of several microprocessors taken during active operation. Trends describing the percentage of test period consumed by short-term jitter versus test frequency for the microprocessors of Table II are presented in Fig. 2. These measurements indicate that the percentage of clock period consumed by short-term jitter is increasing with clock frequency at a rate higher than that exhibited by skew. This trend is attributed to the presence of heightened power supply noise within increasingly complex microprocessor architectures. It should be noted that each microprocessor incorporates a significant amount of decoupling capacitance to minimize the effect of power supply noise (see references of Table II).

Limits on global interconnect performance are reflected in the projections of the 2000 ITRS through the separation of chip-level clock frequency into global and local components. Implicit to the notion of local and global clock frequencies is the use of clock multipliers to generate the local clock frequency for individual locally synchronous regions. Conventionally, PLLs have been used to perform board-to-chip-level frequency multiplication. The use of *local* PLLs to generate the local clock frequency implies embedding each analog PLL within local islands of digital logic, and represents a challenging extension to the practice of locating the chip-level PLL within areas more easily made

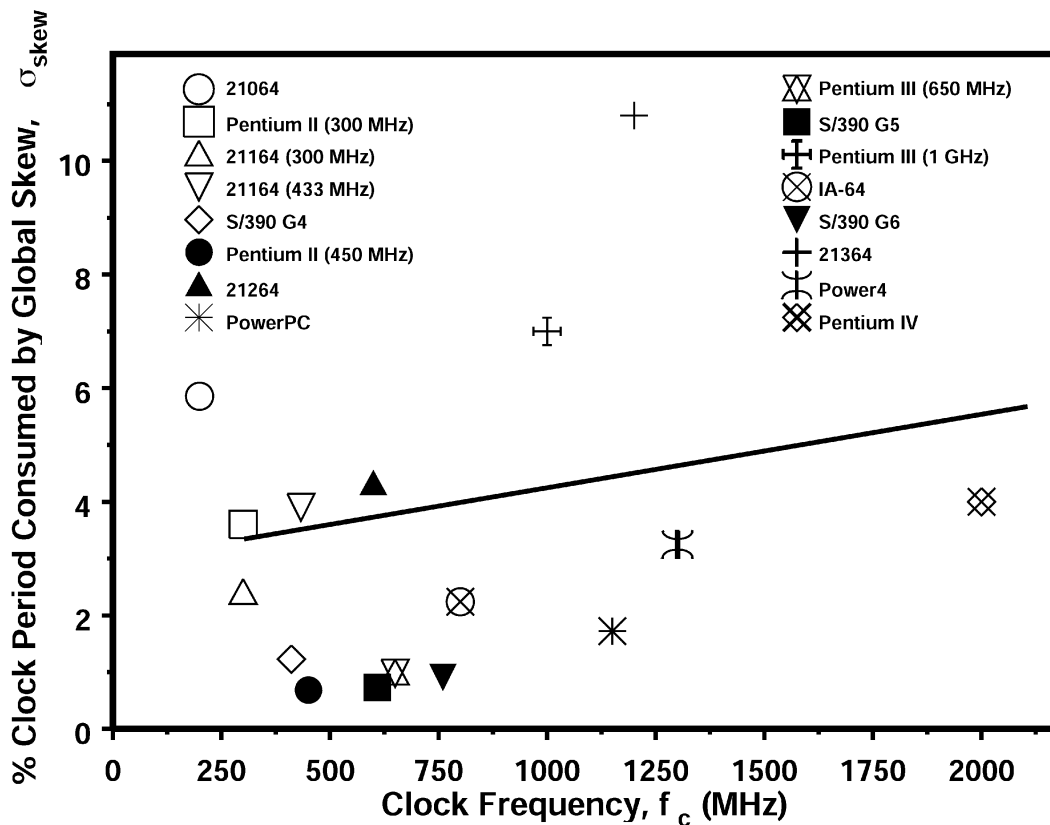


Fig. 1. Percentage of clock period consumed by global skew versus clock frequency for microprocessors of Table I. Straight line fit to data is shown.

TABLE II
PERCENTAGE OF TEST PERIOD CONSUMED BY SHORT-TERM JITTER FOR REPORTED MICROPROCESSORS

Microprocessor	Ref.	Test Frequency (MHz)	Peak-to-peak Short-term Jitter (ps)	Consumption of Test Period (%)
Pentium II	[13]	320	154	4.93
Pentium III	[24]	1000 ⁽¹⁾	92	9.20
Pentium IV	[25]	2000 ⁽¹⁾	35	7.00
Power3	[32]	180	80	1.44
Power4	[31]	1300 ⁽¹⁾	30	3.90
PowerPC	[33]	1046	72	7.53
PowerPC	[30]	1360	36	4.90

(1) Test frequency not specified. Frequency listed is nominal clock frequency of microprocessor.

immune to power supply and substrate noise, such as the corner [16] or center [26] of the chip. The design of PLLs with respect to jitter involves the optimization of the PLL loop bandwidth to minimize the contributions of input jitter and internal jitter generated within the VCO to output jitter [34]. Jitter present at the input of a PLL experiences a low-pass transfer function, while jitter generated internally within the VCO experiences a high-pass transfer function at the PLL output. Although the attenuation of short-term input jitter provides a solution to the trends depicted in Fig. 2, short-term jitter generated within the VCO of a local PLL will propagate to the local clock distribution network following the PLL, resulting in the negation of performance gains achieved through the elimination of jitter from the

global clock signal. Given the anticipated increase in transistor count and power-supply/substrate noise, the use of local PLLs, therefore, may not meet the challenge of reducing the consumption of available clock period by short-term jitter. The use of local pulse generators to generate the local clock frequency (as seen in the Pentium IV clock distribution, for example) relies on the efficient filtering of noise from the power supply, as the broadband nature of pulse generators is such that any jitter present at the input is transferred to the output [25]. Barring the advent of novel local frequency multiplication schemes that provide immunity to jitter, the global-to-local clock distribution scheme projected by the 2000 ITRS may prove insufficient for future gigahertz clock distribution networks. An estimate for the

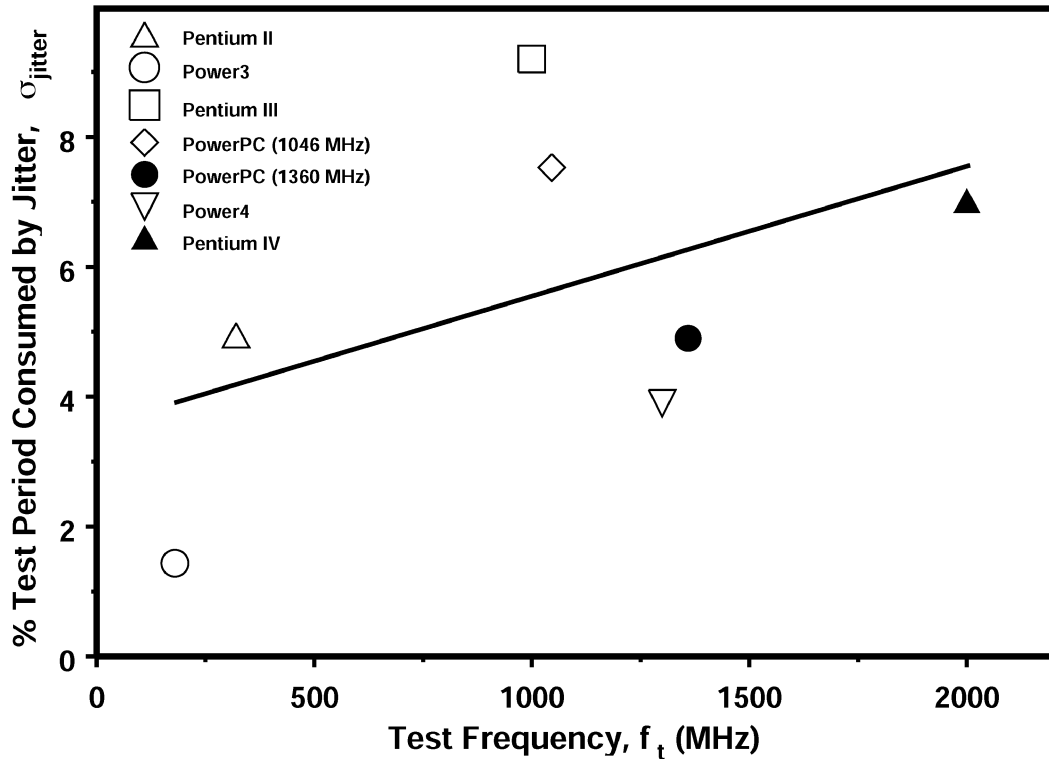


Fig. 2. Percentage of test period consumed by short-term jitter versus test frequency for microprocessors of Table II. Linear fit to data is shown.

TABLE III
PROJECTIONS FOR SHORT-TERM JITTER AND SKEW VERSUS TECHNOLOGY GENERATION

2002 ITRS Technology Generation, TG (Year / nm)	Local Clock Frequency (GHz)	% Local Period Consumed by Skew/Jitter, $\sigma_{skew}/\sigma_{jitter}$ (extrapolated)	% Local Period Available for Critical Path Logic
2004 / 90	4.0	6.0 / 11.0	83.0
2007 / 65	6.7	9.0 / 16.0	75.0
2010 / 45	11.5	13.5 / 26.0	60.5
2013 / 32	19.3	21.0 / 41.0	38.0
2016 / 22	28.8	30.0 / 60.0	10.0

initial technology generation within which an alternate global clocking methodology is merited is provided in the next section.

IV. 2002 ITRS 32-nm (2013) TECHNOLOGY GENERATION: INITIAL TECHNOLOGY GENERATION MERITING ALTERNATE GLOBAL CLOCKING METHODOLOGY

To predict the initial technology generation within which an alternate global clocking methodology is merited, the trends depicted in Figs. 1 and 2 are extrapolated to encompass the range of local clock frequencies predicted by the 2002 ITRS. Table III summarizes the extrapolated values for local skew and short-term jitter for each technology generation. To predict the amount of local clock period consumed by skew and short-term jitter, it is assumed that local clock multipliers operate in a manner similar to that seen in the Pentium IV, where any variations at multiplier input are passed unattenuated to the output. As Table III indicates, the percentage of local period

available for critical path gate delays reduces from 60.5% to 38% in transitioning from the 45-nm technology generation to 32-nm technology, respectively. Assuming the consumption of 39.5% of the local clock period by timing uncertainties is acceptable within the 45-nm generation, *the 32-nm (or 2013) technology generation of the 2002 ITRS represents the first generation within which an alternate global clock distribution methodology may be warranted.*

One approach for overcoming constraints on the clock frequency of distribution imposed by global clock interconnect and jitter is global optical clock distribution. As both guided-wave and free-space propagation media impose no restrictions on the maximum frequency of modulation of an optical signal, *global distribution of the local clock frequency is possible.* Optical methods of clock distribution, therefore, allow for the elimination of global electrical clock interconnect, distributed global buffers, and local clock multipliers. General properties of three different approaches to optical clock distribution are provided in the next section.

V. GENERAL PROPERTIES OF OPTICAL CLOCK DISTRIBUTION NETWORKS

Optical clock distribution was first suggested by Goodman in 1984 [35]. The main components of an optical clock distribution network for interboard, interchip, or intrachip synchronization are as follows:

- the photon source, either on-chip or off-chip;
- the propagation medium of the optical path, either free-space or guided-wave;
- a passive diffractive optics device for light redirection (depending on the assumed optical propagation medium);
- a photodetector for optical-to-electrical conversion that is included through monolithic or hybrid attachment methods;
- a receiver circuit for amplification of low-level photocurrent.

In general, three different approaches exist for optical clock distribution based upon the propagation medium between the photon source and chip-level detectors: a) unfocused free-space, b) focused free-space, and c) guided wave [35]. In unfocused free-space communication, an off-chip photon source broadcasts to the entire chip, with detectors placed at desired points for optical-to-electrical conversion. The main weaknesses to this approach are the need for global masking to avoid unwanted photoelectron generation in nondetector regions (depending on the wavelength of light propagating within the distribution), the reduction in the amount of optical power incident to each detector, and the need for a three-dimensional propagation volume. Focused free-space communication sends optical power to the desired detector locations through the use of a focusing element, reducing the need for global masking of regions where no incident signal is desired.

A focused free-space optical clock distribution employing holographic redirection of the optical signal is illustrated in Fig. 3 [36]. Focused free-space propagation requires a transmissive or reflective diffractive optical element (DOE) as the focusing medium, precise alignment of the source-DOE-chip optical system to ensure proper detection of the focused beams, and a three-dimensional propagation volume. Perhaps most importantly, such distributions are not conducive to heat removal and supply-of-power techniques that require the use of area on both sides of a fully packaged chip.

Substrate-mode guided-wave distribution represents a compact approach to focused free-space distribution by confining the optical signal to a dielectric substrate. A substrate-mode guided-wave distribution is illustrated in Fig. 4 [37]. Redirection of the optical signal within the dielectric substrate can occur through multiple gratings, mirrors, and/or microlenses located along individual optical pathways. High-performance surface-relief gratings for beam-splitting and reflection require the use of advanced fabrication techniques such as direct e-beam writing. In addition, the presence of multiple microoptic components along each signal pathway results in increased optical loss. Guided-wave distribution employing fiber-optic or integrated optical waveguides allows for compact, planar packaging of the optical/electrical system without the need

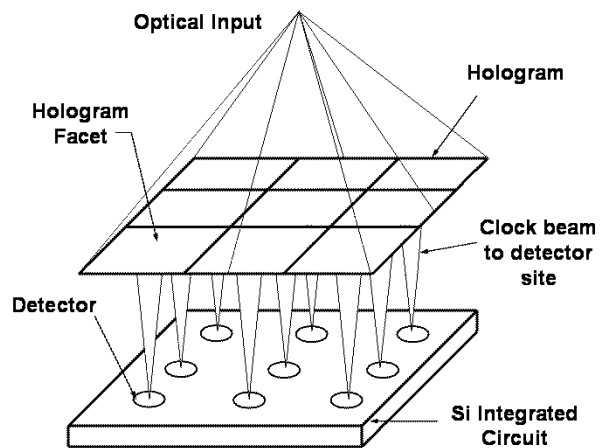


Fig. 3. Focused free-space optical clock distribution configuration [36].




Fig. 4. Substrate-mode guided-wave optical clock distribution configuration [37].

for multiple diffractive components along individual optical pathways. An integrated optical guided-wave distribution is illustrated in Fig. 5. Unless a low loss guided-wave optical interconnect technology employs a sufficient contrast in refractive index between core and cladding materials, optical loss inherent in bent waveguide arms can become significant for small radii of curvature. In addition, efficient coupling of optical power into a guided-wave distribution is inherently more difficult, as only a finite number of optical modes are involved.

Since Goodman [35], the majority of research on optical clock distribution has been directed toward interboard and interchip synchronization, with limited research focused on intrachip distribution. To the authors' knowledge, no work has been reported in the literature on unfocused free-space distributions. Representative examples of optical clock distributions for interboard, interchip, and intrachip synchronization based on focused free-space and guided-wave distributions are discussed in the following section.

VI. OPTICAL CLOCK DISTRIBUTION NETWORKS REPORTED IN THE LITERATURE

A. Focused Free-Space Distributions

Research on the use of diffractive optical elements for transmission or reflection of optical signals for interboard and interchip *data* interconnection is abundant in the literature [36]. Distribution of an optical *clock* signal by means of a diffractive




Fig. 5. Optical clock distribution using integrated optical waveguides.

optical element can be thought of as a specialized case of optical data interconnection, where the explicit need for chip-level optical photon sources is eliminated and fewer chip-level detector/receiver pairs are required. In both cases, a diffractive optical element is responsible for directing the incident optical power to the desired detector locations.

Holograms are a common diffractive optical element investigated in the literature for optical clock distribution networks based on focused free-space distribution. Investigation of the skew properties of an array of optical transimpedance receivers associated with a hologram-based focused free-space distribution has been presented by Goodman and Clymer [38]. In this work, each receiver consists of a silicon p-i-n detector to convert the incident optical power into a low-level photocurrent. The generated photocurrent is subsequently converted to a voltage by a transimpedance receiver, whereby it is amplified by a chain of analog inverter amplifiers. Test circuits comprised of 18 receivers each were fabricated in 3- μm Mosis technology. Each p-i-n detector was $20 \times 20 \mu\text{m}^2$, surrounded by a $100 \times 100 \mu\text{m}^2$ aluminum sheet derived from the topmost interconnect metal layer to prevent unwanted optical carrier generation within the silicon area surrounding each detector. A mean maximum input frequency of 15.1 MHz was experimentally observed, along with an average of 13.5-ns skew using a novel skew measurement circuit. A major contribution to the measured skew was the wide variation in metal-oxide-semiconductor field-effect transistor (MOSFET) threshold voltages inherent in the 3- μm digital fabrication process.

Focused free-space optical clock distribution requires intricate alignment of the source-DOE-receiver system [39]. The main drawback, however, is the need for a large three-dimensional propagation volume. *Large volume distributions prohibit compact packaging of the optical/electrical system, thereby inhibiting high volume manufacture.* Focused free-space distribution can be achieved in a compact, planar manner by coupling optical power into a dielectric substrate and subsequently redirecting it through monolithically incorporated microoptic components. In this approach, the optical signal is confined within the substrate along the transverse direction, defined as normal to both the direction of propagation and the surface of the substrate. Diffraction, refraction, and/or reflection of the optical signal at desired points within the substrate can occur through monolithic integration of surface-relief or

volume gratings, microlenses, and/or dielectric or metallic mirrors, respectively, within the top and/or bottom surface(s) of the substrate. This approach to focused free-space distribution is also known as substrate-mode guided-wave distribution. In contrast to integrated optical waveguide propagation, confinement of the optical signal does not occur in the lateral direction, defined as normal to the direction of propagation and parallel to the surface of the substrate.

Research investigating substrate-mode wave propagation is presented in [37] and [40]. The optical substrate representing the central component of the distribution is depicted in Fig. 4. Within this substrate, a normally incident optical beam is diffracted by a beam-splitting surface-relief binary phase grating into two separate beams. Each diffracted beam propagates within the substrate by repeatedly impinging on metallic mirrors until reaching a multiphase reflection grating, where it is redirected normally to the next beam-splitting diffraction grating. This process repeats until each optical wave terminates onto a final reflection grating, where it is directed to a chip-level detector for optical-to-electrical conversion.

Other examples of optical clock distributions employing substrate-mode wave propagation can be found in [41]–[43]. Challenges facing substrate-mode guided-wave distribution include the need for multiple diffraction gratings, mirrors, and/or microlenses along individual signal pathways, and the need for integration of an appropriate substrate technology into microprocessor systems. With respect to the latter challenge, the integration of an appropriate substrate technology implies either hybrid integration with or outright replacement of conventional printed wiring board technology. The former approach implies placing the substrate on either the top or bottom of a printed wiring board, which, depending on the density of board-level electrical components and/or the integration of components on both sides of the printed wiring board, may not be feasible. The latter approach, where the optical substrate would serve to interconnect both optically and electrically the printed wiring board components, represents a significant shift from conventional low-cost printed wiring board technology.

B. Guided-Wave Distribution

Both fiber-optic and integrated optical guided-wave distributions have been reported in the literature for interboard and interchip applications. An integrated optical guided-wave distribution network is depicted in Fig. 5. Guided-wave distribution of an optical clock signal in an interboard fashion mandates the use of fiber-optic waveguides to enable flexibility in routing long-distance optical communication between source and detector locations. The most aggressive approach reported in the literature involves communication to 1024 separate ports via fiber-optic waveguides for board-to-board synchronization within telecommunications switching machines [44]. In this approach, a mode-locked femtosecond semiconductor laser diode system feeds a single optical fiber through a collimating lens with an input coupling efficiency of approximately 40%. Splitting of optical power is performed using an optical fiber splitter. The system, operating at 302 MHz, delivers 1 μW of optical power to each fanout port with < 12 ps jitter.

Research investigating interchip synchronization using fiber-optic waveguides is reported in [45], where distribution of an optical clock signal to 64 individual sites on a printed wiring board is realized through a bundle of optical fibers routed in an H-tree topology. In this approach, each receiver site is hand-fed a single optical fiber through a via hole in the printed wiring board. By passing individual fibers through the board at specific via locations, the need for light-redirection through diffractive optics components is eliminated. A measured skew of 23 ps was observed within the distribution. Right-angle bend radii of 580 μm experienced measured optical losses of 0.12 dB. Uniformity of the optical output power across the 64 fibers communicating with each point of fanout is distributed over an approximate 3.5-dB range due to the butt-coupling scheme used in coupling optical power into the fiber bundle.

Interchip optical clock distribution can also be achieved using integrated optical waveguides. A system-level design for an integrated optical waveguide H-tree distribution achieving interchip optical synchronization is presented in [46]. To distribute the clock signal, a silica glass material system is proposed for the waveguide distribution with SiON introduced within desired areas of optical power coupling. Optical power feeds the distribution via butt-coupling of an optical fiber to the input of a multimode waveguide. Fanout of the clock signal is performed via 3-dB directional couplers for power splitting at individual 1×2 junctions within the H-tree distribution. Outcoupling of the optical clock signal is performed at each fanout through a binary surface-relief grating with a theoretical coupling efficiency of 60%.

Research investigating distribution of a multigigahertz optical clock signal within a Cray supercomputer multiprocessor board is presented in [47] and [48]. In this approach, a polyimide optical waveguide distributes the optical clock signal in an H-tree distribution topology to 48 clock nodes over a $14.5 \times 27 \text{ cm}^2$ printed wiring board area. Key features of this approach are:

- a) the proposition of integrating thin-film vertical-cavity surface-emitting-laser sources and metal–semiconductor–metal silicon detectors within the printed wiring board;
- b) multimode polyimide waveguides and y-junction power splitters;
- c) the use of tilted binary surface-relief grating couplers and/or 45° total internal reflection (TIR) mirrors etched into optical waveguide arms for coupling of light into and out of the distribution.

Fabrication of tilted surface relief couplers and 45° TIR mirrors is achieved through reactive ion etching (or laser writing for TIR mirror fabrication), where grating coupler fabrication involves a Faraday cage to achieve a tilted profile. Measured efficiencies of the tilted binary surface-relief gratings and TIR mirrors are reported as 35% and approximately 100%, respectively. The use of TIR mirrors embedded within the printed wiring board allows for highly efficient coupling of light in a direction normal to the direction of propagation. The main drawback to this method of outcoupling, however, is the lack of control over output beam shape. The half-width at half-maximum is reported as 60 μm in [48], requiring in turn a large detector area for full beam cov-

erage. The use of a TIR mirror requires that both source and detector to be placed in close proximity of the mirror to ensure efficient coupling.

A research effort featuring silicon-based microphotronics is reported in [6] and [49]–[54]. In this approach, every element of the optical system is fabricated at the chip level from silicon-based materials. Monolithic light emission at an optical wavelength of 1.54 μm is achieved using erbium-doped silicon light-emitting diodes. Monolithic chip-level waveguides are fabricated using polysilicon and silicon dioxide as the waveguide core and cladding materials, respectively. Submicrometer waveguide dimensions ($0.2 \times 0.5 \mu\text{m}^2$) and right-angle bend radii of 2 μm radius have been fabricated with less than 1 dB insertion loss due to the high refractive index contrast Δn between polysilicon and silicon dioxide ($\Delta n \geq 2$) [49]. Optical-to-electrical conversion is achieved through heterojunction germanium-silicon detectors operating in the 1.3–1.55 μm wavelength range with responsivities of 0.3 and 0.2 A/W at $\lambda = 1.3$ and 1.55 μm , respectively [50]. Integration with chip-level waveguides is achieved by butt-coupling of waveguide outputs to individual detectors [51]. To date, the maximum measured optical output power of the Er:Si diodes is tens of μW s at $< 1\%$ output efficiency [52]. Design specifications enabling reduction of sidewall scattering losses within optical waveguides to 0.1 dB/cm are presented in [53], with the lowest reported values of measured propagation loss measuring 0.8 dB/cm [54].

Table IV summarizes the target application, location of the photon source, propagation medium(s) with respect to global signal distribution, diffractive optical element(s), method of photodetector integration, optical wavelength, distribution fanout, system clock frequency, and reported values of skew and/or jitter for select optical clock distribution networks reported in the literature. Performance limitations facing optical clock distribution networks summarized in Table IV are presented in the next section.

VII. STRENGTHS AND WEAKNESSES OF OPTICAL CLOCK DISTRIBUTION

The two fundamental obstacles challenging the adaptation of optical clock distribution within high-performance microprocessor systems are: 1) realization of a manufacturable, efficient optical distribution and 2) design of an appropriate chip-level receiver. From Section VI, focused free-space distributions employing diffractive optical elements located in a manner separate from the optical/electrical system require a significant three-dimensional propagation volume. Such architectures are not compatible with volume microprocessor manufacture and do not provide for concurrent use of projected heat removal and supply-of-power techniques. Compact, planar packaging is achieved within focused free-space distributions employing substrate-mode propagation. It is unclear, however, whether production of an appropriate substrate technology in a manner integrated within or parallel to printed wiring board manufacture represents a viable alternative to conventional practice.

Guided-wave distribution provides compact packaging of the optical distribution in a manner that minimizes the use of

TABLE IV
SUMMARY OF INTERBOARD THROUGH INTRACHIP OPTICAL CLOCK DISTRIBUTION NETWORKS

Ref.	Target Application	Location, Photon Source	Propagation Medium	DOE(s)	Photodetector Incorporation	λ (nm)	Fanout	f_{clk} (GHz)	Skew (ps)
[38]	IAC	Off-chip	FS	CGH	Monolithic	632.8	NA	0.015	13.5E3 ⁽¹⁾
[37]	IEC	Off-board	SMWG	SRDG	NA	850	8, 64	NA	~ 0 ⁽²⁾
[41]	IEB, IEC	Off-board	FS (IEB), SMWG(IEC)	VDG	(Stand-alone Module)	670	4	0.622	10 ⁽¹⁾ /36 ⁽¹⁾
[42], [43]	IEB, IEC	Off-board	FS (IEB), SMWG (IEC)	VDG	NA	850	42	NA	110 ⁽²⁾
[44]	IEB	Off-board	GW	None	NA	830	1024	0.302	12 ⁽¹⁾
[45]	IEC	Off-board	GW	None	NA	670/830	64, 128	NA	23 ⁽²⁾
[46]	IEC	Off-board	GW	SRDG	Hybrid	1300	16	0.500	~ 0 ⁽²⁾
[48]	IEC	In-board	GW	SRDG, TIRM	Hybrid	850	48	6.000	NA
[6]	IAC	On-chip	GW	None	Monolithic	1300, 1550	NA	NA	NA

1) Measured 2) Simulated

IEB: Interboard

FS: (Focused) Free-space

SRDG: Surface-Relief Diffraction Grating

TIRM: Total Internal Reflection Mirror

IEC: Interchip

SMWG: Substrate-Mode Waveguide

CGH: Computer Generated Hologram

NA: Not Available

IAC: Intrachip

GW: Guided-Wave

VDG: Volume Diffraction Grating

diffractive or reflective components. In addition, waveguide technologies exist that offer immediate compatibility with standard methods of printed wiring board [55], [56] and chip [6] manufacture. The lack of an efficient means by which optical power can be coupled into and out of a guided-wave distribution was cited in [35] as the central argument against waveguide-based optical clock distribution. As summarized in Section VI, a variety of methods have been investigated for achieving high-efficiency input and output coupling, including butt-coupling of optical sources to waveguide input/output regions, the use of TIR mirrors, or the use of diffractive couplers. The optimum choice depends primarily on constraints imposed by device manufacture and the location of optoelectronic devices with respect to coupling elements within a particular system. Butt-coupling of source-to-waveguide or waveguide-to-detector regions is possible only by locating each component within the same physical plane, which implies the integration of efficient chip level monolithic optical sources. In addition, placement of optical waveguides at local levels may be difficult due to routing constraints imposed by local electrical device and interconnect and by via blockage. Total internal reflection mirrors, typically defined through reactive ion etching or laser ablation [57], require immediate proximity between active device and waveguide to avoid excessive beam diffraction. Surface-relief couplers rely on wet-chemical or reactive ion etching for the definition of low-to-moderate-performance devices, or e-beam writing for high-performance devices. High-efficiency volume diffractive couplers, however, are produced through a holographic process, thereby avoiding fabrication-related errors in grating profile common to wet and dry etch chemistries. In addition, the ability to integrate both preferential coupling (i.e., directing the majority of optical power into a single or finite number of orders in the substrate or cover region) and focusing capabilities within a single device is possible using standard holographic fabrication techniques. The structure of a preferential-order volume focusing output coupler reported in [58] and [59] that achieves 98% preferential coupling into the cover region and 95% overall coupling effi-




Fig. 6. Preferential-order volume focusing grating coupler [58], [59].

ciency is depicted in Fig. 6. The importance of incorporating diffractive grating coupler technology (either surface-relief or volume) increases with the distance between active device and waveguide. In the case of power coupling from waveguide to detector, this allows for reduced detector area, and hence higher operating frequency of the distribution.² For high-efficiency input coupling into a guided-wave distribution, a volume focusing coupler with a Gaussian intensity profile could provide enhanced phase matching with Gaussian beams typically generated by semiconductor lasers [59]. Depending on the system-level architecture and associated design constraints, application of one or more of the above technologies to an optical waveguide clock distribution network can address the challenge of optical power coupling.

One architectural uncertainty with respect to an optical waveguide distribution is the degree to which the distribution must communicate with local latches. Implicit to each approach described by Goodman [35] is the extension of the global optical distribution to all latches via electrical interconnection. Assuming a local optoelectronic latch structure could in fact be designed to perform the required optical-to-electrical conversion in a power and area efficient manner at gigahertz frequencies, optically distributing a clock signal directly to each latch would

²The volume grating preferential-order focusing coupler reported in [58] produces a near-diffraction-limited focal line with a full width at half-maximum of 10.49 μm

require a highly nonsymmetrical routing geometry. For either focused free-space or guided-wave distributions, the required distribution fanout would need to correspond to the number of local latches, which for architectures including ≤ 100 million transistors already amounts to approximately 10^4 – 10^6 elements [11]. This fanout exceeds that of the highest fanout reported by any optical clock distribution by one to three orders of magnitude [44]. In addition, short electrical interconnection offers high density, high speed, and low switching energy with respect to the local distribution of a clock signal. The ability to place and route automatically local clock interconnection using computer-aided design tools additionally compliments the advantages of short electrical interconnection [60]–[62]. Thus, the local portion of any optical clock distribution network will likely be realized through electrical interconnect, and the requirement for a global fanout on the order of the number of local latches does not exist.

Optical waveguide technology does not allow for arbitrarily high levels of distribution fanout, however. In contrast to electrical interconnection, where a signal can be regenerated at periodic points through the use of a repeater, an optical signal must rely on the fixed amount of optical power available at the distribution input. This power must be sufficient for all individual propagation paths stemming from the optical source to ensure that adequate power is available at each receiver. In addition, this power must be distributed in a uniform manner to all points of fanout to ensure comparable receiver performance across an entire fanout array. For any global waveguide distribution, the maximum fanout is dictated by the available optical input power, optical power loss mechanisms inherent in guided-wave distribution, the minimum optical power that can be detected by chip-level receivers, and the area over which the distribution must communicate. Optical input power is dictated by that available from the optical source. Optical loss mechanisms include:

- a) input coupling loss from the optical source to the guided-wave distribution network;
- b) loss due to the tapering of waveguide dimensions from wide multimode to narrow single-mode waveguide dimensions (if necessary);
- c) power-splitting loss at each split junction due to reflection and scattering losses;
- d) bending loss incurred along curved waveguide arms due to the difference in phase velocity of the guided mode(s) and radiation mode(s) in the inner core and outer cladding regions, respectively;
- e) waveguide propagation loss due to absorption and scattering losses;
- f) output coupling loss going from the waveguide to detector.

The incorporation of volume couplers and/or TIR mirrors (depending on system architecture) can address the challenge of optical power coupling. Loss due to tapering of waveguide dimensions can be minimized, provided that a sufficiently long input taper can be realized. Assuming a fanout of two per power-splitting junction, each junction incurs some amount of optical loss due to reflection and scattering. Bending loss incurred along curved waveguide paths within curved waveguide arms decreases with an increase in refractive index contrast for small radii of curvature due to the strong confinement

within the core region. To realize the benefits of a high-contrast waveguide technology, however, optically-smooth sidewalls must be realized. Hence, if a low loss, high contrast waveguide technology is available, optical loss inherent in the fanout of optical power can be minimized, and the maximum global fanout of an intrachip guided-wave distribution is left to depend on the optoelectronic source and receiver technology.

Once an appropriate waveguide technology has been selected, a system-level design advantage afforded through optical interconnection is the elimination of global clock interconnect redesigns following enhancements in system operating frequency, as no changes in cross-sectional waveguide geometry are required with an increase in the frequency of modulation. In addition, the integrity of clock pulse transitions is decoupled from propagation delay along global optical clock lines communicating over chip-level and board-level distances [63]. This is in contrast to electrical clock interconnection, where the design of interconnect delay and the clock edge rate are strongly interdependent due to constraints imposed by RLC parasitics and jitter [10]. By removing the impact of interconnect delay on clock edge integrity, global clock interconnect design can be simplified. Design considerations with respect to operating conditions are also simplified using global optical waveguides, as currently available waveguide technologies (e.g., polymer or poly-Si/SiO₂) exhibit acceptable fluctuations in refractive index with changes in operating temperature over the temperature range typically encountered during microprocessor operation [63]. Three design considerations paramount to global electrical clock interconnect design can therefore be eliminated with the use of global optical waveguides.

It should be noted that early incorporation within the design cycle of both waveguide and receiver placement is important, as optical waveguides cannot be arbitrarily routed to the same degree as can electrical interconnections. In a similar fashion, the optimal placement of global clock repeaters in microprocessor design is not always possible due to area constraints imposed by large functional units and memory arrays [10]. In this case, nonoptimal buffer placement is compensated through tailoring of the electrical interconnect distribution. Such tailoring would be possible, although limited in comparison, within the context of an optical waveguide distribution, with low-loss, high-contrast waveguide technology offering the greatest degree of flexibility.

The second challenge facing mainstream integration of optical clock distribution techniques is the design of a high-performance optoelectronic receiver. Although a global waveguide distribution network is capable of delivering a skew- and jitter-free clock signal to each receiver through the use of equal or near-equal length global pathways free from electrical clock repeaters, optoelectronic receivers must be designed such that negation of these gains does not occur through poor receiver performance. A vast amount of research has been devoted to the design of high-performance smart-pixel receivers for optical data interconnection. Example CMOS smart-pixel receivers include both NMOS [65] and CMOS [66]–[68] transimpedance preamplifiers followed by single or multiple gain stages. The majority of smart-pixel-based optical receiver research reported in the literature consists of a photodetector to provide the

optical-to-electrical conversion and an integrated amplifier, which in its simplest form consists of a transimpedance amplifier for conversion of the detector-generated photocurrent to an analog output voltage. The transimpedance front end is followed by a series of small-signal voltage amplifiers to raise the analog signal amplitude to a level sufficient for a final digital voltage-level decision stage. A transimpedance receiver was investigated by Clymer and Goodman for clock distribution, which, as the authors discovered, did not result in robust, stable performance when implemented within a digital CMOS process [38]. Whether such high-gain receivers can operate with acceptable levels of performance within future gigascale microprocessors remains open to debate [63].

In addition, power dissipation associated with the receiver array must be considered. It must be noted that the majority of power dissipation in electrical clock distribution networks resides within the local clock distribution, which in certain microprocessor families consumes at least an order of magnitude more power than the global portion of the distribution network [10]. As such, the use of optical clock techniques may not offer significant power savings with respect to an intrachip clock distribution network unless a sufficient fanout can be achieved. To date, the development of an optical receiver capable of high-speed, low-jitter, and low-power operation that can drive a significant clock load capacitance is still pending, and remains a fundamental challenge to the mainstream integration of optical clock techniques.

VIII. CONCLUSION

A summary of electrical clock distribution networks implemented within high-performance microprocessors has been presented. Intelligent design of global distribution geometries and the use of active compensation techniques have controlled global skew, leaving short-term jitter and global interconnect performance as the main obstacles to achieving multigigahertz operation in future technology generations. Extrapolation of present trends describing consumption of clock period by both skew and short-term jitter predicts the need for an alternate global clocking methodology for microprocessors fabricated in the 32-nm technology generation of the 2002 ITRS.

Two fundamental challenges face the successful integration of optical clock techniques into gigascale microprocessors. The first challenge, realization of a compact, manufacturable optical distribution scheme, can be addressed with the use of an optical waveguide distribution employing volume diffractive couplers and/or total internal reflection mirrors in conjunction with an appropriate low loss, high-contrast waveguide technology. Electrical interconnection will play a role in optical clock distributions by connecting global optical pathways to local latching elements, thereby reducing global fanout requirements. Optical waveguide clock distribution allows for global propagation of the local clock frequency to optoelectronic receivers in a manner free from power-supply-induced jitter, thereby allowing for higher aggregate distribution performance and increased margin with respect to the design of local clock domains. Optical clock distribution mitigates global clock interconnect design considerations for operating frequency, clock

edge integrity, and temperature of operation. Considerations for reduced flexibility in waveguide routing in comparison to electrical interconnection must be accounted for during design. The second challenge, realization of a robust chip-level receiver, remains an active topic of research.

ACKNOWLEDGMENT

The authors would like to thank D. Bailey, I. Young, and P. Restle for clarifications regarding published information on microprocessor clock distribution networks.

REFERENCES

- [1] A. Naeemi, P. Zarkesh-Ha, C. S. Patel, and J. D. Meindl, "Performance improvement using on-board wires for on-chip interconnects," in *Proc. IEEE 9th Topical Meeting Electrical Performance of Electronic Packaging*, Scottsdale, AZ, Oct. 2000, pp. 325–328.
- [2] Q. Zhu and S. Tam, "Package clock distribution design optimization for high-speed and low power VLSI's," *IEEE Trans. Comp., Packag., Manuf. Technol. B*, vol. 20, pp. 56–63, Feb. 1997.
- [3] B. Floyd, K. Kim, and K. O, "Wireless interconnection in a CMOS IC with integrated antennas," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 2000, pp. 328–329.
- [4] M. Josephs, S. Nowick, and C. Berkel, "Modeling and design of asynchronous circuits," *Proc. IEEE*, vol. 87, pp. 234–242, Feb. 1999.
- [5] V. Gutnik and A. Chandrakasan, "Active GHz clock network using distributed PLL's," *IEEE J. Solid State Circuits*, vol. 35, pp. 1553–1560, Nov. 2000.
- [6] L. C. Kimerling, "Silicon microphotonics," *Appl. Surf. Sci.*, vol. 159–160, pp. 8–13, June 2000.
- [7] "The Intel MCS-4/4004," *Automat. Informat. Industriell.*, pp. 45–47, Apr. 1975.
- [8] D. Deleagnes, J. Douglas, B. Kommandur, and M. Patrya, "Designing a 3 GHz, 130 nm, Intel Pentium 4 processor," in *Proc. IEEE Symposium on VLSI Circuits*, Honolulu, HI, June 2002, pp. 130–133.
- [9] G. F. Taylor and G. Geannopoulos, "Microprocessor clock distribution," in *Proc. IEEE Electrical Performance of Electronic Packaging*, Napa, CA, Oct. 1996, pp. 28–30.
- [10] P. J. Restle, T. G. McNamara, D. A. Webber, P. J. Camporese, K. F. Eng, K. A. Jenkins, D. H. Allen, M. J. Rohn, M. P. Quaranta, D. W. Boerstler, C. J. Alpert, and C. A. Carter, "A clock distribution network for microprocessors," *IEEE J. Solid State Circuits*, vol. 36, pp. 792–799, May 2001.
- [11] B. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [12] P. E. Gronowski, W. J. Bowhill, R. P. Preston, M. K. Gowan, and R. L. Allmon, "High-performance microprocessor design," *IEEE J. Solid State Circuits*, vol. 33, pp. 676–686, May 1998.
- [13] I. A. Young, M. F. Mar, and B. Bhushan, "A 0.35 μm CMOS 3-880 MHz PLL N/2 clock multiplier and distribution network with low-jitter for microprocessors," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 1997, pp. 330–331.
- [14] D. W. Dobberpuhl, R. T. Witek, R. Allmon, R. Anglin, D. Bertucci, S. Britton, L. Chao, R. A. Conrad, D. E. Dever, B. Gieseke, S. M. N. Hassoun, G. W. Hoepfner, K. Kuchler, M. Ladd, B. M. Leary, L. Madden, E. J. McLellan, D. R. Meyer, and J. Montanaro, "A 200-MHz 64-b dual-issue CMOS microprocessor," *IEEE J. Solid State Circuits*, vol. 27, pp. 1555–1567, Nov. 1992.
- [15] B. J. Benschneider, A. J. Black, W. J. Bowhill, S. M. Britton, D. E. Dever, D. R. Donchin, R. J. Dupcak, R. M. Fromm, M. K. Gowan, P. E. Gronowski, M. Kantrowitz, M. E. Lamere, S. Mehta, J. E. Meyer, R. O. Mueller, A. Olesin, R. P. Preston, D. A. Priore, S. Santhanam, M. J. Smith, and G. M. Wolrich, "A 300 MHz 64-b quad-issue CMOS RISC microprocessor," *IEEE J. Solid State Circuits*, vol. 30, pp. 1203–1214, Nov. 1995.
- [16] P. E. Gronowski, W. J. Bowhill, D. R. Donchin, R. P. Blake-Campos, D. A. Carlson, E. R. Equi, B. J. Loughlin, S. Mehta, R. O. Mueller, A. Olesin, D. J. W. Noorlag, and R. Preston, "A 433 MHz 64-b quad-issue RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1687–1696, Nov. 1996.
- [17] D. W. Bailey and B. J. Benschneider, "Clocking design and analysis for a 600 MHz Alpha microprocessor," *IEEE J. Solid State Circuits*, vol. 33, pp. 1627–1633, Nov. 1998.

- [18] T. Xanthopoulos, D. W. Bailey, A. K. Gangwar, M. K. Gowan, A. K. Jain, and B. K. Prewitt, "The design and analysis of the clock distribution network for a 1.2 GHz Alpha microprocessor," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 2001, pp. 402–403.
- [19] M. R. Choudhury and J. S. Miller, "A 300 MHz CMOS microprocessor with multi-media technology," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 1997, pp. 170–171.
- [20] J. Schutz and R. Wallace, "A 450 MHz IA32 P6 family microprocessor," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 1998, pp. 236–237.
- [21] G. Geannopoulos and X. Dai, "An adaptive digital deskewing circuit for clock distribution networks," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 1998, pp. 400–401.
- [22] R. Senthinathan, S. Fischer, H. Rangchi, and H. Yazdanmehr, "A 650-MHz, IA-32 microprocessor with enhanced data streaming for graphics and video," *IEEE J. Solid-State Circuits*, vol. 34, pp. 1454–1465, Nov. 1999.
- [23] 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 J. Solid-State Circuits*, vol. 35, pp. 1545–1552, Nov. 2000.
- [24] P. Green, "A 1 GHz IA-32 microprocessor implemented on 0.18 μm technology with aluminum interconnect," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 2000, pp. 98–99.
- [25] N. A. Kurd, J. S. Barkatullah, R. O. Dizon, T. D. Fletcher, and P. D. Madland, "A multi-GHz clocking scheme for Pentium 4 microprocessor," *IEEE J. Solid State Circuits*, vol. 36, no. , pp. 1647–1653, Nov. 2001.
- [26] , C. J. Anderson, L. Sigal, K. L. Shepard, J. S. Liptay, J. D. Warnock, B. Curran, B. W. Krumm, M. D. Mayo, P. J. Camporese, E. M. Schwarz, M. S. Farrell, P. J. Restle, R. M. Averill III, T. J. Slegel, V. V. Huott, Y. H. Chan, B. Wile, T. N. Nguyen, P. G. Emma, D. K. Beece, C. T. Chuang, and C. Price, "A 400 MHz S/390 microprocessor," *IEEE J. Solid-State Circuits*, vol. 32, pp. 1655–1675, Nov. 1997.
- [27] G. Northrop, R. Averill, K. Barkley, S. Carey, Y. Chan, Y. H. Chan, M. Check, D. Hoffman, W. Huott, and J. H. Wuorinen, "609 MHz G5 S/390 microprocessor," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 1999, pp. 88–89.
- [28] T. McPherson, R. Averill, D. Balazich, K. Barkley, S. Carey, Y. Chan, Y. H. Chan, R. Crea, A. Dansky, R. Dwyer, A. Haen, D. Hoffman, A. Jatkowski, M. Mayo, D. Merrill, T. McNamara, G. Northrop, J. Rawlins, L. Sigal, T. Slegel, D. Webber, P. Williams, and F. Yee, "760 MHz G6 S/390 microprocessor exploiting multiple V_t and copper interconnects," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 2000, pp. 96–97.
- [29] R. M. Averill, K. G. Barkley, M. A. Bowen, P. J. Camporese, A. H. Dansky, R. F. Hatch, D. E. Hoffman, M. D. Mayo, S. A. McCabe, T. G. McNamara, T. J. McPherson, G. A. Northrop, L. Sigal, H. H. Smith, D. A. Webber, and P. M. Williams, "Chip integration methodology for the IBM S/390 G5 and G6 custom microprocessors," *IBM J. Res. Develop.*, vol. 43, pp. 681–706, Nov. 1999.
- [30] P. Hofstee, N. Aoki, D. Boerstler, P. Coulman, S. Dhong, B. Flachs, N. Kojima, O. Kwon, and K. Lee, "A 1 GHz single-issue 64 b PowerPC™ processor," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 2000, pp. 92–93.
- [31] P. J. Restle, C. A. Carter, J. P. Eckhardt, B. L. Krauter, B. D. McCredie2, K. A. Jenkins, A. J. Weger, and A. V. Mule', "The clock distribution of the Power4 microprocessor," in *Proc. IEEE Int. Solid State Circuits Conf.*, San Francisco, CA, Feb. 2002.
- [32] D. W. Boerstler, "A low-jitter PLL clock generator for microprocessors with lock range of 340–612 MHz," *IEEE J. Solid-State Circuits*, vol. 34, pp. 513–519, Apr. 1999.
- [33] D. W. Boerstler and K. A. Jenkins, "A phase-locked loop clock generator for a 1 GHz microprocessor," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, June 1998, pp. 212–213.
- [34] B. Razavi, *Design of Analog CMOS Integrated Circuits*. New York: McGraw-Hill, 2001.
- [35] J. W. Goodman, F. J. Leonberger, S. C. Kung, and R. A. Athale, "Optical interconnections for VLSI systems," *Proc. IEEE*, vol. 72, pp. 850–866, July 1984.
- [36] S. K. Tewksbury and L. A. Hornak, "Optical clock distribution in electronic systems," *J. VLSI Signal Process*, vol. 16, pp. 225–246, June-July 1997.
- [37] S. J. Walker and J. Jahns, "Optical clock distribution using integrated free-space optics," *Opt. Commun.*, vol. 90, pp. 359–371, June 1992.
- [38] B. D. Clymer and J. W. Goodman, "Timing uncertainty for receivers in optical clock distribution for VLSI," *Opt. Eng.*, vol. 27, pp. 944–954, Nov. 1988.
- [39] S. K. Patra, J. Ma, V. H. Ozguz, and S. H. Lee, "Alignment issues in packaging for free-space optical interconnects," *Opt. Eng.*, vol. 33, pp. 1561–1570, May 1994.
- [40] B. Lunitz and J. Jahns, "Tolerant design of a planar-optical clock distribution system," *Opt. Commun.*, vol. 134, pp. 281–288, Jan. 1997.
- [41] J.-H. Yeh, R. K. Kostuk, and K.-Y. Tu, "Board level H-tree optical clock distribution with substrate mode holograms," *J. Lightwave Technol.*, vol. 13, pp. 1566–1578, July 1995.
- [42] S. Tang and R. T. Chen, "1-to-42 optoelectronic interconnection for intra-multichip-module clock signal distribution," *Appl. Phys. Lett.*, vol. 64, pp. 2931–2933, May 1994.
- [43] C. Zhao and R. T. Chen, "Performance consideration of three-dimensional optoelectronic interconnection for intra-multichip-module clock signal distribution," *Appl. Opt.*, vol. 36, pp. 2537–2544, Apr. 1997.
- [44] P. J. Delfyett, D. H. Hartman, and S. Z. Ahmad, "Optical clock distribution using a mode-locked semiconductor laser-diode system," *J. Lightwave Technol.*, vol. 9, pp. 1646–1649, Dec. 1991.
- [45] Y. Li, J. Popelek, L. Wang, Y. Takiguchi, T. Wang, and K. Shum, "Clock delivery using laminated polymer fiber circuits," *J. Opt. A: Pure Appl. Opt.*, vol. 1, pp. 239–243, Mar. 1999.
- [46] S. Koh, H. W. Carter, and J. T. Boyd, "Synchronous global clock distribution on multi-chip modules using optical waveguides," *Opt. Eng.*, vol. 33, pp. 1587–1595, May 1994.
- [47] B. Bihari, J. Gan, L. Wu, Y. Liu, S. Tang, and R. T. Chen, "Optical clock distribution in supercomputers using polyimide-based waveguides," in *Proc. Optoelectronic Interconnects VI*, San Jose, CA, Jan. 1999, pp. 123–133.
- [48] R. T. Chen, L. Lin, C. Choi, Y. J. Liu, B. Bihari, L. Wu, S. Tang, R. Wickman, B. Picor, M. K. Hibbs-Brenner, J. Bristow, and Y. S. Liu, "Fully embedded board-level guided-wave optoelectronic interconnects," *Proc. IEEE*, vol. 88, pp. 780–794, June 2000.
- [49] J. S. Foresi, D. R. Lim, L. Liao, A. M. Agarwal, and L. C. Kimerling, "Small radius bends and large angle splitters in SOI waveguides," in *Proc. SPIE Silicon-Based Monolithic and Hybrid Optoelectronic Devices*, vol. 3007, San Jose, CA, Feb. 1997, pp. 112–118.
- [50] G. Masini, L. Colace, G. Assanto, H.-C. Luan, and L. C. Kimerling, "High-performance p-i-n Ge on Si Photodetectors for the near infrared: From model to demonstration," *IEEE Trans. Electron Devices*, vol. 48, pp. 1092–1096, June 2001.
- [51] L. Giovane, L. Liao, D. Lim, A. Agarwal, E. Fitzgerald, and L. C. Kimerling, " $\text{Si}_{10.5}\text{Ge}_{0.5}$ relaxed buffer photodetectors and low-loss polycrystalline silicon waveguides for integrated optical interconnects at $\lambda = 1.3 \mu\text{m}$," in *Proc. SPIE Silicon-Based Monolithic and Hybrid Optoelectronic Devices*, San Jose, CA, Feb. 1997, pp. 74–80.
- [52] T. D. Chen, A. M. Agarwal, L. M. Giovane, J. S. Foresi, L. Ling, D. R. Lim, M. T. Morse, E. J. Ouellette, S. H. Ahn, D. Xiaoman, J. Michael, and L. C. Kimerling, "Erbium-doped silicon light emitting devices," in *Proc. SPIE Light-Emitting Diodes: Research, Manufacturing, and Applications II*, San Jose, CA, Jan. 1998, pp. 136–145.
- [53] K. K. Lee, D. R. Lim, L. Hsin-Chiao, A. M. Agarwal, J. Foresi, and L. C. Kimerling, "Effect of size and roughness on light transmission in a Si/SiO₂ waveguide: Experiments and model," *Appl. Phys. Lett.*, vol. 77, pp. 1617–1619, Sept. 2000.
- [54] K. K. Lee, D. R. Lim, and L. C. Kimerling, "Fabrication of ultralow-loss Si/SiO₂ waveguides by roughness reduction," *Opt. Lett.*, vol. 26, pp. 1888–1890, Dec. 2001.
- [55] E. Griese, "A high-performance hybrid optical-electrical interconnection technology for high-speed electronic systems," *IEEE Trans. Adv. Packag.*, vol. 24, pp. 375–383, Aug. 2001.
- [56] F. Mederer, R. Jager, H. J. Unhold, R. Michalzik, K. J. Ebeling, S. Lechmacher, A. Neyer, and E. Griese, "3-Gb/s data transmission with GaAs VCSEL's over PCB integrated polymer waveguides," *IEEE Photon. Technol. Lett.*, vol. 13, pp. 1032–1034, Sept. 2001.
- [57] H. Franke and T. Sterkenburgh, "Patterning polymer surfaces by laser ablation for integrated optics," in *Proc. 4th Int. Conf. Properties and Applications of Dielectric Materials*, Brisbane, Qld., Australia, July 1994, pp. 208–210.
- [58] S. M. Schultz, E. N. Glytsis, and T. K. Gaylord, "Volume grating preferential-order focusing waveguide coupler," *Opt. Lett.*, vol. 24, pp. 1708–1710, Dec. 1999.
- [59] S. M. Schultz, "High-efficiency volume grating coupler," Ph.D. dissertation, Georgia Inst. of Technology, 1999.
- [60] S. Ganguly, D. Lehter, and S. Pallela, "Clock distribution methodology for PowerPC™ microprocessors," *J. VLSI Signal Process.*, vol. 16, pp. 181–189, June-July 1997.
- [61] K. M. Carrig, A. M. Chu, F. D. Ferraiolo, J. G. Petrovick, P. A. Scott, and R. J. Weiss, "A clock methodology for high-performance microprocessors," *J. VLSI Signal Process.*, vol. 16, pp. 217–224, June-July 1997.
- [62] D. J. Hathaway, R. R. Habra, E. C. Schanzenbach, and S. J. Rothman, "Circuit placement, chip optimization, and wire routing for IBM IC technology," *J. VLSI Signal Process.*, vol. 16, pp. 191–198, June-July 1997.

- [63] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, pp. 728–749, June 2000.
- [64] K. Glukh, J.-H. Lipian, R. Mimna, P. S. Neal, R. Ravikiran, L. F. Rhodes, R. A. Shick, and X.-M. Zhao, "High-performance polymeric materials for waveguide applications," in *Proc. SPIE Linear, Nonlinear, and Power-limiting Organics*, San Diego, CA, Aug. 2000, pp. 43–53.
- [65] C. L. Schow, J. D. Schaub, R. Li, J. Qi, and J. C. Campbell, "A 1-Gb/s monolithically integrated silicon NMOS optical receiver," *IEEE J. Select. Topics Quantum Electron.*, vol. 4, pp. 1035–1039, Nov. 1998.
- [66] T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, and L. M. F. Chirovsky, "Optical receivers for optoelectronic VLSI," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 106–115, Apr. 1996.
- [67] A. V. Krishnamoorthy and D. A. B. Miller, "Scaling optoelectronic-VLSI circuits into the 21st century: A technology roadmap," *IEEE J. Select. Topics Quantum Electron.*, vol. 2, pp. 55–76, Apr. 1996.
- [68] D. A. V. Blerkom, C. Fan, M. Blume, and S. C. Esener, "Transimpedance receiver optimization for smart pixel arrays," *J. Lightwave Technol.*, vol. 16, pp. 119–126, Jan. 1998.



Anthony V. Mule (S'97) received the B.S.E.E. degree from the University of Illinois at Urbana-Champaign in 1996. He is currently pursuing the Ph.D. degree in electrical engineering at the Georgia Institute of Technology, Atlanta.

He is a Member of the Gigascale Integration (GSI) group. His research interests include integrated optics, diffractive optics, optoelectronics, optoelectronic packaging, and optical materials. He has participated in internship activities at both Intel Corporation and the IBM T. J. Watson Research

Center.

Mr. Mule is a student member of the Optical Society of America. He received the President's Fellowship at the Georgia Institute of Technology.



Elias N. Glytsis (S'81–M'81–SM'91) joined the Faculty of the School of Electrical and Computer Engineering, Georgia Institute of Technology, as an Assistant Professor in January 1988 and has been a Professor since 2000. His current research interests are in electromagnetic theory of holographic and diffractive grating couplers, diffractive optical interconnections, optoelectronic devices, semiconductor quantum electron wave devices and applications, intersubband emitters and detectors, high-spatial frequency grating surfaces, ferroelectric liquid

crystal waveguides, design of binary optical elements, optimization, integration software, and electromagnetic problems in power systems. He has published more than 90 journal publications and more than 75 conference papers. He has received eight U.S. patents. He was Co-Guest Editor of two special issues of the *Journal of the Optical Society of America* on grating diffraction. He was a Topical Editor of the *Journal of the Optical Society of America A* on Scattering and Grating Diffraction (1992–1997). He also was a Guest Editor of the *Microelectronics Journal* of a special issue on quasi-bound states in quantum heterostructure devices.

Dr. Glytsis is a Fellow of the Optical Society of America and a member of LEOS, Sigma Xi, Eta Kappa Nu, and the Greek Society of Professional Engineers.



Thomas K. Gaylord (S'65–M'70–SM'77–F'83) received the B.S. degree in physics and the M.S. degree in electrical engineering from the University of Missouri-Rolla and the Ph.D. degree in electrical engineering from Rice University, Houston, TX.

Presently, he is with the Georgia Institute of Technology, where he is Julius Brown Chair and Regents' Professor of electrical and computer engineering. He is the author of some 350 technical publications and 25 patents in the areas of diffractive optics, optoelectronics, and semiconductor devices.

Dr. Gaylord is a Fellow of the Optical Society of America and the American Association for the Advancement of Science. He received the Curtis W. McGraw Research Award from the American Society for Engineering Education; the IEEE Centennial Medal; the IEEE Graduate Teaching Award; the Georgia Tech Outstanding Teacher Award; and the Engineer of the Year Award from the Georgia Society of Professional Engineers.



James Meindl (M'56–SM'66–F'68–LF'97) received the Bachelor's, Master's, and Doctor's degrees in electrical engineering from Carnegie Institute of Technology, Carnegie-Mellon University, Pittsburgh, PA.

He is Director of the Joseph M. Pettit Microelectronics Research Center and the Joseph M. Pettit Chair Professor of Microelectronics at the Georgia Institute of Technology, Atlanta. He is also Director of the Interconnect Focus Center, a multiuniversity research effort managed jointly by the Microelectronics

Advanced Research Corporation and the Defense Advanced Research Projects Agency for the U.S. Department of Defense. His current research interests focus on physical limits on gigascale integration. He was Senior Vice President for Academic Affairs and Provost of Rensselaer Polytechnic Institute, Troy, NY, from 1986 to 1993. He was with Stanford University, Stanford, CA, from 1967 to 1986 as the John M. Fluke Professor of Electrical Engineering, Associate Dean for Research in the School of Engineering, Founding Director of the Center for Integrated Systems, Director of the Electronics Laboratories, and Founding Director of the Integrated Circuits Laboratory.

Prof. Meindl is a Fellow of the American Association for the Advancement of Science. He is a member of the American Academy of Arts and Sciences and the National Academy of Engineering and its Academic Advisory Board. He received the Benjamin Garver Lamme Medal from ASEE in 1991, the IEEE Education Medal in 1990, and the IEEE Solid-State Circuits Medal in 1989. He has also received the IEEE Electron Devices Society's J.J. Ebers Award, the 1997 Hamerschlag Distinguished Alumnus Award from Carnegie-Mellon University, and five outstanding paper awards from the IEEE ISSCC. He also received the 1999 SIA University Research Award, the IEEE Third Millennium Medal, and, most recently, the Georgia Institute of Technology 2001 Distinguished Professor Award.