

Invited Paper

Trends in Emerging On-Chip Interconnect Technologies

SUDEEP PASRICHA^{†1,*1} and NIKIL DUTT^{†1}

In deep submicron (DSM) VLSI technologies, it is becoming increasingly harder for a copper based electrical interconnect fabric to satisfy the multiple design requirements of delay, power, bandwidth, and delay uncertainty. This is because electrical interconnects are becoming increasingly susceptible to parasitic resistance and capacitance with shrinking process technology and rising clock frequencies, which poses serious challenges for interconnect delay, power dissipation and reliability. On-chip communication architectures such as buses and networks-on-chip (NoC) that are used to enable inter-component communication in multi-processor systems-on-chip (MPSoC) designs rely on these electrical interconnects at the physical level, and are consequently faced with the entire gamut of challenges and drawbacks that plague copper-based electrical interconnects. To overcome the limitations of traditional copper-based electrical interconnects, several research efforts have begun looking at novel interconnect alternatives, such as on-chip optical interconnects, wireless interconnects and carbon nanotube-based interconnects. This paper presents an overview and current state of research for these three promising interconnect technologies. We also discuss the existing challenges for each of these technologies that remain to be resolved before they can be adopted as replacements for copper-based electrical interconnects in the future.

1. Introduction

Multi-processor system-on-chip (MPSoC) designs today are becoming more and more complex, with rapidly increasing levels of component integration. Emerging MPSoCs today typically have tens to hundreds of components (microprocessors, memories, peripherals, etc.) on a single chip. These components invariably need to communicate with each other during application execution. It is the responsibility of the on-chip interconnect architecture to ensure that the multiple, co-existing data streams on the chip are correctly and reliably routed from the

source components to their intended destinations.

With the rising number and variety of components being integrated into MPSoC designs, communication between on-chip components is playing an increasingly critical role in ensuring that application performance constraints are satisfied. According to ITRS 2005 predictions^{1),2)}, the gap between interconnection delay and gate delay will increase to 9:1 at the 65 nm technology. This is in sharp contrast to the 2:1 gap between interconnection delay and gate delay at the 180 nm technology. This indicates that communication, and not computation, will be the key performance bottleneck in deep submicron (DSM) technologies³⁾. In addition, total wire length on a chip is expected to grow to 2.22 km/cm² by the year 2010²⁾. Another observation is the increase of power dissipation due to the charging and discharging of interconnection wires on a chip. According to Refs. 1), 2), the interconnect will consume about 50 times more power than logic circuits. All of these trends indicate that for MPSoC designs in the DSM era, the performance, power consumption, cost and area will be much more influenced by the on-chip interconnect architecture than the gates on the chip.

To cope with the increasing MPSoC performance requirements in the DSM era, on-chip interconnect architectures have also undergone an evolution in complexity — from shared buses, to hierarchical shared buses, and on to bus matrix (or crossbar bus) architectures. Several bus-based interconnect architecture standards exist today, such as AMBA AHB⁴⁾ and AXI⁵⁾, IBM CoreConnect⁶⁾, Sonics Smart Interconnect⁷⁾ and STMicroelectronics' STBus⁸⁾. These standard architectures provide designers with a variety of configuration options and find widespread use in MPSoC designs today. Researchers have recently proposed using Networks-on-Chip (NoC)^{9),10)} as interconnect architectures for future large-scale MPSoC designs. Unlike traditional bus-based on-chip interconnect architectures, NoCs use packets to route data from the source to the destination component, via a network fabric that consists of switches (routers) and interconnection links (wires).

Both bus-based and NoC architectures rely on copper based electrical interconnects at the physical level to transfer information between components. Unfortunately, in deep submicron (DSM) VLSI technologies, it is becoming increasingly harder for copper based electrical interconnects to satisfy the design require-

^{†1} University of California, Irvine
^{*1} Presently with Colorado State University

ments of delay, power, bandwidth, and delay uncertainty¹¹⁾. Indeed, the situation is likely to become worse for future giga- and tera-scale electronic systems. The resistance of copper interconnects, in current and imminent technologies is increasing rapidly under the combined effects of enhanced grain boundary scattering, surface scattering, and the presence of a highly resistive diffusion barrier layer¹²⁾. Copper interconnects also constitute up to 70% of the total on-chip capacitance, and are major sources of power dissipation. The semiconductor industry has made major research and development investments in alternative (e.g., low- k dielectric) materials in response to this urgent need. However, low- k materials suffer from poor mechanical and thermal properties²⁾. The steep rise in parasitic resistance and capacitance of copper interconnects poses serious challenges for interconnect delay (especially at the global level), power dissipation, and interconnect reliability⁷⁰⁾.

Consequently, it has become imperative to look beyond copper interconnects and explore different interconnect technologies. According to the ITRS roadmap¹⁾, interconnect innovation is the key to satisfying performance, reliability, and power requirements in the long term. Future interconnect technologies must support ultra-high data rates (e.g., greater than 100 Gbps/pin), be scalable enough to support tens to hundreds of concurrent communication streams, and involve fabrication techniques that are compatible with mainstream MPSoC and system-in-package (SiP) technologies. In this paper, we present an overview of three such emerging interconnect technologies that promise to overcome the limitations of copper interconnects. Section 2 describes optical interconnects, that make use of light and an on-chip optical medium to transfer data. Section 3 presents RF/wireless interconnects that forego the need for metal wires, and instead transfer data on a chip wirelessly, using transmitting and receiving antennas integrated on the same chip. Finally, Section 4 discusses carbon nanotubes, which have been proposed as an evolutionary replacement for copper interconnects, just like copper previously replaced aluminum as the interconnect material of choice in electronic systems.

2. Optical Interconnects

While board-to-board and chip-to-chip optical interconnects have been

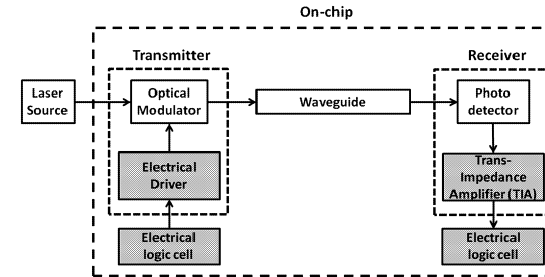


Fig. 1 Block diagram of on-chip OI systems.

proposed and are actively under development^{13),14)}, the feasibility of using on-chip optical interconnects (OIs) is an open research problem that is today beginning to attract immense interest from academia and industry alike. It is claimed that OIs will be suitable replacements for global on-chip interconnects, where they can be advantageous because of their inherent low propagation delay, low crosstalk, and a near constant power profile over long distances¹⁵⁾. OIs offer many advantages over traditional electrical (copper-based) interconnects: (i) they can support enormous intrinsic data bandwidths in the order of several Gbps using only simple on-off modulation schemes, (ii) they are relatively immune to electrical interference due to crosstalk, and parasitic capacitances and inductances, (iii) their power dissipation is completely independent of transmission distance at the chip level, and (iv) routing and placement is simplified since it is possible to physically intersect light beams with minimal crosstalk. Once a path is acquired, the transmission latency of the optical data is very small, depending only on the group velocity of light in a silicon waveguide: approximately 6.6×10^7 m/s, or 300 ps for a 2 cm path crossing a chip. After an optical path is established, data can be transmitted end to end without the need for repeating or buffering, which can lead to significant power savings.

2.1 OI System Overview

Figure 1 shows a block diagram of an on-chip OI system. With the exception of the laser light source, all of the OI components are integrated on the chip. Despite a lot of achievements in the area of optical gain in silicon over the last few years^{16),17)}, a high-speed, electrically driven, on-chip monolithic laser light

source still remains to be realized.

The transmission of data using an on-chip OI requires an electro-optical modulator and an electrical driver circuit. The laser source provides light to the modulator, which transduces electrical data supplied from the electrical driver into a modulated optical signal. Several high-speed electro-optical modulators^{18),19)} have been proposed that can change the refractive index or the absorption coefficient of an optical path when an electrical signal is injected. The two most popular modulators in literature are the Mach-Zehnder interferometer-based silicon modulators¹⁸⁾ and microresonator-based P-I-N diode type modulators¹⁹⁾. MOS capacitor structures such as the Mach-Zehnder interferometer-based silicon modulators have higher modulation speeds (several GHz) but a large power consumption and greater silicon footprint (around 10 mm). On the other hand, microresonator based P-I-N diode type modulators are compact in size (10–30 μm) and have low power consumption, but possess low modulation speeds (several MHz). The performance of a modulator is dependent on the on-to-off light intensity ratio, referred to as the *extinction ratio*, which depends on the electrical input signal strength. A higher extinction ratio is desirable for proper signal detection, while a poor one may cause transmission errors. An extinction ratio greater than 10 dB has been recently reported with high input signal swing²⁰⁾, which is high enough to enable proper signal detection without causing any transmission errors. Modulator size is also an important consideration for integrated applications, and significant efforts have been made to realize compact-sized modulators, such as the circular shaped 10 μm ring-modulators²⁰⁾. The modulator driver consists of a series of inverter stages that drive the modulator's capacitive load.

The light signal from the transmitter is routed to the destination through a waveguide. The refractive index of the waveguide material has a significant impact on the bandwidth, latency, and area of an optical interconnect. *Silicon* and *polymer* are two of the most promising materials for on-chip waveguide realization. These waveguides involve a trade-off in propagation speed and bandwidth. The smaller refractive index of polymer waveguides results in higher propagation speed. However, the polymer waveguide requires a larger pitch than silicon, which reduces bandwidth density (i.e., number of bits transmitted/unit surface area). Modulators for polymer waveguides are also typically bulkier (requiring

higher voltage drive and higher frequency operation) than modulators for silicon waveguides. This makes it harder to use polymer waveguides in on-chip optical interconnects. Polymer waveguides are however feasible in a transmission system based on VCSELs (Vertical Cavity Surface Emitting Laser)²¹⁾, where the modulator is not required. A VCSEL based solution may however lead to an increase in on-chip power consumption as a result of the complex on-chip flip-bonded laser sources. Additionally, light is emitted vertically and needs to be transferred to the horizontal chip surface in this scenario, requiring integrated mirrors and sophisticated lithographic techniques.

The optical receiver is responsible for converting the optical signal received from the waveguide into an electrical signal at the destination. It consists of a photo-detector and a trans-impedance amplifier (TIA) stage. In applications involving simultaneous transmission at different wavelengths per waveguide using wave division multiplexing (WDM), the receiver also requires a wave selective filter for each received wavelength. The P-I-N diode²²⁾ is an example of one of the more popular photo-detectors in literature. The *quantum efficiency* of a photo-detector is an important figure of merit for the system. A high value for quantum efficiency means lower losses when converting optical information into an electrical form. One of the key trade-offs in the design of a photo-detector is between detector speed and quantum efficiency. Recently, inter-digitated metal-semiconductor-metal (MSM) receivers have attracted attention due to their fast response and excellent quantum efficiency. High-speed, low-power inter-digitated MSM Ge and SiGe photo-detectors operating at telecommunication wavelengths were described in Ref. 23). Detector size is also an important criteria for both compactness and next stage capacitance. Typically, the detector has large base capacitance, which poses a design challenge for the high speed gain stages following it. The TIA stage converts photo-detector current to a voltage which is thresholded by subsequent stages to digital levels²⁴⁾. To achieve high-gain and high-speed detection, an analog supply voltage higher than the digital supply voltage may be required, thereby dissipating higher power.

2.2 Comparison between on-chip EI and OI

Figure 2 shows a comparison of signal propagation delay in copper-based electrical interconnects (EI) and the two common OI waveguides — silicon and




Fig. 2 Propagation delay of silicon and polymer waveguides as compared to EIs¹⁵⁾.

polymer. Low refractive index polymer and high refractive index silicon waveguides are selected for the comparison, as they represent two opposite types of optical waveguides in terms of signal propagation delay and crosstalk. Regardless of the waveguide material, optical interconnects provide a lower propagation delay than electrical interconnects. This is because optical signal propagation is intrinsically faster than electrical signal propagation due to the absence of RLC impedances. In order to exploit the propagation delay advantage offered by optical waveguides, an electrical signal must however first be converted into an optical signal and then back into an electrical signal. This conversion has a fixed delay associated with it, which is nearly independent of the interconnect length for a given technology. OIs will therefore have a delay advantage over EIs if the waveguide propagation delay dominates the overall delay. It has been estimated¹⁵⁾ that the combined transmitter and receiver delay should be lower than 280–370 ps for polymer waveguides and 180–270 ps for silicon waveguides, to have an advantage over EIs. In addition to delay, power is another metric for which OIs need to have an advantage over EIs. For OIs that span the chip length, the total power consumption should be less than 17–18 mW to have an advantage over EI power consumption.

Finally, if OIs are to eventually replace EIs, then in addition to lower signal delay and lower power consumption, they must also possess superior bandwidth density compared to EIs. Bandwidth density is a metric that characterizes

information throughput through a unit cross section of an interconnect. Experiments have shown that with advances in process technology, EIs will exploit more efficient repeaters, resulting in a single wavelength OI being inferior to a delay-optimized EI in terms of bandwidth density. So the question arises: why would anyone use OIs? To improve the bandwidth density in OIs, wavelength division multiplexing (WDM)^{15),25)} can be used in the optical waveguides. It has been shown that polymer core waveguides require a higher WDM to match the EI bandwidth density, but allow for a larger conversion delay overhead. Silicon-core waveguides on the other hand permit lower WDM but require faster conversion (i.e., faster transmitters and receivers).

2.3 On-chip OI Research

In recent years, a few researchers have studied how on-chip optical networks can be used for clock-tree networks^{26)–30)}. Clock networks are characterized by long interconnect lengths spanning the entire chip. Using optical links as a replacement for electrical clock networks can reduce clock distribution skew on the chip in the multi-GHz operating range. Since the clock source (i.e., the light source) can be external to the chip, it can remove energy consumption constraints and also alleviate the difficult problem of integrating III-V optoelectronic emitters on top of Si CMOS circuits²⁶⁾. A full CMOS-compatible process including Si-photo-detectors is feasible even if the quantum efficiency of CMOS-compatible silicon photo-detectors is small^{31),32)}. Intel researchers have claimed²⁷⁾ that WDM can enable optics to achieve high bandwidth and low latency for global signaling. However, they concluded that until efficient high-speed and low capacitance CMOS-compatible modulators and detectors, and practical schemes for implementing WDM are realized, there is little power, jitter, or skew improvement from using optics in clock distribution. A similar claim was made by Chen, et al.²⁸⁾, where the authors argued that since most of the skew and power of clock signaling arises in local clock distribution, there is no significant skew and power advantages in using an optical solution. Ackland, et al.²⁹⁾ presented a study showing how an H-tree electrical clock network does not scale well, resulting in unacceptable levels of skew and jitter, compared to an optical tree. The authors claimed that while there are many technical challenges in implementing an optical tree, in principle the best solution might be a hybrid




Fig. 3 Optical Ring Bus (ORB) on-chip communication architecture for MPSoCs³⁵⁾.

tree network, in which the front end is implemented optically, while the backend consists of a large number of small electrical trees.

A few recent works have proposed using optical links in an on-chip network-on-chip (ONoC)^{30),33),34)}. O’Conner³⁰⁾ and Briere, et al.³³⁾ gave a high level overview of a 4×4 multi-stage optical network on chip. Shacham, et al.³⁴⁾ presented an analysis of topology, routing algorithms, path-setup and teardown procedures, and deadlock avoidance for optical NoCs, and described simulation results for a 2D folded torus optical NoC architecture, that can theoretically result in high bandwidth, low power intra-chip communication. However, the high power overhead of electrical routers and opto-electric/electro-optic conversion at the interface of each component, as well as a lack of availability of wideband photonic switching elements makes realizing these hybrid opto-NoC architectures a challenging proposition in the near future.

Recently, in Ref. 35) we proposed the ORB (Optical Ring Bus) on-chip communication architecture for emerging MPSoC designs. ORB is a novel opto-electric communication architecture that uses an optical ring bus as a global interconnect between computation clusters (i.e., islands of multiple cores performing dedicated tasks; **Fig. 3**) and traditional copper-based local interconnect within clusters. The optical ring bus transfers data between clusters on a chip, while preserving the standard bus protocol interface (e.g., AMBA AXI) for inter- and intra-cluster communication. ORB consists of four major building blocks: (i) an *off-chip laser* from which light is coupled onto the chip using optical fibers; (ii) *transmitters*

that convert electrical signals into optical waves and are made of hybrid Mach-Zehnder interferometer/micro resonator based P-I-N diode modulators, driven by a series of tapered inverters (i.e., drivers); (iii) an *optical waveguide*, made from a low refractive index polymer material (that has lower signal propagation delay than a silicon waveguide), having a ring shape (to avoid sharp turns that can lead to significant signal loss) and implemented on a dedicated layer, that is responsible for transporting data via light signals from the source modulator to the destination receiver; and (iv) *receivers*, consisting of a photo-detector to convert the light signal into an electrical signal, and trans-impedance amplifier (TIA) circuits to amplify the resulting analog electrical signal to a digital voltage level. To improve bandwidth density of the optical interconnect, wavelength division multiplexing (WDM) was used to transmit data on multiple wavelength channels on the same waveguide. Simulation studies have shown that the ORB communication architecture dissipated significantly lower power (more than a $10\times$ reduction) and improved overall performance (more than $2\times$) compared to traditional pipelined, all-electrical global interconnects, across the 65–22 nm CMOS technology nodes, for several networking MPSoC design case studies.

2.4 Challenges

While recent research efforts have shown that optical interconnects can be viable candidates for global on-chip interconnects in future MPSoC designs, several open problems remain that need to be resolved in the coming years:

(i) *Efficient transmitter and receiver components*: High speed, low power, and small feature-size electro-optical modulators and photo-detector receivers need to be developed, that have a combined delay and power dissipation which is lower than the threshold required to be advantageous over EIs, and replace them in future technologies. There is immense research interest in designing efficient electro-optical modulators and receivers, and some interesting results are emerging. For instance, recently, Mach-Zehnder electro-optic modulators with an ultra-compact length of 100 to 200 μm , having low power consumption and high modulation efficiency were presented in Ref. 36).

(ii) *Integrated on-chip light source*: The number of materials and processes available for optical interconnect fabrication is limited to those technologies that are compatible with microelectronics. Currently, this limitation results in the

absence of a monolithic on-chip light source. While several exciting scientific achievements have been published in the area of optical gain in silicon, high speed, electrically driven monolithic light sources have remained elusive. Innovative solutions, such as the Indium Phosphide Hybrid Silicon Laser from Intel and UCSB³⁷⁾ may however solve this problem in the future.

(iii) *Temperature management*: On-chip optical interconnect modules are very sensitive to temperature variations. Designers need to ensure that viable operating temperatures for components are maintained, or design new OI structures that are not so sensitive to temperature. Either an active or passive optical control method³⁸⁾ will be required to maintain stable device operation.

3. RF/Wireless Interconnects

Another revolutionary approach to overcome the limitations of traditional electrical interconnects in future DSM technologies is to make use of *active* RF/wireless interconnects. The main idea in such a system is to replace on-chip wires with integrated on-chip antennas communicating via electromagnetic waves. Data in such a scheme would be converted from baseband (i.e., digital) to RF/microwave signals and transmitted either through free space or guided mediums³⁹⁾. RF/wireless interconnects have the potential to support high data bandwidths, effectively handle concurrent communication streams from multiple cores, and provide online firmware network reconfigurability (i.e., dynamic rewiring using software instructions). These wireless interconnects can not only reduce the wires in integrated circuits, but also be used to replace I/O pins.

Free space signal broadcasting and reception is a common practice in modern wireless systems due to its low cost implementation and excellent channeling capabilities. However, free space transmission and reception of RF/microwave signals requires an antenna size that is comparable to its wavelength. This is a problem because even at near 100 GHz operating and cut-off frequencies in the future¹⁾, the aperture size of the antenna for efficient transmission will be in the order of 1 mm², which is too large to implement in future VLSI designs. Microwave transmission in guided mediums such as microstrip transmission line (MTL) or coplanar waveguide (CPW) has a low attenuation up to at least 200 GHz, requires a smaller antenna size and is therefore a more viable alternative




Fig. 4 RF/wireless interconnect with multiple transmitters and receivers³⁹⁾.

to free space transmission/reception for intra-chip communication. Simulation results³⁹⁾ have shown that a 1 cm long CPW experiences extremely low loss (-1.6 dB at 100 GHz), and low dispersion (less than 2 dB for a frequency range 50–150 GHz). This is in contrast to a -60 dB and -115 dB loss per centimeter at 100 GHz, and a frequency dispersion of 30–40 dB across the same frequency range for conventional electrical interconnects. Thus microwave transmission over MTL or CPW has a clear advantage over conventional electrical wire based transmission, especially for global interconnects in future VLSI designs running in the multi-GHz range.

The choice of transmission and receiving components for an RF/wireless interconnect system is an important one, in order to ensure compatibility with future VLSI designs³⁹⁾. Since the distance of global on-chip interconnects is relatively short (a few centimeters), the conventional *far field* antenna can be substituted for much smaller *near field* capacitive couplers. The width of the center/top conductor of the CPW/MTL is typically 10–100 μm , depending on its characteristic impedance (25–100 Ω) and other signal transmission requirements⁴⁰⁾. This size makes CPW or MTL more likely to be used as an off-chip but in-package transmission medium, which is shared by multiple VLSI I/Os. Based on these considerations, **Fig. 4** shows the structure of a typical RF/wireless interconnect system. The system acts as a miniature wireless LAN (local area network), consisting of VLSI I/Os as users, capacitive couplers as near field antennas, RF circuits for transceivers, and a uniform and homogeneous MTL or CPW channel (with characteristic impedance Z_0) as a shared broadcasting medium. Output signals are *up-linked* to MTL or CPW via transmission capacitive couplers (TX_i),

then down-linked via receiving capacitive couplers (RX_i) to input ports. Since the channel is bidirectional, both its ends are terminated with Z_C to avoid signal reflections. With orthogonal-coded and/or frequency-filtered RF transceivers, a passive MTL or CPW becomes suitable for relaying ultra broadband signals up to 150 GHz^{40),41)}.

3.1 Simultaneous Communication in RF/wireless Interconnects

Frequency division multiple access (FDMA), code division multiple access (CDMA), or a combination of the two (FDMA/CDMA) can be used to achieve simultaneous communications in RF/wireless interconnects.

In an FDMA interconnect, frequency bands of I/O channels can be allocated between 5–105 GHz with a bandwidth of approximately 5–20 GHz in each channel and sustaining a minimum data rate of 5–40 Gb/s depending on the modulation scheme. Whereas in a traditional electrical interconnect, only the lowest frequency band (i.e., baseband) is occupied by the signals, RF modulated frequency bands can improve data bandwidth by transmitting data over multiple frequency bands. Each data stream at the transmitter is multiplied by a sinusoidal carrier, and the resulting signal is filtered through a band pass filter (BPF), and then eventually coupled into the CPW (or MTL). At the receiver, the received signals are boosted by a preamplifier and then demodulated. Subsequently, the signal is fed into a threshold comparator to recover the original data bits from the transmitter⁴²⁾. The FDMA transmitter and receiver architectures consist of I/O transceivers, frequency synthesizers, threshold comparators, frequency/code mixers, and bandpass filters. The receiver is made up of preamplifiers, mixers, and frequency synthesizers. Preamplifiers with 20–30 dB gains are required for input signal amplification. Balanced or double balanced active mixers, such as the Gilbert cell, may be used for modulation and demodulation functions. To simplify receiver implementation, a non-coherent detection scheme such as frequency shift keying (FSK) can be a suitable choice. Low loss and high selective band pass filters that are needed in FDMA interconnects require tunable and high Q ^{*1} inductors that are hard to realize due to their high

*1 The quality factor (or Q) of an inductor is the ratio of its inductive reactance to its resistance at a given frequency, and is a measure of its efficiency. The higher the Q factor of the inductor, the closer it approaches the behavior of an ideal, lossless, inductor.




Fig. 5 Types of on-chip antennas (a) linear, (b) meander, (c) zig-zag, (d) folded⁴⁷⁾.

energy loss to the conductive silicon substrate. This can be partially resolved by using a transformer-type inductor design where the lost energy is recovered via a secondary inductor with delayed phase angles to attain high inductance and tunability⁴³⁾. Recent progress in MEMS (micro-electro-mechanical systems) have shown promise for high Q silicon resonators and filters in micrometer wave frequencies⁸³⁾.

In a CDMA interconnect^{39),44),45)}, baseband CDMA signals are typically modulated with RF carriers. The data stream from a transmitter is first converted into a spread spectrum signal by orthogonal Walsh codes (or PN codes) and then modulated with its RF sinusoidal carrier. Subsequently the resulting signals from the transmitters are capacitively coupled into a superposed multilevel signal on the shared CPW (or MTL) and transmitted to receivers. At the receiver end, coherent demodulation⁴²⁾, sequence timing acquisition and tracking⁴⁶⁾, and de-spreading by the same Walsh code are applied to the received superposed signal to recover the original data. Unlike the hardware-oriented FDMA interconnect, the CDMA interconnect can be easily reconfigured by changing spreading codes through software commands.

To improve data rate, a combined CDMA/FDMA³⁹⁾ access system (also referred to as multicarrier CDMA) is also a possibility, with frequency bands being divided by using different carriers while data is spread within each frequency band by using orthogonal codes.

3.2 On-chip Antenna and Transmission Path

The on-chip antennas fabricated on substrates are categorized as printed antennas⁴⁷⁾, which include microstrip, dipole and loop antennas. **Figure 5** shows the popular linear dipole, and folded dipole antennas that have a compact implementation and are thus suitable as on-chip integrated antennas. The zig-zag




Fig. 6 Antenna transmission paths⁵¹⁾.

dipole antenna emerged after Nakano, et al.⁴⁸⁾ showed that a halfwave dipole antenna whose arms were bent rectangularly at its central points, maintained the basic characteristics of conventional linear dipole with the same length. The loop antenna is another suitable candidate for on-chip integrated antennas. It was shown to have an isotropic radiation pattern and good signal-to-noise ratio when the receiver is placed close to the loop antenna⁴⁹⁾. Wang, et al.⁴⁷⁾ showed that combining different antenna structures such as folded and meander can provide a higher power gain and a more compact on-chip antenna structure.

There has been some recent research on using the on-chip silicon substrate/dielectric layer as a transmission path, instead of using off-chip and in-package CPW or MTL guided mediums^{50),51)}. At 24 GHz, the wavelength of electromagnetic waves in silicon is 3.7 mm. This implies that a quarter wave antenna needs to be only about 0.9 mm in silicon. This, in conjunction with increased chip sizes of $2\text{ cm} \times 2\text{ cm}$, has made the integration of antennas for wireless on-chip communication feasible. **Figure 6** shows the possible paths for signal propagation between two on-chip integrated antennas. There is a direct path through air and paths through the silicon substrate. The paths through silicon substrate include a path formed by refraction through the SiO_2 layer and reflection at the interface between the silicon substrate and the underlying dielectric layer (AlN), and a path refracted through the SiO_2 and silicon layers and reflected by the metal chuck that emulates a heatsink in the back of a die. The signals propagating on these paths constructively and destructively interfere. There are also

multiple reflected paths, but because of longer path lengths, the signals propagated on these paths suffer from greater attenuation and therefore these paths are not included. It was found in Ref. 51) that by increasing the AlN thickness, destructive signal interference is significantly reduced. Benech, et al.⁵²⁾ conducted a similar feasibility study and found that higher resistivity substrates are better suited for wireless communication. They showed that antennas on lower resistivity silicon substrate present gains of -30 dB , while antennas on SOI (Silicon On Insulator) substrates of higher resistivity present gains of -15 dB at frequencies around 30 GHz.

3.3 On-chip RF/wireless Interconnect Research: Clock Networks

It has been proposed that RF/wireless interconnects can be used in clock networks to reduce signal skew⁵³⁾. Test chips have been created to demonstrate the feasibility of RF/wireless on-chip interconnects and their use in clock networks. For instance, on a wafer, a 15 GHz transmitted signal 2.2 cm away from a clock receiver with an on-chip antenna was shown to have been successfully picked up by the receiver and amplified to generate a digital output signal⁵⁴⁾. The receiver and transmitting antennas were fabricated using a $0.18\text{ }\mu\text{m}$ CMOS process⁵⁰⁾. Another demonstration of wireless clock was presented by Floyd, et al.⁵⁵⁾ at the frequency of 7.4 GHz and with a transmission gain of -49 dB between the transmitting and the receiving antenna placed at a distance of 3.3 mm from each other.

3.4 Challenges

Unlike on-chip optical interconnects (Section 2), on-chip RF/wireless interconnect technology is still very much in its early stages of evolution. Much research is still needed to resolve the open problems in the area^{39),51)}. Some of these issues are presented below:

(i) *Packaging and Interference Issues:* The most difficult problem anticipated for this technology is dealing with the packaging effects, which can add numerous metal structures that could potentially interfere with RF/wireless operation on a chip. Metal structures near antennas can change input impedances and phase of received signals. Design guidelines to exclude the interference structures which significantly change the input impedance and techniques to correct the phase changes are being developed⁵⁶⁾. Another concern is the interference

effects between the transmitted/received signal and switching noise of nearby circuits^{57),58)}. Some promising initial results were presented by Branch, et al.⁵⁰⁾, where a sine wave generated in a transmitter was transmitted through an on-chip antenna and the wave was picked up by a receiver on the same chip about 4 mm away, for an on-chip silicon substrate/dielectric transmission path. It was concluded that on-chip wireless interconnects can function correctly for more complex systems, provided the interference concerns are addressed.

(ii) *Ultra High Frequency Requirements:* For antenna sizes (and related RF interconnect circuit sizes) to be feasible enough for on-chip fabrication, RF circuits must operate in the ultra high frequency domain, i.e., in the hundreds of GHz range. This makes them unsuitable for applications in the very near future that will not be able to achieve such high frequencies. However, due to technology scaling, CMOS circuits operating in the multi GHz and higher ranges are becoming gradually feasible. It is estimated^{1),51)} that by the year 2015, it will become possible to implement RF circuits operating at 200–250 GHz, which will enable low-cost wireless interconnect implementation.

(iii) *Power Overhead:* If network flexibility is the highest concern (i.e., bi-directional communication and arbitrary transceiver distribution are top priorities) then each transceiver should have its own dedicated RF and CDR (clock and data recovery) circuits. This of course will impose heavy circuit overhead as well as large power consumption to the RF/wireless interconnect implementation. To reduce the power consumption, multiple transmitters within a synchronous access range may need to share a common RF transmitter, while multiple receivers may end up sharing a common RF receiver and even a shared CDR without sacrificing significant channel reconfigurability.

(iv) *On-chip Antennas:* While there has been a lot of past research on fabricating antennas on lossless or lower loss substrates such as polytetrafluoroethylene (PTFE), quartz, duroid, and GaAs in the millimeter wave range applications, there has not been sufficient research in the area of fabricating printed antennas on silicon substrate which is much more lossy than other types of substrates. The conductivity of silicon substrate will reduce the antenna efficiency, which necessitates choosing either low loss type of substrates (that may not be compatible with mass-scale fabrication techniques) or high transmission power to improve

radiation efficiency⁵⁹⁾.

(v) *Reference Crystal Oscillator:* Another issue for the adoption of RF/wireless interconnects is the reference crystal oscillator required for FDMA, which cannot be easily implemented on-chip and has a relatively large size compared with future ULSI. Consequently, this crystal oscillator will have to be implemented off-chip. It may however be possible to reuse a reference clock to minimize off-chip overhead.

(vi) *Security:* As with any wireless transmission system, RF/wireless interconnects are susceptible to hackers that intend to snoop on transmitted data and compromise the system. A lot of research is needed to identify security issues with RF/wireless interconnects and, if needed, to develop new (or adapt previously developed) techniques that prevent wireless signal decryption and protect transmitted data from malicious entities.

4. Carbon Nanotube (CNT) Interconnects

Carbon nanotubes (CNTs) have been recently proposed as a replacement for metal interconnects in future technologies^{12),60),61),84)}. CNTs are sheets of graphite rolled into cylinders of diameters varying from 0.6 nm to about 3 nm. Depending on the direction in which they are rolled (called *chirality*), they can demonstrate either metallic (i.e., conducting) or semiconducting properties. CNTs possess high mechanical and thermal stability, high thermal conductivity, and large current carrying capacity, making them promising candidates as on-chip interconnects in future technologies⁶²⁾. Due to their covalently bonded structure, they are highly resistant to electromigration and other sources of physical breakdown⁶³⁾. They can support very high current densities with very little performance degradation. For instance, it was shown in Ref. 64) that the current carrying capacity of CNTs did not degrade even after 350 hours at current densities of about 10^{10} A/cm² at 250°C. CNTs possess high thermal conductivity in the range of 1700–3000 W/m.K⁶⁵⁾. They also have much better conductivity properties than copper owing to longer electron mean free path lengths (MFP) in the micrometer range, compared to nanometer range MFP lengths for Cu⁶⁶⁾.

Figure 7 shows the different CNT alternatives that are being investigated as replacements for copper interconnects. It is predicted that isolated single walled




Fig. 7 Carbon nanotube (CNT) types: (a) single-walled CNT (SWCNT), (b) multi-walled CNT (MWCNT), and (c) SWCNT bundle.

CNTs (SWCNTs) can replace copper at the local interconnect level because of their much lower lateral capacitance which improves latency for very short distances⁶⁷⁾. However, the high intrinsic resistance associated with an SWCNT (greater than $6.45 \text{ K}\Omega$) necessitates the use of a bundle of SWCNTs conducting current in parallel to form longer on-chip interconnections^{60),61)}. SWCNT bundle structures have recently been demonstrated and their metallic conducting properties reported in Ref. 68). Due to the lack of control over chirality in current fabrication techniques, any bundle of SWCNTs consists of metallic nanotubes that contribute to current conduction, as well as semi-conducting nanotubes that do not contribute to current conduction in an interconnect. Multi-walled CNTs (MWCNTs) are comprised of multiple concentric SWCNT shells. Even though MWCNTs are primarily metallic, preliminary research has shown that it is relatively more challenging to achieve transmission over long lengths with them⁶⁹⁾, compared to SWCNTs.

4.1 Comparison between on-chip EIs and CNTs

Srivastava, et al.¹²⁾ presented a performance comparison between a copper (Cu) electrical interconnect and a SWCNT bundle having the same dimensions as the Cu interconnect. **Figure 8** shows how the ratio of the propagation delay for




Fig. 8 Ratio of local interconnect propagation delay with sparse CNT bundle interconnect having (a) perfect contacts and (b) imperfect ($120 \text{ K}\Omega$) contacts, to that with Cu interconnect as a function of interconnect length¹²⁾.

CNT-bundle interconnects $t_p(\text{CNT})$ with the propagation delay for Cu interconnects $t_p(\text{Cu})$ varies with wire line length, at the local interconnect level. A ratio value of greater than 1 indicates that the propagation delay for Cu interconnect is lower than that for CNT-bundle interconnect, while a value of less than 1 indicates that propagation delay for the CNT-bundle interconnect is lower than that for the Cu interconnect. Figure 8(a) shows that the performance of CNT bundles with perfect contacts is better than Cu wires, assuming the distance between adjacent metallic CNTs forming a bundle is large (i.e., assuming sparse bundles). However, with realistic contacts between metal and CNTs (Fig. 8(b)), the CNT performance is lower than Cu interconnect performance. It is worth noting that the CNT density can be reduced only up to a small extent beyond

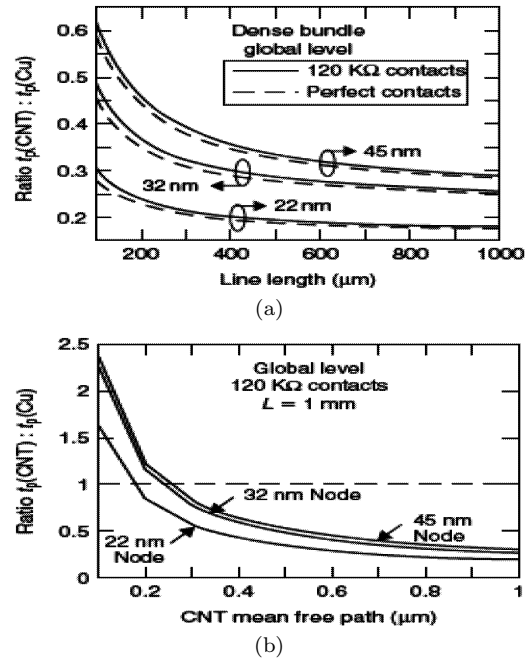


Fig. 9 Ratio of global interconnect propagation delay with dense CNT bundle to that with Cu as a function of (a) as a function of long interconnect length ($L_0 = 1 \mu\text{m}$), (b) as a function of L_0 for 1 mm long interconnect¹²⁾.

which the improvement in performance is lost due to increasing resistance of the bundle. For the case of densely packed CNT bundles, it was shown¹²⁾ that the propagation delay of local interconnects with CNT bundles is higher than that with Cu interconnects across all technology generations, even if contacts are perfect and a mean free path (L_0) as large as $10 \mu\text{m}$ can be achieved. This is because the higher capacitance of CNT bundles and the high resistance of minimum sized drivers at the local interconnect level overshadow the advantage from low CNT-bundle resistance.

Figure 9 (a) shows how propagation delay for densely packed CNT-bundle interconnects depends on the length at the global level. It can be seen that global interconnects implemented with CNT-bundles can achieve significantly better

performance than copper. In the case of global interconnects, the improvement in performance with CNT-bundles is larger for longer interconnects and saturates beyond a certain length. These observations hold if the mean free path has an ideal value of $1 \mu\text{m}$. Unfortunately, the presence of defects in a nanotube leads to the mean free path being much less than the typical $1 \mu\text{m}$ ⁷¹⁾. The impact of reduced mean free path lengths on global interconnect propagation delay with CNT bundles is shown in Fig. 9 (b). It can be seen that it is critical to maintain mean free path lengths in the range of a micrometer by ensuring freedom from such defects. The improved performance of global level CNT-bundle interconnects is because of the much lower resistance compared to Cu global interconnects. While Cu interconnect resistance increases linearly with length, in the case of CNT bundle interconnects it is only the scattering resistance (proportional to $h/4e^2$) that increases linearly. The additional resistance arising from imperfect metal-nanotube contacts (which dominates the resistance in the case of short local interconnects) does not increase with length. Hence Fig. 9 shows only a minor difference between the performance of CNT-bundles with perfect contacts and with imperfect contacts. The performance improvement in global interconnects as a result of using CNT-bundles, decreases when the CNT packing density is decreased. This is because when CNT packing density is decreased, the effective resistance increases and hence performance degrades. This is in contrast to short local interconnects for which the effect of interconnect capacitance dominates and a slightly lower CNT density (and hence lower interconnect capacitance) actually leads to improved performance.

Multi-wall carbon nanotubes (MWCNTs) can be considered to be a coaxial assembly of SWCNT cylinders (shells), one within another. While SWCNTs have diameters in the few nanometer range, MWCNTs may have diameters in a wide range varying from a few to hundreds of nanometers. It has been recently shown that if properly connected to contacts, all the shells in a MWCNT can conduct⁷²⁾. The conductance of a (SWCNT) graphene shell in a MWCN increases as the diameter of a shell becomes larger. The ratio of the diameters of the inner and outer shells varies in different MWCNTs, ranging from 0.35 to 0.8⁷²⁾. Preliminary results shown in **Fig. 10**⁷³⁾ indicate that for long lengths (a few hundreds of micrometers) MWCNTs can have conductivities several times larger than that

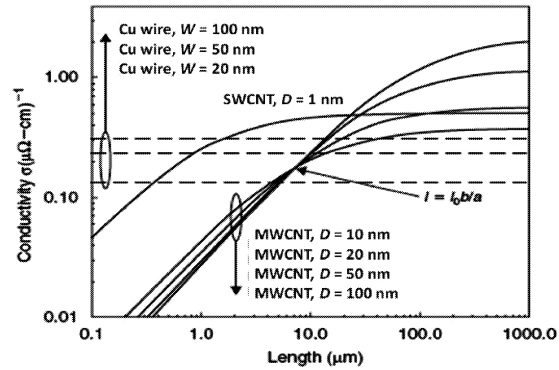


Fig. 10 Conductivity of MWCNTs with various diameters, SWCNT-bundle and Cu interconnect versus length⁷³⁾.

of copper or SWCN bundles.

4.2 On-chip CNT Interconnect Research

Carbon nanotube (CNT) interconnects offer a promising alternative to traditional Cu based interconnects that are reaching their limits with scaling technology, especially at the global interconnect level^{73)–76),84)}. While an isolated SWCNT is not a viable option to replace a Cu wire for global on-chip communication^{12),74),75),84)}, SWCNT bundles and MWCNTs can be a viable replacement for Cu at the intermediate and global interconnect levels. Isolated SWCNTs can however replace Cu at the local interconnect level because of their much lower lateral capacitance which improves latency for short distances⁶⁷⁾.

Researchers have developed RLC circuit models for various CNT interconnect alternatives and compared their performance with Cu interconnects. RLC circuit models have been developed for individual SWCNTs^{67),74),75)}, SWCNT bundles^{12),77)} and MWCNTs^{77),78)}. Other studies have presented some interesting discussions on the impact of process variations on CNT performance⁷⁹⁾ and the possibility of CNTs replacing Cu interconnects in future FPGA fabrics⁸⁰⁾. All of these studies have shown that CNT interconnects are a viable alternative to Cu interconnects in future technologies.

4.3 Challenges

Much work however still needs to be done to develop and improve fabrication

techniques, and resolve the open problems in the area of CNT-based interconnects, some of which are enumerated below:

(i) *Inefficient metal-nanotube contacts*: It has been observed⁷¹⁾ that imperfect metal-nanotube contacts that are fabricated today give rise to an additional contact resistance which makes propagation delay on CNT interconnects higher than with Cu interconnects. Making a reliable contact to CNT is very challenging⁷¹⁾, but is a critical issue that needs to be addressed. A few studies^{81),82)} have shown that state-of-the-art fabrication techniques can make it possible to reduce contact resistance down to a very small value (a few hundred Ω).

(ii) *Small mean free path length*: It is important to maintain mean free path (MFP) length for CNTs in the range of a micrometer. Due to limitations in nanotube fabrication technology today and the presence of defects in nanotubes, the mean free path is much less than a micrometer. This leads to propagation delays in CNT bundles that are larger than for Cu interconnects in future technologies. Better fabrication techniques are needed to reduce these defects and keep the mean free path length in the micrometer range.

(iii) *Density of nanotube bundles*: Densely packed nanotubes are needed for global CNT interconnects that perform better than Cu interconnects. However, due to the lack of control on chirality during nanotube fabrication, any bundle of CNTs consists of metallic as well as semi-conducting nanotubes. The semi-conducting nanotubes do not contribute to current conduction in CNT bundles, which results in a lower effective density for the CNT bundle (creating a sparse bundle). The sparse CNT bundle that is created has a higher resistance, and consequently poor propagation delay, compared to Cu interconnects. CNT-bundles currently fabricated do not have a very high density of CNTs⁷⁰⁾, but improvements are needed if CNT interconnects are to become viable replacements for Cu interconnects.

(iv) *Inductive effects at high frequencies*: Currently, inductive effects have been ignored while calculating propagation delay in CNTs. Inductive effects are expected to become significant at very high frequencies (greater than 10 GHz). At such frequencies, these inductive effects can negatively influence performance on CNTs, which needs to be addressed when comparing CNTs to Cu interconnects in future technologies.

5. Concluding Remarks

Traditional on-chip copper-based electrical interconnects are beginning to show their limitations due to a number of reasons: high resistivity (due to electron grain and grain boundary scatterings, leading to large propagation delays and poor performance), low reliability and high susceptibility to electromigration (at high current densities) with technology and interconnect scaling. Consequently, novel on-chip interconnection schemes need to be explored for future high frequency (giga- and tera-scale) electronic systems. In this paper, we presented three of the most promising on-chip interconnection technologies that have the potential to mitigate the difficulties faced by current on-chip metallic interconnects. Optical interconnects convert data into light and transmit them over an optical waveguide on a chip. Wireless interconnects convert data in RF/wireless signals that can propagate between the transmitting and receiving antennas on the chip, without the need for any metal wires. Carbon nanotubes utilize rolled sheets of graphene to transmit current, with less crosstalk and lower propagation delay than metallic interconnects. All three of these emerging technologies have several issues and open problems (such as the need for improvements in fabrication technology) that must be resolved before they can be adopted as part of on-chip interconnect fabrics. With the rapid advances in technology, it is however only a matter of time before one or more of these technologies becomes feasible and advantageous to use in tomorrow's VLSI designs.

Acknowledgments This research was partially supported by grants from NSF (CCF-0702797).

References

- 1) Semiconductor Industry Association: International Technology Roadmap for Semiconductors (2005).
- 2) Naeemi, A., Sarvari, R. and Meindl, J.D.: On-Chip Interconnect Networks at the End of the Roadmap: Limits and Nanotechnology Opportunities, *International Interconnect Technology Conference (IITC)*, pp.201–203 (2006).
- 3) Ho, R., Mai, K.W. and Horowitz, M.A.: The future of wires, *Proc. IEEE*, Vol.89, No.4, pp.490–504 (Apr. 2001).
- 4) ARM AMBA Specification and Multi layer AHB Specification (rev2.0) (2001). <http://www.arm.com>
- 5) ARM AMBA 3.0 AXI Specification. www.arm.com/armtech/AXI
- 6) IBM CoreConnect Specification. http://www.ibm.com/chips/techlib/techlib.nsf/productfamilies/CoreConnect_Bus_Architecture
- 7) Sonics SMART Interconnect. <http://www.sonicsinc.com>
- 8) STBus Communication System: Concepts and Definitions, Reference Guide, STMicroelectronics (2003).
- 9) Pasricha, S. and Dutt, N.: *On Chip Communication Architectures*, Systems on Silicon Series, Morgan Kaufman (2008).
- 10) De Micheli, G. and Benini, L.: *Networks on Chips*, Systems on Silicon Series, Morgan Kaufman (2006).
- 11) Chen, G., et al.: On-Chip Copper-Based vs. Optical Interconnects: Delay Uncertainty, Latency, Power, and Bandwidth Density Comparative Predictions, *International Interconnect Technology Conference (IITC)*, pp.39–41 (2006).
- 12) Srivastava, N. and Banerjee, K.: Performance Analysis of Carbon Nanotube Interconnects for VLSI Applications, *Proc. IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pp.383–390 (2005).
- 13) Savage, N.: Linking with light, *IEEE Spectrum*, Vol.39, No.8, pp.32–36 (Aug. 2002).
- 14) Plant, D.V.: System design of chip and board level optical interconnects, *Proc. 2004 Meeting Bipolar/BiCMOS Circuits and Technology*, pp.72–78 (Sep. 2004).
- 15) Haurylau, M., et al.: On-Chip Optical Interconnect Roadmap: Challenges and Critical Directions, *IEEE J. Selected Topics in Quantum Electronics*, Vol.12, No.6, pp.1699–1705 (Nov.-Dec. 2006).
- 16) Rong, H., Jones, R., Liu, A., Cohen, O., Hak, D., Fang, A. and Paniccia, M.: A continuous-wave Raman silicon laser, *Nature*, Vol.433, pp.725–728 (Feb. 2005).
- 17) Pavesi, L.: Routes toward silicon-based laser, *Mater. Today*, Vol.8, No.1, pp.18–25 (Jan. 2005).
- 18) Liu, A., et al.: A high-speed silicon optical modulator based on a metal-oxide-semiconductor capacitor, *Nature*, Vol.427, No.6975, pp.615–618 (Feb. 2004).
- 19) Xu, Q.F., et al.: Micrometre-scale silicon electro-optic modulator, *Nature*, Vol.435, No.7040, pp.325–327 (May 2005).
- 20) Almeida, V.R., et al.: All-optical switching on a silicon chip, *Optics Letters*, Vol.29, No.24, p.2867 (Dec. 2004).
- 21) Tatum, J.: VCSELs for 10 GB/s optical interconnects, *IEEE Emerging Technologies Symposium on BroadBand Communications for the Internet Era*, Richardson, TX, pp.58–61 (Sep. 2001).
- 22) Yin, T., Pappu, A.M. and Apsel, A.B.: Low-cost, high efficiency, and high-speed SiGe phototransistors in commercial BiCMOS, *IEEE Photonics Technology Letters*, Vol.18, No.1, pp.55–57 (Jan. 2006).
- 23) Reshotko, M.R., Kencke, D.L. and Block, B.: High-speed CMOS compatible pho-

- photodetectors for optical interconnects, *Proc. SPIE*, Vol.5564, pp.146–155 (Oct. 2004).
- 24) Pappu, A.M. and Apsel, A.B.: A low power, low delay TIA for on-chip applications, *Conference on Lasers and Electro-Optics*, pp.594–596 (May 2005).
 - 25) Kobrinsky, M., et al.: On-chip optical interconnects, *Intel Technology Journal*, Vol.8, No.2 (May 2004).
 - 26) Collet, J.H., Caignet, F., Sellaye, F. and Litaize, D.: Performance constraints for onchip optical interconnects, *IEEE J. Selected Topics in Quantum Electronics*, Vol.9, No.2, pp.425–432 (Mar.-Apr. 2003).
 - 27) Kobrinsky, M.J., Block, B.A., Zheng, J.-F., Barnett, B.C., Mohammed, E., Reshotko, M., Robertson, F., List, S., Young, I. and Cadien, K.: On-chip optical interconnects, *Intel Technol. J.*, Vol.8, No.2, pp.129–141 (May 2004).
 - 28) Chen, K.-N., Kobrinsky, M.J., Barnett, B.C. and Reif, R.: Comparisons of conventional, 3-D, optical, and RF interconnects for on-chip clock distribution, *IEEE Trans. Electron Devices*, Vol.51, No.2, pp.233–239 (Feb. 2004).
 - 29) Ackland, B., Razavi, B. and West, L.: A comparison of electrical and optical clock networks in nanometer technologies, *Proc. IEEE 2005 Custom Integrated Circuits Conference, 2005*, pp.779–782 (Sep. 2005).
 - 30) O'Connor, I.: Optical solutions for system-level interconnect, *Proc. 2004 International Workshop on System Level interconnect Prediction (SLIP)*, pp.79–88 (2004).
 - 31) Heide, T., Ghazi, A., Zimmermann, H. and Seegebrecht, P.: Monolithic CMOS photoreceivers for short-range optical data communications, *El. Lett.*, Vol.35, pp.1655–1656 (1999).
 - 32) Zimmermann, H., Heide, T. and Ghazi, A.: Monolithic high-speed CMOS photoreceiver, *IEEE Photon. Technol. Lett.*, Vol.11, pp.254–256 (Feb. 1999).
 - 33) Briere, M., Drouard, E., Mieyeville, F., Navarro, D., O'Connor, I. and Gaffiot, F.: Heterogeneous modelling of an optical network-on-chip with SystemC, *16th IEEE International Workshop on Rapid System Prototyping (RSP)*, pp.10–16 (2005).
 - 34) Shacham, A., Bergman, K. and Carloni, L.P.: On the Design of a Photonic Network-on-Chip, *First International Symposium on Networks-on-Chip*, pp.53–64 (2007).
 - 35) Pasricha, S. and Dutt, N.: ORB: An On-chip Optical Ring Bus Communication Architecture for Multi-Processor Systems-on-Chip, *IEEE Asia & South Pacific Design Automation Conference (ASPDAC 2008)*, Seoul, Korea (Jan. 2008).
 - 36) Green, W.M., Rooks, M.J., Sekaric, L. and Vlasov, Y.A.: Ultra-compact, low RF power, 10 Gb/s silicon Mach-Zehnder modulator, *Opt. Express*, Vol.15, No.25, pp.17106–17113 (Dec. 2007).
 - 37) www.intel.com/research/platform/sp/hybridlaser.htm
 - 38) Weiss, S.M., Molinari, M. and Fauchet, P.M.: Temperature stability for silicon-based photonic band-gap structures, *Appl. Phys. Lett.*, Vol.83, No.10, pp.1980–1982 (Sep. 2003).
 - 39) Chang, M.F., Roychowdhury, V.P., Zhang, L., Shin, H. and Qian, Y.: RF/wireless interconnect for inter- and intra-chip communications, *Proc. IEEE*, Vol.89, No.4, pp.456–466 (Apr. 2001).
 - 40) Gupta, K.C., et al.: *Computer-Aided Design of Microwave Circuits*, p.69, Artech House, Norwell, MA (1981).
 - 41) Frankel, M.Y., et al.: Terahertz attenuation and dispersion characteristics of coplanar transmission lines, *IEEE Trans. Microwave Theory Tech.*, Vol.39 (June 1991).
 - 42) Feher, K.: *Digital Communications*, Prentice-Hall, Englewood Cliffs, NJ (1981).
 - 43) Pehlke, D.R., Burstein, A. and Chang, M.F.: Extremely high-Q tunable inductor for Si-based RF integrated circuit applications, *Dig. 1997 Int. Electron Devices Meeting*, Washington, DC, pp.63–66 (1997).
 - 44) Chang, R.T., Talwalkar, N., Yue, C.P. and Wong, S.S.: Near speed-of-light signaling over on-chip electrical interconnects, *IEEE J. Solid-State Circuits*, Vol.38, No.5, pp.834–838 (May 2003).
 - 45) Chang, M.F., et al.: Advanced RF/Baseband Interconnect Schemes for Inter- and Intra-ULSI communications, *IEEE Trans. Electron Devices*, Vol.52, No.7, pp.1271–1285 (July 2005).
 - 46) Viterbi, A.J.: *CDMA Principles of Spread Spectrum Communication*, Communication Series, Addison-Wesley, Reading, MA (June 1995).
 - 47) Wang, Y., Makadia, D. and Margala, M.: On-Chip Integrated Antennas — The First Challenge for Reliable on-Chip Wireless Interconnects, *Canadian Conference on Electrical and Computer Engineering*, pp.2322–2325 (May 2006).
 - 48) Nakano, H. and Yamauchi, J.: A Consideration of Curved Halfwave Dipoles, *Trans. IECE Japan*, Vol.J61-B, pp.210–213 (1978).
 - 49) Pozar, D.M.: *Microwave and RF Design of Wireless Systems*, p.137, John Wiley & Sons, Inc., New York (2000).
 - 50) Branch, J., et al.: Wireless communication in a flip-chip package using integrated antennas on silicon substrates, *Electron Device Letters, IEEE*, Vol.26, No.2, pp.115–117 (Feb. 2005).
 - 51) K.K. O., et al.: The Feasibility of On-Chip Interconnection Using Antennas, *IEEE International Conference on Computer Aided Design*, San Jose, CA, pp.976–981 (Nov. 2005).
 - 52) Benech, P., Ndagijimana, F., Triantafyllou, A., Farcy, A. and Torres, J.: Design and performance of integrated antennas for wireless intra chip interconnections, *IEEE 32nd Annual Conference on Industrial Electronics (IECON)*, pp.2953–2957 (Nov. 2006).
 - 53) K.K. O., Kim, K., Floyd, B.A., Mehta, J. and Yoon, H.: Inter and Intra-Chip Wireless Clock Signal Distribution Using Microwaves: A Status of an Initial Feasibility Study, *Government Microcircuit Applications Conference*, Monterey, CA, pp.306–309 (Mar. 1999).
 - 54) Guo, X., Caserta, J., Li, R., Floyd, B. and K.K. O.: Propagation Layers for Intra-

- Chip Wireless Interconnection Compatible with Packaging and Heat Removal, *2002 Symposium on VLSI Technology*, pp.36–37 (June 2002).
- 55) Floyd, B.A., et al.: Wireless interconnection in a CMOS IC with integrated antennas, *Proc. ISSCC*, pp.328–329 (2000).
 - 56) Seok, E.-Y. and K.K. O.: Design Rules for Improving Predictability of On-chip Antenna Characteristics in the Presence of Other Metal Structures, *2005 International Interconnect Technology Conference*, San Francisco, CA, pp.120–123 (June 2005).
 - 57) Mehta, J. and K.K. O.: Switching Noise of Integrated Circuits (IC's) Picked up by a Planar Dipole Antenna Mounted Near the IC's, *IEEE Trans. Electro-Magnetic Compatibility*, Vol.44, No.5, pp.282–290 (May 2002).
 - 58) Dickson, T.O., Bravo, D. and K.K. O.: Noise Coupling to On-Chip Antennas, *2002 IEEE International Symposium on EMC*, Vol.1, pp.340–344 (2000).
 - 59) K.K. O., et al.: On-Chip Antennas in Silicon ICs and Their Application, *IEEE Trans. Electron Devices*, Vol.52, No.7, pp.1312–1323 (July 2005).
 - 60) Kreupl, F., et al.: Carbon Nanotubes in Interconnect Applications, *Microelectronic Engineering*, pp.399–408 (2002).
 - 61) Li, J., et al.: Bottom-up Approach for Carbon Nanotube Interconnects, *Applied Physics Letters*, Vol.82, No.15, pp.2491–2493 (Apr. 2003).
 - 62) Wei, B.Q., et al.: Reliability and Current Carrying Capacity of Carbon Nanotubes, *Applied Physics Letters*, Vol.79, No.8, pp.1172–1174 (2001).
 - 63) Dresselhaus, M.S., Dresselhaus, G. and Avouris, P.: *Carbon nanotubes: synthesis, structure, properties and applications*, Springer (2001).
 - 64) Wei, B.Q., Vajtai, R. and Ajayan, P.M.: Reliability and current carrying capacity of carbon nanotubes, *Appl. Phys. Lett.* (2001).
 - 65) Hone, J., Whitney, M., Piskoti, C. and Zettl, A.: Thermal conductivity of single-walled carbon nanotubes, *Phys. Rev., B*, Vol.59, No.4 (1999).
 - 66) Kreupl, F., Graham, A.P., Liebau, M., Duesberg, G.S., Seidel, R. and Unger, E.: Carbon nanotubes for interconnect applications, *IEDM Tech. Dig.* (Dec. 2004).
 - 67) Naeemi, A. and Meindl, J.D.: Design and Performance Modeling for Single-Walled Carbon Nanotubes as Local, Semiglobal and Global Interconnects in Gigascale Integrated Systems, *Trans EDL* (Jan. 2007).
 - 68) Thess, A., et al.: Crystalline Ropes of Metallic Carbon Nanotubes, *Science*, Vol.273, No.5274, pp.483–487 (1996).
 - 69) Schonenberger, C., et al.: Interference and Interaction in Multi-wall Carbon Nanotubes, *Applied Physics A*, Vol.69, pp.283–295 (1999).
 - 70) Srivastava, N. and Banerjee, K.: A Comparative Scaling Analysis of Metallic and Carbon Nanotube Interconnections for Nanometer Scale VLSI Technologies, *Proc. 21st Intl. VLSI Multilevel Interconnect Conf.*, pp.393–398 (2004).
 - 71) Hunger, T., Lengeler, B. and Appenzeller, J.: Transport in Ropes of Carbon Nanotubes: Contact Barriers and Luttinger Liquid Theory, *Physical Review B*, Vol.69 (2004).
 - 72) Nihei, M., Kondo, D., Kawabata, A., Sato, S., Shioya, H., Sakaue, M., Iwai, T., Ohfuti, M. and Awano, Y.: Low-resistance multi-walled carbon nanotube vias with parallel channel conduction of inner shells, *Proc. IEEE Int. Interconnect Technol. Conf.*, pp.234–236 (Jun. 2005).
 - 73) Naeemi, A. and Meindl, J.D.: Compact Physical Models for Multiwall Carbon-Nanotube Interconnects, *Electron Device Letters, IEEE*, Vol.27, pp.338–340 (2006).
 - 74) Raychowdhury, A. and Roy, K.: Modeling of Metallic Carbon Nanotube Interconnects for Circuit Simulation and a Comparison with Cu Interconnects for Scaled Technologies, *IEEE Trans. Computer-Aided Design of IC's*, pp.58–65 (Jan. 2006).
 - 75) Raychowdhury, A. and Roy, K.: A circuit model for carbon nanotube interconnects: Comparative study with Cu interconnects for scaled technologies, *Proc. 2004 IEEE/ACM international Conference on Computer-Aided Design (ICCAD)*, pp.237–240 (2004).
 - 76) Naeemi, A. and Meindl, J.D.: Monolayer metallic nanotube interconnects: Promising candidates for short local interconnects, *Electron Device Letters, IEEE*, Vol.26, pp.544–546 (2005).
 - 77) Naeemi, A., et al.: Performance Modeling and Optimization for Single- and Multi-Wall Carbon Nanotube Interconnects, *DAC* (2007).
 - 78) Haruehanroengra, S. and Wang, W.: Analyzing Conductance of Mixed Carbon-Nanotube Bundles for Interconnect Applications, *IEEE EDL* (Aug. 2007).
 - 79) Nieuwoudt, A. and Massoud, Y.: Assessing the Implications of Process Variations on Future Carbon Nanotube Bundle Interconnect Solutions, *ISQED* (2007).
 - 80) Eachempati, S., et al.: Assessing Carbon Nanotube Bundle Interconnect for Future FPGA Architectures, *DATE* (2007).
 - 81) Sato, S., Nihei, M., Mimura, A., Kawabata, A., Kondo, D., Shioya, H., Iwai, T., Mishima, M., Ohfuti, M. and Awano, Y.: Novel approach to fabricate carbon nanotube via interconnects using size-controlled catalyst nanoparticles, *Proc. International Interconnect Technology Conference (IITC)*, pp.230–232 (2006).
 - 82) Hjortstam, O., Isberg, P., Söderholm, S. and Dai, H.: Can we achieve ultra-low resistivity in carbon nanotube-based metal composites?, *App. Phy Materials Science & Processing*, pp.1175–1179 (2004).
 - 83) Katehi, L.P.B. and Herrick, K.: Si-micromachining in MM-wave circuits, *Proc. 1997 Topical Symp. Millimeter Waves*, New York, pp.85–88 (1997).
 - 84) Pasricha, S., Kurdahi, F. and Dutt, N.: System Level Performance Analysis of Carbon Nanotube Global Interconnects for Emerging Chip Multiprocessors, *IEEE/ACM NanoArch* (2008).

(Received March 17, 2008)

(Released August 27, 2008)

(Invited by Editor-in-Chief: *Hidetoshi Onodera*)



Sudeep Pasricha received the B.E. degree in electronics and communication engineering from Delhi Institute of Technology, Delhi, India, in 2000, and his M.S. and Ph.D. degrees in computer science from the University of California, Irvine in 2005 and 2008. He became an assistant professor at Colorado State University in 2008, with joint appointments in the ECE and CS departments. His current research interests are in the areas of on-chip communication architectures, emerging interconnect technologies, embedded systems, electronic design automation, system-level modeling languages and design methodologies, and computer architecture. He has presented several tutorials in the area of on-chip communication architecture design at leading conferences and recently authored a book titled ‘On-chip Communication Architectures’ (Morgan Kaufman/Elsevier 2008). He has received a Best Paper Award at ASPDAC 2006, a Best Paper Award nomination at DAC 2005, and several fellowships and awards for excellence in research. He is currently in the program committee of the ISQED and VLSID conferences.



Nikil D. Dutt received the Ph.D. degree in computer science from the University of Illinois at Urbana-Champaign, Urbana, in 1989, and is currently a Chancellor’s Professor at the University of California, Irvine, with academic appointments in the CS and EECS departments. His research interests are in embedded systems, electronic design automation, computer architecture, optimizing compilers, system specification techniques, and distributed embedded systems. He has coauthored 6 books and has received 6 Best Paper Awards at premier CAD and Embedded Systems conferences. He currently serves as Editor-in-Chief of ACM Transactions on Design Automation of Electronic Systems (TODAES) and as Associate Editor of ACM Transactions on Embedded Computer Systems (TECS), and of IEEE Transactions on VLSI Systems (IEEE T-VLSI). He was an ACM SIGDA Distinguished Lecturer during 2001–2002, and an IEEE Computer Society Distinguished Visitor for 2003–2005. He has served on the steering, organizing, and program committees of several premier CAD and Embedded System Design conferences and workshops, including ASPDAC, CASES, CODES+ISSS, DATE, ICCAD, ISLPED and LCTES. He serves or has served on the advisory boards of ACM SIGBED and ACM SIGDA, and previously served as Vice-Chair of ACM SIGDA and of IFIP WG 10.5. Professor Dutt is a Fellow of the IEEE, an ACM Distinguished Scientist, and an IFIP Silver Core Awardee.