Optical Clock Distribution in Electronic Systems

S. K. Tewksbury and L.A. Hornak
Microelectronic Systems Research Center
Dept. of Electrical and Computer Engineering
West Virginia University
Morgantown, WV 26506

Abstract

Techniques for distribution of optical signals, both free space and guided, within electronic systems has been extensively investigated over more than a decade. Particularly at the lower levels of packaging (intra-chip and chip-to-chip), miniaturized optical elements including diffractive optics and micro-refractive optics have received considerable attention. In the case of optical distribution of data, there is the need for a source of optical power and a need for a means of modulating the optical beam to achieve data communications. As the number of optical data interconnections increases, the technical challenges of providing an efficient realization of the optical data interconnections also increases. Among the system signals which might be transmitted optically, clock distribution represents a substantially simplified problem from the perspective of the optical sources required. In particular, a single optical source, modulated to provide the clock signal, replaces the multitude of optical sources/modulators which would be needed for extensive optical data interconnections. Using this single optical clock source, the technical problem reduces largely to splitting of the optical clock beam into a multiplicity of optical clock beams and distribution of the individual clocks to the several portions of the system requiring synchronized clocks. The distribution problem allows exploitation of a wide variety of passive, miniaturized optical elements (with diffractive optics playing a substantial role). This article reviews many of the approaches which have been explored for optical clock distribution, ranging from optical clock distribution within lower levels of the system packaging hierarchy through optical clock distribution among separate boards of a complex system. Although optical clock distribution has not yet seen significant practical application, it is evident that the technical foundation for such clock distribution is well established. As clock rates increase to 1 GHz and higher, the practical advantages of optical clock distribution will also increase, limited primarily by the cost of the optical components used and the manufacturability of an overall electronic system in which optical clock distribution has been selectively inserted.

1 Introduction

This paper provides a broad overview of the several approaches which have been investigated for use of optical clock distribution within electronic systems. Such a review can proceed from two starting points. One starts at the higher level (optical connections among racks of electronics or among locally distributed computers) and proceed to optical connections among printed circuit boards (PCBs) through interconnections on a PCB to low level interconnections at the multichip module (MCM) and integrated circuit (IC) levels. Such a review would reflect an evolutionary extension of fiber-optics techniques such as the optical ribbon cable approaches currently being developed for commercial use (e.g., early exploratory approaches [1] through representative, recent parallel optical links [2, 3]).

A second approach, used here, starts with the more risky optical interconnect technologies needed at the lowest levels of packaging (e.g., the intra-IC interconnections) and proceeds to successively higher levels of packaging (intra-MCM, inter-MCM, inter-PCB, etc.). This approach highlights several of the exciting innovations which have been investigated by a large number of researchers over more than a decade. The importance of diffractive micro-optics (see, for example, the special issues [4, 5]) in miniaturized forms suitable for low levels of packaging clearly emerges as a primary direction for the optical "wires." Moving from the IC to the intra-MCM level, folded diffractive optics becomes a clear contender, avoiding the need to place the optical lenses, beam splitters, couplers and other elements well above the plane of the MCM. At the inter-PCB interconnection level, longer distances and less precise alignment of the electrical components (relative to the intra-IC devices and ICs mounted on MCMs) lead to significant adaptions of the folded

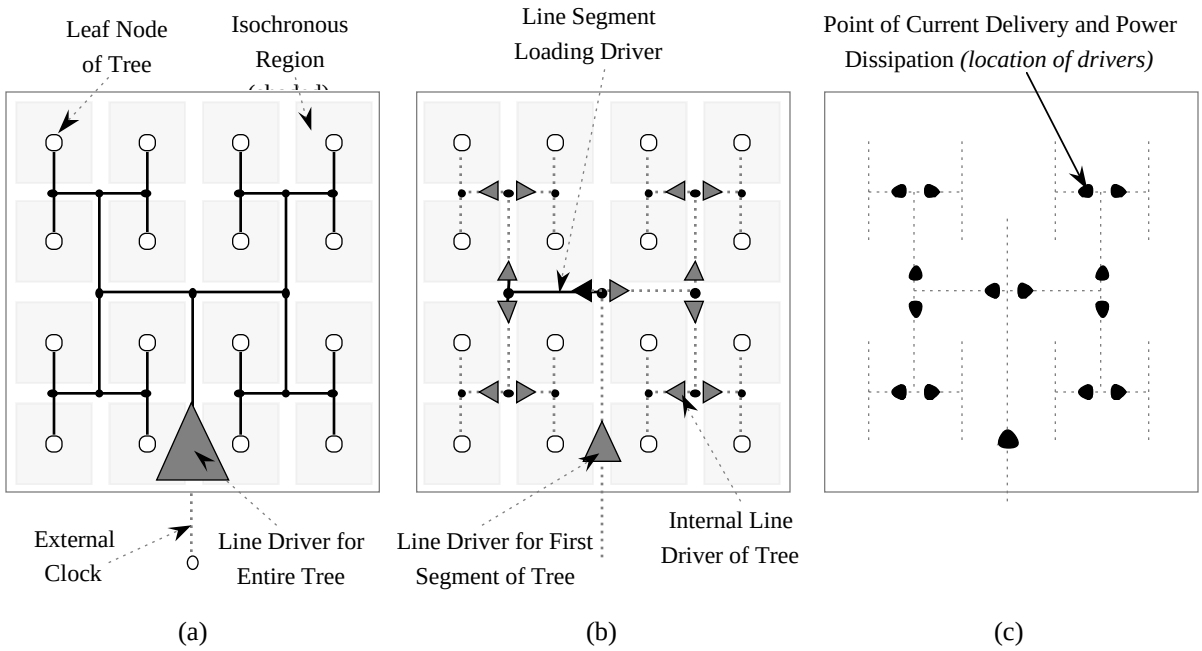


Figure 1: H-Tree distribution of optical clock. (a) H-Tree with single driver driving entire clock net. (b) Distributed drivers, separating net into shorter line segments. (c) Distribution of clock net driver power for case in (b).

diffractive optical techniques and the possibility of using more conventional optical elements (GRIN lenses rather than diffractive lenses) appears. In addition, as one progresses from the lowest level to the higher levels of packaging, the importance of implementing parallel data connections between components (rather than serial data connections) becomes an increasing priority.

Much of the research and exploratory development of optical interconnections at these lower levels of packaging have focussed on data interconnections, requiring an optical source (e.g., laser or LED) and a corresponding optical receiver for each data link provided. For this reason, several of the techniques discussed in this review are drawn from such optical data interconnections. However, the basic components (source, optical elements for the path taken by light, and receiver) are common also to distribution of an optical clock. The main differences between optical data and optical clock distribution are (i) the simplification of eliminating the many optical sources, using instead a single internal or external source, and (ii) the complication of having to convert that single optical clock signal into the multiplicity of clock signals to be distributed to various points in the system.

The technologies and design issues associated with optical interconnections (particularly at the lower levels) are not only rather complex but also depend critically on the specific set of elements selected for the interconnections. For this reason, the review does not attempt to provide an overview of all the detailed issues which must be treated when designing optical interconnections. Instead, the focus is on the diversity of approaches which have been extensively explored, providing a large set of references where more detail regarding technology and design issues can be found by the interested reader. The topic of optical interconnections at the lower level of packaging is a source of several innovative directions in the underlying technologies and in the architectures which might benefit from such interconnections. In addition to the considerable literature on optical interconnections, there has been considerable study of technologies and optical interconnections in the area of optical computing. In fact, several of the advances in optical interconnects were developed by those who were also exploring optical computing. In this review, the optical interconnection is taken as a passive function, and techniques using switching, such as reconfigurable optical interconnection networks, are not discussed here.

Section 2 provides some general background on the limitations of clock distribution within VLSI integrated circuits, followed by a discussion of techniques through which optical clock distribution might be

provided. Several of the techniques discussed for intra-IC interconnections are also applicable to intra-MCM interconnections, which are considered in Section 3. Section 4 considers optical interconnections among MCMs. Section 5 discusses optical backplanes and direct free space optical interconnections between PCBs. A general overview of electrical and optical interconnections in electronic systems is provided in [6], with a collection of papers on such interconnections in [7]. Overviews of the considerable early work related to optical interconnections are provided in [8, 9].

2 Clock Delivery to Multiple Sites on an Integrated Circuit

2.1 Limitations of Electrical Distribution of Clock Edges

Clock signals in silicon VLSI circuits must be distributed for an external clock connection to each flip-flop within the IC. The dense population of flip-flops throughout the IC leads to a complex clock distribution network. The small feature size of interconnections used in submicron silicon VLSI imposes a significant resistance per unit length on the clock lines In addition, the minimum capacitance of the electrical interconnections is bounded by a lower limit of about 0.6 pF/cm due to fringing fields (i.e., as the line width is reduced, the electric fields tend toward the fields of a zero width line above the ground plane). The high resistance of the interconnection combined with the minimum limit on line capacitance imposes a significant *RC* delay factor for the long clock lines, limiting clock rates. The large current required to rapidly charge and discharge the total capacitance of the clock network, which is by far the longest signal interconnection in the IC, limits the maximum clock rate due to power dissipation (power dissipated for clock distribution can be a substantial portion of the overall IC power dissipation). The above effects limit the maximum data rate. In principle, if all clock paths to flip-flops are of equal length, the clock skew would be zero. However, the clock signal received at different flip-flops is influenced by the routing of the clock interconnections among a dense set of data interconnections, switching increasingly rapidly as data rates evolve to Gb/s and higher data rates on ICs. These neighboring, high speed data lines couple a significant amount of crosstalk onto the clock lines, the specific crosstalk appearing at any single flip-flop depending on the specific path of the clock line to that flip-flop and the specific activity of data lines coupled to that path.

The H-Tree approach shown in Figure 1a has become an increasingly useful approach for clock distribution. The equal length of each path from the external connection to each terminal point of the net provides the potential for zero clock skew under idealized conditions (i.e., no data lines on the IC and equal loading of the line segments). Addition of drivers along the H-Tree network, as illustrated in Figure 1b, provides several performance improvements.

- A faster rise time of the clock at the terminal end is obtained by limiting the *RC* delay term to only that of the line segment between the drivers (which regenerate a fast rise time signal).
- The amount of crosstalk introduced by nearby data lines is limited to the length of the line segment between drivers, with the drivers restoring the noisy signal to a clean signal. Decreasing the length of the line segment allows the crosstalk noise to be reduced to negligible levels, with virtually full elimination of the crosstalk added along the segment.
- By introducing drivers, each driver supplies only the current necessary to drive the next leg of the H-Tree network. As a result, the maximum current appearing on any clock line is greatly reduced (and the crosstalk coupled to neighboring data lines is correspondingly reduced).

The H-Tree network is not extended until each flip-flop of the VLSI circuit is included at a separate leaf node of the network. Instead, delivery of a clock signal to a region (isochronous region shown shaded in Figure 1) of the IC is adequate, as illustrated by the modest number of terminal points of the H-Tree network in Figure 1a. Within each isochronous region, clock signals can be distributed without precise equalization of interconnection lengths to flip-flops due to the small propagation delays and rise-time delays within that region. At sufficiently short line lengths, given the clock rate, the interconnection behaves as a static RC interconnection, with an *RC* delay τ_{RC} given by $\tau_{RC} = (R_{dr} + R_l \cdot L)(C_l \cdot L + C_{ld})$, where R_{dr} is the driver resistance, R_l and C_l are the line resistance and capacitance, respectively, per unit length, L is the line length, and C_{ld} is the load capacitance imposed by the gates driven by the line. The propagation delay ($\tau_{pr} = c/\sqrt{\epsilon}$, where c is the speed of light and ϵ is the relative dielectric constant of the insulator

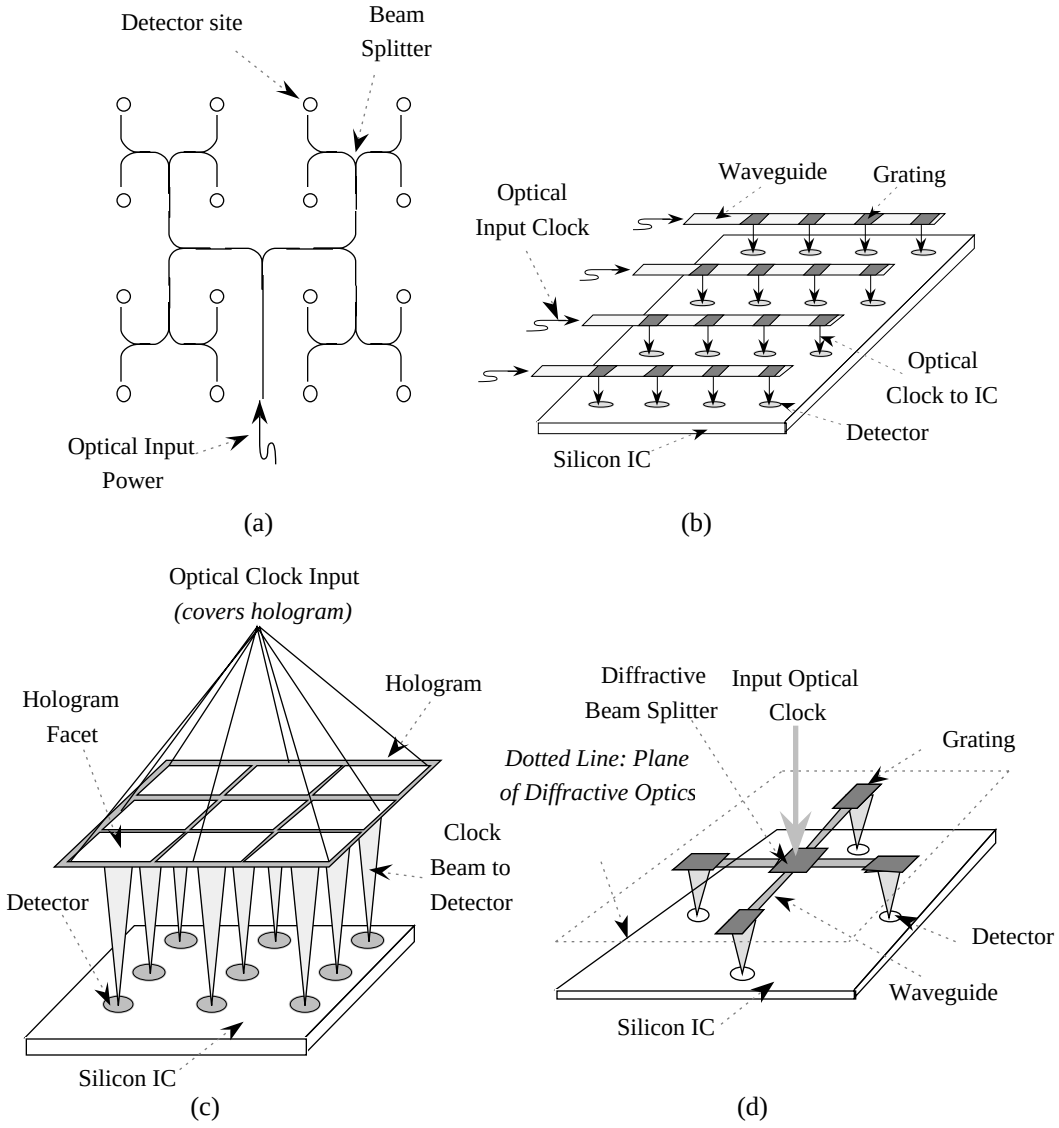


Figure 2: General approaches for optical clock distribution over the area of an IC. (a) Waveguide emulating H-Tree. (b) Tapped waveguides. (c) Hologram distribution of clock. (d) Planar diffractive optics providing vertically optical clocks.

material – here SiO_2) is negligible over the short distances encountered in the isochronous region (e.g., for $\epsilon \approx 4$, $\tau_{pr} \approx 60$ psec/cm). It is therefore the RC delay which dominates the area of the isochronous regions. Similarly, it is the RC delay along the H-Tree network's line segments which dominates the delay of the clock to the leaf nodes of the tree.

As clock rates in VLSI ICs increase, it is necessary to provide not only higher speed clocks but also smaller clock skews, a difficult challenge since the cross sectional area of the interconnections continue to decrease in smaller feature VLSI technologies while the overall length of the clock distribution network remains proportional to the chip dimensions, rather than becoming smaller as in the case of specific digital functions on the IC. Optical clock distribution has, for several years, been considered a particularly important application of optical interconnections to provide the higher clock rates and smaller clock skews needed.

2.2 General Approaches for Optical Clock Distribution

Optical clock distribution to ICs has been studied for several years, with early work described in [10, 11, 12]. General approaches are illustrated in Figure 2 (approaches which are also applicable to inter-chip clock distribution within MCMs, where ICs are mounted unpackaged on a substrate).

- In Figure 2a, optical waveguides replicate the H-Tree topology, with externally applied optical power distributed to the isochronous regions of the IC where, following conversion to electrical clock signals, electrical interconnections continue the distribution of the clock.
- In Figure 2b, a set of parallel waveguides traverse the IC, with diffractive elements redirecting a portion of the light onto underlying detectors. The resulting array of detector sites mimics the arrangement in Figure 2a, but with the condition of precisely equalized paths removed. A related approach uses a single waveguide plate covering the IC and having grating couplers to direct light to selected sites on the IC. Within the waveguide, light travels at a speed $v = c/n_r$, with typical indices of refraction $n_r \approx 1.5$ leading to a clock skew between the two ends of the waveguide of about 50 psec for a 1 cm long waveguide.
- In Figure 2c, the optical clock signals are applied vertically (avoiding the need for fabricating waveguides on the IC surface) through use of a hologram. A related approach (more realistic for inter-IC connections within an MCM) places the optical source on the same plane as the optical detectors. As discussed later, the hologram must be separated from the IC surface by approximately 1 cm, adding significant volume to the basic IC circuit.
- In Figure 2d, optical clock signals are again applied vertically. However, this example uses a planar optical unit (with diffractive optics, mirrors, etc.) through which the light is first distributed in the plane of the unit and then redirected vertically to the IC (with zero skew along the optical paths).

The approaches illustrated in Figure 2 assume a significant number of clock lines into the IC, favoring the use of silicon photodetectors compatible with the VLSI technology. Silicon is transparent to longer wavelengths (e.g., 1.3 μm and 1.5 μm used for lightwave communications), making optical power detection difficult. For this reason, shorter wavelengths (e.g., about 850 nm) are often used. Representative silicon photodetectors include the photoresistor, the photodiode (particularly the PIN diode), the avalanche photodiode, and the phototransistor. The PIN diode has been most extensively used for optical interconnections terminating on silicon ICs. At sufficiently high optical powers, the photodetector's output voltage can drive CMOS circuitry directly. However, receiver amplifiers associated with the detector allow the optical power to be greatly reduced (for a given clock rate).

2.3 H-Tree Optical Clock Distribution using Waveguides

Two distinct classes of waveguide are multi-mode waveguides (characterized by cross section dimensions substantially larger than the wavelength) and single mode waveguides (with cross sections comparable to the wavelength).

Multi-mode waveguides: Several polymers (e.g., polyimide) can be easily deposited or spun onto a substrate and patterned to form waveguides which have sufficiently low loss for optical interconnections (e.g.,

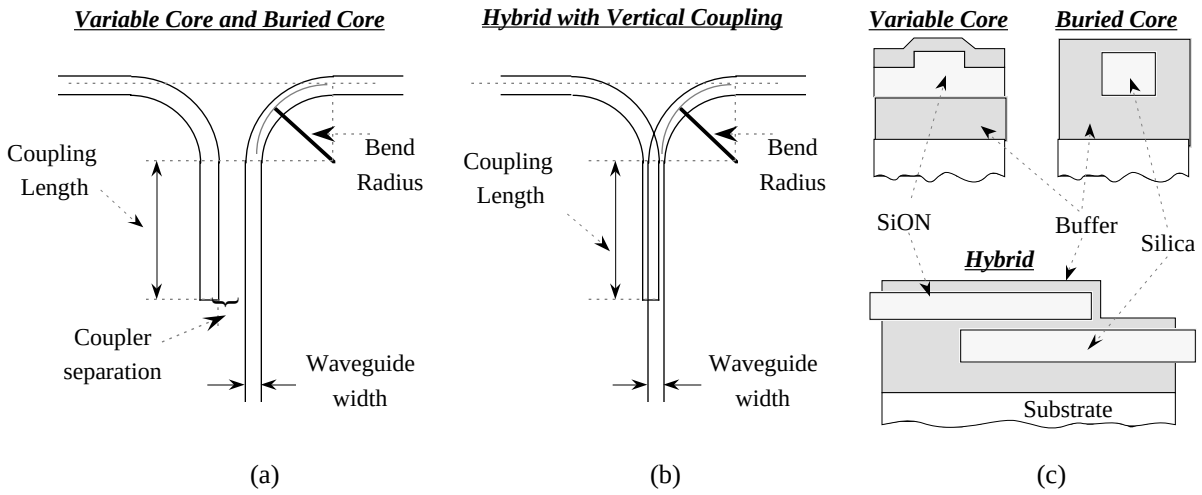


Figure 3: Single-mode waveguide implementing a H-Tree network. (a) Planar waveguide layout with coupler and bend. (b) Vertical coupler hybrid approach. (c) Cross section of SiON waveguide, buried channel silica waveguide, and hybrid approach. After [13].

losses less than about 0.5 dB/cm). Polyimide, for example, is often used as the inter-metal dielectric for MCMs and provides a convenient waveguide material for such components. However, such waveguides generally have cross sectional dimensions substantially larger than the wavelength of semiconductor lasers or LEDs. Such large cross section waveguides support propagation of a multiplicity of modes (wavefront propagation angle relative to direction of waveguide). As a result, most polymer waveguides are multi-mode waveguides. As will be clear later, diffractive optical elements are sensitive to the incident angle of light beams, making such multi-mode waveguides less efficient for waveguide diffractive optics.

Single-mode waveguides: Non-organic materials such as SiO_2 (used as the inter-metal dielectric on ICs) can also be used as the foundation material for low loss waveguides. Cross sections comparable to the wavelength of semiconductor lasers are readily fabricated. These waveguides can be designed to propagate only a single mode, convenient for integration with diffractive elements embedded in the waveguide. For example, SiO_2-based single-mode waveguides could be fabricated directly on a silicon VLSI wafer prior to fabrication of electronic circuitry on that wafer. However, such an approach would require significant adaptation of the VLSI technologies, a significant barrier. Alternatively, assuming a sufficiently planar surface, such single mode waveguides may be fabricated above the metal layers of an IC. In the case of an MCM using a silicon substrate, single mode waveguides on the substrate are a realistic approach.

Koh et al. [13] describe an approach, illustrated in Figure 3, to create an H-Tree using optical couplers and waveguide bends in a single-mode SiO_2 waveguide. The approach was developed for MCM-based clock distribution (considered in the next section) but is included here to illustrate both the potentials and limitations of single-mode waveguides for optical clock distribution. Figure 3a illustrates a planar coupler used at each node of the H-Tree whereas Figure 3b illustrates a two-level waveguide approach using vertically coupled waveguides (the two sections are shown slightly separated for clarity). The variable core (SiON core surrounded by an SiO_2 cladding) and buried core (silica core surrounded by an SiO_2 cladding) structures in Figure 3c are used for the planar coupled waveguides whereas the hybrid structure (using both a SiON core and a silica core) in Figure 3c is used for the vertically coupled waveguides.

The bend regions in Figures 3a and 3b, consists of a coupler through which incident optical power P_0 is equally divided (3-dB coupler) between the left and right branches followed by corners with a gradual bend (to minimize losses due to the evanescent component of the propagating optical beam). The waveguide dimensions (width w_{wg}, waveguide separation s_{cp} in the coupler, length L_{cp} in the coupling region, and radius R_{bnd} of the bend) determine the surface area for the nodes of the H-Tree. Table 1 summarizes the

dimensions of the examples reported in [13].

Table 1: Waveguide parameters for coupler/bend design of [13]. The coupling loss for the hybrid design includes only the bending loss. An additional loss (< 2 dB) occurs at the vertical coupler of the hybrid design.

OPTICAL ELEMENT OF H-TREE	SiON Waveguide	Silica Waveguide	Hybrid Design
Channel Waveguide			
Upper buffer thickness (μm)	0.3	0.1	0.1
Thickness at channel (μm)	0.6	3.0	3.0
Thickness out of channel (μm)	0.48	NA	-
Lower buffer thickness (μm)	5.0	5.0	5
Channel width (μm)	4.0	2.5	2.5
Grating Coupler			
Output coupling angle (degrees)	80	80	80
Grating period (μm)	0.99	0.99	0.99
Grating index	1.55	1.48	1.55
Grating depth (μm)	0.28	0.27	0.28
Grating length (80% efficiency)	392	17,057	333
3-dB Coupler			
Channel separation s_{cpl} (μm)	4.0	4.0	0.1
Bending radius R_{bnd} (μm)	3026	2000	2000
Coupling length L_{cpl} (μm)	1022	1208	70 or 210
Coupler bending loss (dB)	< 0.01	< 0.01	< 0.01

The power $P_L(L)$ and $P_R(L)$ into the left and right branches, respectively, is related to the input power P_0 by [13]

$$\begin{aligned} P_L(L_{cpl}) &= P_0 \cdot \sin^2(C_c L_{cpl}) \\ P_R(L_{cpl}) &= P_0 \cdot \cos^2(C_c L_{cpl}) \end{aligned} \tag{1}$$

where C_c is the coupling coefficient of the coupler. For a 3 dB coupler, $C_c L_{cpl} = \pi/4$. The coupling coefficient is related to the properties of the waveguide by

$$C \approx \frac{2h^2 p \exp(-p s_{cp})}{k n_{eff} \left(w_{wg} + 2/p\right)\left(h^2 + p^2\right)} \tag{2}$$

where

$$h^2 = k^2\left(n_{core}^2 - n_{eff}^2\right) \quad a \tag{3}$$

$$p^2 = k^2\left(n_{eff}^2 - n_{buffer}^2\right) \quad b. \tag{4}$$

In (4), $k = 2\pi/\lambda$ (λ the wavelength of the optical signal), n_{eff} is the effective refractive index. n_{core} in the refractive index of the core, and n_{buffer} is the refractive index of the buffer layer (SiO_2 in Figure 3). For the planar coupler, Table 1 shows a relatively long coupling length (about 0.1 cm), contributing significantly to the area of the H-Tree node. The close spacing between the vertically coupled waveguides in Figure 3b leads to a considerably shorter coupling distance (see Table 1).

The minimum radius R_{bd} of the bend region is limited by the allowed optical loss around the bend. For a loss of about 0.01 dB, Table 1 shows a bending radius between 0.2 and 0.3 cm, dominating the area of the H-Tree node. The total area A_{nd} is $N_{nd} = 2R_{bd} \cdot (R_{bd} + L_{cp})$, large for ICs using the values from Table 1 but reasonable for use on MCMs (for which the study in [13] was performed).

The relatively large bending radius in the example above results from the relatively small difference in the refractive index of the core and cladding of the $SiON$ waveguide, a limit which may be encountered if the waveguide is buried under overlaying layers of metal and insulator. If the waveguide is on the surface of

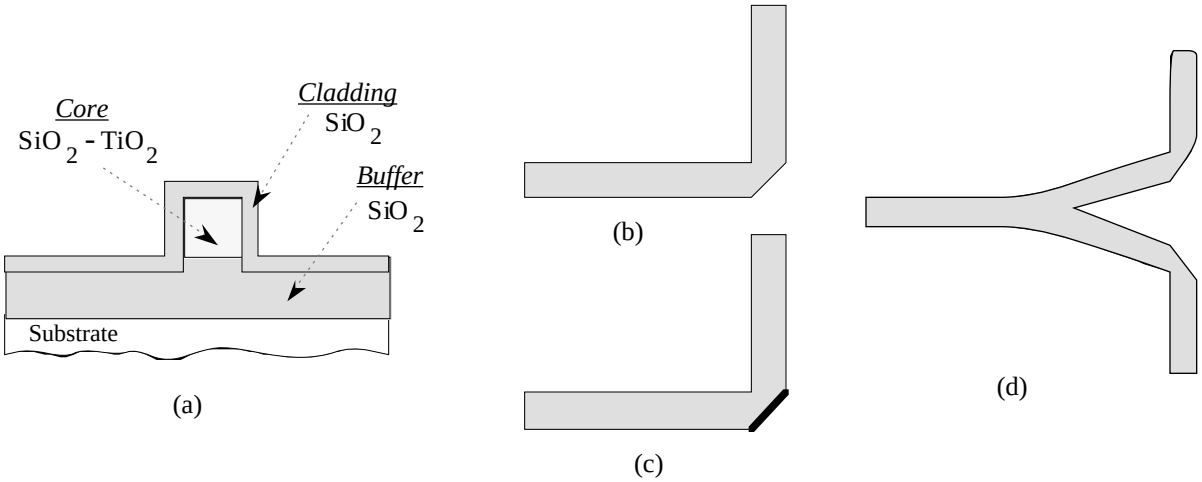


Figure 4: Right angle bends using corner mirror on waveguide. (a) Cross section of high-silica single-mode waveguide. After [14]. (b) Right angle bend with air at reflection region. (c) Right angle bend with aluminum selectively placed on reflection sidewall. (d) Illustration of a beam splitter followed by bends (less than 90°) to complete node of H-Tree.

the IC or MCM, then right angle bends can be obtained by drawing on the larger index change between the waveguide and air. For example, Himeno et al. [14] investigated sharp bends (up to 90°) using the reflecting mirror structure illustrated in Figure 4. For a reflecting corner (air interface at mirror – Figure 4b), the bending loss was found to be in the range 0.9 – 2.5 dB per corner up to 90° but increased sharply at bend angles slightly greater the 90°. Metalized corners (Figure 4c) showed a bending loss less than 3 dB, higher than for the air-terminated mirror but extending to quite large bend angles (e.g., up to 140°). The 90° bend would significantly reduce the area of the H-Tree node (compared to the results in [13], though a sharper bend would be obtained in that case if increased loss was allowed). High losses at the bend are problematic in the H-Tree network since there are potentially many bends between the root and leaf nodes of the tree. The mirrored bends in [14] could be combined with a "Y-junction" beam splitter, as shown in Figure 4d. However, the gradual bend required for a low loss splitter would impose a lower limit on the area of the node.

At the terminal points of the H-Tree network in Figure 3a, a grating coupler is etched in the waveguide, redirecting the beam toward the underlying surface. Suitable detectors and receiver circuitry completes the conversion of the optical clock to an electronic clock signal. Table 1 provides the grating length for the various waveguides used in [13], showing the considerable variation in length required for different waveguide structures.

Larger cross section, multi-mode waveguides can also be used to route the optical clock signal in an H-Tree configuration. The bends themselves would probably be right-angle splitters (e.g., as in Figure 4d) but with a substantially shorter Y-splitter length. Diffractive gratings to couple light into the IC would be less efficient and the multi-mode waveguides may favor a mirror reflection at terminal points of the network. The lower performance of this multi-mode waveguide approach may be compensated by the more convenient waveguide fabrication technology, depositing and patterning polymers such as polyimide on the surface of a completed IC or MCM substrate.

2.4 Planar Diffractive Units for Optical Clock Distribution

Rather than placing optical structures for planar routing of optical beams directly on the IC (or MCM substrate), a separate optical planar distribution structure can be used, as illustrated in Figure 2b (using optical waveguides with gratings) and Figure 2d (using optical waveguides, a beam splitter, and coupling gratings). Figure 5 illustrates two adaptions of the beam splitter approach. Both examples support equal length optical paths to multiple destinations, similar to the H-Tree discussed earlier (with Figure 5b directly illustrating a four-leaf H-Tree using two optical planes).

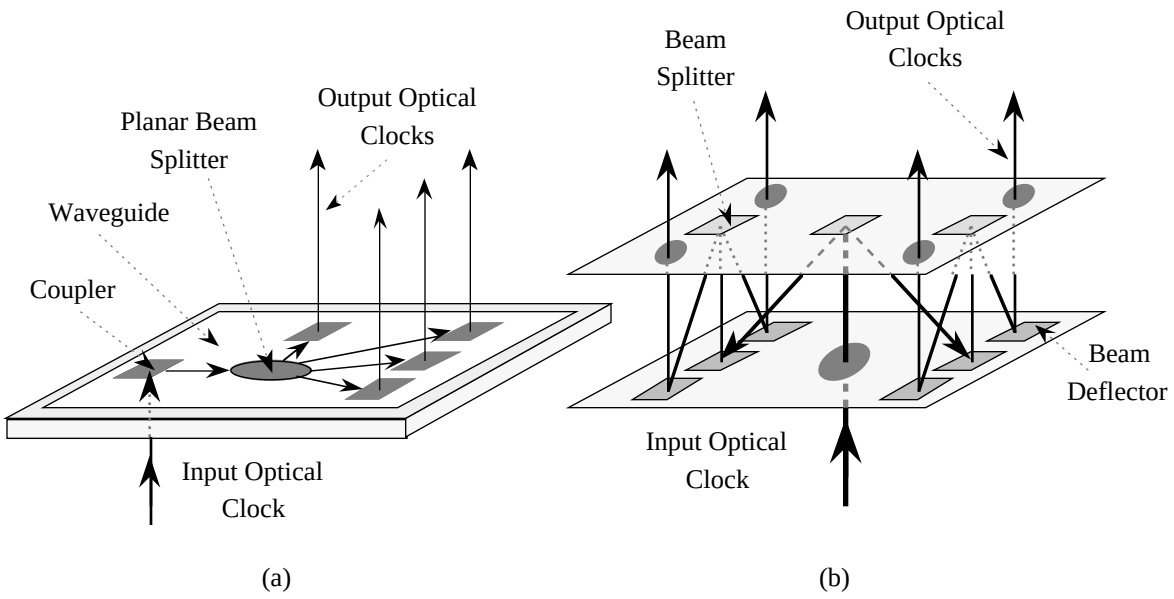


Figure 5: Example of planar diffractive optics for clock distribution. (a) Approach using waveguide holograms (after [15]). (b) Approach using free space optics (after [17]).

In the experimental version of the structure in Figure 5a reported in [15], the waveguide was a low-loss, ion-exchanged glass planar waveguide, with the planar holograms formed in a layer of dichromated gelatin deposited on the glass place. 1-to-3 interconnection fanout was demonstrated at $\lambda = 633$ nm at selected angles between 30° and 60°, using an coupler interaction length of 700 μm. Like other diffractive optical systems, wavelength shifts modify the expected behavior of the optical elements. Lin et al. [16] consider the effect of such wavelength dispersion on the input grating coupler, the waveguide hologram, the output grating coupler, and together the overall end-to-end connection.

Figure 5b illustrates a "planar micro-optic system" described by Walker et al. [17]. This example is designed to split an incident beam into four beams, the four beams exiting from the opposite surface as the entering beam. By following the three-dimensional paths, the optical power distribution is seen to be a 3-D H-Tree structure. Such a folded, 3-D approach reduces the vertical height of the overall unit, collapsing several diffractive planes into a pair of planes. The experimental example was constructed on a 3 mm thick, silver-coated glass plate. Beam splitters were formed using binary gratings while the beam deflectors used 4-level gratings, both designed for 850 nm light. Performance was evaluated for 1 and 2 μm grating periods.

Figure 6a illustrates an approach developed by Kubota and Taneda [18] which, like the general approach in Figure 2b, uses a wide waveguide, covering the full area of the IC (or MCM). Grating couplers are formed in openings in the reflective layer covering the waveguide. Optical power is input through a prism element (shown at the left hand side of the wide waveguide) which injects light at an incident angle Θ_{in} (measured relative to the normal to the waveguide direction), with Θ_{in} greater than the critical angle for total internal reflection. On striking the phase grating (Figure 6b), part of the light is coupled out of the waveguide. Letting t_{wgd} be the thickness of the waveguide, p_{grt} the period of the grating, n_{core} the index or refraction of the core, n_{top} the index of refraction of the medium above the grating, and $k_{opt} = 2\pi/\lambda$ (λ the wavelength of the optical signal), the exit angle Θ_{out} and Θ_{in} are related by [18]

$$\begin{aligned} k_{opt} n_{top} \sin(\Theta_{out}) &= k_{opt} n_{core} \sin(\Theta_{in}) \\ &\quad -2\pi m p_{grt} \end{aligned} \tag{5}$$

with m is an integer. Given Θ_{in}, λ, and the waveguide materials, the only variable parameter in (5) is p_{grt}, which is selected in this example to provide output beams normal to the surface of the plate.

Holograms have been studied to generate a regular array (linear or 2-D array) of optical beams from a single beam. To achieve efficient generation of high contrast spots, multi-level holograms (i.e., the surface contours are defined by surface relief quantized to several levels) rather than binary holograms (surface relief

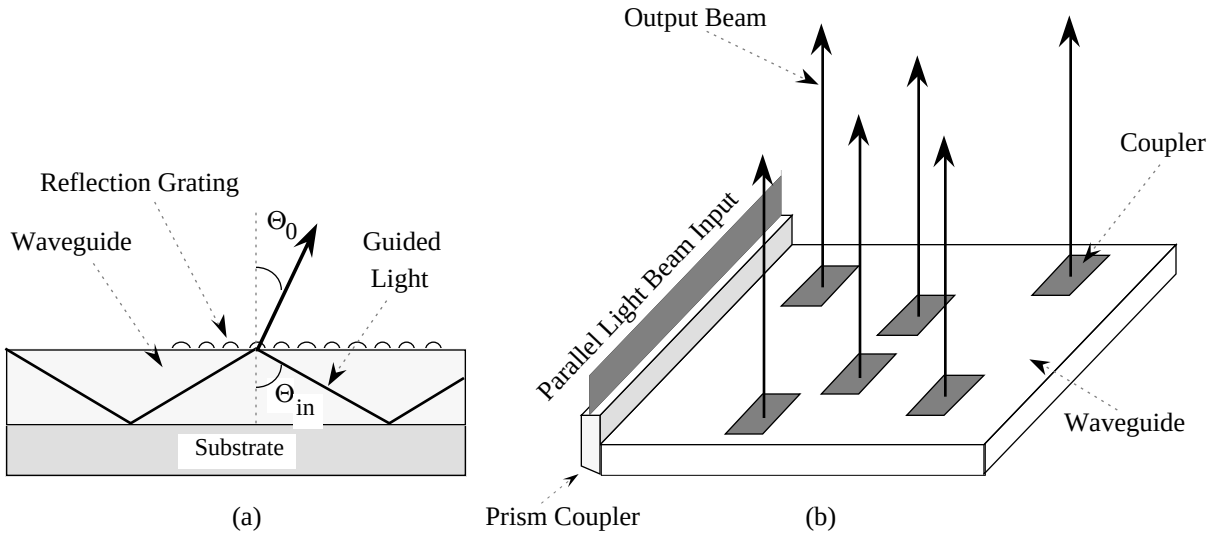


Figure 6: Parallel beam generation using a wide waveguide and grating couplers. (a) Basic coupler/waveguide structure (after [18]). (b) Beam generator with input beam coupled to waveguide using a prism and couplers placed arbitrarily on the waveguide (after [18]).

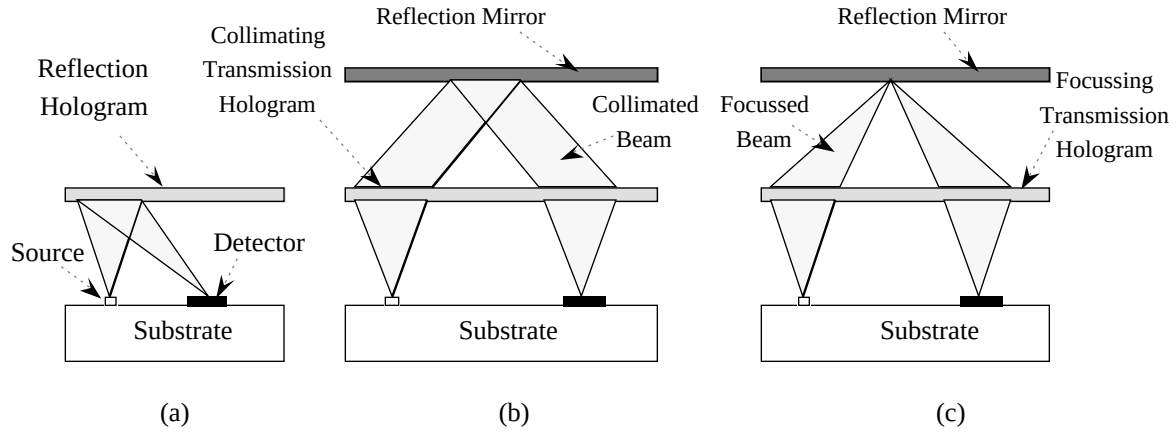


Figure 7: Basic holographic optical distribution approaches. (a) Single reflection hologram. (b) Two-pass approach using reflection mirror and collimated beams between hologram and mirror. (c) Two-pass approach using reflection mirror and optical beam focussed on mirror.

quantized to two levels, on and off) are normally used. Such multi-level relief structures provide a closer approximation to the continuously varying surface relief of an "analog" hologram by coding the surface relief into a binary coded approximation of that continually varying surface relief. The resulting binary coded surface relief is fabricated through a sequence of etching and masking steps. For example, starting with the smallest etch depth, successively deeper etches (increasing by a factor of 2) are made, with the total etch depth at any point on the surface the sum of the individual etch depths. Constraints on the fabrication of high efficiency, multi-level holograms for generation of arrays of optical beams is discussed, for example, in [19, 20]. High efficiency is important to avoid excessive background light accompanying the optical beams.

2.5 Holographic Distribution of Optical Beams to IC clock nodes

There is a very rich literature on the potential application of holograms for optical interconnections at the IC and MCM levels. Most of the investigations have considered providing a multiplicity of optical point-to-point data interconnections, with the optical signal both originating and terminating on the circuit module.

Holographic distribution of optical clocks is a specialized case, using a single source but requiring the splitting of a single optical beam into several beams.

2.5.1 Basic Holographic Optical Data Interconnects

Three basic hologram schemes for transmitting light from a source on the IC to a detector on the IC are illustrated in Figure 7. The reflection hologram in Figure 7a gathers the diverging beam from the optical source, reflecting it back to a detector while focusing the beam on the detector. The two-pass approaches in Figures 7a and b use the hologram to first collect the light from the source and to focus the light beam returning from the mirror on the detector. Figures 7b and c differ in the use of a collimated beam or a focusing beam between the hologram and the mirror. A variety of hologram structures can be used. For distribution of light from several sources, each source's output directed to one or more locations, a multifacet hologram is generally used. In this case, the area of the hologram is divided into regions (facets), each facet illuminated by a single laser source.

The hologram planes in Figure 7 must be about 1 cm from the surface of the IC due to the f/# of the hologram facet and the divergence angle of light from a semiconductor laser. The f/# of a lens is the ratio of its focal length f_{lens} to its diameter (aperture) a_{lens} (i.e., f/# = f_{lens}/a_{lens}). Semiconductor lasers generally produces an output beam with divergence angles of about 30° (8° for VCSELs – Vertical Cavity Surface Emitting Lasers). To collect all the light from the source, the aperture of the hologram facet redirecting the light is related to the divergence angle Θ_{div} of the laser beam by

$$tan(\Theta_{div}) = \frac{a_{lens}}{2 \cdot f_{lens}} = \frac{1}{2 \cdot f/\#}. \tag{6}$$

A lens with an f/# ≈ 1 therefore collects light originating at a distance equal to its focal length with a divergence angle of about 26.5°, reasonably well matched to the output characteristics of representative semiconductor lasers.

2.5.2 Alignment and Speed Issues for Holographic Interconnections

Alignment requirements for such holographic approaches have been widely discussed (e.g., [21, 22, 23]). Patra et al. [21] provide a recent theoretical treatment of the effects of mechanical misalignments (lateral shift in plane of hologram, longitudinal shift in separation between IC and hologram, and tilt of the hologram relative to the IC), misalignments due to thermal effects (e.g., coefficient of thermal expansion acting on physical distances in the hologram), and wavelength shifts of the optical source from the wavelength for which the hologram was designed. In [23], such alignment errors were evaluated to determine the relative advantages of the two approaches shown in Figure 7b and c. Representative misalignment mechanisms and beam area issues, including those illustrated in Figure 8, are as follows.

- A small tilt angle $\delta\theta$ of the mirror (Figure 8a), shifts the position of the collimated beam at the collecting lens by [23]

$$\Delta x = \frac{2h \cdot \delta\theta}{cos^2(\theta) - 2\delta\theta \cdot \cos(\theta)\sin(\theta)} \tag{7}$$

  This lateral shift can become substantial for large nominal angles θ approaching $\pi/2$, as needed to connect to a detector well separated from the source.

- In Figure 8b, the diffractive grating (with grating period p_{grt}) has a nominal reflection angle θ_{grt} for the design wavelength λ. However, the reflection angle varies with wavelength λ. Bradley, et al. [24] and others have discussed the resulting shift in the position of an optical beam.. Using the formulation of [24], $d\theta_{grt}/d\lambda = 1/[p_{grt} \cdot \cos(\theta_{grt})]$, leading to a horizontal displacement Δx of the beam (after a total propagation distance L_{prop}) due to a wavelength shift $\Delta\lambda$ given by

$$\Delta x = \left(\frac{d\theta_{grt}}{d\lambda}\right) \cdot \Delta\lambda \cdot L_{prop} \tag{8}$$

  For $p_{grt} = 1\ \mu m$, $L_{prop} = 1$ cm, and $\theta_{grt} < 45°$, each 1 Å shift in wavelength displaces the beam position by about 1 μm.

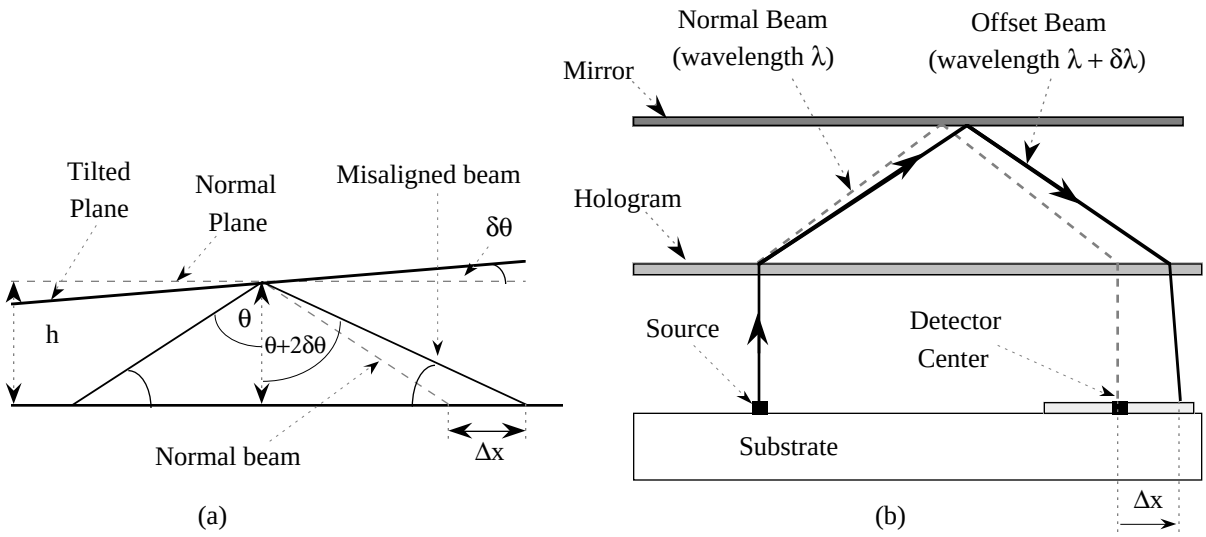


Figure 8: Examples of beam misalignment. (a) Tilt of mirror for two-pass approaches in Figure 7b and c. (b) Wavelength shift $\delta\lambda$ in source wavelength, leading to changed diffraction angle at hologram and beam shift Δx at detector.

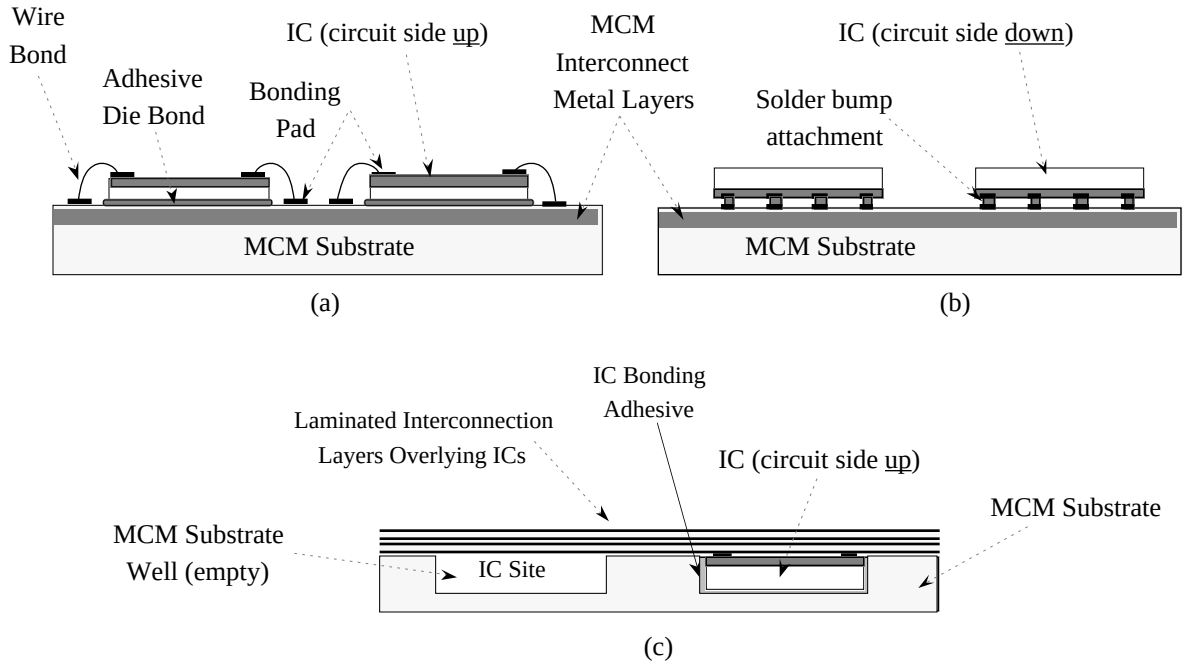


Figure 9: Basic IC mounting in MCM substrates. (a) "Chips last" epoxy bonding of IC (circuit side up) to MCM with wire bonding (or TAB bonding) to MCM substrate. (b) "Chips last " flip chip mounting (circuit side down) of ICs using solder bump connections to MCM substrate. (c) "chips first" approach with ICs placed (circuit side up) in recessed wells of MCM substrate and MCM interconnections placed on metallization layers running over ICs.

There may also be a significant temperature dependent shift in beam position. For example, temperature changes can lead to changes in the grating periods due to thermal expansion effects, leading to changes in the focal length of a diffractive lens and changes in the diffraction angle of a linear grating. Jahns et al. [26] considered such temperature changes in an optical interconnection scheme and, for the materials use in the study, the effect of a 100K temperature change was very small (e.g., lateral shift and defocusing well below 1 μm). Behrmann and Bowen [27] describe the general influence of temperature on diffractive lenses. In addition, the wavelength of a semiconductor laser shifts by about 1 Å per degree (Celcius). In an environment with temperatures changing substantially (e.g., 100C), the temperature-dependent wavelength shift can become a serious practical limit. Achromatic holograms may reduce this limitation [28].

- A collimated beam's width $w_{bm}(z)$ increases with increasing propagation distance z due to diffraction. Assuming a Gaussian beam with initial width $w_{bm}(0)$, [23]

$$w_{bm}(z) = w_{bm}(0)\left\{1 + \frac{\lambda^2 z^2}{\pi^2 n^2 w_{bm}(0)^4}\right\}^{1/2} \tag{9}$$

  where n is the index of refraction of the propagating medium. Assuming a 100 μm diameter hologram facet and a 1 μm wavelength, the beam width doubles at a distance $z = 85$ mm when propagating in air.

- Due to diffraction effects, light emitted from an aperture of dimension a in a medium with refractive index n diverges at an angle $\phi = \lambda/(n \cdot a)$. In the case of a perfectly collimated beam incident on a collector lens (acting as a source with aperture a defined by the facet diameter), the diffraction limited focused spot size s_{spot} at a detector a distance d from the aperture is

$$s_{spot} = \frac{\lambda \cdot d}{n \cdot a}. \tag{10}$$

2.5.3 Timing Errors in Optical Interconnections

Delay, jitter, and timing errors across an end-to-end optical interconnection for electronic systems includes contributions from a variety of sources. In the case of optical clock distribution, fixed delays introduced by each of the active elements and by propagation delay merely impose an upper limit on the rate at which clock pulses can be delivered to the system.

Timing jitter related to the electrical clock generated from the received optical clock signal includes jitter associated with electrical noise in the laser driver, with electrical noise in the laser bias, with the turn-on jitter of the laser, with the detector's opto-electric conversion timing jitter, with electrical noise in the detector, and with electrical noise in the receiver circuitry. Generally, timing jitter in semiconductor lasers is rather small (see, for example, [29, 30]). Timing jitter associated with the overall source end is not a source of timing skew, since all end points receive their clock from a common optical clock source. However, jitter associated with the source end produces a statistical distribution in the arrival time of a given clock pulse, with the potential of causing occasional clock pulses to arrive before electrical signals have settled to steady values if sufficient timing margin is not provided. On the other hand, timing jitter at the detector/receiver end does produce a dynamic clock skew effect (i.e., the arrival times of a single clock pulse differs for that specific pulse among the regenerated electrical clock signals).

Crosstalk can appear in optical signals as they propagate between source and receiver, though no interactions occur between beams intersecting in free space. Crosstalk can result from coupling between waveguides and from higher-mode effects in a diffractive element for one signal coupling light into a diffractive element or detector of another signal. In the case of optical clocks, such crosstalk is another source of timing jitter, qualitatively similar to electrical noise at the detector end in impacting clock integrity..

If the response time of a detector/receiver depends on the optical power received, then variations in the simultaneously received optical power at each of the detectors produces timing skew among the regenerated electrical clocks from those detectors.

A general review of time skew, timing jitter, and turn-on delay for holographic interconnections is provided in [31]. However, the comments above suggest that the precision timing often presumed for optical clock distribution may be compromised in real-world systems.

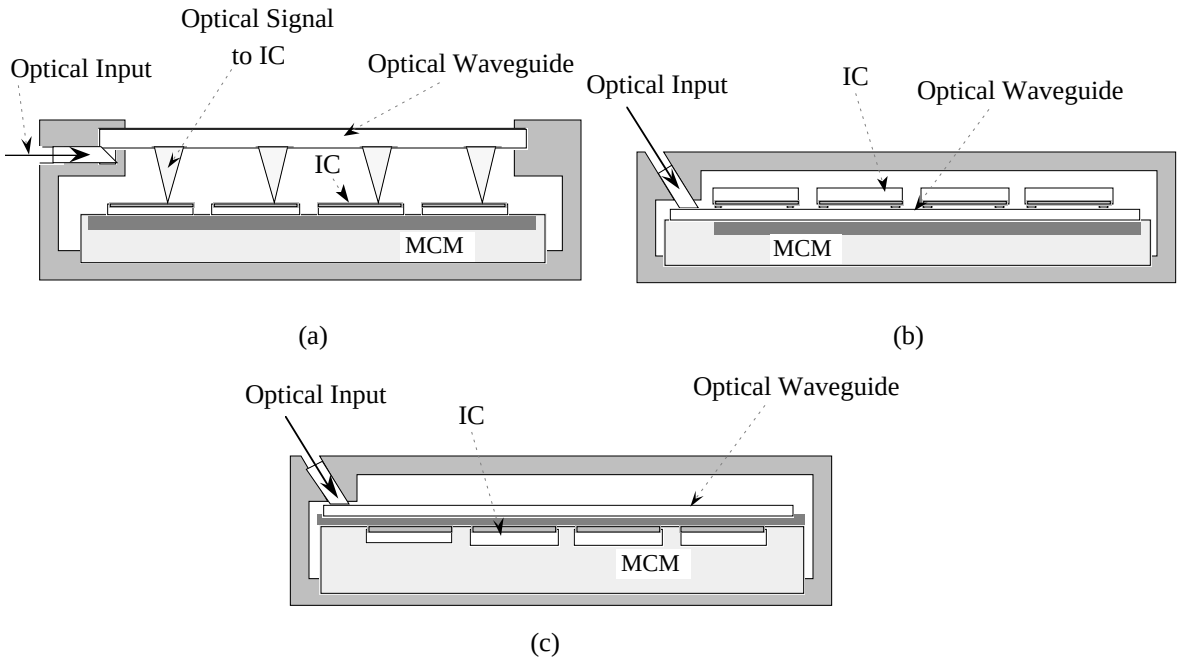


Figure 10: Examples of possible optical waveguide distribution of clock signals for approaches in Figure 9. (a) Waveguide plate in top cover of MCM package for wire-bonded IC approach of Figure 9a. (b) Waveguide distribution on top of planarized MCM substrate running under solder bumped ICs in approach of Figure 9b. (c) Waveguide distribution in laminate layers overlaying ICs in chips first approach of Figure 9c.

3 Optical Clock distribution to ICs within an MCM

3.1 General MCM Packaging Approaches and Related Optical Interconnection Approaches

The previous section reviewed techniques for intra-IC optical interconnections. Many of those techniques are also applicable to inter-IC interconnections within an MCM, where distances are longer and alignment among ICs on the MCM is substantially less precise than among devices on an IC. Discussions in Section 2 of approaches also applicable for MCMs are not repeated here. MCMs are between ICs and PCBs in terms of the distance between components (transistors and ICs, respectively) and the precision of component alignment. For this reason, the new techniques discussed in this section for optical clocks to ICs in an MCM are often adaptable for clock distribution to packaged ICs on a PCB. However, the larger area of the MCM limits application of reflection or diffraction/reflection holograms such as illustrated in Figure 7.

MCMs provide an efficient technology for modules containing not only unpackaged VLSI ICs but also other ICs, including III-V components such as semiconductor lasers and III-V detector/receiver circuits. The placement of lasers within the MCM allows generation of the optical clock within the MCM package, rather than requiring input of an external optical clock. For advanced, MCM-level systems, communications between different MCMs may be increasingly asynchronous data/packet transfers, avoiding the need to distribute a common clock to all the MCMs of a system. Compared to PCB assemblies using packaged ICs, the MCM simplifies the delivery of optical clocks to the ICs (since no package intervenes). For such reasons, the MCM level of packaging is considered here a quite different environment for optical clocks than either the IC or PCB environment. Cinato and Young provide a broad overview of optical interconnection approaches suggested for use within MCM modules [32]. A view of inter-MCM optical interconnections, arguing that this is the lowest level for optical data interconnections, is given in [33, 34].

Figure 9 illustrates the three primary approaches for placing and interconnecting unpackaged ICs on a multichip module. The approaches with the circuit side up (Figures 9a and c) generally require that the optical interconnections be placed above the ICs. The circuit-side down approach in Figure 9b generally requires that the optical signals be routed on the MCM substrate, again assuming detectors on the silicon ICs.

These conditions are relaxed by using "through-wafer" optical interconnects [44, 35], discussed in Section 4

In the case of wire bonding (Figure 9a), the alignment accuracy is modest (based on the accuracy of eutectic bonding of the IC to the substrate). The chips last approach in Figure 9c also has only modest alignment accuracy (chips must fit without difficulty into the wells of the substrate), with the general requirement that interconnect wires from the MCM substrate to the IC be custom drawn for each MCM (accounting for the IC misalignment). The approach in Figure 9b provides good IC-to-MCM alignment since the solder bumps provide a self-alignment mechanism (e.g., [36, 37]). In particular, forces exerted by the many solder bumps tend to move the overall IC into an equilibrium position in which the solder bumps are vertical, even though the starting position (before the solder is melted) may be offset.

3.2 Waveguide Distribution of Clock

Optical waveguides can be used for clock distribution to the ICs within an MCM in much the same manner as discussed in the previous section. Early work on polymer waveguides on MCMs was reported by Selvaraj et al. [38], in that case applied to wafer-scale integrated circuits. The results of [13] discussed earlier as an example of possible H-Trees on ICs was actually addressing such H-Trees on MCM substrates (where the area of the nodes of the tree is less constraining).

Figure 10 illustrates representative approaches for addition of waveguide to MCM modules. In Figure 10a, the optical plate containing the waveguides is shown as being part of the top cover of the MCM package, suitable for circuit-side-up MCM modules. By providing a mechanical alignment feature within the MCM package to precisely position the MCM substrate, alignment relative to the optical "cover plate" can be improved. Alternatively, the optical "cover plate" can perhaps be actively aligned to the MCM substrate during final assembly of the packaged MCM component.

In Figure 10b, the optical waveguides are fabricated directly on a chips-last, MCM substrate, with the waveguides aligned to the IC bonding pads on the MCM substrate. The waveguides might be polymer waveguides residing on top of the electrical connections of the MCM substrate, though surface topology may present a limitation. In the case of silicon MCM substrates (or some other MCM substrate materials), SiO_2-based waveguides can be fabricated on the silicon substrate prior to fabrication (deposited and etched layers or laminated layers) of the MCM interconnection layers.

Figure 10c illustrates optical waveguides formed in one of the laminate layers placed on the MCM substrate and over the ICs. The HDI (High Density Interconnect) MCM technology discussed in [39], a chips-first MCM technology extensively studied by General Electric Company, is a good example of the use of laminates overlaying the ICs (in recessed wells) with a planar surface well matched to the needs of waveguides.

In the case of single-mode waveguides, diffractive beam splitters, bends, gratings or mirrors can be used much as described in Section 2. However, MCMs may favor multi-mode waveguides, with their larger cross sectional areas more compatible with the larger features sizes seen in MCM interconnections (e.g., electrical line thicknesses of about 10 μm and line widths of $25 \rightarrow 60$ μm) and the correspondingly larger surface topographies (for the chips last approaches).

3.3 Planar Diffractive Optical Modules for Optical Signal Distribution among ICs

As discussed in Section 2.5.2, the reflection and diffraction/reflection hologram approaches shown in Figure 7 confront increasingly severe tolerance requirements as the reflection angle increases to route a signal to a larger distance from the source. The angle can be reduced by increasing the height of the hologram from the substrate, but this creates a large volume component from a large, planar area component such as an MCM.

To provide a compact MCM with holographic distribution of optical signals, the optical beam path of the approach in 7 can be folded, as illustrated in Figure 11. Unfolding of the end-to-end optical path (maintaining a single, reflection point near the center of the path) in Figure 11 illustrates its qualitative equivalence to a reflective/diffractive approach with a large separation between the mirror and the hologram. Alignment issues (e.g., [23, 24, 55]) for such folded systems are similar to those encountered in the non-folded approach.

The folded, planar diffractive optical element in Figure 11a has been widely discussed (e.g., [25, 26, 41, 42, 43]). In Figure 11b and c, the holograms are placed on the lower surface of the plate, which is placed

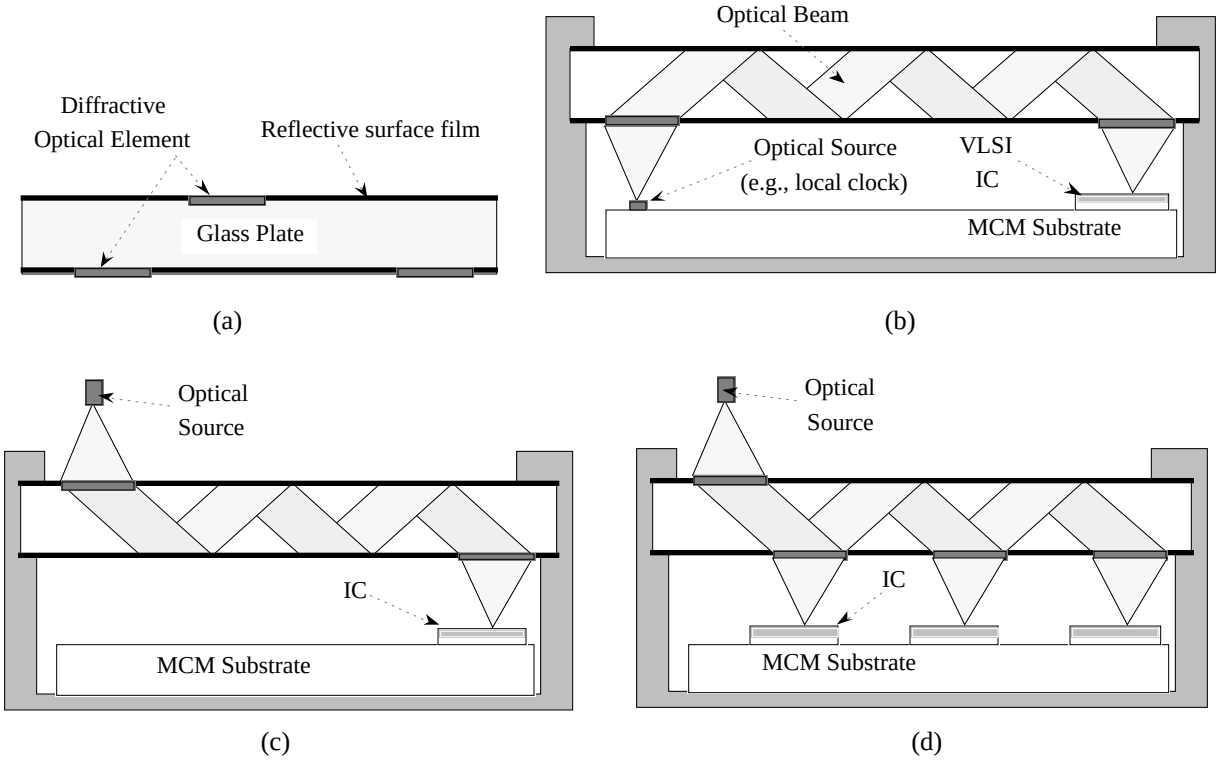


Figure 11: Folded geometry diffractive optics. (a) Thick optically transparent plate (e.g., glass) with reflective film (typically metal) with openings in the metal holding diffractive optical components (gratings, couplers, etc.) (b) Illustration of folded geometry approach for chip-to-chip interconnection within an MCM (including use of a locally generated clock). (c) Similar to (b) but using an externally provided optical signal, rather than an internally generated optical signal. (d) 1-to-many optical connection, as might be used to distribute an optical clock to several ICs in an MCM.

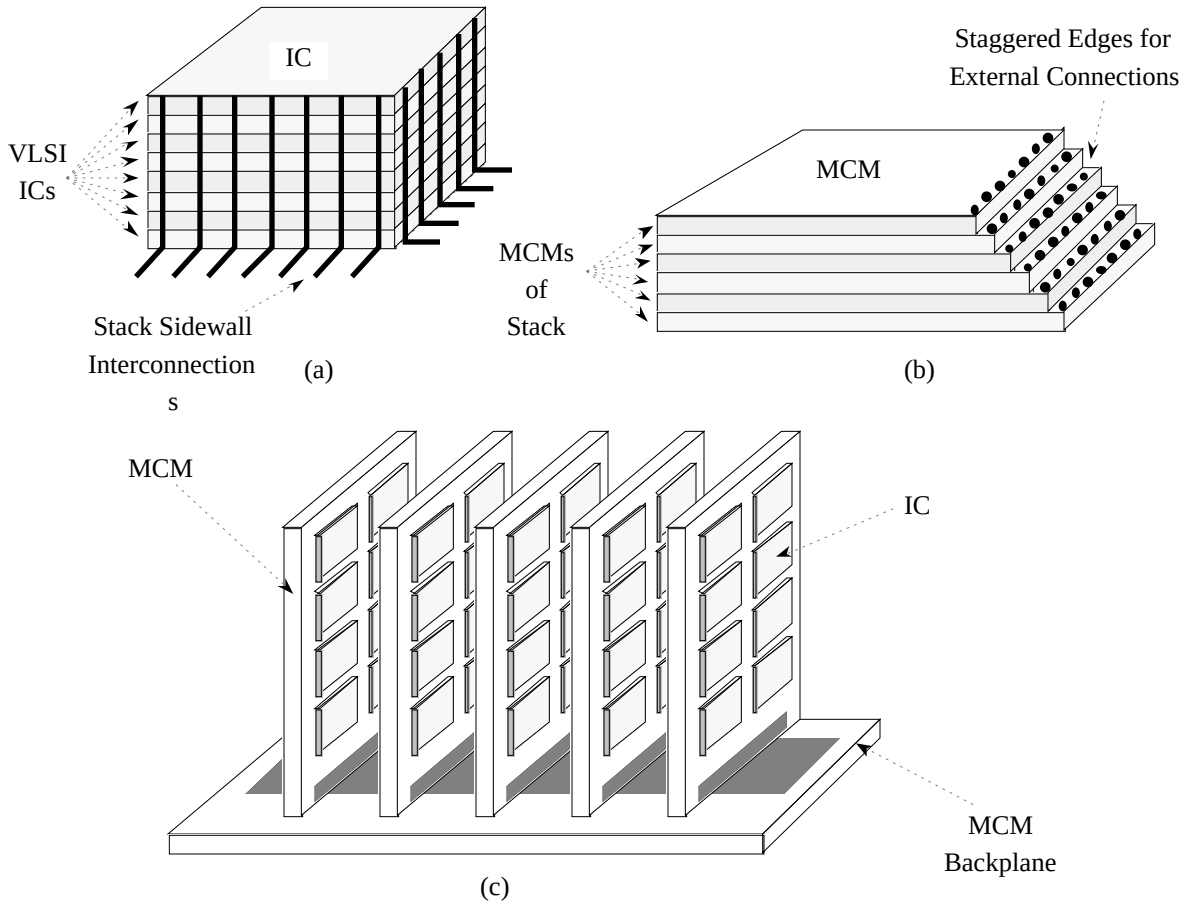


Figure 12: Three-dimensional, low level packaging approaches. (a) Stacked ICs (e.g., memory ICs). (b) stacked MCMs. (c) vertically oriented MCMs on MCM backplane "board".

about 1 cm above the MCM substrate for the same reasons discussed earlier regarding 7. In Figures 11c and d, a hologram is also placed on the top surface of the plate to receive optical clocks from outside the MCM. The hologram at the source end of the optical path collimates the optical beam and outputs that beam at an angle giving totally internal reflection within the plate (which typically has a mirrored surface), leading to the bouncing beam illustrated. At each point where the beam encounters a surface (top or bottom) of the plate, micro-optical elements can be placed allowing a variety of manipulations of the beam. In the example shown in Figure 11b and c, the beam is simply reflected between the two surfaces, terminating on a second hologram which focuses the beam onto the detector. Figure 11d illustrates multiple detector sites, as would occur for optical clock distribution. At each reflection point on the lower surface, the optical clock can be "tapped" to provide an optical clock to a detector lying below that point. The separation between tap points can be varied by changing the reflection angle of the light beam within the plate, so long as the angle meets the requirement for total internal reflection.

4 Optical Interconnections for Compact, Multi-MCM Modules

Electronic systems almost exclusively use planar packaging for the various levels from the IC itself to the rack level of system packaging. Low-level, 3-D packaging has recently received increasing attention. Examples include stacked ICs, with individual ICs thinned and then physically assembled into a compact stack.[1] Such 3-D ICs structures exploit fine dimension wiring along the edges of the stack, such as shown in Figure 12a.

[1]Both Irvine Sensors, Corp and Texas Instruments, Corp., for example, have developed such 3-D chip stacks, particularly for the example of high-density memory modules.

Compact MCM-based systems may also benefit from 3-D stacking, in this case a stack of MCMs such as illustrated in Figure 12b. 3-D MCM stacks, for example, have been investigated using the chips first approach which has advantages for stacking due to the flatter surface and more protected ICs, compared to the chips last approaches. Figure 12c illustrates another compact multi-MCM structure (using unpackaged MCMs in a backplane, rather than the 3-D, configuration) with the thin MCM modules mounted vertically on a miniature "backplane."

Although Figure 12 illustrates electrical interconnections among elements of the "stack", optical interconnections provide some interesting opportunities. Figure 13 illustrates representative approaches which might be used for optical connections among layers of the 3-D stacked MCMs in Figure 12b. Figure 13a shows optical fiber ports into an optical distribution place between pairs of MCM substrates, though such optical fiber connectors will be difficult to implement efficiently. In Figure 13b, a sidewall optical plate is shown, receiving optical signals through an optical input port and redirecting the optical signals to optical distribution planes between pairs of MCMs. Figure 13c illustrates routing of external optical signals through via holes through the MCMs, with optical power intercepted by optical distribution planes serving adjacent MCMs. In Figures 13a through c, only one optical plane is needed for each pair of MCM substrates. Figure 13d illustrates a through-wafer optical interconnection approach first reported in [44] for stacked WSI circuits. In this example, long wavelength light passes directly through the MCM substrates (assumed silicon) and is focussed using diffractive lenses etched directly on the backside of the MCM (WSI) substrates (e.g., SiN Fresnel zone plate lenses in [44]).

Figure 14 illustrates an approach suggested in [45], mounting a preassembled optical module within an MCM package to provide parallel optical free space connections through optical ports in the MCM package. To eliminate the need for precision mechanical alignment of the MCM substrate to the optical elements, wire (or TAB) bonding connects the optoelectronics of the optical I/O module to the fully electrical MCM substrate. This approach is only conceptual but was developed to illustrate that precision alignment of electronic components (or the need to adapt of standard electronic components) for insertion of optical interconnections can be avoided. In this example, the optical signals connect adjacent MCMs on a plane by positioning those MCMs very close to one another (an approach suggested in [33]).

5 Optical Clock Distribution Among PCBs

Several investigators (e.g. [46, 47, 48, 49, 50, 51, 52, 53, 28]) have addressed optical interconnections between and on PCBs, addressing the difficulty of achieving very high speed electrical signal transfers at these packaging levels.

Optical waveguides for intra-PCB interconnections have been investigated by a few groups (e.g., [48, 54]). To handle the larger surface topography of PCBs (relative to MCMs and ICs) and the increased errors in alignment, large cross section, polymer waveguides were routinely used in the studies. Separate packaged components generally were assumed for the optoelectronic devices and the silicon IC circuitry. Though several interesting approaches were investigated, the emergence of MCMs has significantly changed the conditions for integration of optical interconnections on PCBs (e.g., easy placement of optoelectronics within the same package holding the MCM electronics).

Optical interconnections between PCBs moves the discussion closer to the optical fiber ribbon connections noted briefly in the introduction for high-level system interconnections. Despite the expected commercial development of such optical fiber ribbon connections, there has been considerable investigation of the extension of techniques discussed in earlier sections to the inter-PCB case. The discussion below is restricted to such extensions.

5.1 Optical Buses for Vertically Oriented PCBs

Figure 15 illustrates a planar diffractive optical backplane for use with PCBs (e.g. [55, 56, 46, 47, 50, 57]). These examples extend the application of planar diffractive optical elements from the use with planar mounted components discussed earlier to vertically mounted components. Figure 15a illustrates a direct, point-to-point optical link between two boards. Figure 15b provides a bidirectional broadcast capability, better reflecting the functionality of a conventional bus. In both examples in Figure 15, the optical paths shown typically reflect parallel data paths (e.g., each optical path shown can be, in fact, a parallel array of

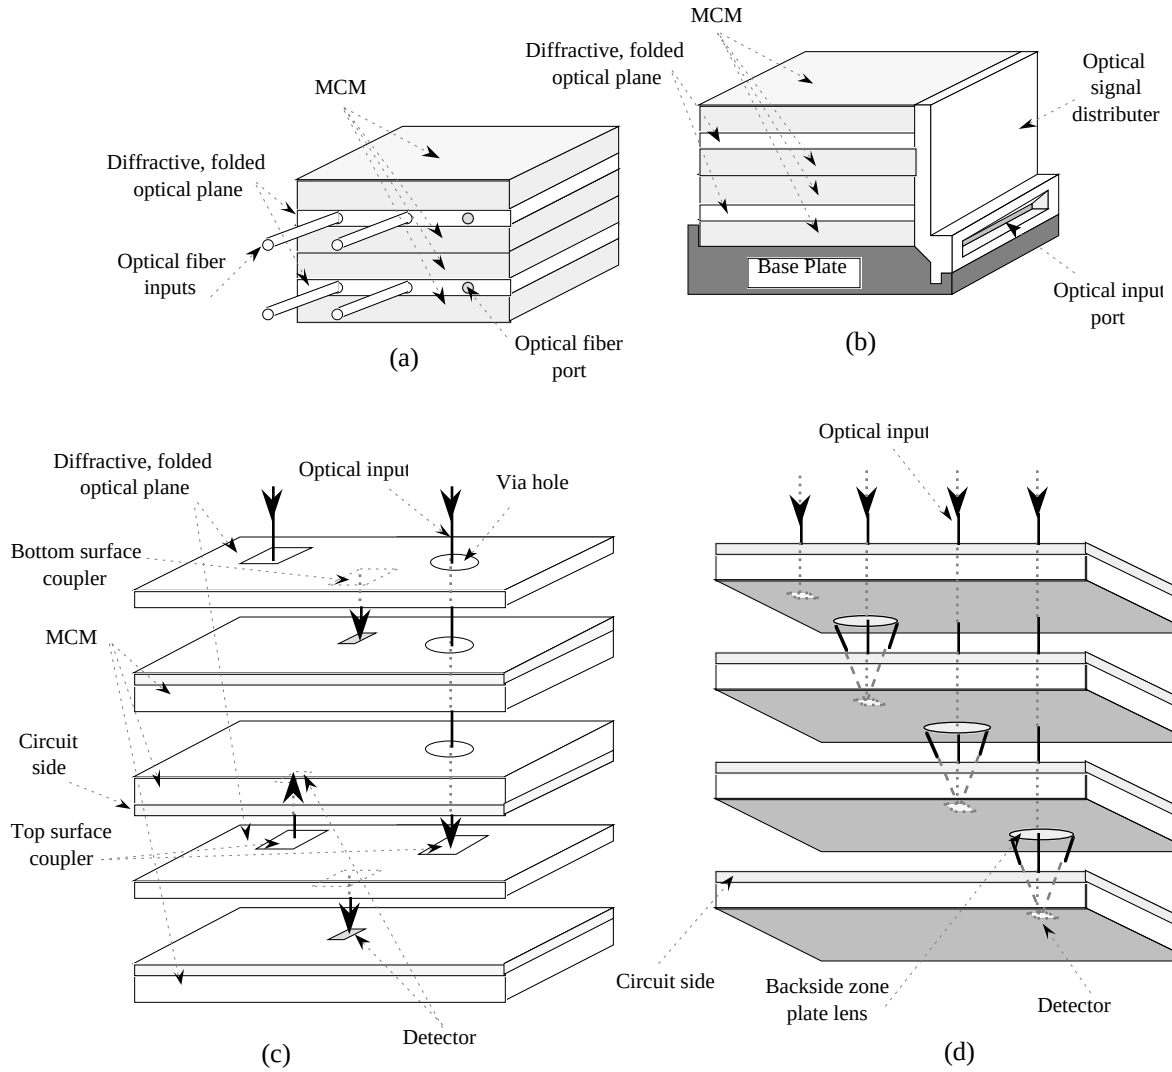


Figure 13: General approaches for optical interconnections to stacked MCMs. (a) Optical fiber connections to waveguides through sidewall of MCM. (b) Free space optical connection to waveguides through sidewall of MCM. (c) Vertical optical interconnections through optical "vias" in MCM stack. (d) Through wafer optical interconnections for MCM stack.

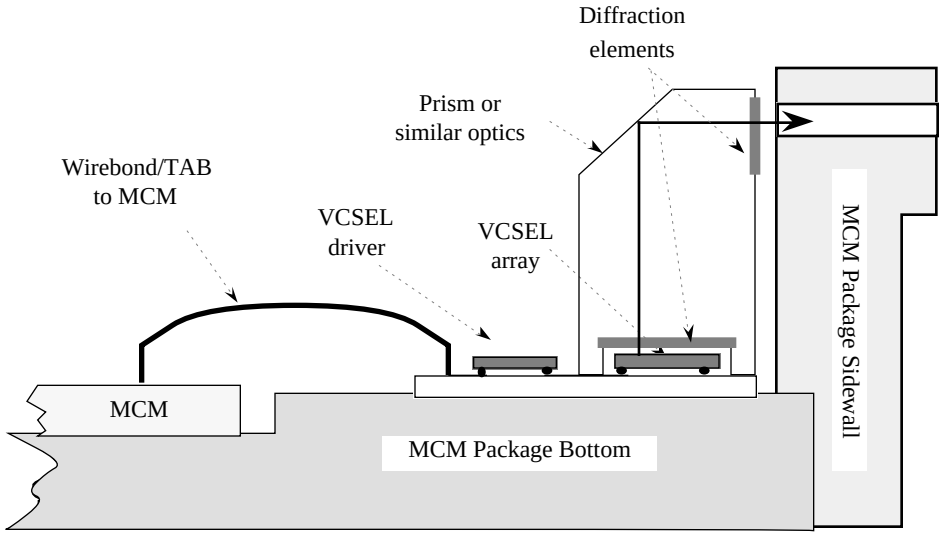


Figure 14: Direct free space optical connections between adjacent MCMs. A preassembled optical module is mounted on the package and includes optoelectronic conversions, diffractive optical elements to interface the package sidewall ports to the source/detector arrays, with wire bonding used to connect the optical elements to the fully electronic MCM substrate. (After [45].)

optical beams). High density parallel optical beams have been extensively studied for several applications, including optical computing, and the backplane bus considered here is a particularly appropriate application of such approaches. Such optical backplanes (and the direct board-to-board free space connections discussed later) are capable of substantially higher numbers of parallel beams than the optical fiber ribbon cables presently under development for commercial products.

Beech and Ghosh [55] address the issue of misalignments for point-to-point optical interconnections between pairs of optical backplane connected boards such as shown in Figure 15a, with results similar to those already discussed for related folded diffractive optics.

Natarajan et al. [57] describes the more interesting case of a bidirectional optical bus shown in Figure 15b, with each PCB reading optical data placed on the bus by any PCB. The bidirectionality is achieved through use of a holographic beam splitter on the top plate of the optical distribution plate, the hologram splitting the output optical beam from the PCB into two beams (one propagating to the right and the other propagating to the left, thereby reaching the connection sites of all other PCBs). For light propagating along the plate, the hologram also extracts a portion of the light (received from either the right or left of the hologram) and directs that portion to a detector on the PCB. As optical power is extracted from the propagating optical beam in the plate, the remaining optical power propagating to subsequent sites is reduced. For a backplane supporting several PCBs, the decrease in optical power can be considerable if the fraction α_{out} of power extracted from the beam is significant. For an initial optical power P_0, such extraction leads to a remaining optical power P_n after N_r reflections of $P_n = (1-\alpha_{out})_r^N \cdot P_0$ with the optical power $P_{in}(N)$ provided to the $N_r + 1$st PCB being $P_{in}(N_r + 1) = \alpha_{out} \cdot (1 - \alpha_{out})^{N_r} \cdot P_0$. For $N_r = r$, and $\alpha_{out} = 0.7$ ($\alpha_{out} = 0.9$), the ratio of the optical power into the 1st PCB to the optical power into the last PCB is 0.04 (0.39). For efficient optical signal transfers, it is advantageous to choose α_{out} to be sufficiently small that most of the optical power entered into the backplane is extracted to detectors on the PCBs, though the disparity in the amount of power received increases as α_{out} increases. The prototype described in natarajan.95 used a bouncing angle of 45° and demonstrated 1.2 Gb/s data transfers per optical "wire" at 1.3 μm wavelength.

An approach similar to that in Figure 15b was implemented by Yamanaka et al. [47]. In this example, the sensitivity of the diffraction angle of the grating beam splitters to wavelength variation was reduced by passing the light through two identical grating beam splitters (important since the multiple beams must be accurately focussed on the individual detectors of the receiver array). Using a 2 mm coupling lens to the PCB and assuming typical board spacings of about 2 cm, it is suggested in [47] that such a lens can handle

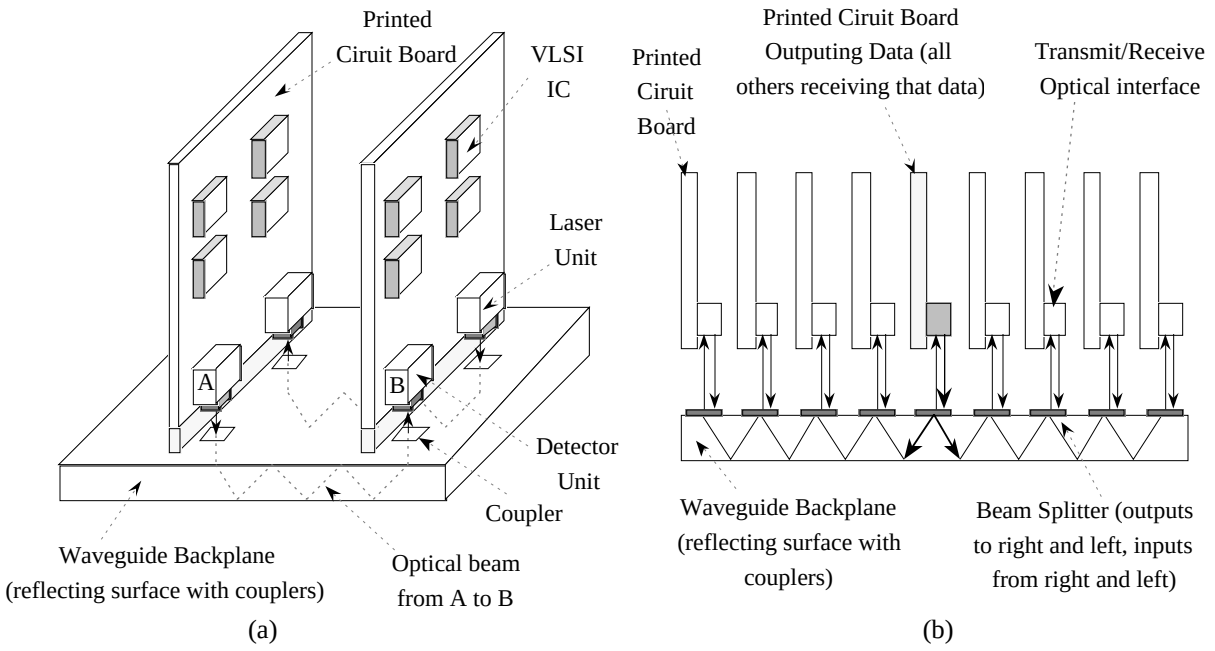


Figure 15: Optical backplane approach. (a) PCB backplane with direct point-to-point connections between pairs of PCBs (e.g., as in [55]). (b) PCB backplane with bidirectional bus (after [57]).

100 parallel connections. The experimental unit used pnpn vertical-to-surface transmission electrophotonic (VSTEP) devices [49], providing the combination of light emission, light detection, thresholding, and latching.

5.2 Direct Optical Interconnections between Vertically Oriented PCBs

Figure 16 illustrates direct free space interconnections between PCBs. In both cases shown, the length of the direct optical interconnection between the two boards can be substantially less (e.g., about 2.5 cm) than the length of interconnections routed from the outer edge of one board to the backplane, across the backplane, and then to the outer edge of the receiver board (e.g., a total distance of about 25 cm for 10 cm wide PCBs).

Figure 16a illustrates the approach investigated at MIT/Lincoln Labs [51, 52, 53]. In this approach [53], graded index (GRIN) lenses are used, one collimating the beam from the source laser on one board and the other focusing the beam on the detector of the second board. In such arrangements, the lens on the source board can be aligned with good precision to the source laser and the lens on the receiver board can similarly be aligned with good precision to the detector. As a result, if the free space beam from the transmitter board strikes the lens of the receiver board, the transmitted light will be efficiently collected by the detector. In this manner, the sensitivity of the optical efficiency (end-to-end) can be made less sensitive to lateral misalignments of the two boards (that alignment tolerance being less well controlled than alignment of components on a circuit board) by using the collimated beam between the two boards and by using an oversized receiver lens. Angular misorientations of the source board relative to the transmitter board must be small to ensure that the transmitted beam is directed toward the receiver lens (that angular misalignment having to decrease as the board separation increases). In [53], precision stops associated with the card cage holding the boards assist in establishing the required lateral alignment of the two boards (important since a small angular tilt of the board will lead to substantial displacements of points at the outer edge of the board for their nominal positions).

In the prototype [53] of the example in Figure 16a, the boards were separated by 3 cm, the transmitter's GRIN lens had a 1 mm diameter, the receiver's GRIN lens had a 2 mm diameter, and the focal length of the GRIN lenses was 1.7 mm. Lateral misalignments as great as ± 0.7 mm produced less than 20% loss of light at the detector. Similarly, only 20% light loss was produced by angular misalignments (of the two PCB planes) by $\pm 2°$.

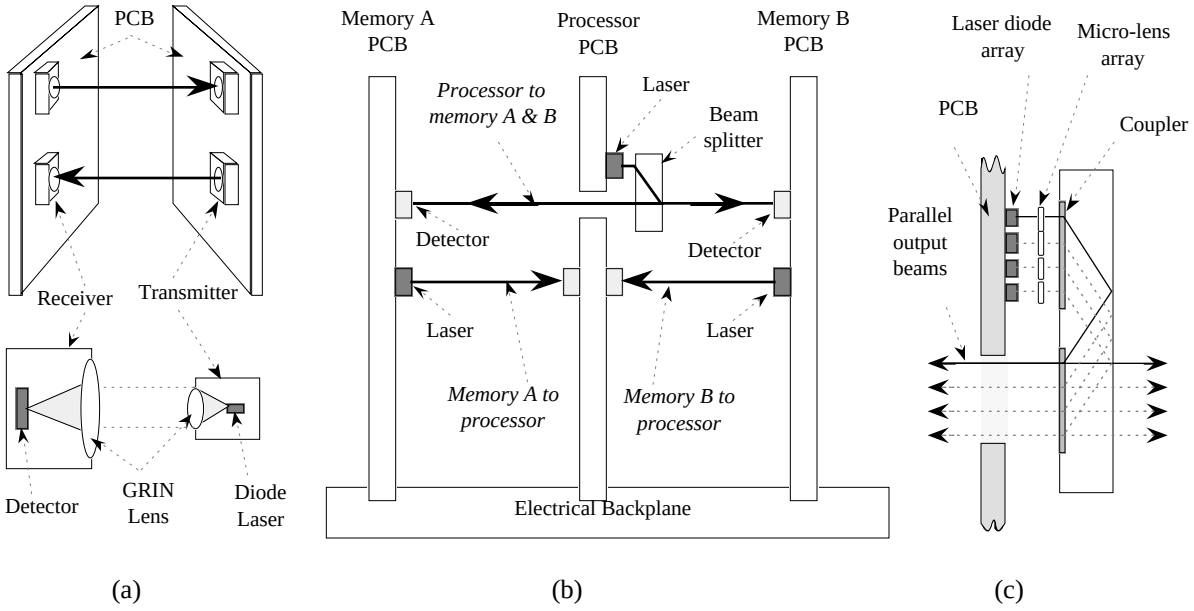


Figure 16: Direct free space optical interconnections between Vertically Oriented PCBs (a) Direct point-to-point, unidirectional connection between adjacent PCBs (after [53]). (b) Direct connections between a PCB and multiple other PCBs (after [58]). (c) Parallel beam generator and splitter on processor board in (b) (after [58]).

Figure 16b illustrates a more aggressive example, in this case using short optical interconnections between a central processor board and two memory boards to support fast memory access between the processor board and memory boards [58]. Rather than a single optical beam as in the example above, this example used parallel optical beams to implement a parallel data bus (each optical "path" in Figure 16b is, in fact, an array of beams). The transmit end at the processor board communicates the same data to both memory boards, requiring that the processor board's output optical beam be split into two beams (see optoelectronic source module of processor board shown in Figure 16c), one propagating to the left side memory board and the other propagating to the right side memory board. Several practical issues related to this general approach are addressed in [58], providing a good reference for the general technique of direct, inter-PCB optical interconnects. Assuming a 2.5 cm distance between boards in Figure 16b, the propagation delay of light between the two boards is about 83 psec, substantially shorter than would be encountered for an electrical signal traversing the two boards and connected to the backplane. However, as noted in [58], the total delay imposed by the optoelectronic elements greatly exceeds this propagation delay. The various electrical and optoelectronic components interposed between the electronic devices which would transmit and receive the signal in the absence of the optical interconnection include the line driver between the digital logic and the laser driver, the laser driver circuitry, the laser itself (with a given turn-on delay), the optical detector (e.g., photodiode), the optical receiver circuitry, and the line driver from the receiver to the logic circuitry. In the example implemented in [58], the laser driver delay (about 100 psec), the laser diode turn-on delay (about 50 psec), and the receiver (high performance design) delay (about 500 psec) combine to greatly exceed the propagation delay.

Alignment issues are also addressed in [58]. In this prototype, the focal length of the collimator and collector lens is 1.5 mm, the board separation is 2.5 cm, and the optical wavelength is 850 nm. The collimator len's radius was 165 μm. Under perfect alignment conditions, the radius of the the optical beam at the collector lens is 40 μm, leading to a beam radius of $10\mu m$ at the detector. A longitudinal misalignment of 500 μm increased the beam width at the collector lens by only about 2 μm. Chromatic dispersion of ± 5 nm increased the beam radius at the collector lens by about 20 μm. The total beam radius (in the absence of lateral and angular misalignments) is therefore 62 μm. With this nominal beam size, the size of the collector lens sets the bounds on lateral and angular misalignments. For example, a collector lens with

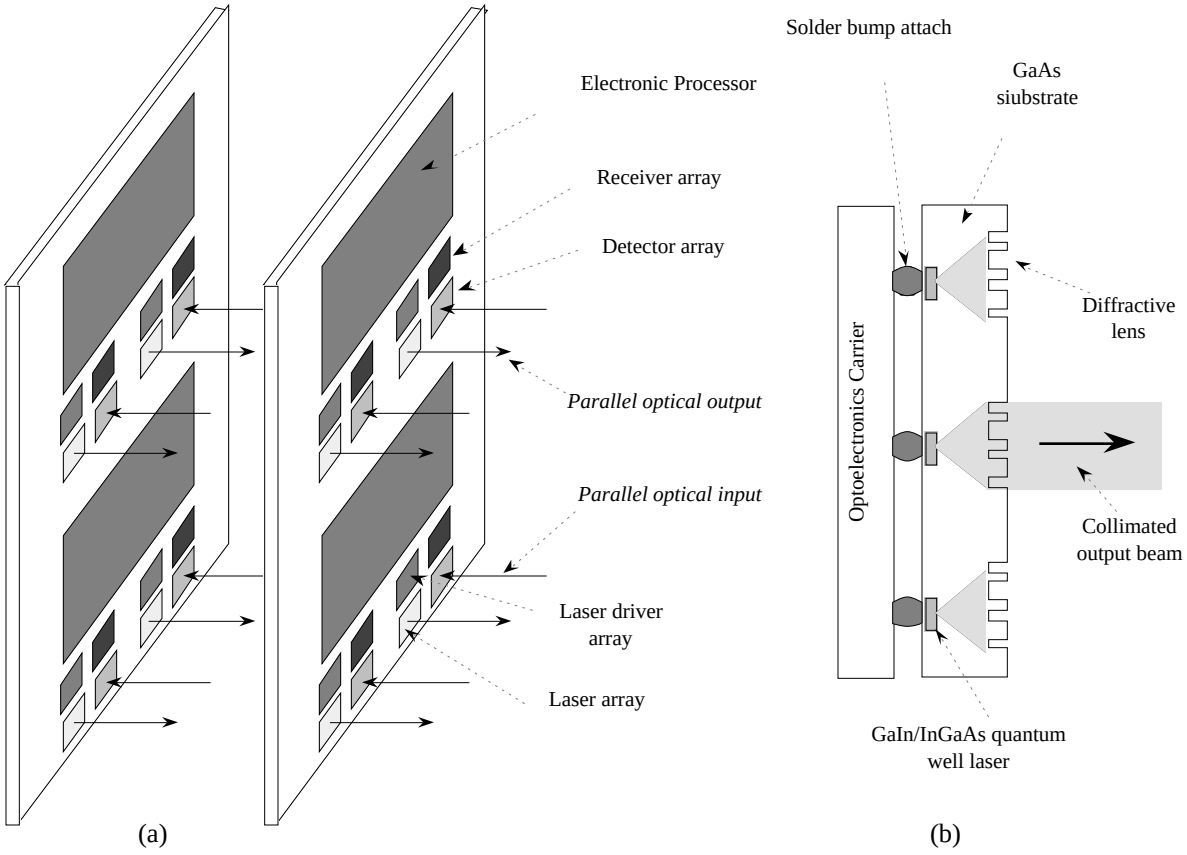


Figure 17: Cointegrated LEDs and diffractive lens for source arrays, after [60]. (a) Illustration of multiple parallel beams extending between adjacent PCBs. (b) Illustration of LED with backside integrated lens, attached by solder bumps to the module carrier.

radius 250 μm would allow $\pm(250-62) = \pm 188$ μm, allowing lateral misalignments of ± 188 μm or angular misalignments (tolerance at lens divided by distance between boards) of $\pm 0.43°$. Such strict requirements for angular alignment are typically encountered, with increasing length of the beam path magnifying the effect of a given angular misalignment.

Design of optical interconnects such as shown in Figure 16 is complex, with several competing design constraints. Computer-aided design tools promise to reduce the design complexity, with an early example of such a tool provided in [59].

The examples above have used distinct elements for the optical source (detector) and the lens arrays associated with the source (detector). However, lenses have been directly integrated on the optoelectronic devices directly. For example, Dhoedt et al. [60] describe the fabrication of the lens array directly on the back side of the GaAs substrate in which the optical sources (in this case LEDs) are fabricated. The optical wavelength is restricted to those for which the substrate is transparent (in this example, 925 nm optical beans were used). Figure 17a illustrates the PCB interconnection architecture investigated in [60]. Figure 17b illustrates the overall structure of the source array module, with the InAs/InGaAs quantum well LEDs toward the solder bump side and Fresnel lenses fabricated on the back of the substrate. Arrays of such devices allow compact source array and detector array units to be mounted at various points on the PCB, with parallel optical data beams providing direct point-to-point connections for each pair of source/detector sites. Design rules for the binary lens are given in [60]. In this prototype, solder bump attachment to the optical carrier's substrate provides close alignment of the arrays to the carrier, as discussed earlier.

6 Summary

An overview of the many approaches and studies of optical interconnections, with a focus here on optical clock distribution using optical interconnections, has been provided, starting with the intra-IC level of connection and extending to the inter-PCB level of connection. A common thrust in the various approaches studied is the importance of diffractive optical elements, waveguide structures, and free-space paths as elements with which the overall connection can be created. Miniaturization of the optical interconnection elements is a continuing priority in such applications. The techniques presented in this review are largely drawn from approaches which provide such miniaturization, consistent with the decreasing size of high performance electronic systems as the VLSI electronics technologies advance. Another priority for optical interconnections will be manufacturable optical interconnection modules and their practical insertion into real electronic systems. Several of the examples described here (e.g., the folded optical path approach) are attractive in moving toward mass producible optical interconnections. Despite the progress to date, practical insertion of optical interconnections into low levels of electronics has been elusive to date. This is to a large extent a result of the high cost of providing the initial optical interconnections, before a manufacturing environment and equipment have been developed for mass production. However, another barrier, and perhaps a more profound barrier, lies in the interface between the optical interconnections and the standard VLSI circuits which are mass produced in vast numbers. Any adaptions of the manufacturing environment for such VLSI circuits to handle optical interconnections as a specialty item would face considerable barriers to being supported. For such reasons, despite the considerable research and exploratory development and despite the increasing need for a cost effective alternative for critical electrical interconnections, there remains considerable research and exploratory development ahead.

References

[1] Y. Ota, R.C. Miller, S.R. Forrest, D.A. Kaplan, C.W. Seabury, R.B. Huntington, J.G. Johnson, and J.R. Potopowicz, "Twelve-channel individually addressable InGaAs/InP p-i-n photodiode and InGaAsP/InP LED arrays in a compact package," IEEE J. Lightwave Technol., vol. 5(4), 1987, pp. 1118-1122.

[2] Y.M. Wong, et al., "Technology development of a high-density 32 channel 16 Gb/s optical data link for optical interconnections for the Optoelectronic Technology Consortium (OETC)," IEEE J. Lightwave Technol., vol. 13(6), 1995, pp. 995-1016.

[3] H. Karstensen, C. Hanke, M. Honsberg, J-R. Kropp, J. Wieland, M. Blaser, P. Weger, and J. Popp, "Parallel optical interconnection for uncoded data transmission with 1 Gb/sec-per-channel capacity, high dynamic range, and low power consumption," IEEE J. Lightwave Technol., vol. 13(6), 1995, pp. 1017-1030.

[4] Special Issue: Diffractive Optics: Design, Fabrication, and Applications, Appl. Optics,1993, vol 32(14).

[5] Special Issue: Micro-Optics, Optical Eng., vol. 33(11), 1994, pp. 3504-3669.

[6] S.K. Tewksbury, "Interconnections within Microelectronic Systems," in Microelectronic System Interconnections: Performance and Modeling, S.K. Tewksbury (Ed), pp. 1-49, 1994, IEEE Press: Piscataway.

[7] S.K. Tewksbury (Ed), *Microelectronic System Interconnections: Performance and Modeling,*, IEEE Press: Piscataway, NJ (1994).

[8] J.W. Goodman, F.I. Leonberger, S-Y. Kung, and R.A. Athale, "Optical interconnections for VLSI systems," Proc. IEEE, vol. 72, 1984, pp. 850-866.

[9] Special Issue: Optical Interconnects, Appl. Opt., vol. 29, 1990, pp. 1067-1177.

[10] D.B. Clymer and J.W. Goodman, "Optical clock distribution to silicon chips," Opt. Eng., vol. 25, 1986, pp. 1103-1108.

[11] L.A. Bergman, W.H. Wu, A.R. Johnston, R. Nixon, S.C. Esener, C.C. Guest, P. Yu, T.J. Brabik, M. Feldman, and S.H. Lee, " Holographic optical interconnects for VLSI," Opt. Eng., vol. 25, 1996, pp. 1109-1118.

[12] R.K. Kostuk, L. Wang, and Y-T. Huang, "Optical clock distribution with holographic optical elements," in *Real-Time Signal Processing XI,* J.P. Letellier (Ed), Proc. SPIE, vol. 977, 1988, pp. 24-36.

[13] S. Koh, H.W. Carter, and J.T. Boyd, "Synchronous global clock distribution on multichip modules using optical waveguides," Optical Eng., vol. 33(5), 1994, pp. 1587-1595.

[14] A. Himeno, H. Terui, and M. Kobayashi, "Loss measurement and analysis of high-silica reflection bending waveguides," IEEE J. Lightwave Technol., vol. 6(1), 1988, pp. 41-46.

[15] F. Lin, E.M. Strzelecki, and T. Jann son, "Optical multiplanar VLSI interconnects based on multiplexed waveguide holograms," Appl. Optics, vol. 29(8), 1990, pp. 1126-1133.

[16] F. Lin, C. Nugyen, J. Zhu, and B.M. Hendrickson, "Dispersion effects in a single-mode holographic waveguide interconnect system," Appl. Optics, vol. 31(32), 1992, pp. 6831-6835.

[17] S.J. Walker, J. Jahns, L. Li, W.M. Mansfield, P. Mulgrew, D.M. Tennant, C.W. Roberts, L.C. West, and N.K. Ailawadi, "Design and fabrication of high-efficiency beam splitters and beam deflectors for integrated planar micro-optic systems," Appl. Optics, vol. 32(14), 1993, pp. 2494-2501.

[18] T. Kubota and M. Takeda, "Array illuminator using grating couplers," Optics Letts., vol. 14(12), 1989, pp. 651-652.

[19] J.M. Miller, M.R. Tagizadeh, J. Turunen, and N. Ross, "Multilevel-grating array generators: fabrication error analysis and experiments," Appl. Optics, vol. 32(14), 1993, pp. 2519-2525.

[20] E. Sidick, A. Knoesen, and J.N. Mait, "Design and rigorous analysis of high-efficiency array generators," Appl. Optics, vol. 32(14), 1993, pp. 2599-2605.

[21] S.K. Patra, J. Ma, V.H. Ozguz, and S.H. Lee, "Alignment issues in packaging for free-space optical interconnects," Optical Eng., vol. 33(5), 1994, pp. 1561-1570.

[22] J. Schwider, W. Stork, N. Streibl, and R. Völkel, "Possibilities and limitations of space-variant holographic optical elements for switching networks and general interconnects," Appl. Optics, vol. 31(35), 1992, pp. 7403-7410.

[23] K-H. Brenner and F. Sauer, "Diffractive-reflective optical interconnects," Appl. Optics, vol. 27(20), 1988, pp. 4251-4254.

[24] E. Bradley, P.K.L. Yu, and A.R. Johnston, "System issues relating to laser diode requirements for VLSI holographic optical interconnections," Optical Eng., vol. 28(3), 1989, pp. 201-211.

[25] J. Jahns and R.A. Brumback, "Integrated-optical split-and-shift module based on planar optics," Opt. Commun., vol. 76, 1990, pp. 318-323.

[26] J. Jahns, Y.H. Lee, C.A. Burrus, Jr., and J. Jewell, "Optical interconnects using top-surface-emitting microlasers and planar optics," Appl. Optics, vol. 31(5), 1992, pp. 592-597.

[27] G.P. Behrmann and J.P. Bowen, "Influence of temperature on diffractive lens performance," Appl. Optics, 1993, vol. 32(14), pp.2483-2489.

[28] J. Schwider, "Achromatic design of holographic optical interconnects," Optical Eng., vol. 35(3), 1996, pp. 826-831.

[29] T.M. Shen, "Timing jitter in semiconductor lasers under pseudorandom word modulation," IEEE J. Lightwave Technol., vol. 7, 1989, pp. 1394-1399.

[30] A. Weber, W. Ronghan, E. Bottcher, M. Schell, and D. Bimberg, "Measurement and simulation of the turn-on delay time jitter in gain-switched semiconductor lasers," IEEE J. Quantum Electronics, vol. 28, 1992, pp. 441-446.

[31] V.N. Morozov and W. Thomas Cathey, "Practical speed limits of free-space global holographic interconnects: time skew, jitter and turn-on delay," Appl. Optics, vol. 33(8), 1994, pp. 1380-1390.

[32] P. Cinato and K.C. Young, Jr., "Optical interconnections within multichip modules," Optical Eng., vol. 32(4), 1993, pp. 852-860.

[33] S.K. Tewksbury, *Wafer Level System Integration: Implementation Issues,* Boston: Kluwer Academic Publishers, 1989.

[34] S. K. Tewksbury and L. A. Hornak, "Multichip Modules: A platform for optical interconnections within microelectronic systems," Int. J. Optoelectronics, Devices, and Technologies, MITA Press (Japan), vol. 9(1), 1994, pp. 55-80.

[35] D.S. Wills, W.S. Lacy, C. Camperi-Ginestet, B. Buchanan, H.H. Cat, S. Wildinson, M. Lee, N.M. Jokerst, and M.A. Brooke, "A three-dimensional high-throughput architecture using through-wafer optical interconnect," IEEE J. Lightwave Technol., vol. 13(6), 1995, pp. 1085-1092.

[36] M.J. Wale, et al., "A new self-aligned technique for the assembly of integrated optical devices with optical fiber and electronic interfaces," Proc. ICIC '89, paper ThA19-7, 1989, p. 368.

[37] M.J. Goodwin, A.J. Mosely, M.G. Kearly, R.C. Morris, D.J. Groves-Kirkby, J. Thompson, R.C. Goodfellow, and I. Bennion, "Optoelectronic component arrays for optical interconnection of circuits and systems," IEEE J. Lightwave Technol., vol. 9(12), 1991, pp. 1639-1645.

[38] R. Selvaraj, H.T. Lin, and J.F. McDonald, "Integrated optical waveguides in polyimide for wafer scale integration," IEEE J. Lightwave Technol., vol. 6, 1988, pp. 1034-1037.

[39] J.C. Lyke, R. Wojnarowski, G.A. Forman, E. Bernard, R. Saia, and B. Gorowitz, "Three dimensional patterned overlay high density interconnect (HDI) technology," Journal of Microelectronic Systems Integration, vol. 1(2), 1993, pp. 99-141.

[40] R.S. Beech and A.K. Ghosh, "Optimization of alignability in integrated planar-optical interconnect packages," Appl. Optics, vol. 32(29), 1993, pp. 5741-5749.

[41] F. Sauer, "Fabrication of diffractive-reflective optical interconnects for infrared operation based on total internal reflection," Appl. Optics, vol. 28, 1989, pp. 386-388.

[42] J. Jahns and A. Huang, "Planar integration of free-space optical components," Appl. Optics, vol. 28, 1989, pp. 1602-1605.

[43] H.J. Haumann, H. Kobolla, F. Sauer, J. Schmidt, J. Schwider, W. Stork, N. Streibl, and R. Völkel, "Optoelectronic interconnections based on a light-guiding plate with holographic coupling elements," Opt. Eng., vol. 30, 1991, pp. 1620-1623.

[44] L.A. Hornak and S.K. Tewksbury, "On the feasibility of through-wafer optical interconnects for hybrid wafer-scale integrated architectures," IEEE Trans. Elect. Dev., vol. 34(7), 1987, pp. 1557-1563.

[45] L.A. Hornak, S.K. Tewksbury, J.C. Barr, W.D. Cox and K.S. Brown, "Optical interconnections and cryoelectronics: Complimentary enabling technologies for emerging mainstream systems," SPIE Photonics West Symposium, Optical Interconnections III Conference, San Jose, CA, Feb 1995.

[46] J.W. Parker, "Optical interconnection for advanced processor systems: A review of the ESPRIT II OLIVES program," IEEE J. Lightwave Technol., vol. 9, 1991, pp. 1764-1773.

[47] Y. Yamanaka, K. Yoshihara, I. Ogura, T. Numai, K. Kasahara, and Y. Ono, "Free-space optical bus using cascaded vertical-to-surface transmission electrophotonic devices," Appl. Optics, vol. 31(23), 1992, pp. 4676-4681.

[48] A. Guha, J. Briston, C. Sullivan, and A. Husain, "Optical interconnects for massively parallel architectures," Appl. Optics, vol. 29, 1990, pp. 1077-1093.

[49] K. Kasahara, Y. Tashiro, N. Hamao, M. Sugimoto, and T. Yanese, "Double heterostructure optoelectronic switch as a dynamic memory with low-power consumption," Appl. Phys. Lett., vol. 52, 1988, pp. 679-681.

[50] K. Rastani and W.M. Hubbard, "Alignment and fabrication tolerances of planar gratings for board-to-board optical interconnects," Appl. Optics, vol. 31, 1992, pp. 4863-4870.

[51] "High-speed optical interconnects for digital systems," Lincoln Labs. Journal, vol. 4, 1991, pp. 31-43.

[52] D.Z. Tsang, "Free-space board-to-board optical interconnections" in *Optical Enhancements to Computing Technology*, J.A. Heff (ed), SPIE vol. 1563, 1991, pp. 66-71.

[53] D.Z. Tsang and T.J. Goblick, "Free-space optical interconnection technology in parallel processing systems," Optical Eng., vol. 33(5), 1994, pp. 1524-1531.

[54] C.T. Sullivan, "Optical waveguide circuits for printed wire-board interconnections," Proc. SPIE, vol. 994: Optoelectronic Materials, Devices and Packaging, 1988, p. 92.

[55] R.S. Beech and A.K. Ghosh, "Optimization of alignability in integrated planar-optical interconnect packages," Appl. Optics, vol. 32(29), 1993, pp. 5741-5749.

[56] R.C. Kim, E. Chen, and F. Lin, "An optical holographic backplane interconnect system," IEEE J. Lightwave Technol., vol. 9, 1991, pp. 1650-1656.

[57] S. Natarajan, C. Zhao, and R.T. Chen, "Bi-directional optical backplane bus for general purpose multiprocessor board-to-board optoelectronic interconnects," IEEE J. Lightwave Technol., vol. 13(6), 1995, pp. 1031-1040.

[58] R.K. Kostuk, J-H. Yeh, and M. Fink, "Distributed optical data bus for board-level interconnections," Appl. Optics, vol. 32(26), 1993, pp. 5010-5021.

[59] R.K. Kostuk, "Simulation of board-level free-space optical interconnects for electronic processing," Appl. Optics, vol. 31(14), 1992, pp. 2438-2445.

[60] B. Dheodt, P. De Dobbalaer, J. Blondelle, P. Van Daile, P. Demester, and R. Baets, "Monolithic integration of diffractive lenses with LED arrays for board-to-board free space optical interconnect," IEEE J. Lightwave Technol., vol. 13(6), 1995, pp. 1065-1073.