

Optoelectronic Integrated Systems Based on Free-Space Interconnects with an Arbitrary Degree of Space Variance

TIMOTHY J. DRABIK

Invited Paper

*It is appealing to contemplate how VLSI or wafer-scale integrated systems incorporating free-space optical interconnection might outperform purely electrically interconnected systems. One important dimension of this question concerns the limits that optical physics imposes on the interconnect density of these systems: what bounds can be placed on the physical volume of an optical system that implements a particular interconnect? Another primary consideration arises when optical sources and detectors are integrated with circuit substrates, and the substrates are interconnected optically. Because the input and output planes of the optical interconnect are coincident with the substrates themselves, the underlying physical optical constraints are coupled with new considerations driven by circuit and packaging technology; namely, speed, power, circuit integration density, and concurrency of operation. Thus the overall **optoelectronic integrated system** must be treated as a whole and means sought by which optimal designs can be achieved.*

This paper first provides a uniform treatment of a general class of optical interconnects based on a Fourier-plane imaging system with an array of sources in the object plane and an array of receptors in the image plane. Sources correspond to data outputs of processing "cells," and receptors to their data inputs. A general abstract optical imaging model, capable of representing a large class of real systems, is analyzed to yield constructive upper bounds on system volume that are comparable to those arising from "3-D VLSI" computational models. These bounds, coupled with technologically derived constraints, form the heart of a design methodology for optoelectronic systems that uses electronic and optical elements each to their greatest advantage, and exploits the available spatial volume and power in the most efficient way. Many of these concepts are embodied in a demonstration project that seeks to implement a bit-serial, multiprocessing system with a radix-2 butterfly topology, and incorporates various new technology developments. The choice of a butterfly for a demonstration vehicle highlights the benefits of using free-space optics to interconnect high-wire-area topologies.

Manuscript received February 2, 1994; revised August 23, 1994. This work was supported by the National Science Foundation under Grant MIP-9110276 and by the Georgia Institute of Technology Microelectronics Research Center.

The author is with the School of Electrical and Computer Engineering and the Microelectronics Research Center, Georgia Institute of Technology, Atlanta, GA 30332-0250 USA.

IEEE Log Number 9405791.

The practice of wafer-scale integration has been hindered by the large area overhead required by sparing strategies invoked against the inevitable fabrication defects. If wafers are partitioned into functional circuit cells which are then interconnected strictly optically, then defective cells can be accommodated in a wafer-specific diffractive interconnect, with no substrate area lost to inter-cell connectivity. The task of this interconnect is to map a desired network topology onto the physical set of functional cells. Results of mapping regular processor topologies onto wafers with defective cells are given in terms of asymptotic volume complexity.

I. INTRODUCTION

Computational complexity theory seeks to determine how the resources necessary to perform a computational task grow as some measure of the "size" of the task increases [1], [2]. When integrated circuit technology became sufficient for the incorporation of entire systems on a single chip (*integrated systems*) [3], curiosity about the performance limitations of such systems led to a generalization of complexity theory that set formal bounds on the simultaneous chip area and time required to implement solutions to various problems [4], [5]. Further advances in very large scale integration (VLSI) technology that increased the number of interconnection levels available and improved the prospect of vertical integration of active devices [6]–[9], inspired further generalization to "3D VLSI" computational models that permit significant or unlimited stacking of active (transistor) or passive (interconnect) layers [10]. Recent advances in technologies for integration of self-luminous sources, passive modulators, and high-performance detectors with complex circuitry have brought closer the implementation of what can be called *optoelectronic integrated systems*, Fig. 1, based on electronic circuitry and free-space optical interconnection.

It is a pressing question whether VLSI or wafer-scale integrated (WSI) systems incorporating free-space optical interconnection might not in some ways outperform purely electrically interconnected systems. An important dimen-

0018-9219/94\$04.00 © 1994 IEEE

PROCEEDINGS OF THE IEEE, VOL. 82, NO. 11, NOVEMBER 1994

1595

FNC 1016




Fig. 1. Optoelectronic integrated systems concept: electronic substrate with electrical-optical transduction elements, interconnected optically.

sion of this question concerns the limits that optical physics imposes on the physical size of these systems: what bounds can we place on the volume of an optical system that implements a particular interconnect? We expect intuitively that optoelectronic systems with many optical signal paths will take up more space than those with fewer optical connections, just as planar VLSI layouts with many distinct electrical nets require more area than those with fewer wires. Further, VLSI complexity theory has revealed a strong connection between the communication efficiency of a network, as quantified in terms of graph-theoretical measures of connectivity, and the area required for its planar layout. Thus we also expect the required optical system volume to depend on topological aspects of the interconnect pattern.

It is desirable to find interconnect structures that offer high computational performance and lead as well to optical systems with low complexity. A treatment of optical interconnect capacity that applies to a large class of systems could serve as a tool to identify such structures and perhaps determine what attributes of interconnect structure lead to small volume requirements.

Analyses have investigated the capacity of numerous optical interconnect configurations in contexts of varying generality. Particularly, Feldman and Guest [11] established upper bounds on the area of computer-generated holograms required to implement the type of arbitrary connection patterns with fan-out that characterize signal paths in VLSI. Barakat and Reif [12] placed optical interconnect systems into the context of computational complexity theory by using arguments based on Gabor's theorem [13] to derive fundamental bounds on interconnect capacity of systems with optical sources and detectors lying on a bounding surface, and thereby enabled the direct theoretical comparison of optical interconnects with 3D VLSI circuitry. This analysis was generalized by Ozaktas and Goodman [14] to accommodate sources and detectors placed arbitrarily within a system's volume.

Along with these fundamental treatments have arisen various studies that incorporate aspects of structure or regularity of the interconnect, to derive constructive or engineering bounds [15]–[24], and recent, ever more comprehensive and general analyses have deepened understanding of the utility of optical interconnection in computing [14], [25]–[27].

This paper introduces an approach to the study of optical interconnect complexity, that yields practically applicable criteria to facilitate the *design* of optoelectronic integrated systems that efficiently use the primary resources of power and physical volume.

The first half of the paper provides a uniform treatment of a general class of optical interconnects based on a Fourier-plane imaging system with an array of sources in the object plane and an array of receptors in the image plane. The sources correspond to data outputs of processing "cells," and the receptors to their data inputs. Interconnection may be thought of as a mapping—one-to-one, or possibly one-to-many—from the set of sources to the set of receptors. In Sections II and III, an abstract optical imaging model is defined that is sufficiently general to represent a large class of real systems. Section IV relates elements of this *canonical system* to practical imaging systems, diffractive elements, detectors, modulators or self-luminous sources, and microlens arrays. In Section V, bounds are computed for the overall system volume and for other important parameters, in terms of certain characteristics of the interconnect pattern, and for an arbitrary number of cells. These are not fundamental lower bounds, but rather upper bounds determined by constructive example, that are consistent with such practical systems constraints as maintaining a constant signal-to-noise ratio and numerical aperture. Section VI uses these results to derive meaningful bounds on the physical volume of optical implementations of specific regular interconnection networks.

In the second half of the paper, results of the first half are applied to the design of optoelectronic integrated systems. Section VII introduces a methodology for the design of optoelectronic integrated systems based on planar arrays of electronic processing cells. Cells are interconnected optically by means of integrated light modulators or sources, and detectors, and an external optical routing network. This methodology establishes requirements on cell area and on the speed–power characteristics of optical sources, modulators, and detectors, and gives rise to guidelines for optimizing designs of regularly interconnected arrays. Many of the concepts discussed above are embodied in a design for a prototype system described in Section VIII, for a bit-serial, pipelined processor array. The design incorporates electrooptic and optomechanical technology building blocks recently developed through the author's research collaborations. Finally, Section IX demonstrates how the functioning physical processing cells on an optoelectronic substrate with defects can be interconnected as a regular array with no need for on-wafer redundancy or reconfiguration circuitry. The effect of wafer defects on the complexity of the optical interconnect required to accommodate them is determined for a useful topology: the 2D mesh.

II. GENERAL ASPECTS OF THE TWO-HOLOGRAM FOURIER-PLANE INTERCONNECT

Before proceeding with the rigorous descriptions and definitions needed to support the complexity analysis of Section V, it is beneficial to discuss informally the optical




Fig. 2. Optical interconnection with a Fourier-plane imaging system and beam-deviating elements.

system on which this work is based. The system is a refinement of the hybrid “basis-set” interconnection originally described by Jenkins and Strand [28] and Jenkins *et al.* [29], [30]. The functional objective is to communicate optical signals in a space-variant way from a plane \mathcal{O} containing sources to a plane \mathcal{I} containing receptors, Fig. 2. In order that light from a given source be directed to the desired receptors, it must be possible to change the direction of a beam of light as it leaves the object plane \mathcal{O} , and as it passes through the Fourier plane \mathcal{F} of the imaging system. These deviations can be realized with diffractive elements. For convenience, all elements causing beam deviation or splitting will be referred to as *holograms*, in the spirit of [31]¹.

The high interconnect capacity of the system described in this work derives from the shift invariance (or space invariance) of Fourier-plane spatial filtering: angular deflection of a beam as it passes through the Fourier plane results in equal displacements of each feature in the image plane relative to that feature’s location in the absence of the deflection. A single region of $H^{\mathcal{F}}$ that causes a specific angular beam deflection can thus be used to send light from many different sources through identical *relative* displacements. In this way, a large number of regular interconnections (involving a small total number of distinct, relative displacements) can be realized optically in a small volume.

Figure 2 illustrates how the required beam direction is accomplished with two segmented holographic elements. Beams of light leave the sources and are, in general, deflected or split upon passing through holographic element $H^{\mathcal{O}}$, which is in direct contact with the sources in the object plane \mathcal{O} . Holograms that perform deflection act on the light traversing them as do triangular prisms. The paths through the lenses of the rays representing the beams’ centers, are described by the theory of geometric optics. Both $H^{\mathcal{O}}$ and $H^{\mathcal{F}}$ are partitioned into independently programmed regions. A region in $H^{\mathcal{O}}$ serves to “point” the light leaving a source toward one or more of the regions in $H^{\mathcal{F}}$. A region of $H^{\mathcal{F}}$ receiving light imparts an angular deflection (or splits the light into multiple, differently deflected beams) that causes the resulting spot in the image plane \mathcal{I} to suffer a particular relative displacement (or to be replicated into

multiple, differently displaced spots). This way, light from different sources can be differently directed to receptors in the image plane. The behavior of this optical system is described comprehensively by Fourier optics².

As Ozaktas and Goodman have shown [14], the optical interconnect capacity of a system with sources and receptors on the convex hull of its volume is generally less than that of a system whose sources and receptors are allowed to lie anywhere within its volume. However, the former may serve as a building block for the latter in an obvious way [26] and, we will show, can attain desirable performance levels with low optomechanical complexity.

The elegance of the basis-set system of Fig. 2 lies in that an arbitrary overall degree of shift-variance can be realized by varying the number of segments in $H^{\mathcal{F}}$. If only one segment, representing a particular sum of linear phase factors, occupies the whole aperture of $H^{\mathcal{F}}$, then $H^{\mathcal{O}}$ serves no purpose and the interconnect is shift-invariant. At the other extreme, light from each source in \mathcal{O} can be directed by the segment of $H^{\mathcal{O}}$ in front of it to a distinct segment of $H^{\mathcal{F}}$. In this case each source can be mapped independently, which corresponds to total shift variance.

The degree of shift variance, represented by the number of segments into which $H^{\mathcal{F}}$ is divided, is directly linked to the physical volume required to contain the system. Each segment of $H^{\mathcal{F}}$ functions as a pupil in the Fourier plane, whose diameter determines the spatial frequency response of the imaging system [32]. This pupil diameter must be large enough to image spots of light from the sources with sufficient acuity to prevent crosstalk between adjacent sources or adjacent receptors. Thus all other things being held equal, an increase in degree of shift-variance necessitates an increase in total Fourier-plane area. If the focal lengths of lenses L_1 and L_2 are held constant, then an increase in their numerical aperture is prescribed. Numerical aperture is a useful measure of the cost of fabricating a lens, however, and it is often desirable to hold it constant and vary some other parameter to accommodate the increased demands on the imaging system. It is plain that the interplay of all optical system parameters must be

¹“‘When I use a word,’ Humpty Dumpty said, in a rather scornful tone, ‘it means what I choose it to mean—neither more, nor less.’” [31].

²A thorough treatment of this elegant theory and its applications is found in Goodman [32]. The Appendix of this paper presents essential results of the theory, establishes notational conventions, and discusses the nature of diffractive elements.




Fig. 3. Canonical Fourier-plane optical system configuration and notational conventions. Sense of the coordinate system in \mathcal{I} is opposite that in \mathcal{O} because of the image inversion that occurs upon propagation through the system.

considered methodically such that the system “grows” in a reasonable way. This is the goal of the next section.

III. A CANONICAL REPRESENTATION OF FOURIER-PLANE-BASED OPTICAL INTERCONNECT SYSTEMS

Because many physical embodiments of optical interconnects are based on a variant of the system of Fig. 2, it is sensible to define an abstract model sufficiently general to represent them all, both for simplicity of analysis and so that different systems can be evaluated comparatively. This *canonical system*, shown in Fig. 3, comprises two ideal lenses of equal focal length f and diameter D , arrays of sources in \mathcal{O} and receptors in \mathcal{I} , and two *holographic elements*: $H^{\mathcal{O}}$ in \mathcal{O} , and $H^{\mathcal{F}}$ in \mathcal{F} . All elements are centered on a common system axis, which also pierces the origins of \mathcal{O} , \mathcal{F} , and \mathcal{I} .

A. Sources and Receptors

The nature of the physical sources and receptors being interconnected can vary widely across different applications. Sources with emitting areas only a few wavelengths in extent, such as edge-emitting laser diodes, exhibit greater beam divergence than do large-area passive modulator cells or vertical-cavity lasers, for example. Physical receptors can differ as well: an integrated p-i-n photodetector might be made 10 μm in diameter or smaller to reduce parasitic capacitance and increase speed, whereas metal–semiconductor–metal (MSM) detectors $> 100 \mu\text{m}$ in diameter can perform satisfactorily [33]. Numerical aperture of the receptor elements must be taken into account, and spacing of sources and receptors on a substrate is coupled to the size of the cells being interconnected.

Different types of physical sources and receptors are accommodated by defining a simple object plane specification to which light from any type of emitter or modulator can be made to conform, and an image plane specification compatible with any type of detector: “sources” and “receptors” in the abstract model are simply regions that emit light or are sensitive to it. For ease of analysis, these regions will be understood to be circularly symmetrical, although elliptical or other shapes would better suit some device technologies.

We will consider n processing elements or *cells* P_0, P_1, \dots, P_{n-1} , each of which contains a collection of

sources and receptors. These may be arranged identically within all cells, but this is neither essential to the analysis nor always desirable. Figure 4(a) shows a typical cell P_j populated with r sources, or outputs $O_j^0-O_j^{r-1}$ and m receptors, or inputs $I_j^0-I_j^{m-1}$. Sources and receptors have diameter d , and the cell has area a^2 . Although pictured as squares in Fig. 4, cells may assume any appropriate shape.

With these conventions established, an interconnection can then be characterized by a (possibly multivalued) mapping from the sources to the receptors

$$\mathcal{M} : \left\{ \mathcal{O}_j^\ell \right\} \left| \begin{array}{l} \ell \in \{0, \dots, r-1\} \\ j \in \{0, \dots, n-1\} \end{array} \right. \rightarrow \left\{ \mathcal{I}_j^i \right\} \left| \begin{array}{l} i \in \{0, \dots, m-1\} \\ j \in \{0, \dots, n-1\} \end{array} \right. \quad (1)$$

Because of the underlying physical context of electronic substrates populated with emitters and detectors, the sources and receptors in a given cell are physically tied to it. However, we will conceptually separate all sources into plane \mathcal{O} and all receptors into plane \mathcal{I} , to conform to the linear nature of the canonical imaging system. The source and receptor planes thus appear as in Fig. 4(b). The mutual spatial relationships among the sources themselves, and among the receptors themselves, are undisturbed. Folding the system back on itself with reflective elements upon translation to a physical embodiment (Section IV-E), will reunite sources and receptors corresponding to each P_j . Sources emit monochromatic, collimated light of free-space wavelength λ normally from plane \mathcal{O} ; sources and receptors may not overlap, and the cell area a^2 must be large enough to accommodate them. It is useful to define the dimensionless quantity

$$\eta := \frac{d}{a} < 1 \quad (2)$$

where “:=” indicates definition or assignment; η can be thought of as a spatial duty cycle. Let the *object plane diameter* $w_{\mathcal{O}}$ be defined as the maximum center-to-center distance between two sources, and the *image plane diameter* $w_{\mathcal{I}}$ as the maximum center-to-center receptor separation, as in Fig. 4.

B. Holographic Elements

$H^{\mathcal{O}}$ is an array of $r \times n$ subholograms abutted directly to the sources in \mathcal{O} . Each subhologram $H_{j\ell}^{\mathcal{O}}$ intercepts all light from exactly one O_j^ℓ . $H^{\mathcal{F}}$ is an array of K subholograms located in the Fourier plane. Beams intercepted by the $H_k^{\mathcal{F}}$ undergo angular deflections that result in the spatial image plane shifts required to implement the mapping \mathcal{M} . The function of each $H_{j\ell}^{\mathcal{O}}$ is to broadcast the spot of light from O_j^ℓ to the correct subset of the K available subholograms $H_k^{\mathcal{F}}$. The required complex transmittance functions are equivalent to sums of linear phase functions

$$t_{j\ell}^{\mathcal{O}}(x, y) = \delta_{j\ell} \sum_k e^{j2\pi(\varphi_{j\ell}^k x + \psi_{j\ell}^k y)} \quad (3)$$

for (x, y) belonging to the region occupied by $H_{j\ell}^{\mathcal{O}}$, where constant $\delta_{j\ell}$ accounts for the decimation of light power




Fig. 4. (a) Source and receptor configuration in canonical cell. (b) Separation of all sources into object plane \mathcal{O} and receptors into image plane \mathcal{I} .

caused by fan-out. Similarly

$$t_k^{\mathcal{F}}(x, y) = \delta_k \sum_l e^{j2\pi(\varphi_k, x + \psi_k, y)} \quad (4)$$

for (x, y) within $H_k^{\mathcal{F}}$. In the canonical system, $H^{\mathcal{F}}$ and $H^{\mathcal{O}}$ have zero physical thickness.

Source plane subholograms $H_{j\ell}^{\mathcal{O}}$ need only have diameter d because they are in direct contact with the O_j^{ℓ} . Spots from the O_j^{ℓ} spread as they propagate to the Fourier plane because of diffraction, and an $H_k^{\mathcal{F}}$ must be able to intercept some large fraction of the power in a beamlet directed toward it from the source. Light falling outside the intended $H_k^{\mathcal{F}}$ can fall on other Fourier-plane subholograms and be misdirected, and is therefore a systematic source of crosstalk even under the assumption of ideal holographic elements. Let the $H_k^{\mathcal{F}}$ be circular with diameter h . If h is given the functional dependence

$$h = \frac{q\lambda f}{d} \quad (5)$$

where q is a form factor, then the fraction of intercepted power is independent of focal length [see (A6b)]. If q is chosen as a constant, a fixed fraction of the power intended for a given $H_k^{\mathcal{F}}$ is actually intercepted and properly directed. However, as explained in Section V, q must sometimes grow as n increases in order to maintain a constant overall signal-to-noise ratio (SNR). Thus in general

$$q \equiv q(n). \quad (6)$$

C. Fourier Plane

The K subholograms $H_k^{\mathcal{F}}$ can be placed arbitrarily in \mathcal{F} , because they need not align with other physical components. Let the *Fourier-plane diameter* $w_{\mathcal{F}}$ be the diameter of the smallest disk centered at the origin of \mathcal{F} that covers all the $H_k^{\mathcal{F}}$. If circles are packed hexagonally, as suggested in Fig. 3, then K circles of diameter h fit inside a circle of diameter

$$w_{\mathcal{F}} = h \left\lceil \sqrt{(4K-1)/3} \right\rceil \quad (7)$$

where $\lceil x \rceil$ denotes the smallest integer not less than a real number x .

D. Lenses

L_1 and L_2 are ideal Fourier transforming lenses of focal length f . In order to capture all light from the source plane and avoid *vignetting* [34], the diameter D_1 of L_1 must at least equal the sum of the object plane and Fourier plane diameters

$$D_1 \geq w_{\mathcal{F}} + w_{\mathcal{O}}. \quad (8)$$

Likewise, the diameter of L_2 must at least equal the sum of $w_{\mathcal{F}}$ and $w_{\mathcal{I}}$

$$D_2 \geq w_{\mathcal{F}} + w_{\mathcal{I}}. \quad (9)$$

For simplicity, both lenses are assigned diameter D , where

$$D_1 \equiv D_2 \equiv D := w_{\mathcal{F}} + \max(w_{\mathcal{I}}, w_{\mathcal{O}}). \quad (10)$$

Finally, the f -number F is declared to be a fixed parameter

$$\frac{f}{D} \equiv F. \quad (11)$$

The f -number is linked closely to the cost and difficulty of making a lens and also relates to the minimum fringe spacing required in the holographic elements. Fixing F permits evaluation of system growth at a fixed level of technological effort.

The above definitions of abstract sources, holograms, lenses and receptors, their attributes and capabilities, and their admissible configurations, collectively constitute the *canonical representation* or *canonical model* of all optical systems in this paper. This idealized abstraction is amenable to analysis of its volume complexity as well as to the establishment of correspondences between its constituent parts and real optical components. It is therefore a vehicle through which optically interconnected systems may easily be described.

IV. PHYSICAL EMBODIMENTS OF THE CANONICAL MODEL

The canonical model is meaningful to the extent that the demands on its idealized constituents can be met by the physical sources, receptors, beam-deviating elements, and focusing elements that comprise real systems. Details of realizing these components are treated in this section. Investigation of physical elements and their nonidealities establishes the conditions under which conclusions drawn from the canonical model have practical meaning.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.