DISSERTATION

INTERFACES BETWEEN PARALLEL OPTICAL MEMORIES AND ELECTRONIC COMPUTERS

Submitted by

Leonidas Ioannis Irakliotis

Department of Electrical and Computer Engineering

In partial fulfillment of the requirements

for the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2000

UMI Number: 9981338

COLORADO STATE UNIVERSITY

29 February 2000

WE HEREBY RECOMMEND THAT THE DISSERTATION PREPARED UNDER OUR SUPERVISION BY LEONIDAS IOANNIS IRAKLIOTIS ENTITLED INTERFACES BETWEEN PARALLEL OPTICAL MEMORIES AND ELECTRONIC COMPUTERS BE ACCEPTED AS FULFILLING IN PART REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY.

Committee on Graduate Work

Adviser

Department Head

ABSTRACT OF DISSERTATION

INTERFACES BETWEEN PARALLEL OPTICAL MEMORIES AND ELECTRONIC COMPUTERS

Three-dimensional optical memories are considered in the context of information processing architecture with the objective of identifying their suitability for general and special-purpose computation. Their suitability is better delineated by the memories' limitations rather than their abilities.

Optical memories have the potential to offer inexpensive and fast access to large sets of data. Furthermore, they are immune to electromagnetic interference and they exhibit fault-tolerance that is unusual in other types of storage including magnetic disks and capacitive solid-state devices.

Data recording in optical memories is a rather slow process. This fact renders optical memories more suitable for archival storage application, where large storage volumes are required and high data densities are desirable. The question for a suitable optoelectronic interface emerges next.

An optoelectronic interface for three-dimensional optical memories should meet certain design specifications. These specifications are derived from many considerations including data error minimization, performance capacity, hardware complexity, etc.

Data error minimization, specifically reduction of intersymbol interference at the interface regime led to some geometrical design specifications. These specifications are consistent with regular patterns that reduce hardware complexity and enable system scalability.

Performance capacity, in the context of available electronic computer technology, has also led to design specifications regarding the size of data pages and timing of interface's functions. Again, these specifications are scalable to the performance characteristics of modern electronic computers.

Finally, architectural and functional considerations led to the investigation of algorithms for improved performance in data processing. We demonstrated a functioning auto-associative processor and we proposed a fast class of algorithms for data processing in bit-slice mode.

Our findings support the notion that optical memories do have the potential for becoming a leading technology in archiving applications. It may also be possible that with advances in recording techniques, optical memories may one day be suitable for secondary storage, random access, and cache applications. When this occurs, optical memories will contribute to the flattening of the conventional memory hierarchy generally accepted today. But in their current form, optical memories matched with a suitable interface may be ideal architectures for applications involving associative access to data.

Leonidas Irakliotis
Department of Electrical and Computer Engineering
Colorado State University
Fort Collins, CO 80523
Spring 2000

Acknowledgements

My late mother Urania, my father Ioannis, and my brother Aristidis made my graduate education possible through their support and faith in me.

It was the support, friendship, guidance, patience and understanding of my mentor, Professor Pericles Mitkas that sustained me during the years in graduate school at Colorado State. Friendship and support was offered generously by Pericles' wife Sofia Mardiri and their children Alexander-Akyllas and Danae-Zoe.

Professor Carl W. Wilmsen, a co-adviser for all practical purposes, was a source of inspiration and a resourceful guide at many instances. Professor Anura P. Jayasumana was a valuable consultant and a skillful teacher during my stay at Colorado State University. Professor Yashwant Malaiya brought his Computer Science expertise in my committee and provided valuable information at frequent occasions.

I would like to acknowledge the support of the following individuals: Tom Aurand, Sunny Bains, Dr. George A. Betzos, Dr. Fred R. Beyette Jr., Dr. Steward A. Feld, Terry Garrett, Eleni Fytoka, Haralampos Hadjimichael, Eric Hayes, John Hutton, Dimitris Kazis, Eva Konstantinidou, Joel W. Lurkins, Keith Richling, Carol Sarantos, Maureen Schaffer, Cathy Smith, Patrick Stanko, Carol Story, June Thompson, Georgia Triantafyllidou, and Deborah Boucoyannis.

I would also like to acknowledge the friendship and support of Professor Paul D. Scholten at the Department of Physics, Miami University.

The friendship and hospitality of Professor John Sofos and his family during my stay in Fort Collins, Colorado has been a treasured haven.

During my stay at Colorado State I was fortunate to meet Professor Evan Vlachos who has been a friend and a mentor ever since.

Parts of the work reported here were completed under the support of the National Science Foundation, the NSF Engineering Research Center for Optoelectronic Computing Systems Center, and the Colorado Advanced Technology Institute. The support of these agencies is hereby acknowledged.

To the loving memory
of my mother
Urania
(1941–1995)

«ΑΛΛΑ ΠΡΩΤΑ θὰ δεῖς τὴν ἐρημιὰ καὶ θὰ τῆς δώσεις τὸ δικό σου νόημα, εἶπε
Πρὶν ἀπὸ τὴν καρδιά σου θά 'ναι αὐτὴ
καὶ μετὰ πάλι αὐτὴ θ' ἀκολουθήσει
Τοῦτο μόνο νὰ ξέρεις
Ὅ,τι σώσεις μὲς στὴν ἀστραπὴ
καθαρὸ στὸν αἰώνα θὰ διαρκέσει»
(Ὀδυσσέας Ἐλύτης, Ἄξιον Ἐστι.)

BUT FIRST you will see the wildnerness and give it your own meaning," he said.
The wildnerness will precede your heart
and then again the wildnerness will precede your heart
Know this above all:
whatever you salvage under lightning
will remain pure eternally"
(Odysseus Elytis, Axion Esti, Transl. Keely/Savides).

List of Symbols

α, α_0, absorption coefficients

γ, homogeneous width of spectral transition

$\Delta\alpha$, spectral hole depth

δ, displacement (misalignment) vector

ϵ_0, dielectric constant in vacuum

λ, wavelength

ν, photon frequency

ρ, charge density

σ, selectivity (or standard deviation)

Φ_k, RAM clock phase

$\chi^{(n)}$, n-th order electric susceptibility

ω, molecular frequency

d, data pattern

E, electric field

$\mathbf{E}^{sc}$, spece-charge electric field

f, frequency

h, selection or locator vector

h, Plank's constant

I, beam intensity

k, wave vector

k, wave number

N, *M*, number of pixels

n, index of refraction

P_{event}, propability for *event*

Q, input/output rate for electronic computer

R, R_{eff}, data rates

R_{ij}, molecular transition matrix elements

r, distance vector

r, code rate

r_{ijk}, Pockels tensor

s, search vector

T_r, t_{access}, $t_{retrieval}$, time parameters

TABLE OF CONTENTS

LIST OF FIGURES

These lectures are about what we can and can't do with machines today, and why. I have attempted to deliver them in a spirit that should be recommended to all students embarking on the writing of their Ph.D. thesis: Imagine that you are explaining your ideas to your former smart, but ignorant, self, at the beginning of your studies.

(Richard P. Feynman, in "Feynman lectures on Computation")

1

Introduction

On a daily basis, numerous organizations produce large sets of data as part of their daily operations. Astronomical, meteorological, and reconnaissance satellites orbiting the Earth transmit detailed images and other information. The Chandra X-ray Observatory alone is producing approximately 40–50 GBytes (GB) of imagery data daily. This could lead to an annual accumulation of approximately 15,000 GB of data that need to be stored and retrieved periodically for processing and analysis.

Hospitals produce tremendous amount of data on a daily basis as well. Computed Axial Tomography (CAT) and Magnetic Resonance Imaging (MRI) scans result in large image files that need to be stored and retrieved frequently. Financial institutions process transactions and store the related information for long periods of time. The internet infrastructure carries between 10–100 Gbits of data per second. This number alone is staggering. On an annual basis the internet hauls approximately 180 PB of data (1 Peta Byte = 2^{20} GB).

It is necessary to store and archive all these data for a number of reasons. First, because it may not always be possible to process data in real-time. Second, because further processing may be required in the future. Third, because the law may require that certain data be maintained for specific periods of time (e.g., medical records, financial transactions etc). High-density storage techniques are required for storage and archival of large sets of data.

Optical memories were introduced in the early 1970s as a competitive candidate to capacitive storage technologies. The interest in optical storage gained strength in the late 1980s. In the 1990s amble funding for research in viable optical storage systems led to proof-of-concept systems and demonstrated prototypes. These early systems supported the notion that optical memories are indeed a viable technology. The development of these systems presented the research community with numerous challenges about the enabling technologies and the architectural scope of optical memories.

Broadly speaking optical memories rely on three enabling technologies: optics, optoelectronics, and the ubiquitous electronics.

Optics, in this context, include storage material, light guiding tools (e.g., lenses), light sources, etc. Storage materials come in various shapes but are rather small in size. Therefore light must be guided precisely in the active areas and also be manipulated carefully as it emerges from the storage material carrying data. To achieve the required precision the imaging subsystem must be reliable, accurately controlled, and of the highest quality. The optical sources must be stable and should not overload the system with excessive power.

Optoelectronics are at the heart of interfacing the optical memory with the electronic world. These subsystems should meet the power, resolution, and device density requirements that optical memories dictate.

Electronics include standard interfaces to computers and should be readily available as well. Underlying these technologies, there is a set of *software* requirements that enables error detection and correction techniques, data formats, etc. This is the *firmware* of optical storage systems and its components are distributed among the three enabling technologies mentioned above.

Defining the architectural scope of optical memories is an important topic that has been addressed in the recent years [Chiarulli, 96a; Ozaktas 97a, 97b]. Optical storage is already an integral part of today's computing environments. Compact disks (CD's) and related technologies including recordable and rewritable CD's (CD–R's and CD–RW's) are readily available to consumers at attractive prices. These successful representatives of optical memory systems produce linear (one-dimensional) output, i.e., bits are retrieved one at a time. They combine technology that is being employed by conventional magnetic disk drives for a long time and the advantage of high-density recording offered by lasers to become commodity equipment. However, we are interested in optical memories that produce two-dimensional output, specifically pages of data that are retrieved (or recorded) one at a time. The advantage of this characteristic is clearly its parallelism and suitability for numerous applications. For example optical memories with two-dimensional formats can be used to store and recall images and video data. A video frame, complete with its accompanied audio data and indexing information, can be stored in a single page of sufficient capacity. Throughout this dissertation we use the term **optical memories** as synonymous to **optical memories with two-dimensional output.**

Table 1.1 presents some of the advantages and disadvantages of using optical memories for different levels of the memory hierarchy. The capability for associative processing and the potential for high-density and high-capacity recording make optical memories extremely appealing for application across the entire memory hierarchy. At the same time optical memories are limited by the abilities of their enabling technologies. These limitations are discussed next.

Table 1.1. Advantages and disadvantages of optical memories.

Hierarchy Level	Advantages	Disadvantages
Cache	Associative processing	Slow access time Bulky
RAM	Associative processing High density	Slow recording time
Storage	Associative processing High density	Slow recording time
Archiving	High density	Poor data persistence

While optical memories offer the advantage for full and parallel associative retrieval that makes them ideal for cache application, they are falling short on access time and real estate expectations for this particular niche. The promise for high-density storage makes optical memories a suitable competitor to magnetic hard disks. But while data retrieval can be performed fast enough, it is the recording of data that is rather slow by comparison to magnetic media. Because of this characteristic optical memories are better suited for archival applications and services including high-performance delivery of multimedia content.

Optical memories are also suitable for archiving because of their potential for high-volume storage. However most material do not allow for persistent recording and data will degrade faster when compared to conventional archiving media such as data tapes and high-capacity movable magnetic disks (e.g. Jazz drives). While all the shortfalls of optical memories mentioned here are the subject of intense research aiming into improving the performance in those areas, there are a number of successful systems that employ optical storage techniques and clearly demonstrate the promise of the technology. Compact disks (CD's) and their derivatives are consumer products whose performance and cost characteristics fit into the model of economies of scale.

1.1 Statement of the Problem

The potential for the integration of optical memories in the data processing fabric creates a challenging problem that could be stated as following. Optical memories rely on light to record and retrieve data and because of that are formidable elements in parallel architectures. Furthermore they are tolerant to channel cross-talk and electromagnetic interference-induced noise and capable of high data densities, thus superior to capacitive memories that employ electric charges or magnetism to store data.

Because of their advantages optical memories are a desired component in any computer system. However we need to understand better not only their advantages but also their limitations in order to determine their architectural scope and build appropriate systems by using the most compatible resources and components that the enabling technologies have to offer.

1.2 Objective of the Dissertation

The main objective of the present study is to identify advantages of optical memories not by looking directly at their characteristics that are superior to other forms of data storage but by delineating their limitations. The investigation is coupled with the study of the interface between optical memories and electronic computers. This becomes evident if we rephrase the objective of our study as following:

> *Given the optical memories that have been developed (or proposed) already and the existing architectures that dominate the computer market, what kind of interface between these two technologies is the most likely to turn optical storage into a viable peripheral?*

Ultimately we will determine if there are feasible interface architectures that will allow computer users to integrate optical memories into their systems and benefit from their numerous advantages without experiencing the impact of any limitations that are characteristic of optical storage.

To propose a good — if not the best — interface requires the identification of what optical memories do best. Without a doubt their strength lies in their ability to store two-dimensional arrays of data in a single address. This advantage promises very fast access to large volumes of data thus making optical memories suitable for associative architectures and for parallel computing. It becomes clear that our investigation should focus on applications that can be enhanced by associative processing and parallelism and propose interfaces to implement the desired enhancements.

1.3 Methodology and Outline of the Dissertation

Chapter 2 reviews the physical characteristics of optical memories and examines the recording and retrieval mechanisms for a number of *principal* methods of optical data storage. Here *principal* refers to the set of physical phenomena that characterize each of the methods examined here. Specifically we direct our attention to photochromic, photorefractive, and spectral hole-burning methods. The review of these methods demonstrates optical memories' ability for high-density data storage and helps the reader understand that optical memories can yield data fast (thus overcoming any access time related limitations) if the appropriate technology is used.

Chapter 3 proposes how optical memories should be evaluated and compared to competing technologies. Criteria and performance metrics are derived from existing rules used to evaluate magnetic disks and solid-state capacitive storage devices. The criteria proposed here can be used to determine an optical memory's placement in the memory hierarchy (shown in Table 1.1) and its degree of suitability for a particular class of applications. The latter also suggests the nature of interface that is more appropriate for the particular memory.

Chapter 4 presents a modelling study on the geometric characteristics of optical memories. We are interested in establishing acceptable tolerance levels of cross-talk noise at the interface stage. Although optical memories have very small cross-talk noise levels problems may arise when their output is directed on the interface. Imaging systems could easily distort the output and create some channel and or symbol interference. Errors due to interference can be corrected if detected. An analysis of the fundamental spatial arrangement of data could lead to a better understanding of these errors' nature.

Chapter 5 presents the interfacing concepts that are the focus of this dissertation. The performance analysis included in our discussion of the proposed interfaces suggests that these proof-of-concept designs have the potential to grow into realistic computing (sub)systems that can enhance specific classes of applications. In Chapter 6 we extend our discussion on the interfacing concepts to analyze our findings and proposed suitable inferface architectures.

This dissertation concludes with Chapter 7 where we summarize our findings and we present a critical review of them.

In my beginning is my end. In succession
Houses rise and fall, crumble, are extended,
Are removed, destroyed, restored, or
in their place
Is an open field, or a factory, or a by-pass.
(T. S. Eliot: Four Quartets)

2

Optical Memories

The concept of using light to store information certainly is not a new idea. If we associate the optical data storage with visual representation of information we may argue that painting and photography are techniques for or applications of optical data storage. By extending this analogy we can further assume that engraving, sculpture, writing, cinematography and typography are all optical data storage mechanisms because light is required to reproduce the information stored. To a further and debatable extend navigating a ship 200 years ago required the application of optical storage, because at night the navigator had to interpret the relative locations of stars on the sky in order to determine the location of his vessel. Although one may argue that these examples are extreme and open to debate, the importance of using light for storing information is well demonstrated and established.

Regardless of the level of abstraction and information encoding, humans rely heavily on optical information storage and processing in each and every aspect of life. This should not come as a surprise, if you consider the complexity of the human eye and of the visual neural structure of our brains. The eye, the optical nerve, and the visual cortex of the mammals' brain are some of the most elaborate and complex information processing structures we have ever dealt with.

The field of optical storage began to develop in its present form in the 1960s with the advent of coherent light sources. Before that, optical storage was limited to permanent recording of images such as photography, where the light's intensity was used as the information carrier. The introduction of lasers and solid state lasers in 1960's [Bromberg88a] and of the microlaser arrays in 1980s were the landmarks in the evolution of optical storage.

Optical data storage depends on light-matter interactions. In holographic storage with photorefractive media, for example, light-matter interactions photo-excite carriers to drift to dark regions of the material and recombine there, thus producing a space-charge electric field distribution which modulates the index of refraction via the linear electrooptic effect (Pockels effect). Two-photon three-dimensional storage, light-matter interactions force a molecule to undergo a photochromic transition which is stable and reversible [Hunter, 90a; Dvornikov, 94a]. Other storage techniques rely on resonance heating of microvolumes in thin films or crystals, modulating populations between ground and excited states, etc.

Optical data storage can be viewed as the technological progression from magnetic and electrical/electronic storage. However it is not just technology evolution which leads to optical storage. We should take also into account the increasing demand for storage capacity. Although the vast majority of information is still stored in paper (which, to some extend, is considered optical storage) the transfer rate from printed paper to digital storage is increasing, and the production rate of digital stored information increases as fast.

Data storage owes its existence to the invention of digital electronic computers. In the beginning there was a need to store limited amount of information, mostly source code for programs and initial input data. Magnetic based memories such as cores were used extensively. Archiving was based on printed paper, punch cards and paper tape, and later on magnetic drums, disks, and tapes. In the 1970s, newly introduced semiconductor memories took over magnetic core memories. At approximately the same time frame, massive storage on magnetic disks improved dramatically but, until recently, failed to follow the increasing performance of solid state memories and CPUs. As a result, most of today's computer applications are input/output bounded and the need for large and fast memories becomes pressing.

The cost per megabyte of storage seems to be following Moore's law. In fact we can argue now that the cost of megabyte on hard disks is upper bounded by Moore's law. Moore's law is a good criterion to separate between commodity and special components in a computer architecture. Apparently, memory has become a commodity item since it follows (bounded by) Moore's law. An example of special component is high quality displays. Despite the technology evolution in glass cathode ray tubes (CRT) their prices remain high and nearly flat over time.

Optical storage appears to be a promising solution for several reasons:

(a) it exhibits a high immunity to electromagnetic interference and it is not affected by external magnetic fields in most cases;

(b) the unit area (or volume) per bit of an optically stored bit is extremely small, in the order of λ^2 (or λ^3), where λ is the wavelength of the light source used for recording;

(c) light beams can intersect each other in free space without interference;

(d) **light propagates in parallel, when collimated, and a high density parallel multi-beam channel can be built this way.**

Soon after the laser was invented, researchers discovered different ways of using this light source to store and retrieve information. Over the years several mechanisms were suggested and designed, and some were implemented. Perhaps the most successful is the optical disk either in the magnetooptical format or the write-once read-multiple times (WORM) format. Optical disks are not part of this dissertation because they constitute an individual topic for study. In this review we focus mostly on thick optical storage media where data can be stored throughout the volume of a material.

This section is a brief discussion on the physical principles of various methods used for optical storage. Some of the methods presented in this review have their roots in the beginning of this century. However optical storage evolved into a coherent concept in the 1960s when the lasers were made available to the scientific and engineering community [for example: Chen, 67a; Glass, 78a]. Laser sources played a key role in the development of optical storage methods. A good and brief historic review on the birth of lasers can be found in [Bromberg, 88a].

The advent of laser sources enabled studies of non-linear interaction of light with matter. In classical electrodynamics the polarization vector **P** of a material can be expressed as:

$$\mathbf{P} = \epsilon_0\chi^{(1)}\mathbf{E} + \epsilon_0\chi^{(2)}\mathbf{EE} + \epsilon_0\chi^{(3)}\mathbf{EEE} + \ldots \tag{2.1}$$

where ϵ_0 is the dielectric constant of vacuum, $\chi^{(i)}$ are the electric susceptibilities of the material, and **E** is the electric field component of the incident light beam [Jackson 75a; Wangsness, 86a]. For weak electric fields, the higher order terms of Eq. (2.1) are negligible and the polarization vector can be approximated by:

$$\mathbf{P} = \epsilon_0\chi^{(1)}\mathbf{E} \tag{2.2}$$

Here, $\chi^{(1)}$ is the electric susceptibility of the material. Eq. (2.2) is valid for relatively low intensities of incident light. For example, sunlight's electric field is only a few V/cm. Laser beams however have electric field components in the order of 10^6 V/cm and the high-order terms of Eq. (2.1) cannot be neglected.

These nonlinearities give rise to phenomena that are being used to store information in materials. Of these phenomena, the photorefractive effect, the two-photon process, and the spectral hole-burning are the most interesting ones with potential applications to optical storage.

2.1 Recording Media

The non-linear optical mechanisms that we discussed in the previous section are used in certain media to modulate a material's index of refraction or absorption in a reversible way. Optical recording is thus possible.

2.1.1 Photochromic Media

Photochromic media exploit the photochromic effect, i.e., the change in the molecular structure and color of a material after it absorbs light of certain frequency. The apparent change in color is because of the absorption profile of the new molecular structure [Dvornikov, 94a]. Some materials undergo the change when subjected to a two photon process. Two-photon absorption studies date back to the early-1930s [Göppert, 31a]. In principle, the first photon excites the molecule to a virtual state. While the molecule is still at that virtual state, the second molecule will excite it from there to a real excited state (for example, its first excited state). If A and B are the ground and excited states of the molecule respectively, then the two photon process can be described by:

$$E_A + h\nu_1 \rightarrow E_{\text{virtual}};\ E_{\text{virtual}} + h\nu_2 \rightarrow E_B \tag{2.3}$$

where ν_1, ν_2 are the two photon frequencies and they can be interchanged.

It should be noted here that neither of the two photons can excite the molecule to a real state alone. Also if the time difference between the two photons is long enough, the molecule will revert from the virtual state to the ground state before the second photon arrives. Then the second photon will reexcite the molecule to another or the same virtual state and the molecule will again relax to its ground state. Therefore it is imperative that the following conditions be met:

(a) the two photons must be guided precisely at the location of the molecule, and

(b) **the time difference between the two photons arrivals must be less than the relaxation time of the molecule from any of its virtual states between its ground and excited states.**

The probability of a two-photon driven transition to occur is given by:

$$P_{i\to f} = P_{sp} \times P_{tp} \times P_{12} \tag{2.4}$$

where i denotes the initial state (which in most cases is the ground state) and f is the final state which is an excited state of the molecule. Here P_{sp} is the spectral profile of a two photon transition:

$$P_{sp} = \frac{\gamma_{i\to f}}{\left[\omega_{i\to f} - \omega_1 - \omega_2 - v \cdot (k_1 + k_2)\right]^2 + \left(\frac{\gamma_{i\to f}}{2}\right)^2} \tag{2.5}$$

This is the profile of a single photon transition at frequency $\omega_1 + \omega_2 + v \cdot (k_1 + k_2)$. ω_1 and ω_2 are the frequencies of the two photons involved in the process. The homogeneous width of the transition is $\gamma_{i\to f}$ The second factor in Eq. (2.4) is the transition probability for the two photon process:

$$P_{tp} = \left|\sum_j \left(\frac{R_{ij} \cdot e_1 \cdot R_{jf} \cdot e_2}{\omega_{ji} - \omega_1 - k_1 \cdot v} + \frac{R_{ij} \cdot e_2 \cdot R_{jf} \cdot e_1}{\omega_{ji} - \omega_2 - k_2 \cdot v}\right)\right|^2 \tag{2.6}$$

The elements R_{ij} and R_{jf} are the elements of the transition matrix between the initial (i) and final (f) states through and intermediate molecular level (j).

Finally, the third term in Eq. (2.4) is the product of the intensities of the two beams:

$$P_{12} = I_1 I_2 \tag{2.7}$$

There are several ways to increase the probability $P_{i\to f}$ The most obvious one is to maximize the P_{12} probability by using very short laser pulses. Another way is to use the first photon ω_1 to excite the molecule to a virtual state very close to a real state and from there use the second photon ω_2 to reach the real excited state. In this case the term P_{tp} is greatly increased.

Let us use the spiropyran molecule as an example. This is one of the most promising candidates for optical storage based on two-photon processes on photochromic material [Dvornikov, 94a].

If the spiropyran molecule in Figure 2.1(a) is illuminated with light at 532 nm, a two-photon process will be initiated because the energy of the light is equal to the gap between the ground and the first excited state of the molecule. The excitation of the molecule will result in the opening of thecarbon-oxygen bond at the pyran ring. This cleavage of the oxygen bond will lead to an isomerization of the molecule to that of Figure 2.1(b). This isomer is called merocyanine. The merocyanine now absorbs light at 560 nm.

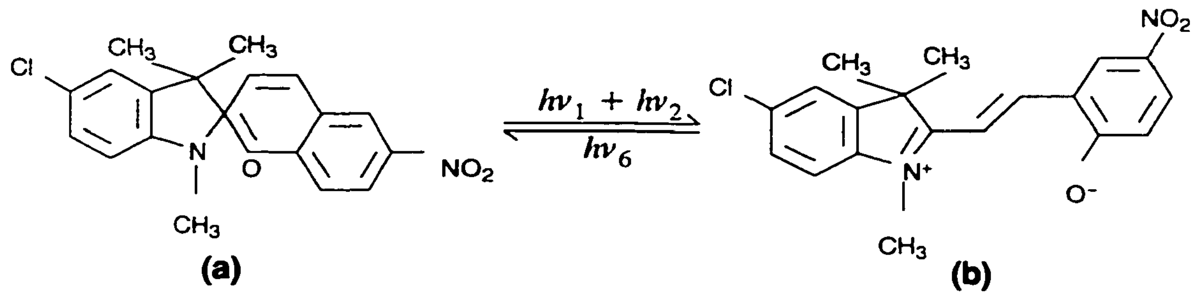

Figure 2.1. Photochromic process in the spiropyran molecule.

The process or writing, reading, and erasing information on a single molecule is outlined in the following equation:

$$\left.\begin{array}{c} A \xrightarrow{h\nu_1 + h\nu_2} A^* \\ A^* \leftrightarrows I \\ I \leftrightarrows B \end{array}\right\} \text{write}$$
$$\left.\begin{array}{c} B \xrightarrow{h\nu_3 + h\nu_4} B^* \\ B^* \rightarrow B + h\nu_5 \end{array}\right\} \text{read} \qquad (2.8)$$
$$\left. B \xrightarrow[Q]{h\nu_6} A \right\} \text{erasing}$$

In Eq. (2.8) A is the ground state of the molecule, A^* is the electronic excited state, I is an intermediate state with a short lifetime, B is the ground state of the merocyanine, and B^* is the excited state of the merocyanine; $h\nu_1$ and $h\nu_2$ are the *write* energies, $h\nu_3$ and $h\nu_4$ are the energies which correspond to the *read* probe, $h\nu_5$ is the luminescence detected as a result of the read probe, if there is a bit recorded, and $h\nu_6$ is the *erasing* energy which in some cases could be the same as the write energy; finally Q denotes heat.

Now let us try to explain Eq. (2.8) with the help of Figure 2.2. The write cycle is shown in Figure 2.2(a). The energy $2h\nu_1$ from Eq. (2.8) here corresponds to the transition caused by the two photons at 1064 nm and 532 nm. Alternatively the molecule can be driven to its second excited state by two photons at 532 nm each.

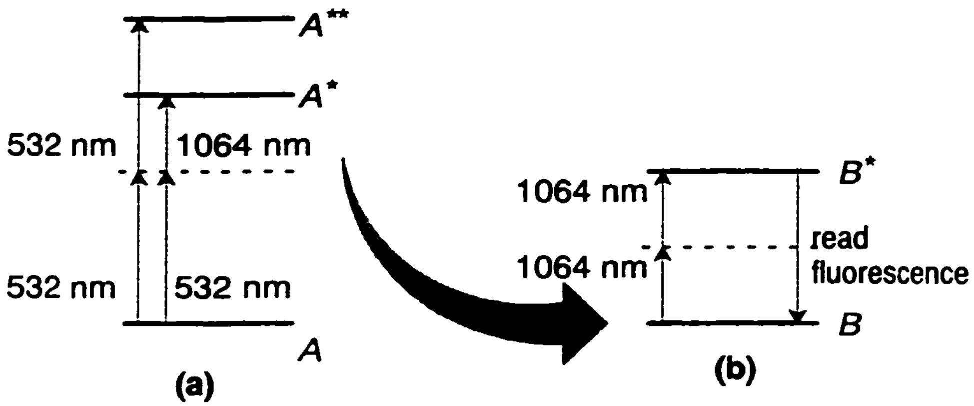

Figure 2.2. Energy transitions in the photochromic process. (a) write cycle; (b) read cycle. The dashed lines mark virtual states during the two-photon excitation.

The photoexcitation of the molecule from A to A^* causes a structural transition as we mentioned earlier. Once the molecule enters its isomer phase, it relaxes at the ground state B of its new structure. We assign the binary value '1' to the isomer structure. Therefore when a value of '0' is desired we simply do not expose the molecule to the write beams.

During the read cycle, the molecule is probed by two photons at 1064 nm. If a value of '0' is present, i.e. if the molecule is at state A, then the two photons at 1064 nm are not enough to cause any change. If a value of '1' has been recorded however, the two probing photons will drive the molecule from B to B^* which has a lifetime in the order of a few pico-seconds. From there, the molecule will relax to its ground state by emitting a detectable fluorescence.

We can selectively erase information by forcing the recording of a ‘0’ on a open (isomer) molecule. This is achieved by providing enough energy ($h\nu_4$) to the molecule so that it could reestablish the previously broken carbon-oxygen bond at the pyran ring. If selective erasing is not required, the entire memory, which is an ensemble of such photochromic molecules, can be erased by raising the system’s temperature. It has been found that 100 ps is enough time to write information on the write-form of the spiropyran molecule [Dvornikov, 94a].

In order for a material such as the spiropyran to be useful in optical storage the following requirements must be met:

(a) when a molecule is probed to determine its phase, the emitted fluorescence —if any— should be detectable, and

(b) (the written form of the molecule must be stable for long periods of time at room temperature.

To achieve a detectable level of fluorescence emission from spiropyran molecules the quantum efficiency of the process must be high enough. It has been reported that the efficiency ranges from 10^{-3} to 10^{-2} [Dvornikov, 94a].

2.1.2 Photorefractive media

Much as photochromic materials change their transmissivity because of the molecular structure transition that characterizes them, photorefractive materials change their transmissivity because of an induced modulation of the medium’s index of refraction [Hesselink, 93a]. A steady-state theory describing the phenomenon was developed in the late 1970s [Kukhtarev, 79a, 79b]. This modulation is described by the following equation:

$$\Delta\left(\frac{1}{n^2}\right)_{ij} = r_{ijk}E_k^{sc} \qquad (2.9)$$

Equation (2.9) describes the linear (or first order) electro-optic effect where the index of refraction n is modulated by a space-charge electric field $\mathbf{E}^{sc}$ according to the components of the Pockels’ tensor r_{ijk}. Although the effect itself in nonlinear because it corresponds to the second term of Eq. (2.1), it is called "linear" in the context of the mathematic formulation for electro-optic phenomena. [Yeh, 93a, pg. 21]. The linear electro-optic effect simply refers to the modulation of the index of refraction in the presence of an applied electric field. In the linear regime, where light intensities are relatively low and there is no manifestation of the higher order terms of Eq. (2.1), the index of refraction n is considered a constant for a given material:

$$n^2 = 1 + \chi^{(1)} \tag{2.10}$$

When a photorefractive material is exposed to a high intensity light beam, like the one produced by a laser source, then the higher order terms of Eq. (2.1) are not negligible anymore.

Figure 2.3 illustrates the formation of the space-charge field. Figure 2.3(a) shows the state diagram of an electro-optic crystal with impurities when the crystal is at an electric equilibrium. Electronic states in the valence band are filled. The impurities are responsible for the presence of trap states in the bandgap. These states are empty when the crystal is at an electric equilibrium. If we use two laser beams to illuminate the crystal, their electric field components will form the following interference pattern:

$$\mathbf{E} = \mathbf{E}_A \exp(i\omega t - i\mathbf{k}_A \cdot \mathbf{r}) + \mathbf{E}_B \exp(i\omega t - i\mathbf{k}_B \cdot \mathbf{r}) \tag{2.11}$$

where $\mathbf{E}_A$ and $\mathbf{E}_B$ are the electric fields of the two beams and $\mathbf{k}_A$, $\mathbf{k}_B$ are their respective wave vectors. This field is a direct result of the interference pattern shown in Figure 2.3(b):

$$I(\mathbf{r}) = I_0 + \mathrm{Re}\{I_1 \exp(-i\mathbf{K} \cdot \mathbf{r})\} \tag{2.12}$$

The dc-component of the intensity pattern is $I_0 = |\mathbf{E}_A|^2 + |\mathbf{E}_B|^2$ and the amplitude of the ac-component is $I_1 = 2\mathbf{E}_A\mathbf{E}_B^*$. The wave vector of the interference pattern is $\mathbf{K} = \mathbf{k}_A - \mathbf{k}_B$.

This pattern will produce a space-charge distribution in the crystal because at the points where Eq. (2.12) reaches a maximum value, electrons in the valence band will be photoexcited and move to the conduction band. They will drift for a while there and after they lose their energy they will be trapped in the bandgap, thus creating a space charge distribution with a density:

$$\rho = \rho_0 \cos(\mathbf{K} \cdot \mathbf{r}) \tag{2.13}$$

Since $\nabla \cdot \epsilon \mathbf{E} = \rho(\mathbf{r})$, by integration we get the space charge field:

$$\mathbf{E}^{sc} = \rho_0 \frac{\mathbf{K}}{\mathbf{K} \cdot \epsilon \mathbf{K}} \sin(\mathbf{K} \cdot \mathbf{r}) \tag{2.14}$$

which modulates the index of refraction according to Eq. (2.9)

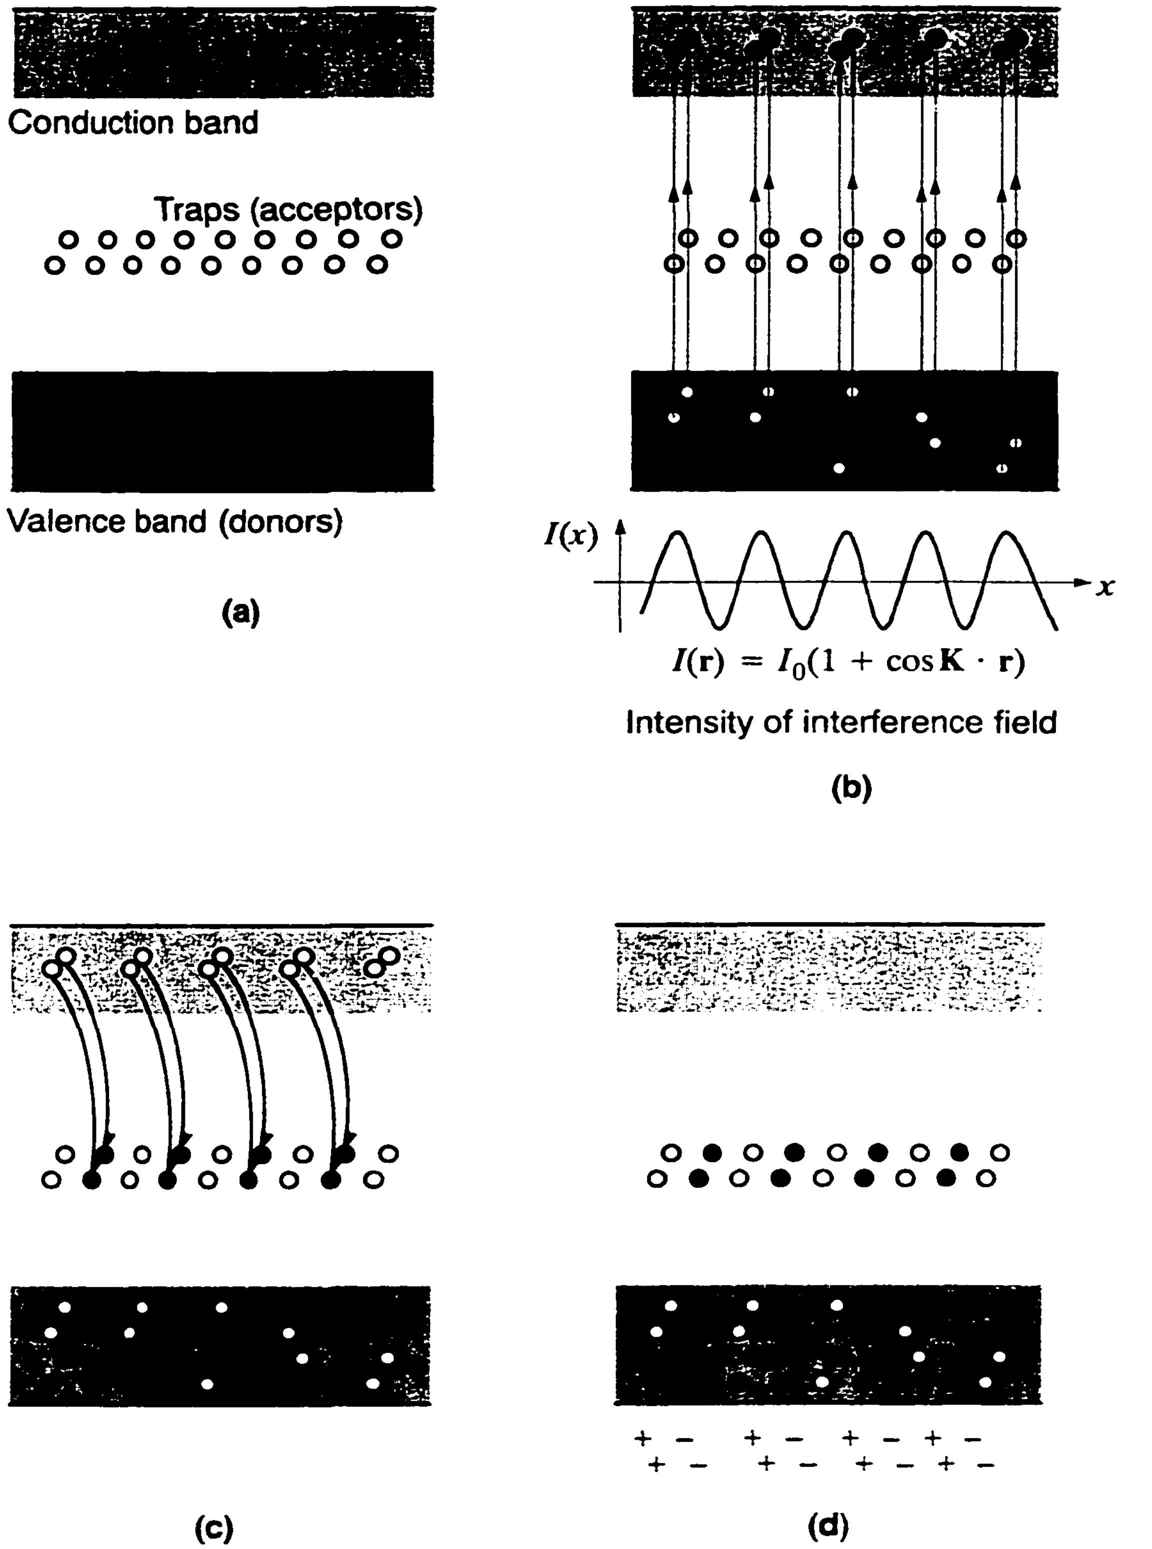

Figure 2.3. Band transport model for the photorefractive effect. (a) band model; (b) photoexcitation of electrons; (c) relaxation of electrons in traps; (d) formation of space-charge distribution.

2.1.3 Spectral hole-burning media

Another way of controlling the optical transmissivity of storage media is by modulating the absorption profile of inhomogeneously broadened materials. These are host crystals with high concentrations of impurity ions in their matrices [Moerner 88a, 88b]. The impurity ions are grouped in small ensembles and each ensemble has a narrow absorption profile located around a center frequency as shown in the inset of Figure 2.4. There are multiple such ensembles in the crystal and their center frequencies are distributed over a broader range as shown in Figure 2.4. The macroscopic effect of this broader distribution is that the crystal appears to exhibit a broader absorption profile. This inhomogeneously broadened absorption profile follows the distribution of the ensembles' profiles.

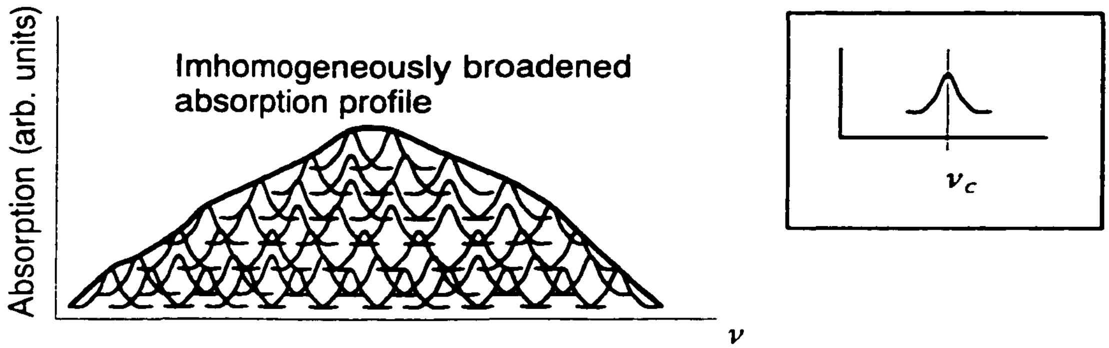

Figure 2.4. Absorption band broadening in an inhomogeneous medium. The inset shows the homogeneous absorption band for one of the constituent systems of the inhomogeneous medium.

Broadening of the absorption band was a fundamental problem in the spectroscopic study of inhomogeneous materials because it was impossible to probe and investigate the finer structure of the system. The invention of laser introduced a way by which spectroscopists were able to probe in selected narrow sets of the broadened band by selecting the laser's frequency to match with the transition energy (i.e. the centered frequency) of the desired subsystem in the medium. Soon it became evident that in certain systems, the selection of a narrow sub-band of the broad band can be made permanent, though the process is reversible, thus creating a spectral hole around the specific frequency. This effect is known as *spectral hole burning*.

Spectral holes were first observed by Bloembergen et al. [Bloembergen, 48a] in their studies of inhomogeneously broadened magnetic resonance absorption lines. The invention of the laser opened the road to the observation of spectral holes in optical transitions which where predicted by Schawlow [Schawlow, 61]. In 1970 Szabo reported that a *fluorescence line narrowing* for centers in inhomogeneously broadened spectra using laser sources [Szabo, 70a]. Fluorescence line narrowing refers to the narrowing of fluorescence excitation in a inhomogeneous medium. Later on, Szabo reported a transient spectral hole burning [Szabo, 75a] and at the same time persistent hole burning was reported [Kharlamov, 74].

Two fundamental mechanisms have been identified as the leading effects which could cause spectral hole burning: photochemical and non photochemical processes [Moerner, 88a]. The photochemical mechanism may cause either bond-breaking, isomerization, tautomerization, ionization etc. in the inhomogeneous system thus causing the frequency center of a subsystem to change. It should be noted here that a subsystem is the guest in a guest-host system such as ions scattered in a crystal. In non photochemical processes, the change in the center frequencies is caused by changes in the environment of the guest systems, i.e. the processes affect mostly the host than the guest. A more complete description of the mechanisms leading to persistent spectral hole burning can be found in [Moerner, 88a].

Spectral hole-burning turned to be a very efficient spectrographic technique for high-resolution spectroscopy. In the late-1980's Meixner *et al.* [Meixner, 89a] and Renn *et al.* [Renn, 90a] produced a series of comprehensive studies investigating the relationship between spectral hole-burning and holography. They successfully argued that holographic detection of spectral holes is free of background noise thus making the detection more efficient than classical methods involving transmission and fluorescence excitation.

In transmission and fluorescence detection schemes the signal obtained is described in general by:

$$I_t = \left\{1 - I_0 \exp\left[-\alpha_0 + \Delta\alpha L'(\omega - \omega_b)\right]d\right\} \tag{2.15}$$

for the transmission excitation, and

$$I_f \propto \left\{1 - I_0 \exp\left[-\alpha_0 d + \Delta\alpha L'(\omega - \omega_b)d\right]\right\} \tag{2.16}$$

in the case of fluorescence excitation. α_0 is the initial absorption coefficient, ω_b is the burning frequency (or in other terms the spectral location of a homogeneous subsystem), $\Delta\alpha$ is the hole's depth and $L'(\omega - \omega_b)$ is the amplitude of the Lorentzian shape function describing the shape of the hole and given by:

$$L'(\omega - \omega_b) = \frac{(\Gamma/2)^2}{(\omega - \omega_b)^2 - (\Gamma/2)^2} \tag{2.17}$$

where Γ is the width of the spectral hole at ω_b.

From Eqs. (2.15) and (2.16) it becomes apparent that both I_t and I_f suffer from background noise which makes the detection of the hole more difficult. In addition to the background noise due to the Lorentzian, classical detection of spectral holes requires extremely stable light sources.

2.2 Systems

In this section we discuss three representative optical memory systems that we consider promising technologies for massive storage. The two-photon three-dimensional memory is based on the two-photon effect and the holographic memory employs the photorefractive effect. Optical disks can employ any non-linear mechanism (or even linear mechanisms) to record and reproduce information and they are not considered to be true volumetric storage systems. However they deserve attention because commercially they are the most successful optical memory systems available.

2.2.1 Two-photon Three-dimensional Memory

Two-photon three-dimensional memory (TPTDM) has been introduced in the late 1980s as a potential system for high density optical storage (Hunter, 90a; Dvornikov 94a). The underlying mechanism is a two-photon process which causes the molecules of a photochromic material to undergo a structural phase transition. This structural transition is reversible and detectable and therefore can be used for information recording.

Data are stored in planes within the crystal, as shown in Figure 2.5. A plane contains a two-dimensional array of volume elements, each corresponding to a bit. A volume element is addressed by two intersecting beams. During the write cycle these beams are called data beam and address beam respectively. Data beams offer the first photon in the two-photon process and they originate from a binary page composer like a spatial light modulator (SLM) and they are focused through an autofocus mechanism (not shown in Figure 2.5) at the selected data plane. At the same time the data plane is illuminated by a multitude of address beams which provide the second photon required for the two-photon process. The two beams intersect in a volume element and cause the photochromic molecules to undergo a molecular structure transition. If there are no data beam present, i.e., the intended bit value corresponds to a zero, then the photochromic material in the volume element does not undergo the molecular transition because the address beam's photons cannot initiate a two-photon process alone.

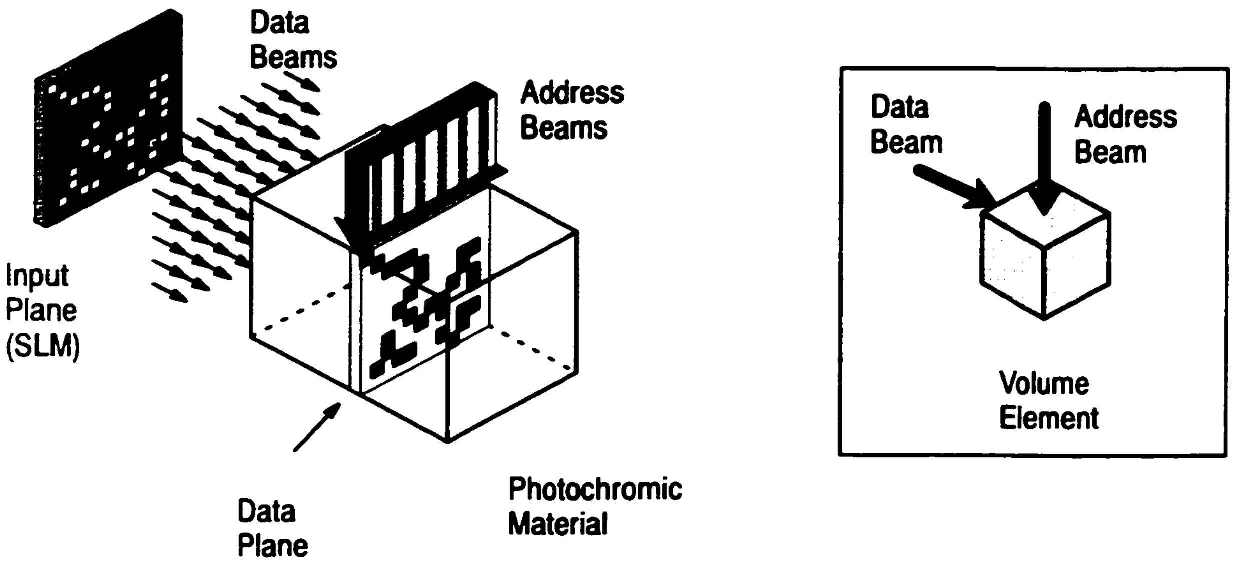

Figure 2.5. Two-photon three-dimensional memory system. The inset at right shows the addressing of independent bits.

When reading a plane, each volume element is intersected by an address beam and a readout beam. All volume elements in the plane receive beams from the SLM, i.e., all the pixels on the SLM are being turned on. If a volume element has undergone a molecular structure transition during the read cycle, photons from the address and read beams will excite it and make it fluoresce. If it has not undergone a transition, then the energy offered by the two intersecting beams is not enough to excite it and therefore no fluorescence will be detected from that plane.

Data stored in photochromic materials via the two-photon process can be maintained for several weeks. Erasure is possible by raising the temperature of the crystal thus providing enough energy to restore the original molecular structures. Selective erasure is also possible by re-writing a selected bit location.

2.2.2 Holographic Memory

Photorefractive volume holographic storage (PVHS), or holographic memory, uses electro-optic crystals such as lithium niobate ($LiNbO_3$), which exhibit the photorefractive effect. This is the spatial modulation of the index of refraction in proportion to the intensity of incident light. If the incident light corresponds to the interference pattern between a reference and an object beam, a hologram is formed and it is stored in the crystal volume. If the object beam is the scattered light from a data page (such as picture, tabulated data etc) then a holographic record is created. Multiple holograms can be stored this way using different reference beams. By illuminating the holographic volume with a reference beam we can reconstruct the corresponding hologram [Kogelnik, 69a]. In reverse, if we illuminate the holographic volume with the image of a recorded object (or part of it) we can reconstruct the corresponding reference beam [Burr, 94a].

The volume data density for PVHS is in the order of λ^{-3} [Yeh, 93a]. In fact, this is the maximum possible volume storage density for all kinds of optical storage machinery that employ three-dimensional geometric space.

We indicated already that multiple holograms can be written in the same volume. Although the physics of multiplexing techniques is quite intriguing, it is beyond the scope of this chapter and relevant information can be found in References [Yeh, 93a; Hesselink, 93a]. A generic system description however is in order. Figure 2.6 shows a basic multiplexing technique known as angular multiplexing. By changing the angle between the reference beam and the crystal, we can record multiple holograms in the crystal. Each hologram is identified by a unique angle and therefore random access is possible by selecting the appropriate incidence angle for a reference beam.

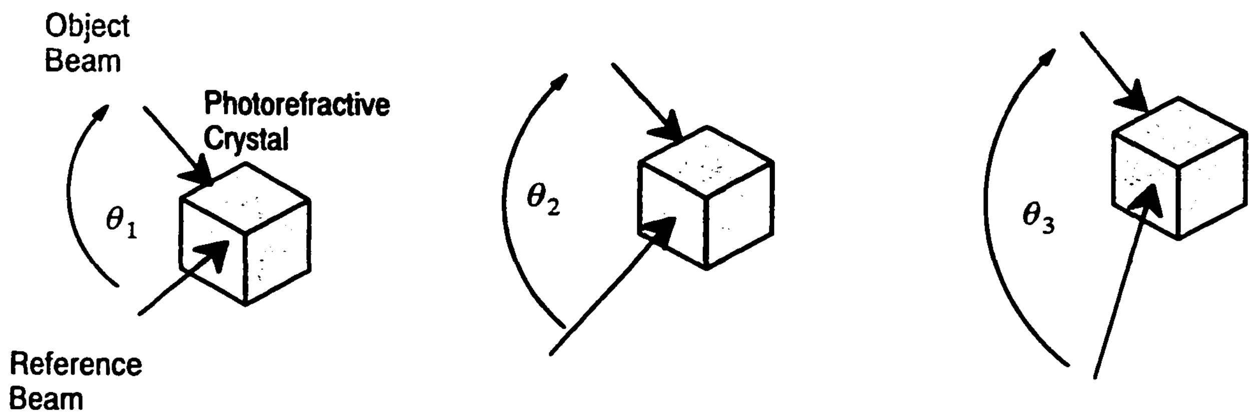

Figure 2.6. Angular multiplexing in photorefractive volume holographic storage.

Holograms have some interesting properties that are very useful when building an associative memory system. A hologram can still reconstruct an image stored in it even if the medium is damaged or broken in several pieces. Assuming that we are dealing with Fourier holograms i.e., with interference patterns between a reference beam and the Fourier transform of the object beam, we spread information about the object all over the hologram. Even when a subset of the original hologram is available, there are probably enough spatial frequency components in the remaining hologram to reconstruct the original image at somewhat lower quality. This fault tolerance of Fourier holograms is bi-directional in the sense that if a partially damaged on incomplete version of an object beam is imaged on the hologram, the corresponding reference beam will be reconstructed at a lower SNR. Since the reference beam is essentially the address of that hologram, one can see the associative nature of holography.

Angular multiplexing has a dynamic range of approximately 500–5,000 holograms per cm^3. Additional capacity is obtained by employing other addressing mechanisms like spatioangular [Burr, 94a] and peristrophic multiplexing [Pu, 95a]. Using spatioangular multiplexing up to 10,000 holograms have been recorded [Burr, 94]. At that time the number of holograms was expected to reach 160,000 holograms in the near future [Burr, 95a, 99a]. Recent reports indicate capacities in the order of 10^9 bits/cm^3 [Shelby, 97a; Psaltis, 98a].

Recording in photorefractive crystals is quite persistent and reversible. Persistence time varies from a few minutes to several weeks depending on the host and the impurities. Thermal fixing of the crystal can also increase the storage time. Data can be erased in bulk if the crystal is illuminated by intense white light which restores the electric field equilibrium in it.

Recently there has been an increasing interest in the use of polymer based materials are recording media in holographic memories. It has been demonstrated [Lessard, 98a; Steckman, 98a; Trout, 98a] that commodity polymers such as polymethylmethacrylate (PMMA) when properly doped can act as an ideal recording medium. Production of these materials is much easier and less expensive compared to the growth of non linear optical crystals. Furthermore, polymer-based media are suitable for read-only memories because they can be permanently fixed.

2.2.3 Optical Disks

Optical disks were the next step in planar storage, following magnetic disks. Data retrieval from optical disks occurs by reflecting a laser beam off the surface of the recording medium which is typically enveloped in hard plastic in order to prevent scratches and other damages [Psaltis, 90a]. The incident beam is modulated by the surface properties at the specific spot and the reflected beam carries that signal to a photodetector which determines the bit value at that spot, as illustrated Figure 2.7.

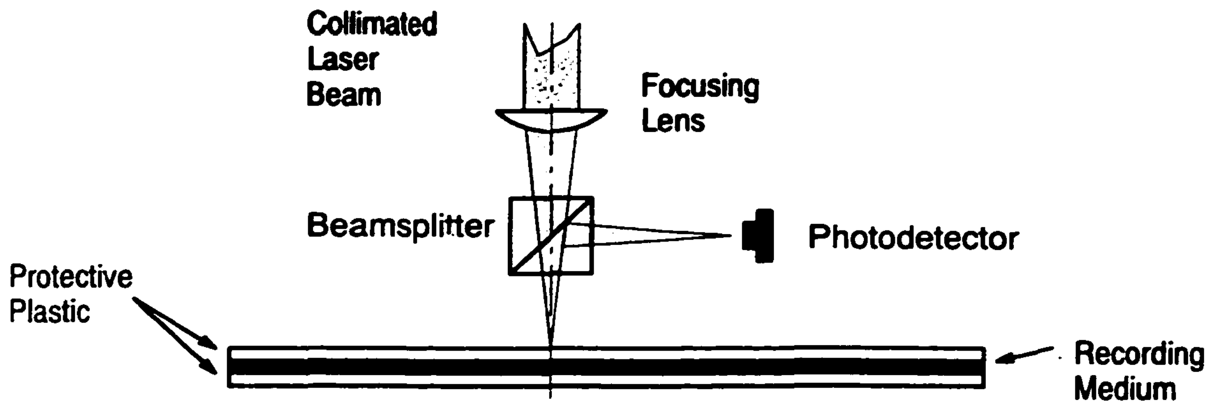

Figure 2.7. Optical readout of an optical disk.

Modulating properties at the surface of the disk include polarization rotation, reflectance of light, etc [Marchant, 1990]. Polarization rotation is used on magnetooptical disks where different magnetization on the surface will cause different rotation of polarization through the Kerr effect. Reflectance modulation is the mechanism of choice for compact disks (CD) and write-once/read-many (WORM) disks. Maximum storage density on a disk is approximately $\lambda^{-2}/2$ bits/cm^2 where λ is the wavelength used to record the data.

A monochromatic light beam can be focused on a spot of diameter λ, i.e. the spot cannot be made smaller than the wavelength of beam without having diffraction-limited distortion come into effect. The area of such a spot is $\pi\lambda^{-2}/4$ and it can be approximated by λ^{-2}. If we leave a blank space in between successive bits, interleaving this pattern as shown in Figure 2.8, then the data density is in the order of λ^{-2} as we mentioned already.

For red light emitted at approximately 600 nm the data density reaches approximately 1.2 Gbits/cm^2 or 33 Mbytes/cm^2. For magnetic disks the areal data density varies between 40 KBytes/cm^2 to 52 MBytes/cm^2 [Esener, 98a].[1]

1. It should be noted here that IBM demonstrated data densities in the order of 2×10^{10} b/in^2; i.e., approximately 370 MB/cm^2. It is expected (per the company's public statements at www.storage.ibm.com) that systems will become commercially available in the second half of 2002. Furthermore we should mention that the National Institute for Standards and Technology is challenging the research community to engineer planar magnetic storage systems that will exhibit densities between 185 and 18,482 MB/cm^2.

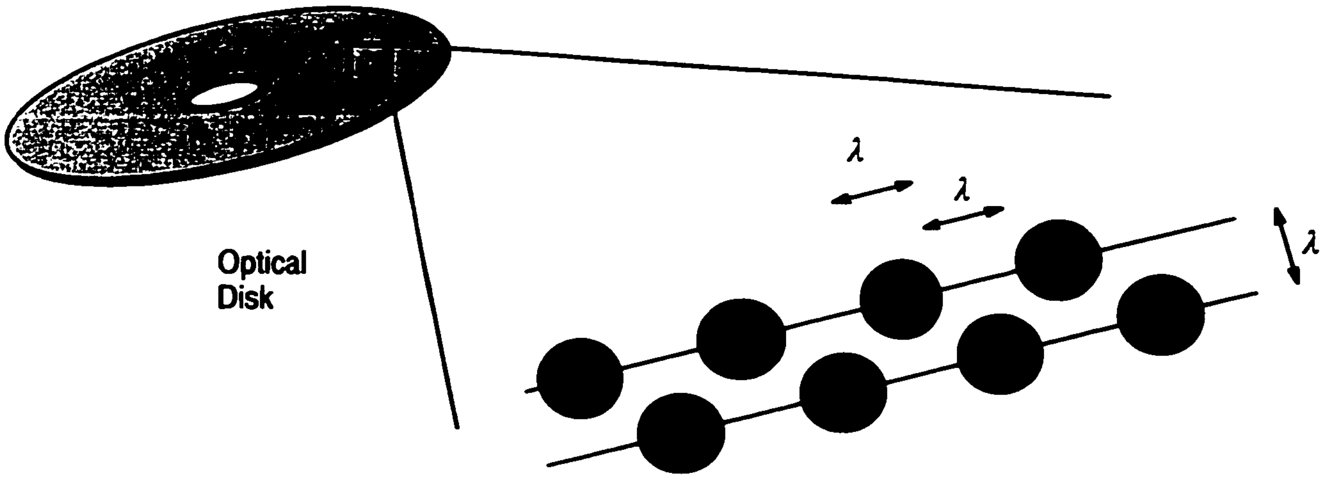

Figure 2.8. ***Bit arrangement along the tracks of an optical disk.***

3

Performance of Optical Memories

Optical memories are the most mature component of optical computing because they offer a viable alternative to data storage problems. Most notably, optical memories offer the potential for high data densities because of the volumetric confinement of a light spot. In a diffraction-limited system, i.e. an optical system which has perfect refractive and reflective components, a data spot can be confined on a surface of area $\sim 1/\lambda^2$ or in a space with volume $\sim 1/\lambda^3$. It is unlikely that a realistic system will achieve the theoretical data spot sizes in the near future, however it is possible that optical storage system can offer higher storage densities than capacitive and magnetic memories. We prefer the term *capacitive* to *electronic* memory in the context of this discussion for two reasons. First, the term electronic refers to something related with electrons. As we saw in our review on optical storage machinery, some mechanisms do rely on electronic interactions to store a bit value optically. Second, modern VLSI dynamic random access memories (DRAMs) store a bit value in form of a charge at a capacitive junction of the micro-circuit. Therefore, the term *capacitive* describes better the DRAMs which are the current best choice for computers' memory.

Modern data storage systems are characterized by the following criteria:

(i) storage capacity;
(ii) areal density;
(iii) volumetric storage density;
(iv) data transfer rate (DTR) or simply bit rate (BR);
(v) bit-error rate (BER);
(vi) probable life time (PLT);
(vii) cost.

The storage capacity of the system and the volumetric storage capacity are closely related. For example, a hard disk drive with 10 GByte capacity comes into a small package with approximate dimensions $5 \times 15 \times 15$ cm^3. Therefore, the volumetric density is approximately 27.3 MBytes/cm^3 or about 218 Mbit/cm^3.

The areal density varies significantly for different media. For example a hard disk unit that uses giant magnetoresistive materials can hold up to 166 Mbits/cm^2. At the same time a 64 MByte SDRAM on a 2×3 cm^2 dual in-line memory module (DIMM) has a storage density of 85.3 Mbit/cm^2.

The data transfer rate of a memory indicates how fast we can retrieve data from it and channel them to a bus. A typical—by today's standards—DRAM with access times in the order of 50 ns can output data at a rate of 20 Mbits per second (Mbps) [SIA, 98a]. The data transfer rate by itself is not a very useful measure of a memory's quality, because we also need to know how many of these bits are transmitted error free. The bit-error rate measures also the quality of a memory system by providing an expected time between two erroneous transmissions. The inverse product of the bit-error rate and the bit rate gives the mean time between errors (MTBE):

$$\text{MTBE} = (\text{BER} \times \text{BR})^{-1} \tag{3.1}$$

For example, a memory with a BR of 20 Mbps and a BER of 10^{-12} is *expected* to produce an erroneous state every 50,000 seconds or approximately every 14 hours of continuous operation.

The probable life time and the cost are two related parameters. The cost of a memory should be kept low while the probable life time of the system should be kept long. In addition it is desired that small increments of cost should yield longer lifetimes

In this chapter we discuss the performance of optical memories based on the criteria we mentioned earlier and we provide a comparative analysis with respect to capacitive and magnetic memories. The objective of the present analysis is to assert the competitiveness of optical memories with respect to currently available memories. We begin our discussion with a comparative analysis of densities and capacities in Section 3.1. A comparison between the transfer rates of optical memories and those of electronic and capacitive memories is presented in Section 3.2. Arguments about the lifetime and cost of optical memories versus their competing technologies are discussed in Section 3.3. Finally, we conclude this chapter with an overall discussion of our evaluation.

3.1 Densities and Capacities

It was mentioned earlier that optical memories show the potential for higher storage densities than capacitive and magnetic storage. In this section we offer a comparison basis between optical memories and competing technologies, including magnetic disks and capacitive storage.

At present, capacitive memories offer 512 and 1,204 Mbit synchronous dynamic random access memory (SDRAM) dual in-line memory modules (DIMM's). Two and 4 Gbit SDRAM chips are expected within the next few years [SIA, 98a].

The state of the art in magnetic storage is giant magnetoresistive materials (GMR) which exhibit densities in the order of 1 Gbit per square inch, or approximately 166 Mbit/cm^2. IBM has demonstrated systems with densities in the order of 2×10^{10} bits/in^2, i.e., 2.88 Gb/cm^2. These systems are expected to reach the market in 2002.

The best volumetric optical memory demonstrated at present offers a density of approximately 34 Gbit/cm^3 [Shelby, 97a; Psaltis, 98a]. The maximum possible density for these memories is in the order of 8×10^{12} bits/cm^3, assuming a wavelength at $?$ = 500 nm.

These figures do not provide a lot of information about the overall system. It is useful to consider a *system volumetric density*, i.e., the volume of the actual hardware component. For our estimates we assume that a hard disk drive can be housed (and actually is) in 10^3 cm^3. Further we will assume that a holographic storage system can also be housed in a similar volume.

SDRAM's with capacity of 1 Gbit require a chip size of approximately 10×10 mm^2 [SIA, 98a]. The areal density of the chip is approximately 1 Gbit/cm^2. These chips are placed into a DIMM package measuring $2 \times 3 \times 0.5$ cm^3 and the volumetric density becomes approximately 342 Mbits/cm^3.

First we compare existing systems. A 1 Gb SDRAM, a 20 GB hard disk, and a holographic memory with a total of 34 Gbits stored in it. For simplicity we will be using bits instead of bytes.

Table 3.1. Existing system volumetric densities.

System	*Capacity (Gb)*	*Volume (cm^3)*	*Volumetric Density (Mb/cm^3)*
SDRAM	1	3	342
Hard disk	160	1,000	163.84
Optical Memory	34	1,000	34.82

Table 3.1 shows that optical memories are an order of magnitude behind their competing technologies. Now let us consider the best that each technology promises to offer. The results are summarized in Table 3.2.

Table 3.2. Best system volumetric densities.

System	*Capacity (Gb)*	*Volume (cm^3)*	*Volumetric Density (Mb/cm^3)*
SDRAM	4	3	1,368
Hard disk	800	1,000	819
Optical Memory	7,451	1,000	7,630

For Table 3.2 we used the Semiconductor Industry Association's (SIA) best estimate for SDRAM chips [SIA, 98a] and IBM's projections for high-density hard disk performance. We also used the theoretical limit λ^3 to obtain the optical storage density at 8×10^{12} bits/cm^3. The path from Table 3.1 to Table 3.2 is mapped in Figure 3.1. From this figure it becomes clear that optical memories have a significant density advantage over competing technologies.

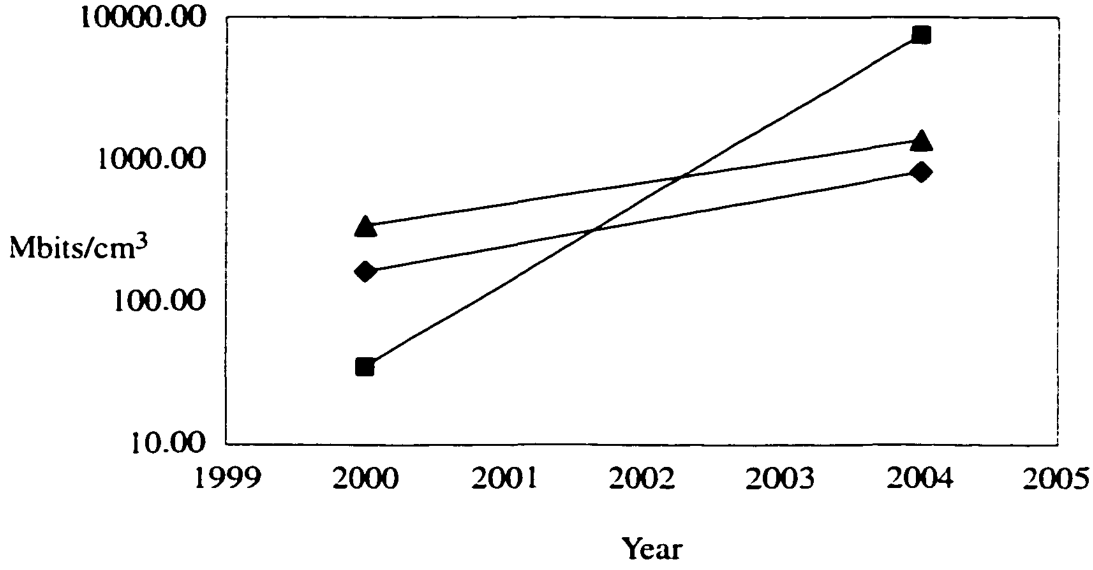

Figure 3.1. *System volumetric density projection.* ▼ *SDRAM;* ◆ *Hard disks;* ■ *Optical memory.*

To make an optical memory system more competitive the following conditions must be met:

(i) reduce the operational volume of the system;
(ii) increase the volume of the storage medium;
(iii) increase the storage density of the medium.

The first condition, making the operational volume smaller, will be achieved gradually as optical storage systems move from academic laboratories into industrial R&D projects. From a user-end point of view a realistic system requires electrical power input and connects electronically to a computer via a bus such as PCI, SCSI etc. If the optical memory system is compact enough, as is an external hard disk drive, with a total volume of 1,000 cm^3 and a storage capacity of 100–200 Gbytes (800–1,600 GBits), then the volumetric density becomes approximately 820–1,640 $Mbits/cm^3$.

The same storage capacity based on 1 Gbit DRAMs will require 800–1,600 DIMM's, each measuring 3 cm^3. The volumetric density of the DRAM system will remain approximately 341 $Mbit/cm^3$ or perhaps will become a bit lower than this, if we take into consideration the spacing of the chips on the board. The required volume for that many chips ranges between 2,400–5,000 cm^3 without taking into consideration room for air circulation for cooling purposes. However we should mention here that we offer this comparison only for reference. In all fairness optical memories and capacitive devices should never compete in the same memory hierarchy.

To improve the performance of optical memories further it is necessary to increase the volume of the storage medium. A larger medium, either occupies more in the physical domain or is the result of superposition techniques such as wavelength division multiplexing. It has been already demonstrated that a photorefractive crystal measuring 1 cm^3 can store approximately 34 Gbits/cm^3 (or approximately 4 GBytes/cm^3). Assuming a linear extrapolation, a 64 cm^3 crystal should be able to store approximately 2.125 Tbits which corresponds to 640 GBytes. Confining this system into a 1,000 cm^3 volume yields a system volumetric density of approximately 2.176 Gbits/cm^3. It will take 2,176 DIMM's of 1 Gbit to achieve the same capacity. The volume required for that many chips is at least 6,000 cm^3. At a nominal cost of approximately $250 per DIMM the cost of the capacitive alternative will be $544,000. A volume holographic storage system costs approximately $100,000 (including the salary of a graduate research assistant, optoelectronic interfacing, lasers, lenses, etc). One dollar can buy 4 Mbits of the SDRAMs used in our example, or 22.3 Mbits of optical memory.

Again, one may argue here that DRAMs are not intended for secondary storage while optical memories are, and therefore the comparison above is not quite fair. To complete the comparison, we shall take a brief look at high-density magnetic storage which employs the giant magnetoresistive effect. For this technology, the best data density is in the order of 1 Gbit/in^2 or 166.5 Mbits/cm^2. At the time of writing a 37 GByte disk would sell for approximately $315, yielding $0.001/Mbit. It would cost approximately $2,230 to store 2.125 Tbits. This cost is considerably lower than the cost of a volume holographic system. In all fairness we should mention that the cost of the disk-based system is based on street-pricing of a commercially available product while the cost of the optical system is based on R&D estimates. A commercially available optical storage system, if and when it hits the market, should be significantly less expensive.

Finally, the third condition has a bifold solution: (a) improve the optical design of the system and (b) increase the dimensionality of the storage space. The first solution should lead, ideally, to a diffraction-limited system. This system will exhibit a data density in the order of 1 $Tbit/cm^3$. Based on the availability of optical memories with density approximately 33 $Gbits/cm^3$ we should expect one or two more orders of magnitude increased density in the next few years.

The second solution mentioned above, provides more space for storing data. In addition to the physical three-dimensional space it is possible to perform wavelength division multiplexing and bias electric field multiplexing as well, thus storing more data in the same physical volume. A factor of 100 in wavelength multiplexing and a factor of 100 in electric field multiplexing can increase the storage density of a system by four orders of magnitude. Such breakthroughs are expected from the development of spectral hole-burning memories. Table 3.3 summarizes the discussion on densities and capacities andFigure 3.2 shows the ratio of MB/S versus system volumetric densities from Table 3.3.

Table 3.3. Trends in digital storage technologies.

System	*Volume (cm^3) for 2.125 Tbit system*	*System Volumetric Density (Mb/cm^3)*	*Cost*	*Cost Density (Mb/$)*
DRAM	6,000	~370	$544,000	4
Hard disk	60,000	~ 37	$2,230	1,000
Optical Memory	1,000	~2,230	$100,000	22
Optical Memory (best case)	1,000	~2,230	$1,000	2,230

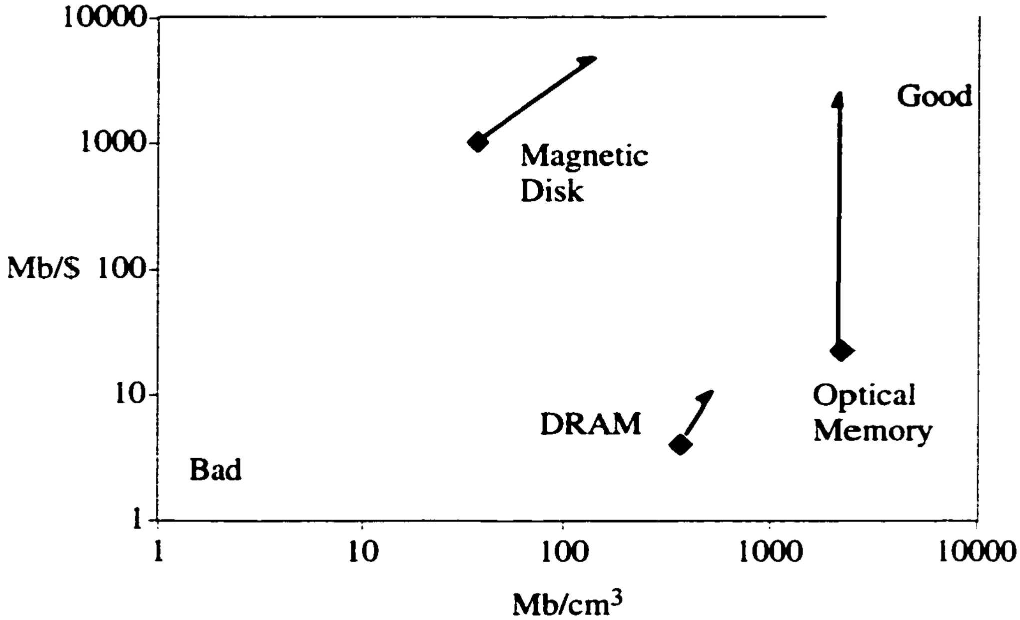

Figure 3 2. Expected trends in dollar power versus volumetric densities. The arrows indicate the most probable direction and speed of the evolution of the systems discussed. The terms "good" and "bad" are relative and they are used with respect to secondary storage standards. Data are taken from Table 3.3 and are based on [Esener, 98a; SIA, 98a].

3.2 Transfer Rates

In the previous section we suggested a correlation between the system volume densities of optical memories and the respective *cost densities*. By cost densities here we mean the volume of information we can fit in a currency unit, which in our study is one U.S. dollar. As a figure of merit cost densities and volume densities provide an evaluation scale for optical memories but they do not offer a complete picture in a systemic context. To complete the picture and provide the ground for a more accurate evaluation of optical memories we should introduce the metric of transfer rates and relate it to cost and volume densities.

The transfer rate is the number of bits an optical memory outputs in a unit time, typically one second. As such, the transfer rate R depends on the number c of bits in the output page and the time $t_{\text{retrieval}}$ it takes to retrieve that unit:

$$R = \frac{c}{t_{\text{access}} + t_{\text{retrieval}}} \tag{3.2}$$

In Eq. (3.2) t_{access} is the time required for the system to locate the desired set of data.

Obviously we wish to have a small denominator and a large numerator. However, there are practical limitations. Let us consider a volume holographic storage system as a typical example. In such a system a composing device like a spatial light modulator (SLM) is used to produce an image of the data to be stored in the memory. These devices have a finite number of pixels which are used to compose the image. Assume that there are $M \times N$ pixels in the device and a bit is composed by a block of $m \times n$ such pixels. The bit capacity of the data page is $\lfloor M/m \rfloor \times \lfloor N/n \rfloor$ which is the maximum value for the numerator of Eq. (3.2).

There are similar considerations for the denominator of Eq. (3.2). For example, the access time t_{access} cannot be made less than the time required by the position equipment in a spatio-angular multiplexing system.

We should clarify an important issue now. There are two distinct data rates in an optical memory system. One is related to the write operation and the other to the read. Since optical memories are touted as ideal for archival storage we narrow our attention in this work to the data rate during the read operation.

In Eq. (3.2) the denominator corresponds to the frame period or, inverted, to the frame rate, i.e., the number of pages that we retrieve per unit time. To reach a high frame rate apparently it is necessary to achieve low access and retrieval cycles. For video applications a frame rate between 20-30 frames per second (fps) is expected. This translates to $t_{access} + t_{retrieval} \leq 50$ msec. At that rate (20 fps) the system produces $20 \times \lfloor M/m \rfloor \times \lfloor N/n \rfloor$ bits/sec. Maximum performance is achieved when $m = n = 1$ and the data rate for the system becomes $R = 20 \times M \times N$ bits/sec. For $M = N = 1024$ the rate becomes 20 Mb/s. Considering that 50 msec is a rather slow access and retrieval cycle the estimate of 20 Mb/s should be considered as the worst-case scenario. But if we consider a holographic system that uses no mechanical components for addressing (a system would employ an SLM array for addressing) the access and retrieval cycle could be approximately 1 msec [Burr, 98a] and for $M = N = 1024$ the data rate becomes 1 Gb/s. It is possible to achieve sub-millisecond performances thus reaching rates above 10 Gb/s [Burr, 99a].

A final note should be made on the fact that the data rate of an optical memory (or any memory system for that matter) is affected by the error rate of the system. Storage systems are prone to errors and optical memories are no exception to this rule.

The issue of error control in page-oriented optical memories is in its infancy and it is such a broad subject that we will not discuss it in detail in this work.

3.3 Lifetime and Cost Issues

Any attempt to analyze the cost parameters of an emerging technology is an improbable task. We are capable of analyzing the cost of developing a specific electronic product or device because we have a cogent understanding of the cost issues related to the implementation of a fabrication/assembly facility and we know that with the exception of the negligible fraction of the production volume, all products will be shipped to consumers and they will create a revenue. Economies of scale and marginal cost analysis apply for such products.

This is not the case with page-oriented optical memories yet. There is no established process of mass production. In addition the ongoing research and development is in the basic stages and any estimate of the cost parameters will be severely skewed.

However we should consider the fact that, according to the National Institute for Standards and Technology, in 1996 the demand for rigid disk drives (to be used in PC's and servers) reached a total volume of 160×10^6 TBytes (or 1.4×10^{21} bits). It is expected that the demand will exceed 2,800 $\times 10^6$ TBytes in the year 2000 (or 2.46×10^{22} bits) [source: NIST's Advanced Technology Program, http://www.atp.nist.gov/atp/focus/dds.htm].

The cost of magnetic storage follows Moore's law and it is consistent with economies of scale. Magnetic storage, therefore, is considered a commodity today. It is imperative that volumetric optical storage should exhibit similar economic attributes when it becomes available. Optical disks already follow the rules of economies of scale and are priced on marginal cost.

If volumetric optical memories are introduced as rewritable media for secondary storage to compete with magnetic disks in the market of 2.46×10^{22} bits that NIST expects in the year 2000, they shall demonstrate several advantages. There are two key trade-off figures that we can consider here. One is the number of MBytes that $1 can buy for the user and the other is the number of MBytes/sec that the same $1 buys. Both figures are a function of time in the sense that technology progresses with the advance of time and prices are reduced. In other word more megabytes and faster access to those, can be purchased with the same amount of money as time progresses. The trade-off space is shown in Figure 3.3. Here t_0 denotes the time when optical memories are introduced in the market. At that time both $(MB/\$)_{optical} > (MB/sec/\$)_{magnetic}$ and $(MB/sec/\$)_{optical} > (MB/sec/\$)_{magnetic}$. As generations for both technologies emerge at isochronous[1] steps $t_1, t_2, \ldots$ the gap between the two performance figures narrows.

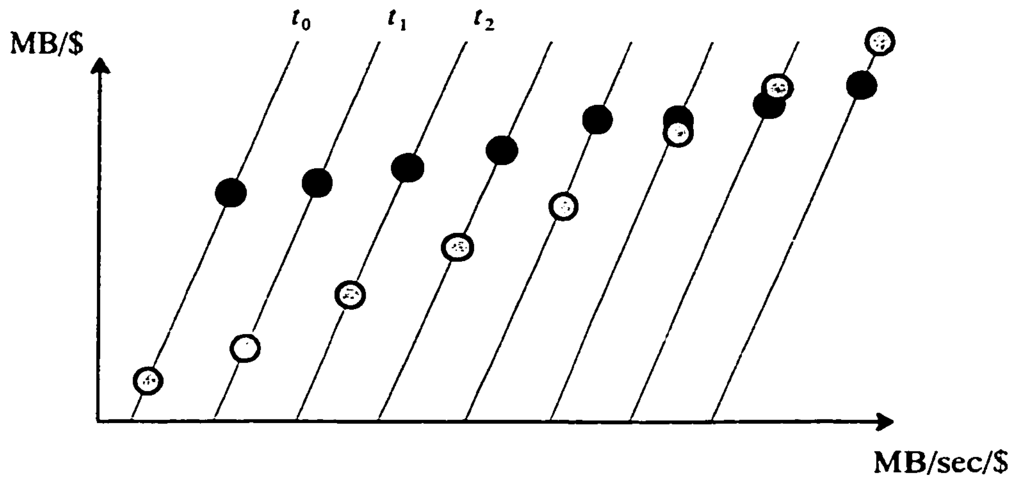

Figure 3.3. Trade-off patterns for magnetic (○) and optical (●) memories.

1. We assume isochronous steps for simplicity.

I was seeking the Venus with favorable tali, but the damned dog always lept out.

(Propertius, 50 B.C. – 15 A.D.)

4

Interface Modelling

Parallel optical memories produce two-dimensional pages of data outputs which are then read by photodetector arrays. In most experiments these photodetector arrays are connected to a real-time monitor, a frame-grabber or a computer for display and storage (in magnetic disks) of the memory output. Little consideration has been given into feeding the output of a parallel optical memory directly to an electronic host for further digital processing at the bit level as raw data. A simple method of transferring data from an optical memory to an electronic computer was proposed [Heanue, 94a]. However, neither a generic model nor a performance analysis were offered. In 1997 Ozaktas produced two papers summarizing the work in the field of optoelectronic computing [Ozaktas, 97a, 97b]. To complement and extend Ozaktas' work we present models and simulations results about the major compatibility issues related to interfacing optical memories and electronic computers.

An optical memory/electronic computer hybrid system has two cycles (or phases) of operation as shown in Figure 4.1. During the *write* cycle data are transferred from the electronic computer to the optical memory and during the *read* cycle data are moved from the optical memory to the electronic computer. In both occasions, an optoelectronic interface acts as an ambassador between the two systems.

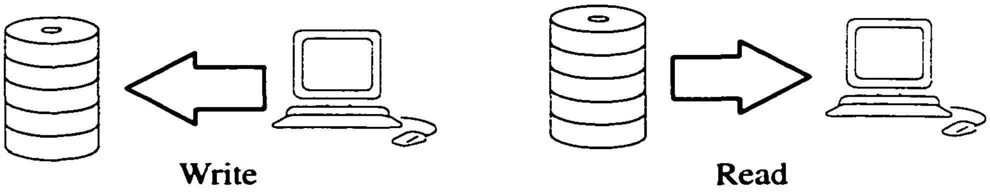

Figure 4.1. Operation cycles for an optical memory/electronic computer system.

A number of factors can decrease the quality of the system operation during both cycles. These effects introduce signal-to-noise degradation and crosstalk phenomena. In an ideal, and thus impractical, system the objective will be to eliminate all sources of such detrimental effects. In a realistic system however, the objective is to determine the system's tolerances to those effects and assert its operation limits. The following two sections describe attempts to establish a quantitative model to describe the system's performance and tolerances during the write and read cycles.

4.1 Signal Degradation and Crosstalk During Write Phase

During the write cycle, data are transferred from the electronic computer to the optical memory. The optoelectronic interface employed by the system, maps the electronic data to a two dimensional array of optical signals which are then imaged in the storage medium and recorded there. Four distinct subsystems can be identified within the path of data from the computer to the optical memory:

(i) electronic data bus;
(ii) two-dimensional data page composing mechanism;
(iii) data imaging subsystem;
(iv) optical medium.

In addition to the components listed above, the path includes an addressing mechanism and related control subsystems which are not directly involved in the data path per se.

The analysis assumes that the electronic data bus is effectively error-free and that data arrive at the page composing mechanism error-free. Therefore, the analysis focuses on the detrimental effects introduced by the last three components of the list above.

4.1.1 Data Page Composing Mechanism and Imaging Optics

The data page composing mechanism and the imaging optics subsystems are very much interdependent components, and we discuss them together in this section. One widely adopted technique is the use of a spatial light modulator (SLM). A collimated beam illuminates the SLM as shown in Figure 4.2. The result is the production of an array of beamlets according to a binary pattern stored on the modulator. This pattern represents the data received from the electronic computer. A number of problems can arise from this particular arrangement. Most notably, SLM's cannot handle high optical powers, they cause significant diffraction side-effects, are relatively slow, and are available in small to medium array sizes. Despite their drawbacks, SLM's are a very convenient tool for data composing mechanisms.

The amount of optical power an SLM can withstand is a serious limitation especially in reflective mode devices where a substrate is expected to absorb much of the incident optical power. Heat is of course the primary concern here but in addition photoexcitation of a conductive substrate can be a serious problem in electrically addressed devices. In such devices the incident optical power may create sufficiently high surface photocurrents which will result to erroneous states at some pixels of the modulator. In transmissive mode devices, photoexcitation and erroneous states are not much of a problem, but heat is, especially when the device has a number of pixels in the off-state and thus absorbing much of the optical power. A number of experiments [Goertzen, 95a; Burr, 95a] have demonstrated that 50–150 mW of optical power incident to a transmissive thin-film transistor (TFT) module does not overload thermally the device.

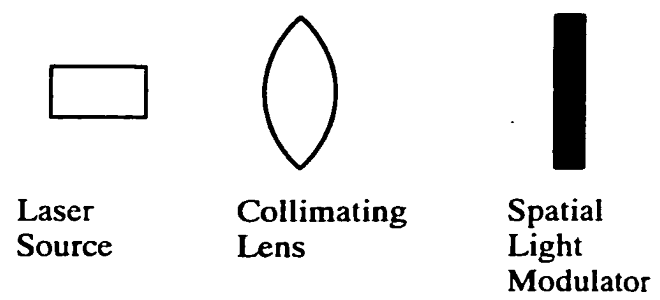

Figure 4.2. SLM used in a data page composition mechanism.

Diffraction is a particularly difficult problem when dealing with SLMs. Ideally an SLM pixel is a square aperture which can be imaged without much distortion. At present SLM technologies offer pixel sizes in the order of $10 \times 10\,\mu m$ up to $50 \times 50\,\mu m$ and therefore diffraction effects from the square aperture itself are not an issue. However, it is expected that advances in SLM technologies will soon offer pixel sizes comparable to the wavelengths used with optical memories, and diffraction from the aperture will affect significantly the imaging system. Nonetheless, sufficient image restoration can be achieved with a simple $2f$ type system or anamorphic lenses. A more important diffraction-based problem arises from the micro-structure of devices such as thin-film transistor flat panel displays which exhibit a regular pattern of transistors, interconnects, and other very large scale integration (VLSI) features in the micrometer scale. This micro-structure causes the image to be diffracted to numerous orders, and Fourier filtering is needed to select the zeroth order pattern. This situation is depicted in Figure 4.3.

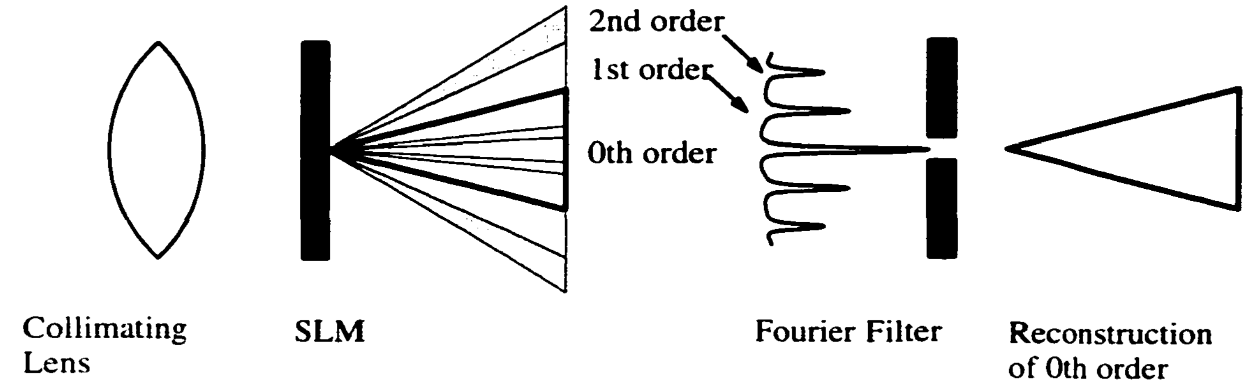

Figure 4.3. Dealing with diffraction of SLMs.

Another consideration when using SLM's for data page composition is their speed. Most of the currently available SLM's offer scanning rates which are less than a few thousand frames per second. Considering that most SLMs were developed for television and computer imaging, these rates are acceptable and satisfactory. In an optical storage system, however, it is desirable to be able to update the contents of the data page composition system as fast as possible. From this perspective, the access rate to the optical memory is given by:

$$\mathcal{R}_{\text{access}} = f_{\text{SLM}} \times N_{\text{SLM}}^2 \qquad (4.1)$$

where f_{SLM} is the frame rate of the SLM and N_{SLM} is the linear size of a square SLM in pixels. Eq. (4.1) holds if we use only one pixel per bit. We will see later that this may not always be possible. Nevertheless, the importance of a fast SLM device is evident.

Another important factor, implied by Eq. (4.1) is the array size of the SLM, N_{SLM}. For simplicity we assume a square SLM with N_{SLM}^2 pixels. The access rate scales quadratically with respect to the array size of the modulator and therefore SLMs with a large number of pixels are desired. An additional reason for having large N_{SLM} is that more than one SLM pixels are used to represent a bit and therefore the larger the number of pixels, the larger the number of bits. Ideally, one would expect a one-to-one mapping between data bits and pixels. Unfortunately SLM's do not provide the necessary picture quality (in terms of contrast ratio) to achieve this mapping but it is expected that eventually these devices will reach sufficient quality levels.

Despite their numerous drawbacks, SLM's are used today because they are readily available and relatively inexpensive which make them ideal for system demonstrations. Alternative technologies for data page compositions include arrays of laser sources such as vertical-cavity surface-emitting lasers (VCSELs). These devices generate coherent radiation which deviates slightly by $\pm 5°$. Therefore, a collimating lens is needed close to the VCSEL's aperture to steer the beam.

Inhomogeneities in the VCSELs' output planes may cause some beams to be out-of-focus with respect to the collimating lens array and some crosstalking will occur. However the deviations from the focal plane are in the order of up to 100 Å and resulting mis-collimation will not create any crosstalk unless the optical storage medium is positioned several meters away.

Small arrays of VCSELs with up to 32×32 devices are available. The size of these arrays is not sufficient for a realistic system. However it is possible to build such a system to demonstrate the potential of VCSEL arrays in volume optical storage.

4.2 Signal Degradation and Crosstalk During Read-out Phase

Typically an optical memory will produce a rather regular pattern of spots as shown in Figure 4.4. Darker spots correspond to logic-low bit values and brighter spots correspond to logic-high bit values. Ideally a logic-low data spot should correspond to zero intensity but in a realistic finite system one should not expect such performance. We should expect that the contrast ratio between a logic-low and a logic-high data spot, no matter how large it could be made, it will always be finite.

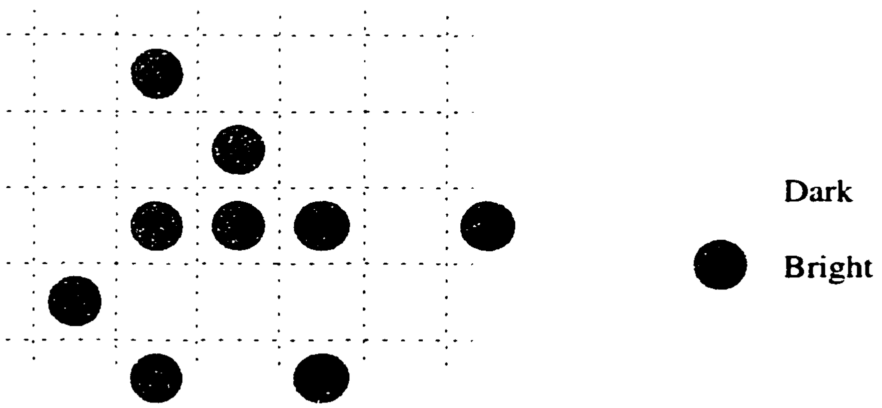

Figure 4.4. Regular pattern output of an optical memory.

In addition to the contrast ratio restriction we should also expect that the arrangement of data spots, although regular at a large scale, it may exhibit some irregularity between nearest neighbors. This irregularity, which is illustrated in Figure 4.5, can cause the following two problems: (a) signal degradation and (b) crosstalk.

Signal degradation can have detrimental effects when a logic-high bit is imaged on a photodetector, assuming bright-true logic. If the data spot is severely misaligned, the incident optical power on the photodetector may not be sufficient to switch the device to the logic-high state and a bit-error state will occur. Cross-talk can affect the detection of a logic-low state. Specifically if a photodetector device receives significant crosstalk from neighboring cells it may be possible to switch to a logic-high state even if the corresponding bit value is logic-low (i.e. dark). The opposite effects will be observed in a dark-true logic. Specifically

crosstalk will cause an erroneous logic-high state and severe misalignment will cause an erroneous logic-low state.

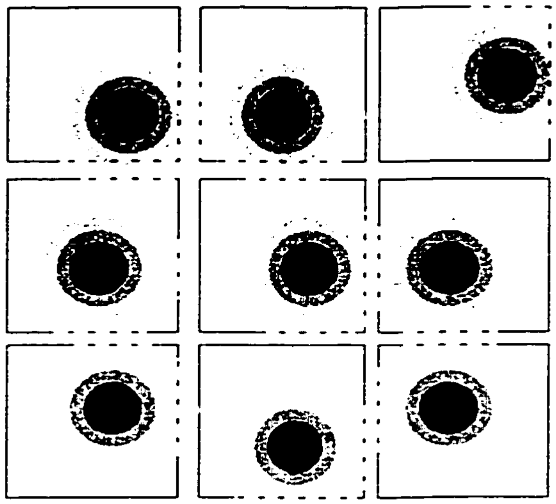

Figure 4.5. Nearest-neighbor pattern irregularity in an optical memory output. The squares correspond to photodetectors and the circles to data spots regardless of their bit values.

To determine the tolerances of a photodetection scheme, we performed a series of computer simulations based on a number of assumptions. Our first assumption was that the light intensity distribution of a data spot is normal, i.e., it obeys the Gaussian probability density function:

$$f(r) = \frac{1}{\sigma\sqrt{2\pi}}\exp\left[-\frac{(r-\eta)^2}{2\sigma^2}\right] \qquad (4.2)$$

where r is the distance of any point within the data spot from the center of the spot, σ is the distribution's standard deviation, and η is the distribution's average. It can be shown [Papoulis, 94a] that

$$\int_{\eta-\sigma}^{\eta-\sigma} f(r)dr \approx 0.7 \qquad (4.3)$$

In other words the interval $[-\sigma, \sigma]$ contains approximately 70% of the intensity's distribution. Therefore, a circle with diameter 2σ centered at η will confine 70% of the incident light from a data spot. A normal distribution centered at $\eta = 0$ and with $\sigma = 1$ and its integral are shown in Figure 4.6. As expected, the integral of a distribution density function over all possible values of its independent variable, is one.

Our study at this point has the following two objectives: first, determine the maximum misalignment possible without degrading the logic-high signal below a certain threshold. Then, use this maximum misalignment as a problem constraint to determine the maximum crosstalk from nearest-neighbor cells.

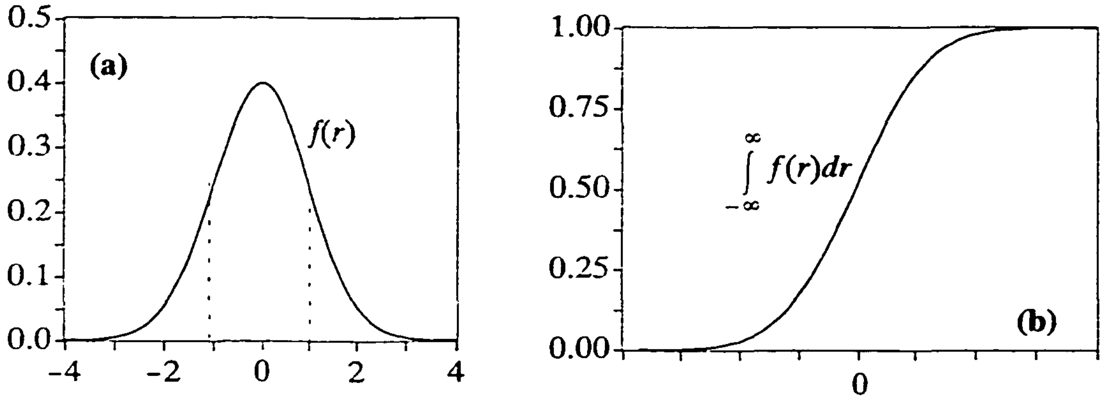

Figure 4.6. (a) Normal Distribution around η=0. (b) Cumulative probability for a normal distribution.

4.2.1 Signal Degradation Study

The geometric model used for this analysis is shown in Figure 4.7. For simplicity we assume that the photodetector has a square shape with side equal to $2a$, setting a=1. Ideally, the data spot should be a circle perfectly inscribed in this square. Therefore, the radius of this circle will be r_b=1, where the subscript b denotes that this circle correspond to a bit value. We further assume that this circle confines approximately 70% of the total optical power corresponding to the specific data spot. This assumption allows us to use Eq. (4.2) with $\eta = 0$ and $\sigma = 1$.

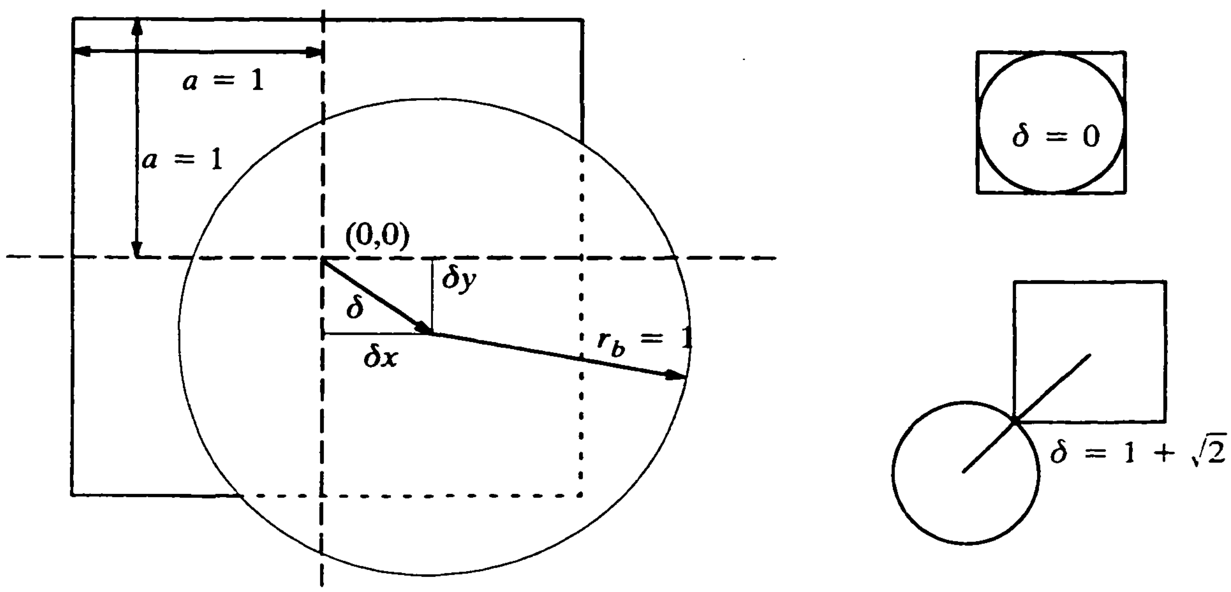

Figure 4.7. Geometric model of signal degradation due to misalignment. The two drawings at the right show the best (top) and worst (bottom) cases.

In a realistic situation, the data spot will be misplaced by δx and δy along the x and y axes respectively. The displacement (or misalignment vector) will then have a length:

$$\delta = \left(\delta x^2 + \delta y^2\right)^{1/2} \tag{4.4}$$

Now the objective of our study is to express the incident optical power on the photodetector I_d as a function of the misalignment vector δ, i.e., $I_d = I_d(\delta)$. If the indefinite integral of

Eq. (4.2) were known, we could evaluate it for different integration limits. However the indefinite integral of the normal distribution is not known explicitly (we recognize it as the error function). Fortunately, it can be evaluated numerically using computer simulations and specifically Monte-Carlo techniques.

The principal idea of the simulation is to randomly distribute points within the displaced circle and then count how may of them are in the square as well. The ratio:

$$R = \frac{\text{number of random points in the square}}{\text{total number of random points}} \tag{4.5}$$

is proportional to the incident intensity on the photodetector. It is easy to determine that for $\delta = 0$ the ratio will be R=1 and for $\delta = 1 + \sqrt{2}$ the ratio becomes R=0.

Our Monte-Carlo algorithm is summarized in the following steps:

Step 1.	pick a random point (x, y) in a unitary circle;
Step 2.	determine the normal probability $f(r)$ of having a point at $r = \left(x^2 + y^2\right)^{1/2}$;
Step 3.	pick a random number ψ in $[0, 1]$;
Step 4.	if $f(r) < \psi$ go back to step 1 and try again;
Step 5.	displace the (x, y) point by $(\delta x, \delta y)$; if the new point $(x + \delta x, y + \delta y)$ is within the square then update the count of points and skip to step 1; otherwise skip to step 1.

Each value of δ is the average of ten different samples for various combinations of $(\delta x, \delta y)$. Each sample points is evaluated after 250,000 trials, i.e. the total number of random points in the circle is 250.000. The results of the simulation are shown in Figure 4.8. For $\delta = 0$ and for $\delta = 1 + \sqrt{2}$ the incident intensity on the photodetector is 1 and 0 respectively, as predicted. One can also see that for $\delta = 1$ which corresponds to a half circle inscribed in the square, $I_d(\delta) = 0.5$

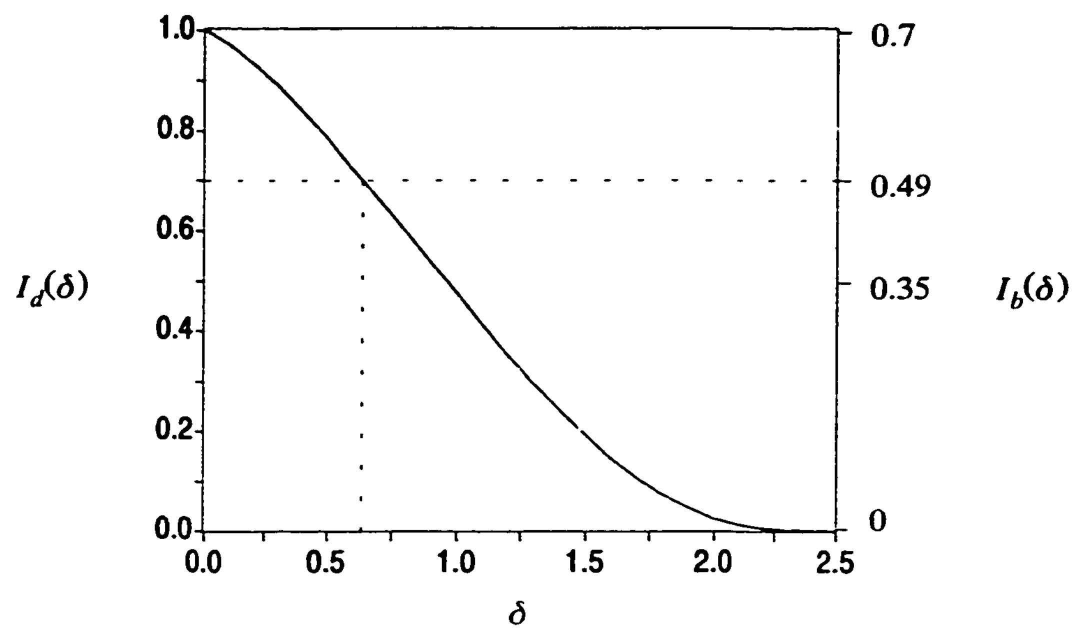

*Figure 4.8. **Simulation results of signal degradation due to misalignment. The scale to the right corresponds to the total intensity of a data spot (of which only 70% is incident on the photodetector).***

At a misalignment distance of $\delta = 0.6$ the incident intensity on the photodetector is only 70% of the total intensity confined in the data spot. Let us keep in mind here that the total intensity of the data spot corresponds to only 70% of the total intensity emitted at the source of the data bit. Therefore at $\delta = 0.6$ we receive only 49% of the optical power that the memory can produce during the readout cycle. Based on this information, the photodetectors can be designed so that $0.7 \leq I_d(\delta) \leq 1$ will correspond to a logic-high state and $0 \leq I_d(\delta) \leq 0.5$ will correspond to a logic-low state. These logic margins are shown in Figure 4.9.

The analysis so far has been focused around the hypothesis that $I_d(\delta) = I_b(\delta)/0.7$, where I_b is the total intensity of the data spot, or in mathematical terms $I_b = \int_{-\infty}^{\infty} f(r)dr = 1$; where the function f is given in Eq. (4.2). The scaling factor was obtained, as already demonstrated, by considering that the useful portion of the data spot is confined within a circle with radius r_b equal to the standard deviation σ of the normal distribution.

It is possible to change the scaling factor and make it closer to unity, by confining more optical power in the circular region. The original assumption was that $r_b = \sigma = 1$. If we keep this value for r_b but make the standard deviation σ smaller, then we can confine more than 70% of the optical power within the area of a circle with radius $r_b = 1$. The problem of reducing the standard deviation of an optical Gaussian profile has several components, each leading to a different solution. One solution is to better focus the emerging data beamlet from the optical memory onto the photodetector. This solution is part of the optical system analysis and will be discussed later.

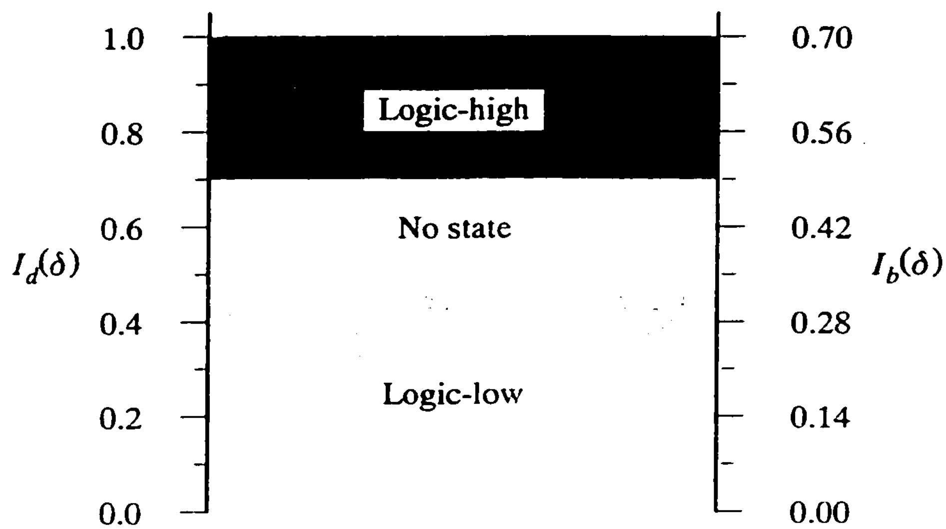

Figure 4.9. Logic state margins.

4.2.2 Crosstalk Studies

In the previous section, we discussed the tolerance of a photodetector with respect to the displacement of the data spot from the center of the device cell. We concluded that for displacements $\delta \le 0.6$ the incident power on the photodetector is sufficient to drive the device to a logic-high state. In the case where $\delta_{max} = 0.6$ it is important to calculate the crosstalk noise received by a photodetector cell from its nearest neighbors.

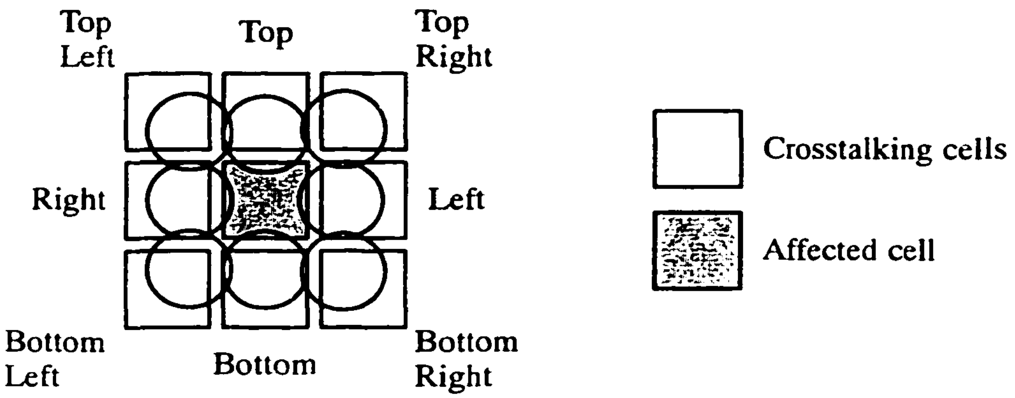

Figure 4.10. Nearest-neighbor crosstalk topology.

Of particular interest is the situation depicted in Figure 4.10, where each neighboring cell has its corresponding data cell misplaced towards the center cell by a vector $\delta_{max} = 0.6$. Figure 4.10 illustrates the worst case were the center cell receives maximum crosstalk from all its nearest neighbors.

In order to simplify our discussion we will argue that the crosstalk terms from the diagonal nearest neighbors can be ignored. Figure 4.11 represents the geometric model used for this argument. The diagonal gap between the two cells is

$$g = \left(g_x^2 + g_y^2\right)^{1/2} \tag{4.6}$$

where g_x and g_y are the horizontal and vertical gaps respectively. If the data spot of the offending diagonal neighbor is misplaced by $\delta_{max} = 0.6$ towards the affected cell along the direction that connects the centers of the two cells, then its perimeter will exceed the lower left corner of the offending cell by $1.6 - \sqrt{2}$. The condition

$$g \geq 1.6 - \sqrt{2} \tag{4.7}$$

ensures that the perimeter of a data spot will not reach a diagonal neighbor, therefore no diagonal crosstalk terms need to be considered. If we assume that $g_x = g_y = g_0$ then Eq. (4.7) yields $g_0 \geq 0.13137$. This gap corresponds to approximately 6.5% the side length of the square photodetector. Such separation between discrete elements on a VLSI chip is almost certain even at the submicron technologies, because of the various layout rules. It is also reasonable to expect at least that much physical separation between photodetector cells, because of the real estate requirements for interconnection lines and additional logic that the system may require.

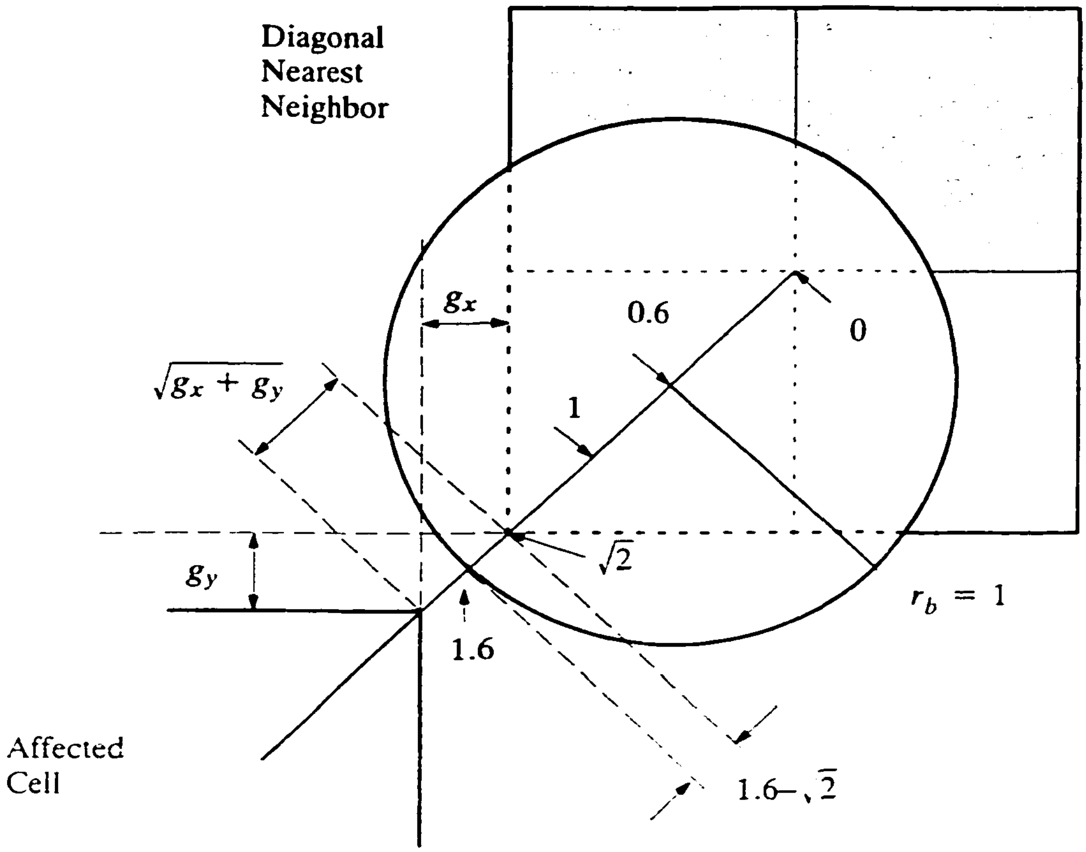

Figure 4.11. Geometric model for the calculation of crosstalk from a diagonal nearest-neighbor.

In fact, the very last argument of the previous paragraph is strong enough to justify a dismissal of any crosstalk from the horizontal and vertical nearest neighbors. But in order to provide a complete discussion on the subject, we use the geometric model of Figure 4.12. Again we assume that the data spot of the offending cell is misplaced by $\delta_{max} = 0.6$ towards the neighboring cell. In this case the crosstalk is due to the shaded area of Figure 4.12. For that area w we can write:

$$w \le \frac{1}{4}\int_{0.6}^{1} f(r)dr \tag{4.8}$$

where f is given by Eq. (4.2). The integral can be rewritten as $\int_{0.6}^{1} = \int_{0}^{0.6} + \int_{0.6}^{1}$ and from numerical calculation we obtain $w \lesssim 0.0464$. Because a cell has four immediate (horizontal and vertical) nearest neighbors, the total crosstalk will be $w \lesssim 0.185607$. In the worst case, assuming $g_0 = 0$, the total crosstalk will be $w \approx 0.185607$. From the logic margins of Figure 4.9 we see that a crosstalk of approximately 0.2 cannot create a false high-state.

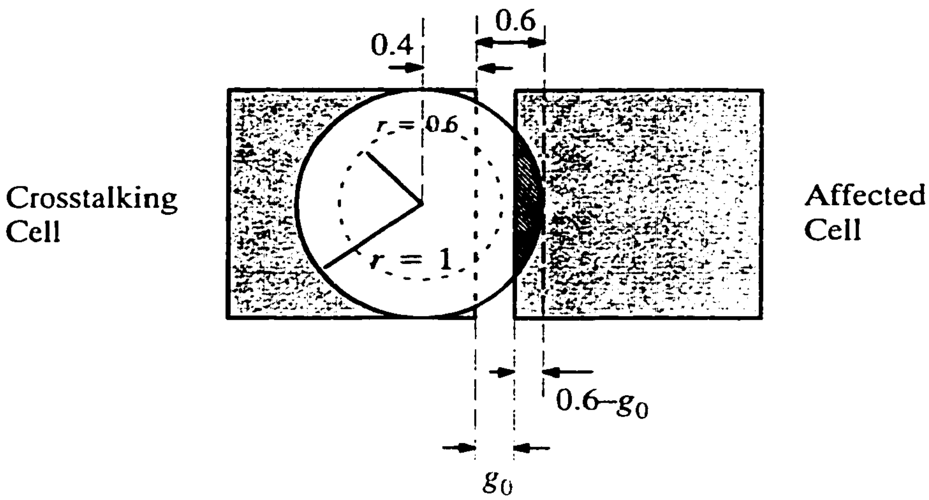

Figure 4.12. Geometric model for horizontal and vertical crosstalk analysis.

4.2.3 Discussion on Signal Degradation and Crosstalk

The analysis of signal degradation showed that a misalignment of up to 0.6 is acceptable for a photodetector with the logic state margins shown in Figure 4.9. For the model presented here, a misalignment of 0.6 corresponds to 30% of the photodetector's edge length or 60% of the data spot radius, assuming that 70% of the total spot's intensity is confined within this radius. A concealed assumption for this analysis was that the photodetector's responsivity is uniform throughout its active surface.

Using the previous result as a constraint in the study of the crosstalk effects we demonstrated that even in the worst case the crosstalk from the four immediate neighbor cells cannot force a false state. We also showed that the crosstalk from the diagonal nearest neighbors can be ignored. However we should point out that in the theoretical limit a cell will receive crosstalk noise from every cell on the photodetector array. Because the normal distribution drops asymptotically to zero we should expect small crosstalk noise from every cell no matter how far away it is on the array. We believe that these crosstalk terms can be safely ignored.

Kawabata spoke no English, Richie no Japanese, and their interpreter stayed home, sick in bed with a cold. So they talked in writers. Richie said "Andre Gide!" Kawabata thought about it, and then replied "Thomas Mann!" They both grinned. They'd go on grinning all afternoon, trading names like Flaubert, Edgar Allan Poe and Stefan Zweig.

(John Leonard: The Last Innocent White Man in America)

5

Interfacing Concepts

5.1 The Optoelectronic Interface

The optoelectronic interface (OEI) between parallel optical memories and electronic computers incorporates all the technology required for bidirectional flow of data and control signals between the two systems. There are two distinct data paths in the OEI. The computer-to-memory path which is the write path, and the memory-to-computer path which is the read path. We examine the system functionality along these two paths in the following sections.

5.1.1 The Write Path

During the write cycle data are transferred from the electronic computer to the OEI where they are encoded according to the data format employed by the memory. Then they are loaded onto the data-page composer, typically a two-dimensional spatial light modulator (SLM) and finally they are recorded into the optical storage medium. The flow of data during this cycle is shown in Figure 5.1.

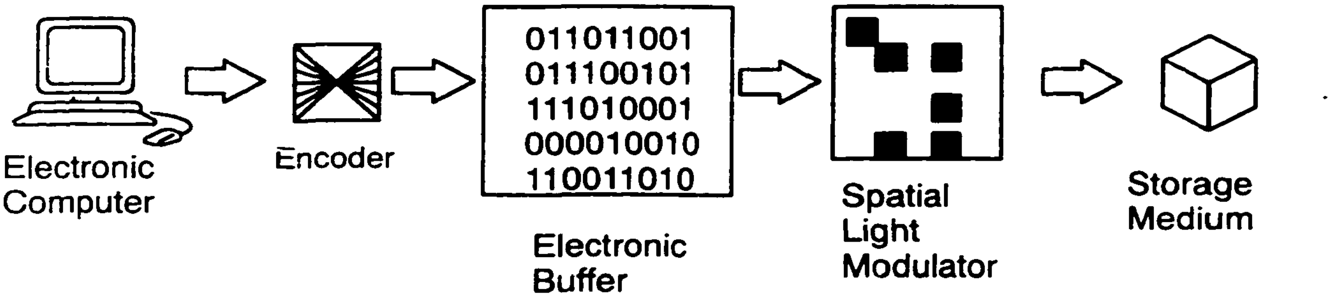

Figure 5.1. Data flow in the write path.

The path from the computer to the spatial light modulator is rather slow compared to the degree of parallelism the latter can offer (which is in the order of 512×512 channels, or higher). In other words we have to load the spatial light modulator in a serial fashion. An alternative is to route the data to an electronic buffer after their encoding and once the buffer is full to upload its contents onto the spatial light modulator. For commercially available spatial light modulators, such as liquid crystal-based devices, this approach appears to be simpler because it bypasses the rather complex electronics required for a sequential loading of image data.

Assuming that data exit the encoder at a rate P (in Mb/s), the total time required to load the page composer is

$$T_{\text{comp}} = MN/P \tag{5.1}$$

For MN=512^2 and P=160 Mb/s, T_{comp} becomes approximately 1.5 msec. If the addressing mechanism is based on an acoustooptic modulator, beam positioning (in a holographic system) can be accomplished while loading the data in the composer. Shutter control is also in the same order of magnitude, but unfortunately the dominant time component is the exposure interval which in current experimental systems can be several seconds long.

Therefore, from a parallel processing system point of view, the write path is a rather simple architecture with its parallelism located at the page composition device. A dense two-dimensional parallel interconnection network is formed optically between this device and the storage medium in order to transfer the data to the memory. The read path which we examine next is actually the real parallel processing component of the optoelectronic interface.

5.1.2 The Read Path

During the read cycle, a two-dimensional page of data bits is retrieved from the storage medium and imaged onto an array of photodetection devices. This array converts the presence or absence of photonic power into electronic logic states and routes these signals to the computer in the appropriate format. In fact, the array of photodetectors may be required to perform several tasks including:

- weighted detection of signals
- error detection and correction
- data decoding
- data filtering and selective routing, etc.

It is implied here that each of these tasks must be executed in parallel because data are received in parallel. As a result, a free-space optical interconnection network is being established between the storage medium and the photodetection array. For the photodetection array to function according to its prescribed role a fine-grain structure is required. Each processing element in this structure processes one bit. The components of the processing element include a photodetector device, logic gates, and signal feed-throughs to and from other cells.

By interconnecting a number of these elements in a two-dimensional array we can receive pages of data from the optical memory, process them in parallel and route the results to the electronic computer for further processing. In doing so, we must address a number of engineering problems such as:

- signal degradation and cross talk
- geometry and topology issues
- optomechanical precision
- embedding intelligence
- data rate balancing
- enabling technologies and implementation, etc.

Based on our introductory description of the OEI and its data paths we can say that the read path is determining the nature of the data processing architecture that the interface constitutes. Therefore we will shift the focus of our attention to the read path for the remaining of this paper where we discuss some of the issues listed above.

5.2 Classes of Interfacing

The functionality of optoelectronic interfaces (OEI) allows us to classify them into three broad classes:

- pass-through only,
- transform-only, and
- transform-and-reduce.

The pass-through interface is the simplest form of interfacing exhibiting minimal complexity and limited functionality, mainly in the form of signal detection and amplification. Specifically its role is to transfer data from the optical memory to the input/output bus of the electronic computer "as are". A typical example of a pass-through interface is the two-dimensional charged-coupled device (CCD) array. Because there is no preprocessing at all, incoming data may create bottlenecks at the input/output channel.

The transform-only (TO) interface has a more complex structure and therefore it offers some additional functionality, i.e., data preprocessing. Data are imaged from the optical memory onto the TO interface and they are subject to data decoding, error detection and correction, and other transformation operations.

The transform-and-reduce (TR) interface performs additional functions such as data filtering, i.e., its output is only a fraction of the input data. There is a fine distinction between the last two kinds of interfacing. One may argue that data decoding operations performed by the TO interface are data reduction operations. From a raw bit point of view this is certainly true. If, for example, data are encoded using an eight-out-of-fifteen code, decoding will reduce the 15-bit strings to 8-bit strings. However, from a data content point-of-view this reduction does not alter the information contained in these bit sequences. Based on this observation we say that data decoding is not a reduction operation.

TO and TR interfacing require an array of processing elements suitably interconnected with each other in order to perform the necessary operations. In addition, the TR interface requires external control signals. This characteristic of the TR makes it different from the TO. The processing elements comprise photodetection devices and combinational and/or sequential logic and they are arranged into two-dimensional arrays that match the geometry of stored data pages. The presence of these devices in the processing elements raises several issues about the potential performance of the interface. These issues involve misalignment tolerances, timing estimates, power budget, etc.

5.2.1 Timing and Rate Issues

We defined earlier (Eq. (3.2)) the readout data rate as $R = c/(t_{access} + t_{retrieval})$. Here c is the number of pixels of the data page. Let's assume that there are NM pixels and that the total time for the denominator is $T_r = (t_{access} + t_{retrieval})$. Then we can rewrite Eq. (3.2) as:

$$R = MN/T_r \qquad (5.2)$$

In fact, this equation needs to be scaled by a factor r to account for the redundancy introduced by the data code employed in the system, (i.e., r is the code rate). For example, if we use standard 8-bit ASCII code mapped on 15 bits then $r = 8/15$. Thus, the effective readout rate becomes:

$$R_{\text{eff}} = rMN/T_r \tag{5.3}$$

Now let us assume that an electronic computer can absorb data at a rate Q (in bits per second). We need to investigate the conditions under which the electronic computer may be overwhelmed; in other words we need to specify the regions for which the following inequality holds:

$$R_{\text{eff}} > Q \tag{5.4}$$

This inequality leads to:

$$T_r < rNM/Q \tag{5.5}$$

Therefore when the time required to readout a single page with nm bits coded with rate r, is less than rmn/Q, then a bottleneck will occur between the optical memory and the electronic computer. Figure 5.2 shows the rate R_{eff} for various values of $(mn)^{1/2}$, assuming a code rate of 8/15. The horizontal line indicates the 20 MB/sec threshold which is the current best performance for a SCSI-type data bus.

It is clear that the data flow rate at the read path has the potential to surpass the data bus rate of electronic computers. A trivial solution would be to design a faster bus. Although this is a desirable goal we should mention that optical memories will become even faster and perhaps they will overwhelm those fast buses as well. Another solution would be to reduce the amount of data entering the electronic computer. By allowing only a small fraction of the readout data to reach the computer we can balance the rates between the two systems and, hopefully, avoid a bottleneck.

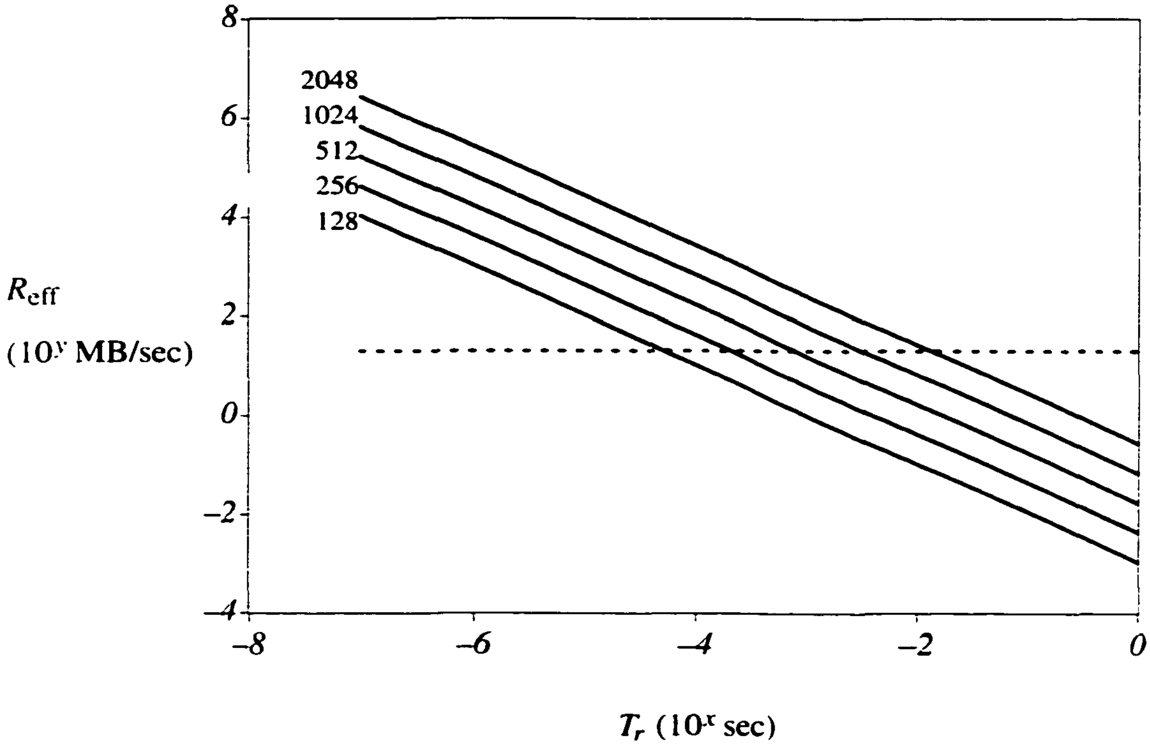

Figure 5.2. Data rates from the optical memory to the electronic computer for various sizes of the data page. The horizontal line marks the performance of a fast SCSI bus (approx. 20 MB/s).

To discuss this idea further we will briefly describe the relational database operation called *selection*. An example of the selection operation is shown in Figure 5.3. An original set of data containing the names of higher education institutions, their home states, and their Internet domain names is reduced in size by applying a search criterion to select only colleges located in Ohio.

A selection will produce a *selection vector*, as shown in Figure 5.4, which determines which data matched the search criterion. We proposed a CMOS array of photodetectors, called the optoelectronic random access memory (OERAM), which can output only the selected data, thus conditioning the data transfer rate between the optical memory and the electronic computer.

School Name	State	Domain Name
Miami University	OH	muohio.edu
Colorado State University	CO	colostate.edu
Tufts University	MA	tufts.edu
University of Arizona	AZ	arizona.edu
University of Dayton	OH	dayton.edu
Duke University	NC	duke.edu

State=OH

School Name	State	Domain Name
Miami University	OH	muohio.edu
University of Dayton	OH	dayton.edu

Figure 5.3. Example of selection.

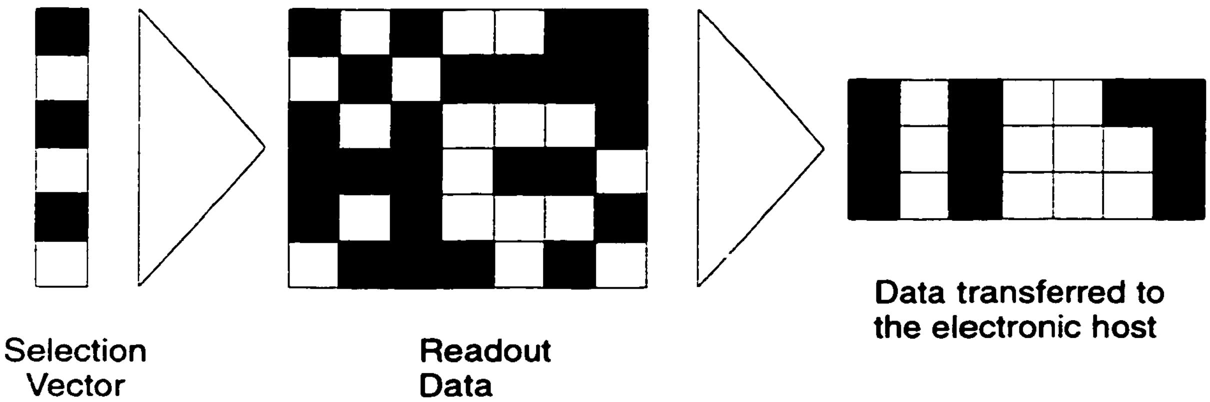

Figure 5.4. Functionality of selective row decoder-based read out of a data set.

The intelligent feature of this array of photodetectors is a *selective row decoder.* We use the term decoder because our photodetectors resemble the architecture of a static random access memory (RAM) chip. The functionality of the selected row decoder is shown in Figure 5.5(a). Using standard library cells the selective row decoder can be fabricated using a register and a shift register as shown in Figure 5.5(b). The register contains the bits of the selection vector while the shift register cascades a flag to indicate which row is under consideration. The expected logic states are shown in Figure 5.5(c) where we assumed that only rows 1 and 3 have been selected.

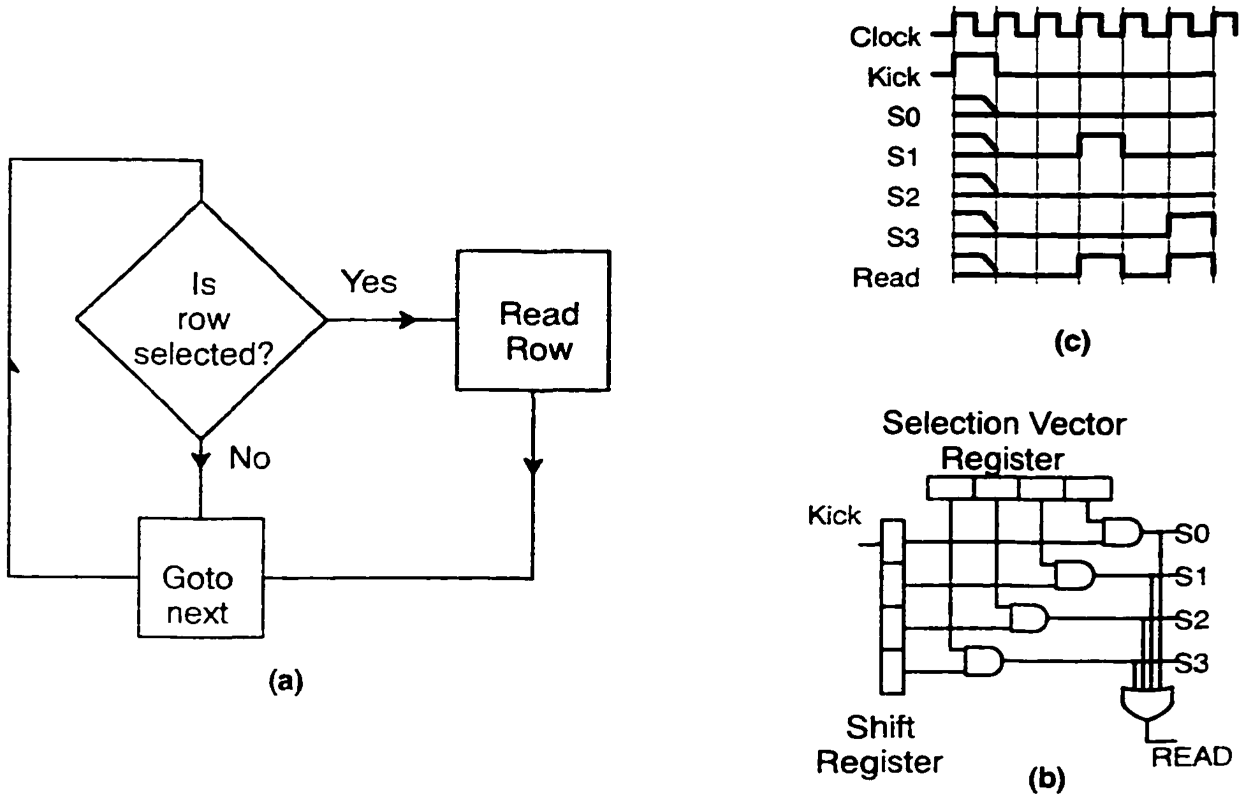

Figure 5.5. The selective row decoder. (a) flow chart; (b) simple logic implementation, (c) signal pattern.

The operation of the selective row decoder may appear trivial in a system where one word is read in one clock cycle. However, in a realistic situation where the photodetector array receives over 100 bits per row, several clock cycles may be required before the entire row is read. To demonstrate the importance of the selective decoder, let us introduce the following parameters. Consider a selectivity factor σ which is defined as the ratio of the number of rows matching the search criterion over the total number of rows. Also consider that the optical memory has a total of $\mathcal{M}_b$ rows stored. The total number of selected rows is $\sigma\mathcal{M}_b$ and the unselected rows are $(1 - \sigma)\mathcal{M}_b$. Now, let us assume that it takes one clock cycle to determine if a row is selected or not (based on the corresponding bit value of the selection vector) and $\lceil N/p \rceil$ clock cycles to read one row, where N is the number of bits in a row and p is the number of bits read per clock cycle. The total time required to complete filtering the data received from the optical memory is then:

$$T_{srd} = \left[(1 - \sigma)\mathcal{M} + \sigma\mathcal{M}(\lceil N/p \rceil + 1)\right]f^{-1} = \mathcal{M}(1 + \sigma\lceil N/p \rceil)f^{-1} \qquad (5.6)$$

where f is the frequency of the clock. Without a selective row decoder, the total time would have been:

$$T = \left[\mathcal{M}(\lceil N/p \rceil + 1)\right]f^{-1} \qquad (5.7)$$

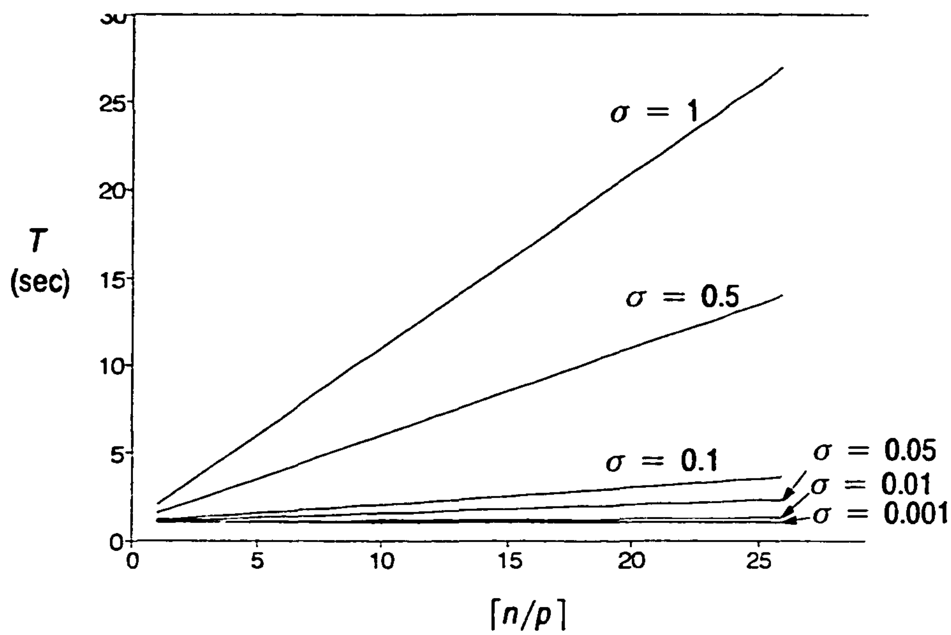

Figure 5.6. Performance of the selective row decoder-based interface for different values of the selectivity factor.

Obviously Eq. (5.6) is simplified to (5.7) for $\sigma = 1$. Figure 5.6 shows the time performance of a system with $\mathcal{M} = 10^7$ and f=10^7 Hz for six different values of σ.

The improvement that the selective row decoder introduces is bound by $1/\sigma$ because:

$$\frac{T_{without}}{T_{with}} = \frac{1 + \lceil N/p \rceil}{1 + \sigma\lceil N/p \rceil} \qquad (5.8)$$

It is reasonable to assume that $1 + \sigma\lceil N/p \rceil \approx \sigma\lceil N/p \rceil$ and $1 + \lceil N/p \rceil \approx \lceil N/p \rceil$ and the ration of Eq. (5.8) is reduced to $1/\sigma$.

The importance of the selective row decoder is now evident especially for large values of $\lceil N/p \rceil$. In addition to this important function, the selective row decoder plays another, equally important, role: it allows only "useful" data to enter the electronic computer thus reducing the probability for bottlenecks between the computer's electronic memory and its central processing unit.

We stated earlier that the electronic computer is overwhelmed by the incoming data from the optical memory at the limit

$$R_{\text{eff}} = Q \qquad (R_{\text{eff}} > Q) \tag{5.9}$$

Assuming a very fast SCSI bus running at 20 MB/s this limit becomes:

$$5 \times 2^{25} = rMN/T_r \Rightarrow T_r = 0.2rMN \times 2^{-25} \tag{5.10}$$

We examine the dependency of T_r for 25 different situations corresponding to all possible cases for $(MN)^{1/2} = 128, 256, 512, 1024, 2048$ and $r = 0.2, 0.4, 0.6, 0.8, 1$. Figure 5.7 summarizes our scenarios. In order to meet a fixed data rate, T_r increases with respect to MN and r as expected from Eq. (5.10) which is actually a proportionality relation scaled by Q. Depending on the parameters involved, T_r can vary from a few hundred microseconds to almost one second. Existing systems store data in pages with approximately 512^2 bits. For this page size, T_r could be 0.3–1.5 ms, depending on the code rate, before R_{eff} exceeds Q. This value is reasonable with respect to the various equipment timing parameters involved in an experimental optical memory system.

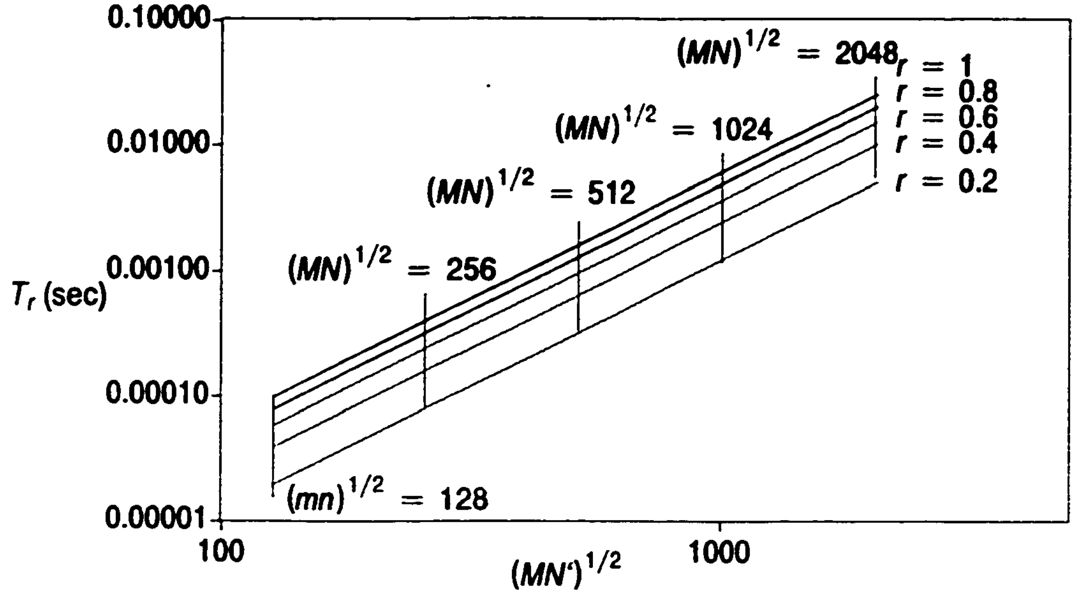

Figure 5.7. T_r constrained by a fast SCSI driver, for different code rates. Access times below the lines corresponding to various code rates will overwhelm a typical SCSI.

6

Analysis and Directions

The distinction between TO and TR interfaces, according to their complexity and functionality, suggests that their physical implementation will vary accordingly. In fact, we introduce the terms *smart photodetectors* and *intelligent photodetectors* to describe the complexity of the optoelectronics involved in each interface implementation. Smart photodetectors are processing elements with some resident logic which may (or may not) be interconnected to the logic of other elements. The complexity of interconnectivity between these elements determines their functionality. Intelligent photodetector arrays exhibit an interconnectivity pattern similar to smart arrays but they also accept external signals (i.e., signals that do not originate from the optical memory) to control additional operations which enhance their functionality. Figure 6.1 shows the functionality of smart and intelligent photodetectors with respect to the complexity of the interface implementation. We feel necessary to clarify here that the terms "smart" and "intelligent" are far from being based on any anthropomorphic attributes; in most contexts they have the same meaning. We arbitrarily assign more complexity (thus more processing power) to intelligent photodetectors than to smart photodetectors for the purpose of our discussion.

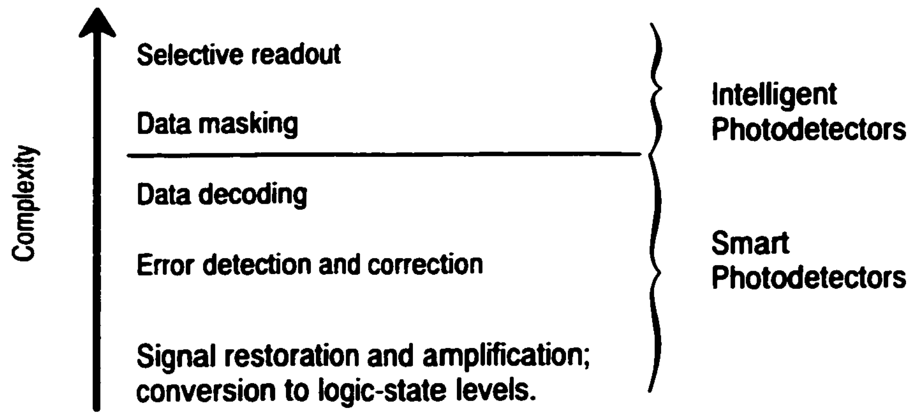

Figure 6.1. Descriptive task assignment for smart and intelligent photodetectors.

Smart photodetectors are used for TO interfacing. For example a TO interface could decode an eight-out-of-fifteen code. This would require a specific interconnection pattern between sets of 15 elements to provide the 8 bit output. At the same time a smart photodetector array can also perform error detection and correction. The difference between a smart photodetector array and an intelligent photodetector array is the degree of interconnectivity between the processing elements and the fact that an intelligent interface requires external control signals to perform its function (namely data reduction) as shown in Figure 6.2. The functionality of the static RAM-styled photodetector array with the selective row decoder which we discussed in the previous section is an example of an intelligent interface.

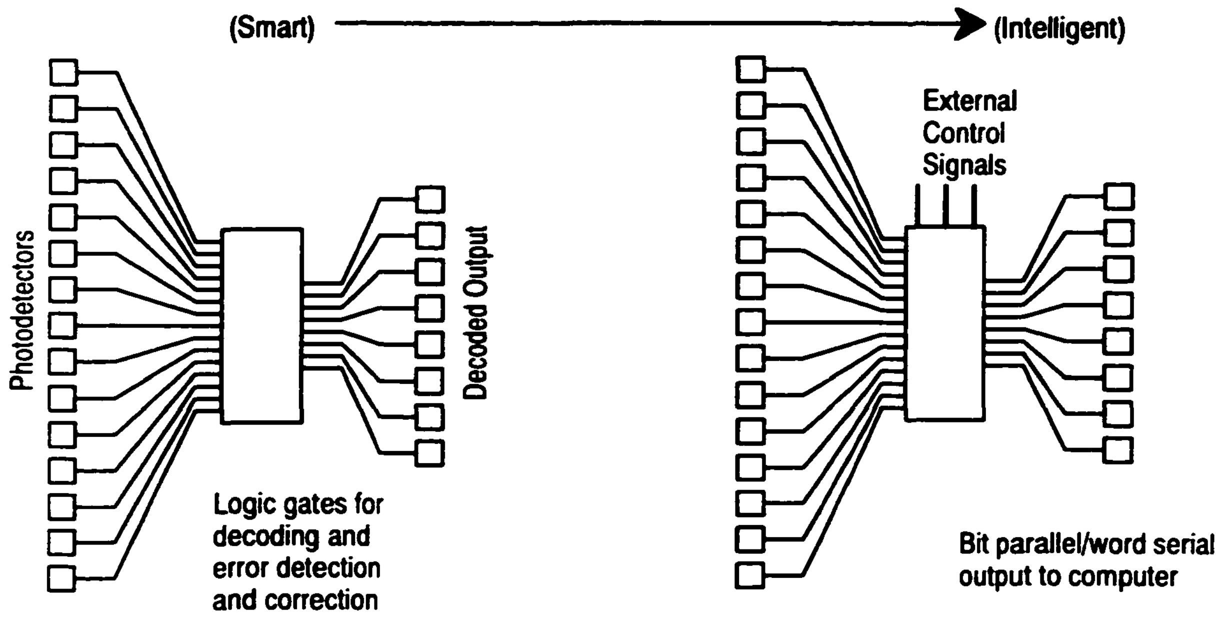

Figure 6.2. Smart and intelligent interfacing. In contrast to smart interfaces, intelligent interfaces can accept external control signals.

Figure 6.2 suggests that we can combine smart and intelligent photodetector arrays in order to achieve higher preprocessing power and therefore relax the workload of the electronic computer. These combinations yield two different scenarios: (a) transformed data are transferred from the TO component of the interface to the TR via optical interconnects and (b) transformed data are transferred from the TO component to the TR via electronic interconnects. The first scenario suggests that the two components are located on different stages and that multiple stages optically interconnected can be cascaded. The second scenario suggests (not exclusively) that the two components are integrated on the same stage and therefore cascading is not possible given the degree of planar interconnectivity required. Based on these scenarios, there is a choice to be made between single vs. multiple stage processing.

Single stage processing requires an extensive degree of planar interconnects in order to provide communications between the TO and TR components and the external control signals. The design rules associated with VLSI implementation of interconnects may effectively decrease the areal density of processing elements. In addition, having multiple interconnect traces close to each other increases the risk of cross talk and makes the entire system very sensitive to electromagnetic interference. A system like this will not exhibit an appreciable fault tolerance. However, it offers the advantage of a monolithic compact design, with all communication signals located on the chip. Additionally it requires relatively minimal effort for optomechanical alignment with respect to the optical memory's output only.

On the other hand, multiple stage processing allows us to build pipelined models with a large degree of parallelism at fine-grain levels. Their advantage is a higher processing element areal density because the complexity of the single stage processing element is now shared (and enhanced) by multiple stages. Higher processing element areal density means that it will be easier to image the output of the optical memory onto the photodetection grid because the geometric features of the two grids will have a higher degree of similarity than before, thus requiring simpler and less distorting imaging systems. In addition the risk of cross talk in significantly lower because of the employment of optical interconnects and for the same reason the system is less prone to electromagnetic interference. The disadvantages associated with the multiple stage approach include strict requirements for mechanical alignment, the need for hybrid integration of some components such as light sources and photodetectors (which favor the material system III–V), and logic elements which are silicon-based. The indirect bandgap nature of silicon-based devices prohibits us from fabricating efficient light sources and therefore necessitate the hybrid approach. A solution to this problem would be a monolithic fabrication approach using GaAs foundries.

In the next two sections we present two different approaches to the implementation of TR interfaces. The optoelectronic data filter is a hardware implementation of fundamental relational operations, employing vertical-cavity surface-emitting lasers and heterojunction phototransistors to form optical interconnects; the tomographic storage scheme is a bit-slice-based algorithmic approach to the same problem.

6.1 Optoelectronic Data Filter for Selection and Projection

The optoelectronic data filter (ODF) is a TR interface which has been designed for multiple stage implementation [Irakliotis, 95a]. It can selectively transfer data based on horizontal (selection) and vertical (projection) filtering. It performs projection by masking out the undesired data fields of a page by bringing it, along with a projection argument, on an AND optoelectronic gate array. Selection is performed by comparing the data page to a two-dimensional matrix formed by replicating the selection argument onto each row of the matrix. The comparison is completed optically in parallel as a bitwise XOR operation between the data bits and the argument bits. The results of all the XOR operations in each row are ORed together in order to determine matching rows. The selection operation is described by the following equation:

$$\begin{bmatrix} h_1 \\ h_2 \\ \vdots \\ h_m \end{bmatrix} = \begin{bmatrix} d_{11} \oplus s_{11} + d_{12} \oplus s_{12} + \ldots + d_{1n} \oplus s_{1n} \\ d_{21} \oplus s_{21} + d_{22} \oplus s_{22} + \ldots + d_{2n} \oplus s_{2n} \\ \vdots \\ d_{m1} \oplus s_{m1} + d_{m2} \oplus s_{m2} + \ldots + d_{mn} \oplus s_{mn} \end{bmatrix} \tag{6.1}$$

where h_j are the elements of a vector identifying matches (if h_i=0, there is a match) and $d_{ij} = b_{ij} \cdot f_{ij}$. Here, f_{ij} are the bits of the selection mask pattern, b_{ij} are the data bits received from the optical memory, and s_{ij} are the bits of the selection argument replicated to the two-dimensional array, (therefore $s_{1j} = s_{2j} = \ldots = s_{nj}, j = 1, 2, \ldots, n$). The presence of a selection mask allows us to perform selection only on a specified attribute. The functionality of the selection mask is to alleviate the need for "don't care" values which are not available in the strictly binary encoding scheme employed by this architecture.

Figure 6.3 shows a system view of the ODF. In this picture several optical components such as lenses and microlens arrays have been omitted for clarity. Data enter the system from the lower right beamsplitter and are split in two directions. One copy of the data page propagates towards the selection module and the other is directed to the projection module. Data in the projection module are ANDed with the projection mask and the output is imaged on the optoelectronic RAM array. The optoelectronic RAM is an ordinary RAM array with an array of photodetectors integrated on top of the memory cells. This allows us to write the entire RAM array in a single step by imaging a binary pattern on the photodetector array.

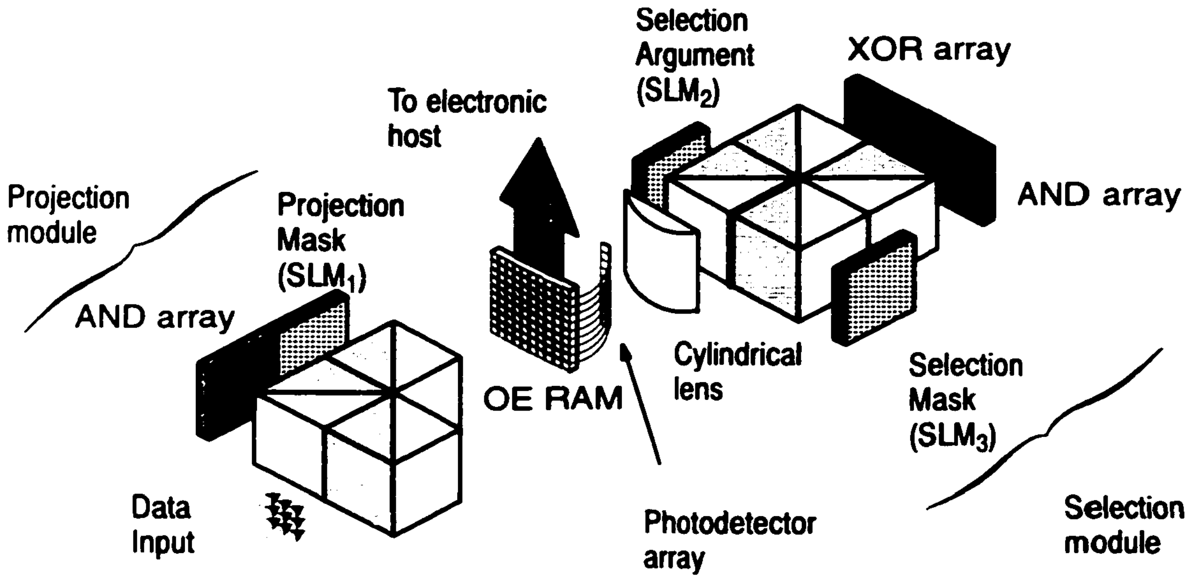

Figure 6.3. A system view of the optoelectronic data filter prototype.

Data entering the selection module will be ANDed with the selection mask which marks the don't care bits as 0-value bits. Then this masked page will be XORed with the two-dimensional replica of the selection argument. The rows of the XOR array output are focused by a cylindrical lens along a linear photodetector array. The cylindrical lens functions as an optical OR gate. As shown in Figure 6.4, light beams from all the elements in a row are focused by the cylindrical lens onto a point. The light intensity at that point corresponds to the sum of the light intensities of the row elements. Assuming a binary encoding, the focal point represents the OR of the row elements.

The photodetector array determines which rows of the optoelectronic RAM will be readout and subsequently transferred to the electronic host computer. The combination of the photodetector array and the optoelectronic RAM allows us to transfer only those portions of the projection output which match with the selection argument, thus reducing the effective transfer rate between the parallel optical memory and the electronic host. Figure 6.5 shows the transformations of a data page through its route in the selection module.

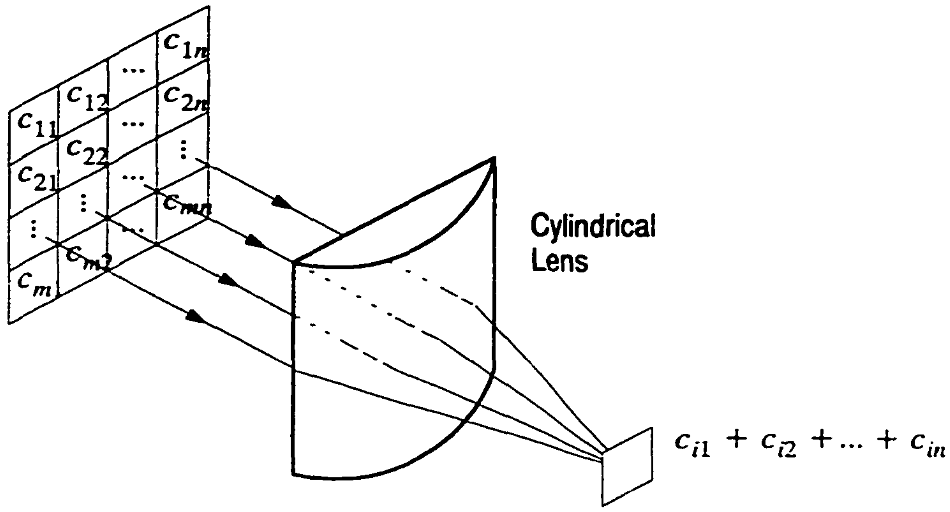

Figure 6.4. Cylindrical lens acting as an OR gate.

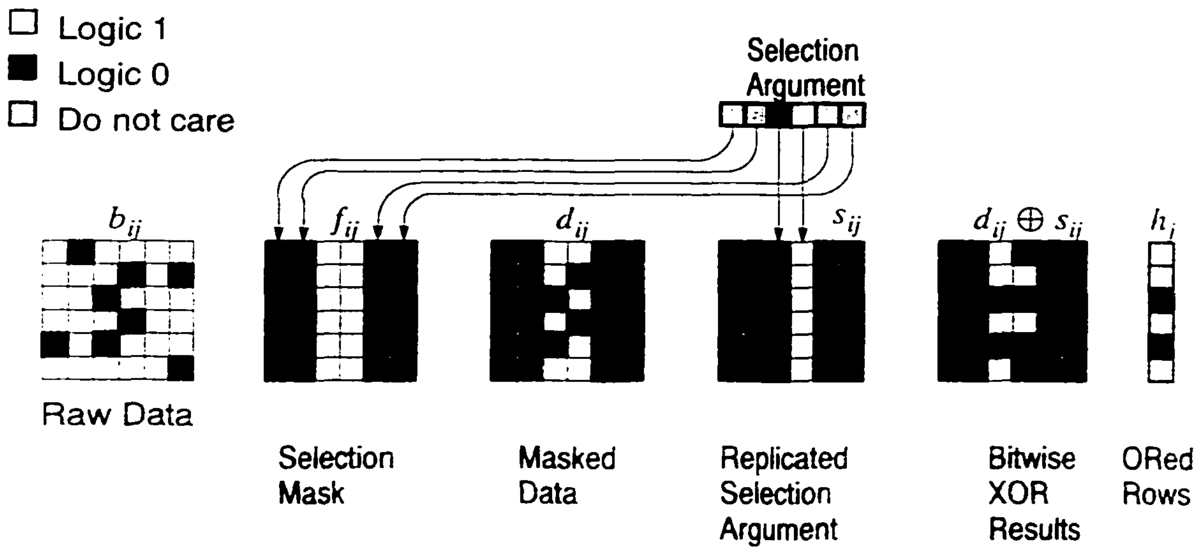

Figure 6.5. Block-oriented associative processing for the optoelectronic filter. "Do not care" bits are converted to logic-0 bits because the selection mask masks them out.

Logic operations are carried out on two-dimensional arrays of optoelectronic logic gates, also known as smart pixels. Smart pixels can be implemented as combinations of detectors, emitters, and electronic logic. We have fabricated our smart pixels using heterostructure phototransistors (HPT) as detectors and vertical cavity surface emitting lasers (VCSEL's) as emitters [Snyder, 94a; Hayes, 97a].

6.2 Three-dimensional Tomographic Storage Scheme

The three-dimensional tomographic storage scheme (TSS) has been proposed as an alternative data format suitable for optoelectronic implementation of database operations [Irakliotis, 94a]. At the same time the TSS can be used for other applications where massive parallelism at the bit-slice level is desired (e.g. image processing) The TSS exploits the ability of volume optical memories to store and retrieve two-dimensional pages of data, each page at a single logical and physical address. Its name has been derived by the Greek words *tomé* and *graphé* which mean view of a slice. In the TSS the data pages correspond to bit slices of all the tuples in the relation as shown in Figure 6.6(a).

6.2.1 Selection

In the case of selection we need to identify the tuples which match the selection criterion. An example of a selection is shown in Figure 6.6(b). The tuples are identified by their Cartesian coordinates on the tuple plane. If the selection criterion s is:

$$\left[s_{a+1}\, s_{a+2} \cdots s_{a+q}\right];\ 1 \le a+1 \le a+q \le L \qquad (6.2)$$

where s_i are bit values, L is the length of the tuple and q is the length of the selection criterion, then the qualified tuples are marked by the table:

$$M(i,j) = \prod_{k=1}^{q}\left[\overline{S_{a+k}(i,j) \ominus B_{a+k}(i,j)}\right] = \prod_{k=1}^{q}\left[\overline{S_{a+k}(i,j)} \ominus B_{a+k}(i,j)\right] \qquad (6.3)$$

where S is a matrix with all its elements set to the current value of the selection argument bit, i.e.

$$S_l(i,j) = s_l; \quad 1 \le i \le m; \ 1 \le j \le n; \quad a+1 \le l \le a+q \tag{6.4}$$

and $B_k(i,j)$ is the two-dimensional data bit slice of k-th order.

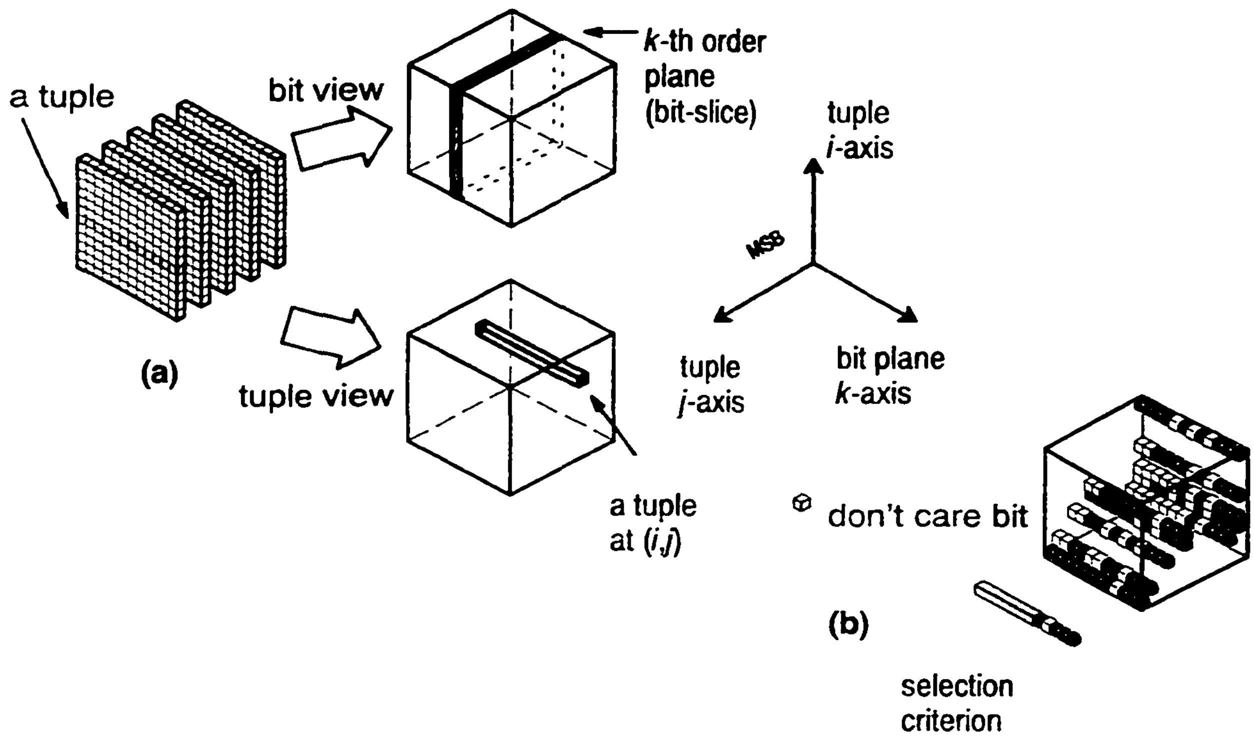

Figure 6.6. Tomographic storage scheme. (a) Tomographic mapping, (b) selection operation in the tomographic scheme.

TSS can be used to perform selections in a bit-serial word-parallel associative processing scheme. The sequence of steps in the tomographic selection is illustrated in Figure 6.7. The bits of the selection argument are replicated in two-dimensional arrays of equal size to the bit slices of the relation, one at a time. Then, the two-dimensional selection bit replica is bitwise XORed with the corresponding bit slice from the optical memory. Next, the result of the XOR operation is ANDed with the result of the same operation from the previous slice. To denote the sequential time of these AND operations, Equation (6.3) is written as:

$$M_k(i,j) = M_{k-1}(i,j) \cdot \left[\overline{S_{a+k}(i,j) \oplus B_{a+k}(i,j)}\right]; \quad 0 < k \le q \tag{6.5}$$

with the initial condition $M_0(i,j) = 1, \forall i,j$. From Equations (6.3) and (6.5) we can see that the time complexity of the tomographic operation is $O(q)$. After q+1 steps, the matrix $M_q(i,j)$ identifies the Cartesian coordinates (i,j) of those tuples matching the selection argument. For the example of Figure 6.7, these tuples are located at points (1,1) and (4,3).

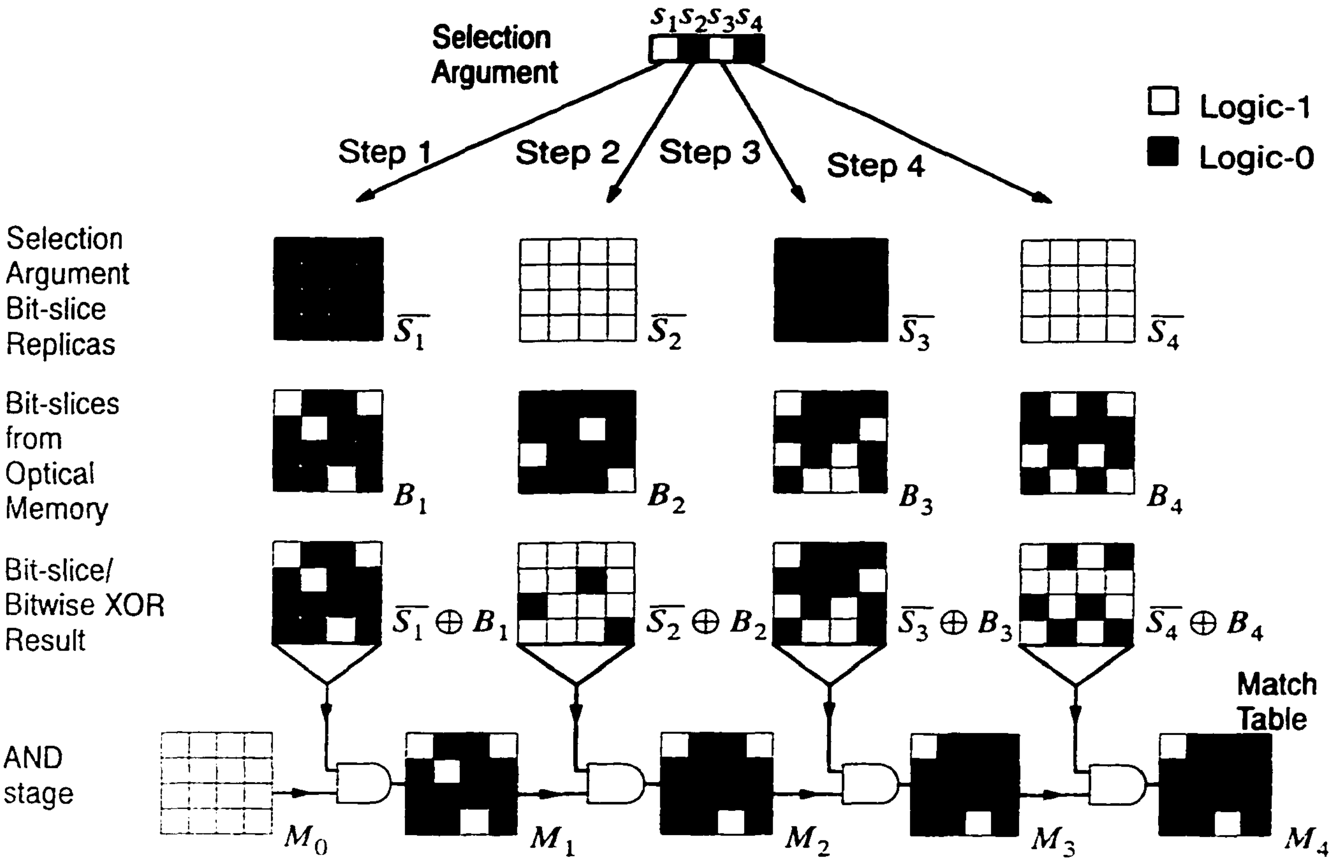

Figure 6.7. Selection operation under the tomographic storage protocol.

From Equation (6.5) it becomes apparent that we need to hold the results of an XOR operation until a similar operation at the next bit-slice level is completed. Using electronics for the processing, this can be done in a straight forward way by employing flip-flops in a buffer array.

Exact matches can be accomplished with the tomographic protocol. When a "don't care" value is encountered during tomographic selection, it is possible to ignore the corresponding bit-slice and skip the XOR operation at that time step.

The TSS appears to be highly suitable for implementation in combination with a two-photon three-dimensional memory because we can address its contents both in a word-parallel/bit-serial or word-serial/bit-parallel fashion relatively easy.

6.2.2 Join

The join operation is significantly more time consuming than selection and projection. Its computational complexity is $O(C_1 \times C_2)$ where C_1 and C_2 are the cardinalities of the two operand relations R_1 and R_2. Relations are joined over a common attribute. Let us assume here that the join attribute has a length of q bits and that $C_1 < C_2$. In this case we need to compare only q bit planes with the join attribute bits and this process must be repeated C_1 times until we exhaust all the tuples of the first relation. Therefore, the time complexity becomes $O(C_1 \times q)$.

In a compact mathematical formulation, the tomographic join operation produces a set of join matrices M_t which identify the tuples from both relations to be joined as follows:

$$1 \le t \le C_1: \quad M_t(i,j) = 1 \Leftrightarrow \text{tuple } (i,j) \text{ of } R_2 \text{ join with tuple } (\mu,\nu) \text{ of } R_1$$

where $\mu = \lceil t/N_1 \rceil$ and $\nu = t - (\mu - 1)N_1$. N_1 is the number of tuples per row in the bit-slice view of R_1. Matrix M_t is formed by a series of AND functions described by:

$$M_t(i,j) = \prod_k f_{ijkt}$$

Here. f_{ijkt} are the matching patterns for the k-th bit slice of the q bit slices of the join attribute. The formal definition of the f pattern is:

$$f_{ijkt} = B_{1,k}(\mu,\nu) \cdot B_{2,k} \tag{8.1}$$

where $B_{1,k}(\mu,\nu)$ is the k-th significant bit of the tuple located at (μ,ν) in relation R_1 and $B_{2,k}$ is the k-th bit plane of the relation R_2.

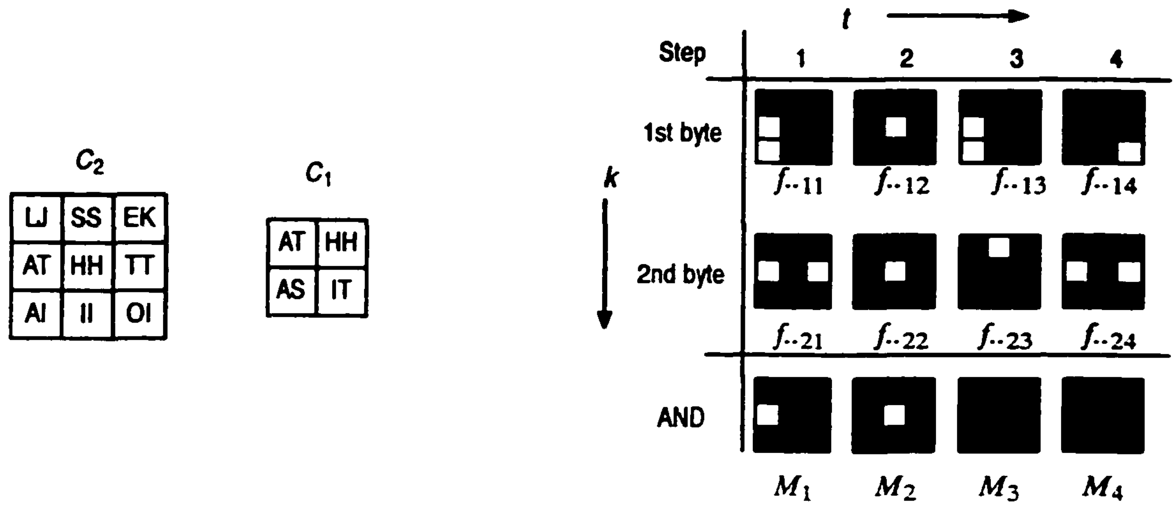

Figure 6.8. Outline of the join operation in the tomographic storage scheme.

Because we can have only one f term at a time, it takes a total of $t \times k$ steps to complete the process. Since $1 \le k \le q$ and $1 \le t \le C_1$, the time complexity is $O(C_1 \times q)$ as mentioned earlier. The k and t indices in f map the outer relation (in this case R_1) to the inner relation (R_2) as follows: at the first time step the first byte of the first tuple of C_1 is broadcasted to C_2 and the f_{ij11} pattern is recorded as shown in Figure 6.8 where we use two relations with cardinalities C_1=4 and C_2=9. Note that in this example we deal with bytes rather than with bits to simplify the illustration. At the second time step the second byte of the first tuple of C_1 is broadcasted to C_2 and the f_{ij21} pattern is recorded. Then patterns f_{ij11} and f_{ij21} are ANDed together to obtain M_1 which indicates the join locations, if any. This process continues with the next tuple from C_1 until all the tuples are scanned. At the end of the scan of each tuple from C_1 the matching table M_t is obtained. Using this table we can reconstruct the qualified tuples from the tomographic scheme into the relational scheme assuming the table has at least one non-zero element.

6.2.3 Sorting algorithm

Sorting can be also performed using the tomographic scheme. We choose a combination of the odd/even transposition and the radix sort algorithms. The process is initiated by looking at the most significant bit slice and trying to partition it into two groups: the 0-group at the top and the 1-group at the bottom.

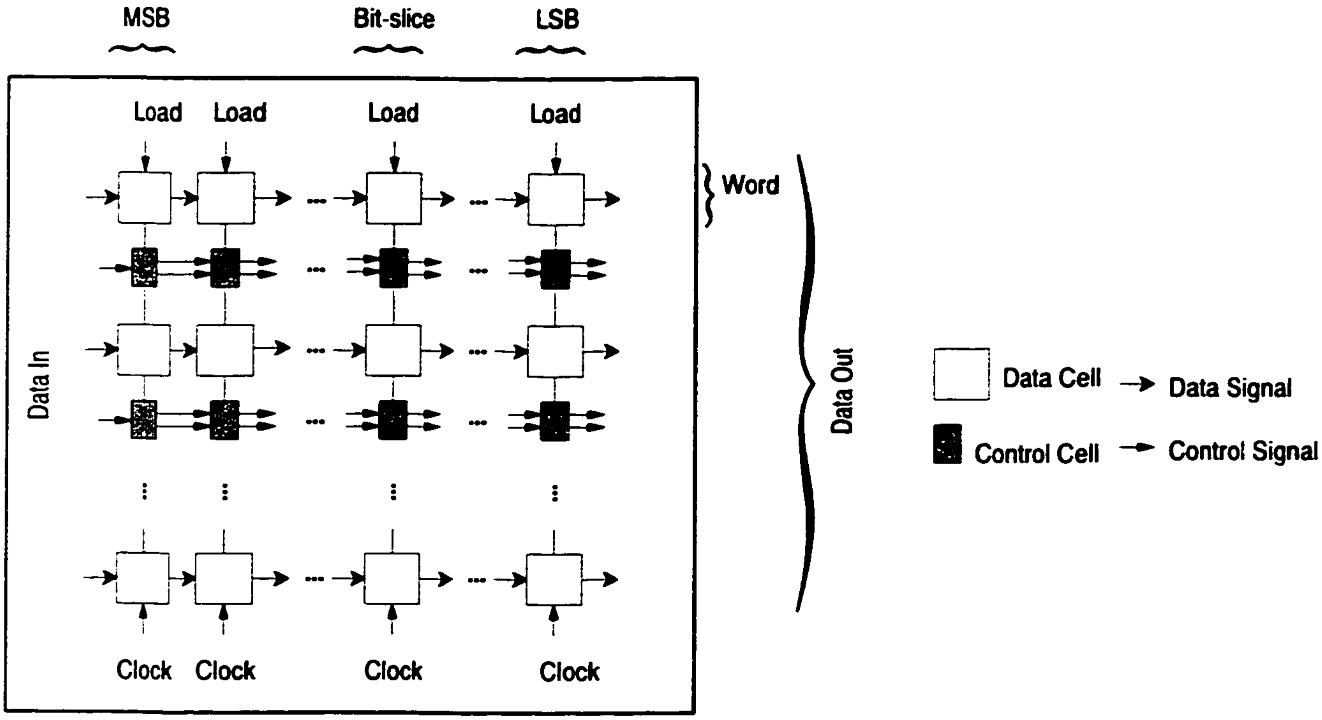

Figure 6.9. Outline of the tomographic sorting in two-dimensions

Every pair of adjacent bits that are not equal allows a decision on the relative magnitude of two words to be reached. If the two words must be exchanged, the process is completed in a pipelined fashion with an unconditional exchange control signal gradually forcing all the lower significance bits of the two words to exchange positions. If a decision cannot be reached at the MSB slice because the two bits under consideration are equal, the decision is deferred to the next most significant bit slice by sending a conditional exchange control signal to the next stage. This will force the next control cell to perform a bit comparison to determine the relative magnitude of the two words. While an unconditional or conditional exchange signal ripples through all the stages of the network, a new set of tests can be initiated at the MSB slice. The outcome of these tests and the subsequent execution of any action required will not affect the completion of the previously initiated exchanges that may still continue to occur towards the final stages of the pipeline [Beyette, 95a].

Figure 6.9 shows the processing array for the sorting algorithm. The complexity of the processing elements here is higher compared to that of the processing elements for the relational operations discussed earlier.

6.2.4 Tuple retrieval

As we mentioned earlier, it is important to be able to reconstruct the flat files from the tomographic storage space. The relational database operations we are considering produce a match-table $M(i,j)$ which indicates whether a tuple at the location i, j has qualified for further processing. Because $M(i,j)$ can be either 1 or 0, the number of qualified tuples is given by $Z = \sum_{i=1}^{N} \sum_{j=1}^{N} M(i,j)$. Therefore Z bits must be read out of each bit plane and then rearranged in a tuple format. From each bit plane we read only the pixels indicated by 1s in $M(i,j)$. It is possible to build a two-dimensional photodetector array which will compress the empty pixels from the image and produce a serial output with the qualified pixels only. By repeating this process for each bit plane we obtain all the bits corresponding to the qualified tuples. Tuple-mode access is then a matter of properly partitioning that ensemble of bits in a random access memory.

6.3 Toward an Optoelectronic Digital Associative Processor

In associative processing, an element from a set of information is identified based on input which is a partial image of the desired element, as illustrated in Figure 6.10. Associative access is based on the content, not on the address (logical of physical) of the desired information. Access to the desired information can be achieved in various degrees of parallelism. The optimal case is where the contents of the entire memory are searched and desired data are accessed in a single step. Volume holographic memories are capable of such full and parallel associative processing. In lower degrees of parallelism, pages of data can be scanned in a single step. The optoelectronic digital associative processor (OEDAP) that we discuss here is capable of scanning a two-dimensional page of data in a single step and identify the location of the desired data based on partial input.

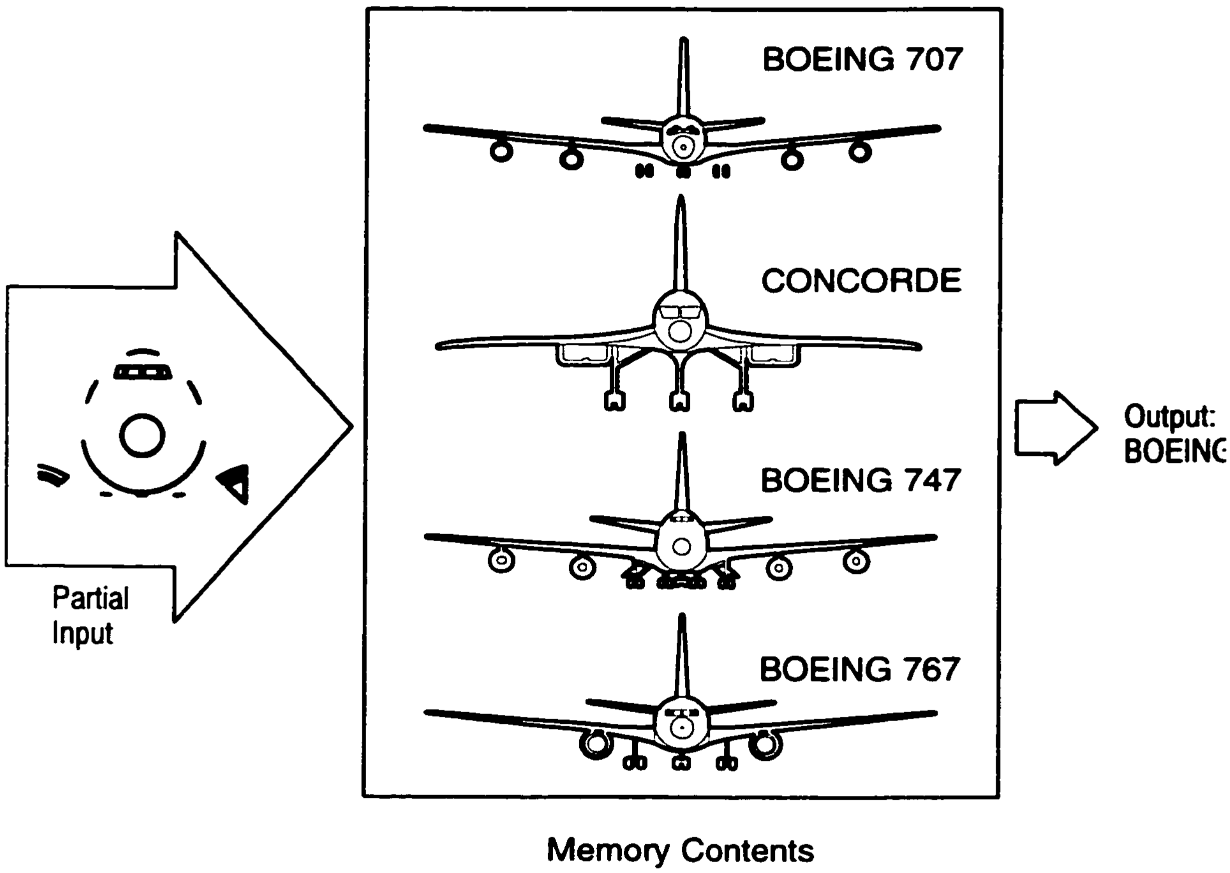

Figure 6.10. Illustration of hetero-associative processing. .

bold The theory for page-oriented associative processing is rather simple. The data page is represented by a two-dimensional matrix **d**:

$$\mathbf{d} = \begin{bmatrix} d_{11} & d_{12} & \cdots & d_{1n} \\ d_{21} & d_{22} & \cdots & d_{2n} \\ \vdots & \vdots & \vdots & \vdots \\ d_{m1} & d_{m2} & \cdots & d_{mn} \end{bmatrix}, \; d_{ij} \in \{0,1\} \tag{8.2}$$

and the search argument by a row vector **t**:

$$\mathbf{t} = \begin{bmatrix} s_1 \; s_2 \; \cdots \; s_n \end{bmatrix}, \; s_j \in \{0,1,X\} \tag{8.3}$$

where X represents the "do not care" bit. A search table **s** is made by vertically replicating the search vectors as many times as the number of rows in **d**:

$$\mathbf{s} = \left.\begin{bmatrix} \mathbf{t} \\ \mathbf{t} \\ \vdots \\ \mathbf{t} \end{bmatrix}\right\} m \text{ rows} \tag{8.4}$$

Using matrices **d** and **s** we can produce a column vector which identifies the rows of matrix **d** containing the desired data. This locator vector **h** is obtained by:

$$\mathbf{h} = \begin{bmatrix} \sum_{i=1}^{n} d_{1i} \oplus s_i \\ \sum_{i=1}^{n} d_{2i} \oplus s_i \\ \vdots \\ \sum_{i=1}^{n} d_{mi} \oplus s_i \end{bmatrix} \tag{8.5}$$

where the operator $\oplus$ corresponds to the XOR operation.

The summations in Equation (8.5) are taken over the elements of the rows of the matrix $\mathbf{d} \hat{\ominus} \mathbf{s}$: here the operator $\hat{\ominus}$ denotes a bitwise XOR operation.

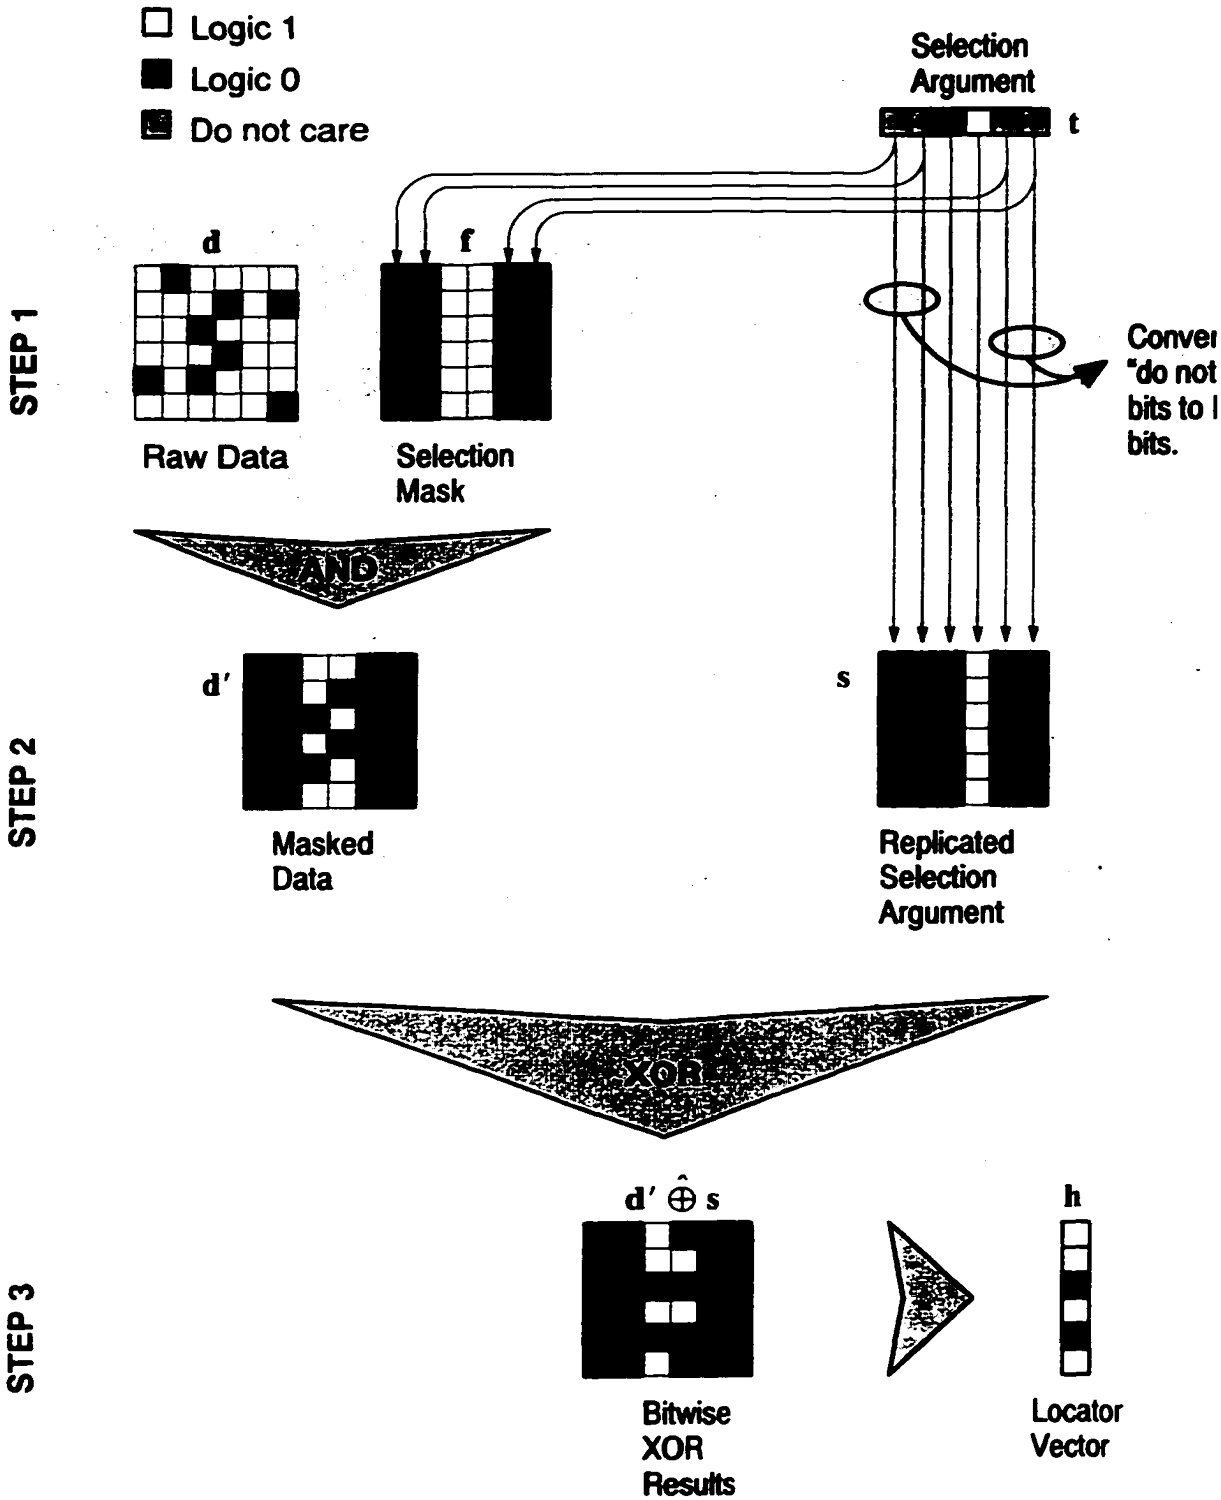

Figure 6.11. Visualization of the process described by Equations (8.2)-(8.5).

The process described by Equations (8.2)-(8.5) is illustrated in Figure 6.11 with a few additional steps that have been introduced in order to eliminate the "do not care" bits. To keep our algorithm strictly binary, we use the following technique to reduce the need for "do not care" bits. First, we create a two-dimensional matrix **f** such as that:

$$f_{ij} = \begin{cases} 0, & \text{if } t_j = X \\ 1, & \text{if } t_j \neq X \end{cases} \tag{8.6}$$

which we then bitwise AND with matrix **d**. As a result, we obtain matrix **d'**:

$$d'_{ij} = \begin{cases} 0, & \text{if } t_j = X \\ d_{ij}, & \text{if } t_j \neq X \end{cases} \tag{8.7}$$

Now, in creating the matrix **s** we can replace the "do not care" bits of **t** with logic-0 bits and then perform the bitwise XOR operation $\mathbf{d}' \hat{\oplus}\, \mathbf{s}$.

6.3.1 Implementation

To perform the bitwise XOR operation between **d'** and **s** we must bring their images together on the same plane. Since we are interested in an optoelectronic architecture, this can be done quite easily using conventional optical components such as a beamsplitter, as shown in Figure 6.12. Now that all the signals are on the same plane we can look for a planar implementation of an optoelectronic XOR array.

A direct implementation of the XOR function would require the present of each bit and its complements, thus adding unnecessary complexity to the system.

To avoid this complexity, we used the optoelectronic XOR design shown in Figure 6.13. When both signals *A* and *B* are absent, there will be no activity at all and the vertical-cavity surface-emitting laser (VCSEL) will not emit. If both signals are present, current will flow through the path parallel to the VCSEL (i.e., through the heterostructure phototransistors transistors (HPT) HPT3 and HPT4), and again the VCSEL will not emit. However, when one of the two signals (either *A* or *B*) is present, current will flow through the VCSEL because the path through HPT3 and HPT4 is open, and therefore the VCSEL will emit. In other words, the circuit of Figure 6.13 will perform the XOR operation without the negated version of the input signals.

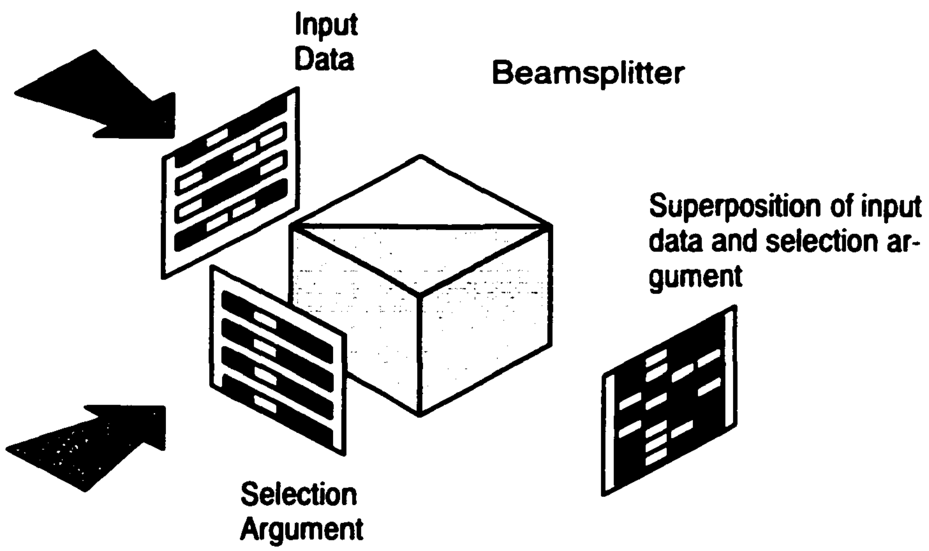

Figure 6.12. Imaging the input data and the selection argument together.

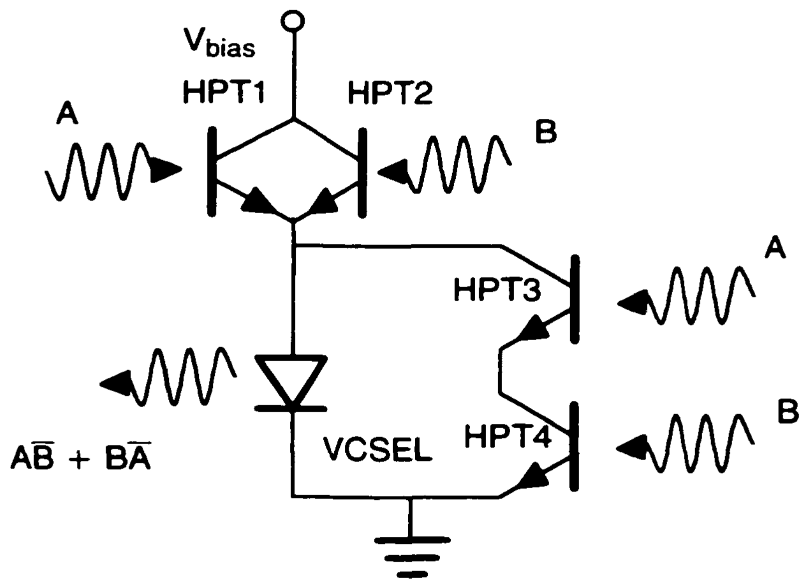

Figure 6.13. Design of the optoelectronic XOR gate with using heterostructure phototransistors (HPT) and vertical-cavity surface-emitting laser (VCSEL) [Snyder, 94a].

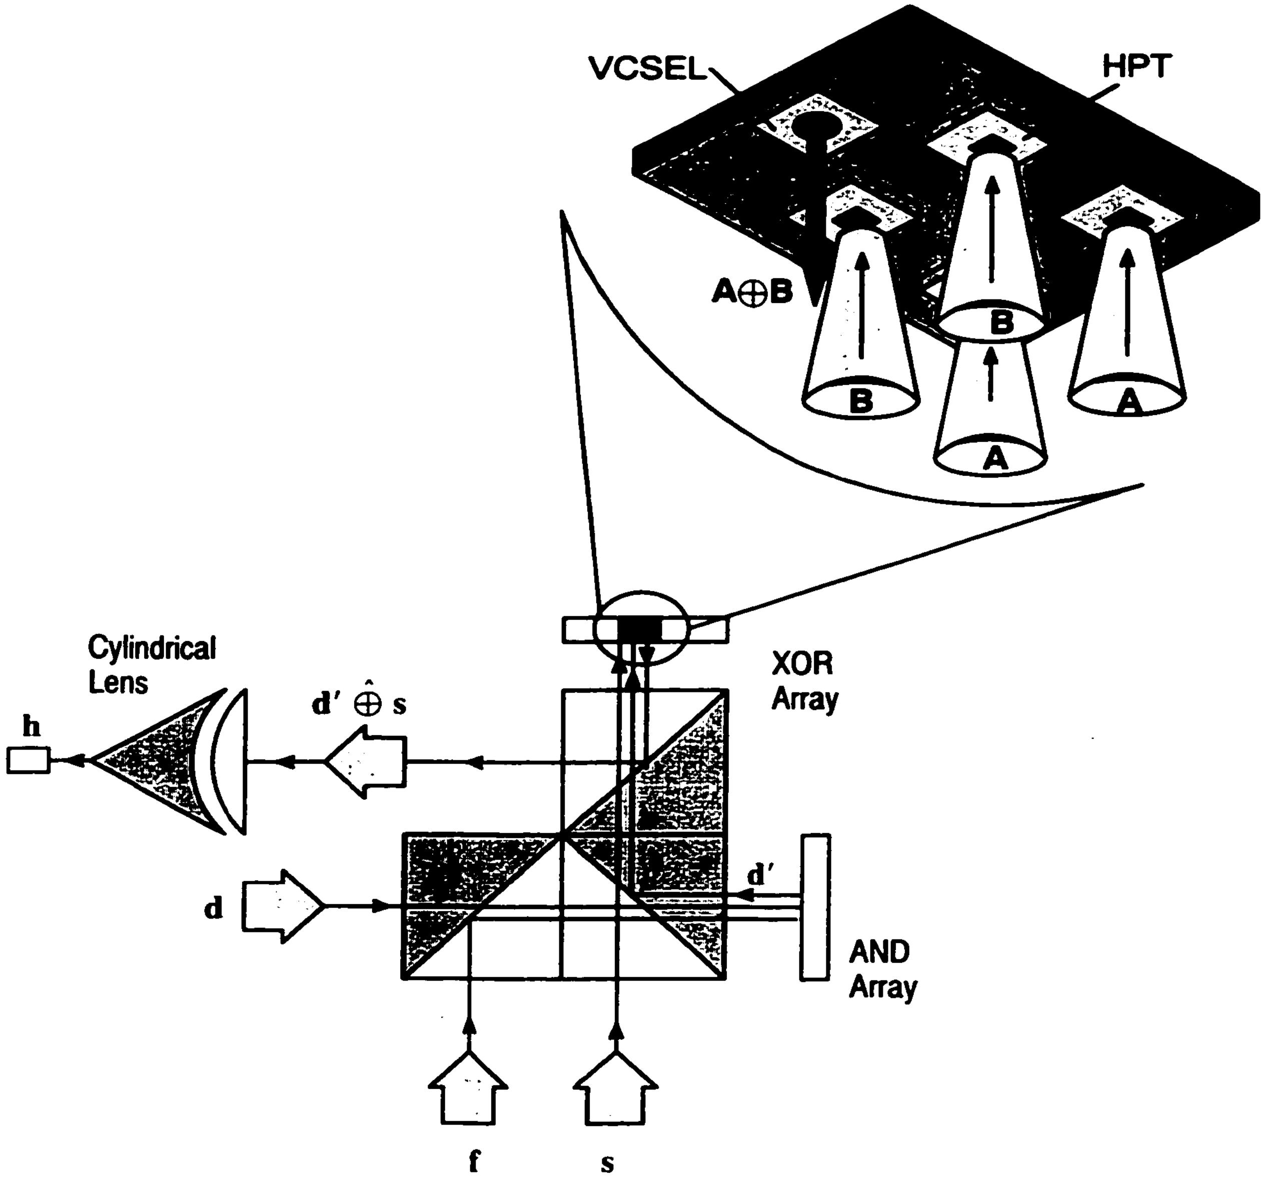

Figure 6.14. Complete schematic of the optoelectronic digital associative processor and a detailed representation of a smart pixel.

The integration of VCSEL's and HPT's is a convenient way to perform optoelectronic logic functions. Being the optical output device, the VCSEL will perform the data transportation between stages. It was chosen because it is a source of coherent, collimated light that is easily coupled to other devices using simple optics. In addition, large two-dimensional arrays of VCSEL's are feasible and VCSEL structures fabricated by implantation of protons provide a planar surface which allows the fabrication of the logic circuit.

The HPT's were chosen as the detectors because they have high-sensitivity, with responsivities that can reach several hundred amperes of output current per watt of input optical power. In addition, the spectral sensitivity range of the HPT is quite broad when compared to other detector technologies such as quantum-well type devices. HPTs are also excellent current sources and are ideal for driving the VCSEL. Finally, the fabrication of HPTs is straight-forward, requiring a minimum of additional processing steps beyond the processing of the VCSEL. A detailed discussion on these issues can be found in [Snyder, 94a].

Figure 6.14 shows the complete schematic of the proposed optoelectronic digital associative processor (OEDAP) but a simpler version was implemented and demonstrated. The simpler version is shown in Figure 6.15. In the demonstration version we use only the XOR array. The function of the AND array is replaced by a shadow mask which produces the selection matrix s. Our objective was to demonstrate the core function of the OEDAP, namely the look-up table operation.

We implemented a look-up table with a capacity of 4×4 bits per page. To simulate the search matrix s we used a programmable array of VCSEL's, interfaced to a computer. The simulated data page arriving from an optical memory was a fixed pattern on a gold-plated mask which was used as a screen for a beamlet array produced by a binary phase grating. The image of the screened beamlet array is shown in Figure 6.16. The optoelectronic XOR gates we used require two copies of each signal as input, therefore the following encoding was employed for each bit value:

$$\begin{aligned} 0 &\mapsto 00 \\ 1 &\mapsto 11 \end{aligned} \tag{8.8}$$

and a total of 64 signals were mixed at the XOR plane.

With reference to Figure 6.15, input data are interleaved through the first beamsplitter (BS1) with the selection argument which is produced by the computer-driven VCSEL array thus forming a new array of data as shown in Figure 6.17. Following this interleaving, the new array of input and search bits is imaged onto the XOR array through the second beam-splitter (BS2). The results of the bitwise XOR operations are then produced on the VCSEL array of the XOR gate array and the emerging beams are focused by the cylindrical lens — which acts as an OR gate — onto the CCD camera. The image formed corresponds to the locator-vector **h** of Equation (8.5).

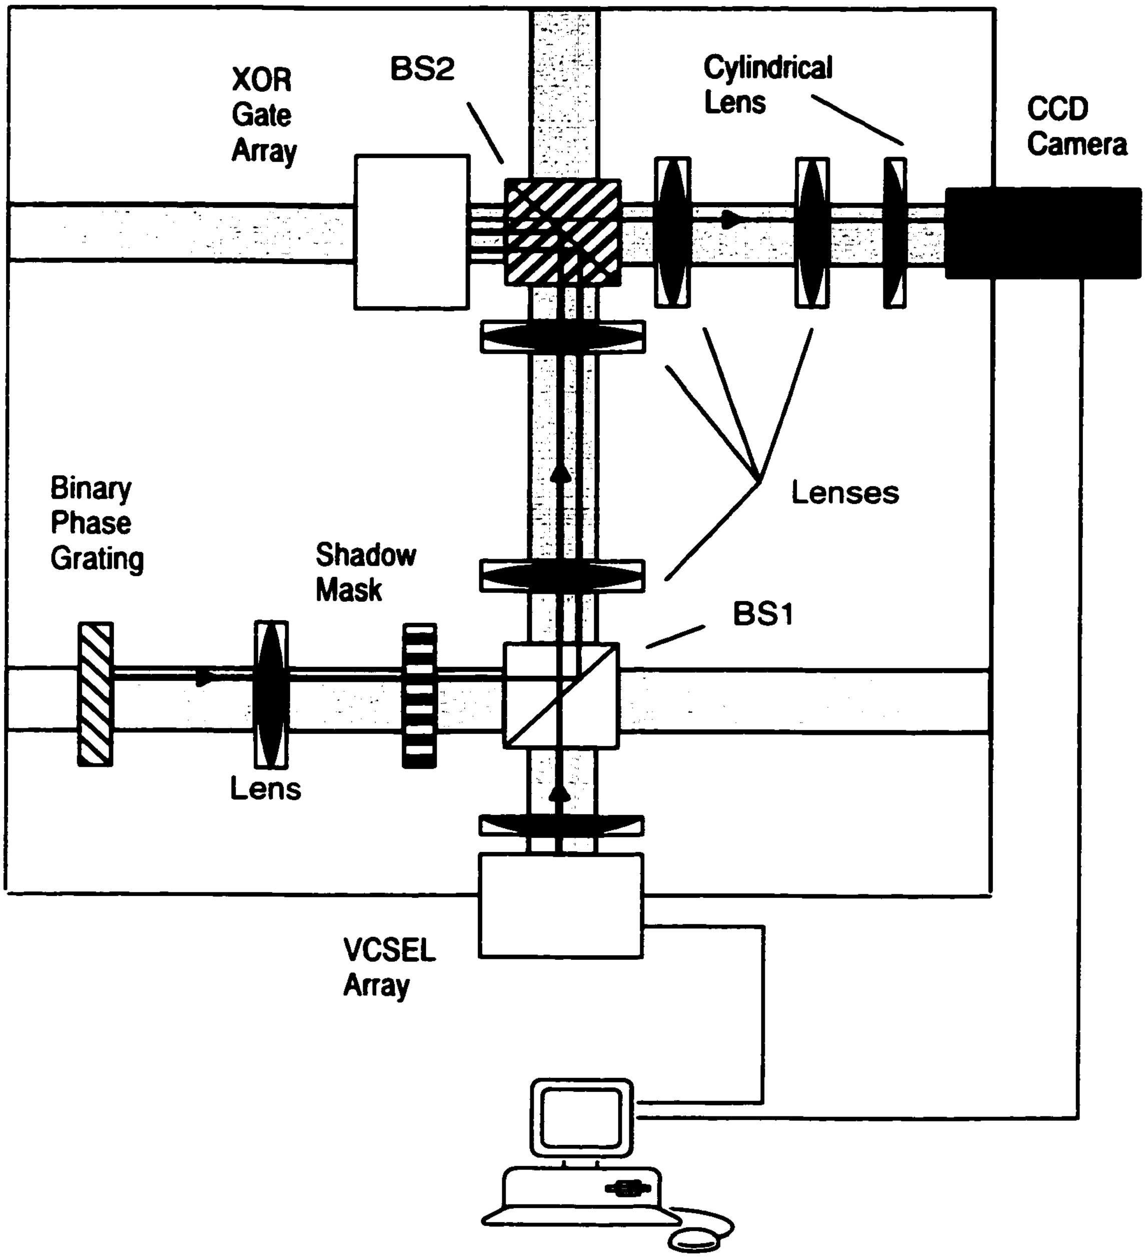

Figure 6.15. Overall architecture of the demonstrated optoelectronic digital associative processor.

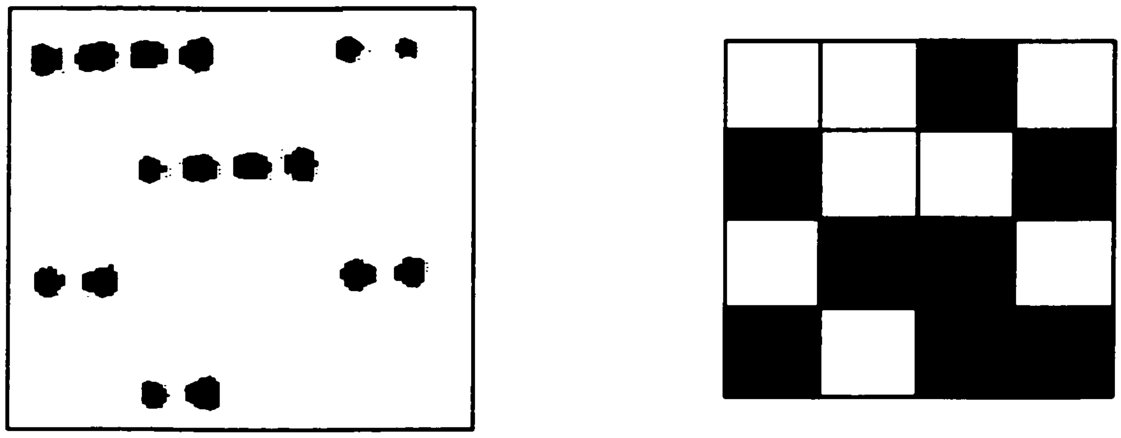

(a) (b)

Figure 6.16. The search matrix s used in the demonstration of the look-up table. (a) actual image; (b) corresponding logic pattern after decoding of the dual bit pattern.

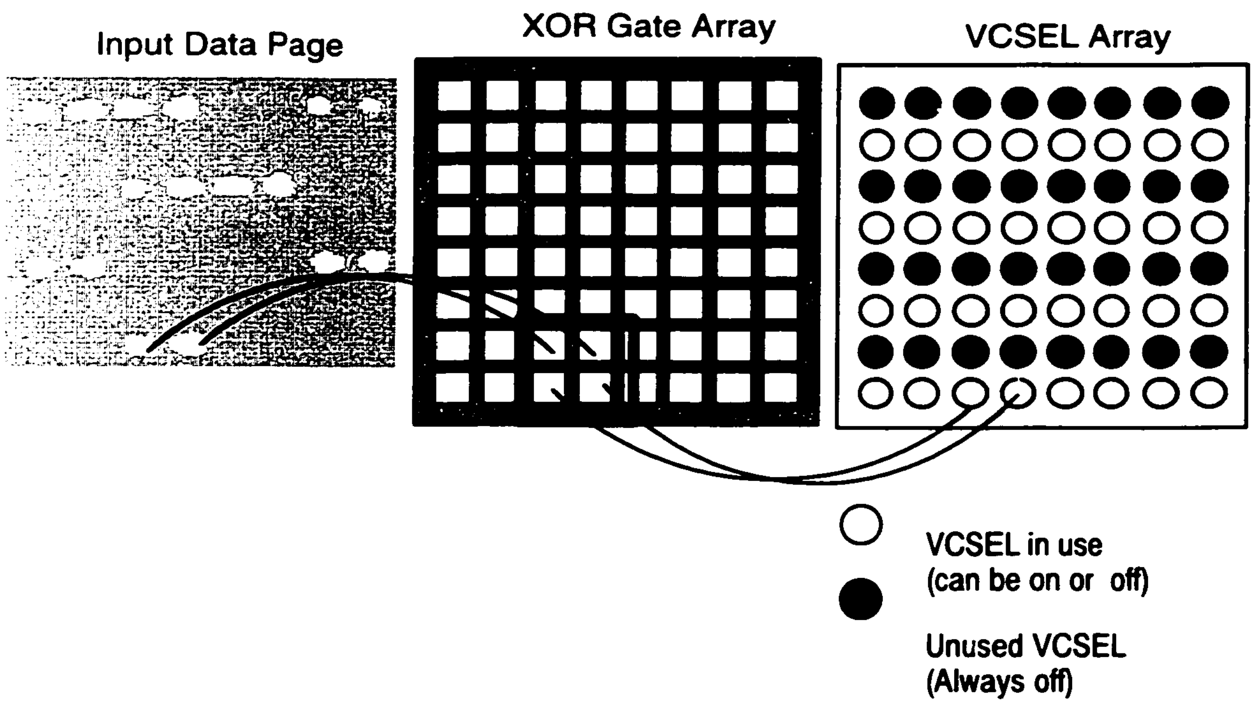

Figure 6.17. Interleaving of the input data page and of the search matrix s formed on the programmable VCSEL array.

6.3.2 Experimental results

For the experimental demonstration each of the four rows of the data pattern was used as a search argument. In addition, we used a search argument that does not exist in the data page. The outline of the experiments is shown in Figure 6.18. The results from each search are shown in Figure 6.18 (a)–(e). The patterns shown are the outputs of the XOR gate array as seen by the CCD camera through the cylindrical lens. A match is identified by a blank spot indicating that all the VCSELs in the corresponding row remained disabled. Figure 6.18 (e) shows the result obtained when we used the string `1010` as a search argument. This string of data is not present in the data page, therefore at least one VCSEL per row at the bitwise XOR output matrix will be active.

Each locator vector in Figure 6.19 is imaged on an area of approximately 40×250 pixels. By measuring the intensity of each pixel along the center line of this pixel matrix, we get an estimate of the contrast ratio (CR) performance of our system. The intensity plots are shown in Figure 6.20. Based on peak intensities we calculated the CR to be in the range of 140:1–210:1. The lowest CR corresponds to a single active VCSEL in a row, and the highest CR corresponds to all four VCSELs in a row being active. In the evaluation of CR using the following formula:

$$\mathrm{CR} = \frac{\text{peak intensity}}{\min\{\text{pixel intensity}\}} \tag{8.9}$$

the minimum pixel intensity was taken to be always 1 due to the dynamic range of the graphics file format used in the experiment (8-bit GIF).

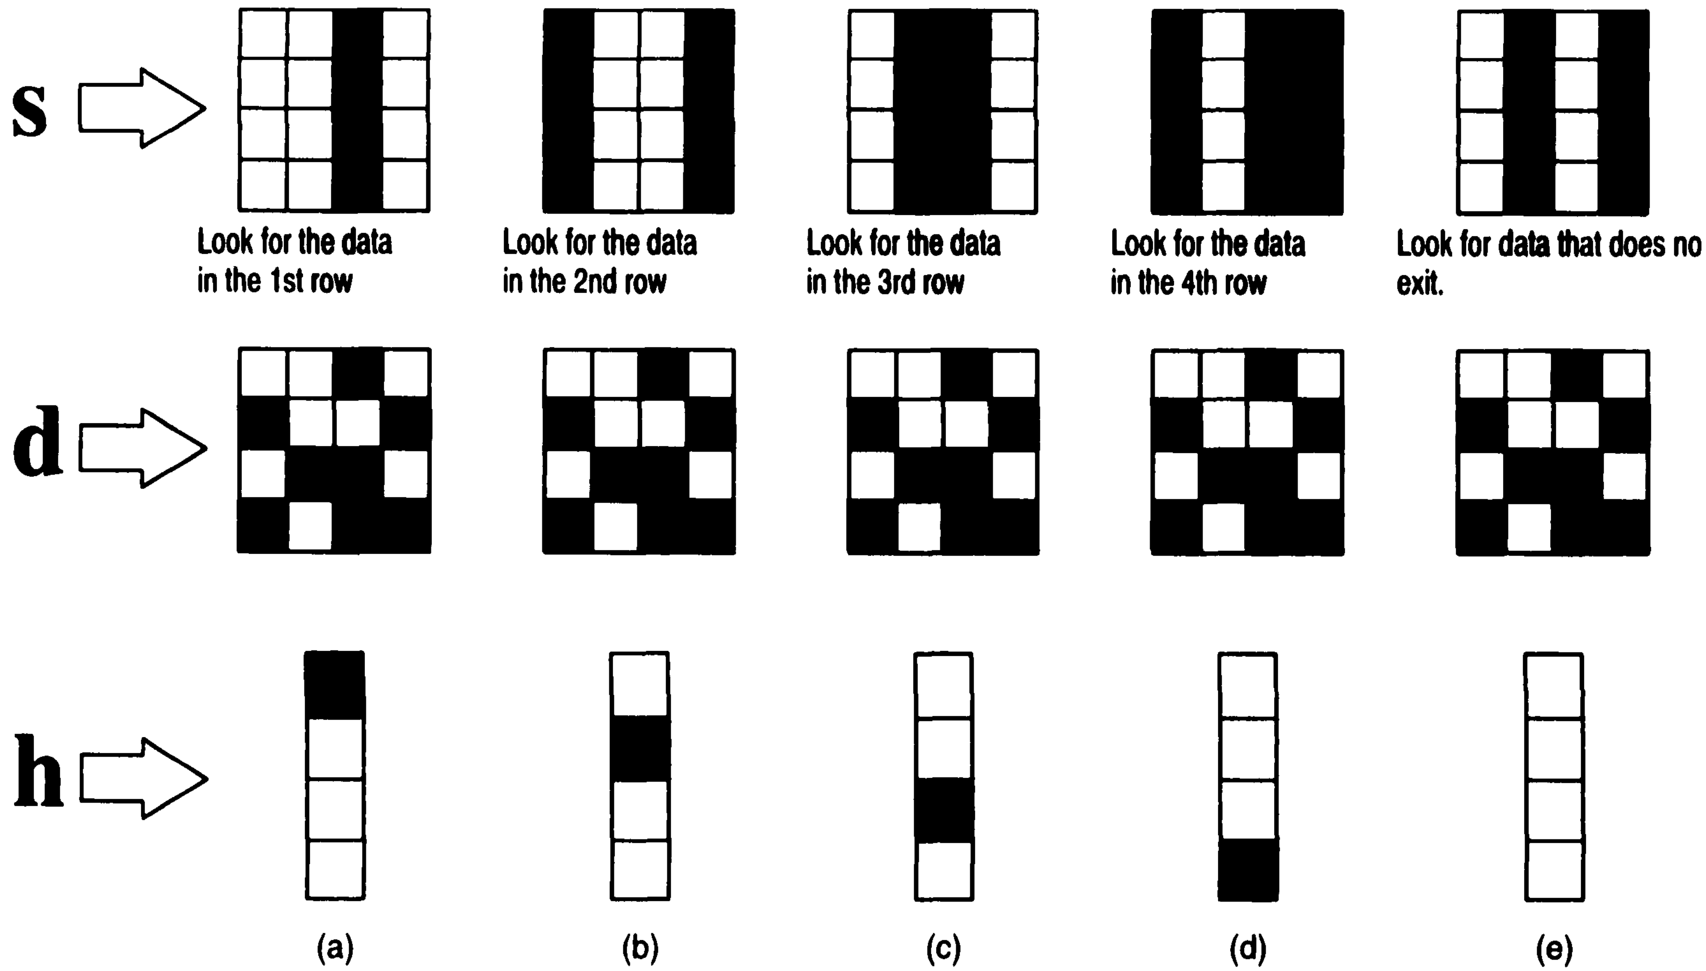

Figure 6.18. Outline of the five experiments performed to verify the core function of the optoelectronic digital associative processor (look-up table operation). Each column above contains the search matrix s, the data page d and the locator vector h produced by the bitwise XOR of s and d.

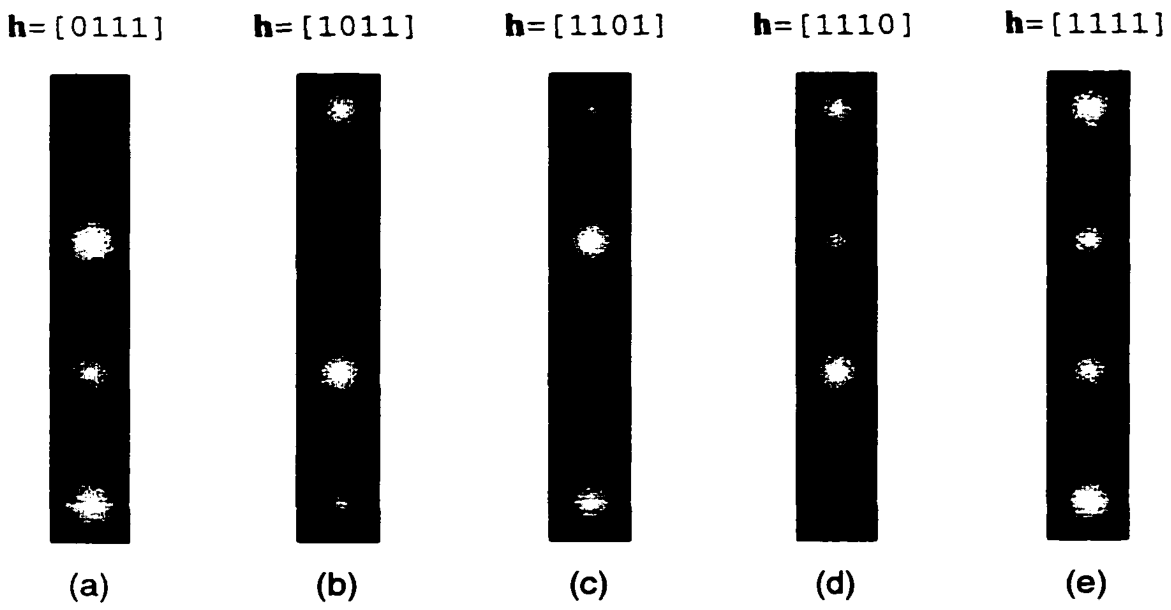

*Figure 6.19. **Experimental results verifying the look-up table functionality.***

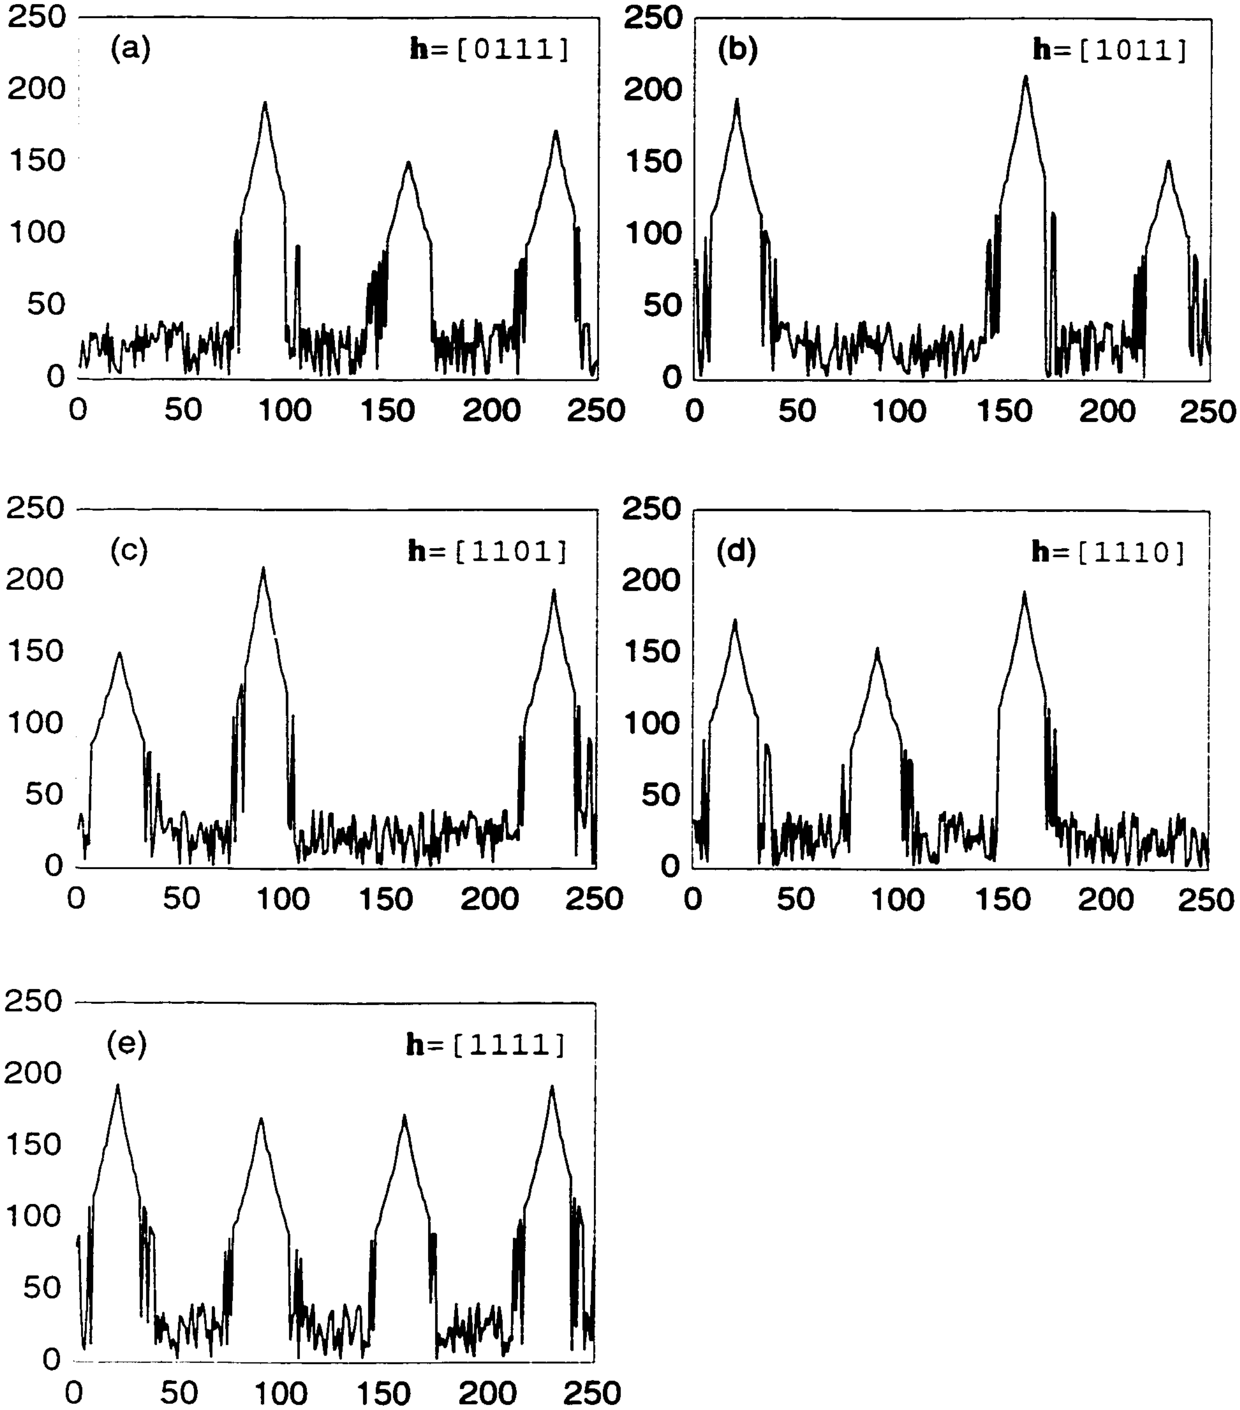

Figure 6.20. Intensity plots along the center line of each locator vector.

6.4 Multi-phase OERAM

In the original design for the OERAM (page 74) , a single-phase clock (SPC) was used to control the entire processes. Based on that design, the number of clock cycles required to scan a page with M rows is:

$$t_s = (M - s_p) + s_p(b + 1) \\ = s_p b + M \tag{8.10}$$

where s_p is the number of selected rows in a page and b is the number of word-blocks each of which is readout in a single clock cycle. (Eq. (8.10) is slightly different than Eq. (5.6) in that the former refers to the number of pages in the entire memory while the latter refers to a single page). For the original design of the OERAM we scan the selection bit for each row. When we find a selected row we pause scanning the selection bits and we proceed with reading out the row. After completion of reading out the row we resume the scanning of the selection bits for subsequent rows.

We have also nvestigated the possibility of a two-phase clock design. The rationale for this investigation is straight forward: in the original design the system remains idle while clock cycles are consumed to read out a selected row. We could take advantage of these idle states by searching for the next selected row while reading out a selected row. This is particularly useful for larger pages with a few selected rows. To do so, we need to introduce a second phase in the clock distribution. Each time a selected row is encountered, two things happen: using the first phase we read the selected row, while at the same time we scan the subsequent rows using the second phase, to locate the next selected row.

The time required to complete this operation, based on the flow chart of Figure 6.21 is:

$$t_d = [\rho_1 + (b + 1)] + \sum_{k=2}^{s_r} \Phi_k \tag{8.11}$$

$$\Phi_k = \begin{cases} b + 1, & if \rho_{k-1} - \rho_k \le b \\ \rho_{k-1} - \rho_k, & if \rho_{k-1} - \rho_k > b \end{cases}$$

where ρ_j are the locations of the selected rows. Eq. (8.11) is simple to follow. It takes ρ_1 clock cycles to locate the first selected row. Once the first selected row has been found, it takes $b + 1$ steps to read it out; this is the first term of Eq. (8.11).

The second term is a sum whose terms Φ_j depend on the distance between pairs of subsequent selected rows. If this distance is less or equal than b (i.e., the number of clock cycles it takes to read one row), then the next selected row is found while the previous selected row is being read. If, however, the distance is greater than b then we need to go a few extra steps to find the next selected row; $\rho_{k-2}-\rho_k-b$ steps to be precise. Add this number of steps to the b steps it took to read the previous selected row, and the number of clock steps required to read a selected row and find the next selected one when their distance is greater than b, is given by the lower branch of Φ_j.

Numeric simulations produced encouraging results, depending on the number of rows per page, and the number b of clock cycles it takes to read out one row. The results of two simulations are shown in Figure 6.22. The first simulation is for a data page with 100 rows. Each row has 1,024 bits grouped in 32 blocks for read out. The plot in Figure 6.22(a) shows the ratio t_s/t_d which we call improvement, over various values of selectivity. As selectivity reaches 1, the improvement ratio reaches a value of 1; this is a confirmation of the model's validity. For a selectivity of 1 both models should exhibit the same behavior.

The plot in Figure 6.22(b) shows the same ratio for a data page with 1000 rows of 1,024 bits each. The improvement here is dramatic.

The data shown here support the original hypothesis that a multi-phase clock distribution will improve significantly the performance of an intelligent optoelectronic interface.

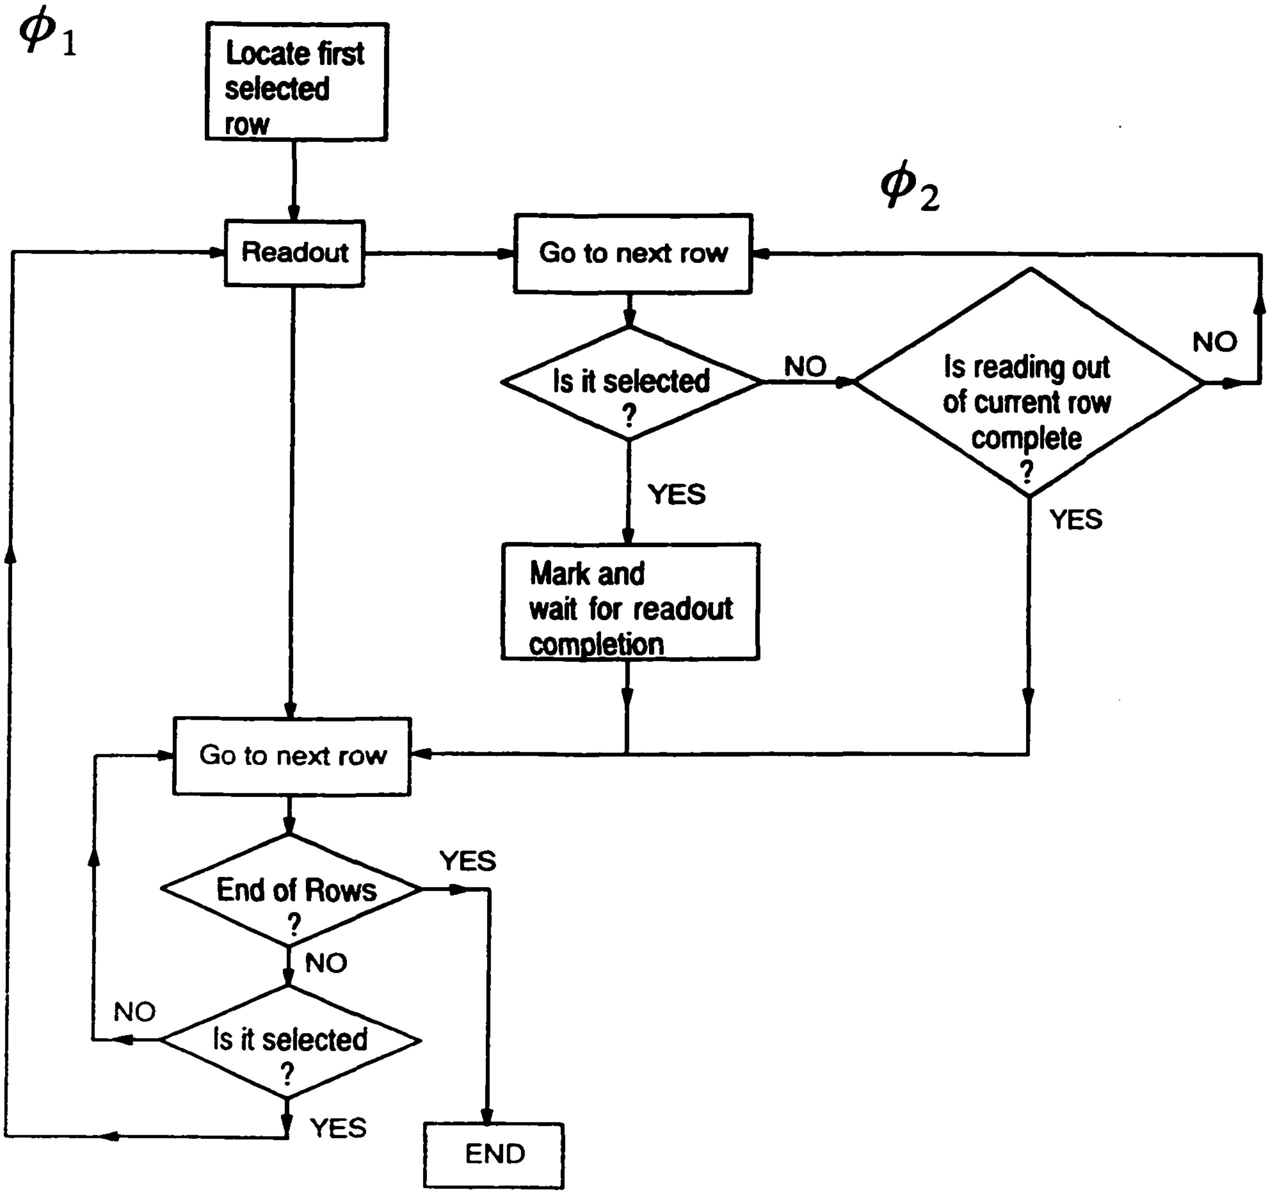

Figure 6.21. Flow chart for the two-phase clock operation. The areas of the two different phases ϕ_1, ϕ_2 are marked by the shaded boxes.

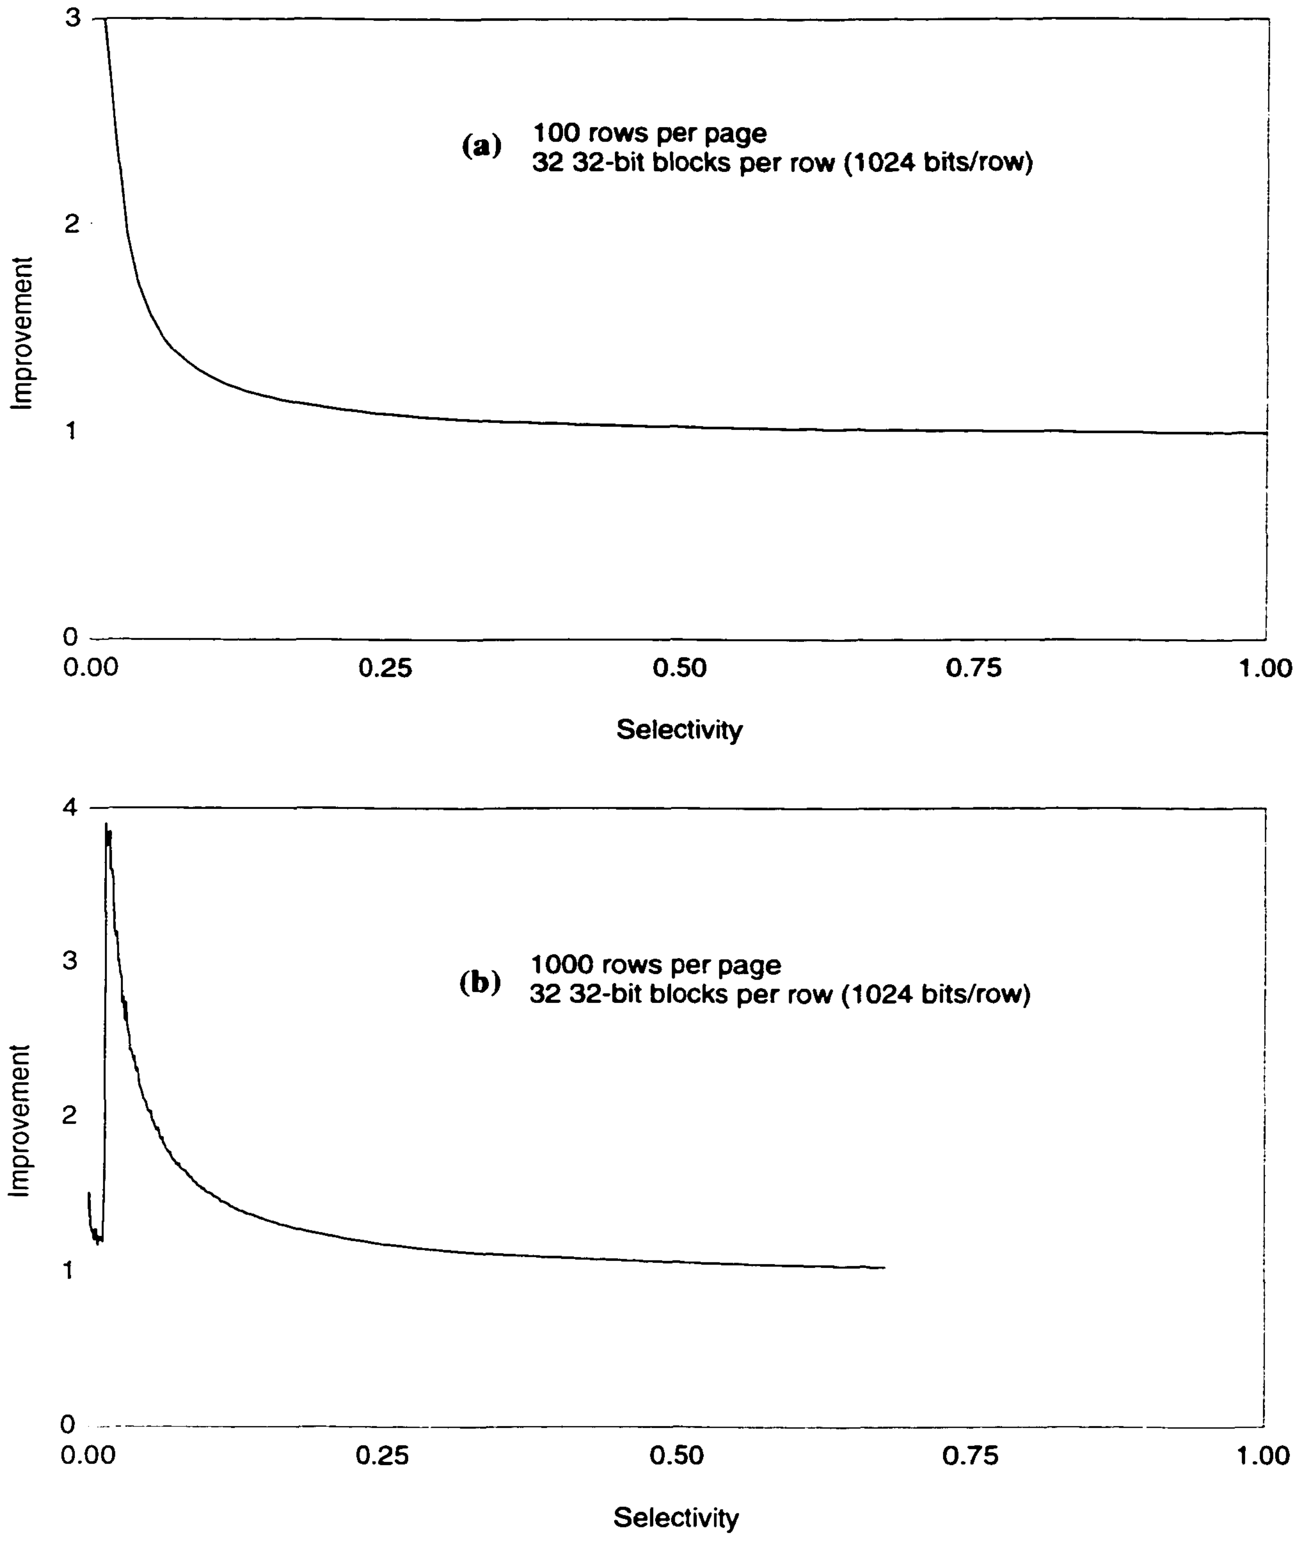

Figure 6.22. Simulation results for multiphase OERAM.

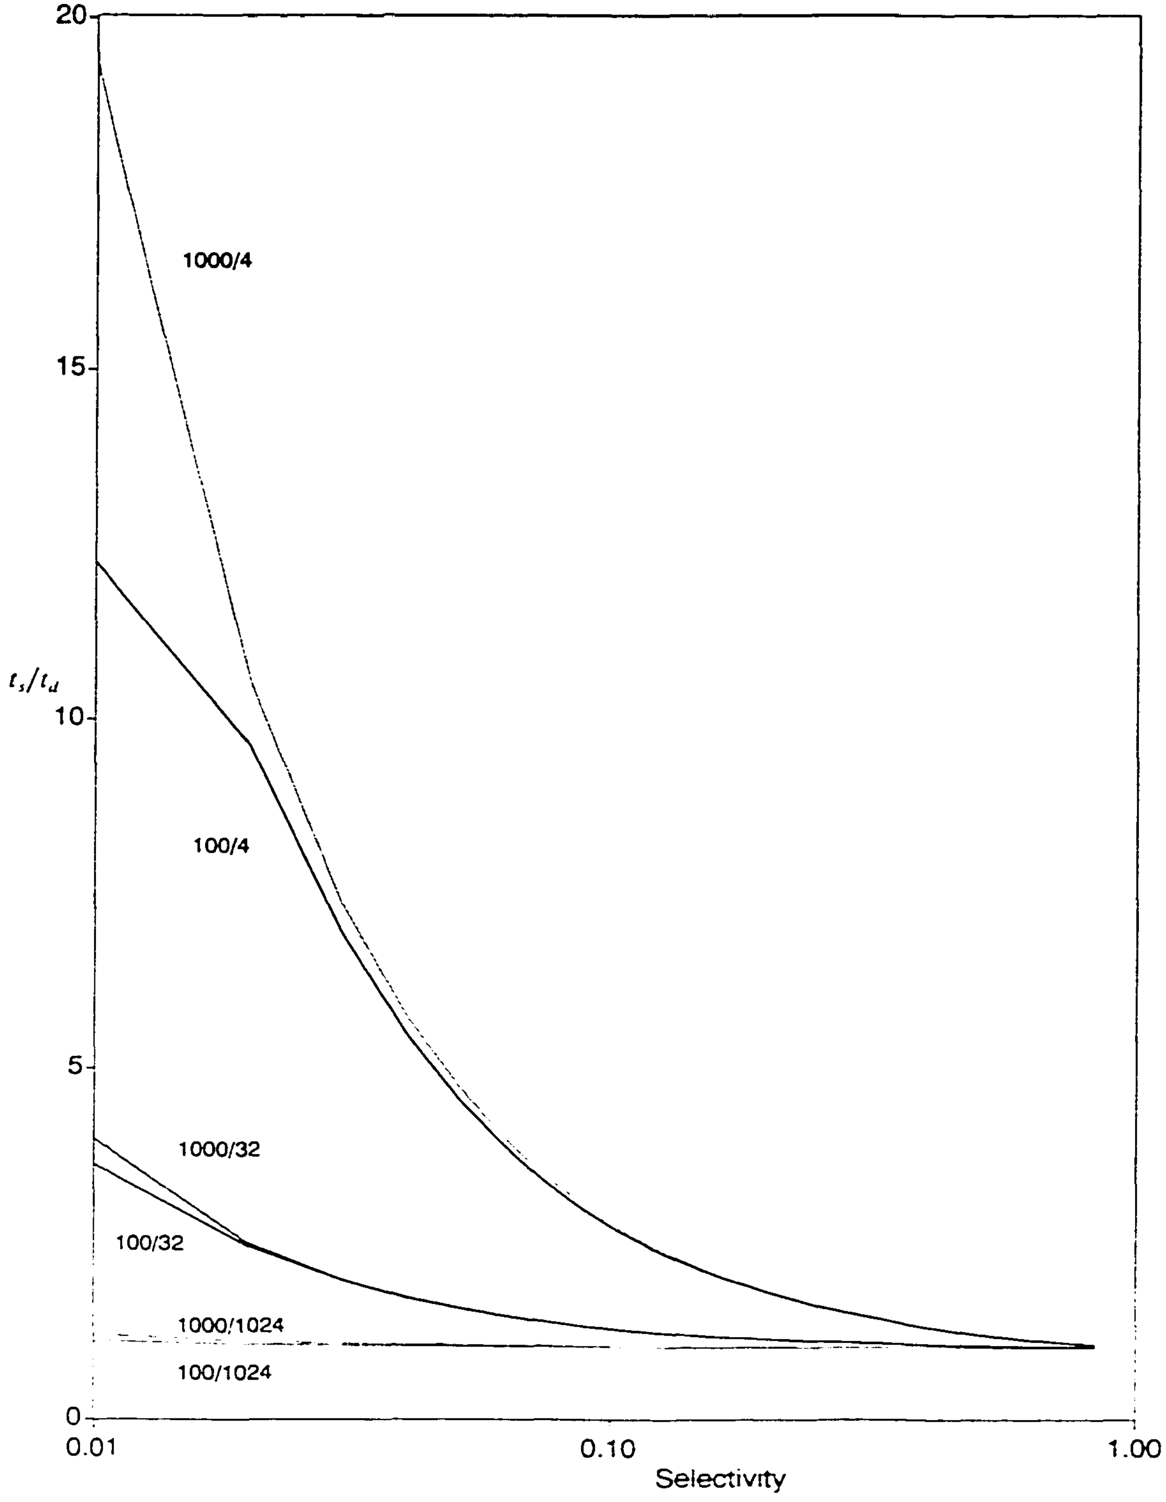

Figure 6.23. Improvement ratio in multiphase OERAM (t_s/t_d) versus selectivity for various combination of number of rows per page M, and number of readout blocks per row b. The fractional numbers indicate the pairs M/b.

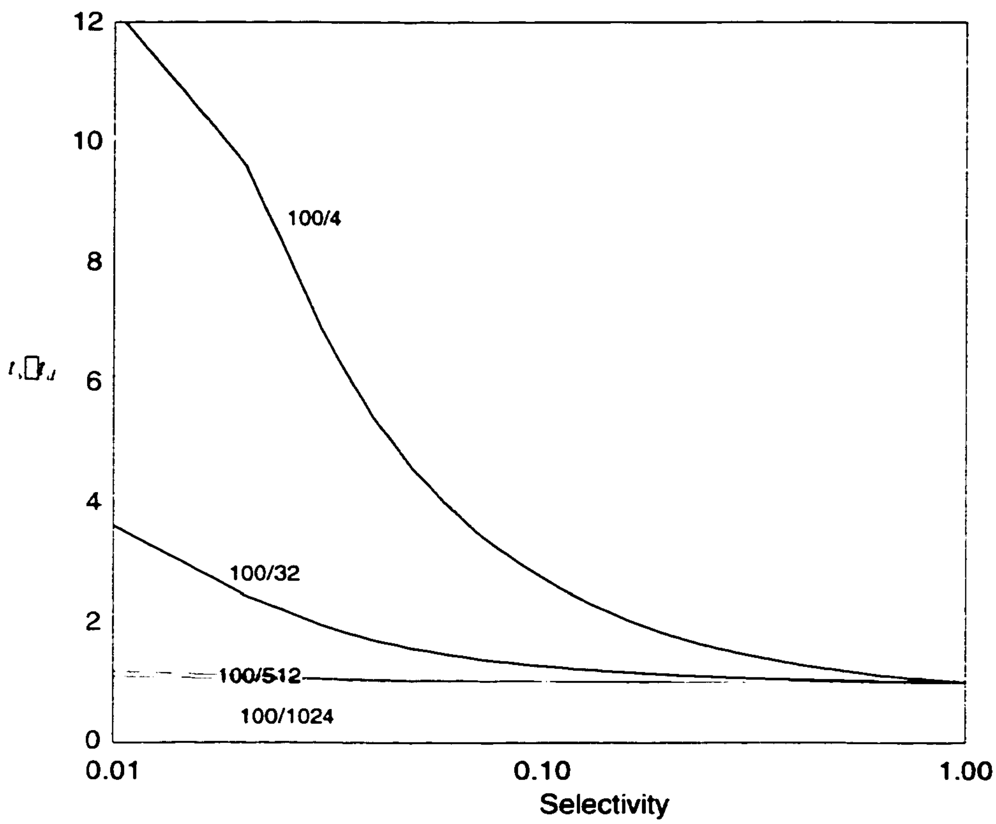

Figure 6.24. Multiphase OERAM simulations for b=4,32,512,1024 and M=100

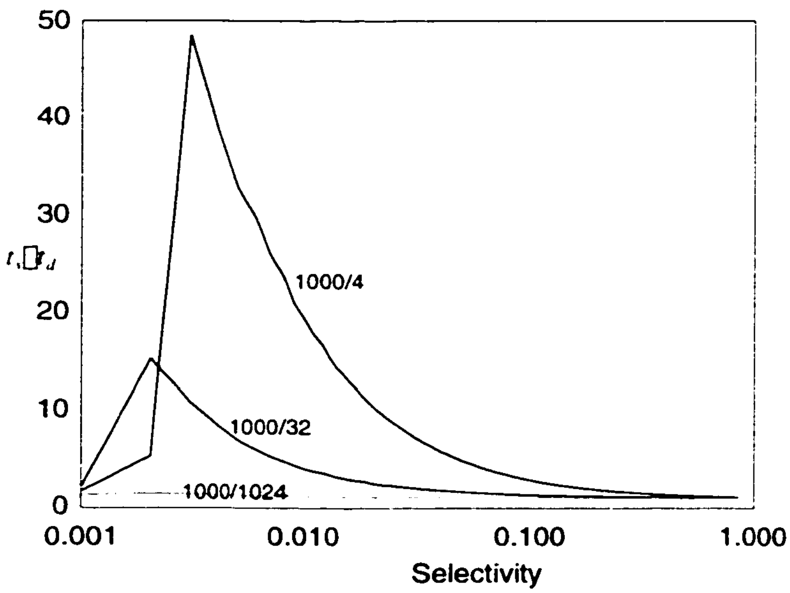

Figure 6.25. Multiphase OERAM simulations for b=4,32,1024 and M=1000.

7

Conclusion

In this dissertation we established that optical memories are a feasible and promising technology for archiving applications. For these applications large volumes of storage space is required because of the increasing quantities of data produced daily. Optical memories can only be integrated in the information processing fabric if they are matched by an opto-electronic interface (OEI). An OEI should meet certain design specifications. These specifications are derived from many considerations including data error minimization, performance capacity, hardware complexity, etc.

To derive the specifications mentioned above the following, broader question had also to be considered:

> *Given the optical memories that have been developed (or proposed) already and the existing architectures that dominate the computer market, what kind of interface between these two technologies is the most likely to turn optical storage into a viable peripheral?*

To answer this question we looked into the system description of optical memories and compared them to other storage technologies. This study led to some interesting findings regarding the suitability of optical memories for various applications. From this investigation we reached the conclusion that optical memories are better suited for archival applications. If for no other reason, at least for the fact that recording information in an optical memory is a rather slow process that prohibits their use as secondary storage, random access memories, or cache.

The discussion on the modeling and design of the interface between optical memories and electronic computers suggests, quite emphatically, that this is an ideal area to implement new parallel processing architectures and to improve traditional techniques. The introduction of the tomographic storage scheme for relational databases in Section 6.2 and the selection module for the optoelectronic digital associative processor in Section 6.3 are examples of medium to fine-grain parallelism that can be done nicely with optics and optoelectronic architectures.

However there are limitations that affect the implementation of such architecture. The most notable limitation is the following. The architectures that we propose here aim in enhancing the transfer of data from an optical memory to an electronic computer through optoelectronic filters. We design these filters for bit-wise operations and therefore a one-to-one correspondence between the data bits emerging from the memory and the processing elements on the filters is required. For a data page with $1{,}024 \times 1{,}024$ bits we will need an array of $1{,}024 \times 1{,}024$ processing elements. At the minimal configuration these elements (be it a smart pixel or a smart photodetector) comprise two photodetecting devices and a few logic gates. The second photodetector allows us to input command bits in the processing elements and thus perform simple logic operations with the datum. Furthermore, bits emerging from an optical memory need to be imaged on the processing elements with high-precision and minimum distortion and crosstalk.

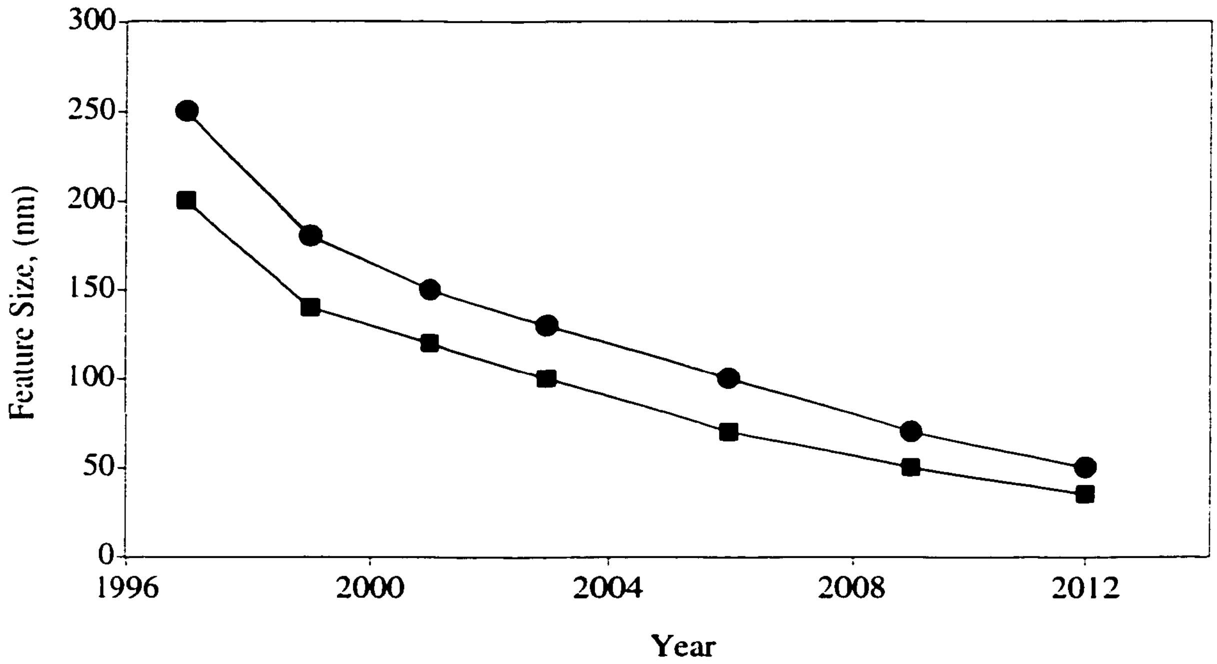

Figure 7.1. Feature size for current and new generations of DRAM (•) and MPU (▪) chips. Adopted from the 1997 National Technology Roadmap for Semiconductors, © Semiconductor Industry Association.

The real-estate requirements for the processing elements in the optoelectronic filter prohibit us from 1:1 imaging of the data bits as they emerge from the optical memory. Data bits are stored in high-densities in an optical medium. Their pitch size is much smaller than any conceivable pitch size on an array of optoelectronic processing elements available today. Therefore $1 : x,\ x > 1$ imaging will be necessary. An optical function like that will most likely increase distortion and reduce optical power delivered to the photodetectors because of the additional optical components, thus requiring larger photodetectors. This, in turn will increase the zoom factor x, and by the time we reach an acceptable factor the architecture will be very difficult to implement just because of the size requirements.

Realistically we should not expect arrays with more than 10^4 or 10^5 processing elements unless a major (and quite unexpected) R&D effort yields some results. Figure 7.1 shows the reduction of feature sizes in VLSI technologies as projected for the next few years. If that kind of integration becomes available for the manufacturing of optoelectronic processing elements it may be possible to achieve reasonable densities in the order of 10^6. However it seems very unlikely.

On the other hand it is quite reasonable to expect that page-oriented optical memories will produce pages with more than 10^6 bits per page. If we cannot use arrays of optoelectronic processing elements to transfer these pages to the electronic computer then the best alternative would be to use CCD devices and then process the data electronically as a digital image. This approach that is currently being supported by IBM [Shelby, 97a] offers a way of transferring data from the memory to the electronic computer but offers no means for on-the-fly processing and enhancement of computing performance. With the increasing performance and decreased cost of electronic processors however the argument of optoelectronically enhanced computing may become obsolete [Ozaktas, 97a; 97b].

7.1 Future Directions

Let us now return to the question what we can do with optoelectronic processing architectures. Given the technology limitations we should focus our attention to arrays of low and medium complexity with no more than 10^4 or 10^5 processing elements. Clearly these arrays cannot be a match for the projected densities of optical memories. However the size of these arrays is just right for networking switching applications. In particular, arrays of optoelectronic processing elements can be very useful in highly localized switching fabrics ranging from backplane interconnects to a small switch supporting a high-performance computing cluster. To that extend we are in agreement with [Ozaktas, 98a].

Furthermore, it is necessary to understand better the computability issues related to the problems that optoelectronic computing tries to solve. Bottlenecks in database processing are a real problem; however with increase network bandwidth, faster I/O systems, and more efficient algorithms, database transactions seem to complete fast enough [Ullman, 99a]. Associative processing seems to be an area where optics have the most competitive edge.

The future of optoelectronic processing architectures seems to be tied to that of parallel optical memories. If volumetric optical storage is successful in competing with magnetic systems in the near future it is certain that optoelectronic architectures for data processing will develop and mature further.

References

[Bartzos, 98a] Bartzos, D., P. Panayotatos, S. R. McAfee, and D. G. Daut, 1998, "Parameter extraction for a figure-of-merit approach for the comparison of optical interconnects," *Optical Engineering,* Vol. **37**, pg. 2772–2783.

[Bernal, 98a] Bernal, M., G. W. Burr, H. Coufal, and M. Quintanilla, 1998, "Balancing interpixel cross talk and detector noise to optimize areal density in holographic storage systems," *Applied Optics,* Vol. **37**, pg. 5377–5385.

[Berra, 89a] Berra, P. B., A. Ghafoor, P. A. Mitkas, S. J. Marcinkowski, and M. Guizani, 1989, "The impact of optics on data and knowledge base systems," *IEEE Transactions on Knowledge and Data Engineering,* Vol. **1**, pg. 111–132.

[Betzos, 98a] Betzos, G. A. and P. A. Mitkas, 1998, "Performance evaluation of massively parallel processing architectures with three-dimensional optical interconnections," *Applied Optics,* Vol. **37**, pg. 315–325.

[Beyette, 95a] Beyette Jr., F. R., *Design and analysis of an optoelectronic recirculating sorter,"* Ph.D. Dissertation, Colorado State University, Fort Collins, Colorado, 1995.

[Blombergen, 48a] Bloembergen, N., E. M. Purcell, and R. V. R. V. Pound, 1948. *Phys. Rev.* Vol. **71**, pg. 679.

[Brackenbury, 98a] Brackenburry, L. E. M., 1998, "Optical solution to the Lee algorithm by use of symbolic substitution," *Applied Optics,* Vol. **37**, pg. 3705–3716.

[Bromberg, 88a] Bromberg, J. L., 1988 (October), "The birth of laser," *Physics Today,* pg. 46.

[Burr, 94a] Burr, G. W., F. H. Mok, and D. Psaltis, 1994, "Large scale volume holographic storage in the long interaction length architecture," in

Photonics for Processors, Neural Networks, and Memories II, J. L. Horner, B. Javidi, and S. T. Kowel (editors), Proceedings of SPIE, Vol. **2297**, pg. 402–414.

[Burr, 95a] Burr, G. W., 1995, *private communication.*

[Burr, 98a] Burr, G., 1998, *private communication.*

[Bursky, 95a] Bursky, D., 1995, "Gigabit DRAMs, 64-bit CPUs and More at ISSCC," *Electronic Design*, ppg. 61–62.

[Chen, 67a] Chen, F. S., 1967, "A laser-induced inhomogeneity of refractive indices in KTN," *Journal of Applied Physics*, Vol. **38**, pg. 3418–3420.

[Chiarulli, 96a] Chiarulli, D. M. and S. P. Levitan, 1996, "Optoelectronic–cache Memory System Architecture," *Applied Optics*, Vol. **35**, (14), pg. 2449–2456.

[Cruz-Rivera, 98a] Cruz-Rivera, J. L., D. S. Wills, T. K. Gaylord, and E. N. Glytsis, 1998, "Optimal usage of available wiring resources in diffractive-reflective optoelectronic multichip modules," *Applied Optics*, Vol. **37**, pg. 233–253.

[Dvornikov, 94a] Dvornikov, A. S., S. E. Esener, and P. M. Rentzepis, 1994, "Three-dimensional optical storage memory by means of two-photon interaction," in *Optical Computing Hardware*, AT&T and Academic Press, pg. 287–325

[Esener, 98a] Esener, S. E. and M. H. Kryder, (chairs), *Future of Data Storage Technologies*, online proceedings of the World Technology Evaluation Center, http://itri.loyola.edu/hdmem

[Fey, 98a] Fey, D., B. Kasche, C. Burkert, and O. Tschache, 1998, "Specifications for a reconfigurable optoelectronic VLSI processor suitable for digital signal processing," *Applied Optics*, Vol. **37**, 284–295

[Gao, 98a] Gao, Q. and R. Kostuk, 1998, "Cross-talk noise and storage capacity of holographic memories with a $LiNbO_3$ crystal in the open-circuit condition," *Applied Optics*, Vol. **37**, pp. 929–936.

[Glass, 78a] Glass, A. M., 1978, "The phororefractive effect," *Optical Engineering*, Vol. **17**, pg. 470–479.

[Göppert, 31a] Göppert-Mayer, M., 1931, "Über Elementerakte mit zwei quantensprüngen," *Ann. Phys. (Leipsig)*, Vol. **9**, pg. 273.

[Goertzen, 95a] Goertzen, B. J., 1995, *private communication.*

[Haney, 98a] Haney, M. and M. P. Christensen, 1998, "Performance scaling comparison for free-space optical and electrical interconnection approaches," *Applied Optics*, Vol. **37**, pg. 2886–2894.

[Hayes, 97a] Hayes, E. M., Jurrat, R., Pu, R., Snyder, R.D., Feld, S.A., Stanko, P., Wilmsen, C.W.. Choquette, K.D.. Geib, K.M., Hou, H.Q., "Foundry fabricated array of smart pixels integrating MESFETs/MSMs and VCSELs," *International Journal of Optoelectronics*, Vol. **11** (3), 1997, pg.229–237.

[Heanue, 94a] Heanue, J. F., M. C. Bashaw, and L. Hesselink, 1994, "Volume Holographic Storage and Retrieval of Digital Data," *Science*, Vol. **265**, 5 August 1994, pg. 749–752.

[Heerden, 63a] Heerden, P. J. van, 1963, "Theory of optical information storage in solids," *Applied Optics*, Vol. **2**, pg.393–400.

[Hesselink, 93a] Hesselink, L. and M. C. Bashaw, 1993, "Optical memories implemented with photorefractive media," *Optical and Quantum Electronics*, Vol. **25**, pg. S611–S661.

[Hoanca, 98a] Hoanca, B. and A. A. Sawchuk, 1998, "Cellular interconnects optimization algorithm for optoelectronic single-instruction multiple data," *Applied Optics*, Vol. **37**, pg. 871–883.

[Huang, 98a] Huang, C.-C., B. K. Jenkins, and C. B. Kuznia, 1998, "Space-variant interconnections based on diffractive optical elements for neural networks: architectures and cross-talk reduction," *Applied Optics*, Vol. **37**, pg. 889–911.

[Hunter, 90a] Hunter, S., F. Kiamilev, S. Esener, D. A. Parthenopoulos, and P. M. Rentzepis, 1990, "Potentials of two-photon-based three-dimensional optical memories for high performance computing," *Applied Optics*, Vol. **29**, No. 14, pg. 2058–2066.

[Irakliotis, 94a] Irakliotis, L. J. and P. A. Mitkas, "Tomographic Storage Scheme for Three-dimensional Implementation of Database Operations," in *Proceedings of SPIE*, vol. **2297,** 1994, pg. 359–367.

[Irakliotis, 95a] Irakliotis, L. J., S. A. Feld, F. R. Beyette Jr., P. A. Mitkas, and C. W. Wilmsen, 1995, "Optoelectronic parallel processing with surface-emitting lasers and free-space interconnects," *Journal of Lightwave Technology*, Vol. **13**, No. 6, pg. 1074–1084.

[Irakliotis, 98a] Irakliotis, L. J. and P. A. Mitkas, 1998, "Optics: a maturing technology for better computing," *Computer*, Vol. **31,** pg. 36–37.

[Kajita, 98a] Kajita, M., K. Kasahara, T. J. Kim, D. T. Neilson, I. Ogura, I. Redmond, and E. Schenfeld, 1998, "Wavelength-division multiplexing free-space optical interconnect networks for massively parallel processing systems," *Applied Optics*, Vol. **37**, pg. 3746–3755.

[Jackson, 75a] J. D. Jackson, *Classical Electrodynamics*, 2nd ed., Wiley, New York, 1975.

[Kharlamov, 74a] Kharlamov, B. M., R. I. Personov, and L. A. Bykovskaya, 1974, "Stable gap in absorption spectra of solid solutions of organic molecules by laser irradiation," *Optics Communications*, Vol. **12**, pg. 191–193.

[Kogelnik, 69a] Kogelnik, H., 1969, "Coupled wave theory for thick hologram gratings," *The Bell System Technical Journal*, Vol. **48**, pg. 2909–2947.

[Kukhtarev, 79a] Kukhtarev, N. V., V. B. Markov, S. G. Odulov, M. S. Soskin, and V. L. Vinetskii, 1979, "Holographic storage in electrooptic crystals. I. Steady state," *Ferroelectrics*, Vol. **22**, pg. 949–960.

[Kukhtarev, 79b] Kukhtarev, N. V., V. B. Markov, S. G. Odulov, M. S. Soskin, and V. L. Vinetskii, 1979, "Holographic storage in electrooptic crystals. II. Beam coupling — light amplification," *Ferroelectrics*, Vol. **22**, pg. 961–964.

[Kurokawa, 98a] Kurokawa, T., S. Matso, T. Nakahara, K. Tateno, Y. Ohiso, A. Wakatsuki, and H. Tsuda, 1998, "Design approaches for VCSEL's and VCSEL-based smart pixels toward parallel optoelectronic processing systems," *Applied Optics*, Vol. **37**, pg. 194–204.

[Lessard, 98a] Lessard, R. A. (ed), *Photopolymer Device Physics, Chemistry, and Applications IV*, SPIE Proceedings, Vol. 3417, 1998.

[Li, 98a] Li, Y., T. Wang, and S. Kawai, 1998, "Distributed crossbar interconnects with vertical-cavity surface-emitting laser-angle multiplexing and fiber image guides," *Applied Optics*, Vol. **37**, pg. 254–263.

[Li, 98b] Li, Y., K. Fasanella, and T. Wang, 1998, "Lost-cost planar star-coupling structure for large-core polymer optical fibers," *Applied Optics*, Vol. **37**, pg. 884–888.

[Louri, 97a] Louri, A. and C. Neocleous, 1997, "Incrementally scalable optical interconnection network with a constant degree and constant diameter for parallel computing," *Applied Optics*, Vol. **36**, pg. 6594–6604.

[Marchant, 90a] Marchant, A. B., 1990, *Optical Recording, A Technical Overview*, Addison-Wesley.

[Meixner, 89a] Meixner, A. J., A. Renn, and U. P. Wild, 1989, "Spectral hole burning and holography. I. Transmission and holographic detection of spectral holes," *J. Chem. Phys.*, Vol. **91**, 6728.

[Mitkas, 93a] Mitkas, P. A. and P. B. Berra, 1993, "PHOEBUS: An optoelectronic database machine based on parallel optical disks," *Journal of Parallel and Distributed Computing*, Vol. **17**, pg.230–244.

[Mitkas, 96a] Mitkas, P. A. (ed.), 1996, *Data Encoding for Page-oriented Optical Memories 1996*, held in Phoeniz, Arizona, March 1996; proceedings published by Colorado State University.

[Mitkas, 98a] Mitkas, P. A., G. A. Betzos, and L. J. Irakliotis, 1998, "Optical processing paradigms for electronic computers," *Computer*, Vol. **31**, pg. 45–51.

[Miyazaki, 98a] Miyazaki, D., S. Masuda, and K. Matsushita, 1998, "Self-alignment with optical microconnectors for free-space optical interconnections," *Applied Optics*, Vol. **37**, pg. 228–232.

[Moerner, 88a] Moerner, W. E., 1988, *Persistent Spectral Hole-Burning: Science and Applications*, edited by W. E. Moerner, Topics in Current Physics No. 44, (Springer, Berlin), pg. 1

[Moerner, 88b] Moerner, W. E., W. Lenth, and G. C. Bjorklund, 1988, "Frequency domain optical storage and other applications of persistent spectral hole-burning," in *Persistent Spectral Hole-Burning: Science and Applications*, edited by W. E. Moerner, Topics in Current Physics No. 44, (Springer, Berlin), pg. 251

[Mori, 98a] Mori, M., Y. Yagai, T. Yatagai, and M. Watanabe, 1998, "Optical learning neural network with a Pockels readout optical modulator," *Applied Optics*, Vol. **37**, pg. 2852–2862.

[Neifeld, 98a] Neifeld, M. A. and R. K. Kostuk, 1998, "Error correction for free-space optical interconnects: space-time resource optimization," *Applied Optics*, Vol. **37**, pg. 296–307.

[Ozaktas, 96a] Ozaktas, H., 1996, *private communication*.

[Ozaktas, 97a] Ozaktas, H., 1997, "Toward an optimal foundation architecture for optoelectronic computing. Part I. Regularly interconnected device planes," *Applied Optics*, Vol. **36**, pg. 5682–5696.

[Ozaktas, 97b] Ozaktas, H., 1997, "Toward an optimal foundation architecture for optoelectronic computing. Part II. Physical construction and application platfroms," *Applied Optics*, Vol. **36**, pg. 5696–5705s.

[Papoulis, 94a] Papoulis, A., 1994, *Probability, Random Variables, and Stochastic Processes*, 3rd edition, McGraw-Hill, New York.

[Peng, 98a] Peng, C. and M. Mansuripur, 1998, "Sources of noise in erasable optical disk data storage," *Applied Optics*, Vol. **37**, pg. 921–928.

[Pinkston, 98a] Pinkston, T. M., M. Raksapatcharawong, and Y. Choi, 1998, "WARP core: optoelectronic implementation of network-router deadlock-handling mechanisms," *Applied Optics*, Vol. **37**, pg.276–283.

[Psaltis, 90a] Psaltis, D., M. A. Neifeld, A. Yamamura, and S. Kobayashi, 1990, "Optical memory disks in optical information processing," *Applied Optics*, Vol. **29**, pg. 2038–2057.

[Psaltis, 98a] Psaltis, D. and G. W. Burr, 1998, "Holographic data storage," *Computer*, Vol. **31**, pg. 52–60.

[Liu, 98a] Liu, Y., B. Robertson, G. C. Boisset, M. H. Ayliffe, R. Iyer, and D. V. Plant, 1998, "Design, implementation, and characterization of a hybrid optical interconnect for a four-stage free-space optical backplane demonstrator," *Applied Optics*, Vol. **37**, pg. 2895–2914.

[Pu, 95a] Pu A., R. Denkewalter, and D. Psaltis, 1995, "Robot navigation using a peristrophic holographic memory," in *Optical Computing 1995*, OSA 1995 Technical Digest Series, Vol. **10**, pg. 137–139, Salt Lake City March 13–16.

[Sahin, 98a] Sahin, A., M. A. Kutay, and H. M. Ozaktas, 1998, "Nonseparable two-dimensional fractional Fourier transform," *Applied Optics*, Vol. **37**, pg. 5444–5453.

[Sayles, 91a] Sayles, A. H. and J. P. Uyemura, 1991, "An optoelectronic CMOS memory circuit for parallel detection and storage of optical data," *IEEE Journal of Solid-State Circuits*, Vol. **26**, pg. 1110–1115.

[Schaffer, 98a] Schaffer, M., and P. A. Mitkas, 1998, "Smart Photodetector Array for Page-Oriented Optical Memory in 0.35-μm CMOS," *IEEE Photonics Technology Letters*, Vol. **10**, No. 6, pg. 866–867.

[Schaffer, 98b] Schaffer, M., and P. A. Mitkas, 1998, "Requirements and Constraints for the Design of Smart Photodetector Arrays for Page-Oriented Optical Memories," *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. **4**, No. 5, pg. 856–865.

[Schawlow, 61a] Schawlow, A. L., 1961, "Fine structure and properties of chromium fluorescence in aluminum and magnesium oxide," in *Advances in Quantum Electronics*, edited by J. R. Singer, (Columbia University, New York), pg. 50.

[Shelby, 97a] Shelby, R. M., J. A., Hoffnagle, G. W. Burr, C. M. Jefferson, M.-P. Bernal, H. Coufal, R. K. Grygier, H. Gunther, R. M. Macfarlane, and G. T. Sincerbox, "Pixel-matched holographic data storage with megabit pages," *Optics Letters*, Vol. **22**, pg. 1509–1512.

[SIA, 98a] Semiconductor Industry Association, 1998, *The National Technology Roadmap for Semiconductors*.

[Snyder, 94a] Snyder, R. D., Beyette, F.R. Jr., Feld, S.A., Geib, K.M., Irakliotis, L.J., Mitkas, P.A., and Wilmsen, C.W., "Phototransistor/surface Emitting Laser-Smart Pixel Array Implementation of an Opto-electronic Data Filter," in *LEOS Summer Topical Meeting on Smart Pixels*, 1994, pg. 45–46.

[Steckman, 98a] Steckman, G. J., Solomatine I., Zhou G., and Psaltis D., "Characterization of phenanthrenequinone doped PMMA for holographic memory," *Optics Letters*, Vol. **23,** (16) pg 1310–1312.

[Strzelecka, 98a] Strzelecka, E., D. A. Louderback, B. J. Thibeault, G. B. Thompson, K. Bertilsson, and L. A. Coldren, 1998, "Parallel free-space optical interconnect based on arrays of vertical-cavity lasers and detectors with monolithic microlenses," *Applied Optics*, Vol. **37**, pg. 2811–2819.

[Szabo, 70a] Szabo, A., 1970, "Laser-induced fluorescence line narrowing in ruby,," *Phys. Rev. Lett.*, Vol. **25**, 924.

[Szabo, 75a] Szabo, A., 1975, "Observation of hole burning and cross relaxation effects in ruby," *Phys. Rev. B*, Vol. **11**, 4512.

[Szymanski, 98a] Szymanski, T. H., A. Au, M. Lafreniere–Roula, V. Tyan, B. Supmonchai, J. Wong, B. Zerrouk, and S. T. Obenaus, 1998, "Terabit optical local area networks for multiprocessing systems," *Applied Optics*, Vol. **37**, pg. 264–275.

[Trout, 98a] Trout, T. J. (ed), *Holographic Materials IV*, SPIE Proceedings, Vol. 3294, 1998

[Ullman, 99a] J. Ullman, 1999, *private communication.*

[Uyemura, 92a] Uyemura, J. P. and B. L. Austin, 1992, "Hybrid optical-electronic logic gates in complementary metal-oxide semiconductor very-large-scale integration," *Applied Optics*, Vol. **31**, pg. 1774–1782.

[Wangsness, 86a] R. K. Wangsness, *Electromagnetic Fields*, 2nd ed., Wiley, New York, 1986.

[Yayla, 98a] Yayla, G. I., P. J. Marchand, and S. C. Esener, 1998, "Speed and energy analysis of digital interconnections: comparison of on-chip, off-chip, and free-space technologies," *Applied Optics*, Vol. **37**, pg. 205–227.

[Yeh, 93a] Yeh, P., 1993, *Introduction to Photorefractive Nonlinear Optics*, John Wiley and Sons.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

V

W

Y