

Optical Parallel Database Management System for Page Oriented Holographic Memories

Jian Fu, Marius P. Schamschula, and H. John Caulfield*

Physics Department, Alabama A. & M. University, Normal, AL 35762,
Tel. (205) 851-5871, Fax. (205) 851-5622, e-mail marius@caos.aamu.edu

*Physics Department, Fisk University, Nashville, TN 37208
Tel. (615) 329-8785, Fax. (615) 329-8634, e-mail hjc@dubois.fisk.edu

Abstract

Current data recall rates from page oriented holographic memories far exceed the ability of electronics to even read or transmit the data. For database management, we must not only read those data but also query them-computationally, a far more complex task. That task is hopeless for electronics. We show here the rudiments of an optical system that can do most of the query operations in parallel in optics, leaving the burden for electronics significantly less. Even here, electronics is the ultimate speed limiter. Yet we can query data far faster in our optical/electronic system than any purely electronic system.

Keywords: Holographic memory; Fourier correlator; Database management

I Introduction

Most massive database and knowledge base systems store data on magnetic or optical disks and employ indexing techniques to minimize disk accesses. Various clustering and accessing techniques are used to reduce response time. Even so, when the joint requirements of very large database (VLDB) and very short response times are imposed, existing technologies degrade considerably. Holographic memories can: (1) store hundreds of billions of bytes of data, (2) transfer them at a rate of a billion or more bits per second and (3) select a randomly chosen data element in 100 microseconds or less because of the 3-D recording and the simultaneous readout of an entire page of data at one time. No other memory technology that offers all three advantages is as close to commercialization [1].

Disks and tapes store bits of information which can be accessed sequentially. A page oriented holographic memory (POHM) stores and retrieves 2D data arrays (pages) in parallel from a store (a thick hologram) containing many such pages. In the recording of each page, a unique reference beam characterized by a parameter P is used. We then change P and record a different page. In replay, we choose the page by selecting P as suggested in Fig. 1. The first POHMs featured P as spatial position. Most current work uses P as angle. Using P as the index of a

coded reference beam has great potential. With tunable lasers, the use of P as wavelength is also attractive. For the most part, those details are not very relevant to the question of how to do POHM database management system (DBMS) [2].

Each page can contain B bits (usually in the 10^5 - 10^6 range). We can store and recall P pages (usually in the 100 to 10,000 range). The random access time in P is T (usually in the 10^{-3} to 10^{-7} range). The total storage capacity is $C = BP$ pages and the data random access rate is $\hat{C} = C/T$.

Let us illustrate with some achievable values such as $B=10^5$ bits, $P=1000$, and $T=10^{-6}$ seconds. Then $\hat{C}=10^{14}$ bits/s.

These are achievable, fairly conservative numbers. The critical number is \hat{C} . Ultra-fast electronics may approach bit rates of 10^9 bits/sec, a factor of 10^5 away from keeping up. Electronics is not going to bridge that gap soon.

If we hope to understand information almost as fast as we can recall it, we must abandon serial electronics in favor of parallel optics.

Driven primarily by the military need to find uncooperative target in noisy backgrounds, Fourier optical pattern recognition has improved immensely over the last 35 years. We now have compact systems with fast components implementing optimum masks on line.

The POHM DBMS problem will require all of those advances and more. This time, the targets are cooperative. We can design their signatures. Instead of roughly TV frame rates ($\sim 20/\text{second}$), we seek page decisions much faster ($\sim 10^5/\text{sec}$ to $10^6/\text{sec}$).

The purpose of this paper is to describe optical DBMSs for POHMs. With POHMs we can randomly access data many of orders of magnitude faster than we can serialize and digest them electronically. Parallel optical DBMSs may offer the only hope of catching up with the data rate. In this paper we discuss the application of Fourier optical correlators in POHMs and focus on how construct a relational or object oriented DBMS for POHMs with low access time, and high transfer rate.

II Background

2.1 DBMSs

A DBMS is a software program that is concerned with the task of controlling and managing the database as a resource somewhat independently of the computer hardware that hosts it and the application programs that interface with it. The DBMS must have the facility to establish the database within the system in response to the database designers. The DBMS make the data available to a wide variety of users ranging from external application programs to a casual user

posing a particular query. Inevitably, the database must be updated. That is, new data must be added, while old data must be deleted, and existing data must be changed. Thus, the DBMS must have the capability of performing these updates. In fact, many databases have almost as much update activity. Of course there are many types of databases that have limited or controlled update activity. The DBMS must provide the facility for insuring the integrity of the database. This is obtained through various consistency checks and backup and recovery systems. Finally, the DBMS must regulate access to the database to protect it, the system itself, and the privacy of users.

We are concerned with systems that must deal with a VLDB with a real time requirement. VLDBs are at the heart of many existing information systems and will play an even more prominent role in the future [3]. Since the DBMS is just another application program, it must adhere to normal execution procedures, just as other programs. The database user interacts with the DBMS through a query language (or other language) to accomplish a task. The DBMS must interact with the operating system to obtain data from the database which is stored on the computing system's secondary memory. Since the operating system must satisfy a large number of types of user, the size of the block of data retrieved from disk is optimized for all users and is thus fixed. The block of data is placed in main memory and turned over to the DBMS

that sifts through it to find what it wants. There may be little data of interest to the DBMS due to the organization of the data and type of query. Thus, the DBMS may have to ask the operating system for many pages of data to satisfy a query. This repeated access to secondary storage considerably degrades the performance of the DBMS since the access time to the disk is about one million times slower than access to main memory.

Two of the most important types of DBMSs are relational and object oriented. They differ in how and where data are organized. In a relational DBMS, data are indexed by what is deemed to be the most important component of an n-tuple of data sets. A telephone book is an excellent example. It is indexed by family name. This makes looking up Joseph Drake very easy. We can then look elsewhere in the n-tuple

Family Name	Given Name	4-Digit Location
-------------	------------	------------------

for his telephone number. On the other hand, if we ask the address corresponding to the distress call just received from the number 369 0462, the search of the telephone book is more difficult. The specification of which data to look for is called a query. In object oriented databases, the data are organized in terms of logical groupings called objects that may exist at various levels. There may also be links between objects, e. g.

ROBIN	is a	BIRD	is a	CHORDATE	is a	ANIMAL
-------	------	------	------	----------	------	--------

Once we know the object is a robin, we know it inherits the properties of all animals, e. g. DNA/RNA.

2.2 Prior work

The only previous paper addressing POHM DBMSs as a whole was by Berra *et al.* [4]. This paper outlined the problem and suggested the use of optical correlation as a key technology. This paper is a significant extension and development of that one Berra *et al.* first consider the ways that optics can play an important role in the increase in performance of massively parallel optical DBM and present a series of storage strategies, access strategies, and processing of data prior to conversion to electronic form. For the POHM, Fig. 2 shows their suggestion of how to select and readout pages. The deflector addresses a single hologram on the POHM and a page of data is written in parallel onto the output laser beam. An optically addressed SLM is placed at the output as an image amplifier to read the page into the optical system in parallel. A modulation pattern (intensity, phase, polarization, etc.) is produced by the SLM. The output beam comes from the portion of the SLM that is illuminated. If we illuminate the SLM addressed by the POHM with light from a second SLM (electrically addressed, intensity

modulated), we can restrict entry into the optical system to those portions of the page of direct interest as shown in Fig. 3. A Joint transform correlator (JTC) is used to search the illuminated portion of the retrieved page in parallel for space-invariant pattern recognition. In Fig. 4, electronically addressed SLM 3 is used to write a reference pattern that is matched with the light coming from POHM-driven SLM 1. The lens 1 jointly Fourier transforms both images that must be illuminated by the same laser in such a manner that they are mutually coherent on SLM 4. They strike SLM 4 which is read out by another laser beam. The laser beam, after reflection from the SLM 4, is Fourier transformed to produce an output that resembles the input page but is bright only where the reference pattern appears in the page.

One of the approaches of selecting one of an array of holograms is to allow each hologram to store multiple images. For this method presented here, the JTC is a key factor that affects the system greatly. But we will show later in this paper that JTC may not correctly be as useful for ODBMSs as sequential transform correlator (STC).

In the mean time another group has begun to design smart pixel systems for doing logical operations in a page parallel fashion [5,6]. These researchers designed the system that performs the selection function before the data is converted to an electronic signal, thus saving time and preventing bottlenecks. The system shown in figure 5 consists

of two distinct parts. (1) The projection system filters out any information that is definitely unnecessary. Both the projection mask and data input are laid out in arrays: rows of records separated into columns of fields. The mask contains 1's across each field that is to be included in the search, and 0's across the unwanted fields. Beamsplitters combine the input and mask data on their own detectors on an array of AND gates: the integrated VCSEL for that pixel will only be turned on if both mask and data show a 1. The output from the VCSELS, which contains the pertinent fields from the database records, is then stored in an optoelectronic random access memory. (2) The selection system works in a similar way, and in parallel with the projection module. First the input data is filtered via another mask and an AND gate array so that it only contains the fields being searched. The output from the AND array VCSELS is combined with the selection argument on an array of XOR or "exclusive OR" pixels. If the argument is 1 and the data is 0 or the data is 1 and the argument is 0, then the VCSEL will turn on. If the argument and data are the same, then the VCSEL will remain off. The cylindrical lens combines the outputs from all of the VCSELS in each row onto a single photodetector. If the photodetector detects little or nothing, then the row represents a match, and associated data is captured through the projection part of the system is turned into an electronic signal and sent into the computer. If the detector receives a large amount of light, then

there are many non-matching pixels in the row and the information is discarded. After all of the data from a particular page has been used or discarded, then the next page can be read into the system.

A major barrier to make such a system practical is how to produce the smart pixel arrays, such as 100x100 arrays. If the array size becomes large, electrically connecting the pixels with the remotely placed output devices becomes difficult and defeats the purpose of the optical interconnects. Such components can play a critical role in a POHM DBMS. We understand that prototype versions have been made.

2.3 Fourier optical correlators

Optical pattern recognition, of course, is the core activity in database management for POHMs. We must (1) identify target encrypted patterns, (2) discriminate against other patterns, and (3) locate the pattern(s) so identified. None of those tasks is easy. Clearly we need an operation which is (1) limited in speed only by I/O (input/output), (2) capable of operating on whole pages in parallel (space-invariant filtering), and (3) readily reprogrammable to recognize new targets. The only systems we know which can do this are Fourier optical pattern recognizers.

Fourier optical pattern recognizers can be classified as either sequential or joint transform systems. Although sequential transform

systems occurred earlier than his work, A. B. VanderLugt [7,8] is the individual who popularized this method by introducing holographic matched filters between the two successive Fourier transform systems. There is a physical mask which implements a linear discriminant on the Fourier transform of the input. The other popular form of optical correlator is a JTC. The masking is virtual, being accomplished by interference between the Fourier transforms of the input object wavefront and the reference object wavefront. Here is a comparison of the masking in the two methods.

	Sequential Transform	Joint Transform
Mask Implementation	Physical	Virtual
Preparation of Mask	Prior to Correlation	Simultaneously with Correlation

The mask (real or virtual) can vary from a simple matched filter to an on-line evolved optimum (by any arbitrary measure) filter.

We prefer the sequential transform system, because it has both speed and signal-to-noise advantages over the joint transform system in practice. Probably an electronically addressed SLM is the best choice as a mask. The matched filter is easy to calculate electronically by fast Fourier Transform (FFT). The electronic addressing allows either fast, local optimization by steepest descent or slow, global optimization by genetic algorithms. Electronic addressing also makes rapid recall of

masks from memory convenient. The details of our proposed sequential transform correlator are given in Sec. III.

III SYSTEM OVERVIEW

The optical DBMS for POHMs is shown in Fig. 6. Users require access the POHM for querying, updating, and generating reports. The DBM primarily exists for their use. The query is provided in electronic format. Of course, the POHM can not execute such a search. A dedicated computer works as a query controller. It has three tasks:

1. Recognize the query as a special case of something it knows how to do. That is, it must generate a virtual algorithm.
2. Convert the virtual algorithm into a physical algorithm, a series of physical acts to be performed.
3. Carry out the physical algorithm.

The query controller controls page accessing. The output light pattern must be operated upon by the DBMS hardware. As light does not operate on light, some electronics must be involved here. It seems unlikely that we can work out a complex query in parallel. We can not do complex queries on multiple pages in parallel. Thus some sort of scratchpad memory is required. Finally, the answer to the query must be provided in a convenient electronic format. We imagine that this

information will be read out from the scratchpad memory where it has been accumulated and processed.

It checks the cache. The output of query is normally kept in cache, on a temporary basis. When we need a particular piece of information, we first check whether it is in the cache. If it is, we use the information directly from the cache; if it is not, the output of searching the POHM will be put in the cache. Then we can obtain what we want in the cache. A copy of the output will be put in the cache under the assumption that there is a high probability that it will be needed again. If the particular information is not in the cache, the computer will send three signals to POHM, flash addressed random access memory (RAM) and optical pattern recognizer (OPR), respectively.

The POHM has been described in Sec. 1. A POHM will be able to store gigabits of information in the form of paper-2D arrangements of data (usually bits). For example, a 0 might be represented by the pair 0,1; and a 1 by the pair 1,0. The POHM can call forth any of its stored parts in random access mode very quickly, 10^{-5} sec to 10^{-9} sec, depending on the system. The whole page appears optically in parallel, with the 1 positions represented by light and the 0 positions by the absence of light. There are many forms of data organization in current use. The POHM will inspire new ones [9]. We could certainly group similar items on a page if we could guess which variable will be searched most

frequently. The query controller could consult its page index electronically and access only the useful pages. The signal to POHM is to start search of the POHM. Two morphological discriminants among search procedures can be defined here. First, we have the degree of parallelism. Disk based systems produce a bit at a time (for each head). A POHM produces an entire page. Some of our methods will search a bit at a time, a page at a time, and perhaps even a POHM at a time. In general, the greater the parallelism, the faster will be the rate of search. Second, there is a strategy choice for complex queries. We can do the full complex search on each page before moving to the next. Alternatively, we can do a less-than-full search on all of the pages sequentially. There is no general better choice. It depends largely on the time needed to prepare to search one page. If the time is long compared with the time to switch between pages and if the whole POHM must be searched, scanning through the pages is probably faster. The output light pattern will be used in OPR.

In a RAM, any addressable location in the memory can be accessed at random. That is, the process of reading from and writing into a location in a RAM is the same and consumes an equal amount of time no matter where the location is physically in the memory. With an optically flash addressed RAM, we can enter the whole page in parallel and read only the preselected data into the cache memory.

After receiving the signal from the computer, the OPR will prepare to realize the pattern selection. A physical recognition mask will be used. If the mask can be an electronically addressed SLM, it can be evolved on line. Masks can be stored electronically and called into the SLM from memory quite quickly. In other cases we will prefer an optically addressed SLM. We must avoid serializing the data on a page. An SLM is used to write the page onto a light beam. SLMs can be optically addressed, so serialization can be avoided. The output of OPR is then projected onto a 2-D charge-coupled device (CCD) array that senses light-and-dark patterns and produces currents in response to the patterns, thereby reading all retrieved data. Now the data is in electronic form and send to cache through application specific integrated circuit (ASIC).

In optical processing there are also virtual algorithms and physical algorithms. We sometimes call the latter “archirithms” or “algotecturers” as they are virtual algorithms embodied in an architecture. For parallel search of an optically represented data page, the virtual algorithm of choice is Fourier optics. The virtual algorithm is to take two successive optical Fourier transforms of an input coherently illuminated pattern. This is an imaging system which turns the input amplitude pattern $f(x,y)$ into an output image $f(-x,-y)$. In the middle of this imaging system is a spatial display of

$$\begin{aligned}
F(u, v) &= [f(x, y)] \\
&= \int f(x, y) e^{-2\pi i(xu + yv)} dx dy
\end{aligned}$$

the Fourier transform of $f(x, y)$. We now use the shift theorem

$$[f(x - x_0, y - y_0)] = e^{-2\pi i(x_0 u + y_0 v)} F(u, v)$$

That is, the shape and phase of the Fourier transform is independent of (x_0, y_0) -the displacement of the input. The displacement information retained in a simple phase factor. Suppose now we multiply $F(u, v)$ by a pattern $M(u, v)$ which factors the Fourier transform of the pattern we seek and attenuates other patterns. Then the image plane will contain bright points everywhere the desired pattern occurs.

There are two physical algorithms. Both use a lens to perform the Fourier transforms as shown in Fig. 7(a) and in Fig. 7(b).

$M(u, v)$ is a physical mask. This system is widely used. The mask can be an electronically addressed SLM which can be evolved on line. Masks can be stored electronically and called into the SLM from memory quite quickly. We will call this the physical mask system. The other approach places the search-for pattern next to the input pattern and Fourier transforms them both onto a detector plane that nonlinearly detects the sum of those patterns. This creates product terms-the ones needed in the virtual algorithm. This detected pattern then Fourier transformed to

get the desired output pattern. This is called the joint transform method.

We can now compare the effectiveness of the two methods. The physical mask method and the joint transform method can both be optimized on line. The two methods both require two Fourier transform lenses. Both involve a detector array and a system to locate the maxim. Both need two SLMs. There are major problems with the joint transform method when there are many inputs, in that the dynamic range must be shared among all correlations. There are major problems with the physical mask system if the mask is to be evolved quickly. There are, as of today, no good algorithms for this. On balance and at this moment, the dynamic range problem with the joint transform method appears more troublesome than the mask evolution problem with the physical mask method.

There are primitive physical algorithms for pages

Algorithm 1 is called “page search.” We evolve a physical mask to recognize the desired pattern wherever it occurs.

Algorithm 2 is a variation called “limited page search.” We limit the search part of page by blocking the other parts with an SLM or restricting our attention to the key regions of the detector array.

Algorithm 3 incorporates limited page search into a very powerful and very general algorithm which we call “logical merge.” Suppose we

detect two sets of data A and B. We can draw, symbolically, the relationship between A and B as a Venn diagram. See Fig. 8.

We can then form all sorts of logical merges such as A OR B, A AND B, A AND NOT B, etc. As A and B are found as 2D maps sequentially. We can do the merging in various ways. Although there are various ways to do this optically, it is probably easier to scan both maps synchronously and do the logical operations point-by-point electronically. It is easy to use one pass through both masks (stored on scratchpads) to mark patterns that satisfy the required logical pattern. That pattern can be written back into the scratchpad pattern to allow the formation of more complex patterns such as (A OR B) AND NOT C or to store for subsequent output. To store outputs efficiently, we must recall all of the indicated information selected as an answer to the query. Of course, we want to compact this data-leaving no blank spaces corresponding to data not selected. Again, this is a task probably best done electronically.

The operations just described are for single pages. Now we consider the operations on multiple pages. Suppose we want to find A AND B page by page. Alternatively, we could find all of the A-matching items on all pages and remember their locations electronically. We can then search only those regions for B. This reduces the search time significantly only if there are pages on which A items are absent. If we have a relational data base, for example, this can be assured ahead of

time. A phone book is a relational data base. If we want to find all the Smith with phone numbers beginning with the prefix 369. It makes sense to locate the Smiths first, then search the prefixes. It would make no sense at all to search for the 369 first. We have reason to believe that the 369s may be fairly randomly distributed. In general, especially when dealing with ANDs, it is desirable to search in this order:

- items likely to be concentrated on a small fraction of the pages,
- items expected to occur at an average rate of less than one per page,
- all other “fairly-unlikely” events, and
- all other events.

Another approach is to accumulate all the As into a compact set on one “page” of cache memory. We can then search this page for B in parallel optically. Depending on the physical speeds of the operations, this is likely to be the fastest multiple page approach.

IV Critical Component analysis

4.1 Optical Hardware Choices

The heart of the ODBM system is an optical pattern recognizer. There are two such generic systems: (a) the sequential transform correlator

(STC) often also called the $4f$ correlator, even though a shorter $2f$ version is known and (b) the JTC discussed earlier.

The STC optically performs two Fourier transforms in sequence. The first lens Fourier transforms the input amplitude imposed by the mask (the data to be searched). In the Fourier domain the transformed data is multiplied by a second mask (containing the complex conjugate of the Fourier transform of the pattern to be identified). The product is then Fourier transformed yet again to produce the correlation.

The JTC differs in that there is only one Fourier transform lens. The data and the search pattern are introduced side-by-side in the input plane. The first Fourier transform is performed. The intensity pattern in the Fourier plane is captured into a computer that may perform nonlinear operations on it. The resulting pattern is then reintroduced into the optical system. It is again transformed to produce the final output.

The advantages of a STC over a JTC are: (1) no intermediate processing, higher potential speed and (2) a higher signal-to-noise ratio. In a JTC, we must record the virtual masks for each object in the same medium. For multiple objects this places great strains on the available recording media.

The advantages of a JTC over a STC are: (1) More compact, and (2) no pre-stored filters.

We prefer a system that combines the advantages of the JTC and STC. A single step system, similar to the JTC can be used to generate filters near real time. These filters can then be inserted into the STC system for pattern recognition.

The POHM serves as the source of the database data, while the search pattern is usually composed on the computer. The addressing speed of the POHM is more important than that of the filter SLM, since the search pattern does not change as frequently as the data to be searched. The method of addressing individual pages in the POHM depends on the way the data is initially stored. For example, if it is angle multiplexed, an acousto-optic modulator would be a good choice. The speed of filter insertion is limited by the (serial) electronics to (parallel) optics bottleneck. The only way this could be avoided is to pre-store filters. This option is unattractive, since it comes at a loss of generality.

4.2 Encryption

To make the most efficient use of electro-optical input transducers and the storage capacity of volume holographic memory, we need to apply an efficient encoding scheme. Pattern recognition has long strived to “read” human writing, be it hand-written or machine generated form. This approach to encoding works well, most of the time, for humans, but

electronic and optical processors still struggle to read. Text has a number of problems:

1. A number of similar symbols, F is often contained within E, O within Q, etc.,
2. Symbols that are similar to upside down letters, A and V, and u and n are very difficult for optical correlators to distinguish, and
3. It takes a lot of information to form a letter ($7 \times 9 = 63$ pixels minimum).

Any encoding scheme we choose has to be able to cleanly tell apart any symbol in the character set, while reducing the data bandwidth required of the system.

We want to reduce the number of false identifications, even in the presence of noise. There are numerous schemes to encode symbols into orthogonal one- and two-dimensional spaces. We can use these as a starting point. We then stipulate that an auto-correlation should yield a higher peak than any cross-correlation. In effect, we require a Hamming distance proportional to the true-to-false ratio desired. There is a trade-off, however. The greater the Hamming distance, the more data have to be stored and transferred for a given amount of information.

Most text-based data is in 7- or 8-bit character sets (ASCII and extended ASCII), let us assume an 8-bit format. The Hamming distance should be chosen relative to the size of this set. In this case a Hamming

distance of 3 is the minimum functional value, while 6 is quite adequate. We used a direct computer search to select a 7-bit character set for a Hamming distance of 6. This requires a 16 bit code sequence. In Table 1 we show some examples:

The best way of arranging these 1D codes in 2D is to use linear code strings separated by blank space, combined with 1D Fourier transform lenses.

4.3 Computer and Bus Considerations

The speed of the ODBM system is a function of several factors:

1. the electro-optic transducers,
2. the optical components,
3. the opto-electronic transducers, and
4. the front end host computer.

The electro-optic transducers, usually SLMs, are becoming increasingly fast with rates in the kHz regime. The speed of the optical passive components is never an issue. The opto-electronic transducers, CCD arrays, etc., can also be operated at high frame rates.

These high frame rates show the real limitation of a useful optical processor: the drive electronics. The data bandwidth of the optical system components is much greater than that of the electronic system components. If we assume a 512 X 512 array with a bit-depth of 8,

each frame is 0.25 MB. At the normal video frame rate of 30, this is still a manageable 7.5 MB/sec, but at a frame rate of 1 kHz the required bandwidth becomes 250 MB/sec. One can use data compression to reduce the required bandwidth somewhat, but even for a 1 kHz frame rate the resulting bandwidth is well above what is available on present-day PCs. Some buses such as EISA, ISA, NuBus and VME cannot handle the bandwidth of video rate data, not to mention 1 kHz frame rates. The PCI bus (32 bit/33 MHz) is the best we can do with non-proprietary technology. It can handle simultaneous read and write of 40 MB/sec each. With PCI bus systems one can hope to achieve frame rates between 200 and 400 per second. Proprietary standards such as found Silicon Graphics multiprocessor servers, can exceed 1 GB/sec. These proprietary standards, however, cannot be used as a basis for interface cards.

V Conclusions

In this paper we have considered DBMSs for POHM, which is the only practicable storage for ultra high speed VLDB. Parallel optical processing is used for database management of parallel optical data from a POHM. There are two features in database operations: (1) retrieval, update, and process of a large collection of data to generate the query result; (2)

comparison of the data fields with a search argument. In electronic computers, the segments of the database are loaded from the secondary storage to the main memory and the query criterion is applied. Because secondary storage is typically orders of magnitude slower than main memory, the retrieval of data constitutes a major bottleneck that makes database processing an input-output-bound task and causes significant delays to the overall system response. We have shown that the optical DBMS for POHM have considerable potential to solve the bottleneck by employing Fourier optical pattern recognition. We choose to use STC because it has both high speed and high signal-to-noise advantages.

Since the data are already in optical form, the optical parallel DBMS for POHM processes it optically before conversion to sequential electronic form. This will have significant performance advantages, especially as data can be read from storage and queried at hundreds of megabytes per second.

VI Acknowledgments

This work is supported by SEA-DOE under contract no. DE-FG05-94Er25229.

VII References

- [1] N. N. Vyukhina, I. S. Gibin, V. A. Dombrovsky, S. A. Dombrovsky, B. N. Pankov, E. F. Pen, A. N. Potapov, A. M. Sinyukov, P. E. Tverdokhleba, and V. V. Shelkovnikov, "A review of aspects relating to the improvement of holographic memory technology," *Opt. & Laser Tech.* **28**, 269-276 (1996)
- [2] C. Denz, T. Dellwig, J. Lembcke, and T. Tschudi, "Parallel optical image addition and subtraction in a dynamic photorefractive memory by phase-code multiplexing," *Opt. Lett.* **21**(4), 278-280 (1996)
- [3] T. Baer, "Relational Technology in the land of the giants," *Data Management/DBMS Software Magazine*, 61, Feb. (1996)
- [4] P. B. Berra, K.-H. Brenner, W. T. Cathey, H. J. Caulfield, S. H. Lee, and H. Szu, "Optical database/knowledgebase machines," *Appl. Opt.* **29**(2), 195-205 (1990)
- [5] F. R. Beyette, Jr., K. M. Geib, C. M. St. Clair, S. A. Feld, and C. W. Wilmsen, "Optoelectronic Exclusive-Or Using Hybrid Integration of Phototransistors and Vertical Cavity Surface Emitting Lasers," *IEEE Photonics Tech. Lett.* **5**(11), 1322-1324 (1993)

- [6] P. A. Mitkas, L. J. Irakliotis, F. R. Beyette, S. A. Feld, and C. W. Wilmsen, "Optoelectronic data filter for selection and projection," *Appl. Opt.* **33**, 1345-1353 (1994).
- [7] A. B. VanderLugt, "Signal Detection by Complex Spatial Filtering," Radar Lab., Rept. No. 4594-22-T, Institute of Science and Technology, The University of Michigan, Ann Arbor (1963)
- [8] A. B. VanderLugt, "Signal Detection by Complex Spatial Filtering," *IEEE Trans. Inform. Theory* **IT-10**(2), 139-145 (1964).
- [9] A. Louri, and J. A. Hatch Jr., "An optical associative parallel processor for high-speed database processing," *Computer*, 65-72, Nov. (1994)

VIII Figure Captions

Fig. 1. The POHM concept is to vary a parameter P and record a hologram of a page with each P . Then when we restore the P value, we recall the corresponding page.

Fig. 2. Selection and readout of pages

Fig. 3. Output Selection

Fig. 4. Pattern Selection


Fig. 5. Optoelectronic data filter

Fig. 6. Optical DBMS for POHMs

Fig. 7. Two physical algorithms for Fourier transforms

Fig. 8. Intersection of data A and B

Table 1 Examples of 1-D Hamming codes



P(PARAMETER)

NAME OF POHM SYSTEM

Spatial Position
 Angle
 Wavelength
 Phase Pattern
 On Input SLM

Space Division Multiplexing
 Angle Division Multiplexing
 Wavelength Division Multiplexing
 Code Division Multiplexing

Fig. 1.

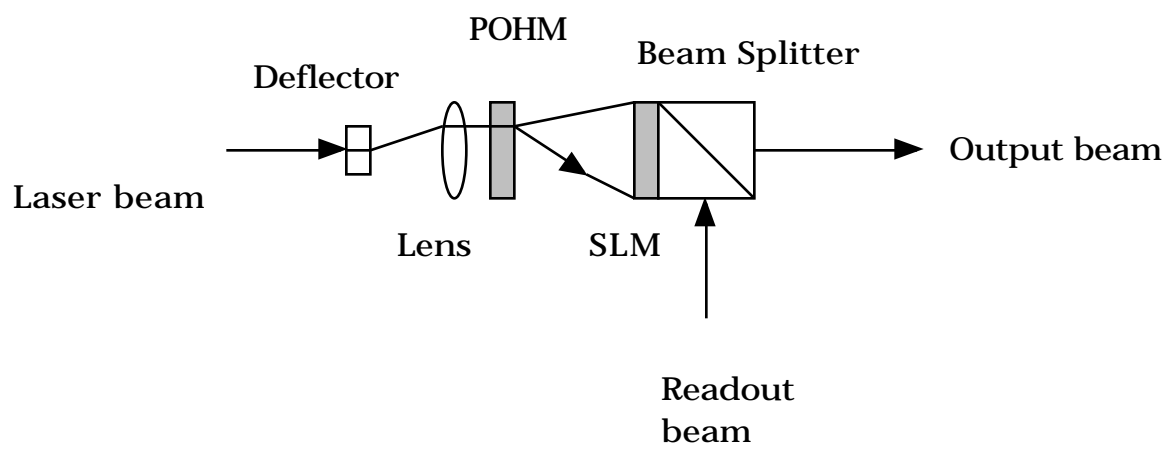


Fig. 2.

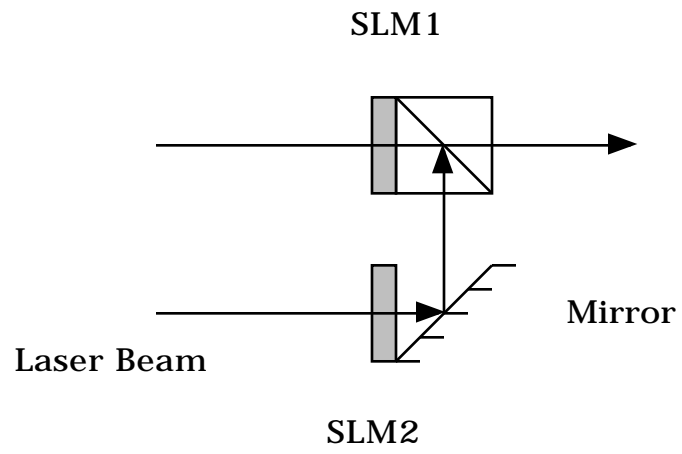


Fig. 3.




Fig. 4.

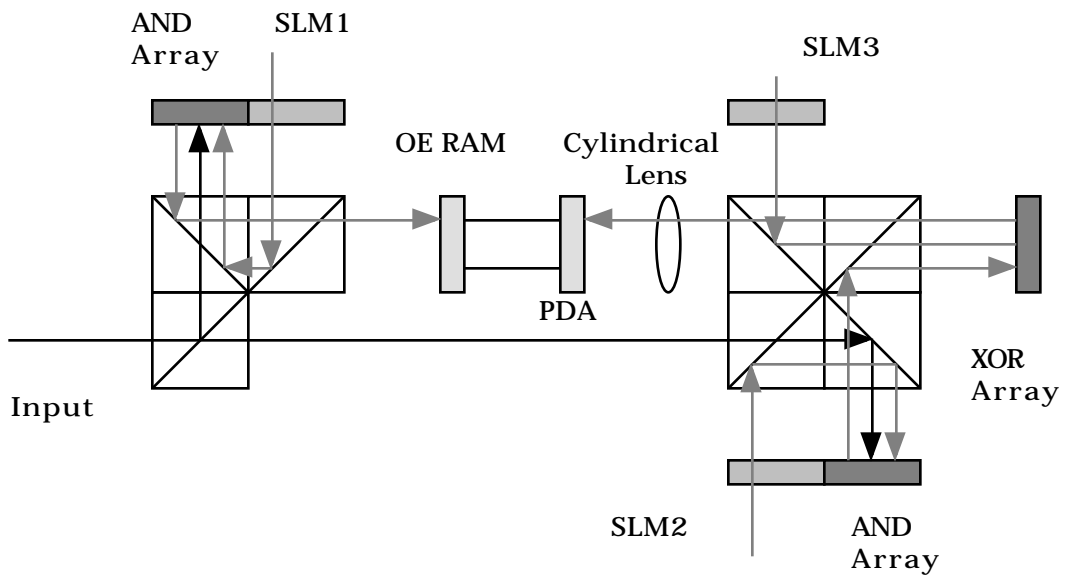


Fig. 5.




Fig. 6.




Fig. 7.




Fig. 8.

0)	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1)	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1
2)	0	0	0	0	0	0	0	1	1	1	0	0	0	1	1	1	1
3)	0	0	0	0	0	0	0	1	1	1	1	1	1	1	0	0	0
4)	0	0	0	0	0	1	1	0	0	1	0	0	1	0	1	1	1
5)	0	0	0	0	0	1	1	0	0	1	1	1	0	1	0	0	0
6)	0	0	0	0	0	1	1	1	1	0	0	0	1	1	0	0	0
.																	
.																	
.																	
121)	1	1	1	1	1	0	0	0	0	0	1	0	0	1	1	0	0
122)	1	1	1	1	1	0	0	1	1	1	0	1	1	1	1	0	0
123)	1	1	1	1	1	0	0	1	1	1	1	0	0	0	0	1	0
124)	1	1	1	1	1	1	1	0	0	1	0	1	0	0	1	0	0
125)	1	1	1	1	1	1	1	0	0	1	1	0	1	1	0	1	0
126)	1	1	1	1	1	1	1	1	1	0	0	1	0	1	0	1	0
127)	1	1	1	1	1	1	1	1	1	0	1	0	1	0	1	0	0

Table 1