Multiwavelength optical content-addressable parallel processor for high-speed parallel relational database processing

Peng Yin Choo, Abram Detofsky, and Ahmed Louri

We present a novel, to our knowledge, architecture for parallel database processing called the multiwavelength optical content-addressable parallel processor (MW-OCAPP). The MW-OCAPP is designed to provide efficient parallel data retrieval and processing by means of moving the bulk of database operations from electronics to optics. It combines a parallel model of computation with the many-degrees-of-processing freedom that light provides. The MW-OCAPP uses a polarization and wavelength-encoding scheme to achieve a high level of parallelism. Distinctive features of the proposed architecture include (1) the use of a multiwavelength encoding scheme to enhance processing parallelism, (2) multicomparand word-parallel bit-parallel equality and magnitude comparison with an execution time independent of the data size or the word size, (3) the implementation of a suite of 11 database primitives, and (4) multicomparand two-dimensional data processing. The MW-OCAPP architecture realizes 11 relational database primitives: difference, intersection, union, conditional selection, maximum, minimum, join, product, projection, division, and update. Most of these operations execute in constant time, independent of the data size. We outline the architectural concepts and motivation behind the MW-OCAPP's design and describe the architecture required for implementing the equality and intersection–difference processing cores. Additionally, a physical demonstration of the multiwavelength equality operation is presented, and a performance analysis of the proposed system is provided.

OCIS codes: 200.2610, 200.3050, 200.4540, 200.4560, 200.4860, 200.4960.

1. Introduction

Databases are emerging as the most important ingredients in information systems. They have penetrated all fields of human endeavor and are no longer limited to business-oriented processing. Searching, retrieving, sorting, updating, and modifying nonnumeric data such as databases can be significantly improved by the use of content-addressable memory (CAM) instead of location-addressable memory.[1–4] CAM-based processing is not only more akin to the way database users address their data (in parallel), but it is also faster than location-addressing schemes, since the overhead cost of address computations is completely eliminated. An electronic CAM, however, faces several serious obstacles. First, an electronic CAM has a serious limitation in terms of moving data in and out of the processor. If a search request results in multiple matches, each of these matched words must be read out serially. This input–output (I/O) bottleneck becomes increasingly problematic as the CAM array size scales upward. Second, the traditional electronic CAM faces many technologically daunting issues, such as high bit cell complexity, low storage densities, clock skew, and interconnect latencies. Because of these weaknesses, electronic CAM's have been included in computers systems only as small auxiliary units.[4,5]

Optics can alleviate the cell complexity of CAM-based storage cells by migrating their interconnects into the third dimension.[6] The high degree of connectivity available in free-space and fiber-based optical interconnects and the ease with which optical signals can be expanded (which allows for signal broadcasting) and combined (which allows for signal funneling) can also be exploited to solve the interconnection problems. The multidimensional

The authors are with the Department of Electrical and Computer Engineering, the University of Arizona, Tucson, Arizona 85721. A. Louri's e-mail address is louri@ece.arizona.edu.

Received 27 October 1998; revised manuscript received 7 June 1999.

0003-6935/99/265594-11$15.00/0

nature of optical systems presents an additional degree of freedom (spatial) for the design of parallel I/O than do pure electronic I/O systems.[3,7] Data can be retrieved and stored as pages in optical storage systems such as page-oriented holographic memory, volume holograms, or optical disks.[8,9] The use of optics greatly reduces the disparity between CAM-based processing and data I/O found in electronic systems.

Several optoelectronic database processing systems have been proposed in the literature. The optical data filter and related systems[10–14] utilize smart pixel technology that merges the parallel interconnectivity of optics with the fast switching speed of electronics. Highlights of the system include simplicity of design and speed of execution. Restricting optics to mere interconnections, however, limits the amount of spatial and functional parallelism that can be extracted from the system. A novel, to our knowledge, architecture called the optical content-addressable parallel processor (OCAPP) was demonstrated that implemented a limited set of high-speed database operations.[15–17] The implemented operations were divided into parallel equivalence and relative magnitude searches and were accomplished with an intensity and polarization-encoding scheme.[18,19] Although superior to electronic approaches, the OCAPP is limited to the parallel matching of only a single word against a database. To match several words in parallel, a different approach is needed.

An architecture is now proposed that we call the multiwavelength OCAPP (MW-OCAPP).[20–22] It utilizes polarization-division multiplexing introduced in the OCAPP combined with wavelength-division multiplexing to achieve an even higher degree of parallelism and system integration. The ability to propagate multiwavelength light planes through the same space without mutual interference allows for true two-dimensional (2-D) operations to take place. This enables the processing of multiple data arguments at the same time and within the same space, thus greatly increasing the parallelism of the processor over single-wavelength designs. Additionally, utilizing these additional degrees of freedom allows for a more compact system to be created, thus simplifying the manufacturing process and relaxing the alignment and power constraints.

2. Overview of the Multiwavelength Optical Content-Addressable Parallel Processor System Architecture

Figure 1(a) illustrates a structural organization of the MW-OCAPP processing model. The optical processing logic of the MW-OCAPP consists of six modules: an optical selection unit, an optical match-compare unit, an optical equality unit, an optical magnitude comparison unit, an optical relational operations unit, and an optical output unit. The architecture is designed such as to implement a total of 11 database primitives. Most of these execute in a time span that is independent of the problem size. The MW-OCAPP can realize difference, intersection, union, conditional selection, join, maximum, minimum, product, projection, division, and update.

The inputs to the MW-OCAPP are the comparand array (CA), which contains the database search arguments, and relational array (RA), which contains the database to search against. Each row (tuple) of the CA in the selection unit is polarization logic encoded with different wavelengths by a multiwavelength source array and an electronically addressable spatial light modulator (EASLM). This form of encoding allows for the superposition and parallel processing of multiple comparands as they propagate through the match-compare and equality units. The selection unit produces the selection register (SR), a light plane that holds the multiple tuples in the comparand array to be matched. The optical match-compare unit produces the match-compare register (MCR), a wavelength–polarization encoded light plane that holds the locations of all the matched and mismatched bits. The magnitude comparison unit takes the CA and the RA and performs a magnitude comparison (greater than and less than) between CA and RA tuples and outputs the less-than register (LR) and the greater-than register (GR). The optical equality unit takes the MCR and produces an output called the equality register (ER) that represents the intersection locations of the CA and the RA tuples. The ER, LR, and GR light planes pass through the optical relational operation processing unit where they are operated on to produce the relational operation registers. Lastly, the relational operation register and the MCR light planes are routed through the optical output unit that interfaces with a host computer.

The equality operation, which is rooted in the exclusive-OR (XOR) operation, is one of the most basic operations that the MW-OCAPP can perform. If a CA word and a RA word are a match, a bitwise XOR operation will produce a resultant word containing only logical 0 bits. If the CA and the RA do not match, the resultant XOR word will be a mixture of 0 and 1 bits. Determining equivalency results from simple logical ORing of all of the bits together in this intermediate word. The two words are mismatched if the result is a 1, and likewise the words are equivalent if the result is a 0. The optical equality unit operates under these same basic principles, just in a much more parallel fashion.

Magnitude comparison is the second fundamental operation that the MW-OCAPP implements, and it is shown schematically in Fig. 1(b). Operations of this type include the greater-than, less-than, inbound, out-of-bound, and extremum tests. The input to the optical magnitude comparison unit is the ER from the equality unit, and the output is the LR and the GR.[4] The algorithm that the MW-OCAPP uses to perform this operation can be decomposed into four steps: (1) Compute and store the com-

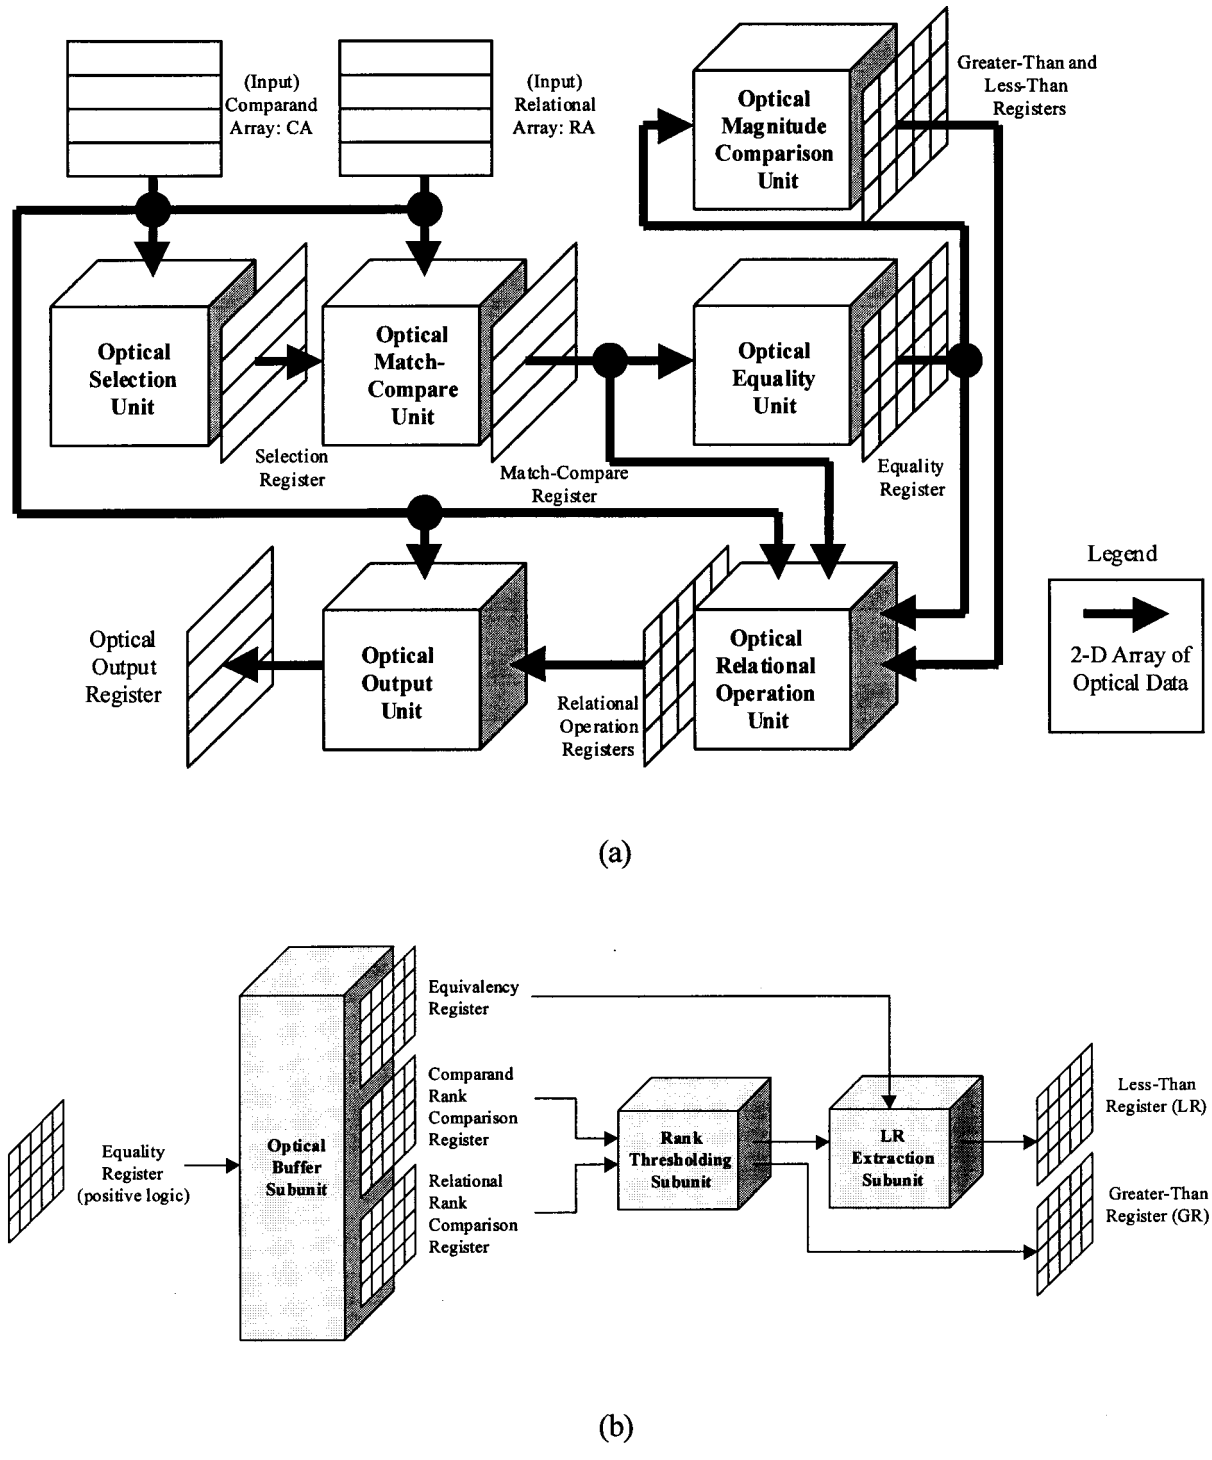

Fig. 1. (a) MW-OCAPP and (b) magnitude comparison schematic organizations.

parand rank comparison register, comparing the CA with a rank table; (2) compute and store the relational rank comparison register, comparing the RA with a rank table; (3) compute and store the equivalency register, comparing the CA and the RA; (4) compute and output the LR and the GR. Tasks (1)–(3) are processed by the optical selection, match-compare, and equality units, and the results (the ER's) are stored sequentially in the optical buffer subunit. Step (4) of the algorithm requires that these three registers (termed the equivalency register, the comparand rank comparison register, and the relational rank comparison register) be presented simultaneously to both the rank thresholding subunit and the LR extraction subunit for final processing. Details of the optical implementation of the magnitude comparison unit can be found in Ref. 23.

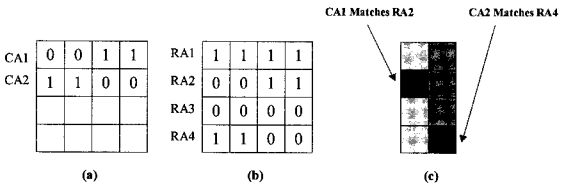

Fig. 2. Equality operation example matches two tuples in (a) the CA with the four tuples found in (b) the RA. The ER result (c) indicates that there is a match between CA1 and RA2 as well as between CA2 and RA4. Nonilluminated (black) pixels indicate an exact match.

3. Multiwavelength Optical Content-Addressable Parallel Processor

A. Data Encoding

The MW-OCAPP uses several methods for encoding a data plane on a light plane. Binary patterns are represented by spatially distributed orthogonally polarized locations on a 2-D pixilated grid.[19,20] Logical 1 is defined as vertically polarized light and logical 0 as horizontally polarized light. The presence or absence of light (intensity threshold) within a light plane indicates the selection or deselection of tuples or attributes in the system. Individual tuples are differentiated from one another by polarization encoding of each on a unique wavelength.

B. Implementation of the Equivalence Operation

The example to follow is based on the data planes found in Fig. 2. Two words in the CA are matched simultaneously against two words in the RA by use of the optical selection, match-compare, and equality units. The following descriptions of the three optical processing units are based on this example.

The optical selection, match-compare, and equality units are shown combined in Fig. 3. Together they perform the equality operation. The selection unit's purpose is to encode a pixilated 2-D optical wave front with the CA to be processed. The rows in the wave front, each encoded on a unique wavelength, represent tuples in the CA. Polarization encoding of the desired data pattern is employed to differentiate the binary states of each of the pixilated bits.

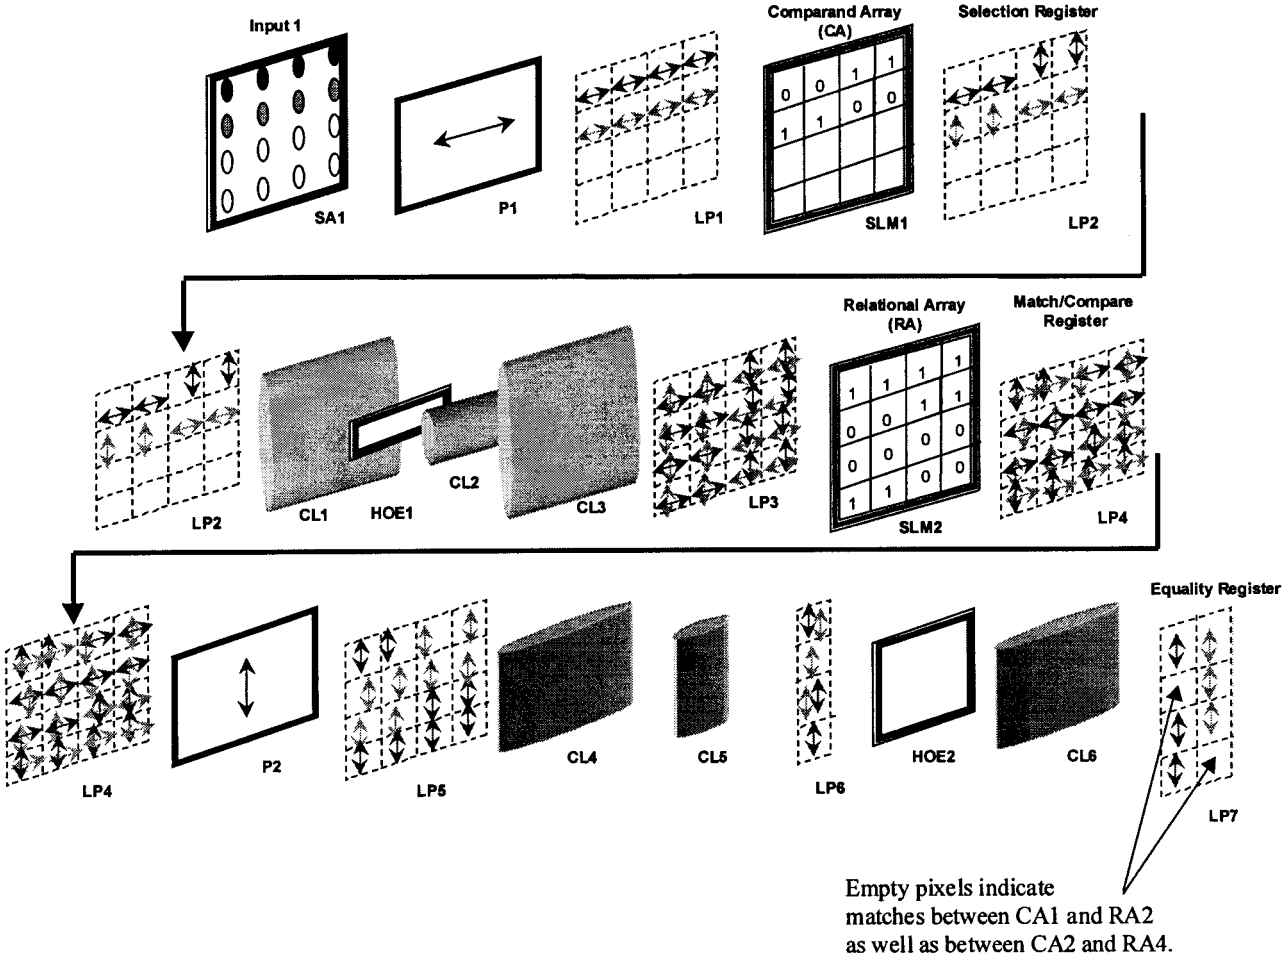

Fig. 3. Selection (first row), match-compare (second row), and equality (third row) units that implement the equality operation. Note that the system correctly predicts a match between CA1 and RA2 as well as between CA2 and RA4 [Fig. 2(c)]. (SA, source array; CL, cylindrical lens; P, polarizer; HOE, holographic optical element; LP, light plane).

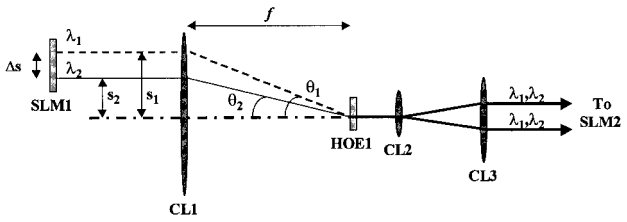

Fig. 4. HOE1 grating multiplexer from Fig. 3 takes each row from SLM1 that is encoded on different wavelengths and spatially multiplexes it. HOE1 is placed at the focal plane of CL1. CL2 and CL3 are an afocal doublet that takes this multiplexed channel and expands it such that it fills an aperture equal to the diameter of SLM2.

The selection unit begins with a multiwavelength source array (SA1) in which each row (which corresponds to a separate tuple) radiates at a different wavelength. Two wavelengths are required for the example; so only the first two rows are selected to radiate. This wave front passes through a horizontally oriented polarizer (P1) to reset all the bit positions to the 0 logical state (LP1). Light plane LP1 impinges on an EASLM (SLM1), which polarization encodes the light passing through it with the bit patterns 0–0–1–1 and 1–1–0–0. The resultant light plane, LP2, is called the SR and represents the optically encoded version of the comparand array.

The match-compare unit is illustrated in the second row of Fig. 3. Its purpose is to bitwise (XOR) each tuple in the comparand array with each tuple in the RA. This produces a logical 1 at every bit position where there is a CA and RA mismatch. The SR, LP2, passes through a holographic optical element and through lenses CL1, CL2, and CL3, which duplicate each of the rows corresponding to different wavelengths over the full surface of an EASLM (SLM2). HOE1 is a transmissive planar grating multiplexer[24] and is illustrated with SLM1, CL1, CL2, and CL3 in Fig. 4. This structure can be analyzed with the following grating equation:

$$d \sin \theta = m\lambda, \tag{1}$$

where d is the grating period of HOE1, θ is the diffraction angle from normal incidence, m is the diffraction order, and λ is the wavelength of the light. When the analysis is restricted to the first diffraction order, the ratio of the lateral offset distance s to the focal length f as a function of λ/d can be found with

$$\frac{s}{f} = \frac{\lambda}{d\left[1 - \left(\frac{\lambda}{d}\right)^2\right]^{1/2}}. \tag{2}$$

Given two wavelengths, λ_1 and λ_2, the change in lateral spacing (Δs) can be found with

$$\Delta s = s_1 - s_2 = \frac{f}{d}\left\{\frac{\lambda_1}{\left[1 - \left(\frac{\lambda_1}{d}\right)^2\right]^{1/2}} - \frac{\lambda_2}{\left[1 - \left(\frac{\lambda_2}{d}\right)^2\right]^{1/2}}\right\}. \tag{3}$$

The wave fronts converging on HOE1 will have an approximate spherical curvature that will produce an undesirable position-dependent divergence angle following HOE1. Given a SLM1 pixel diameter q, the full divergence angle can be approximated with

$$\theta_{\text{divergence}} = 2 \tan^{-1}\left[\frac{qf}{2(s^2 + f^2)}\right]. \tag{4}$$

As a design example, suppose that f is equal to 10 cm and that the wavelength extremes, λ_1 and λ_2, are 780 and 790 nm, respectively. If the SLM1 array diameter is 0.5 cm, we can substitute this for Δs in Eq. (3) and, using the extremum wavelengths for λ_1 and λ_2, we obtain a value of 0.97 μm for the required grating period d and $\sim$0.8 for λ/d. With Eq. (2) the ratio of s/f is 1.3, the maximum diffracted angle is 53° [Eq. (1)], and CL1 requires an equivalent $f/\#$ of 0.4. With Eq. (4) and a pixel diameter of 960 μm, the full divergence angle is a mere 0.2°. Since only a maximum of half of the lens CL1 interacts with LP2, for large diffraction angles CL1 appears like a prism. Therefore, under these circumstances CL1 may be replaced with a prism of the appropriate wedge. Also, note that the above analysis assumes that CL1 has negligible chromatic and spherical aberration. As the diffraction angle becomes large, however, spherical aberration may become significant. One can either expand CL1 into multiple elements to correct for these aberrations or introduce a lateral offset to each row in the source array such that all ray bundles converge to the same line focus.

Returning now to the match-compare unit in Fig. 3, light plane LP3 passes through SLM2, which is encoded with the RA to be searched. The EASLM rotates the polarization(s) of the incident light according to the logic states of its pixels, effectively generating the result of the logical XOR operation in LP4. Light plane LP4 is called the MCR and contains all the bit match and mismatch locations of each of the CA and RA tuple combinations (designated by horizontally polarized and vertically polarized light, respectively).

The equality unit is shown in the third row of Fig. 3. It identifies which combinations of CA and RA tuples are matches by converting the MCR to a pixilated map called the ER, which represents the equivalency of all of the CA and RA tuple combinations. The MCR (LP4) enters and passes through a vertically oriented polarizer (P2) to form LP5. LP5 contains an illuminated pixel corresponding to all bit mismatch positions. LP5 is funneled down to a single column by CL4 and CL5 and is wavelength demultiplexed into a plane that has a pixel count width equal to the number of tuples in the CA. This wavelength separation is accomplished with a transmission grating structure (HOE2). Cylindrical lens (CL6) focuses the light exiting HOE2 and produces the ER (LP7). The ER light plane is a 2-D representation of the intersection of the CA and

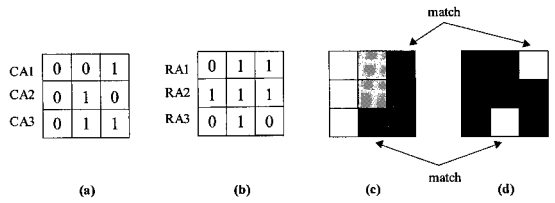

Fig. 5. As a precursor to the intersection–difference example, the ER (c) is generated from three tuples in (a) the CA and (b) the RA. The intersection–difference optical hardware requires that the ER be in a positive logic format (d), which can be accomplished with trivial hardware.

the RA. If n represents the number of tuples in the CA and m represents the number of tuples in the RA, then the ER must consist of $m \times n$ pixels. It is encoded in negative logic, meaning that nonilluminated pixels correspond to exact matches. For an $m \times n$ ER grid, pixel_{mn} is illuminated such that tuple RA_m is not equal to tuple CA_n.

C. Implementation of Selected Higher-Order Database Operations

As stated above, the MW-OCAPP is capable of performing several operations that go beyond the basic equality or magnitude comparison algorithms. Two such operations are illustrated here: intersection and difference. The intersection operation forms a new relation that consists of all tuples appearing in both of two specified relations (tables or records). The difference operation constructs a relation comprising all tuples contained in the first but not the second of two specified relations. Both of these results are generated simultaneously and with the same hardware.

Figure 5 illustrates the CA [Fig. 5(a)] and the RA [Fig. 5(b)] light planes that are processed by the hardware in Fig. 3 to form the ER [Fig. 5(c)]. It can be seen that there are two tuples that match between the CA and the RA. The intersection–difference subunit (IDS) hardware takes the positive logic version of the ER [Fig. 5(d)] and generates the intersection and difference results, namely 011, 010, and 111, respectively.

Figure 6 shows the IDS optical implementation with the ER as the input [Fig. 5(d)]. The ER light plane (LP1) enters the IDS and is funneled down to a single column (LP2). LP2 impinges on the write side of an optically addressable spatial light modulator (SLM1). A multiwavelength source array (SA1) consisting of two columns (each column emits at a different wavelength) produces a light plane that passes through polarizers P1 and P2, respectively. The polarizers set each of the light planes to orthogonally oriented polarization states. The two columns are funneled down to a single column by CL3 and CL4 to form light plane LP3 that reflects off of SLM1. LP3 contains the tuple selection information for both the intersection and the difference results, each on an individual wavelength. Light plane LP4 containing the tuple selection information is spread

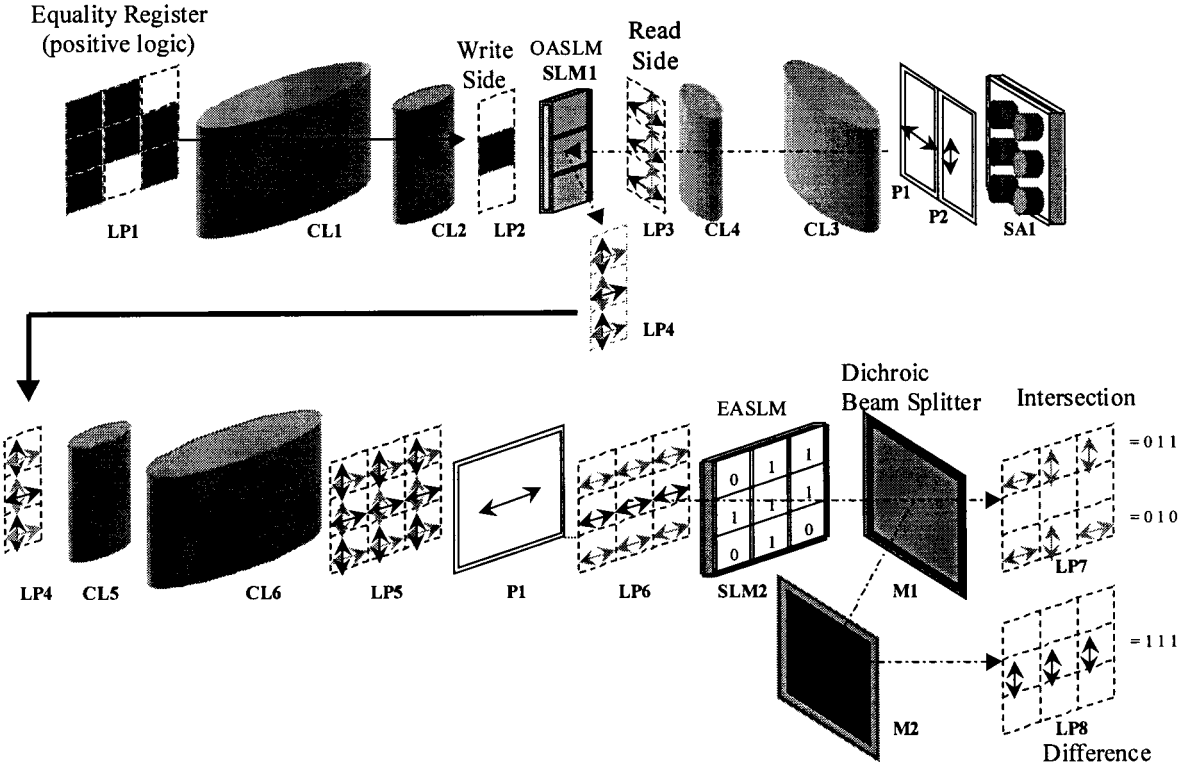

Fig. 6. Intersection–difference operation subunit. This optical unit operates on the positive logic version of the ER (LP1) and extracts the appropriate tuples that result from the intersection (LP7) and the difference (LP8) relational operations. OASLM, optically addressable SLM.

Table 1. Symbols Used in MW-OCAPP Performance Analysis

Symbol	Definition
r	Number of tuples in the RA
s	Number of tuples in the CA
m	Number of tuples (rows) in a SLM
t_{su}	Setup time
t_{out}	Data unload time
t_{load}	Time to load and display a page on a SLM

laterally by CL5 and CL6 to fill the full word length of the aperture (LP5). LP5 passes through a horizontally oriented polarizer that removes all of the vertically oriented information to produce LP6. LP6 reveals that the first and the third rows are encoded on the same wavelength and correspond to the intersection result. The second tuple, which is encoded on the second wavelength, corresponds to the difference result. These three rows in LP6 are encoded with the RA by use of SLM2 and are separated into intersection (LP7) and difference (LP8) registers with a dichroic beam splitter (M1). Logically, LP7 reads as 011 and 010, whereas LP8 reads as 111. Therefore the desired intersection result is present in LP7, whereas LP8 holds the desired difference result.

4. Algorithmic Performance Analysis

There are several metrics that could be used for evaluating the MW-OCAPP performance. One method is to examine the time complexity of each of the database operations. This estimates the number of steps involved in an operation and reports it in big O notation. For example, a serial adder that sums n numbers would have a time complexity of $O(n)$. A second method is to extend the time complexity estimate to a real system, taking into account component response times and problem size. Results are reported in temporal units. Both of these methods are discussed here.

Table 1 shows the various symbols to be used in the analysis to follow. Symbols r, s, and m describe the number of rows in the RA, the CA, and the SLM's, respectively. The type of technology chosen determines the values for the various time-based symbols. The setup time, t_{su}, includes the time for the electronic host unit to load and correctly format the data for input into the optical system. The SLM update time, t_{load}, includes the time it takes a SLM to fully update its display. The data unload time, t_{out}, includes the time it takes for a detector to capture the optically processed registers and store them in the electronic host.

The intersection and difference operations are executed simultaneously and essentially by the same hardware. As a result, their time complexities are identical. If the entire RA and CA can fit on their respective SLM's, then these operations can be completed in a single step, or $O(1)$ time. If this assumption is not valid, then database page swapping is required. Comparison of each of the RA and CA pages against one another requires a time complexity equal to $O(\lceil r/m \rceil \times \lceil s/m \rceil)$. The introduction of the time-based symbols allows for the execution time to be computed with

$$t_{\text{intersection}} = t_{\text{difference}} = t_{su} + \left\{ \left\lceil \frac{r}{m} \right\rceil \times \left\lceil \frac{s}{m} \right\rceil \times (t_{load} + t_{out}) \right\}. \quad (5)$$

The union operation relies on the concatenation of the first of two relations with the results from the difference operator acting on both relations. Since the time complexity closely follows that of $t_{\text{difference}}$, the time complexity for t_{union} is $O(\lceil r/m \rceil \times \lceil s/m \rceil)$. The corresponding execution time can be computed with

$$t_{\text{union}} = 2t_{su} + \left\{ \left\lceil \frac{r}{m} \right\rceil \times \left\lceil \frac{s}{m} \right\rceil + 1 \right\} \times (t_{load} + t_{out}). \quad (6)$$

The projection operation takes a relation and forms a second relation from it that contains a subset of the original attributes. The required duplicate tuple removal process involves performing a self-equivalency search identical in time complexity to the intersection operation, $O(\lceil r/m \rceil \times \lceil s/m \rceil)$. The execution time is, in addition, identical to the intersection operation and can be found with

$$t_{\text{projection}} = t_{su} + \left\{ \left\lceil \frac{r}{m} \right\rceil \times \left\lceil \frac{s}{m} \right\rceil \times (t_{load} + t_{out}) \right\}. \quad (7)$$

The product operation builds a relation consisting of all possible concatenated pairs of tuples from two specified relations. The time complexity of this two-level nested update operation is found to be $O(\lceil r/m \rceil \times \lceil s/m \rceil)$. The execution time of this operation can be found with

$$t_{\text{product}} = t_{su} + \left\{ \left\lceil \frac{r}{m} \right\rceil \times \left\lceil \frac{s}{m} \right\rceil \times (t_{load} + t_{out}) \right\}. \quad (8)$$

Figure 7 illustrates projected execution time estimates per tuple for the above operations. The number of rows in the SLM's (m) is set to 1024, the number of tuples in the CA (s) is 1024, t_{load} and t_{out} are each 10 μs, and t_{su} is 100 μs. The plot reveals a region where the execution time drops as the database size increases. The curve levels out predictably at the point where the database size is equal to the capacity of the SLM's. The graph also reveals a minimum execution time per tuple comparison of approximately 5×10^{-11} s. This corresponds to an execution rate of 2×10^{10} tuple comparisons/s when the specified system parameters are used. Since many relational operations are processed in parallel, and if r different operations are required for the same data set, the effective speed up of the system would increase by a factor of r.

Table 2 compares the execution complexity for each of the relational operations for various architectures. The systems compared include a serial processor, a single-comparand CAM-based processor, and a MW-OCAPP. The MW-OCAPP is able to achieve an $O(1)$

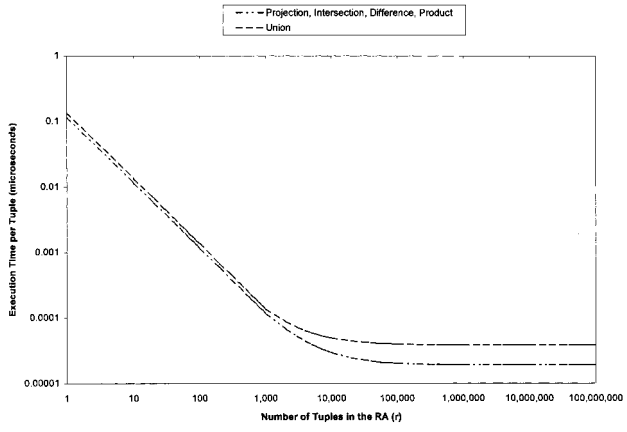

Fig. 7. Relational operation execution time per word comparison as a function of the number of tuples in the RA. Sample projected execution times for the projection, intersection, difference, product, and union operations ($m = 1024$, $s = 1024$, $t_{su} = 100$ μs, $t_{load} = 10$ μs, $t_{out} = 10$ μs).

execution complexity for all operations with the exception of join and update. This translates into a substantial speed up over traditional designs, ranging from an $O(m)$ to an $O(nm)$ factor of improvement.

5. Physical Demonstration of a Preliminary Version of the Multiwavelength Optical Content-Addressable Parallel Processor Equality Operation

An experimental implementation was performed that served as a proof-of-concept design to support further investigation. Figure 2 shows the bit patterns contained in the example RA and CA as well as the expected ER that results from the operation. The CA contains two tuples, 0–0–1–1 and 1–1–0–0. These are compared with four tuples in the RA.

The experimental setup required for implementing the proposed system is displayed in Fig. 8(a). Since an appropriate holographic element (HOE1 in Fig. 3) was unavailable, a right-prism retroreflector and mirror combination was substituted. The drawback of the substitution is that the design is not expandable beyond two optical wavelengths (or two CA words) but is sufficient for demonstration purposes. Second, reflection-mode ferroelectric liquid-crystal (FLC) SLM devices were used instead of the transmissive active elements previously described. This changes the system layout slightly but causes no significant changes in overall operation. The vertically polarized source radiation comes from a 2-W argon-ion laser in its multiline configuration. This beam passes through an afocal beam expander lens system that broadens the beam to a 1.5-cm width. This multiline beam is now filtered to extract the 488.0-nm (blue) and 514.5-nm (green) spectral lines. One wavelength will hold the results from CA2, and the other will hold the results from CA1. These two purified beams are recombined with a cube beam splitter and illuminate opposite halves of a FLC SLM. This SLM is a reflection-mode device that rotates the polarization states by 90° of each of the addressed pixels. It is encoded with the two tuples contained in the comparand array. This passes through a second beam-splitter–prism–mirror assembly that superimposes each of the CA tuples over the remaining tuples. This multiplexed light plane impinges on a second SLM that contains the RA words, and this is imaged onto a screen through a polarizer and grating arrangement.

Figure 8(b) shows the ER that is projected onto a screen located at best focus. Note that there is a dark pixel in the second row of the first column and in the fourth row of the second column. Both these locations correspond correctly to tuple matches predicted in Fig. 2.

6. Experimental System Performance Analysis

The peak modulation frequency of the SLM devices used in the demonstration system is 3 kHz, which is

Table 2. Time Complexity Comparison between MW-OCAPP and Other Systems[a]

Operation	Execution Complexity of a Serial Processor	Execution Complexity of a Single-Comparand CAM	Execution Complexity of MW-OCAPP	Speed Up Over Serial Processor	Speed Up over CAM
Union	$O(n \log n)$	$O(m)$	$O(1)$	$O(n \log n)$	$O(m)$
Intersection	$O(n \log n)$	$O(m)$	$O(1)$ Parallel execution	$O(n \log n)$	$O(m)$
Difference	$O(n \log n)$	$O(m)$	$O(1)$	$O(n \log n)$	$O(m)$
Projection	$O(n \log n)$	$O(m)$	$O(1)$	$O(n \log n)$	$O(m)$
Maximum	$O(n)$	$O(q)$	$O(1)$ Parallel execution	$O(n)$	$O(q)$
Minimum	$O(n)$	$O(q)$	$O(1)$	$O(n)$	$O(q)$
Product	$O(nm)$	$O(m)$	$O(1)$	$O(nm)$	$O(m)$
Divide	$O(n \log n)$	$O(m)$	$O(1)$	$O(n \log n)$	$O(m)$
Limit selection	$O(n)$	$O(1)$	$O(1)$	$O(n)$	$O(1)$
Multiple limit selection	$O(nm)$	$O(m)$	$O(1)$	$O(nm)$	$O(m)$
Join	$O(n \log n)$	$O(m)$	$O(m)$ Parallel execution	$O(n/m \log n)$	$O(1)$
Update	$O(n \log n)$	$O(m)$	$O(m)$	$O(n/m \log n)$	$O(1)$

[a] n = number of tuples in RA, m = number of tuples in CA, q = number of bits in tuple, $n \gg m$.

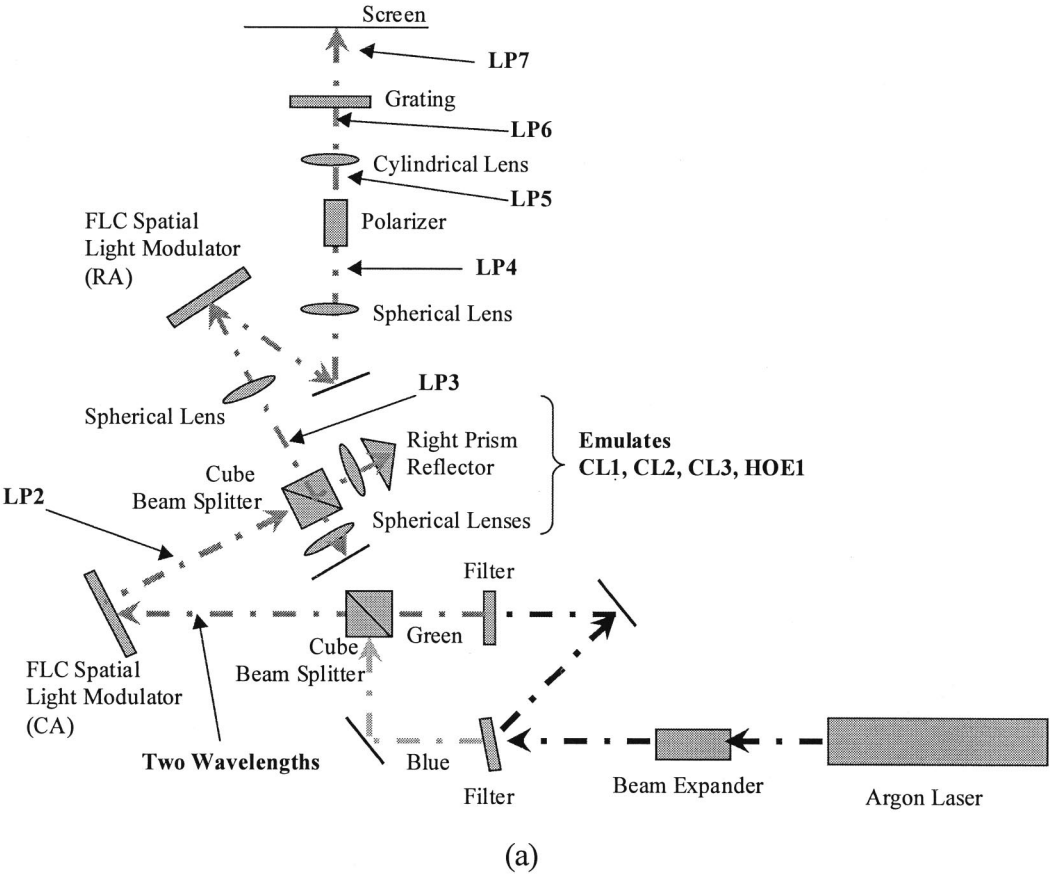

Fig. 8. Demonstration system diagram for (a) the equality operation and (b) the ER result. Component labels correspond to those in Fig. 3. Note that the nonilluminated pixels correctly correspond to the pattern predicted in Fig. 2(c).

limited primarily by the electronic I/O control hardware. This produces a peak bit comparison rate of 96,000 bit comparisons/s for our experimental system. The input and the output bandwidth requirements are 72 and 24 kbit/s, respectively.

The pixel diameter used in the experimental system for SLM1, SLM2, and the detector is 960 μm. If one approximates the pixel apertures as circular, the far-field diffraction pattern has the form of an Airy disk. The diffraction-limited spot size is approximately equal to the diameter of the pattern's central lobe:

$$d = \frac{2.44\lambda f}{D}, \tag{9}$$

where D is the aperture diameter, d is the diffraction-limited spot diameter, f is the distance between the aperture and the diffracted spot, and λ is the wavelength of the light. The diffraction-limited spot size

at SLM2 created by the pixel aperture on SLM1 is 635 μm (λ = 500 nm, f = 0.5 m, D = 960 μm). Optical cross talk can be estimated with[25]

$$P_{\text{cross talk}} = 2\tan^{-1}\left(\frac{d}{2L}\right)\left[P_{\text{enc}}\left(\frac{2L+3d}{4}\right) - P_{\text{enc}}\left(\frac{2L+3d}{4} - d\right)\right], \quad (10)$$

where d is the pixel diameter, L is the center-to-center pixel spacing, and $P_{\text{enc}(\,)}$ is

$$P_{\text{enc}}(r) = 1 - J_0^2\left(\pi\frac{rm}{\lambda f}\right) - J_1^2\left(\pi\frac{rm}{\lambda f}\right), \quad (11)$$

where $J(\,)$ is the Bessel function, λ is the optical wavelength, m is the aperture diameter, and f is the distance from the aperture to the detector pixel. Figure 9(a) illustrates the amount of cross talk received at SLM2 from a neighboring pixel as a function of the pixel gap separation ($L - d$). With zero separation it can be seen that ~2% of the power launched from a neighboring aperture falls on the pixel. If we allow for four nearest-neighbor pixels, the worst-case cross talk added incoherently for a pixel in SLM2 is ~8%. Likewise, the diffraction-limited spot size at the detector is 1.33 mm. With similar analysis as described previously for SLM2, the cross talk at the detector is shown in Fig. 9(b). Unlike SLM2, the 960-μm detector pixel size does not completely enclose the 1.3-mm spot size at the detector. Therefore the cross talk with no pixel gap separation is 8%. The worst-case cross talk with four nearest neighbors is 32%. The experimental optical path length of 1.6 m is fairly long for a free-space optical processor, and this places a lower limit of ~1 mm on the minimum pixel size. Shrinking the optical path lengths would reduce this minimum spot size. However, reducing the longitudinal dimensions of the system results in a larger numerical aperture for the optical components, thereby increasing the amount of aberrations in the optical wave fronts.

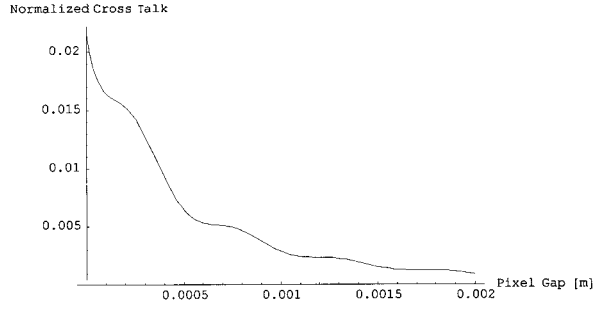

(a)

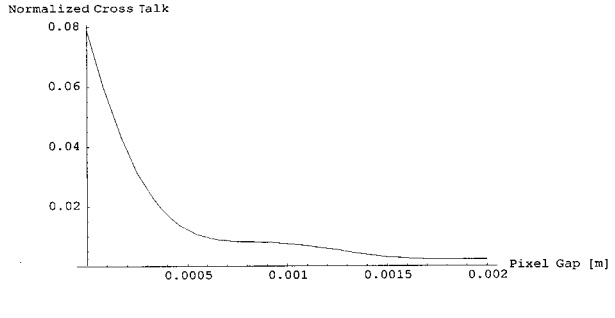

(b)

Fig. 9. Estimated optical diffraction cross-talk power as a function of pixel gap ($L - d$) (a) at SLM2 and (b) at the detector. The cross-talk power is normalized to the total transmitted power from a neighboring aperture (P/P_o).

Table 3. Power Budget

Parameter	Value
Total beam path length (m)	1.6
Loss (dB)	
Beam expander	−1.1
Beam splitter	−6.0
SLM diffraction	−10.4
Prism reflector	−0.4
Lens	−2.6
Polarizer	−0.5
Total loss (dB)	−21
Power fraction (%)	0.8

Power loss in the system is an important consideration when we evaluate choices in the optical sources and the thermal dissipation limits of components. Table 3 shows a summary of the power losses in the system. Supposing that the minimum optical power required for registering a digital 1 is 10 μW, a minimum initial beam power of 1.25 mW is required. In addition to the component losses, significant losses at the source exist. The experimental system uses an argon-ion laser source in its multiline configuration. All optical power that lies outside of the 514- and 488-nm spectral lines is rejected, thus reducing the overall efficiency of the system. Replacing the argon laser with a multiwavelength VCSEL array would eliminate the production of unwanted spectral lines and would improve the system's overall efficiency.

7. Conclusions

In this paper an optical content-addressable processor called the MW-OCAPP is presented. It harnesses a unique method of wavelength multiplexing and polarization multiplexing to achieve a high level of parallelism. Optical implementation is made possible by exploitation of the noninteractive behavior of coincident light planes of differing wavelengths. This architecture offers database systems constant-time parallel equality and magnitude comparison of multiple comparands with multiple tuples in a relational array.

In general it was found that most relational operations were completed with an $O(1)$ time complexity and as much as an $O(nm)$ speed improvement over previous architectures, where n is the database size and m is the number of comparands to be operated on. This performance suggests a substantial performance im-

provement over previous designs for database operations such as sorting, which typically repetitively use the magnitude comparison operation, among others. Interfacing the MW-OCAPP with existing popular environments such as the structured query language should be relatively easy, owing to the structured query language's modularity, simplicity, and inherent portability. With the increasing importance of information management and data mining, flexible and highly parallel symbolic processors such as the MW-OCAPP will become increasingly desirable.

The equality operation demonstration system exhibited a peak processing rate of 96,000 bit comparisons/s, which was limited by the SLM's page update rate of 3 kHz. Increasing the number of wavelengths utilized or optimizing the SLM electronic I/O connection to accommodate high-speed page updates in the megahertz range would improve the practical limits of this system. The worst-case cross talk at the detector plane with no pixel separation was analytically determined to be $\sim$32%. We can improve this by either shrinking the optical path lengths, increasing the detector separation, or increasing the detector pixel diameters.

There are several hurdles that have not yet been addressed in this architecture. The first issue is one of database size. For the proposed examples it was assumed that the databases were small such that they could be encoded in their entirety onto a SLM. In reality, databases will far exceed the spatial resolution of the light modulators. Intelligent iteration algorithms through the CAM array need to be developed to bridge this discrepancy. The second issues are those surrounding the demultiplexing of polychromatic light. Current technologies can handle hundreds of channels (separate wavelengths), but the task of combining and separating each of these channels with minimal loss in signal power can be daunting. The choice in tunable VCSEL's and grating structures, and efforts to keep optical path lengths short, will all contribute to the usable bandwidth of the system.

This research was supported by National Science Foundation grant MIP-9505872.

References

1. S. Y. Su, *Database Computers, Principles, Architectures, and Techniques* (McGraw-Hill, New York, 1988), Chap. 1.
2. L. Chisvin and R. J. Duckworth, "Content-addressable and associative memory: alternatives to the ubiquitous RAM," IEEE Comput. **22**(7), 51–64 (1989).
3. P. B. Berra, A. Ghaffor, P. A. Mitkas, S. J. Marchinkowski, and M. Guizani, "The impact of optics on data and knowledge base systems," IEEE Trans. Knowledge Data Eng. **1,** 111–132 (1989).
4. A. Louri, "Optical content-addressable parallel processor: architecture, algorithms, and design concepts," Appl. Opt. **31,** 3241–3258 (1992).
5. R. Elmasri and S. B. Navathe, *Fundamentals of Database Systems*, 2nd ed. (Addison-Wesley, New York, 1994).
6. K. Giboney, L. Aronson, and B. Lemoff, "The ideal light source for datanets," IEEE Spectrum **35**(2), 43–53 (1998).
7. P. B. Berra, K. Brenner, W. T. Cathey, H. J. Caufield, S. H. Lee, and H. Szu, "Optical database/knowledgebase machines," Appl. Opt. **29,** 195–205 (1990).
8. D. Psaltis and G. W. Burr, "Holographic data storage," IEEE Comput. **21**(2), 52–59 (1998).
9. A. V. Krishnamoorthy, P. J. Marchand, G. Yayla, and S. C. Esener, "Photonic content-addressable memory system that uses a parallel-readout optical disk," Appl. Opt. **34,** 7621–7638 (1995).
10. P. A. Mitkas, L. J. Irakliotis, F. R. Beyette, Jr., S. A. Feld, and C. W. Wilmsen, "Optoelectronic data filter for selection and projection," Appl. Opt. **33,** 1345–1353 (1994).
11. L. J. Irakliotis, S. A. Feld, F. R. Beyette, Jr., P. A. Mitkas, and C. W. Wilmsen, "Optoelectronic parallel processing with surface-emitting lasers and free-space interconnects," J. Light. Tech. **13,** 1074–1084 (1995).
12. P. A. Mitkas, S. A. Feld, F. R. Beyette, Jr., and C. W. Wilmsen, "Optical digital comparison unit for equal-to, less-than and greater-than determination," Appl. Opt. **33,** 806–814 (1994).
13. A. D. McAulay, *Optical Computer Architectures* (Wiley-Interscience, New York, 1991).
14. R. D. Snyder, S. A. Feld, P. J. Stanko, E. M. Hayes, G. Y. Robinson, C. W. Wilmsen, K. M. Geib, and K. D. Choquette, "Database filter: optoelectronic design and implementation," Appl. Opt. **36,** 4881–4889 (1997).
15. A. Louri and J. A. Hatch, "Optical content-addressable parallel processor for high-speed database processing," Appl. Opt. **33,** 8153–8164 (1994).
16. A. Louri and J. A. Hatch, "Optical content-addressable parallel processor for high-speed database processing: theoretical concepts and experimental results," IEEE Comput. Special Issue on Associative Processors **27**(11), 65–72 (1994).
17. A. Louri and J. A. Hatch, "Optical implementation of a single-iteration thresholding algorithm with applications to parallel database/knowledge-base processing," Opt. Lett. **18,** 992–994 (1993).
18. K. W. Wong, L. M. Cheng, and M. C. Poon, "Design of digital-optical processor by using both intensity and polarization-encoding schemes," Appl. Opt. **31,** 3225–3232 (1992).
19. A. W. Lohmann, "Polarization and optical logic," Appl. Opt. **25,** 1594–1597 (1990).
20. P. Y. Choo, A. Detofsky, and A. Louri, "A multiwavelength optical content-addressable parallel processor (MW-OCAPP) for high-speed parallel relational database processing: architectural concepts and preliminary experimental system," in *Parallel and Distributed Processing*, J. Relim, ed., Lecture Notes in Computer Science, Vol. 1586 (Springer-Verlag, Heidelberg, 1999), pp. 873–886.
21. P. Y. Choo, A. Detofsky, and A. Louri, "A multiwavelength optical content-addressable parallel processor (MW-OCAPP) for high-speed parallel relational database processing: architectural concepts," in Digest of Topical Meeting on Optics in Computing (Optical Society of America, Washington, D.C., 1999), pp. 66–69.
22. P. Y. Choo, A. Detofsky, and A. Louri, "An optical architecture based on multiwavelength and polarization for parallel and high-speed relational database processing," in *Optics in Computing '98*, P. Chavel, D. A. B. Miller, and H. Thienpont, eds., Proc. SPIE **3490,** 139–143 (1998).
23. A. Detofsky, P. Y. Choo, and A. Louri, "Optical implementation of a constant-time multicomparand bit-parallel magnitude-comparison algorithm using wavelength- and polarization-division multiplexing with application to parallel database processing," Opt. Lett. **23,** 1372–1374 (1998).
24. R. Watanabe, K. Nosu, and Y. Fujii, "Optical grating multiplexer in the 1.1–1.5-m wavelength region," Elec. Lett. **16,** 108–109 (1980).
25. A. S. Miller and A. A. Sawchuk, "Capabilities of simple lenses in a free-space perfect shuffle," in *Optical Enhancements to Computing Technology*, J. A. Neff, ed., Proc. SPIE **1563,** 81–92 (1991).