Photorefractive adaptive resonance neural network

Donald C. Wunsch II, David J. Morris, Rick L. McGann, and Thomas P. Caudell

We describe a novel adaptive resonance theory (ART) device that is fully optical in the input–output processing path. This device is based on holographic information processing in a photorefractive crystal. This sets up an associative pattern retrieval in a resonating loop that uses angle-multiplexed reference beams for pattern classification. A reset mechanism is used to reject any given beam, permitting an ART search strategy. The design is similar to an existing nonlearning optical associative memory, but ours permits learning and makes use of information that the other device discards. It is a suitable response to the challenges of connectivity, learning, and reset presented by ART architectures. Furthermore, the design includes an efficient mechanism for area normalization of templates. It also permits the user to capitalize on the ability of ART networks to process large patterns. This new device is expected to offer higher information storage density than alternative ART implementations.

1. Introduction

Neural networks based on adaptive resonance theory[1–7] (ART) offer a number of implementation challenges, the greatest of these being the resent mechanism, the massive interconnectivity requirements, and the requirement that these interconnections be adaptive. Previous research on holographic associative memory[8] has inspired us to create a related device capable of learning and of normalizing memories. We have been further encouraged by successful implementation of massive information storage in a photorefractive crystal.[9] Our research represents a synthesis of these results.

In this section we present the motivation for our research by discussing ART and its associated challenges for hardware implementation. In Section 2 we describe our design, and in Section 3 we discuss the system layout in more detail. Section 4 describes our (non-novel) laboratory progress toward realization of a hardware implementation of an optical ART network by means of incremental component-level hardware evaluations. Section 5 discusses future research, and Section 6 summarizes our conclusions.

When this research was being performed, the authors were with the Boeing Company, P.O. Box 24346, Seattle, Washington 98124-0346. D.C. Wunsch is now with the Department of Electrical Engineering, Texas Tech University, Lubbock, Texas 79409.

Received 9 June 1992.
0003-6935/93/081399-09$05.00/0.

ART is attractive because of its stable unsupervised learning properties and its ability to process large input-pattern fields. Recent successful applications of ART[3] used input fields in excess of 10^7 nodes. It appears that ART's scaling properties are limited only by hardware and software implementations. Thus we have significant motivation for the present research. We present only a rough outline of the operation of ART, which is covered extensively elsewhere.[1–7] We merely want to point out some properties of ART networks that are essential to their operation, yet present hardware challenges: reset and massive connectivity. An example illustrates these points.

One type of ART network is depicted in several layers in Fig. 1. The layers are as follow: R, the recognition layer; C, the comparison layer; I, the input layer; V, the vigilance layer; Re, the reset layer. This grouping of layers is similar to that in Ref. 6 and, while it does not follow Carpenter's and Grossberg's original description[1] exactly, it is functionally equivalent. Looking from left to right, we see the ART unit in action. First, the input is registered at the comparison layer and fed up to the recognition layer [Fig. 1(a)]. In the second frame the recognition layer's winner-takes-all property finds the node corresponding to the initial best guess [Fig. 1(b)]. This guess is tested by replaying the winning node's previously learned template (weight values) onto the comparison layer. The input layer, which still contains the original input, and the comparison layer, containing the winning template, compare these two patterns by

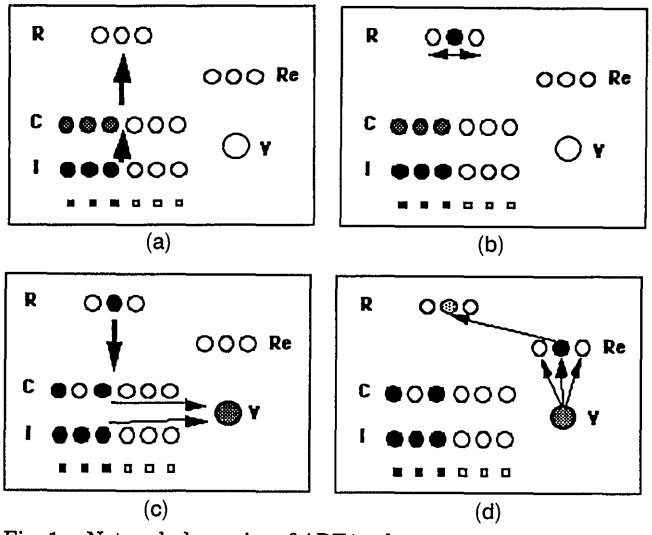

Fig. 1. Network dynamics of ART1: bottom-up pattern matching [(a), (b)] is balanced by top-down feedback expectancy [(c), (d)].

sending competing signals to the vigilance node [Fig. 1(c)]. The final frame shows an example of what happens when the match is not good enough. The vigilance node is here able to activate the reset layer. The reset layer suppresses only nodes at the output that have been recently active, and has no effect on the rest [Fig. 1(d)]. In this example, only the prior winner is affected. With this node removed, the network reclassifies the pattern and continues to do so until it finds a good match. When a good match is found, the winning node's activation is permitted to continue. The simultaneous activation of the input nodes and a winning recognition node creates a resonance. This resonance is frustrated by the reset mechanism until a good match is found. The key is that the pattern classification takes place in a feedback loop and that learning does not set in until resonance is permitted to persist. Until then, the reset mechanism permits a search for a better pattern match, removes all classifications considered previously, and suspends learning until the best answer is found.

The communication between the comparison layer and the recognition layer requires a massive number of interconnects; furthermore, these connections must be capable of real-time adaptation. (Many neural hardware implementations involve hard-wired connection weights; clearly not an option for ART networks.) An ART unit of the type shown in Fig. 1 having N inputs and M outputs requires the following types of connections: $2\,MN$ adaptive, $M^2 + 2M$ inhibitory, $5M$ excitatory, N inhibitory, $3N$ excitatory, and, of course, N-input and M-output connections. This is a compelling argument for the importance of optical implementations.

The advantage of optics lies in the fact that implementing interconnects electronically is difficult because of electromagnetic interference and the necessity of wires to carry the signals. A basic assumption of the artificial neural systems paradigm is that the importance of interconnects is dominant for a certain class of problems. Psaltis and Abu-Mostafa[10] assessed the relationship of computing power to interconnects, independent of the power of the individual processing elements. They suggest that an interconnection-dominated problem "... has the property that local decisions cannot be made until essential information has been communicated from basically the entire input data. Thus useful computation can progress only when all the input information has been considered by the individual elements. For a parallel processor, this implies that all partial results need to be globally communicated. It is this notion that forms the basis for our conviction that communication capability becomes the dominant factor. . . ."[10] Holographic implementations offer superior communication capability because of the degrees of freedom inherent in the medium; it is this capability that we exploit in this implementation.

2. Optical Implementation

The design of the optical ART unit is shown in Fig. 2. It is a modification of the resonating loop reported by Soffer *et al.*,[8] which is shown in Fig. 3. A key difference between the figures is the nonlinear storage crystal of Fig. 2 in place of the hologram of Fig. 3. In our optical ART implementation the crystal is photorefractive barium titanate ($BaTiO_3$), which is capable of recording multiple holograms in real time.[9,11] It acts as part of a resonant cavity designed to converge on the correct images so that it behaves as an ART unit. Contrast this use of a crystal with Fig. 3, in which a fixed hologram is used. The latter design permits the device to be capable of associative pattern retrieval, but not of learning. We also introduce two components between the storage crystal and the right-hand phase-conjugate mirror (PCM1) in

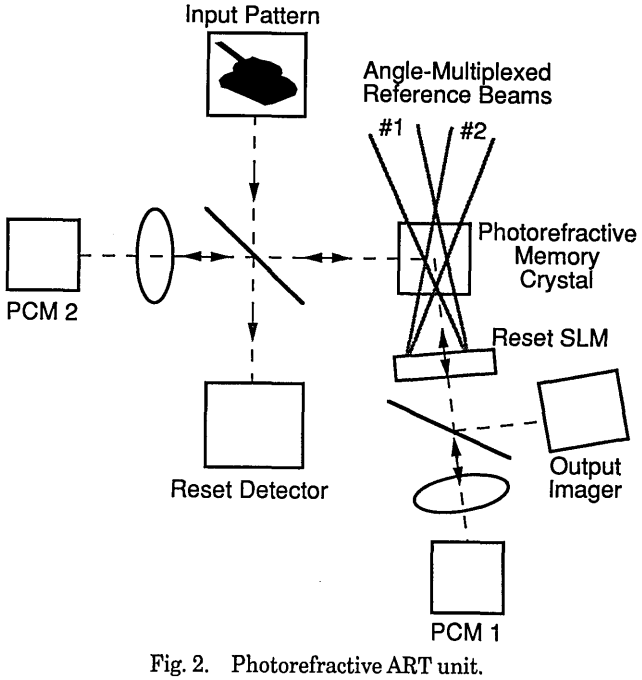

Fig. 2. Photorefractive ART unit.

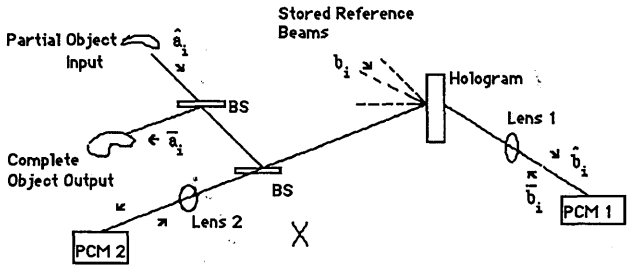

Fig. 3. Holographic optical resonator: BS's, beam splitters.

Fig. 2; these are a reset spatial light modulator (the reset SLM) and a beam splitter to permit imaging of the angle-multiplexed reference beams to identify the active mode.

A brief review of the operation of the holographic resonator shown in Fig. 3 helps to explain the optical ART unit shown in Fig. 2. The holographic resonator is primed with a partial input pattern $\hat{a}_i$, which is shown in the upper left of Fig. 3. This is reflected off the second beam splitter toward the fixed, previously recorded hologram, which excites several reference beams $\hat{b}_i$. These are retroreflected by a thresholding PCM1 back toward the hologram, setting up a resonant loop between PCM1 and PCM2. The loop is biased by the presence of the hologram and by the injected signal $\hat{a}_i$, and it suppresses all stored patterns (and their reference beams) except for the one most closely matching $\hat{a}_i$. This causes a readout of the stored image $\bar{a}_i$ closest to the input $\hat{a}_i$. The device can be considered as a pattern classifier by considering the output reference beam b_i to be an angularly multiplexed classification code. Also, light that contains information from both $\hat{a}_i$ and $\bar{a}_i$ is present at the point marked X in Fig. 3, but this information is not used. Reference 8 contains a more complete description of the device illustrated in Fig. 3.

The optical ART unit of Fig. 2 was inspired by the Soffer *et al.* device, but it uses a $BaTiO_3$ crystal instead of a hologram to permit learning. This simple change is key in designing the device to behave as an ART unit, but providing a reset mechanism to implement ART correctly is also necessary. This is provided by placing a detector in a position to record the overlap of the input pattern and the recorded template (precisely the information that is discarded at location X of Fig. 3) and by controlling and modulating the overall phase in the optically resonant loop. The reset detector of Fig. 2 is an integrating photodetector. During initial setup, with a single template stored in the memory crystal and with the same (matching) pattern input to the resonator, the phase of one of the two pump beams in PCM2 of Fig. 2 is adjusted[12]; this is done while the total power at the reset detector is monitored to establish the phase at which maximum constructive interference between the input pattern and the (matching) stored template occurs. (This implementation of phase adjustment requires that PCM2 be operated in the four-wave-mixing mode, while the nonlinear mode competition ncessitates that PCM1 in Fig. 2 be operated in the thresholding photorefractive self-pumped mode.)

Subsequently, during pattern-search mode (when one or more templates are stored in the memory crystal and a new pattern is presented) the adjustable PCM2 pump-beam phase is left at the pre-established setting, and the system is permitted to resonate until a single mode dominates. During this pattern-search mode the power used in the readout beams is insufficient (for the time scales involved) to rewrite the memory crystal. When a single template is dominant in the resonator (as indicated by a constant signal at the reset detector), the reset decision is made. To make this decision, one sweeps the phase of the pump to PCM2 through several cycles while looking for the presence of power modulation in each pixel of the reset detector's field of view. Each pixel for which an overlap between input and template exists exhibits this modulation. The quality of the pattern match is determined from the number of overlap pixels detected by the reset detector.

Thus far we have explained the operation of the reset detector to measure the overlap of the input pattern and template by the use of constructive or destructive interference. This obtains the inner product between the input and the template. As discussed in Refs. 1 and 2, however, we must normalized these measurements by input-pattern size. This can be accomplished in two ways. The simpler but slower way is to shut off the pump beams to the four-wave-mixing PCM2 (see Fig. 4). This is done only when the input pattern is first presented. The reset detector then measures the total input pattern intensity and saves this number electronically. The pump beams are then switched back on, and the reset detector reads the constructive interference pattern representing the overlap between the input and the template. The reset decision is made from these two pieces of information. The faster way to do this is to add an extra beam splitter and to bleed off a copy of the input pattern before injecting it into the resonator. This copy is then measured by an extra detector and is used for the reset decision.

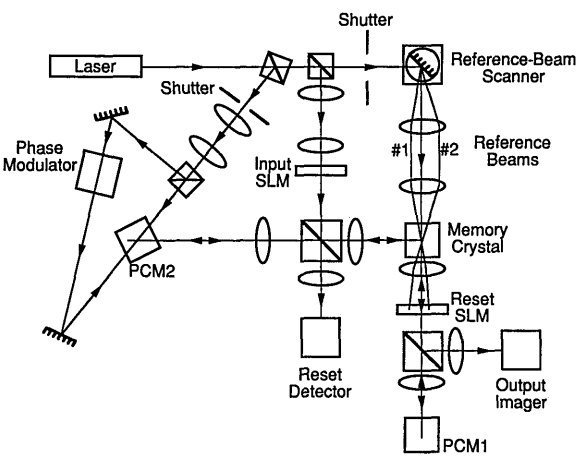

Fig. 4. Lab layout of the optical ART unit.

If the reset detector's match measurement is too small, reset is indicated, and the steps described below are performed. If the reset detector's match measurement is above a user-determined threshold, reset if not triggered, and the system is permitted to resonate in its preferred mode, causing learning of the new pattern. This learning can be accelerated by increasing the power in the readout beam for PCM2.

In the event that a reset condition is indicated, the following actions take place. The optical system viewing the reset SLM (off the beam splitter on the right-hand side of Fig. 2) is used to identify the location of maximum intensity at the reset SLM plane, which corresponds to the location of the reference beam associated with the dominant system mode (say, the kth reference beam). The (formerly inactive) rest SLM is then activated to produce attenuation at the known location at which the kth reference beam focuses through the reset SLM. This significantly reduces the gain of the kth system mode and permits continued pattern search, in which the kth mode is not permitted to compete. This process continues, with poorly matching modes excluded sequentially from competition by the reset SLM, until a dominant mode that sufficiently matches is found. If no such mode is found, the input pattern is stored as a totally new template by using a new reference beam aligned along a direction previously unused.

The reset detector also can play a useful role in normalizing patterns by template size, which is a requirement of the ART neural network.[1,2] This is an important problem, not just for this specific neural-network implementation, but for holographic pattern recognizers in general. Some type of area-based normalization of the templates needs to be performed in order for the pattern classification to be meaningful. However, the reset detector is in effect measuring the size of the pattern that becomes the new recorded template. Therefore, when reset is not indicated and a new template is recorded, the constructive interference measurement is used to determine the desired grating strength associated with that recorded template. This is done by providing the detector with enough electronic memory to store all the template sizes up to the capacity of the system. This uses a small amount of memory, even for a system of large capacity, because only one number must be stored for each template. This information is used to modify the recording schedule for all templates. Small templates are recorded longer; big ones are recorded a shorter time. Alternatively, recording power can be increased by different amounts, depending on template size. This makes the recording schedule slightly more complicated, but the complexity is compensated by improved performance of the system. A normalizing scheme is necessary, and this one is useful because it capitalizes on system components that already exist.

In this implementation the output of the neural network is the classification provided by the reference-beam identity, which can be read off by use of the output imaging optics viewing the reset SLM off the beam splitter, as shown on the right-hand side of Fig. 2. It is also possible to read out the stored template information in a manner similar to that for the complete object output shown on the left in Fig. 3. In other words, the device can be used as either a heteroassociative or an autoassociative ART-based memory. Furthermore, it may be possible to replace the angle-multiplexed reference beams with a third SLM. This could be used for associating input patterns with known output patterns for supervised ART-based learning.[5] The image on the reset SLM would be the inverse of the image on this third training SLM. We do not address the idea further here, we simply mention the possibility.

The optical layout schematic of the device is shown in Fig. 4. For experimental verification of the design, human observation of the reset signal and activation of the reset SLM is acceptable. This role culd ultimately be automated. The angle multiplexing of reference beams is carried out in a manner identical to that described by Mok *et al.*[9] The details of the nonlinear competition between accessible modes in this type of device is described by Soffer *et al.*[8]

3. Experimental Setup

Three photorefractive $BaTiO_3$ crystals are called for in the optical ART demonstrator implementation illustrated in Fig. 4. One is the memory crystal, used to record templates against which input patterns are classified. The other two are used for phase-conjugate feedback of laser light within the ART resonator.

The memory crystal has two modes of operation, template storage and readout. The template storage mode calls for the use of relatively intense waves in the reference beams and the object input wave (through the input SLM) in order to quickly create or modify phase gratings within the photorefractive crystal. The readout mode involves much lower intensity waves in order to permit a search process before overwriting of the stored holograms occurs. In other words, the competitive pattern-matching mechanism of ART can be realized at low power levels to minimize erasure of templates recorded previously. Once the system has made a classification, a higher power level is used to write the new updated template quickly.[13] Over time, both the high-power writing of new templates and the integrated effects of low-power readout degrades templates learned previously. This can be circumvented by rereferencing the storage crystal with the template image(s). This could be compared with the refresh cycles necessary in a dynamic random-access memory. Such rereferencing could take place each time a new template is written after a classification search has terminated. However, such frequent rereferencing may not be necessary. It is possible to stagger power levels,[9,14,15] anticipating template erasure by writing earlier templates at higher power and gradually decreasing the power.

For the feedback PCM's, we employ four-wave mixing in PCM2 with controllable pump-wave power

in a way such that adequate effective gain is established to overcome transmission and holographic efficiency losses within the resonator.[16] PCM1 is operated in the self-pumped mode so that the necessary thresholding mechanism is provided.

The reference beam used to record a particular template must focus at the plane of the reset SLM after propagation through the memory crystal. Each of the reference beams, which are angle multiplexed at the memory crystal, focuses at a different location on the reset SLM plane.

4. Adaptive-Resonance-Theory Memory-Component Evaluations

Here we consider implementation details necessary for optimal configuration of the $BaTiO_3$ crystal that serves as the holographic memory element. Some of these comments are of general relevance for phase-conjugate holography in neural-net implementations, particularly those involving $BaTiO_3$. They are, however, the result of our own incremental progress towards implementation, and are not presented as original contributions. (For other information in this area, we recommend Refs. 16–19.) The central purpose of the experiments described below is to determine the proper geometry and the feasibility of multiple-learned-template storage in the device.

A. Geometry and Grating Formation

To demonstrate the complete optical ART system, we first need to characterize the performance and to determine the most effective beam and crystal geometries of the $BaTiO_3$ memory crystal. The geometry used for these tests is defined in Fig. 5; $\mathbf{k}$ is the grating spatial wave vector, $\mathbf{c}$ is the crystal c axis, α is measured inside the crystal, and γ and β are measured outside the crystal. All input polarizations are extraordinary (i.e., in the plane of the page). Our crystal is approximately 3 mm × 4 mm × 5 mm with the crystal c axis normal to the 3 mm × 4 mm faces. Because $BaTiO_3$ has an intrinsic phase shift of 90° between the applied interference pattern and the resulting photorefractive phase grating, the process of recording holograms leads to two-beam coupling[16] in the crystal.

In our application, beams 1 and 2 are the reference and signal beams incident upon the memory crystal. Grating strength and holding time both increase with decreasing γ because the diffusion of charges is frustrated by the larger scale of the resulting grating. We selected $\gamma \approx 14°$ as a minimum γ consistent with practical mechanical and optical constraints. The theory for the photorefractive grating formation[16] in $BaTiO_3$ predicts that, for a fixed γ, the greatest grating strength should be obtained with $\alpha \approx 40°$. In our geometry the crystal size and shape and the orientation of the c axis limit the external incidence angle to roughly the interval $35° < \beta < 145°$, so that refraction limits the grating angle α to be considerably less than 40°. In practice, we find that the best performance is obtained at $\beta \approx 50°$. A different specification of crystal dimensions to that permits a better choice of β is clearly recommended for future research in this area.

Grating-formation time and saturation strength were characterized by observing the onset and steady-state strengths for the two-beam coupling effect with input beams of equal irradiance, $\sim 25\ mW/cm^2$ at the crystal. Illumination of the crystal interior by a 150-W optical-fiber white-light illuminator was used to erase stored gratings and to hold off formation of gratings as a way to gain a starting time for the determination of grating formation rates. Figure 6 shows the formation and saturation of gratings under the selected geometry (i.e., with $\beta \approx 50°$ and $\gamma \approx 14°$). The data plotted versus time in the figure are two photodetector output traces representing the power of beams 1 and 2, as well as a third trace that is the sum of the power of beams 1 and 2. Prior to event 1 as shown in Fig. 6, the erasure illuminator was on and both detectors were blocked in order to set zero-signal base lines. At event 1 the detectors were unblocked with the erasure illuminator left on in order to prevent formation of the gratings; the power of the two input beams is matched. At the time of event 2 the erasure illuminator is turned off, and formation of the gratings begins. The strength of the grating at any time is indicated by the power imbalance between the initially balanced output beams. At steady state the ratio of the power in output-beam 2 to the power in output-beam 1 has changed from unity to ~8, which is consistent with a single-beam scattering efficiency of 37.5%. The detector-sum trace remains near 100% throughout the

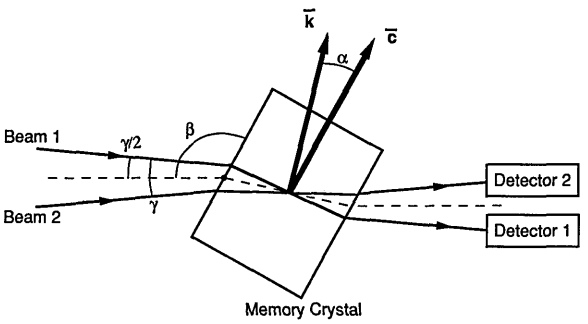

Fig. 5. Geometry for the two-beam-coupling grating-formation tests.

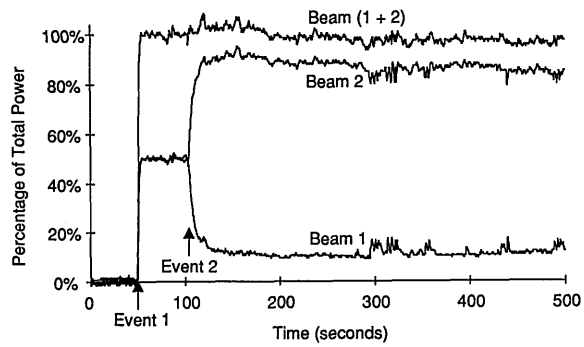

Fig. 6. Two-beam coupling ($\beta = 50°$).

grating formation, indicating that power is being transferred between the two beams with negligible scattering into other directions. The time constant for the grating formation at this power level is ~7 s.

In the ART implementation a number of distinct gratings is formed, each with a signal-beam input from the same (fixed) direction and a reference-beam input along a distinctly different direction. In order to ensure that multiple stored holograms have similar formation time constants and decay rates, all the gratings need to be formed with nearly the same geometry and spatial frequency magnitude $|\mathbf{k}|$.

B. Competition with Self-Phase Conjugation

Since the two-beam coupling theory predicts that, for any angle x, the coupling strength for $\beta = 90° - x$ should be the same as for $\beta = 90° + x$, we attempted to demonstrate the equivalent beam interaction at $\beta \approx 130°$. The result of this experiment is shown in Fig. 7. The two-beam coupling proceeds as in Fig. 6 for several seconds immediately following event 2 (the onset of grating formation). At event 3, however, the interaction abruptly changes from that of Fig. 6, with both beams falling to ~20% of the power levels observed between events 1 and 2. Concurrent with event 3, the total power in both output beams (the upper trace) deviates sharply from 100%, also settling at ~20% in the steady state. This deviation of the upper trace from 100% is indicative that some other process is competing with the two-beam coupling effect, stealing energy from the two output beams. Indeed, by observing the two input beams in the retropropagating direction, we confirmed that the input beams were undergoing phase conjugation. This competing nonlinear effect needs to be minimized in our optical ART memory implementation by selecting reference- and signal-beam geometries that frustrate the occurrence of the undesirable self-phase conjugation.

Figure 8 plots the strength of the self-phase-conjugation response to a single input beam of extraordinary polarization as a function of the external beam–incidence angle β', as illustrated in Fig. 9. For each input β' the crystal was translated laterally

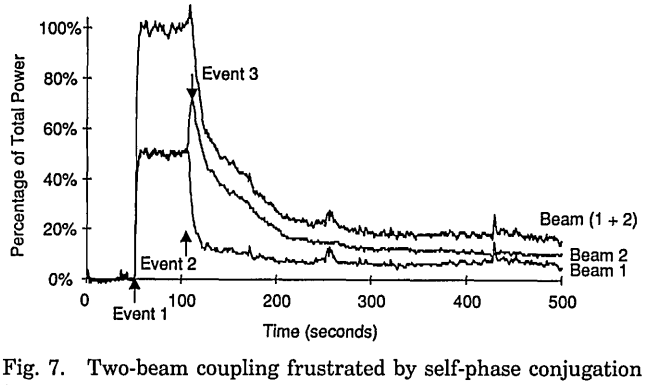

Fig. 7. Two-beam coupling frustrated by self-phase conjugation ($\beta = 130°$).

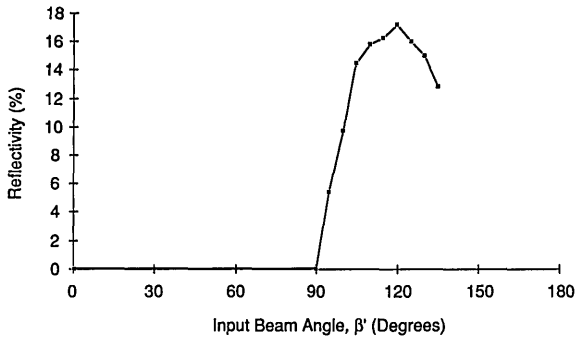

Fig. 8. Input geometries permitting self-phase conjugation.

to find the maximum self-phase-conjugation strength. The plot of Fig. 8 shows that, in order to frustrate the onset of self-phase conjugation, we need $\beta' < 90°$. (This is useful not only because it permits analysis of undesired competition with two-beam coupling, but because it also provides the range of input geometries that permits phase conjugation as desired in PCM1 of Fig. 4.)

Considering the need to avoid self-phase conjugation of the principal beams in the ART memory crystal, we selected the geometry and beam assignments shown in Fig. 10. For the optical ART implementation of the memory crystal there are two operational modes: storage and search. In the storage mode the memory crystal is illuminated by a reference beam and a signal beam simultaneously. In the selected geometry ($\beta \approx 50°$) neither the signal nor the reference beam can self-conjugate during the storage mode. In the search mode the crystal is illuminated by a signal beam in the absence of any reference beam. The hologram forms a virtual image of any reference beam associated with a near-matching stored template. The reference beam(s) generated in the memory crystal propagate out of the crystal and through the reset SLM. They are reflected by PCM1 and propagate back into the crystal, where they scatter from the stored hologram(s) to produce backward propagating beam(s) along the signal beam path. For the memory-crystal geometry we selected and with respect to the self-phase-conjugation geometry illustrated in Fig. 9, the input signal beam is along a path with $\beta' \approx 43°$, the input reference beam is along a path with $\beta' \approx 57°$, the retropropagating signal path has angle $\beta' \approx 137°$, and

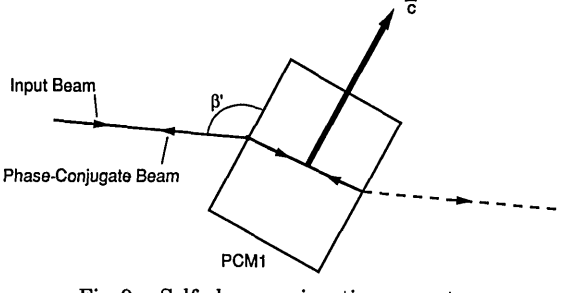

Fig. 9. Self-phase-conjugation geometry.

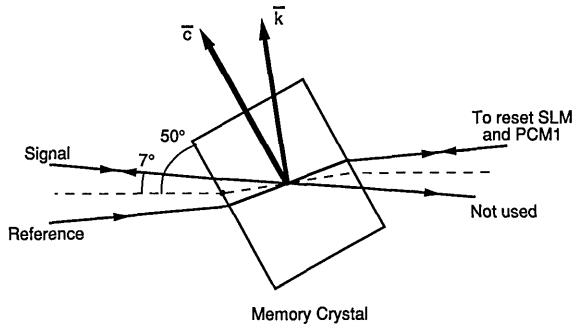

Fig. 10. Selected memory-crystal geometry.

the retropropagating reference path has angle $\beta' \approx 123°$. In this geometry the retropropagating versions of the signal and reference beams both have β' angles that may permit self-phase conjugation. In the event that self-phase conjugation of the retropropagating beams interferes with the optical ART implementation, we attempt to reduce the effect by a variety of means, including translation of the crystal to avoid the total internal reflection at the crystal–air interfaces that can provide optical feedback required for self-pumped phase conjugation.[18]

C. Grating Decay Rate

Experiments were performed to determine the rate at which gratings, formed in the geometry shown in Fig. 5, decay during the process of readout. Figure 11 shows the grating scattering efficiency (i.e., the grating strength) as a function of time under four different illumination conditions. The four curves correspond to the decay in the presence of a single readout beam of four different powers: 1.7 mW and 600, 170, and 20 μW; the highest power is equal to the power in each of the two beams during the grating formation process shown in Fig. 6. In the 20-μW case the decay rate is converging on the intrinsic (zero-illumination) decay rate of the grating. The 1.7-mW curve indicates an exponential decay time constant of ~175 s. For the grating decay caused by the formation of other gratings (requiring input of two 1.7-mW beams) the decay time constant would be ~88 s.

The grating scattering efficiency at zero time for each curve in Fig. 11 is ~37.5%. This represents the maximum (saturation) grating strength achievable with our selected geometry, which plays a role in determining the counterbalancing gain required of the phase-conjugating feedback mirror to permit resonance in the optical ART implementation.

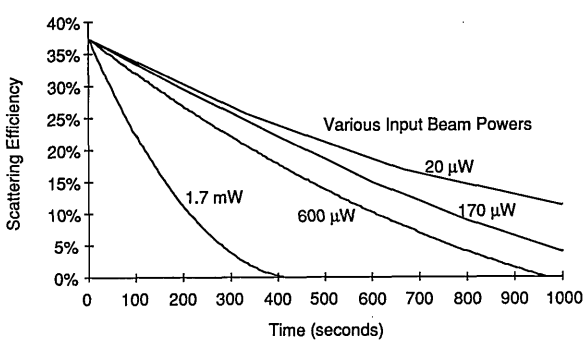

Fig. 11. Grating decay rates.

D. Storage of Multiple Matched Gratings

Given the grating-formation and grating-decay time constants, Mok *et al.*[9] provide the means for determining exposure-time sequences appropriate for the formation of multiple holograms with equal strength at the end of the storage process. We see no reason why $LiNbO_3$, their choice of crystal, could not also be used in this device. This would provide the large template storage capacity they demonstrated, but at the price of requiring higher power levels. Using their technique, we analyzed the storage properties of $BaTiO_3$. Figure 12 shows the maximum matched-grating scattering efficiency that can be achieved as a function of the total number of gratings stored, when this method, with our formation and decay time constants of 7 and 88 s, respectively, is used. For Fig. 12 we assume a 37.5% saturation scattering efficiency (consistent with both the two-beam coupling strength at saturation and the grating scattering efficiency measurements), and we assume that negligible time passes between formation of the sequential gratings.

The results in Fig. 12 can be used to determine the minimum gain required of the phase-conjugating feedback mirror in the optical ART implementation. For example, the figure shows that, if seven gratings are to be stored, the maximum grating scattering efficiency is ~20%. The two-pass efficiency is ~4% for the memory crystal, which, combined with the demonstrated ~20% reflectivity for PCM1, leads to a round-trip efficiency of ~0.8%, which must be counterbalanced by a gain in PCM2 of ~125.

We are encouraged by the high capacity theoretically possible with volume holography and by the actual results achieved by others in this area of research.[9] In the short run we expect that the attainable capacity of the system will be limited by the number of holograms stored, which should be in the hundreds, if not thousands, of memories. In the long run, as the ability to store multiple holograms in

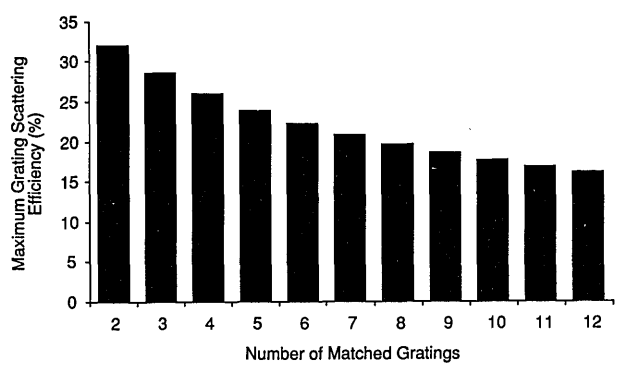

Fig. 12. Maximum scattering efficiency versus the number of matched gratings.

photorefractive crystals improves, we expect to be limited by the reset SLM or by the reference beam steering mechanism. All of these are generous limitations, however, and promise much better storage capacity than alternative ART neural-network implementations.[20–24]

We recently became aware of Ref. 18, in which Chiou describes a strong anisotropy in cross talk in multiplexed volume holograms for both self-pumped photorefractive PCM's and two-wave mixing at the Fourier plane. The increased cross talk reported for angle multiplexing orthogonal to the nominal plane of incidence has implications on the storage capacity of the optical ART neural network. Chiou shows that this increased cross talk may be mitigated for certain classes of input objects by photorefractive mixing at image planes rather than at Fourier planes. This issue requires further study.

5. Future Work

The next stage in our preparations for implementation of the optical ART is the demonstration of a phase-conjugating feedback mirror with gain characteristics suitable for operation with the memory crystal described above. Further experiments will involve determining the geometry and the threshold for the reset detector and demonstrating the ability to defeat a selected template in favor of another. The template selection process can then itself be verified, followed by construction of a full prototype.

6. Conclusions

The optical ART unit is capable of processing large patterns and potentially has a large template capacity. The device's capacity should ultimately approach the capacity of holographic storage systems, which potentially greatly exceed electronic storage capabilities.[9,10] Furthermore, the device is all-optical in the information-processing path; the reset detector's electronics are never used sequentially when an input pattern has already been learned or when it matches an existing template sufficiently. ART has been shown to be useful with extremely large input fields, such as those expected in high-resolution images.[3] The high storage density of this device should enable it to outperform alternative ART implementations,[20–24] which do not offer the same capability for processing the large number of pixels in a high-resolution image. (These implementations are more likely to be used with applications in which speed is more important than the capability to handle large numbers of pixels.) Finally, we have also shown how the architecture provides a convenient means for normalizing the memory templates based on template area, an observation that may be applicable to more general classes of holographic information-retrieval systems.

This work was funded by The Boeing Company, Computer Services Research and Technology Group. We appreciate fruitful discussions with R. Aaron Falk, Gail Carpenter, Robert J. Marks II, Steve Rogers, and Jonathan Kane. We also appreciate several helpful remarks by the reviewers.

References and Notes

1. G. A. Carpenter and S. Grossberg, "A massively parallel architecture for a self-organizing neural pattern recognition machine," Comput. Vis. Graph. Image Process. **37,** 54–115 (1987).
2. S. Grossberg, "Competitive learning: from interactive activation to adaptive resonance," Cogn. Sci. **11,** 23–63 (1987).
3. T. P. Caudell, S. D. G. Smith, G. C. Johnson, and D. C. Wunsch II, "An application of neural networks to group technology," in *Applications of Artificial Neural Networks II,* S. K. Rogers, ed., Proc. Soc. Photo-Opt. Instrum. Eng. **1469,** 612–621 (1991).
4. A. M. Waxman, M. Seibert, R. Cunningham, and J. Wu, "Neural analog diffusion-enhancement layer and spatio-temporal grouping in early vision," in *Advances in Neural Information Processing Systems I,* David Touretzky, ed. (Kaufmann, New York, 1989), pp. 289–296.
5. G. Carpenter, S. Grossberg, and J. Marshall, "ARTMAP: a self-organizing neural network architecture for fast supervised learning and pattern recognition," in *Proceedings of the International Joint Conference on Neural Networks,* (Institute of Electrical and Electronics Engineers, New York, 1991), Vol. 1, pp. 863–868.
6. T. W. Ryan and C. L. Winter, "Variations on adaptive resonance," in *Proceedings of the First International Conference on Neural Networks,* M. Caudill and C. Butler, eds. (Institute of Electrical and Electronics Engineers, New York, 1987), pp. 767–776.
7. B. Moore, "ART1 and pattern clustering," in *Proceedings of the 1988 Connectionist Summer School at Carnegie-Mellon University,* C. Touretsky and G. Hinton, eds. (Kaufmann, New York, 1989).
8. B. H. Soffer, G. J. Dunning, Y. Owechko, and E. Marom, "Associative holographic memory with feedback using phase-conjugate mirrors," Opt. Lett. **11,** 118–120 (1986).
9. F. H. Mok, M. C. Tackitt, and H. M. Stoll, "Storage of 500 high-resolution holograms in a $LiNbO_3$ crystal," Opt. Lett. **16,** 605–607 (1991).
10. D. Psaltis and Y. S. Abu-Mostafa, "Theoretical investigation of parallelism and connectivity in optical computers," in *Optical Processing: Interim Report,* Rep. UDR-TR-85-148 (University of Dayton Research Institute, Dayton, Ohio, 1985), pp. 9.1–9.13.
11. A. Yariv, *Optical Electronics* (Holt, Rinehart & Winston, New York, 1985), pp. 499–525.
12. J. Feinberg, "Optical phase conjugation in photorefractive materials," in *Optical Phase Conjugation,* R. A. Fisher, ed. (Academic, Orlando, Fla., 1983), p. 421. In our ART implementation, phase adjustment and modulation are performed on the reading beam of the photorefractive four-wave PCM. In this way, one can affect the phase of the phase-conjugate beam without directly modifying the phase of the photorefractive grating that is responsible for the phase conjugation.
13. This corresponds to the fast-learning mode of ART, as described in Ref. 1. Slow learning in ART has also been analyzed by G. L. Heileman, M. Georgiopoulos, and C. Abdallah, in "A dynamical adaptive resonance architecture," IEEE Trans. Neural Net. (to be published). This slow learning could be implemented in our device by operating always at the lower power level.
14. D. Psaltis, "Optoelectronic implementations of neural networks," in *Tutorial Notes, International Joint Conference on Neural Networks* (Institute of Electrical and Electronics Engineers, New York, 1990), p. 46.

15. D. Psaltis, D. Brady, and K. Wagner, "Adaptive optical networks using photorefractive crystals," Appl. Opt. **27,** 1752–1759 (1988).
16. J. Feinberg, "Optical phase conjugation in photorefractive materials," in *Optical Phase Conjugation,* R. A. Fisher, ed. (Academic, Orlando, Fla., 1983), pp. 417–443.
17. M. Cronin-Golomb, B. Fischer, J. O. White, and A. Yariv, "Theory and applications of four-wave mixing in photorefractive media," IEEE J. Quantum Electron. **QE-20** (January 1984).
18. A. Chiou, "Anisotropic cross talk in an optical interconnection by using a self-pumped phase-conjugate mirror at the Fourier plane," Opt. Lett. **17,** 1018–1020 (1992).
19. B. Fischer, J. O. White, M. Cronin-Golomb, and A. Yariv, "Nonlinear vectorial two-beam coupling and forward four-wave mixing in photorefractive materials," Opt. Lett. **11,** 239–241 (1986).
20. A. Rao, "VLSI implementation of neural classifiers," Neural Computat. **2,** 35–43 (1989).
21. D. C. Wunsch II, T. P. Caudell, D. Capps, and R. A. Falk, "An optoelectronic adaptive resonance unit," in *Proceedings of the International Joint Conference on Neural Networks* (Institute of Electrical and Electronics Engineers, New York, 1991), Vol. I, pp. 541–549.
22. T. P. Caudell and D. C. Wunsch II, "A hybrid optoelectronic ART1 neural processor," in *Proceedings of the International Joint Conference on Neural Networks* (Institute of Electrical and Electronics Engineers, New York, 1991), Vol. II, p. 926.
23. J. Kane and M. Paquin, "POPART: partial optical implementation of adaptive resonance theory 2," IEEE Trans. Neural Net. (to be published).
24. S. W. Tsay and R. W. Newcomb, "VLSI implementation of ART1 memories," IEEE Trans. Neural Net. **2,** 214–221 (1991).