

# Computing Hough Transforms on Hypercube Multicomputers\*

SANJAY RANKA

*School of Computer and Information Science, 4-116 CST, Syracuse University, Syracuse, NY 13244-4100*

SARTAJ SAHNI

*Department of Computer Science, University of Minnesota, Minneapolis, MN 55455*

(Received January 1989; final version accepted January 1990.)

**Abstract.** Efficient algorithms to compute the Hough transform on MIMD and SIMD hypercube multicomputers are developed. Our algorithms can compute  $p$  angles of the Hough transform of an  $N \times N$  image,  $p \leq N$ , in  $O(p + \log N)$  time on both MIMD and SIMD hypercubes. These algorithms require  $O(N^2)$  processors. We also consider the computation of the Hough transform on MIMD hypercubes with a fixed number of processors. Experimental results on an NCUBE/7 hypercube are presented.

**Key words:** Hough transform, MIMD and SIMD hypercube multicomputers, complexity.

## 1. Introduction

The Hough transform is used to transform edges to another space, called the Hough space, so that the desired group of edges forms a cluster in the transformed space. Let  $I[0 \dots N-1, 0 \dots N-1]$  be an  $N \times N$  image such that  $I[x, y] = 1$  iff the image point  $[x, y]$  is a possible edge point.  $I[x, y] = 0$  otherwise. The  $p$  angle Hough transform of  $I$  to detect straight lines in an image is the array  $H$  such that

$$H[i, j] = |\{(x, y) | i = \lfloor x \cos \theta_j + y \sin \theta_j \rfloor, \theta_j = \frac{\pi}{p}(j+1) \text{ and } I[x, y] = 1\}|. \quad (1)$$

$j$  takes on the integer values  $0, 1, \dots, p-1$ . These correspond to the  $p$  angles  $\theta_j = \frac{\pi}{p}(j+1)$ ,  $0 \leq j < p$ . Hence  $0 < \theta_j \leq \pi$ . For  $\theta_j$  in this range and  $x$  and  $y$  in the range  $0 \dots N-1$ ,  $\lfloor x \cos \theta_j + y \sin \theta_j \rfloor$  is in the range  $-\sqrt{2}N \dots \sqrt{2}N$ . Hence  $H$  is at most a  $2\sqrt{2}N \times p$  matrix.

The general equation of a straight line can be given by the parametric equation

$$x \cos \theta + y \sin \theta = r, \quad (2)$$

where  $\theta$  is the angle that the normal, to the line given by Equation (2), makes with the  $x$  axis and  $r$  is the length of the normal. Any edge point  $(x_i, y_i)$  on this line satisfies the equation

$$x_i \cos \theta + y_i \sin \theta = r. \quad (3)$$

\*This research was supported by the National Science Foundation under grants DCR84-20935 and 86-17374. All correspondence should be mailed to Sanjay Ranka.

The above equation represents a sinusoidal curve in the  $(r, \theta)$  space. Any point on this curve corresponds to a line passing through  $(x_i, y_i)$ . Thus, the curves corresponding to all the points on a line in the  $(x, y)$  space must intersect at the same point in the  $(r, \theta)$  space. Each edge point contributes 1 to the  $(r, \theta)$  cells given by Equation (3) and the cells corresponding to the local maxima give the desired lines.

The generalized Hough transform can be used to recognize curves of arbitrary shapes [Ballard 1981]. It has been used successfully in a wide variety of domains. These include detection of tumors in chest films, recognition of objects in aerial images, and detection of human hemoglobin fingerprints [Ballard and Brown 1982].

The serial algorithm to compute  $H$  has complexity  $O(N^2p)$ . Parallel algorithms to compute  $H$  have been developed by several researchers. Rosenfeld et al. [1988], Cypher et al. [1987], and Guerra and Hambrusch [1987] consider mesh connected multicomputers; Fishburn and Highnam [1987] consider scan line array processors; Ibrahim et al. [1986] consider SIMD tree machines; and Chandran and Davis [1987] consider the use of the Butterfly and NCUBE multicomputers to compute the Hough transform.

In this paper we develop algorithms to compute the above Hough transform on hypercube multicomputers. First, in Section 2, we describe our model for fine-grained MIMD and SIMD hypercubes and how to perform certain fundamental data movement operations on a hypercube. These are used in our subsequent development of hypercube algorithms for the Hough transform. In Section 3 we describe our Hough transform algorithm for the MIMD hypercube. The case of an SIMD hypercube is considered in Section 4. Section 5 considers the computation of the Hough transform on a medium-grained MIMD hypercube. Experimental results on an NCUBE/7 hypercube are also presented in this section.

## 2. Preliminaries

### 2.1. Hypercube Multicomputer

Block diagrams of an SIMD and MIMD hypercube multicomputer are given in Figures 1a and 1b, respectively. The important features of an SIMD hypercube and the programming notation we use follow:

1. There are  $P = 2^p$  processing elements (PEs) connected via a hypercube interconnection network (to be described later). Each PE has a unique index in the range  $[0, 2^p - 1]$ . We shall use brackets  $[ ]$  to index an array and parentheses  $( )$  to index PEs. Thus  $A[i]$  refers to the  $i$ -th element of array  $A$  and  $A(i)$  refers to the  $A$  register of PE  $i$ . Also,  $A[j](i)$  refers to the  $j$ -th element of array  $A$  in PE  $i$ . The local memory in each PE holds data only (that is, no executable instructions). Hence, PEs need to be able to perform only the basic arithmetic operations (that is, no instruction fetch or decode is needed).
2. There is a separate program memory and control unit. The control unit performs instruction sequencing, fetching, and decoding. In addition, instructions and masks are broadcast by the control unit to the PEs for execution. An *instruction mask* is a boolean function used to select certain PEs to execute an instruction. For example, in the instruction

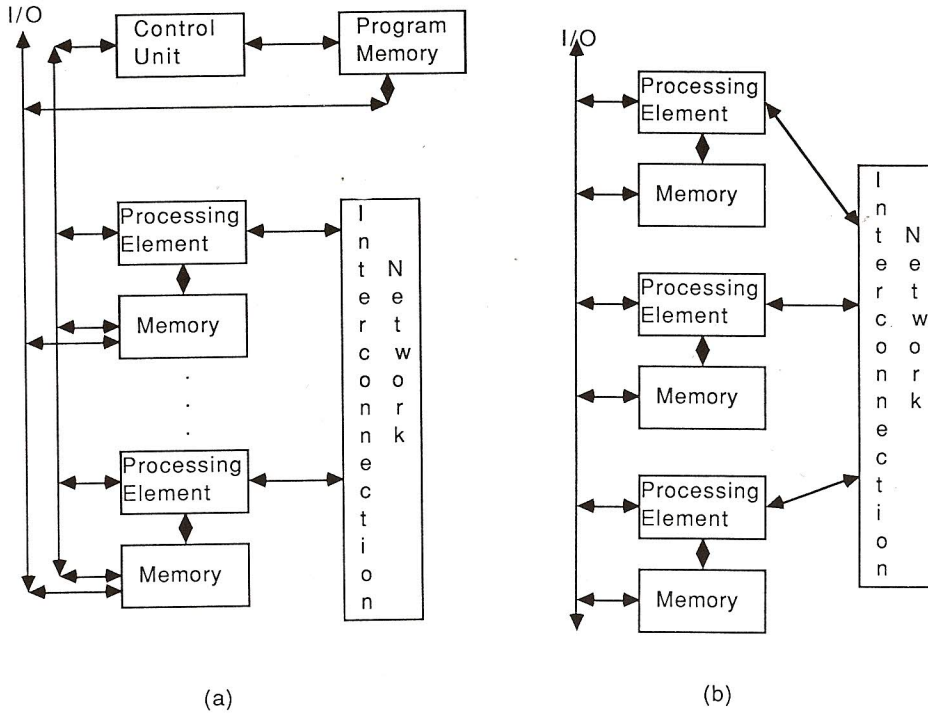


Figure 1. Hypercube multicomputers: a) SIMD hypercube, b) MIMD hypercube.

$$A(i) := A(i) + 1, (i_0 = 1).$$

$(i_0 = 1)$  is a mask that selects only those PEs whose index has bit 0 equal to 1; that is, odd indexed PEs increment their A registers by 1. Sometimes we shall omit the PE indexing of registers. The above statement is therefore equivalent to the statement

$$A := A + 1, (i_0 = 1).$$

3. The topology of a 16-node hypercube interconnection network is shown in Figure 2. A  $p$ -dimensional hypercube network connects  $2^p$  PEs. Let  $i_{p-1}i_{p-2} \dots i_0$  be the binary representation of the PE index  $i$ . Let  $\bar{i}_k$  be the complement of bit  $i_k$ . A hypercube network directly connects pairs of processors whose indices differ in exactly one bit; that is, processor  $i_{p-1}i_{p-2} \dots i_0$  is connected to processors  $i_{p-1} \dots \bar{i}_k \dots i_0$ ,  $0 \leq k \leq p-1$ . We use the notation  $i^{(b)}$  to represent the number that differs from  $i$  in exactly bit  $b$ .
4. Interprocessor assignments are denoted using the symbol  $\leftarrow$ , while intraprocessor assignments are denoted using the symbol  $:=$ . Thus the assignment statement

$$B(i^{(2)}) \leftarrow B(i), (i_2 = 0)$$




Figure 2. A 16-node hypercube (dimension = 4).

is executed only by the processors with bit 2 equal to 0. These processors transmit their B register data to the corresponding processors with bit 2 equal to 1.

5. In a *unit route*, data may be transmitted from one processor to another if it is directly connected. We assume that the links in the interconnection network are unidirectional. Hence, at any given time, data can be transferred either from PE  $i(i_b = 0)$  to PE  $i^{(b)}$  or from PE  $i(i_b = 1)$  to PE  $i^{(b)}$ . Thus the instruction

$$B(i^{(2)}) \leftarrow B(i), (i_2 = 0)$$

takes one unit route, while the instruction

$$B(i^{(2)}) \leftarrow B(i)$$

takes two unit routes.

6. Since the asymptotic complexity of all our algorithms is determined by the number of unit routes, our complexity analysis will count only these.

The features, notation, and assumptions for MIMD hypercubes differ from those of SIMD hypercubes in the following way: There is no separate control unit and program memory. The local memory of each PE holds both the data and the program that the PE is to execute. At any given instance, different PEs may execute different instructions. In particular, PE  $i$  may transfer data to PE  $i^{(b)}$ , while PE  $j$  simultaneously transfers data to PE  $j^{(a)}$ ,  $a \neq b$ .



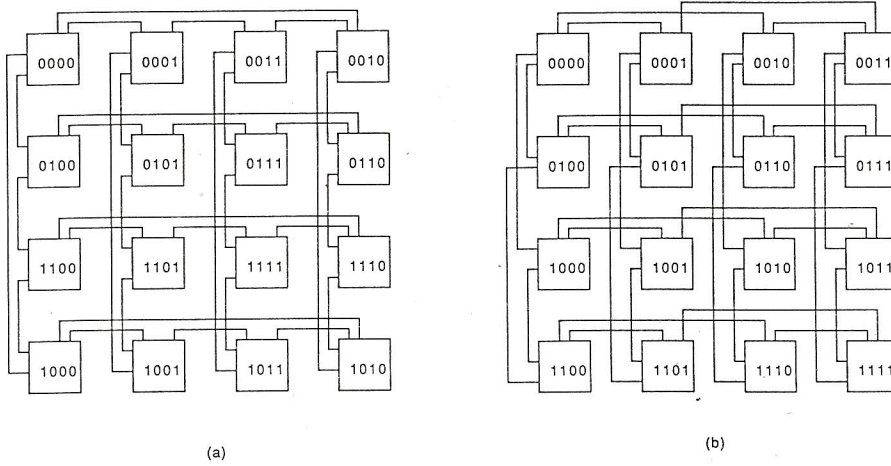


Figure 3. Mapping of the image: a) gray code mapping, b) row major mapping.

## 2.2. Image Mapping

Figure 3a gives a two-dimensional grid interpretation of a four-dimensional hypercube. This is the binary-reflected gray code mapping of [Chan and Saad 1986]. An  $i$  bit binary gray code  $S^i$  is defined recursively as

$$S_1 = 0, 1; S_k = 0[S_{k-1}], 1[S_{k-1}]^R,$$

where  $[S_{k-1}]^R$  is the reverse of the  $k-1$  bit code  $S_{k-1}$  and  $b[S]$  is obtained from  $S$  by prefixing  $b$  to each entry of  $S$ . So,  $S_2 = 00, 01, 11, 10$  and  $S_3 = 000, 001, 011, 010, 110, 111, 101, 100$ .

If  $N = 2^n$ , then  $S_{2n}$  is used. The elements of  $S_{2n}$  are assigned to the elements of the  $N \times N$  grid in a snakelike row major order [Thompson and Kung 1977]. This mapping has the property whereby grid elements that are neighbors are assigned to neighboring hypercube nodes.

Figure 3b shows an alternate embedding of a  $4 \times 4$  image grid into a four-dimensional hypercube. The index of the PE at position  $(i, j)$  of the grid is obtained using the standard row major mapping of a two-dimensional array onto a one-dimensional array [Horowitz and Sahni 1985]. That is, for a  $N \times N$  grid, the PE at position  $(i, j)$  has index  $iN + j$ . With this mapping, a two-dimensional image grid  $I[0 \dots N, 0 \dots N]$  is easily mapped onto an  $N^2$  hypercube (provided  $N$  is a power of 2) with one element of  $I$  per PE. Notice that, in this mapping, image elements that are neighbors in  $I$  (that is, to the north, south, east, or west of one another) may not be neighbors (that is, may not be directly connected) in the hypercube. This does not lead to any difficulties in the algorithms we develop.

We will assume that images are mapped using the gray code mapping for all MIMD algorithms and the row major mapping for all SIMD algorithms.

### 2.3. Basic Data Manipulation Operations

**2.3.1. SIMD Shift.**  $\text{SHIFT}(A, i, W)$  shifts the  $A$  register data circularly counterclockwise by  $i$  in windows of size  $W$ . Thus  $A(qW + j)$  is replaced by  $A(qW + (j - i) \bmod W)$ ,  $0 \leq q < (p/W)$ .  $\text{SHIFT}(A, i, W)$  on an SIMD computer can be performed in  $2 \log W$  unit routes [Prasanna Kumar and Krishnan 1987]. A minor modification of the algorithm given in [Prasanna Kumar and Krishnan 1987] performs  $i = 2^m$  shifts in  $2 \log(W/i)$  unit routes [Ranka and Sahni 1990]. The wraparound feature of this shift operation is easily replaced by an end-off zero fill feature. In this case,  $A(qW + j)$  is replaced by  $A(qW + j - i)$  as long as  $0 \leq j - i < W$ , and by 0 otherwise. This change does not increase the number of unit routes. The end-off shift will be denoted  $\text{EShift}(A, i, W)$ .

**2.3.2. MIMD Shift.** When  $i$  is a power of 2,  $\text{SHIFT}(A, i, W)$  on an MIMD computer can be performed in  $O(1)$  unit routes. An MIMD shift of 1 takes one unit route, of 2 takes two unit routes, of  $N/2$  takes four, and the remaining power of 2 shifts take three routes each. For any arbitrary  $i$  the shift can be completed in  $3(\log W)/2 + 1$  unit routes on an MIMD computer [Ranka and Sahni 1990]. As in the case of the SIMD shift, the MIMD shift is also easily modified to an end-off zero fill shift without increasing the number of unit routes.

**2.3.3. Data Circulation on an SIMD Hypercube.** The data in the  $A$  registers of each of the  $R$  processors in an  $R$  processor subhypercube are to be circulated through each of the remaining  $R-1$  PEs in the subhypercube. This can be accomplished using  $R-1$  unit routes. The circulation algorithm uses exchange sequence  $X_r$ ,  $R = 2^r$  defined recursively as [Dekel et al. 1981]

$$X_1 = 0, X_q = X_{q-1}, q - 1, X_{q-1}(q > 1).$$

This sequence essentially treats a  $q$ -dimensional hypercube as two  $q - 1$ -dimensional hypercubes. Data circulation is done in each of these in parallel using  $X_{q-1}$ . Next, an exchange is done along bit  $q - 1$ . This causes the data in the two halves to be swapped. The swapped data are again circulated in the two half hypercubes using  $X_{q-1}$ . Let  $f(r, i)$  be the  $i$ -th number (left to right) in the sequence  $X_r$ ,  $1 \leq i < 2^r$ . The resulting SIMD data circulation algorithm is given in Figure 4. Here, it is assumed that the  $r$  bits that define the subhypercube are bits  $0, 1, 2, \dots, r - 1$ . Because of our assumption of unidirectional links, each iteration of the *for* loop of Figure 4 takes two unit routes. Hence Figure 4 takes  $2(R - 1)$  unit routes. The function  $f$  can be computed by the control processor in  $O(R)$  time and saved in an array of size  $R - 1$  (actually it is convenient to compute  $f$  on the fly using a stack of height  $\log R$ ). The following lemma allows each processor to compute the origin of the current  $A$  value.

**LEMMA 1.** [Ranka and Sahni 1990] Let  $A_0, A_1, \dots, A_{2^r-1}$  be the values in  $A(0), A(1), \dots, A(2^r - 1)$  initially. Let  $\text{index}(j, i)$  be such that  $A[\text{index}(j, i)]$  is in  $A(j)$  following the  $i$ -th iteration of the *for* loop of Figure 4. Initially,  $\text{index}(j, 0) = j$ . For every  $i, i > 0$ ,  $\text{index}(j, i) = \text{index}(j, i-1) \theta 2^{f(r,i)}$  ( $\theta$  is the exclusive or operator).

```

procedure CIRCULATE(A);
  [data circulation]
  for i := 1 to R - 1 do
     $A(j^{f(r,i)}) \leftarrow A(j)$ ;
  end

```

Figure 4. Data circulation in an SIMD hypercube.

```

procedure ACCUM(A,I,M)
  {each PE accumulates in A, the I values of the next
  M PEs, including itself; P is the window size}
  begin
    A[0] := I;
    for i := 1 to M - 1 do
      begin
        SHIFT(I, -1, P) :
        A[i] := I;
      end
    end {ACCUM}

```

Figure 5. Data accumulation.

**2.3.4. Data Accumulation on an MIMD Hypercube.** For this operation, PE  $j$  has an array  $A[0 \dots M - 1]$  of size  $M$ . In addition, each PE has a value in its  $I$  register. After the data accumulation, the  $M$  elements of  $A$  in each PE  $j$  are such that

$$A[i] \text{ (gray } (j)) = I \text{ (gray } ((j + i) \bmod P)), 0 \leq i < M, 0 \leq j < P.$$

This can be accomplished in  $M - 1$  unit routes (for  $P > 2$ ) by repeatedly shifting by  $-1$  in windows of size  $P$ . The algorithm is given in Figure 5.

**2.3.5. Data Accumulation on an SIMD Hypercube.** After the data accumulation, the  $M$  elements of  $A$  in each PE  $j$  are such that

$$A[i] \text{ (} j) = I \text{ ((} j + i) \bmod P), 0 \leq i < M, 0 \leq j < P.$$

Data accumulation may be done efficiently by modifying the data circulation algorithm. It can be completed in  $2(M - 1) + \log_2(N/M)$  unit routes on an SIMD hypercube.

### 2.4. Initial and Final Configurations

We shall explicitly consider the computation of  $H(i, j)$  only for  $i > 0$ . The computation for the case  $i \leq 0$  is similar. Hence,  $i$  is in the range  $[0, \sqrt{2}N)$  and  $j$  is in the range  $[0, p)$ . We assume that  $N$  is a power of two and that  $2N^2$  PEs are available. These are viewed as an  $N \times 2N$  array, as discussed in Section 2.2 for SIMD and MIMD hypercubes. Actually, only  $N \times \sqrt{2}N$  PEs are needed; however, a hypercube must have a power of 2 processors. Furthermore, it is assumed that  $p$  divides  $N$ .

The image pixel  $I[i, j]$  is initially stored in PE  $[i, j]$   $0 \leq i, j < N$  in the above array view.  $H[i, j]$  is stored in PE  $[j, i]$  on completion.

### 3. MIMD Algorithm

Conceptually, our algorithm is similar to that of Cypher et al. [1987]. It computes the Hough transform in  $O(p + N)$  time on an  $N \times N$  SIMD mesh-connected computer. We show how this algorithm can be mapped onto an MIMD hypercube with  $2N^2$  processors. The complexity of the resulting hypercube algorithm is  $O(p + \log N)$ .

For simplicity, we divide the computation of  $H[i, j]$ ,  $i > 0$ ,  $0 \leq j < p$  into four parts. These, respectively, correspond to the cases  $0 \leq j < p/4$ ,  $p/4 \leq j < p/2$ ,  $p/2 \leq j < 3p/4$ , and  $3p/4 \leq j < p$ . First, consider the case  $p/4 \leq j < p/2$ . Now,  $\pi/4 < \theta_j \leq \pi/2$ . The following lemmas will suggest a computational scheme for this case.

**LEMMA 3.1.** When  $\pi/4 < \theta_j \leq \pi/2$ , two pixels  $(x, y)$  and  $(x, y + z)$ ,  $z > 0$ , can contribute to the count of the same  $H[i, j]$  only if  $z = 1$ .

*Proof.* If  $(x, y)$  and  $(x, y + z)$  both contribute to the count of  $H[i, j]$  then

$$i = \lfloor x \cos \theta_j + y \sin \theta_j \rfloor = \lfloor x \cos \theta_j + (y + z) \sin \theta_j \rfloor$$

for some  $j$ ,  $p/4 \leq j < p/2$ . Hence,

$$(y + z) \sin \theta_j - y \sin \theta_j \leq 1$$

$$\text{or } z \sin \theta_j \leq 1.$$

Since  $\pi/4 < \theta_j \leq \pi/2$ ,  $\sin \theta_j > \sin \pi/4 > 0.5$ . Since  $z$  is a positive integer, only  $z = 1$  can satisfy the relation  $z \sin \theta_j \leq 1$ .

**LEMMA 3.2.** When  $\pi/4 < \theta_j \leq \pi/2$ , two pixels  $(x, y)$  and  $(x + 1, z)$  can contribute to the count of the same  $H[i, j]$  only if  $z \in \{y, y - 1\}$ .

*Proof.* If  $(x, y)$  and  $(x + 1, z)$  contribute to the same  $H[i, j]$ , then  $i = \lfloor x \cos \theta_j + y \sin \theta_j \rfloor = \lfloor (x + 1) \cos \theta_j + z \sin \theta_j \rfloor$ .

$$\text{So, } |(x + 1) \cos \theta_j - x \cos \theta_j + (z - y) \sin \theta_j| \leq 1$$



$$\begin{aligned}
&\text{or } |\cos \theta_j + (z - y) \sin \theta_j| \leq 1 \\
&\text{or } |\cot \theta_j + (z - y)| \leq \operatorname{cosec} \theta_j \\
&\text{or } -\operatorname{cosec} \theta_j - \cot \theta_j \leq z - y \leq \operatorname{cosec} \theta_j - \cot \theta_j.
\end{aligned}$$

Since  $y$  and  $z$  are integers and  $\theta_j$  is in the above range, it follows that  $-1 \leq z - y \leq 0$ . Hence,  $z \in \{y, y - 1\}$ .

The computation of  $H[i, j]$  for  $i > 0$  and  $\pi/4 \leq \theta_j < \pi/2$  can be done in two phases. In the first, subhypercubes of size  $p \times 2N$  compute  $h[i, j] = |\{(x, y) | i = \lfloor x \cos \theta_j + y \sin \theta_j \rfloor, \pi/4 \leq \theta_j < \pi/2, I[x, y] = 1, \text{ and } (x, y) \text{ is in this subhypercube}\}|$ . In the second phase, the  $h[i, j]$  values from the different subhypercubes are summed to obtain

$$H[i, j] = \sum_{\text{subhypercubes}} h[i, j], \quad i > 0, \quad p/4 \leq j < p/2.$$

The phase 1 algorithm for each PE in a  $p \times 2N$  subhypercube is given in Figure 6. In this algorithm,  $[x, y]$  denotes a PE index relative to the whole  $N \times 2N$  hypercube and  $[w, y]$  denotes the index of the same PE relative to the  $p \times 2N$  subhypercube it is in. Note that  $w = x \bmod p$ .

The  $h$  values are computed in a pipeline manner. The PEs in row 0 of a  $p \times 2N$  subhypercube initiate a record  $Z = (i, j, \text{sine}, \text{cosine}, q)$  such that  $h[i, j] = q$  is the number of pixels on this row that contribute to  $h[i, j]$ . This is done by first computing  $i$  for each pixel in row zero (line 7) for a fixed  $j = p/2 - \ell - 2$ . Lemma 3.1 is used in lines 22–24 to combine records that represent the same  $h[i, j]$  entry. This row of  $Z$  records created in row zero moves down the  $p \times 2N$  subhypercube one row per iteration (line 25). Lines 10–21 update the row of  $Z$  values received. Each such row corresponds to a fixed  $j$ . For this  $j$ ,  $\text{PE}[w, y]$  determines the  $h$  entry  $[i', j]$  to which it is to contribute (line 13). If this is the same entry as received from  $\text{PE}[w - 1, y]$ , then the two are added together. If  $i = \phi$  for the received entry, then  $[i', j]$  can occupy this  $Z$  space. If  $i = \phi$ , then from Lemma 3.2 we know that  $Z$  can combine only with the new entry  $[i', j]$  of  $\text{PE}[w, y - 1]$ .

Following the iteration  $\ell = 5p/4 - 1$ , the last initiated row (i.e.,  $j = p/4$ ) has passed through row  $p - 1$  of the  $p \times 2N$  subhypercube. At this time, the PEs in row  $r$  of the subhypercube contain records with  $j = p/4 + r$ ,  $0 \leq r < p/4$ . The records in each row may be reordered such that the record in  $\text{PE}[w, y]$  has  $y = i$  by performing a random access write [Nassimi and Sahni 1981]. Because of the initial ordering of  $i$  values in a row, this random access write can be performed in  $O(\log N)$  time [Ranka and Sahni 1990] rather than in  $O(\log^2 N)$  time as required by the more general algorithm [Nassimi and Sahni 1981].

The phase 2 summing of the  $h[i, j]$  values is now easily done in  $O(\log N)$  time using a window sum. Since the phase 1 algorithm of Figure 6 only shifts by 1 along columns and/or rows, each iteration of this algorithm takes only  $O(1)$  time. Hence, the complexity of the phase 1 algorithm is  $O(p)$ . The overall time needed to compute  $H$  for  $p/4 \leq x < p/2$  is therefore  $O(p + \log N)$ .

The remaining three cases for  $j$  are done in a similar way. Actually, the four cases need not be computed independently as suggested above. In particular, all the computation following phase 1 can be done in parallel for all the cases.

```

1 for  $\ell := 0$  to  $5p/4 - 1$  do
2   if  $(w = 0)$  and  $(\ell < p/4)$  then
3     [{row 0 initiates next  $\theta_j$ }]
4     create a record  $Z = (i, j, \text{sine}, \text{cosine}, q)$ 
5     with
6      $\text{sine} = \sin(\theta)$ ,  $\text{cosine} = \cos(\theta)$ , where  $\theta = \frac{\pi}{p}(p/2 - \ell + m)$ 
7      $i = \lfloor x \text{ cosine} + y \text{ sine} \rfloor$ ,  $j = p/2 - \ell - 1$ 
8      $q = I[x, y]$ 
9   else if  $\max\{1, \ell - p/4 + 1\} \leq w \leq \ell$  and  $y < N$ 
10    then {add in this PE's contribution}
11      [Let  $Z$  be the record received from  $\text{PE}[w - 1, y]$ 
12      Let  $i' = \lfloor x \text{ cosine} + y \text{ sine} \rfloor$  and  $q' = I[x, y]$ 
13      if  $i = i'$  then set  $q = q + q'$ 
14      else if  $i = \phi$  then set  $i = i'$  and  $q = q'$ 
15      else [send  $q$  to  $\text{PE}[x, (y - 1) \bmod 2N]$ 
16            set  $Z = (i', j, \text{sine}, \text{cosine}, q')$ ]
17      if a  $q$  is received from  $\text{PE}[x, y + 1]$  update own  $q$ 
18      to  $q + \text{received } q$ ]
19    else if  $y > N$  and a  $Z$  is received from  $\text{PE}[x, (y + 1) \bmod 2N]$ 
20      then send old  $Z$  (if any) to PE on left ]
21    {combine records with same  $(i, j)$  values}
22    if  $(\lfloor x \text{ cosine} + y \text{ sine} \rfloor = \lfloor x \text{ cosine} + (y - 1) \text{ sine} \rfloor)$  and  $(0 < y < N)$ 
23      then send  $h$  to  $\text{PE}[x, y - 1]$  and set  $i = \phi$ 
24      else if a  $q$  value is received set  $q = q + \text{received } q$ ;
25    send  $Z$  to  $\text{PE}[(w + 1) \bmod p, y]$ 
26 end

```

Figure 6. MIMD algorithm.

#### 4. SIMD Algorithms

We develop two  $O(p + \log N)$  SIMD hypercube algorithms. One uses  $O(\log N)$  memory per PE while the other uses  $O(1)$ . The  $O(1)$  memory algorithm is slightly more complex than that with  $O(\log N)$  memory. Both algorithms are adaptations of our MIMD algorithm. The computations following phase 1 (Figure 6) are easily performed in  $O(\log N)$  time on an SIMD hypercube using  $O(1)$  memory per PE. So we concentrate on adapting phase 1. The

phase 1 algorithm performs  $O(p)$  unit shifts along the rows and columns of  $p \times 2N$  subhypercubes. In an SIMD hypercube, each such row shift takes  $O(\log N)$  time while each unit column shift takes  $O(\log p)$  time. So a direct simulation of phase 1 takes  $O(p \log(Np))$  time.

#### 4.1. $O(\log N)$ Memory per PE

In this case, we divide the  $5p/4$  iterations for the *for* loop of Figure 6 into blocks of  $\log N$  consecutive iterations. In each such block, a  $Z$  record initially in  $PE[x, y]$  can be augmented by pixel values in  $PEs [x + \ell, y - m]$ ,  $0 \leq \ell < \log N$ ,  $-1 \leq m < \log N$ . To avoid unit shifts along the rows, each  $PE[q, r]$  begins by accumulating the pixel value in  $PE[q, r - m]$ ,  $-1 \leq m < \log N$ . Now it is necessary to route the  $Z$  records only down a column; that is, a  $Z$  record initially in  $PE[x, y]$  needs to visit  $PEs [x + \ell, y]$ ,  $0 \leq \ell < \log N$ . These  $PEs$  contain the pixel values needed to update  $Z$  to its values following the block of iterations in Figure 6. This routing is done using the circulation algorithms in windows of size  $\log N$  rather than by unit shifts. The initial pixel accumulation takes  $O(\log N)$  time, and the circulation and  $Z$  updates also take  $O(\log N)$  time. Following the circulation, the  $Z$  records return to their originating  $PEs$  and need to be routed left and down by a distance of  $O(\log N)$ . This can be accomplished in  $O(\log N)$  time on an SIMD hypercube. In this way, we are able to simulate  $O(\log N)$  iterations of the MIMD algorithm in  $O(\log N)$  time on an SIMD hypercube. Hence, the overall asymptotic run time of the SIMD simulation is the same as that of the original MIMD algorithm.

#### 4.2. $O(1)$ Memory per PE

When  $(\log^2 N)/p \leq c$  for some constant, a careful analysis shows that using the strategy employed in the  $O(\log N)$  memory algorithm, the memory requirements can be reduced to  $O(1)$ . In any  $\log N$  block of iterations, two pixels  $[x, y]$  and  $[w, z]$  contribute to the same  $Z$  record only if

$$\lfloor x \cos \theta + y \sin \theta \rfloor = \lfloor w \cos \theta + z \sin \theta \rfloor.$$

Since  $w \leq x + \log N - 1$  during the  $\log N$  iterations, we get

$$\begin{aligned} |(\log N - 1) \cos \theta + (z - y) \sin \theta| &\leq 1 \\ \text{or } -\operatorname{cosec} \theta &\leq (\log N - 1) \cot \theta + z - y \leq \operatorname{cosec} \theta \\ \text{or } -\operatorname{cosec} \theta &\leq (\log N - 1) \cot \theta \leq z - y \leq \operatorname{cosec} \theta - (\log N - 1) \cot \theta. \end{aligned}$$

For any fixed  $\theta \in [\pi/4, \pi/2]$ ,

$$\begin{aligned} z &\in [y - (\log N - 1) \cot \theta - \operatorname{cosec} \theta, y - (\log N - 1) \cot \theta + \operatorname{cosec} \theta] \\ \text{or } z &\in [y - (\log N - 1) \cot \theta - \sqrt{2}, y - (\log N - 1) \cot \theta + \sqrt{2}]. \end{aligned}$$

There is only a constant number of integers in this range. During a  $\log N$  block of iterations,  $Z$  records with  $j$  value differing by up to  $\log N - 1$  may pass through a given  $PE$ . This corresponds to a  $\theta$  variation from  $\theta_1$  to  $\theta_2$  where  $\theta_2 - \theta_1 = \frac{\pi}{p} (\log N - 1)$ .



Hence, the leftmost column from which a contributing pixel is required has a maximum range of

$$\begin{aligned}
 & \operatorname{cosec} \theta_1 + (\log N - 1) \cot \theta_1 - \operatorname{cosec} \theta_2 - (\log N - 1) \cot \theta_2 \\
 & \leq \operatorname{cosec} \theta_1 - \operatorname{cosec} \theta_2 + (\log N - 1)(\cot \theta_1 - \cot \theta_2) \\
 & \leq \operatorname{cosec} \pi/4 + (\log N - 1) \frac{\cos \theta_1 \sin \theta_2 - \cos \theta_2 \sin \theta_1}{\sin \theta_1 \sin \theta_2} \\
 & < \operatorname{cosec} \pi/4 + 2(\log N - 1) \sin(\theta_2 - \theta_1) \\
 & < \operatorname{cosec} \pi/4 + 2(\log N - 1)(\theta_2 - \theta_1) \\
 & = \operatorname{cosec} \frac{\pi}{4} + 2(\log N - 1)(\log N - 1)\pi/p \\
 & < \operatorname{cosec} \pi/4 + 2\pi c.
 \end{aligned}$$

Hence, each PE need accumulate only a constant number of pixels from its row rather than the  $O(\log N)$  pixels being accumulated in the  $O(\log N)$  memory algorithm. This accumulation is done in  $O(\log N)$  time. The run time is the same as that of the  $O(\log N)$  memory algorithm, but the memory requirements are reduced to  $O(1)$ .

## 5. Hough Transform on the NCUBE Hypercube

### 5.1. NCUBE Architecture

In the previous sections we have developed algorithms to compute the Hough transform on a fine-grained hypercube. In such a computer, the cost of interprocessor communication is comparable to that of a basic arithmetic operation. In this section, we shall consider the Hough transform on a hypercube in which interprocessor communication is relatively expensive and the number of processors is small relative to the number of patterns  $N$ . In particular, we shall experiment with an NCUBE/7 hypercube which is capable of having up to 128 processors. The NCUBE/7 available to us, however, has only 64 processors. The time to perform a 2-byte integer addition on each hypercube processor is 4.3 microseconds, whereas the time to communicate  $b$  bytes to a neighbor processor is approximately  $447 + 2.4b$  microseconds. Figure 7 shows the block diagram for the NCUBE/7 hypercube multicomputer.

The size of the image and the Hough array is  $O(N^2)$  and  $O(Np)$ , respectively. These sizes are comparable for typical values of  $N$  and  $p$  ( $N$  is typically 1024 or 2048, while  $p$  varies from 45 to 180). The Hough array may be smaller, if the Hough transform is calculated for a smaller set of angles. Note that in a digitized image,  $p$  has an upper bound of  $O(N)$  under reasonable assumptions. Due to the large amount of memory requirements needed to store the image array and the Hough space, we feel that it is unreasonable to assume that every node has a copy of the image and/or the Hough array. This will be true in any parallel integrated vision system in which information will be saved at the nodes between intermediate stages. Thus, the amount of memory available for performing the Hough transform will only be a part of the total memory. The total amount of memory available on each NCUBE/7 node is only 512K (including the space for system routines, system stacks, message buffers, and program code). Thus, the image has to be initially subdivided among all the nodes. The division of the image among all the nodes will also be required by any parallel edge detection algorithm to be performed before Hough transform computation.






Figure 7. NCUBE/7 hypercube.




Figure 8. A ring (of size 8) embedded in a hypercube of eight nodes.

The Hough array also needs to be distributed among all the nodes due to its size. This will also be preferred by a parallel local maxima finding or clustering algorithm to find the points in the Hough space corresponding to the lines in the image space (after the Hough transform has been performed).

The rest of this section will assume the image and the Hough array have to be distributed because of the above reasons.

## 5.2. Two NCUBE Algorithms

We view the  $P$  hypercube nodes as forming rings. Figure 8 shows this ring for the case  $P = 8$ . For any node  $i$ , let left ( $i$ ) and right ( $i$ ), respectively, be the node counterclockwise and clockwise from node  $i$ . Let logical ( $i$ ) be the logical index of node  $i$  in the ring. The  $N \times N$  image array is initially distributed over the nodes with each node getting an  $N \times N/qp$  block. Logical node 0 gets the first block, logical node 1 the next block, and so on. Similarly, on completion, the  $2\sqrt{2}N \times p$  Hough array  $H$  is distributed over the nodes in blocks of size  $2\sqrt{2}N \times p/P$ . We assume that the number of hypercube nodes  $P$  divides the number of angles  $p$  as well as the image dimension  $N$ . It is further assumed that the thresholding function has already been applied to the pixels and each node has a list of pairs  $(x, y)$  such that  $I[x, y]$  passes the threshold. We call this list the edge list for the node.

```

procedure UpdateHpartition ( $H$ )
  for each  $(x, y)$  in edge list do
    for  $j := j\text{Begin}$  to  $j\text{Begin} + \text{size} - 1$  do
       $\theta = \frac{\pi}{p}(j + 1)$ 
       $i = x \cos \theta + y \sin \theta$ 
      increment  $H[i, \theta]$  by 1
    end;
  end;
end; {of UpdateHpartition}
 $\ell :=$  logical index of this node,  $\text{size} := p/P$ ;
 $j\text{Begin} := \text{size} * \ell$ 
initialize own  $H$  partition to zero;
for  $i := 0$  to  $P - 1$  do
  UpdateHpartition;
  send own  $H$  partition to node on right;
  receive  $H$  partition from node on left;
   $j\text{Begin} := (j\text{Begin} - \text{size}) \bmod p$ ;
end;

```

Figure 9. Nonoverlapping algorithm to compute  $H$ .


Our first algorithm is given in Figure 9. This algorithm is run on each hypercube node. As remarked earlier, each node has an edge list and an  $H$  partition.

The  $H$  partitions move along the ring one node at a time. When an  $H$  partition reaches any node, the edge list of that node is used to update it, accounting for all contributions these edges make to this  $H$  partition. Procedure UpdateH partition does precisely this.  $j\text{Begin}$  is the  $j$  value corresponding to the first angle (column) in the  $H$  partition currently in the node.  $\text{size} = p/P$  is the number of columns in an  $H$  partition.

In the algorithm of Figure 9 no attempt is made to overlap computation with communication. Following the send of an  $H$  partition to its right neighbor, the node is idle until the receive of the  $H$  partition from its left neighbor is complete. Figure 10 shows the activity of a node as a function of time.

During the compute phase, an  $H$  partition is updated. Let  $t_c$  be the time needed to do this. Let  $t_t$  be the time for an  $H$  partition to travel from a sending node to its destination node. So  $t_t$  is the elapsed time between the initiation of the transfer and the receipt of the partition. The time required by the nonoverlapping algorithm of Figure 9 is  $P(t_c + t_t)$ .

Our second algorithm (Figure 11) attempts to overlap much of the transmission time  $t_t$  with computation. This, unfortunately, increases the computation time since some additional work must be done. At the end of each iteration of the *for* loop, the  $H$  partition in a node  $\ell$  is sent to the node on its right. The next iteration proceeds while the  $H$  partition is in

Figure 10. Nonoverlapping algorithm to compute  $H$ .

```

 $\ell :=$  logical index of this node; size =  $p/P$ ;
 $jBegin :=$  size *  $\ell$ ;
for  $i := 0$  to  $P - 1$  do


    if  $i = 0$  then [initialize own  $H$  partition to zero
                   UpdateHPartition ( $H$ )]
    else [initialize  $T$  to zero
          UpdateHPartition ( $T$ )
          Receive  $H$  Partition from left ( $\ell$ )
           $H := H + T$ ]

    send  $H$  to right ( $\ell$ );
     $jBegin := (jBegin - \text{size}) \bmod p$ 
end;
```

Figure 11. Overlapping algorithm for  $H$ .

transit. For this, a temporary space  $T$  of the same size as  $H$  is used to accumulate the contribution of the node's edge list to the  $H$  partition it has yet to receive from its left neighbor. Following this computation, the received  $H$  portion and  $T$  are added as the resulting  $H$  partition is transmitted to the right.

Relative to the nonoverlapping algorithm, the overlapping algorithm does  $P - 1$  initializations of  $T$  and executions of  $H := H + T$  extra computational work. Let  $t_{init}$  be the time to initialize  $T$  and  $t_{add}$  the time to execute  $H := H + T$ . If  $t_t \leq t_{init} + t_c$ , the time diagram has the form shown in Figure 12a. The overall time for the algorithm is  $Pt_c + (P - 1)(t_{init} + t_{add}) + t_t$  when  $t_t \leq t_{init} + t_c$ . So if  $t_{init} + t_{add} < t_t$ , the overlapping algorithm will outperform the nonoverlapping algorithm.

Figure 12a.  $t_t \leq t_{init} + t_c$ .




Figure 12b.  $t_t = t_{init} + t_c \Delta$ ,  $\Delta > 0$ .

When  $t_t = t_{init} + t_c + \Delta t$ ,  $\Delta t > 0$ , the time diagram is as in Figure 12b. In this case, the algorithm run time is  $t_c + (P - 1)t_{add} + Pt_t = Pt_c + (P - 1)(t_{init} + t_{add} + \Delta) + t_t$ . For the overlapping algorithm to outperform the nonoverlapping algorithm, we need  $t_{init} + t_{add} + \Delta < t_t$ .

### 5.3. Load Balancing

The preceding analysis is somewhat idealistic since it assumes that  $t_c$  is the same in each node. Actually, the size of the edge list in each node is different and this difference significantly affects the performance of the algorithm. The node with the maximum number of edges becomes a bottleneck. To reduce the run time, one may attempt to obtain an equal or nearly equal distribution of the edges over the  $P$  nodes. Note that even though the image matrix  $I$  is equally distributed over the nodes, the edge lists may not be, since a different number of pixels in each  $I$  partition will pass the threshold. We shall use the term *load* to refer to the number of pixels in a node that passes this threshold. That is, load is the size of the nodes' edge list. Two heuristics to balance the load are given in Figures 13 and 14.

In both, load balancing is accomplished by averaging over the load in processors that are directly connected. The variables used have the following significance:

- MyLoad = current load in the node processor
- HisLoad = load in a directly connected node processor
- MyLoadSize = size of the load in the node processor
- His Load Size = size of the load in a directly connected node processor
- avg = average size of the load of the two processors

The only difference between the two variations is that in the first a processor transmits its entire work load (including the necessary data) to its neighbor processor, while in the second variation only the amount in excess of the average is transmitted. However, to achieve this reduction in load transmission, it is necessary to first determine how much of the load is to be transmitted. This requires an initial exchange of the load size. Hence, variation 2 requires twice as many message transmissions. Each message of variation 2 is potentially shorter than each message transmitted by variation 1. We expect variation 1 to be faster than variation 2 when the number of bytes in MyLoad and HisLoad is relatively small and the time to set up a data transmission is relatively large. Otherwise, variation 2 is expected to require less time.



```

procedure LoadBalance1();
  for  $i := 0$  to CubeSize do
    Send MyLoad to neighbor processor along dimension  $i$ ;
    Receive HisLoad from neighbor processor along dimension  $i$ 
    and append to Myload;
     $avg = (MyLoadSize + HisLoadSize + 1) / 2$ ;
    if ( $MyLoadSize > Avg$ )  $MyLoadSize = Avg$ ;
    else if ( $HisLoadSize > Avg$ )  $MyLoadSize += HisLoadSize - Avg$ ;
  end;
end;

```

Figure 13. Load balancing (heuristic 1).

```

procedure LoadBalance2();
{
  for  $i := 0$  to CubeSize do
    Send MyLoadSize to neighbor processor along dimension  $i$ ;
    Receive HisLoadSize from neighbor processor along
    dimension  $i$ ;
     $avg = (MyLoadSize + HisLoadSize + 1) / 2$ ;
    if ( $MyLoadSize > Avg$ ) [
      Send extra load ( $MyLoadSize - Avg$ ) to neighbor
      processor along dimension  $i$ ;
       $MyLoadSize = Avg$ ; ]
    else if ( $HisLoadSize > Avg$ ) [
      Receive extra load ( $Avg - HisLoadSize$ ) from neighbor
      processor along dimension  $i$ 
       $MyLoadSize += HisLoadSize - Avg$ ; ]
  end;
end;

```

Figure 14. Load balancing (heuristic 2).

#### 5.4. Experimental Results

The nonoverlapping and overlapping algorithm of Section 5.2, as well as the load balancing heuristics of Section 5.3, were programmed in *C* and run on an NCUBE/7 hypercube with

64 nodes. We experimented with randomly generated images of size  $N \times N$  for  $N = 32, 64, 128, 256$ , and  $512$ . The percentage of pixels in an  $N \times N$  image that passed the threshold was fixed at 5%, 10%, or 20%. The number of edge pixels in each nodes  $I$  partition was determined using a truncated normal distribution with variance being one of 4%, 10%, and 64% of the mean. In all cases, we arbitrarily set  $p = 180$ . We keep  $p$  fixed for all image sizes to observe the effects of increase in the image size on the speedup achieved, with the same number of processors.

Preliminary experiments indicated that the run time of our two load-balancing heuristics was approximately the same, with the second heuristic having a slight edge. Furthermore, the time to load balance is less than 2% of the overall run time (load balance followed by Hough transform computation). The run times of the nonoverlapping algorithm, both with and without load balancing, are given in Tables 1, 2, and 3 for the cases of  $P = 4, 16$ , and  $64$ , respectively. We see that as the load variance increases from 4% to 64%, the run time of the nonoverlapping algorithm without load balancing increases significantly.

Table 1. Run times (in seconds) for nonoverlapping algorithm,  $P = 4$ .

Image Size	%	No Load Balancing			Load Balance 2		
		Variance					
		edges	4%	16%	64%	4%	16%
32×32	5	0.2802	0.3138	0.3940	0.2819	0.2785	0.2804
	10	0.5627	0.6035	1.1563	0.5531	0.5527	0.5527
	20	1.1439	1.3364	1.7874	1.0976	1.0956	1.0967
64×64	5	1.1465	1.3044	1.7575	1.1202	1.1187	1.1176
	10	2.2428	2.4485	3.4152	2.1878	2.1853	2.1818
	20	4.4974	4.7970	8.0548	4.3171	4.3238	4.3190
128×128	5	4.4966	5.0359	7.8626	4.3605	4.3550	4.3564
	10	8.9968	10.0017	15.5813	8.6474	8.6393	8.6423
	20	18.0087	19.1119	31.7456	17.2247	17.2349	17.2108

Table 2. Run times (in seconds) for nonoverlapping algorithm,  $P = 16$ .

Image Size	%	No Load Balancing			Load Balance 2		
		Variance					
		edges	4%	16%	64%	4%	16%
64×64	5	0.2964	0.3494	0.5622	0.2981	0.3012	0.2922
	10	0.5949	0.6803	1.1556	0.5927	0.5830	0.5712
	20	1.1827	1.4140	2.0260	1.1615	1.1574	1.1313
128×128	5	1.2088	1.4113	2.2415	1.1915	1.1798	1.1570
	10	2.3558	2.7429	5.3075	2.3256	2.3065	2.2469
	20	4.6616	5.3293	9.0813	4.5909	4.5600	4.4518
256×256	5	4.6854	5.6429	9.3724	4.6283	4.5810	4.4624
	10	9.3130	11.0024	18.0237	9.1721	9.1270	8.8296
	20	18.4712	21.4809	33.9781	18.2738	18.1359	17.6917

In fact, it almost doubles. With load balancing, however, the run time is quite stable. Furthermore, it is always less than the run time for 4% variance without load balancing. When the variance in load is 64%, load balancing results in a 25% to 53% reduction in run time!

Note that the average load per node when  $P = 4$  and  $N = 128$  is the same as when  $P = 16$  and  $N = 256$  and when  $P = 64$  and  $N = 512$ . From Tables 1-3 we see that run time remains virtually unchanged as  $P$  increases, provided the load per node is unchanged. Hence, the algorithm scales well.

The run times for the overlapping algorithm with load balancing are given in Tables 4 and 5. These times are generally slightly larger than those for the nonoverlapping algorithm with load balancing. So, the computational overhead introduced by the overlapping algorithm more or less balances the positive effects of overlapping computation and communication.

For comparison purposes, the run times on a single hypercube node are given in Table 6 for the cases  $N = 16, 32$ , and  $64$ . The case  $N = 128$  could not be run for lack of sufficient memory.

Table 3. Run times (in seconds) for nonoverlapping algorithm,  $P = 64$ .

Image Size	%	No Load Balancing			Load Balance 2		
		Variance					
	edges	4%	16%	64%	4%	16%	64%
$128 \times 128$	5	0.3462	0.4200	0.6449	0.3416	0.3481	0.3400
	10	0.6512	0.7735	1.3975	0.6313	0.6315	0.6232
	20	1.2692	1.5239	2.8156	1.2051	1.2062	1.1960
$256 \times 256$	5	1.2638	1.5371	2.7062	1.2291	1.2229	1.2057
	10	2.4770	2.9324	5.1395	2.3543	2.3476	2.3288
	20	4.9057	6.0051	11.4614	4.6170	4.6020	4.5470
$512 \times 512$	5	4.9077	5.8094	10.8207	4.6485	4.6232	4.5784
	10	9.7492	11.7256	20.7631	9.1908	9.1611	9.0623
	20	19.3672	23.9020	38.0617	18.2782	18.2166	18.0306

Table 4. Run times (in seconds) for overlapping algorithm,  $P = 4$ .

Block Size	%	$P = 4$		
	edges	4%	16%	64%
$32 \times 32$	5	0.3704	0.3689	0.3708
	10	0.6424	0.6425	0.6925
	20	1.1862	1.1851	1.1857
$64 \times 64$	5	1.3030	1.3011	1.3009
	10	2.3686	2.3680	2.3614
	20	4.4927	4.4967	4.4956
$128 \times 128$	5	4.7045	4.6999	4.7019
	10	8.9787	8.9760	8.9835
	20	17.5374	17.5426	17.5304

Table 5. Run times (in seconds) for overlapping algorithms,  $P = 16$  and  $P = 64$ .

Block Size	%	$P = 16$			$P = 64$		
		Edges	4%	16%	64%	4%	16%
$32 \times 32$	5	0.4081	0.4152	0.4094	0.4578	0.4619	0.4597
	10	0.6843	0.6806	0.6810	0.7328	0.7349	0.7303
	20	1.2211	1.2263	1.2275	1.2705	1.2752	1.2743
$64 \times 64$	5	1.3675	1.3682	1.3685	1.4292	1.4214	1.4173
	10	2.4359	2.4324	2.4307	2.4802	2.4848	2.4860
	20	4.5572	4.5603	4.5570	4.6034	4.6039	4.6100
$128 \times 128$	5	4.8167	4.8148	4.8151	4.8761	4.8691	4.8718
	10	9.0806	9.0850	9.0970	9.1464	9.1467	9.1344
	20	17.6483	17.6388	17.6390	17.6879	17.6730	17.6907

Table 6. Run times for one processor.

Image Size	% Edges	Time in Seconds
$16 \times 16$	5	0.3005
	10	0.5636
	20	1.1016
$32 \times 32$	5	1.1597
	10	2.2209
	20	4.3527
$64 \times 64$	5	4.4399
	10	8.7194
	20	17.2660

Speedup and efficiency are common measures of the quality of a parallel algorithm. Speedup is defined as

$$S_p = \frac{\text{run time}}{\text{time taken by a uniprocessor}},$$

while efficiency,  $E_p$ , is defined as

$$E_p = \frac{S_p}{P}.$$

Table 7 gives the speedup and efficiency figure achieved by our nonoverlapping algorithm with load balancing for the following cases: variance = 64%, % edges = 20, and  $N = 64$  and 128.

## Conclusions

Consider the binary mapping of an image onto a SIMD hypercube. It can be shown that there are at least two pixels,  $(0, 0)$  and  $(0, N - 1)(N = 2^p)$ , which can potentially



Table 7. Speedup and efficiency of nonoverlapping algorithm. No overlap between communication/computation; variance of edges = 64%.

Edges	No. of Nodes	Image = $64 \times 64$			Image = $128 \times 128$		
		Time	Speedup	Efficiency	Time	Est. Speedup	Est. Efficiency
5	1	3.8603	1.0000				
	4	0.9787	3.9440	0.9860	0.3964	3.9440	0.9860
	16	0.2844	13.5728	0.8483	1.0734	14.5640	0.9103
	64	0.1551	24.8754	0.3886	6.3414	45.7945	0.7155
10	1	7.6151	1.0000				
	4	1.91169	3.9724	0.9931	8.3301	3.9724	0.9931
	16	0.5263	14.4682	0.9043	2.2187	14.9288	0.9331
	64	0.2046	37.2515	0.5821	0.6278	52.7058	0.8234
20	1	15.6470	1.0000				
	4	3.9246	3.9868	0.9967	17.0167	3.9868	0.9967
	16	1.0529	14.8604	0.9287	4.4732	15.1660	0.9478
	64	0.3374	46.3741	0.7246	1.1777	56.6400	0.8850

contribute to the Hough array value (0, 0). The distance between (0, 0) and (0,  $N - 1$ ) is  $O(\log N)$ . Thus any algorithm will require at least  $O(\log N)$  unit routes to complete. Since the time complexity of the serial algorithm is  $\Theta(N^2p)$ , any algorithm will require  $\Omega(p + \log N)$  time to complete on an  $N^2$  node SIMD hypercube. Hence, our algorithm, assuming the binary mapping, for the SIMD hypercube is optimal up to a constant factor. By a similar argument it can be shown that our MIMD algorithm, assuming the gray code mapping, is also optimal up to a constant factor.

We have also shown that our algorithms for the medium-grained hypercube exhibit near-optimal speedups when load balancing is done.

### Acknowledgment

The authors are indebted to Elaine Weinman for converting the original file of this text into L<sup>A</sup>T<sub>E</sub>X and for making subsequent changes as they were required.

### References

- Ballard, D.H. 1981. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition*, 12, 2: 111-122.
- Ballard, D.H., and Brown, C.M. 1982. *Computer Vision*. Prentice Hall, Englewood Cliffs, N.J.
- Chan, T.F., and Saad, Y. 1986. Multigrid algorithms on the hypercube multiprocessor. *IEEE Trans. Comps.*, C-35 (Nov.), 969-977.
- Chandran, S., and Davis, L. 1987. The Hough transform on the Butterfly and the NCUBE. Univ. of Md. tech. rept., College Park, Md.
- Cypher, R.E., Sanz, J.L.C., and Snyder, L. 1987. The Hough transform has  $O(N)$  complexity on SIMD  $N \times N$  mesh array architectures. In *Proc., IEEE Workshop on Comp. Arch. for Pattern Analysis and Machine Intelligence*.
- Dekel, E., Nassimi, D., and Sahni, S. 1981. Parallel matrix and graph algorithms. *SIAM J. on Computing*, 10, 4 (Nov.), 657-675.
- Fishburn, A., and Highnam, P. 1987. Computing the Hough transform on a scan line array processor. In *Proc., IEEE Workshop on Comp. Arch. for Pattern Analysis and Machine Intelligence*, pp. 83-87.

- Guerra, C., and Hambrusch, S. 1987. Parallel algorithms for line detection on a mesh. In *Proc., IEEE Workshop on Comp. Arch. for Pattern Analysis and Machine Intelligence*, pp. 99-106.
- Horowitz, E., and Sahni, S. 1985. *Fundamentals of Data Structures in Pascal*. Computer Science Press.
- Ibrahim, H., Kender, J., and Shaw, D.E. 1986. On the application of massively parallel SIMD tree machines to certain intermediate level vision tasks. *Comp. Vision, Graphics, and Image Processing*, 36: 53-75.
- Nassimi, D., and Sahni, S. 1981. Data broadcasting in SIMD computers. *IEEE Trans. Comps.*, C-30, 2 (Feb.), 101-107.
- Prasanna Kumar, V.K., and Krishnan, V. 1987. Efficient image template matching on SIMD hypercube machines. In *Proc., 1987 Internat. Conf. on Parallel Processing* (Chicago, Aug.), pp. 765-771.
- Ranka, S., and Sahni, S. 1990. Parallel algorithms for image template matching. In *Parallel Algorithms for Machine Intelligence* (to appear).
- Rosenfeld, A., and Kak, A.C. 1982. *Digital Picture Processing*. Academic Press, New York.
- Rosenfeld, A., Ornelas, J., and Hung, Y. 1988. Hough transform algorithms for mesh-connected SIMD parallel processors. *Comp. Vision, Graphics, and Image Processing*, 41, 3: 293-305.
- Thompson, C.D., and Kung, H.T. 1977. Sorting on a mesh-connected parallel computer. *CACM*, 20, 4: 263-271.