August 1985
Revised March 1987

LIDS-P-1490

FINDING PARITY IN A SIMPLE BROADCAST NETWORK*

by

Robert G. Gallager**

ABSTRACT

Consider a broadcast network of N nodes in which each binary digit transmitted by each node is received by each other node via a binary symmetric channel of given transition probability. The errors on these channels are independent over transmitters, receivers, and time. Each node has a binary state and the problem is to construct a distributed algorithm to find the parity of the set of states with some given reliability. It is shown that this can be done with O(ln ln N) bits of communication from each node. Communicating all the node states to one node can be accomplished with only marginally more communication.

*This research was conducted at the M.I.T. Laboratory for Information and Decision Systems with partial support provided by the National Science Foundation under Grant NSF-ECS-8310698 and the the Defense Advanced Research Projects Agency under Contract ONR/N00014-84-K-0357.

**Room No. 35-206, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139.

FINDING PARITY IN A SIMPLE BROADCAST NETWORK

by Robert G. Gallager

1) INTRODUCTION

Consider a broadcast network of N+1 nodes in which each binary digit transmitted by each node is received by each other node via a binary symmetric channel of crossover probability $\varepsilon < 1/2$. The errors on these channels are independent over transmitters, receivers, and time. Each node has a binary state, and the problem under consideration is for one special node, called the receiver, to determine the parity of the set of node states. In particular, we want to minimize the number of binary digits that must be sent by each node in order for the receiver to determine parity within some allowable error probability, P*. We assume that whatever algorithm is used for transmitting the required information, all nodes know the algorithm and there is no contention between transmissions; thus each binary digit transmitted is received and identified by all other nodes, subject to the noise introduced by the binary symmetric channels.

The above problem was first formulated by A. El Gamal [1], and is of interest because it is one of the simplest distributed algorithm problems involving noise. A closely related problem that we treat is for the receiver to find the state of each other node. Note that the conventional coding theorems of information theory cannot be used here because each node has only one bit of information to communicate. This situation of communicating a limited amount of information is common in network protocols and more generally in the control of distributed systems. The particular character of the problem here comes from the independence of the noise at each receiver for a given transmission. This means that when a node transmits a single digit, the other nodes collectively could make a good decision on that digit since they have N samples of it with independent noise; unfortunately, the nodes cannot act collectively without using up their own valuable transmissions, which are also noisy.

The straightforward approach to this problem is for each node to broadcast its own state j times for some integer j. The receiver will make an error in decoding a given node's state with a probability ε_j closely upper bounded [2] by

$$\varepsilon_j \leq \alpha^{-j} \quad \text{where } \alpha = [4\varepsilon(1-\varepsilon)]^{-1/2} \qquad (1)$$

The probability P that the receiver will make an error in calculating the parity of the states is then upper bounded by $N\alpha^{-j}$. Since this bound is quite tight for $N\alpha^{-j}$ small, we see that j must grow as ln(N) in this approach for a constant P.

2) FINDING PARITY WITH O[ln(ln(N))] BITS PER NODE

The approach we take here for more efficient communication is to partition the nodes (other than the receiver) into subsets each with approximately the same number of nodes. In particular, it is always possible to partition N nodes into subsets of k or k-1 nodes each for any k satisfying $(k-1)^2 \leq N$. We shall see later that the parameter k is appropriately chosen to be proportional to ln(N). Each node again broadcasts its own state j times (where now j will be chosen proportional to ln(ln N)), but then makes a decision on the state of each of the other nodes in its subset using the j receptions from that node. The node adds these decisions, plus its own state, modulo 2, to estimate parity for its own subset, and then broadcasts this estimated parity exactly once. The receiver will then receive k or k-1 different estimates for the parity of each subset, and, as we shall see, this allows the receiver to obtain a highly reliable decision on a subset's parity. Given the parity of each subset, the parity of the entire set is found by addition modulo 2 of the individual subset parities.

Note that the parity estimate that the receiver obtains from a given node can be wrong either because of noise in the transmission of the parity or because of decision errors at the sending node. Let the random variable I be the ordinary sum of the decision errors at a given sending node plus an additional 1 if an error is made in the receiver's reception of that node's transmission of parity. The received estimate from that node is incorrect if I is odd. In order to find the probability that I is odd, let $f(z) = E\{z^I\}$ be the z transform of I. Assuming for the moment that the given subset contains k nodes, and noting that the k-1 potential decision errors and the one potential transmission error are independent, we see that

$$f(z) = [(1-\varepsilon_j) + \varepsilon_j z]^{k-1}[(1-\varepsilon) + \varepsilon z].$$

It can now be seen that the probability β that I is odd is given by $[f(1)-f(-1)]/2$. Thus

$$\beta = [1 - (1-2\varepsilon_j)^{k-1}(1-2\varepsilon)]/2. \qquad (2)$$

For a subset with k-1 nodes, the exponent k-1 above is replaced with k-2, which yields a smaller probability than β as given above. Similarly, the receiver can form its own preliminary estimate of parity for the given subset from the original j transmissions from each node, and the probability of error in this preliminary estimate is also upper bounded by β.

Finally, the receiver decides on the parity of a subset by taking a majority vote among the received estimates and its own preliminary estimate. Since the errors in these estimates are independent, the probability that half or more of the estimates are erroneous is upper bounded (using the same argument as in Eq. (1)) by

$$P_{\text{subset}} \leq [4\beta(1-\beta)]^{k/2}. \qquad (3)$$

Combining (2) and (3),

$$P_{subset} \leq [1 - (1-2\varepsilon_j)^{2(k-1)}(1-2\varepsilon)^2]^{k/2} \tag{4}$$

The receiver next adds the parities of all the subsets (of which there are at most N) and adds its own state, all modulo 2. The probability P that this decision on parity for the entire set is incorrect is upper bounded by the probability of an error on one or more of the subsets, which is further upper bounded by $NP_{subset.}$ Thus

$$P \leq N[1 - (1-2\alpha^{-j})^{2(k-1)}(1-2\varepsilon)^2]^{k/2} \tag{5}$$

In going from (4) to (5), we have used (1) to upper bound ε_j by α^{-j}. The right hand side of (5) first decreases and then increases with increasing k, and the minimizing integer k is a complicated function of ε and j. It is sufficient for our purposes, however, to simply restrict k to be small enough (or j large enough) to satisfy

$$(1-2\alpha^{-j})^{2(k-1)} \geq 1/4 \tag{6}$$

With this restriction,

$$P \leq N\, Z^{k/2} \qquad \text{where } Z = 1 - (1-2\varepsilon)^2/4 \tag{7}$$

This is essentially the solution we are looking for; we choose k large enough to make the error probability sufficiently small in (7) and then choose j large enough to satisfy (6). As N is varied, we see that k must increase logarithmically with N, and then, as will be seen, j increases logarithmically with k. Since j and k must be integers, however, a little fussing is required to get a valid bound on j. We start by defining real number approximations, k* and j* , to k and j for a given number of nodes N and a given requirement P* on error probability

$$k^* = \frac{2\ln(N/P^*)}{\ln(Z^{-1})} \;;\quad j^* = \frac{\ln(2) - \ln(1-2^{-1/k^*})}{\ln \alpha} \qquad (8)$$

$$k = \lceil k^* \rceil \;; \qquad j = \lceil j^* \rceil$$

where $\lceil x \rceil$ is the smallest integer greater than or equal to x. We now show that with k and j chosen according to (8), the resulting error probability will be at most P*. Note that the equation for j* in (8) can be rearranged to

$$(1-2\alpha^{-j^*})^{2k^*} = 1/4 \qquad (9)$$

Since $k^* > k-1$ and $j^* \leq j$, (6) must be satisfied. This means that (7) must be satisfied, yielding

$$P \leq N Z^{k/2} \leq N Z^{k^*/2} = P^* \qquad (10)$$

The second inequality above is valid because $k^* \leq k$, and the equality is a rearrangement of the definition of k*. One final simplification will now be useful in obtaining our final bound. Assume that k*, as given by (8), satisfies $k^* \geq 1$. Then it is not hard to verify that $1 - 2^{-1/k^*} \geq 1/(2k^*)$. Substituting this into the definition of j*, we obtain

$$j^* \leq \ln(4k^*) / \ln(\alpha) \qquad (11)$$

The number of digits transmitted per node is m = j+1, which is at most j* +2. Substituting this into (11) and using (8) for k*, we obtain our final bound,

$$m \le \frac{\ln[\ln(N/P^*)] + A}{\ln(\alpha)} + 2 \qquad \text{where} \tag{12}$$

$$A = \ln(8) - \ln[\ln(1/Z)] \tag{13}$$

Recall that we have imposed two restrictions on k in deriving this result. First $N \ge (k-1)^2$, which is satisfied if $N \ge k^{*2}$, and second $k^* \ge 1$. From the definition of k*, these restrictions are

$$N \ge \left[\frac{2\ln(N/P^*)}{\ln(1/Z)}\right]^2 \; ; \quad N \ge \frac{P^*}{\sqrt{Z}} \tag{14}$$

The second restriction is always satisfied for N ≥ 2, but the first restriction is more substantive. Note first that (14) is always satisfied for large enough N given any P* and ε, and thus (12) shows that asymptotically, m increases at most as ln(ln(N)) On the other hand, for given N and ε, (14) is always violated for small enough P*. Thus (12) (subject to (14)) does not show that m asymptotically varies with P* as ln(ln(1/P*)). This is reassuring, since even if all nodes other than the receiver knew the parity of the states, the error probability could not decrease faster than ε^{-Nm}. Actually, by changing the strategy somewhat and making all subsets of size k except for one subset of size between 1 and k, the restriction $N \ge (k-1)^2$ could be relaxed to N ≥ k. In this case, some of the nodes would have to transmit an extra digit to help resolve the parity of the small subset, and the bound on m would be somewhat weakened. We omit the details of this since it is tedious and doesn't improve the asymptotic behavior with N.

It is also possible to reduce the value of A for large values of ε by having each node transmit an estimate of parity for several subsets rather than just its own subset. Again we omit the analysis since it is tedious and does not materially improve the result.

It should be noted that if the receiver is placed in one of the subsets, then each node of the network can estimate parity as well as the receiver. This makes one wonder whether the N+1 nodes we have been considering could be considered as a subset in yet a larger set of nodes, with an extra transmission of parity from each node in each such subset serving to find parity for the larger set. The trouble with such a scheme is that the parity estimates in the larger subsets are not independent. It is conjectured that such a multi-tier subset scheme does not improve the asymptotic behavior over the scheme just described.

3) FINDING THE STATE OF ALL NODES

In this section we show that the receiver can determine the state of all nodes with very few more transmissions per node than are required to determine parity. Our strategy in doing this is to form a set of N subsets of the N nodes (not counting the receiver) in such a way that each subset contains k-1 or k nodes for some k and each node is contained in k-1 or k subsets. Furthermore, we constrain the choice of subsets so that no pair of nodes appear together in more than one subset (see Figure 1). In the appendix, we show that such a set of subsets can always be constructed if

$$N \geq 2k(k-1)^2 \qquad (15)$$

We next associate each node with one subset in a one to one fashion so that each node is associated with a subset containing it and each subset has one of its contained nodes associated with it. The appendix also shows how this association can be constructed. Each node then sends its own state j times and then each node estimates the parity of its associated subset in the same way as before. Finally each node sends the parity of its associated subset and the receiver uses this information, plus its own receptions of the node states, to determine the state of each node. Thus each node sends m = j+1 binary digits as before.

$$\begin{matrix}
1&1&1&0&0&0&0\\
1&0&0&1&1&0&0\\
1&0&0&0&0&1&1\\
0&1&0&1&0&1&0\\
0&1&0&0&1&0&1\\
0&0&1&1&0&0&1\\
0&0&1&0&1&1&0
\end{matrix}$$

Example of N=7 nodes and subsets with each subset containing k=3 nodes. Each row corresponds to a subset and the 1's in the row indicate the nodes within the subset. Note that no pair of subsets contain two nodes in common.

Figure 1

Now consider how the receiver can decode the state of each node from the received information. First the receiver makes a preliminary decision on the state of each digit from the j noisy receptions of that digit. The probability of error for each of these preliminary decisions is $\varepsilon_j \leq \alpha^{-j}$ as before. In order to make a final decoding of the state of a given node, say node i, the receiver considers each of the subsets, say $S_{i,1}$, $S_{i,2}$, ..., $S_{i,k}$, or $S_{i,1}$, ..., $S_{i,k-1}$ that contain node i. For a given subset, say $S_{i,\gamma}$, the receiver adds, modulo 2, its preliminary decisions on the nodes in $S_{i,\gamma}$ - {i} to the received parity estimate of $S_{i,\gamma}$. Note that if the receiver's preliminary decisions on these nodes are all correct, and if the node transmitting the parity of $S_{i,\gamma}$ has made correct decisions on the nodes in the subset, and if the transmission of that parity is correct, then this modulo 2 sum is simply the state of node i since all other node states are added twice. Thus this sum is in some sense an estimate of the state of node i. More particularly, the probability that this estimate is incorrect is the probability of an odd number of errors in the transmission of parity, in the receiver's preliminary decisions on the nodes in $S_{i,\gamma}$, and in the transmitting node's decisions on the nodes of the subset other than itself. If the subset contains k nodes, then we are looking at k-

1 decisions at the receiver, k-1 at the transmitting node, and one transmission of parity. The probability that this estimate is incorrect is then

$$\beta' = [1 - (1-2\varepsilon_j)^{2(k-1)}(1-2\varepsilon)]/2 \tag{16}$$

If the subset contains k-1 nodes, then k-1 would be replaced by k-2 in the above equation, so that β' is an upper bound on the probability of an incorrect estimate in that case. Finally, β' also upper bounds the probability of an error in the receiver's preliminary estimate of the state of node i. The receiver now has its preliminary decision of node i's state plus either k or k-1 estimates from the different subsets containing i. Since no two subsets contain more than one node in common, $S_{i,1}-\{i\}$, $S_{i,2}-\{i\}$, ... are all disjoint and all of these estimates are based on mutually independent errors. Thus when the receiver takes a majority vote on its k or k+1 estimates, the probability of error in this final decision on node i is upper bounded by

$$P_i \leq [4\beta'(1-\beta')]^{k/2} \tag{17}$$

Combining (16) and (17) and upper bounding ε_j with α^{-j}, we get

$$P_i \leq [1 - (1-2\alpha^{-j})^{4(k-1)}(1-2\varepsilon)^2]^{k/2} \tag{18}$$

The probability that any node state will be decoded incorrectly is now upper bounded by $P \leq N\,P_i$. We now restrict k to be small enough that

$$(1-2\alpha^{-j})^{4(k-1)} \geq 1/4 \quad \text{yielding} \tag{19}$$

$$P \leq N\,Z^{k/2} \qquad \text{where} \qquad Z = 1 - (1-2\varepsilon)^2/4 \tag{20}$$

As in the analysis of finding the parity of the states, we now consider a required error probability, P*, and choose k and j to meet the requirement.

$$k^* = \frac{2\ln(N/P^*)}{\ln(1/Z)} \;; \qquad j^* = \frac{\ln(2) - \ln[1-2^{-1/(2k^*)}]}{\ln(\alpha)} \tag{21}$$

$$k = \lceil k^* \rceil \;; \qquad j = \lceil j^* \rceil$$

As before, this guarantees that (19) is satisfied and that the error probability is at most P*. If $2k^* \geq 1$, we can upper bound j* by $\ln(8k^*)/\ln(\alpha)$. Since the number of digits per node is at most j*+2, we can use this bound on j* with the definition of k* to obtain

$$m \leq \frac{\ln[\ln(n/P^*)] + A'}{\ln(\alpha)} + 2 \qquad \text{where} \tag{22}$$

$$A' = \ln(\frac{16}{\ln(1/Z)}) \tag{23}$$

Note that the required number of transmissions here exceeds that for finding parity only by $\ln(2)/\ln(\alpha)$. Recall that the construction here required $N \geq 2k(k-1)^2$, which is valid for $N \geq 2k^{*2}(k^*+1)$. From (21), we see that this is satisfied, for any given P* and ε for all sufficiently large N. As before, for fixed N and ε, the restriction is always violated for small enough P*.

APPENDIX

We start with a set of N nodes and want to construct N subsets of k or k-1 nodes each, with the properties that each node is in k or k-1 subsets and no two subsets contain any pair of nodes in common. Assume that $N \geq 2k(k-1)^2$. Let $L = \lceil N/k \rceil$ and let N' = Lk. We first construct N' subsets of k nodes each from a set of N' nodes and then delete N'-N nodes and subsets. The procedure is a special case of an earlier procedure given in [3].

Consider the problem in terms of an N' by N' matrix of 0's and 1's. The columns correspond to nodes and the rows to subsets. The 1's in a row correspond to the nodes in the corresponding subset, so our problem is to construct an N' by N' matrix in which each row and each column contains k 1's and, for any two rows, there is at most one column for which both rows contain 1's. We construct such a matrix by first considering the N' by N' matrix as partitioned into k^2 permutation matrices, each L by L. The fact that each L by L matrix is a permutation matrix guarantees that each row and each column of the entire matrix contains k ones. The figure below shows such a partitioned matrix where L = 3 and k = 3; note that no two rows contain more than a single 1 in common.

$$\left[\begin{array}{ccc|ccc|ccc}
1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\
0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\
0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\
\hline
1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\
0 & 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\
0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 \\
\hline
1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\
0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0
\end{array}\right]$$

Figure 2

Next consider the construction of the L by L permutation submatrices. We choose the submatrix in the upper left corner to be an identity matrix and suppose that after constructing an arbitrary number of the submatrices, we are now constructing a new submatrix. At most k-1 submatrices have already been constructed using the same rows as the new submatrix. Consider a particular row of the new submatrix. There are at most k-1 1's already in that row (one for each previously constructed submatrix using the same rows). A column that contains one of these 1's also contains at most k-1 additional 1's (one for each submatrix using the same column). This means that there are at most $(k-1)^2$ rows that have 1's in the same column as the given row. Each of these rows contains at most a single 1 in one of the columns of the new submatrix. The given row of the new submatrix is forbidden to contain 1's in any of these columns because of the constraint that no two rows can contain more than a single column with 1's in each row. Summarizing, then, there are at most $(k-1)^2$ positions in which any given row of the new submatrix is forbidden to have 1's. Using the same argument on columns, we see that there are at most $(k-1)^2$ row positions in which any given column is forbidden to have 1's. The constraint $N \geq 2k(k-1)^2$ implies that

$L \geq 2(k-1)^2$, and thus each row and each column has at least L/2 positions available for 1's without violating the condition that no two rows have two 1's in common.

Suppose, by using the above procedure, that we mark each position in the new submatrix as available or forbidden for the placement of 1's. We now show how to construct a permutation matrix, placing 1's only in available positions and using the above fact that at least half the positions in each row and each column are available. The procedure to be described is a special case both of [3] and [4]. The procedure is first to place a 1 in an arbitrary available position of the first row, and then to successively attempt to place a 1 in each succeeding row. For each such row, a 1 is placed in an arbitrary position that is both available and not used by any of the previous rows in the submatrix construction. If such a position is found in each of the L rows of the submatrix, then a suitable permutation matrix has been constructed. On the other hand, it is possible that for some given row, each of the available positions has already been used by a previous row. We now show how to change one of the previous rows so that a 1 can be placed in the given row under these circumstances. Consider a given column that contains no 1 within the submatrix. This column position is forbidden in the given row (since no 1 can be placed there), and thus at most L/2 - 1 other entries in the column are forbidden. There are at least L/2 positions marked available in the given row, and each of the corresponding columns contains a 1 in some previously constructed row of the submatrix. Each of these L/2 or more rows intersect the given column and therefore at least one such intersection is at a position marked available. Thus we have the situation shown in Figure 3 in which the given row and given column each have an available position with the property that the intersection between the row position and the column position contains a 1. By moving this 1 to the given column on the same row and then placing a 1 in the available position in the row being constructed, we complete the construction of the new row as desired.

The constraint L ≥ $(k-1)^2$ that we have imposed appears to be stronger than necessary in many cases. For example, if L is prime and L = k, it is possible to construct a matrix with the required properties as follows: let P_i denote the cyclic L by L permutation matrix that is shifted i places from the identity matrix. Then for the m^{th} row and n^{th} column submatrix, use P_i where i = (m-1)(n-1).

The above paragraphs show how to construct an N' by N' matrix where N' = k L ≥ N. By removing the first N'-N rows and columns, we have an N by N matrix. Since N'-N < L, we have removed at most a single 1 from each row and column, so the remaining matrix satisfies the required properties and has either k or k-1 1's in each row and column. Finally, to associate each node with a distinct subset that contains that node, we consider the main diagonal of submatrices in the above construction. The association of one node per subset is then given by the 1's in those submatrices.

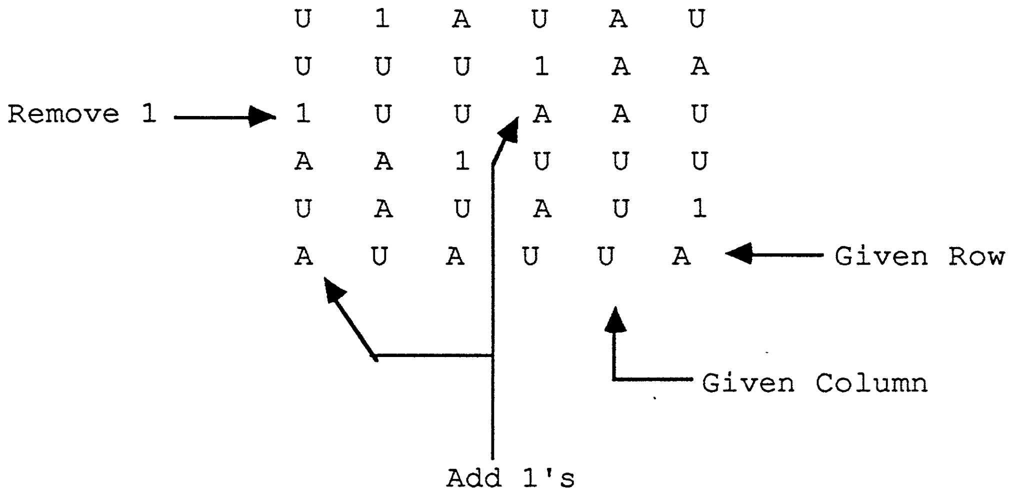

U	1	A	U	A	U
U	U	U	1	A	A
1	U	U	A	A	U
A	A	1	U	U	U
U	A	U	A	U	1
A	U	A	U	U	A

Construction of a submatrix row by changing a previous row.

Figure 3

REFERENCES

1) El Gamal, A., Open problem presented in the 1984 Workshop on Specific Problems in Communication and Computation sponsored by Bell Communication Research.

2) Gallager, R. G., *Information Theory and Reliable Communication,* Section 5.3, John Wiley, N.Y., 1968.

3) Gallager, R. G., *Low Density Parity Check Codes,* Appendix C., M.I.T. Press, Cambridge,Ma. 1963.

4) Hall, M. *Combinatorial Theory,* Page 48, Blaisdell Publishing, Waltham, Ma. 1967.