Noisy Broadcast and Distributed Computation

Final Project Report
ECE 255C

Shouvik Ganguly
Electrical and Computer Engineering
University of California, San Diego
La Jolla, California 92092
shgangul@eng.ucsd.edu

***Abstract*—In this report, I have looked at certain results on the order of the numbers of transmissions required for reliable communication of distributed information.**

I. Introduction

The noisy broadcast problem was first introduced by El Gamal ([1], 1987). It marked a departure from the classical paradigm of Shannon, in which there are a small number of users, each with asymptotically larger and larger amounts of information, which they wish to exchange reliably over a noisy channel. In this framework, there are a large number of users, each with a small amount of information, and, through noisy communication between the users, one or more users have to compute some function reliably.

Gallager ([2], 1988) solved the problem for the special case of finding the parity in a noisy broadcast network. More general versions of the problem have been studied in [10], [3], and [4]. Several asymptotic results have been proved by [12], [5], and [7].

II. Problem Statement and History

Let us consider an undirected graph G with vertex set V and edge set $\mathcal{E}$, where the vertices denote nodes and edges represent noisy communication links. Each node i is given a quantity $x_i \in \mathcal{X}$, and our target is to find a sequence of steps through which, using the least possible number of transmissions, some function f of all the x_i's can be computed by a pre-designated node j, with error probability bounded by some pre-defined $\delta > 0$.

In an early attempt to answer this question, Gallager ([2], 1988) considered the special cases where the setup is a broadcast network with N nodes (i.e., the graph G is complete), $\mathcal{X} = \{0, 1\}$ and the target functions $f : \mathcal{X}^N \to \mathcal{X}$ are just the overall parity of the bits present at each node, or the states of all the nodes. Also in this case, the link between each pair of nodes is a binary symmetric channel with crossover probability $\epsilon < 1/2$, and the links are independent of transmitters, receivers and time. An asymptotic converse to this particular problem was later proved in [3], thus solving this particular problem up to order in the number of nodes.

Kushilevitz and Mansour [12] looked at the same broadcast network, with the only difference being that the function f was now defined as a thresholding function, i.e., $f(x^N) = 1$ if the number of nodes having 1's is greater than some threshold k, and $f(x^N) = 0$ otherwise. They showed a scheme achieving any desired probability of error δ in $)(N)$ operations, even when the crossover probabilities ϵ of the BSCs are not known at the nodes.

[10], [3] and many others have provided protocols for computing different functions of the bits under the same noisy broadcast scenario.

A new direction was opened with the study of this problem under the framework where the graph G is random, and also not complete. As an example, for the random geomteric network, N nodes are placed uniformly at random in a $\sqrt{N} \times \sqrt{N}$ square, and each node is connected through a BSC(ϵ) to all nodes within a radius r of itself. For such a network, Ying, Srikant and Dullerud [4] showed that any symmetric function of the input bits can be computed with probability of error at most ϵ, using $\mathcal{O}(N \log \log N)$ transmissions. As a converse to this, it was shown in [5] that for computing the parity function over a random geomteric network with $r = \mathcal{O}(N^{1/2-\alpha})$ with probability of error less than δ, we need $\Omega(N \log \log N)$ transmissions.

Karamchandani, Appuswamy and Franceschetti [7] considered the problem of finding a symmetric function of the bits, under the framework of deterministic geometric networks, where N nodes are placed on a $\sqrt{N} \times \sqrt{N}$ grid, and each node is connected to all nodes within a radius r of itself. They proved certain strong results in this framework, namely, that for $r > \Theta\Big(\frac{\log N}{\log\log N}\Big)$, the number of transmissions required for computing any symmetric function is $\Theta(N \log\log N)$, while for $r < \Theta\Big(\frac{\log N}{\log\log N}\Big)$, the number of transmissions required is $\Theta\Big(\frac{N \log N}{r^2}\Big)$.

Under the same network model as above, [8] studied the problem where both the number of transmissions and the number of time slots required were studied. They showed that under the setup where nodes not sharing an edge are allowed to transmit simultaneously, the number of time slots required is $\max\{\Theta\Big(\frac{\sqrt{N}}{r}\Big), \Theta(r^2 \log\log N)\}$.

III. Gallager's Scheme

We partition the N nodes (other than the receiver) into subsets of k or $k-1$ nodes each. This can always be done if $(k-1)^2 \le N$. Each node then transmits n times, and then makes a decision on the parity of each of the other nodes in its subsets using the n transmissions from that node. The node then adds these decisions and its own parity, modulo 2, to estimate the parity of its own subset and transmits this decision exactly once. The receiver thus receives k or $k-1$ estimates for the parity of each subset, and makes a decision for the parity of a subset by majority rule. Finally, adding these decisions modulo 2, the receiver estimates the overall parity.

The parity estimate that the receiver receives from a particular node can be wrong, either if the sender node's decision itself is wrong, or if the transmitted decision gets flipped by the channel. Let Z be the random variable denoting the (real) sum of the decision errors at a given sending node, plus an additional 1 if the decision gets flipped in transmission. The estimate received from the node will be wrong if Z is odd. It can be shown [2] that if the subset has k nodes, then (ϵ being the channel parameter of the BSC)

$$\beta = \Pr[Z \text{ is odd }] = \frac{1}{2}[1 - (1-2\epsilon_n)^{k-1}(1-2\epsilon)], \tag{1}$$

where ϵ_n is the probability of decision error at the sender node for estimating the parity of a particular node by majority decision rule, and satisfies

$$\epsilon_n \le [4\epsilon(1-\epsilon)]^{n/2}. \tag{2}$$

The probability of receiving a wrong estimate for the parity of a subset with $k-1$ nodes is easily seen to be upper-bounded by β. Now, by majority decision rule, the probability of the receiver making a wrong estimate of the parity of a subset from the received estimates from all nodes in the subset is upper-bounded by $[4\beta(1-\beta)]^{(k-1)/2}$. Also, there are less than N subsets, so by the union bound, the probability of error in the overall decision is upper bounded by

$$\begin{aligned} &P_e \le N[4\beta(1-\beta)]^{(k-1)/2} \\ &= N\Big[1-(1-2[4\epsilon(1-\epsilon)]^{n/2})^{2(k-1)}(1-2\epsilon)^2\Big]^{(k-1)/2}. \end{aligned} \tag{3}$$

If we can make

$$(1-2[4\epsilon(1-\epsilon)]^{n/2})^{2(k-1)} \ge \frac{1}{4}, \tag{4}$$

we will have

$$P_e \le N\Big[1 - \frac{(1-2\epsilon)^2}{4}\Big]^{(k-1)/2}. \tag{5}$$

Thus we must make k large enough to bound the error probability in (5), and then choose n large enough to satisfy (4). From (5), we see that k has to be $\mathcal{O}(\log N)$ and from (4), we see that n has to be $\mathcal{O}(\log k)$. Thus we have, the number of transmissions per node, n, has to be $\mathcal{O}(\log\log N)$ and the total number of transmissions is $\mathcal{O}(N \log\log N)$.

It turns out that the parities of *all* nodes individually (and not just the overall parity) can also be estimated reliably using $\mathcal{O}(N \log\log N)$ transmissions. In order to do this, we form a set of N subsets of the N nodes (other than the receiver) in such a way that for some fixed k, each subset contains $k-1$ or k nodes, each node appears in $k-1$ or k subsets, and no two different nodes appear together in more than one subset. It can be shown [2] that this construction is possible if $N \ge 2k(k-1)^2$. We then define a one-to-one mapping between nodes and subsets in such a way, that each subset is associated with one node within each, and each node is associated with one subset containing it.

Each node sends its own parity n times, and then each node estimates the parity of its associated subset by majority decision rule. Finally, each node transmits the parity of its associated subset and the receiver uses these transmissions, along with its own previous estimates of the node parities, to make a final decision on the parity of each node.

The receiver first forms a preliminary estimate of the parities of all nodes from the first n transmissions using majority decision rule. In order to use the final round of transmissions, for each node i, the receiver considers all subsets $S_{i,1}$, $S_{i,2}$, $\ldots$ containing i. For each of these subsets $S_{i,j}$, the receiver adds (modulo 2) its preliminary decisions on the nodes in $S_{i,j} \backslash i$ to the received estimate for $S_{i,j}$ to get an estimate for the parity of i. Since i is in $k-1$ or k subsets, the receiver in this way gets $k-1$ or k estimates for the parity of each node i. These, together with the preliminary decision, can be used to estimate the parity of node i using majority decision rule. It can be shown [2] that the probability of error can be bounded in this case, if $k = \mathcal{O}(\log N)$ and $n = \mathcal{O}(\log k)$. Thus, we again have $\mathcal{O}(N \log \log N)$ total transmissions.

IV. Optimality of Gallager's Scheme for Identity

Goyal et. al. [3] proved the following converse for the noisy broadcast problem.

Theorem 1 : $\forall\ \epsilon \in (0, 1/2)$, $\exists\ n_0(\epsilon)$ such that $\forall N \geq n_0$, any protocol for computing identity in the noisy broadcast problem with i.i.d. BSC(ϵ) links requires at least $\dfrac{N \log \log N}{20 \log(1/\epsilon)}$ transmissions. ∎

Thus we see that the scheme for computing identity described in the previous section is order-optimal. However, in the same paper [3], the authors showed a $\mathcal{O}(N)$ protocol for computing *any* symmetric function (including overall parity) of the inputs in the noisy broadcast problem. Thus, the scheme for parity described in the previous section is sub-optimal in order.

V. Beyond Noisy Broadcast : General Networks

The distributed computation problem has been extensively studied over different networks for the past 15 years. In the first significant result outside the domain of broadcast, Ying, Srikant and Dullerud [4] showed that for a random geometric network with $r(N) = \Omega(\sqrt{\log N})$, any symmetric function of the (not necessarily binary) inputs can be computed in $\mathcal{O}(N \log \log N)$ transmissions. Their scheme is similar in principle to Gallager's, in that it divides the network into subsets (or 'cells') in such a fashion that all nodes within a cell are connected to each other, and the decision from each cell is transmitted to the receiver via other, 'closer' cells (in a tree-like fashion).

A partial converse to this result was soon proved, as mentioned in the Section on 'History'. Subsequent years saw the study of this problem under the framework of 'deterministic geometric networks' as well, and finally, in addition to optimizing the number of transmissions required, the latency (number of time-slots required) was also optimized.

References

[1] Abbas El-Gamal. *Reliable communication of highly distributed information. Open Problems in Communication and Computation*, T.M. Cover and B. Gopinath, eds., Springer-Verlag, New York, 1987, pp. 60–62.

[2] R. Gallager. *Finding parity in simple broadcast networks*. IEEE Transactions on Information Theory, 34(2), pp 176–180, Mar 1988.

[3] N. Goyal, G. Kindler and M. Saks. *Lower bounds for the noisy broadcast problem*. SIAM Journal of Computing, 37(6), pp. 1806–1841, March 2008.

[4] L. Ying, R. Srikant and G.E. Dullerud. *Distributed symmetric function computation in noisy wireless sensor networks*. IEEE Transactions on Information Theory, 53(12), pp 4826–4833, Dec 2007.

[5]] C. Dutta, Y. Kanoria, D. Manjunath and J. Radhakrishnan. *A tight lower bound for parity in noisy communication networks. Proc. 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp 1056–1065, 2008.

[6] Y. Kanoria and D. Manjunath. *On distributed computation in noisy random planar networks. Proc. International Symposium on Information Theory*, pp 626–630, 2007.

[7] N. Karamchandani, R. Appuswamy and M. Franceschetti. *Distributed computation of symmetric functions with binary inputs. Proc. IEEE Information Theory Workshop*, pp 76–80, 2009.

[8] N. Karamchandani and M. Franceschetti. *Distributed Function Computation in Networks: a Joint Delay-Energy Perspective Proc. 20th International Symposium on Optimization of Mobile, Ad Hoc and Wireless Networks*, 2011.

[9] A. Giridhar and P.R. Kumar. *Computing and communicating functions over sensor networks*. IEEE Journal on Selected Areas in Communication, 23(4), pp 755–764, Apr 2005.

[10] I. Newman. *Computing in fault tolerance broadcast networks. Proc. IEEE Annual Conference on Computational Complexity*, pp 112–122, 2004.

[11] D. Slepian and J.K. Wolf. *A coding theorem for multiple access channels with correlated sources*. Bell System Technical Journal, 52(7): 1037–1076, 1973.

[12] E. Kushilevitz and Y. Mansour. *Computation in noisy radio networks. Proc. 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 236–243, 1998.

[13] A. Yao. *On the complexity of communication under noise. Proc. 5th Israeli Symposium on Theory of Computing and Systems*, 1997.