A COMMUNICATION ARCHITECTURE FOR REACHING CONSENSUS IN DECISION FOR A LARGE NETWORK

Yao-Win Hong, Anna Scaglione[*]
School of Electrical and Computer Eng.
Cornell University, Ithaca, NY 14853
yh84@cornell.edu; anna@ece.cornell.edu

Pramod K. Varshney
Electrical Eng. & Computer Science Dept.
Syracuse University, Syracuse, NY 13244
varshney@syr.edu

ABSTRACT

One of the most challenging aspects in applying decentralized detection in sensor networks is the efficient exchange of small messages required for data fusion. In this work, we propose a novel communication architecture for a canonical decentralized detection problem where the sensor nodes exchange continuously their local decisions until consensus is reached among all nodes. Our methodology capitalizes on the observation that the information embedded in the exchanged messages decreases to zero as the decisions gradually converge. By using a data-driven multiple access scheme, we show that the number of channel accesses required for each round of message exchange decreases, following the same trend as the aggregate entropy of the sensor decisions. The main contribution is to show that data-driven multiple access strategies can overcome the backlog of communications that many distributed computing algorithms generate in a wireless network setting.

1. INTRODUCTION

Sensor networks typically consist of small sensor devices that are endowed with the capabilities to sense, to compute and to communicate. As opposed to the internet or cellular networks, the sensors often have a common goal, *e.g.* to sense the common environment, to coordinate a common response or to compute a common function with the distributed inputs. For example, to track a vehicle in the network, sensors must exchange their local observations in order to compute the position, the velocity or the direction of the vehicle. Without intercommunication, the sensor network is simply a bunch of "dumb" sensors.

Nonetheless, the capacity limits of the wireless medium [1] have restricted many sensor network applications, especially ones that require a vast amount of communication for distributed computing at each node. In the past, many distributed computing algorithms [2, 3] have been proposed that achieve the computational goals of the sensor network. However, most of these works do not investigate the necessary reduction of the communication complexity, which can be achieved by exploiting the inevitable redundancy in the data exchange.

In the past, distributed source coding (DSC) [4] has been proposed by many authors to solve the sensor communication problem. However, the compression gain achieved by DSC is contingent on the assumption that long blocks of data are available for encoding at each node. Therefore, this method is not applicable to distributed "in network" computations since the messages exchanged at each stage are not only short, *e.g.* one bit per sensor in the binary hypothesis testing problem [see Section 2], but also depend on the input of the previous messages.

The main contribution of this paper is to derive a data-driven communication architecture that extracts the information relevant for computation using the minimum number of channel accesses. The proposed strategy is based on the concept of the Group Testing Multiple Access scheme (GTMA) [5, 6], where the stochastic knowledge of the exchanged data is utilized to facilitate the scheduling of transmissions among the distributed sensors.

In this paper, we discuss specifically the class of decentralized detection problems [2,3] where all the nodes continuously update and exchange their local hard decisions until consensus in decision is reached among all sensors. In this scenario, the statistics of the observations at each sensor are governed by a common phenomenon or hypothesis, therefore, it is reasonable to expect that the messages relevant for computing the network-wide decision should contain low aggregate entropy and are, thus, highly compressible.

2. PROBLEM SETUP

Consider the binary hypothesis testing problem for a network of N sensors $\mathcal{S} = \{s_1, \cdots, s_N\}$ and let $\mathbf{X} = [X_1, \cdots, X_N]$ be the set of observations where the random variable X_i represents the observation made by sensor s_i, and that the observations are mutually independent conditioned on the hypothesis. Initially, each sensor makes a local binary

[*]This work is supported in part by the National Science Foundation under grant CCR-0431077.

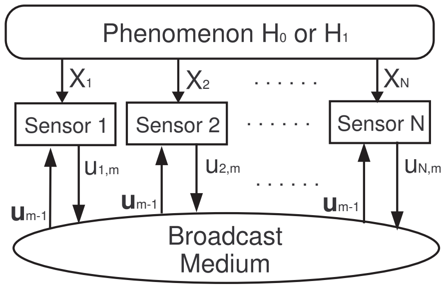

Fig. 1. Illustration of a network of agents reaching a consensus in decision through iterative information exchange.

decision for the hypothesis testing problem:

$$\begin{array}{ll} \mathcal{H}_0: & X_i \sim f_{0,i} \\ \mathcal{H}_1: & X_i \sim f_{1,i} \end{array} \quad (1)$$

for $i = 1, \cdots, N$, where $f_{0,i}$ and $f_{1,i}$ are the probability density functions of X_i given the hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$, respectively. As shown in Fig. 1, each sensor broadcasts its local hard decision to all the other sensors in the network and updates its own decision based on the new information provided by the other sensors. We note that the statistics of the observations at each sensor need not be identical.

Let $u_{i,m}$ be the decision made by sensor s_i after $m-1$ iterations of information exchange, *i.e.* the m-th local decision made at sensor s_i, and let $\mathbf{u}_m = [u_{1,m}, \cdots, u_{N,m}]$. The decisions can be expressed as a function of the local observation and all the previous decisions received from the other sensors, *i.e.*

$$u_{i,m} = \mathcal{D}_i(X_i; \mathbf{u}_0, \cdots, \mathbf{u}_{m-1})$$

where $\mathcal{D}_i$ is the local decision function for sensor s_i[1].

In this paper, we adopt the n-th root decision fusion strategy proposed in [3] to define the messages that are generated during distributed computing. This is referred to as the Parley algorithm [3], as described in the following.

2.1. n-th Root Parley Algorithm

In the n-th root algorithm, the decision taken at sensor s_i after $m-1$ iterations of data exchange is

$$\begin{aligned} u_{i,m} &= \mathcal{D}(X_i; \mathbf{u}_0, \cdots, \mathbf{u}_{m-1}) && (2) \\ &= \begin{cases} 1, & \text{if } \Lambda(X_i) \geq \lambda_{i,m}^{root} \\ 0, & \text{otherwise} \end{cases} && (3) \end{aligned}$$

where $\Lambda(X_i) = f_{1,i}(X_i)/f_{0,i}(X_i)$ is the likelihood ratio obtained at sensor s_i given only the observation X_i. The threshold $\lambda_{i,m}^{root}$ is defined as

$$\lambda_{i,m}^{root} = \left(\lambda \cdot \frac{\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_0)}{\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_1)}\right)^{1/n} \times \frac{\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_1)/\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_1, X_i)}{\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_0)/\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_0, X_i)} \quad (4)$$

where the event $\mathbf{u}_0^m \equiv \bigcup_{n \leq m} \bigcup_k u_{k,n}$ and λ is the optimal decision threshold for centralized hypothesis testing (under either the Bayesian or the Neyman-Pearson criterion). Interestingly, as shown in [3], the global decision obtained with the n-th root algorithm achieves the same performance as the optimal centralized decision rule.

The Parley algorithm is particularly suitable for illustrating the data-driven property of the GTMA strategy since the decisions carry an increasing amount of redundancy as the distributed local decisions progressively converge to consensus. In fact, we show, in the following, that the aggregate entropy of the messages at the m-th iteration decreases as m goes to infinity. This is precisely the property that renders the gain of GTMA over standard sensor polling.

2.2. The Reduced Information Rate

In the Parley algorithm, the conditional probability

$$\mathbf{p}_m = \Pr(\mathbf{u}_m|\mathbf{u}_0, \cdots, \mathbf{u}_{m-1}) \quad (5)$$

can be induced from the data fusion process to facilitate the communication. When the convergence of decisions is reached, one can infer the subsequent decisions with probability 1 based on the previously received messages.

The information embedded in the messages at the m-th iteration can be measured as the conditional entropy given the probability in (5), *i.e.*

$$H(\mathbf{u}_m|\mathbf{u}_0, \cdots, \mathbf{u}_{m-1}) = \sum_{\mathbf{u}_0^m} \Pr(\mathbf{u}_0^{m-1})\mathbf{p}_m \log \mathbf{p}_m$$

where $\mathbf{p}_m$ is as defined in (5). Essentially, at each round of the Parley, the knowledge of the previous decisions can be viewed as the side information that is utilized to encode the new set of messages. As the side information is accumulated over previous rounds of Parley, the innovation within the transmitted messages decreases to zero.

In fact, we prove that the entropy of the set of messages decreases to zero as m goes to infinity. In the proof, we utilize the fact that the decisions converge almost everywhere and that the limit value is the same for all the nodes, *i.e.* $\lim_{m\to\infty} u_{i,m} = u_i$, for all i (convergence) and $u_i = u_j$, $\forall i, j$ (agreement) *a.e.* This has been proved for the Parley algorithm in [3] and more generally in [2, Theorem 6].

Theorem 1 *The entropy of the set of decisions $\mathbf{u}_m$, conditioned on the previous decisions $\mathbf{u}_0, \cdots, \mathbf{u}_{m-1}$ and the*

[1]The broadcast architecture is only a special case of the communication architecture proposed in [2] and the problem is targeted at minimizing a more general cost function than the one considered above. Certainly, the GTMA can also be applied to other decision fusion rules as well.

underlying hypothesis $\mathcal{H}_b$, converges to zero as m goes to infinity, i.e. $\lim_{m\to\infty} H(\mathbf{u}_m|\mathbf{u}_0^{m-1}) = 0$.

Proof:

$$\begin{aligned} &H(\mathbf{u}_m|\mathbf{u}_0, \cdots, \mathbf{u}_{m-1}) \leq H(\mathbf{u}_m|\mathbf{u}_{m-1}) \\ &= \mathbf{E}\left[-\log \Pr(\mathbf{u}_m|\mathbf{u}_{m-1})\right] \\ &= \sum_{\substack{\mathbf{u}_{m-1},\mathbf{u}_m \\ \in\{0,1\}^N}} -\Pr(\mathbf{u}_m, \mathbf{u}_{m-1}) \log \frac{\Pr(\mathbf{u}_m, \mathbf{u}_{m-1})}{\Pr(\mathbf{u}_{m-1})}. \end{aligned} \quad (6)$$

Let us define $\gamma_m(\alpha, \beta) = \Pr(\mathbf{u}_m = \alpha, \mathbf{u}_{m-1} = \beta)$ and $\phi_m(\alpha, \beta) = \Pr(\mathbf{u}_{m-1} = \beta)$. From [2, Theorem 6] and by defining $\lim_{m\to\infty} \mathbf{u}_m = \mathbf{u}$, we have, for $\alpha, \beta \in \{0, 1\}^N$,

$$\begin{aligned} &\lim_{m\to\infty} \gamma_m(\alpha, \beta) \\ &= \lim_{m\to\infty} \sum_{\mathcal{H}_b} p_{\mathcal{H}_b} \Pr(\mathbf{u}_m = \alpha, \mathbf{u}_{m-1} = \beta|\mathcal{H}_b) \\ &= \sum_{\mathcal{H}_b} p_{\mathcal{H}_b} \mathbf{E}\left[\lim_{m\to\infty} \mathbf{1}_{\{\mathbf{u}_m=\alpha\}} \mathbf{1}_{\{\mathbf{u}_{m-1}=\beta\}} |\mathcal{H}_b\right] \\ &= \begin{cases} 1 & \text{if } \alpha = \beta = \mathbf{u} \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

where $p_{\mathcal{H}_b} = \Pr(\mathcal{H}_b)$.
Case I: For $\alpha = \beta = \mathbf{u}$, we have

$$\begin{aligned} &\lim_{m\to\infty} \left|\Pr(\mathbf{u}_m, \mathbf{u}_{m-1}) \log \frac{\Pr(\mathbf{u}_m, \mathbf{u}_{m-1})}{\Pr(\mathbf{u}_{m-1})}\right| \\ &= \lim_{m\to\infty} |\gamma_m(\alpha, \beta) \log \gamma_m(\alpha, \beta) - \gamma_m \log \phi_m(\alpha, \beta)| \\ &\leq \lim_{m\to\infty} 2|\gamma_m(\alpha, \beta) \log \gamma_m(\alpha, \beta)| = 0 \end{aligned}$$

since $\lim_{m\to\infty} \gamma_m(\alpha, \beta) = 1$. The inequality comes from the fact that $|a + b| = |a| + |b|$ and that $\gamma_m \leq \phi_m$.
Case II: Similarly, for $\alpha \neq \beta$, we also have

$$\lim_{m\to\infty} \left|\Pr(\mathbf{u}_m, \mathbf{u}_{m-1}) \log \frac{\Pr(\mathbf{u}_m, \mathbf{u}_{m-1})}{\Pr(\mathbf{u}_{m-1})}\right| = 0,$$

since $\lim_{m\to\infty} \gamma_m \log(\gamma_m) = 0$.

Since the summation in (6) is finite, we have

$$\begin{aligned} &\lim_{m\to\infty} H(\mathbf{u}_m|\mathbf{u}_{m-1}, \mathcal{H}_b) \\ &= -\sum_{\mathcal{H}_b} p_{\mathcal{H}_b} \sum_{\substack{\mathbf{u}_{m-1},\mathbf{u}_m \\ \in\{0,1\}^N}} \lim_{m\to\infty} \gamma_m \log(\gamma_m/\phi_m) = 0. \end{aligned}$$

From Theorem 1, we know that the information embedded in the messages goes to zero as the number of iterations increases. In the following, we describe the data-driven property that we use to extract the relevant information from the messages and show, as an example, the strategy tailored specifically for the problem described in Section 2.

3. DATA-DRIVEN MULTIPLE ACCESS STRATEGY

Conventionally, multiple access protocols aim at providing each user an independent channel to transmit their local messages and the channel assignment is obtained independently of the underlying data structure. This is extremely inefficient in sensor networks where the messages of the distributed sensors are usually highly redundant. For instance, in the decentralized detection scenario that we are considering, it is often the case that the majority of the sensors make the same local decision since the underlying phenomenon is the same for all nodes. In this case, assigning an independent channel to each node would result in a large number of redundant transmissions.

In this paper, we utilize the previously proposed group testing multiple access (GTMA) [5, 6] scheme to schedule the transmission of sensors. The main intuition is to assign a single channel to a group of sensors that have a highly predictable set of messages to transmit, as opposed to assigning a channel to each individual sensor in the conventional case. This is analogous to the blood testing problem in [7] where a group of blood samples are pooled together for each test, instead of testing each sample independently, since it is likely that a group of samples are simultaneously non-defective (*i.e.* the realization with the highest probability). During each channel access, an intelligent guess is imposed on the group of sensors and the response to the guesses will refine the knowledge of the messages at each sensor until the messages at all nodes is resolved (up to a certain distortion).

3.1. GTMA for the Parley Algorithm

During each channel access, the GTMA policy determines a guess on a subset of messages $U \subset \mathcal{U} = \{u_{1,m}, \cdots, u_{N,m}\}$, for some m, and the sensors corresponding to these messages will respond collectively through the cooperative transmission channel. Let us denote by Z_l the output of the channel during the l-th channel access. The sequence of outputs $\mathbf{Z} = [Z_0, \cdots, Z_{L-1}]$ will constitute a data representation of the message $\mathbf{u}$, where L is a random variable depending on the realizations of $\mathbf{u}$, *i.e.* variable length encoding. If the queries take into account the statistics of the data, we expect GTMA to achieve compression gains in $\mathbf{Z}$ and, thus, reducing the number of channel accesses.

One can envision two different network topologies for the GTMA: one is the *centralized topology* where a central node executes the GTMA algorithm and selects the group of messages that are queried; the other one is the *decentralized topology* where the GTMA is executed at each individual node and we assume that all nodes have access to the cooperative transmissions of the sensor groups. In the previous case, the messages are gathered at the central node before they are broadcast back to the sensors; in the second case, all sensors are equally capable of decoding the messages since each cooperative response is accessible to all nodes and each node will transmit accordingly if it falls in the subset U determined by the GTMA.

In general, the output of the channel depends on the cooperative transmission strategy and the receiver structure of each sensor. In this paper, we consider specifically:

(a1) the decentralized network topology;

(a2) the binary *noiseless OR channel* (see below).

In the decentralized detection scenario, it is reasonable to expect that most sensors will make the same local decision which corresponds to the underlying phenomenon. Therefore, during each channel access, it is reasonable to guess that a group of messages $U \subset \mathcal{U}$ all take the same value b, where $b \in \{0, 1\}$ represents the most probable hypothesis. Consider the case where each sensor s_i that corresponds to a message in U, *i.e.* $u_{i,m} \in U$, will respond with a logical 1 if $u_i \neq b$, and it will respond with a logical 0 if $u_{i,m} = b$. The aggregate output Z will be equal to the logical OR of the responses, *i.e.* the output of the channel is

$$Z = \vee_{u_{i,m} \in U} \{u_{i,m} \neq b\}. \quad (7)$$

This is referred to as the *noiseless OR channel* [8]. With this specific channel, one can infer that $u_{i,m} = b$ for all $u_{i,m} \in U$ if and only if $Z = 0$; otherwise, the guess on the set U is incorrect and a smaller subset of U should be chosen for the subsequent transmissions. The ideal assumption that the channel is noiseless makes $\mathbf{Z}$ a lossless representation of the original field.

Remark **1** *The cooperative transmission strategy that achieves the noiseless OR channel in the decentralized setting has been proposed in [9] where a simple pulse transmitter and energy detector is used at each node. Each sensor, e.g.* s_i*, emits a pulse if and only if* $u_{i,m} \in U$ *and* $u_{i,m} \neq b$*. All the other nodes receive with an energy detector and relay the information if and only if a pulse is received and that it was not a source itself. We refer the readers to [6, 9] for further details.*

In Fig. 2, we show the flow chart of the GTMA/Parley algorithm. During the m-th round of the Parley algorithm, a set of messages $\mathbf{u}_m$ is generated at each node and the probability $\mathbf{p}_m$ can be induced from the previous rounds of information exchange. Based on the probabilities $\mathbf{p}_m$, the GTMA determines an *action* A that specifies the group of sensors U that are scheduled to transmit and the guess that is imposed on the sensors. In our case, we guess that all the sensors in U have the message equal to b where b is the optimal detection given the previous rounds of local decisions. The new channel output will be appended to the previous sequence of outputs and a new estimate of $\mathbf{u}_m$ is obtained. The GTMA continues until the estimate $\hat{\mathbf{u}}_m$ falls within the distortion D and the Parley process is terminated when the decisions converge.

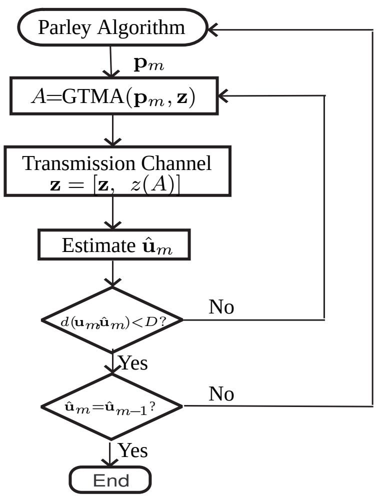

Fig. 2. The flowchart for the GTMA/Parley scheme.

Our goal is to derive the optimal GTMA strategy, or the optimal transmission scheduling policy, that minimizes the expected length of the data representation, *i.e.* $\mathbf{E}[L]$. However, it is easy to see that this is, in fact, an NP-complete problem. Interestingly, as shown in the following, a well-performing heuristic algorithm can be derived using the maximum entropy criterion.

Remark **2** *We should note that, in the multiple access scenario, the guesses and the responses of the sensors are restricted by the fact that sensors have knowledge of only their local data. If this restriction does not exist, one can simply implement the Huffman code by designing the tests to partition the observation space in half after each channel access.*

3.2. Information-Theoretic Approach to GTMA

Specifically, let $\mathbf{Z}_0^{l-1} = [Z_0, \cdots, Z_{l-1}]$ be the output of the channel determined up to the $(l-1)$-th transmission, and let $\mathbf{z}_0^{l-1} = [z_0, \cdots, z_{l-1}]$ be the realization of $\mathbf{Z}_0^{l-1}$. In this paper, we adopt the maximum entropy criterion to choose the best action in the function GTMA$(\mathbf{p}_m, \mathbf{z})$ shown in Fig. 2. In this algorithm, the action A is chosen such that the channel output contains the maximum entropy, *i.e.* the action of the l-th channel access is

$$A_l = \arg\max_A H(Z_l(A) | \mathbf{Z}_0^{l-1} = \mathbf{z}_0^{l-1}). \quad (8)$$

For each round of Parley, the message exchanges are initiated with the action A_0 which is derived with the conditional probability in (5). However, to simplify our analysis and the simulations in Section 4, we derive our strategy using the approximated probability

$$\mathbf{p}_m \approx \Pr(\mathbf{u}_m|\mathbf{u}_0, \cdots, \mathbf{u}_{m-1}, \mathcal{H}_b) \quad (9)$$

$$= \prod_{i=0}^{N-1} \Pr(u_{i,m}|\mathbf{u}_0, \cdots, \mathbf{u}_{m-1}, \mathcal{H}_b) \quad (10)$$

where $u_{i,m}$ is the m-th round decision at sensor s_i and $\mathcal{H}_b$ is the true hypothesis.

We note that the probability we have is actually

$$\Pr(\mathbf{u}_m = b \cdot \mathbf{1}|\mathbf{u}_0^{m-1}) = \frac{\sum_b \Pr(\mathbf{u}_0^{m-1}, \mathbf{u}_m = b \cdot \mathbf{1}|\mathcal{H}_b) p_{\mathcal{H}_b}}{\sum_b \Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_b) p_{\mathcal{H}_b}}$$

where $\mathbf{1}$ is an N-dimensional vector of 1's. However, for any reasonable data fusion strategy, the probability conditioned on the true hypothesis will eventually dominate the terms in both the numerator and the denominator, *i.e.* for m large, $\Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_b) \gg \Pr(\mathbf{u}_0^{m-1}|\mathcal{H}_{\bar{b}})$ when $\mathcal{H}_b$ is the true hypothesis. Therefore, the probability can be approximated as in (9).

With the abuse of notation, let $p_i = \Pr(u_{i,m} = b|\mathbf{u}_0^{m-1}, \mathcal{H}_b)$, $\forall i$ for a particular round of the Parley, *e.g.* the m-th round. Without loss of generality, we enumerate the sensors such that $p_0 \geq p_1 \geq \cdots \geq p_{N-1}$. To avoid searching over all possible groups of messages U in (8), we adopt a suboptimal search where we restrict the set U to consist of the least enumerated messages that have not yet been resolved. For example, if the channel outputs $\mathbf{z}_0^{l-1}$ infer the realization of $u_{1,m}, \cdots, u_{i-1,m}$ with probability 1 but can only infer the messages $u_{i,m}, \cdots, u_{N,m}$ with probability less than 1, then we say that $u_{1,m}, \cdots, u_{i-1,m}$ are resolved and the group of messages $U = \mathbf{u}_m^{i:j} = [u_{i,m}, \cdots, u_{j,m}]$, for some $j \leq N$, will be chosen according to the criterion in (8). With the noiseless binary OR channel described in (7), it is easy to show that the criterion in (8) can be reduced to

$$\mathbf{u}_m^{i:j} = \arg \min_{\mathbf{u}_m^{i:j} \in \mathcal{U}} \left| \Pr[Z(\mathbf{u}_m^{i:j})|\mathbf{Z}_0^{l-1} = \mathbf{z}_0^{l-1}] - \frac{1}{2} \right|. \quad (11)$$

The information theoretic approach described in this section approximates closely the entropy lower bound in many cases. In fact, the optimality is achieved when the probabilities fall in the regions specified by the cutoff probabilities proved in [10] for an *i.i.d.* Bernoulli model where $p_i = p$ for all i, and in [11] for the case where $p_i \neq p_j$ for $i \neq j$ (the cutoff probability defines the conditions for which assigning a channel to each individual node is optimal).

Corollary 1 *Assume the binary noiseless OR channel and the set of messages $u_0, \cdots, u_{N-1}$ that are $i.i.d.$ Bernoulli random variables with probability p. The entropy-based strategy results in testing each node separately if $p < \frac{-1+\sqrt{5}}{2}$ (which coincides with that derived by Ungar in [10]).*

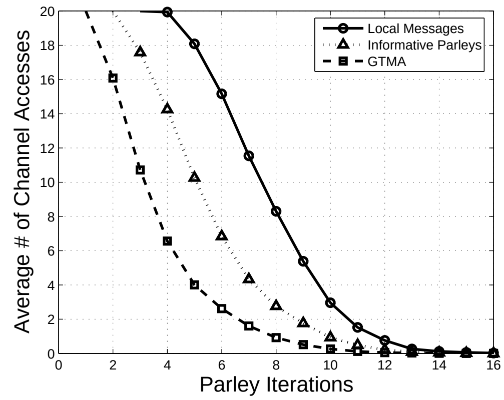

Fig. 3. The average number of channel accesses used for each round of information exchange, where $N = 20$, $\sigma = 2$, $\mu_0 = -1$ and $\mu_1 = 1$.

Proof: From (11), the entropy-based strategy results in testing a single node if the following holds:

$$\left|p - \frac{1}{2}\right| < \left|p^2 - \frac{1}{2}\right| \Rightarrow (p - \frac{1}{2})^2 < (p^2 - \frac{1}{2})^2$$
$$\Rightarrow p(1-p)(p^2 + p - 1) < 0$$

Since $0 < p < 1$, we have $p < \frac{-1+\sqrt{5}}{2}$.

Corollary 2 *Assume the binary noiseless OR channel and the set of messages $u_0, \cdots, u_{N-1}$ such that $p_0 \geq p_1 \geq \cdots \geq p_{N-1}$. The entropy-based strategy results in testing each node separately if $p_0(1 + p_1) < 1$ (which coincides with that derived by Kurtz and Sidi in [11]).*

The proof of Corollary 2 is omitted since it is similar to that shown in the proof of Corollary 1. The cutoff conditions proven for the entropy-based algorithm match exactly the cutoff conditions proven in [10] and [11].

4. SIMULATION RESULTS

In this section, we look specifically at the following binary hypothesis testing problem:

$$\begin{array}{ll} \mathcal{H}_0: & X_i \sim \mathcal{N}(\mu_0, \sigma^2) \\ \mathcal{H}_1: & X_i \sim \mathcal{N}(\mu_1, \sigma^2) \end{array} \quad (12)$$

where $\mu_0 = -1$, $\mu_1 = 1$ and $\sigma = 2$.

In Fig. 3, we show the average number of channel accesses required for each round of message exchanges with the n-th root Parley algorithm described in Section 2. The simulation considers a network of $N = 20$ nodes and the performance is averaged over 1000 realizations of the random phenomenon where $\Pr(\mathcal{H}_0) = \Pr(\mathcal{H}_1) = 0.5$. The

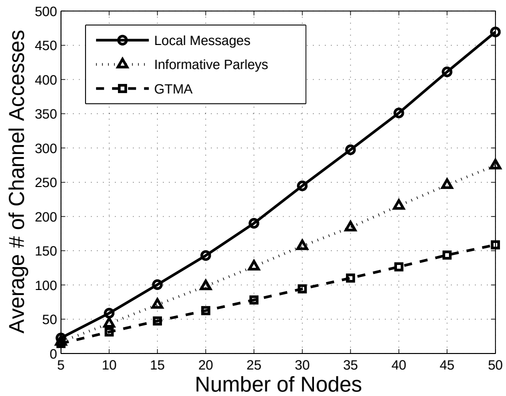

Fig. 4. The average number of channel accesses used before consensus is reached for $\sigma = 2$, $\mu_0 = -1$ and $\mu_1 = 1$.

solid curve represents the average number of local messages generated during that particular round of parley. The maximum entropy heuristic algorithm in (11) is used to determine the group transmissions and the cooperative channel is assumed to be the noiseless OR channel described in (7). We compare the performance of GTMA with the informative parley strategy proposed in [3] where the local decision is not transmitted if it can be inferred from the previously exchanged messages, *i.e.* the messages that are not 'informative' are not exchanged. In Fig. 3, we show that, by utilizing the statistical knowledge of the data to schedule the transmissions, the communication cost to exchange the messages is reduced significantly.

Using the same parameters as in Fig. 3, we show, in Fig. 4, the average number of channel accesses required to reach consensus for different network sizes. It is shown that the rate of increase with respect to the number of nodes is smaller for the GTMA, thus, the savings increase with the network size. The figures imply that even though the number of informative messages are large during the early rounds of parley, the information may be redundant among different messages. In this case, the data-driven multiple access scheme gains over the traditional multiple access scheme.

5. CONCLUSIONS

In this paper, we studied the communication complexity of a specific decentralized detection problem where the sensors in the network continuously update and exchange their local hard decisions until consensus in decision is reached at all nodes. We show that by utilizing the Group Testing Multiple Access strategy, one can reduce significantly the number of channel accesses required by the distributed computing algorithm to reach the computational goal. The GTMA exploits both the compressibility of the distributed messages and the sensors' ability to transmit cooperatively in a wireless medium. This work solves the communication problem for distributed "in network" computation which has been the main obstacle for implementing distributed computing algorithms or distributed data fusion in sensor networks.

6. REFERENCES

[1] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inform. Theory*, vol. IT-46, no. 2, pp. 388–404, Mar. 2000.

[2] John N. Tsitsiklis and Michael Athans, "Convergence and asymptotic agreement in distributed decision problems," *IEEE Trans. Automat. Contr.*, vol. AC-29, no. 1, 1984.

[3] Peter F. Swaszek and Peter Willett, "Parley as an approach to distributed detection," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 31, no. 1, pp. 447–457, Jan. 1995.

[4] S.S. Pradhan, J. Kusuma, and K. Ramchandran, "Distributed compression in a dense microsensor network," *IEEE Signal Proc. Mag.*, vol. 19, no. 2, Mar. 2002.

[5] Yao-Win Hong and Anna Scaglione, "On multiple access for correlated sources: A content-based group testing approach," in *Proc. of IEEE Information Theory Workshop*, San Antonio, TX, Oct. 2004.

[6] Yao-Win Hong and Anna Scaglione, "On multiple access for correlated sources: A content-based group testing approach," *submitted to Trans. on IT*, 2005.

[7] Robert Dorfman, "The detection of defective members of large population," *The Annals of Mathematical Statistics*, vol. 14, no. 4, pp. 436–440, Dec. 1943.

[8] Toby Berger, Nader Mehravari, Don Towsley, and Jack Wolf, "Random multiple-access communication and group testing," *IEEE Trans. Commun.*, vol. 32, no. 7, pp. 769–779, July 1984.

[9] Yao-Win Hong and Anna Scaglione, "A scalable communication architecture for the sensor broadcast problem," in *to be presented at SPAWC*, 2005.

[10] Peter Ungar, "The cutoff point for group testing," *Commun. on Pure and Applied Math.*, vol. 13, 1960.

[11] Doron Kurtz and Moshe Sidi, "Multiple-access algorithms via group testing for heterogeneous population of users," *IEEE Trans. Commun.*, vol. 36, no. 12, pp. 1316–1323, Dec. 1988.