

TOPOLOGY OF SENSOR NETWORKS IN DISTRIBUTED DETECTION

Saeed A. Aldosari, José M. F. Moura

Department of Electrical and Computer Engineering, Carnegie Mellon University
Pittsburgh, U.S.A

{saldosar, moura}@ece.cmu.edu

ABSTRACT

We study the topology of sensor networks. With parallel architectures, the design is trivial—sensors forward their local decisions to a global fusion center. With web architectures, sensors communicate only with ‘neighbors’ and evolve their local decisions to reach a ‘consensus.’ In practice, it is important to reach a consensus with minimal communications and processing cost. The convergence rate of the consensus algorithm depends on 1) the weights assigned to the network links; and 2) the connectivity pattern of the network. We apply concepts from small world networks to design the topology and the weights when the local decisions are quantized and study the impact on network performance when we trade number of links for number of bits per decision.

1. INTRODUCTION

Network topology impacts the performance of distributed decision making in sensor networks. Parallel networks are vulnerable to central node failures and require the sensor nodes to route their information through the complex network. We consider distributed detection in web architectures where sensors communicate only with neighboring sensors. We use an iterative averaging algorithm, recently proposed for estimation [1],[2]. In [3, 4], we studied the design of the weights associated with the network links and the design of the connectivity graph when messages are exchanged among sensors perfectly, with arbitrary precision. Here, we focus on: 1) the effect of quantizing the local decisions on the performance of the distributed detector; 2) the design of the connectivity graph with quantized decisions; and 3) tradeoffs between network parameters (number of links and numbers of bits/ decision).

We show that the distributed detector converges even with quantized messages but the steady state performance may not approach that of the optimal parallel architecture. For fixed network communication cost per iteration, i.e., for fixed Mb where M is the number of links and b is the number of bits/ link/ iteration, our results show that: 1) the detection algorithm converges faster when we increase M , while the steady state performance improves when we increase b ; 2) the sensor network topology affects both the convergence speed of the iterative algorithm and the steady state performance (probability of decision errors); finally, 3) the convergence speed and steady state error probability are significantly improved when the networks exhibit a *small-world* behavior.

We consider a binary hypothesis test H_0 (target absent) versus H_1 (target present): N sensors collect measurements $\mathbf{y} =$

Partially supported by DARPA Integrated Sensing and Processing (ISP) Program under grant ARO DAAD 19-02-1-0180 and by NSF grants # ECS-0225449 and # CNS-0428404.

(y_1, \dots, y_N) that are assumed to be independent and identically distributed (i.i.d.) conditioned on the true state H with known conditional density $f_i(y) = f(y|H_i)$, $i = 0, 1$.

The sensors cooperate over a communication network represented by a graph $G = (V, \mathcal{E})$ where the set of N vertices $V = \{s_1, \dots, s_N\}$ corresponds to the sensor nodes and $\mathcal{E} = \{l_m = (i_m, j_m), m = 1, \dots, M\}$ is the set of M edges, where $(i, j) \in \mathcal{E}$ whenever sensor s_i can communicate with s_j . We consider *connected* graphs for which there exist a *path* (possibly multi-hop) from every sensor s_i to every other sensor s_j , $i, j = 1, \dots, N$.

The goal of the distributed decision algorithm is to reach a global *consensus* (common decision \hat{H}) about the true state H , by *exchanging local* information over the network G : the sensor nodes can only *route* their data to their immediate neighbors, i.e., to nodes they connect to through an edge in the connectivity graph G . We consider a Bayes minimum probability of error network detector whose performance is measured by the average probability of decision error $P_e = \Pr(\hat{H} \neq H) = \pi_0 P_{e0} + \pi_1 P_{e1}$, where $P_{e0} = \Pr(\hat{H} = H_1 | H_0)$, $P_{e1} = \Pr(\hat{H} = H_0 | H_1)$, and π_0, π_1 are the prior probabilities of H_0 and H_1 , respectively.

Section 2 presents the distributed detector when the network topology is fixed and the sensor local messages are unquantized. Section 3 reviews generators for small-world graphs. Section 4 considers that the messages are quantized by b bits to conserve communication and bandwidth resources. Section 5 addresses the impact of quantization on the performance of the detector and studies the tradeoffs between the number of communication links M and the number of quantization bits b per message when the total number of communication bits bM is fixed. It also investigates the implications that the message quantization has on the design of the connectivity graph of the sensor network. Finally, Section 6 concludes the paper.

2. DISTRIBUTED DETECTION

We present a distributed decision algorithm, which preforms global fusion by means of local exchange of information between neighboring sensors. In this section we assume that the connectivity graph over which the sensors communicate is known and fixed. We also assume that the sensors exchange raw messages without quantization. In the proposed distributed decision algorithm, the sensors communicate only with their immediate neighborhood, i.e., the nodes to which they are connected to by an edge in the connectivity graph G . The statistics stored at the sensors are updated in an iterative manner to reach a consensus about the global statistic.

1. Initially, sensors take measurements y_1, \dots, y_N , after which each sensor s_n , $n = 1, \dots, N$, computes the log likelihood ratio $r_n = \log \frac{\Pr(y_n|H_1)}{\Pr(y_n|H_0)}$ of its measurement y_n . These

LLR values serve as the initial states of the consensus decision algorithm. We denote the value stored at sensor s_n at iteration t by $x_n(t)$, and we refer to it as the *state* of the sensor s_n at time t . At $t = 0$, the initial states are given by $x_n(t = 0) = r_n, n = 1, \dots, N$.

- The consensus decision algorithm is then carried out according to the following linear operation [1]

$$x_n(t) = W_{nn}x_n(t-1) + \sum_{i \in \Omega_n} W_{ni}x_i(t-1), \quad (1)$$

where Ω_n is the set of neighbors of sensor s_n and W_{ij} is a weight associated with edge (i, j) (if it exists). The consensus decision procedure is expressed in matrix format as

$$\mathbf{x}(t) = \mathbf{W}\mathbf{x}(t-1), \quad (2)$$

where $\mathbf{x}(t)$ is the $N \times 1$ vector of all current states and $\mathbf{W} = \{W_{ij}\}$ is the matrix of all weights. Entry (i, j) in \mathbf{W} is zero if $(i, j) \notin \mathcal{E}$. In terms of the initial state, the consensus updating rule is $\mathbf{x}(t) = \mathbf{W}^t \mathbf{x}(0)$, where $\mathbf{x}(0)$ is the $N \times 1$ vector of all the initial states $x_n(0) = r_n, n = 1, \dots, N$. The iterative procedure is carried out until $\forall n \lim_{t \rightarrow \infty} x_n(t) = R/N$, i.e., the state $Nx_n(t)$ of each sensor converges to the global log likelihood ratio $R = \sum_{n=1}^N r_n$.

- The final decision is then obtained locally at each sensor according to its current state

$$x_n(t) \begin{cases} \hat{H}(n)=1 \\ \geq v/N, \\ \hat{H}(n)=0 \end{cases}$$

where $\hat{H}(n)$ denotes the decision of sensor s_n , and $v = \log(\pi_0/\pi_1)$ is the fusion threshold.

Design of the Link Weights For a fixed graph topology G , the link weights W_{ij} are designed to ensure convergence of the consensus decision algorithm. Convergence is guaranteed, e.g., [1], whenever: 1) the sum of each row or column of \mathbf{W} is one; and 2) the spectral radius $\rho(\mathbf{W} - \frac{1}{N}\mathbf{1}_{N \times N})$ is less than one. In this paper, we restrict the weights to be symmetric, i.e., $W_{ij} = W_{ji}$. Under this restriction, the convergence speed is maximized by minimizing the spectral norm $\|\mathbf{W} - \frac{1}{N}\mathbf{1}_{N \times N}\|_2$. We use a primal barrier interior-point method to solve this semidefinite programming problem. Details of the weight optimization algorithm can be found in [1] and [3].

Performance Analysis—Unquantized Gaussian Observations Consider a simple Gaussian shift-in-mean model where the observation at sensor s_n under H_i is characterized by $y_n = \mu_i + z$, $z \sim \mathcal{N}(0, \sigma^2)$ where μ_i is the signal mean under H_i , $i = 0, 1$ and z is a zero-mean Gaussian noise with variance σ^2 . Without loss of generality, we let $\mu_1 = -\mu_0 = \mu$. In the consensus decision algorithm (2), decisions are made at all sensor nodes based on a weighted sum of the measured LLRs $r_n, n = 1, \dots, N$. It can be shown that, at iteration t , the probabilities of error at sensor s_n under H_i are given by

$$P_{ei}^n = Q\left(-\frac{v - 2\mu\mu_i N/\sigma^2}{2\mu_i N \sqrt{[\mathbf{W}^{2t}]_{nn}/\sigma}}\right), \quad (3)$$

where $[\mathbf{W}^{2t}]_{nn}$ denotes the n th diagonal element of \mathbf{W}^{2t} and $Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-y^2/2} dy$ is the normal right-tail distribution. As $t \rightarrow \infty$, $[\mathbf{W}^{2t}]_{nn}$ converges to $1/N$ and (3) approaches the performance of an optimal *parallel fusion network* where all sensors deliver their measurements to a single fusion center.

The probabilities of error in (3) followed easily because the sensor measurements are normally distributed and they are forwarded without quantization. When the measurement statistics are non-Gaussian or if they are discrete (e.g., quantized), the error performance is much more involved. In [5], we presented a methodology based on large deviations, more specifically on the saddlepoint approximation, that computes efficiently very accurate approximations to the error probabilities for parallel architectures. This method can also be extended to approximate the error probabilities of the distributed detection algorithm presented here with arbitrary network architectures and under arbitrary observation models (e.g., non-Gaussian, discrete, or quantized observations). We will not pursue this further here due to lack of space.

3. SMALL-WORLD GRAPH GENERATORS

In general, designing the network architecture that maximizes the performance of distributed detection is a very complex problem. Exhaustive search over all possible graph topologies is impractical; randomized search strategies are more useful, but a totally random search will take too long to find good connectivity graphs. Instead, we use a model that generates graphs with certain small-world properties similar to what has been found in many social, biological, and technological networks. These networks exhibit small average path length and strong local clustering, [6]. We use the generator in this reference that can produce graphs ranging from highly structured to completely random. The graph generation is started by constructing a highly structured circular network in which nodes are placed on a circle and each node is connected to its $2k_w$ closest nodes. Then, a random rewiring procedure is conducted on all graph links. With probability a_w , each link is rewired to a different destination chosen uniformly at random. Self and parallel links are prevented in the rewiring procedure, while the number of links is kept constant regardless of a_w . Setting a_w close to zero results in generating highly structured graphs similar to the original circular network, which suffer from high average path length between node pairs. On the other hand, $a_w = 1$ results in a totally random network with low average path length but with weak clustering. Reference [6] observes that, as a_w changes from 0 to 1, a phase-change occurs at which the generated networks start exhibiting a small-world behavior characterized by high clustering coefficient and small average path length.

4. QUANTIZATION EFFECTS

So far, it was assumed that the sensors exchange their state values $x_n(t)$ with high precision. In practice, however, these values are quantized into finite levels (i.e., finite number of bits) in order to conserve power and bandwidth resources. In this section, we examine the effects resulting from quantizing the values exchanged among sensors. In (1), it is assumed that the state values $x_i(t-1), i \in \Omega_n$ of neighboring sensors are received at sensor s_n perfectly without any quantization. When these values are quantized the consensus decision algorithm is carried out according to the following update rule

$$x_n(t) = W_{nn}x_n(t-1) + \sum_{i \in \Omega_n} W_{ni} \Psi[x_i(t-1)], \quad (4)$$

where Ψ denotes quantization. Note that sensor s_n can use its own stored value $x_n(t-1)$ without quantization (i.e., with high precision).




Fig. 1: A small-world network of $N = 51$ nodes, $M = 2N$ links generated using the Watts-Strogatz model with a rewiring probability $a_w = 0.3$.

The quantizer Ψ introduces quantization noise at every iteration of (4). It may appear at first that the accumulation of the quantization noise at every iteration would result in the failure of the distributed detection algorithm at some point, i.e., the divergence of the detection error after the accumulative effects of the quantization errors becomes large. This could happen, e.g., if the quantization noise is modeled as an additive noise with identical statistics at every iteration of the detection algorithm. In [7], the effect of additive noise and its implications on the link-weights design problem are discussed. It was also emphasized that the iterative algorithm may not converge in any meaningful sense under such additive noise scenario. Interestingly, the quantization considered here does not fall under the model of [7] in the sense that the quantization errors at every iteration may not be identically distributed. In fact, simulation results show that the mean of the quantization noise falls rapidly to zero while its variance decays as the iterative algorithm progresses. This means that after a certain number of iterations the effect of quantization becomes very small and the probability of detection error approaches a certain value. It should be noted however that the process of quantization may cause a degradation in the convergence speed and/or the steady state error probability of the distributed detection algorithm. This will be the subject of investigation in the following section.

5. SIMULATIONS

In this section we (1) illustrate the effect of quantization on the convergence of the distributed detection algorithm, and (2) investigate the topology of networks that lead to good convergence properties. We consider a network of $N = 51$ sensors that communicate over a graph G generated using the Watts model with number of links given by $M = k_w N$. Fig. 1 shows a sensor network generated using this model when $k_w = 2$ and the probability of rewiring $a_w = 0.3$. The sensor observations follow a Gaussian shift-in-mean model where the conditional signal means are assumed to be $\mu_1 = -\mu_0 = 1$ and the SNR = -8 dB. The values exchanged between the sensors at every iteration are quantized using a b -bit non-uniform quantizer Ψ . The quantization thresholds are designed with respect to the saddlepoint approximation as described in [5] while the quantization levels are set at the means of the quantization intervals.

Tradeoffs and Quantization Effects The performance (i.e., prob-




Fig. 2: Convergence of the distributed detection algorithm in a network of $N = 51$ sensors. Under quantization, the number of links M and the number of quantization bits b are chosen such that bM is maintained constant. The unquantized cases are represented by $b = \infty$.

ability of error) of the distributed detection algorithm is assessed through Monte-Carlo simulations. Specifically, a large number of random observation vectors (each containing N sensor observations) are generated and the iterative detection algorithm (described in section 2 and (4)) is run for every randomly generated observation vector $\mathbf{x}(0)$. At every iteration, the sensor decision errors are counted and averaged over all sensors and over all Monte-Carlo experiments to produce an estimate of the probability of error after a particular number of iterations. Fig. 2 shows the convergence of the average probabilities of two networks: NET2: $M = 2N$, $b = 8$ quantization bits (shown in Fig. 1), and NET4: $M = 4N$, $b = 4$ bits (not shown here). Note that the parameters of NET2 and NET4 are chosen so that, at every iteration, the total number of exchanged bits bM is the same for both networks. This enables us to study the tradeoffs involved between the number of links M and the number of quantization bits b . Fig. 2 shows that the performance of distributed detection over both networks is close to optimal (i.e., parallel network with unquantized measurements). However there is a clear advantage of NET4 (with higher number of links) in terms of the convergence speed while, on the other hand, NET2 (with higher number of quantization bits) is slightly better in terms of the steady state probability of error. The case of unquantized measurements is also included in Fig. 2 (estimated from (3)) to emphasize the effect of quantization. It can be seen that quantization degrades both the convergence speed and the steady state performance.

Connectivity Graph Design An important issue in sensor networks is the design of the underlying connectivity network that supports the information exchange among sensors. In [4], we showed how the convergence speed of the distributed detection algorithm can be greatly improved by proper design of the network graph. Under quantization, however, the connectivity pattern not only affects the convergence speed but also the steady state performance of the iterative algorithm. Rather than considering the problem of topology-design in general, we focus on networks that exhibit small-world properties. We generate random networks according to the Watts-Strogatz model where the probability of rewiring a_w is varied from 0 (circular network) to 1 (random




Fig. 3: The steady state probability of error to which the distributed detection algorithm converge to under different quantization scenarios. Note that bM is kept constant.

network). For each randomly generated network we estimate the performance (probability of error) of the distributed detection algorithm from which we compute the steady state performance (i.e., average probability of error as the number of iterations $t \rightarrow \infty$) and the convergence speed of the iterative algorithm (i.e., the number of iterations required to approach within 10% of the steady state performance).

Like the previous example, we consider two scenarios $M = 2N$, $b = 8$ quantization bits, and $M = 4N$, $b = 4$ bits where the total number of exchanged bits bM is the same for both in order to study the tradeoff between the number of links M and the number of quantization bits b . Fig. 3 and Fig. 4, respectively, show the steady state probability of error and the convergence speed averaged over 100 randomly generated networks at each value of the rewiring probability a_w . It can be seen that networks with higher number of quantization bits provide a slightly better steady state performance. On the other hand, networks with more links achieve a considerably faster convergence while still providing near optimal performance. More importantly, results illustrate how the different connectivity patterns impact the detection performance. It can be seen that the performance in terms of the steady state error probability and the convergence speed, is worst when $a_w \rightarrow 0$ corresponding to a circular network characterized by a high clustering coefficient and a high average path length. The performance improves as random rewiring is introduced into the network structure with an optimum around $a_w = 0.85$. Notice that even a low ratio of random rewiring, e.g. 30%, can greatly improve the performance over that of a regular network with nearest neighbor connectivity. At such rewiring ratio, the generated networks have the low average path length and high clustering coefficient characteristic of small-world networks.

6. CONCLUSIONS

Results in this paper show that the proposed distributed detection algorithm converges even when the sensor messages are quantized for the purpose of conserving power and bandwidth resources. Results also show that the performance can be greatly improved by introducing a small degree of randomness (small number of short-

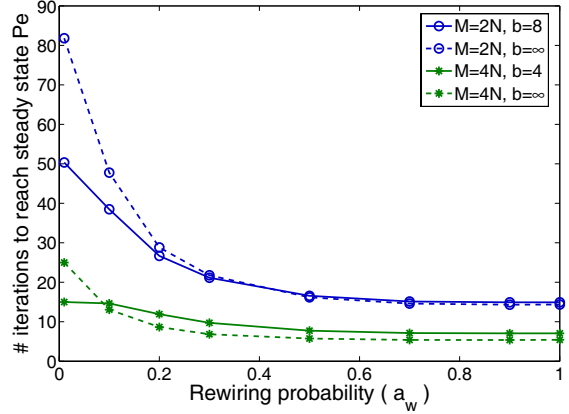


Fig. 4: The number of iterations needed so that the performance of the distributed detection algorithm reaches within 10% of the steady state performance.

cuts) with which the network attains desirable small-world characteristics. There are many avenues to be explored in relation to the current study, in particular, how to design the weights associated with the network links, and the local quantizers. Also, throughout this paper, we assumed equal cost for all communication links regardless of the *physical* length of these links. It would be practically interesting to consider these costs in the design process. Results based on the equal cost assumption can serve as a first-cut study to have insight over what can be characterized as “good” connectivity patterns.

7. REFERENCES

- [1] L. Xiao and S. Boyd, “Fast linear iteration for distributed averaging,” *Syst. Contr. Lett.*, vol. 53, pp. 65–78, Sept. 2004.
- [2] R. Olfati-Saber and R. Murray, “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Trans. Automat. Contr.*, vol. 49, pp. 1520–1533, Sept. 2004.
- [3] S. Aldosari and J. M. F. Moura, “Decision making in sensor networks: Topology issues,” Oct. 2005, Submitted.
- [4] S. Aldosari and J. M. F. Moura, “Distributed detection in sensor networks: connectivity graph and small-world networks,” in *Asilomar Conf. on Sig., Syst., and Comp.*, 2005.
- [5] S. Aldosari and J. M. F. Moura, “Saddlepoint approximation for sensor network optimization,” in *Proc. of ICASSP’05*, Philadelphia, PA, Mar. 2005, vol. 4, pp. 741–744.
- [6] D. Watts and S. Strogatz, “Collective dynamics of small-world networks,” *Nature*, vol. 393, pp. 440–442, June 1998.
- [7] L. Xiao, S. Boyd, and S. Kim, “Distributed average consensus with least-mean-square deviation,” submitted.