Decision-Making Accuracy for Sensor Networks with Inhomogeneous Poisson Observations

Chetan D. Pahlajani, Indrajeet Yadav, Herbert G. Tanner, and
Ioannis Poulakakis

Abstract The paper considers a network of sensors which observes a time-inhomogeneous Poisson signal and has to decide, within a fixed time interval, between two hypotheses concerning the intensity of the observed signal. The focus is on the impact of information sharing among individual sensors on the accuracy of a decision. Each sensor computes locally a likelihood ratio based on its own observations, and, at the end of the decision interval, shares this information with its neighbors according to a communication graph, transforming each sensor to a decision-making unit. Using analytically derived upper bounds on the decision error probabilities, the capacity of each sensor as a decision maker is evaluated, and consequences of ranking are explored. Example communication topologies are studied to highlight the interplay between a sensor's location in the underlying communication graph (quantity of information) and the strength of the signal it observes (quality of information). The results are illustrated through application to the problem of deciding whether or not a moving target carries a radioactive source.

1 Introduction

We currently rely on networks of distributed sensors for triggering a timely response to emergency situations, including natural disasters such as hurricanes [8], earthquakes [10] and tsunamis [6]. Existing work on decision-making over networks of observers that monitor a physical process and have to decide on its state, suggests that the performance of decision-making is affected by the structure of the network. For example, in the context of a network of stochastic evidence accumulators, each

C. D. Pahlajani
Indian Institute of Technology, Gandhinagar, e-mail: cdpahlajani@iitgn.ac.in

I. Yadav, H. G. Tanner, I. Poulakakis
University of Delaware, USA, e-mail: {indragt,btanner,poulakas}@udel.edu

represented by a drift-diffusion process accruing evidence in continuous time by observing a noisy signal, information exchange among individual nodes affects the certainty of each node in a way that is governed by information centrality [14, 15]. Information centrality is a structural property of the underlying interconnection graph that depends on the totality of paths connecting that node with the rest of the network [17]. Heterogeneity in the network has been introduced in [5] by allowing only a limited number of units—termed leaders—to observe the external signal directly. Although these results refer to continuous-time implementations of sequential probability ratio tests inspired by human decision-making models [2], they still highlight the effect of general network topologies on decision-making performance.

Only specific communication topologies for the network have been explored, due to the explosive combinatorial complexity of decision making in distributed sensor systems [20]. There is emphasis on particular directed rooted trees [19, 21–23], where the root is a designated fusion center node that makes the final decision. In this context, and for this limited range of topologies where information from all observers eventually trickles down in some compressed form to the designated fusion center, results show that decision performance is affected by the structure of the network. In addition, to avoid the complexity associated with distributed decision making, the results for more general tree topologies are restricted to an asymptotic analysis, where the order of the graph grows unbounded. For some applications, however—such as nuclear detection—practical and economic considerations preclude the deployment of large-scale networks [18]. In that smaller-scale regime, and particularly for the case where decision makers do not have access to (compressed) information from all other observers, there is not enough knowledge and insight to determine the specific effect of general network topologies on decision-making performance.

The present paper focuses on the case of a sensor network which observes a time-inhomogeneous Poisson process, and has to decide, within a fixed time interval, between two hypotheses concerning the intensity of the observed process. Armed with analytical expressions for bounds on error probabilities for this test in the classical networking case where all sensors, after observing this inhomogeneous process for a time period, submit a statistic to a central decision maker (fusion center) [11], this paper explores alternative cases where the topology of the network can vary, and any sensor can potentially become the decision maker. The goal here is to explore the effect of information sharing on the capacity of individual sensors in the network to make accurate decisions.

The topology of the sensor network is modeled by a directed graph; each sensor is represented by a node, and a directed edge from node i to node j signifies that there is directed flow of information from sensor i to sensor j. Sensors do not need to communicate until time instant T, at which a decision needs to be made by the network. At that time instant each sensor sends its local likelihood ratio $L_T(i)$ along outgoing edges in the communication graph. Next, each sensor—now referred to as a Decision Maker (DM)—implements a Likelihood Ratio Test (LRT) where the product of the available $L_T(i)$'s is compared against a threshold to make a (local) decision. The problem is now to assess the relative accuracy of the DMs.

Intuitively, the factors improving decision accuracy are large numbers of incoming edges in the information-sharing graph, and availability of $L_T(i)$'s collected from sensors which observe strong signals. The principal challenge now lies in formalizing *exactly* how these factors mathematically influence the decision error probabilities, given that these probabilities are, in all but perhaps the simplest cases, intractable to analytic computation. The paper circumvents this difficulty by working with Chernoff upper bounds on the error probabilities for each DM, borrowing the analytic expressions of which from existing work [12]. Treating these explicitly computable bounds as proxies for the true (unknown) error probabilities, an index for decision-making accuracy can be formulated and subsequently used to rank the DMs. The advantage of this approach is that it recasts the problem of ranking incomputable probabilities for each of the DMs in terms of ranking natural computable surrogates, at the expense of some sharpness.

These results find natural application in the problem of detecting, using a distributed network of radiation sensors, illicit nuclear material in transit. The severity of the threat associated with radioactive materials falling into the hands of potential terrorists has been recognized [4]. Possible mitigation strategies include networks of detectors deployed along roads and highways, tasked with detecting illicit material that has slipped through border checks or portal alarms [4].

There are at least two reasons that make nuclear detection—irrespective of whether the sensors are static or mobile—extremely challenging. The first is that radiation detectors pick up signals emitted not just by the illicit radioactive material to be detected, but also from ubiquitous, naturally occurring, background radiation. From a counter's perspective, the two signals are of identical nature and indistinguishable once superimposed. The second reason relates to attenuation: although a kilogram of Highly Enriched Uranium (HEU) can emit as many as 4×10^7 gamma rays per second [4], shielding and attenuation can limit the effective detection range to a few feet, and require detection times that can range from several minutes to hours. To put these in perspective, the gamma-ray emission of nuclear missiles containing HEU becomes comparable to background just 25 cm away from the warhead [16]. The problem is exacerbated by the motion of the signal source. Not only does the mathematical model of the physical phenomenon change (becoming time-inhomogeneous), but now detectors have limited time to decide before the target disappears from sight: the sensors are faced with a problem of detecting in a matter of seconds, a weak time-varying signal, buried inside another signal of same nature.

It is important to note that while the technical approach in this paper is particularized in the context of nuclear detection, the impact of the methods proposed here can also reach other application domains which involve networked decision-making with Poisson process observations.

2 Background

As noted above, our analysis focuses on the effect of information sharing on decision-making by a network of sensors observing a time-inhomogeneous Pois-

son process, and as such, is applicable in a variety of domains. That said, it will be convenient from this point on to frame the discussion largely in terms of applications to nuclear detection. It goes without saying that the treatment can be carried over to other applied problems after suitably reinterpreting various mathematical quantities.

With the aforementioned application domain in mind, consider a network of k radiation sensors that is deployed over a spatial region of interest for the purpose of detecting illicit nuclear materials in transit; see Figure 1. The typical detection scenario involves a vehicle (target) suspected of being a carrier of nuclear material (source) moving through the sensing field of this network. The objective is to decide, at the end of a fixed time interval $[0,T]$, whether the counts recorded at the sensors can be attributed solely to ubiquitous background radiation (hypothesis H_0) or whether they contain, in addition, radiation from a source carried by the moving target (hypothesis H_1). Local decision making can be enhanced through allowing, at the terminal time T, limited communication among sensors according to a suitable communication topology, following which each sensor is enabled to act as a DM operating on the information available to it. The primary goal of the present work is to formulate a metric for ranking the decision-making accuracy of the individual DMs. To wit, it is desirable to develop an index that mathematically captures the interplay between centrality of a DM in the network (quantity of information) and proximity to the suspected source (quality of information) in a way that naturally mimics the true, but incomputable, error probabilities.

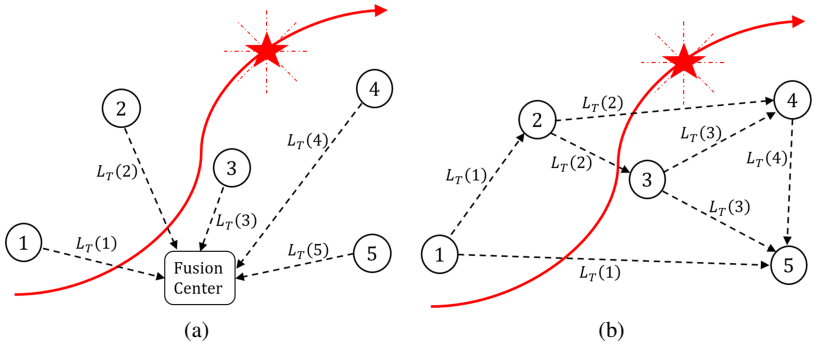

Fig. 1 A source moving (red line) in the sensing field of a network of sensors. Different information sharing scenarios are depicted. (a): All sensors send their L_T's to a fusion center. (b): The sensors share their L_T's according to a directed graph. Communication is allowed only at time T.

A review of some existing notation and results [11] helps set the stage for the decision-making rules considered in this paper. The probabilistic setup is as follows. On the measurable space $(\Omega,\mathscr{F})$, there is a k-dimensional vector of counting processes $\boldsymbol{N}_t \triangleq (N_t(1),\ldots,N_t(k))$, $t\in[0,T]$. Here, $N_t(i)$ represents the number of counts recorded at sensor $i\in\{1,2,\ldots,k\}$ up to (and including) time $t\in[0,T]$. The two hypotheses, H_0 and H_1, regarding the state of the environment correspond, respectively, to two distinct probability measures $\mathbb{P}_0$ and $\mathbb{P}_1$ on $(\Omega,\mathscr{F})$. With respect

to $\mathbb{P}_0$, the $N_t(i)$'s, $1 \le i \le k$, are assumed to be independent Poisson processes with $N_t(i)$ possessing intensity $\beta_i(t)$, while with respect to $\mathbb{P}_1$, the $N_t(i)$'s, $1 \le i \le k$, are assumed to be independent Poisson processes with intensities $\beta_i(t) + \nu_i(t)$, respectively. The functions $\beta_i(t)$, $\nu_i(t)$, $1 \le i \le k$, defined for $t \in [0,T]$ are assumed to be bounded, continuous and strictly positive as in [11]. Here, $\beta_i(t)$ is the (possibly time-varying) intensity at time t due to background radiation at the spatial location of sensor i, while $\nu_i(t)$ represents the intensity due to the source (if present) as perceived by sensor i at time t. Note that the time-dependence of $\nu_i(t)$ arises from relative motion between the source and the sensor; indeed, in the context of radiation measurement it is generally accepted that $\nu_i(t)$ is proportional to the inverse square of the distance $r_i(t)$ between the source and sensor i [9].

A test for deciding between H_0 and H_1 can be thought of as an event B_1, whose occurrence or non-occurrence can be ascertained on the basis of sensor observations over $[0,T]$, and has the following significance: If the outcome $\omega \in B_1$, decide H_1; if the outcome $\omega \in B_0 \triangleq \Omega \setminus B_1$, decide H_0. For such a test, two types of errors can occur. A false alarm occurs if $\omega \in B_1$ with H_0 being the correct hypothesis; this occurs with probability $\mathbb{P}_0(B_1)$. A miss occurs if the outcome $\omega \in B_0$ while H_1 is the true hypothesis; this occurs with probability $\mathbb{P}_1(\Omega \setminus B_1)$.

In this setting, the optimal test for deciding between H_0 and H_1 is an LRT obtained as follows [3, 11]. For $i \in \{1,2,\ldots,k\}$, let $\tau_n(i)$ for $n \ge 1$ denote the n-th jump time, i.e., time at which the sensor generates a count, of $N_t(i)$, and let

$$L_T(i) \triangleq \exp\left(-\int_0^T \nu_i(s)\,\mathrm{d}s\right) \prod_{n=1}^{N_T(i)} \left(1 + \frac{\nu_i(\tau_n(i))}{\beta_i(\tau_n(i))}\right) \tag{1}$$

with the convention that $\prod_{n=1}^{0}(\ldots) = 1$. Assuming that $\mathbb{P}_1$ is absolutely continuous with respect to $\mathbb{P}_0$, the test

$$\{L_T \ge \gamma\} \qquad \text{where} \qquad L_T \triangleq \prod_{i=1}^{k} L_T(i) \text{ and } \gamma \in \mathbb{R} \tag{2}$$

is optimal in the (Neyman-Pearson) sense that if B is any test whose probability of false alarm $\mathbb{P}_0(B) \le \mathbb{P}_0(L_T \ge \gamma)$, then the probability of miss for the test (2) is at least as low as that for B, i.e., we have $\mathbb{P}_1(L_T < \gamma) \le \mathbb{P}_1(\Omega \setminus B)$.

In the context of the test (2), the decision is made at a single network node that receives all sensory information from the network, and processes it by computing the product L_T to issue the decision. This node is called the *fusion center*; see Figure 1(a). Before proceeding further, note that the only information needed from sensor $i \in \{1,2,\ldots,k\}$ includes the functions $\beta_i(\cdot)$ and $\nu_i(\cdot)$ together with the single real number $L_T(i)$. Put another way, once the problem parameters $\beta_i(\cdot)$, $\nu_i(\cdot)$ and the local likelihood ratios $L_T(i)$ are known, there is no increase in accuracy that can be obtained through knowledge of the sample path $t \mapsto N_t(i)$.

3 Deciding without a Fusion Center

One of the drawbacks of the setup of Figure1(a) described above is the vulnerability of the system to targeted attacks or failures: if the fusion center is lost, the entire detection system collapses. This motivates one to study the following problem: Suppose that at the terminal time T (or just after), there is some sharing of the $L_T(i)$'s (and the problem parameters $\beta_i(\cdot)$, $\nu_i(\cdot)$) among the sensors according to a directed graph, as depicted in Figure 1(b). Assume that no single sensor has access to all local likelihood ratios, i.e., there is no obvious choice of fusion center. If each sensor is now enabled to act as a DM operating on the information available to it, one can ask which DM is the most reliable in terms of decision-making accuracy.

To formulate this problem precisely, a directed graph is first specified. This graph encodes the allowed communication among the sensors. Recall that a directed graph is an ordered pair $G = (V,A)$ where V is a set of vertices or nodes, and A is a set of ordered pairs of nodes, referred to as arcs or directed edges. For the problem at hand, the vertex set $V = \{1,2,\dots,k\}$ indexes the set of observers (sensors). The arcs in A correspond to inter-sensor communication in the sense that $(i,j) \in A$ if, and only if, there is directional flow of information *from* sensor i *to* sensor j. Thus, if there is two-way communication between sensors i and j, both (i,j) and (j,i) are included in A. Self-loops are meaningless in this context and therefore excluded. Finally, it will be convenient to assume that information travels exactly one directed edge, and no further. In other words, if there are directed edges (i,j) and (j,ℓ) in A, it is assumed that $L_T(i)$ is sent from sensor i to sensor j and no further, while $L_T(j)$ is sent from sensor j to sensor ℓ and no further.[1]

The probabilistic setup and notation are exactly as in Section 2. The quantities $L_T(i)$ and L_T are defined as in (1), (2); and each sensor $i \in \{1,2,\dots,k\}$ observes $N_t(i)$ over time interval $t \in [0,T]$ and computes the quantity $L_T(i)$ at time T. Now that each sensor is armed with its own $L_T(\cdot)$, inter-sensor communication takes place. For each sensor i, let $\mathscr{S}_i$ comprise the set of sensors whose information is made available to sensor i just after time T. Thus, for $1 \le i \le k$,

$$\mathscr{S}_i \triangleq \{j \in \{1,2,\dots,k\} : \text{Sensor } i \text{ knows } L_T(j) \text{ and } \beta_j(\cdot),\, \nu_j(\cdot) \text{ just after time } T\} \ .$$

Since a sensor always has access to its own information, we have $i \in \mathscr{S}_i$ for all $1 \le i \le k$. Thus, $\mathscr{S}_i$ consists of the index i, together with the indices corresponding to *incoming* edges. Once inter-sensor communication has taken place, each sensor can be considered a DM. Thus, for $1 \le i \le k$, DM(i) refers to sensor i once it has access to the quantities $\{L_T(j), \beta_j(\cdot), \nu_j(\cdot) : j \in \mathscr{S}_i\}$. Letting DM$(i)$ use the test

$$\big\{\mathscr{L}_T(i) \ge \gamma_i\big\} \qquad \text{where} \qquad \mathscr{L}_T(i) \triangleq \prod_{j \in \mathscr{S}_i} L_T(j) \qquad \text{and } \gamma_i > 0$$

the probabilities of false alarm and miss for DM(i) are given by

[1] This entails no loss of generality; indeed, if information is to be sent from sensor i to sensor ℓ, this can be accommodated at the expense of introducing the additional directed edge (i,ℓ).

$$P_{F,i} \triangleq \mathbb{P}_0\left\{\mathscr{L}_T(i) \geq \gamma_i\right\} \ , \qquad P_{M,i} \triangleq \mathbb{P}_1\left\{\mathscr{L}_T(i) < \gamma_i\right\} \ ,$$

respectively. If function $\Lambda_i : \mathbb{R} \to \mathbb{R}$ is defined by $\Lambda_i(p) \triangleq \log \mathbb{E}_0\left[\left(\mathscr{L}_T(i)\right)^p\right]$ for $p \in \mathbb{R}$, then it follows [12, Theorem 8] that

$$P_{F,i} \leq \exp\left(\inf_{p>0}[\Lambda_i(p) - p\eta_i]\right) \ , \quad P_{M,i} \leq \exp\left(\inf_{p<1}[\Lambda_i(p) + (1-p)\eta_i]\right) \ , \quad (3)$$

where $\eta_i \triangleq \log \gamma_i \in \mathbb{R}$, and $\Lambda_i(p)$ is explicitly computable via

$$\Lambda_i(p) = \sum_{j \in \mathscr{S}_i} \int_0^T \left[\mu_j(s)^p - p\mu_j(s) + p - 1\right] \beta_j(s)\, \mathrm{d}s \ . \quad (4)$$

Note that the bounds (3), (4) hold for *any* $\eta_i = \log \gamma_i \in \mathbb{R}$. The reader is referred to [12] for a detailed proof of the bounds in (3).

In order to effectively use these bounds, one would need to know where, if at all, the infima in (3) are realized, and further, whether the infima are negative—to ensure that the bounds are non-trivial. One proceeds here by making repeated use of convexity of the functions $p \mapsto \Lambda_i(p)$. Indeed, it follows [12, Lemmas 16–19] that the function $p \mapsto \Lambda_i(p)$ is C^2 with derivatives given by differentiating under the integral sign; further, $\Lambda_i''(p) > 0$ for all $p \in \mathbb{R}$, implying that the function $p \mapsto \Lambda_i(p)$ is strictly convex; and finally that $\Lambda_i'(0) < 0$, $\Lambda_i'(1) > 0$. It now follows [12, Proposition 13] that if η_i is chosen to lie in $(\Lambda_i'(0), \Lambda_i'(1))$, then the infima in (3) are attained at a unique $p_i^* \in (0,1)$ and the infima are negative. More precisely, there exists a unique $p_i^* \in (0,1)$ given by $\Lambda_i'(p_i^*) = \eta_i$ such that

$$\inf_{p>0}[\Lambda_i(p) - p\eta_i] = \mathscr{E}_{F,i}(p_i^*) < 0 \ , \qquad \inf_{p<1}[\Lambda_i(p) + (1-p)\eta_i] = \mathscr{E}_{M,i}(p_i^*) < 0 \ ,$$

where the *error exponents* $\mathscr{E}_{F,i}(p)$, $\mathscr{E}_{M,i}(p)$ mapping $(0,1)$ to $\mathbb{R}$ are

$$\mathscr{E}_{F,i}(p) \triangleq \Lambda_i(p) - p\,\Lambda_i'(p) \quad \text{and} \quad \mathscr{E}_{M,i}(p) \triangleq \Lambda_i(p) + (1-p)\,\Lambda_i'(p) \ . \quad (5)$$

Thus, if $\eta_i \in (\Lambda_i'(0), \Lambda_i'(1))$, then the tightest error probability bounds for DM(i) are given by

$$P_{F,i} \leq \exp[\mathscr{E}_{F,i}(p_i^*)] < 1 \ , \qquad P_{M,i} \leq \exp[\mathscr{E}_{M,i}(p_i^*)] < 1 \ . \quad (6)$$

To rank the DMs in terms of their capacity to make accurate decisions, one *ideally* solves the following problem: Let $\alpha \in (0,1)$ (acceptable upper bound on the probability of false alarm) be given. Allowing each DM to choose its own threshold to comply with the constraints that $P_{F,i} \leq \alpha$ and $P_{M,i}$ is minimized, rank the nodes in increasing order of $P_{M,i}$. The node with the smallest $P_{M,i}$ is the best DM, at least for the particular α. The challenge, as noted earlier, is that the error probabilities $P_{F,i}$ and $P_{M,i}$ are not amenable to analytic computation. We therefore work with the

corresponding Chernoff bounds (6) and study the problem stated above with $P_{F,i}$ and $P_{M,i}$ replaced by the corresponding tightest upper bounds.

To solve this problem, the threshold selection algorithm [12, Proposition 14] is employed. The algorithm implies that if, for some $1 \leq i \leq k$, we have $\log\alpha > -\Lambda_i'(1)$, then there exists a unique $p_i^\dagger \in (0,1)$ which solves the equation

$$\mathscr{E}_{F,i}(p_i^\dagger) = \log\alpha \ . \tag{7}$$

Moreover, $p_i^\dagger$ minimizes $\mathscr{E}_{M,i}(p)$ over all $p \in (0,1)$ which satisfy $\mathscr{E}_{F,i}(p) \leq \log\alpha$. This minimum value of $\mathscr{E}_{M,i}(p)$ is given by

$$\mathscr{E}_{M,i}(p_i^\dagger) = \log\alpha + \Lambda_i'(p_i^\dagger) \ .$$

Thus, choosing $\eta_i = \Lambda_i'(p_i^\dagger)$, i.e. letting DM$(i)$ select the threshold $\gamma_i = \exp(\Lambda_i'(p_i^\dagger))$, yields

$$P_{F,i} \leq \alpha \ , \qquad P_{M,i} \leq \alpha \exp(\Lambda_i'(p_i^\dagger)) = \alpha\gamma_i \ .$$

To provide some insight on the foregoing expressions—see [12] for details—one can use convexity to show that on $(0,1)$, the functions $\mathscr{E}_{F,i}$, $\mathscr{E}_{M,i}$ are differentiable and negative, with $\mathscr{E}_{F,i}$ being strictly decreasing while $\mathscr{E}_{M,i}$ is strictly increasing. The threshold selection algorithm can be summarized as follows: Take the threshold $\gamma_i = \exp(\Lambda_i'(p))$ and select $p \in (0,1)$ as small as possible to make $\mathscr{E}_{M,i}$ as small as possible, while ensuring that the false alarm constraint is met, i.e., $\mathscr{E}_{F,i}(p) \leq \log\alpha$. The latter is possible only if $\log\alpha$ is strictly greater than the infimum of $\mathscr{E}_{F,i}$ on $(0,1)$, which is easily computed to be $-\Lambda_i'(1)$. Of course, one should now find $p_i^\dagger$ by solving (7).

One can now rank the DMs from most reliable to least reliable by ranking the quantities $\mathscr{E}_{M,i}(p_i^\dagger)$, $1 \leq i \leq k$ (or, equivalently, $\exp(\mathscr{E}_{M,i}(p_i^\dagger)) = \alpha\gamma_i$) from smallest to largest. These findings are summarized as follows.

Theorem 1. *For $\alpha \in (0,1)$ with $\log\alpha > \max_{1\leq i\leq k}(-\Lambda_i'(1))$, define the functions $\mathsf{M}_i(\alpha)$, $1 \leq i \leq k$, by*

$$\mathsf{M}_i(\alpha) \triangleq \exp(\Lambda_i'(p_i^\dagger)) \ , \tag{8}$$

where $p_i^\dagger \in (0,1)$ solves $\mathscr{E}_{F,i}(p_i^\dagger) = \log\alpha$; thus, $\mathsf{M}_i(\alpha)$ is the threshold at DM(i). *Then, $\mathsf{M}_i(\alpha)$ can be used as an index for decision-making accuracy in the following sense: If $(i_1, i_2, \ldots, i_k)$ is an ordering of $\{1,2,\ldots,k\}$ such that*

$$\mathsf{M}_{i_1}(\alpha) \leq \mathsf{M}_{i_2}(\alpha) \leq \cdots \leq \mathsf{M}_{i_k}(\alpha) \ ,$$

then $(i_1, i_2, \ldots, i_k)$ provides a ranking of the DM*'s from most accurate to least accurate, as measured by the Chernoff bounds.*

Remark 1. Note that the accuracy index $\mathsf{M}_i(\alpha)$ is tied to the *specific* $\alpha \in (0,1)$. This naturally prompts the question: for a given network and set of problem parameters, does the sensor which has the smallest $\mathsf{M}_i(\alpha)$ vary with α. It is also natural to ask, for a given network and set of problem parameters, whether the ranking based on

$\mathsf{M}_i(\alpha)$ coincides with the ranking based on minimizing $P_{M,i}$ subject to $P_{F,i} \leq \alpha$, i.e., do the Chernoff bounds *faithfully* capture the relative decision-making accuracies of various sensors. These questions are the subject of ongoing work.

4 Examples

This section presents two examples of networks of nuclear detectors deciding about the observed process. In the first case, it is assumed that the process at each sensor has identical characteristics, and heterogeneity is introduced only through the network topology. In the second example, in addition to the differences among sensors due to their location in the underlying interconnection graph, heterogeneity is introduced in the quality of each sensor's observations. The objective is to highlight the interplay between the network topology and the quality of individual observations, and its impact on the ability of individual sensors to make decisions.

4.1 Example 1: Identical measurement characteristics

This example addresses the following question. If the measurement quality characteristics at all sensors are the same, is it true that the DM with the most $L_T(\cdot)$'s is the most accurate? The answer, as will be seen below, is affirmative, implying that the node with the largest in-degree is better equipped to make a decision.

Indeed, suppose that $\beta_i \equiv \beta$, $\nu_i \equiv \nu$, $\mu_i \equiv \mu = 1+\nu/\beta$. We will use $|\mathscr{S}_i|$ to denote the cardinality of $\mathscr{S}_i$, $1 \leq i \leq k$, that is, the in-degree of node i that corresponds to the number of incoming edges to i. It is easily seen from (4), (5), that $\mathscr{E}_{F,i}(p) = -g(p)\cdot|\mathscr{S}_i|$ where $g(p) \triangleq -\beta T[\mu^p - p\mu^p \log\mu - 1]$ and $\mathscr{E}_{M,i}(p) = -h(p)\cdot|\mathscr{S}_i|$ where $h(p) \triangleq -\beta T[\mu^p + (1-p)\mu^p\log\mu - \mu]$. Since $\mathscr{E}_{F,i}(p)$, $\mathscr{E}_{M,i}(p)$ are negative for $p \in (0,1)$ with the former being strictly decreasing and the latter strictly increasing (see [12, Lemma 19]), it follows that $g(p)$ is positive and strictly increasing for $p \in (0,1)$, while $h(p)$ is positive and strictly decreasing for $p \in (0,1)$.

Fix α as in Theorem 1. Let $i, j \in \{1,2,\ldots,k\}$ with $i \neq j$. The claim is that

1. If $|\mathscr{S}_i| > |\mathscr{S}_j|$, then $\mathsf{M}_i(\alpha) < \mathsf{M}_j(\alpha)$,
2. If $|\mathscr{S}_i| = |\mathscr{S}_j|$, then $\mathsf{M}_i(\alpha) = \mathsf{M}_j(\alpha)$,
3. If $|\mathscr{S}_i| < |\mathscr{S}_j|$, then $\mathsf{M}_i(\alpha) > \mathsf{M}_j(\alpha)$.

The first of these three claims will only be proven here; the arguments can be easily modified to prove the other two. Suppose $|\mathscr{S}_i| > |\mathscr{S}_j|$. By (7), one has $\mathscr{E}_{F,i}(p_i^\dagger) = \log\alpha$ and $\mathscr{E}_{F,j}(p_j^\dagger) = \log\alpha$. Recalling the notation above,

$$g(p_i^\dagger) = -\frac{\log\alpha}{|\mathscr{S}_i|} \quad \text{and} \quad g(p_j^\dagger) = -\frac{\log\alpha}{|\mathscr{S}_j|} \ .$$

Noting that $-\log\alpha > 0$ and $|\mathscr{S}_i| > |\mathscr{S}_j|$, one writes $g(p_i^\dagger) < g(p_j^\dagger)$. Since g is strictly increasing, it now follows that $p_i^\dagger < p_j^\dagger$. If it can be shown that $\mathscr{E}_{M,i}(p_i^\dagger) < \mathscr{E}_{M,j}(p_j^\dagger)$, then it follows (see Theorem 1) that $\mathsf{M}_i(\alpha) < \mathsf{M}_j(\alpha)$. The argument for establishing

the former statement is follows. Since the function h is strictly decreasing, it must be $h(p_i^\dagger) > h(p_j^\dagger)$. Since h is positive and $|\mathscr{S}_i| > |\mathscr{S}_j|$ (by assumption), $|\mathscr{S}_i| \cdot h(p_i^\dagger) > |\mathscr{S}_j| \cdot h(p_j^\dagger)$. Taking negatives, yields $\mathscr{E}_{M,i}(p_i^\dagger) < \mathscr{E}_{M,j}(p_j^\dagger)$ as required.

This example shows that, under the assumption that the measurement process has identical statistical characteristics at each node, the nodes with the largest in-degree are the most accurate DMs. It is interesting to note that this observation is in contrast to the diffusive models studied in [14, 15], in which local centrality measures such as those based on nodal degrees cannot capture the certainty of each unit in terms of the collected evidence.

4.2 Example 2: Non-identical measurement characteristics

Particular examples of directed graphs in the setting of Figure 1(b) help examine the effect of heterogeneity not only in the underlying communication topology, but also in the characteristics of the measurements process at each sensor. Similarly to [9], five sensors are arranged along a straight line at fixed locations with 0.5 m distance from each other and with the first sensor positioned at 0.5 m. A radioactive source moves parallel to this array of sensors with a constant speed of 0.03 m/s. As a result of the source's motion, the measurement characteristics among different sensors are not the same, since certain sensors may spend longer intervals in close proximity to the source than others. The source's straight line path is at a distance of 0.5 m above the array of the sensors, and the initial source position is 0.5 m behind the sensor array. The activity of the source and background radiation are measured in gamma rays emitted per second, i.e., counts per second (cps). For the source, it is taken at $a = 3$ cps, while for the background we assume 0.167 cps. The numerical simulation of the corresponding arrival process with the aforementioned intensities has been generated using a thinning algorithm [7, 13]. The maximum acceptable probability of false alarm is taken as $\alpha = 10^{-3}$ and for numerical purposes the sensor cross-section is assumed to be 1 m^2. Note that these figures are too big for practical applications, and are used here only for reasons of numerical convenience to emphasize the differences in sensor performance; see [12] for details.

Figure 2 depicts three different interconnection topologies with which locally processed information can be disseminated among sensors. Directed edges mark unidirectional flow of likelihood ratios L_T and the histories of the intensities β, ν between (network) adjacent sensors at time T. These three communication topologies (graphs) result in different sensor network performance in terms of accuracy of decision making, at least to the degree that the latter can be reflected on the computed bounds on the probability of missed detection. Various detection scenarios are depicted in Figures 3-4, showing that for a constant background activity, the performance of each sensor as a decision-making unit depends on the strength of the signals perceived by the sensors and the time available to make the decision. In these figures, individual sensors are denoted by S1 to S5, and the motion of the source is indicated by the blue line at the top of each figure. Different decision times T are examined, resulting to different time intervals over which the source remains in the close vicinity of each sensor, thereby introducing heterogeneity in the mea-

surement characteristics of each sensor. In each of the Figures 3-4, the vertical axis corresponds to the logarithm of the bound of the probability of missed detection computed by (6) and the horizontal axis shows the distance between sensors.

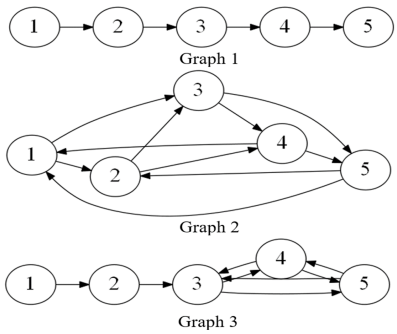

Fig. 2 Different communication topologies studied in Example 2. A directed edge from node i to node j indicates that at time T, the locally computed L_T and μ at node i is transmitted to node j.

Figure 3 shows that the performance of the sensors under more information sharing (graph 2 in Figure 2) is better than their performance under less information sharing (graph 1 in Figure 2). A major advantage of exchanging information among individual sensors is evident from the performance of node 3 in Figure 3 (left), which corresponds to graph 3 of Figure 2. Indeed, due to the larger distance from the source, sensor 3 neither makes good quality measurements nor it receives signals from sensor 1; yet it performs better than 1 or 2 because of its location in the graph. However, more incoming edges to a sensor does not always result into better performance. To see this, Figure 3 (right) describes network performance when the spatial arrangement is scaled up by 2 such that all the distances are doubled (1m instead of 0.5m). Doubling the distances between sensors, results in weakening the observed signal, in the sense that the intensity of the arrival process perceived by each sensor decreases. In this case of larger inter-sensor distance, more sensors do not satisfy the constraint on the probability of false alarm at the end of the decision interval $T = 40$ seconds, and the additional information available to sensor 3 due to its location in graph 3 does not compensate for the overall weak quality of the signals received by all the sensors.

The dependency of node ranking on the available decision time is illustrated in Figure 4 (left) for the case of graph 1. It can be seen that as more time is allowed to make the decision, the location of the best performing sensor shifts in the direction of the motion of the source. This observation reflects the intuitive fact that the sensor that remains closer than the rest to the source during the detection time window performs better. As a final remark it should be mentioned that the aforementioned observations on node ranking in terms of their decision-making capability are based on the Chernoff bounds (6). To compare the prediction of the bounds with that obtained by the actual probabilities, Figure 4 (right) depicts an estimate for the logarithm of the probability of missed detection obtained using Monte-Carlo simulations. In this particular example, the source intensity is assumed to be 1.75

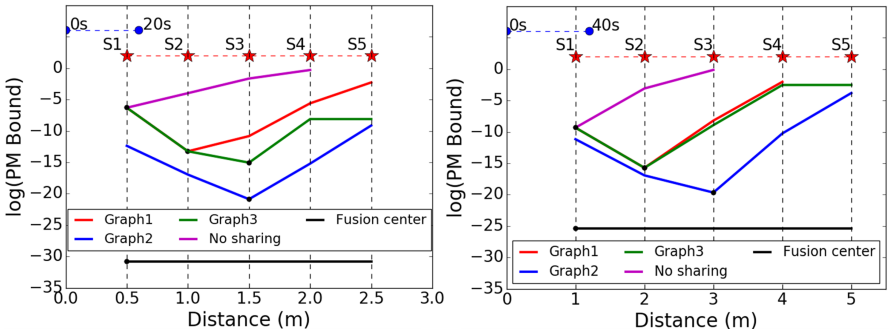

Fig. 3 The communication topologies of Figure 2 result in different bounds for the probability of missed detection for each sensor S1 to S5. These are bounded by the two extreme cases: no information is shared (magenta curve), and fusion center having data from all sensors (black curve). Missing line segments indicate that the corresponding sensors do not satisfy the constraint on the probability of false alarm.

cps and the network topology corresponds to graph 1 of Figure 2. As can be seen, the prediction of the bounds is the same with that of the estimated probabilities of missed detection. It is important to note, however, that although the general trend of actual error probabilities is captured by the bounds [12], one must be aware that the loss of sharpness associated with the use of bounds may result in node rankings that do not faithfully reflect the actual ranking; for example, this is the case when neighboring sensors exhibit comparatively similar decision-making performance in a way that cannot be differentiated by the bounds.

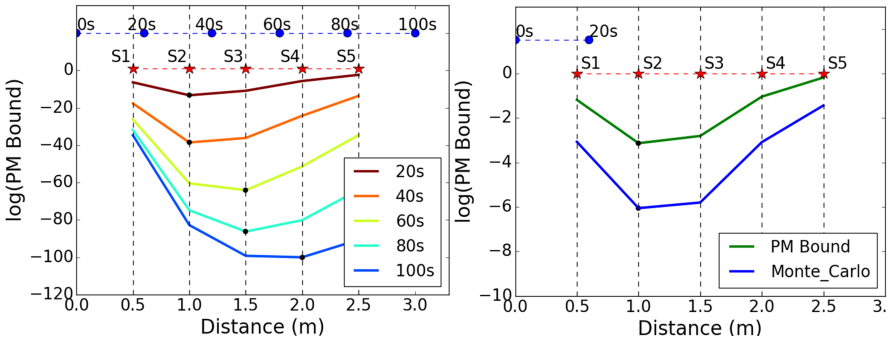

Fig. 4 The left plot shows the logarithm of the bound on the probability on missed detection with decision time T using graph 1 of Figure 2. The right plot shows Monte-Carlo estimates of the probability of missed detection using graph 1 of Figure 2.

5 Conclusions

The paper argues that the decision-making accuracy of individual nodes in a network of radiation detectors can be quantified using explicitly computed Chernoff upper bounds as proxies for true error probabilities. The resulting ranking of nodes can be used to choose which node to make the final decision in a distributed setup

where each individual node can act as a potential decision maker. Examples show how the capability of each sensor to make a decision about the observed process depends on the interplay between the location of the sensor in the underlying network architecture (reflecting the quantity of information available to the sensor) and the spatial location of the sensors with respect to the source (reflecting the quality of individual sensor measurements). Performance comparable to that of decision making with a fusion center can be achieved by allowing partial information sharing between the nodes over some fixed, directional communication topology. This paper lays the foundation for future work in which the network performance is optimized by reconfiguring the underlying communication topology (cf. [1]) given the number of sensors, the strength of the signal, and the time available to make the decision.

Acknowledgment

This work is supported in part by DTRA under award #HDTRA1-16-1-0039.

References

1. Bhambhani, V., Valbuena, L., Tanner, H.G.: Spatially distributed cellular neural networks. International Journal of Intelligent Computing and Cybernetics **4**(4), 465–486 (2011)
2. Bogacz, R., Brown, E., Moehlis, J., Holmes, P., Cohen, J.D.: The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychological Review **113**(4), 700–765 (2006)
3. Brémaud, P.: Point Processes and Queues. Martingale Dynamics. Springer-Verlag (1981)
4. Byrd, R., Moss, J., Priedhorsky, W., Pura, C., Richter, G.W., Saeger, K., Scarlett, W., Scott, S., Jr., R.W.: Nuclear detection to prevent or defeat clandestine nuclear attack. IEEE Sensors Journal **5**(4), 593–609 (2005)
5. Fitch, K., Leonard, N.E.: Information centrality and optimal leader selection in noisy networks. In: IEEE Intenrational Conference on Decision and Control, pp. 7510–7515 (2013)
6. González, F.I., Bernard, E.N., Meinig, C., Eble, M.C., Mofjeld, H.O., Stalin, S.: The nthmp tsunameter network. Natural Hazards **35**, 25–39 (2005)
7. Lewis, P.A.W., Shedler, G.S.: Simulation of nonhomogeneous Poisson processes by thinning. Naval Research Logistics Naval Research Logistics Quarterly **26**(3), 403–413 (1979)
8. Morreale, P., Qi, F., Croft, P.: A green wireless sensor network for environmental monitoring and risk identification. International Journal of Sensor Networks **10**(1-2), 73–82 (2011)
9. Nemzek, R.J., Dreicer, J.S., Torney, D.C., Warnock, T.T.: Distributed sensor networks for detection of mobile radioactive sources. IEEE Transactions on Nuclear Science **51**(4), 1693–1700 (2004)
10. Ogata, Y.: Seismicity analysis through point-process modeling: A review. Pure and Applied Geophysics **155** (1999)
11. Pahlajani, C.D., Poulakakis, I., Tanner, H.G.: Networked decision making for Poisson processes with applications to nuclear detection. IEEE Transactions on Automatic Control **59**(1), 193–198 (2014)
12. Pahlajani, C.D., Sun, J., Poulakakis, I., Tanner, H.G.: Error probability bounds for nuclear detection: Improving accuracy through controlled mobility. Automatica **50**(10), 2470–2481 (2014)
13. Pasupathy, R.: Generating nonhomogeneous Poisson processes. In: Wiley Encyclopedia of Operations Research and Management Science. Wiley (2009)
14. Poulakakis, I., Scardovi, L., Leonard, N.E.: Node classification in collective evidence accumulation toward a decision. In: Proceedings of the IEEE International Conference on Decision and Control (2012)

15. Poulakakis, I., Young, G.F., Scardovi, L., Leonard, N.E.: Information centrality and ordering of nodes for accuracy in noisy decision-making networks. IEEE Transactions on Automatic Control **61**(4), 1040–1046 (2016)
16. Srikrishna, D., Chari, A.N., Tisch, T.: Deterance of nuclear terrorism with mobile radiation detectors. The Nonproliferation Review **12**(3), 573–614 (2005)
17. Stephenson, K., Zelen, M.: Rethinking centrality: Methods and examples. Social Networks **11**(1), 1–37 (1989)
18. Sundaresan, A., Varshney, P.K., Rao, N.S.V.: Distributed detection of a nuclear radioactive source using fusion of correlated decisions. In: Proceedings of the International Conference on Information Fusion, pp. 1–7. IEEE (2007)
19. Tenney, R.R., Sandell Jr., N.R.: Detection with Distributed Sensors. IEEE Transactions on Aerospace and Electronic Systems **17**(4), 501–510 (1981)
20. Tsitsiklis, J., Athans, M.: On the complexity of distributed decision problems. IEEE Transactions on Automatic Control **AC-30**, 440–446 (1985)
21. Tsitsiklis, J.N.: Decentralized detection. Advances in Statistical Signal Processing **2**, 297–344 (1993)
22. Varshney, P.K.: Distributed Detections and Data Fusion. Springer-Verlag, New York (1997)
23. Viswanathan, R., Varshney, P.K.: Distributed detection with multiple sensors: Part I - Fundamentals. Proceedings of the IEEE **85**(1), 54–63 (1997)