Decentralized Decision-Making for Multi-Agent Networks: the State-Dependent Case

Joni Shaska
University of Southern California
shaska@usc.edu

Urbashi Mitra
University of Southern California
ubli@usc.edu

***Abstract*—We consider a new formulation of the decentralized detection problem with parallel agent configuration. In particular, each agent in the network exists in a set of pre-specified states that affects the distribution of their observations as well as the underlying hypothesis. As such, observations are conditionally dependent. Following a person-by-person design methodology, it is shown that the Bayes optimal detection rule for each agent is a likelihood ratio test with a state dependent threshold. Moreover, it is shown that even for statistically identical agents, the optimal rules for the agents may not be the same. Motivated by this, we turn our attention to large networks and find the error exponent, and show that as the number of agents increases there is no loss of asymptotic optimality if the agents use the same rule, dramatically reducing the complexity of computing the decision rules for each agent.**

***Index Terms*—decentralized detection, multi-agent networks, composite hypothesis testing, state-dependent networks, error exponents**

I. Introduction

Decentralized detection in wireless networks has been persistently studied over the years [1], [2]. Despite its long history, the problem remains of interest [3]–[5] as different contexts are considered. In particular, this paper addressed the generalization of decision making wherein each agent observes signals due to a common, unknown hypothesis, but each agent is affected by their individual state. This is a form of composite hypothesis testing. The key challenge is that the introduction of state results in observations being correlated through the state process.

Two applications that fit within this general framework are the presence of anomalous sensors in sensor networks and collective decision making in microbial communities. In the sensor network application, sensors may or may not be functioning properly and this will affect the quality of the observations. We note that our framework removes the need for active anomaly detection [6]–[9] first. Coping with anomalies remains a relevant challenge for modern Internet-of-Things applications (see *e.g.* [10]).

A phenomenon of great interest in microbiology is quorum sensing [11], [12]. In this scenario, bacteria synthesize key molecules that are released into the environment and are also sensed by the bacteria. When the concentration of the molecule exceeds a certain level, new genes are expressed enabling novel collective behaviors such as bacterial infections or the forming of biofilms. However, the entire colony does not express new genes simultaneously and the signals they can transmit and receive are affected by whether they have expressed the gene or not. Thus, the state of each individual bacterium affects the signaling. Quorum sensing has been previously modeled as a decentralized decision making process [13], [14].

Decentralized detection has been investigated under different objective functions and different signal models [2], [15]–[17]. In this paper, we focus on the Bayesian setting wherein each hypothesis has a prior probability and the underlying distributions under all hypothesis-state pairs are known. Under certain assumptions, one can solve decentralized detection problems efficiently. One such assumption is that observations received by different agents are assumed to be *conditionally independent*. This assumption helps reduce complexity and simplifies analysis. Unfortunately, if one removes the conditional independence assumption, as we do here, the problem of optimal decentralized Bayesian detection becomes NP-hard in general [1], [18]. Nevertheless, the problem of optimal decentralized detection with dependent observations [17], [19] has been previously examined. In particular, a hierarchical conditional independence (HCI) model was introduced in [17] for which analysis of the optimal decision making was enabled. We generalize this notion herein.

The contributions of this paper are:

1) The state-dependent decentralized detection problem is formulated.
2) The optimal Bayesian decision rule at each agent is described for networks whose topology is similar to that of Fig. 1, and conditioned on the strategy employed at agent, the *person-by-person* optimal (PBPO) decision rule for a fusion center is also provided.
3) It is shown that the optimal agent rule is *state-dependent* and an example is provided to underscore that even for networks of *identical* agents, common rules across all agents are not optimal for the case of a small set of agents.
4) An asymptotic analysis, as the number of agents goes to infinity, is conducted of the probability of detection error at the fusion center under the assumption of statistically independent states. It is shown that, asymptotically, a common decision rule for each agent is in fact optimal.

This work has been funded in part by one or more of the following grants: ONR 00014-15-1-2550, NSF CCF-1817200, ARO W911NF1910269, Cisco Foundation 1980393, DOE DE-SC0021417, Swedish Research Council 2018-04359, NSF CCF-2008927, ONR 503400-78050.

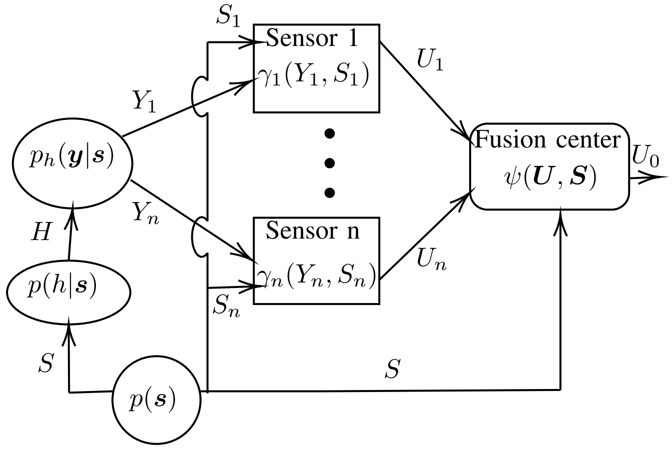

Fig. 1. State-dependent decentralized detection framework.

The new problem framework results in conditionally dependent observations, necessitating new proof strategies over the current state-of-the-art.

5) While finding the optimal set of rules is often intractable, even for small networks, analysis of the asymptotics shows that the optimal common rule can be found by solving a single optimization for *one* agent.
6) Numerical results are provided to show that the asymptotically optimal decision rule is often quite tractable to compute.

The rest of this paper is organized as follows. Section II describes the problem in terms of general agent networks. In Section III we explain our main results. We show that for our problem the PBPO solution is a state-dependent likelihood ratio test. Moreover, for large networks of identical agents, it is asymptotically optimal for each agent to use the same rule. Section IV goes over a specific example, and Section V concludes the paper. Due to space constraints several key derivations and proofs are omitted – this material can be found in [20].

Notation: Random variables are written as capital letters X and realizations as lower case letters x. Vectors are denoted with bold face letters with random vectors denoted as capital bold face letters $\boldsymbol{V}$ and their realizations as lower case $\boldsymbol{v}$. $\boldsymbol{V}^k$ denotes the vector $\boldsymbol{V} \setminus V_k = [V_1, V_2, ..., V_{k-1}, V_{k+1}, ..., V_N]$, i.e., the $\boldsymbol{V}$ vector with the kth term removed. $p(x)$ denotes either the probability mass function (pmf) of a discrete random variable X or the probability density function (pdf) of a continuous random variable X. $p(x, y)$ and $p(x|y)$ denote respectively the joint and conditional pmf or pdf of random variables X and Y.

II. State-Dependent Decentralized Detection

In this section, we introduce the notion of *state-dependent*, decentralized hypothesis testing. A fusion center is connected to n agents via a star network topology as depicted in Fig. 1. We assume that each agent is in one of l states: $S_k \in \{0, 1, ..., l-1\}$, $k = 1, 2, ..., n$. The state vector $\boldsymbol{S}$ has a prior probability $p(\boldsymbol{s})$. The goal of the fusion center is to assess which of m possible hypotheses is true, where $H \in \{0, 1, ..., m-1\}$ with conditional probability $p(h|\boldsymbol{s})$. All agents observe the same underlying hypothesis. Each agent makes a local observation $Y_k \in \mathcal{Y}$, $k = 1, ..., n$ and determines a local decision, $U_k = \gamma_k(Y_k, S_k) \in \mathcal{U} = \{0, 1, ..., b-1\}$. These local decisions are then sent to the fusion center along with the system state for the final decision. The fusion center output is given by $U_0 = \psi(\boldsymbol{U}, \boldsymbol{S}) \in \{0, 1, ..., m-1\}$.

Let the set Γ be the set of all decision rules and Ψ be the set of all fusion rules. We call a collection of decision rules $\gamma_k \in \Gamma$, $k = 1, 2, ..., n$, together with a fusion rule $\psi \in \Psi$ a *strategy* denoted by $\boldsymbol{\psi} \in \Gamma^n \times \Psi$, where Γ^n is the Cartesian product of Γ with itself n times. We denote the conditional pmf (pdf) of a random variable X conditioned on $H = h$, $h \in \{0, 1, ..., m-1\}$, as $p_h(x)$. Also, the pmf (pdf) of X conditioned on $Y = y$ and $H = h$ is denoted as $p_h(x|y)$. Let $c_{u_0,h}$ be the Bayesian cost of deciding $U_0 = u_0$ when $H = h$ is true. We wish to find the strategy $\boldsymbol{\psi}$ that minimizes the expected Bayesian cost J given by

$$J(\boldsymbol{\psi}) = \sum_{u_0=0}^{m-1} \sum_{h=0}^{m-1} c_{u_0,h} p(u_0, h). \tag{1}$$

Note that each agent k, $k = 1, 2, ..., n$, receives the observation (Y_k, S_k), and that, in general, these observations are not conditionally independent. Thus, the scope of the paper goes beyond conditionally independent observations in contrast to other works, [15], [21], [22]. We, however, make several assumptions on the relationships between H, $\boldsymbol{S}$, $\boldsymbol{Y}$, $\boldsymbol{U}$, and U_0.

- Conditioned on S_k and H, Y_k is independent of $\boldsymbol{Y}^k$ and $\boldsymbol{S}^k$, i.e., the joint pmf factors as follows,

$$p(\boldsymbol{y}, \boldsymbol{s}, h) = [\prod_{k=1}^{n} p_h(y_k|s_k)] p(h|\boldsymbol{s}) p(\boldsymbol{s}). \tag{2}$$

- $\boldsymbol{U}$ is a function only of $\boldsymbol{Y}$ and $\boldsymbol{S}$, with U_k being a function only of Y_k and S_k, i.e., the joint conditional pmf is given as

$$p(\boldsymbol{u}|\boldsymbol{y}, \boldsymbol{s}) = \prod_{k=1}^{n} p(u_k|y_k, s_k) \tag{3}$$

- U_0 is a function only of $\boldsymbol{U}$ and $\boldsymbol{S}$. Note that this assumptions says the fusion center has knowledge of the state as depicted in Fig 1.

Furthermore, given the structure of our problem, we have the following Markov chain relationships: $\boldsymbol{S} \to (\boldsymbol{S}, H) \to (\boldsymbol{Y}, \boldsymbol{S}) \to (\boldsymbol{U}, \boldsymbol{S}) \to U_0$. Similarly, we have the following hierarchical conditional independence (HCI) model induced by the following Markov chains: $H \to (\boldsymbol{S}, H) \to (\boldsymbol{Y}, \boldsymbol{S}) \longrightarrow (U, \boldsymbol{S}) \to U_0$. The HCI model coupled with the additional structure of (2) will facilitate analysis in the sequel despite the lack of conditional independence in our observations.

III. Optimal Agent Rules

We present our main results in this section. First, we present the PBPO agent rules and fusion rule for any n. We show that regardless of the rules the other agents are using, agent k, $k = 1, 2, ..., n$, should use a state-dependent likelihood ratio test (SDLRT). This implies that the optimal strategy is in fact to have every agent use a state-dependent likelihood ratio test. This state dependence challenges analysis. The thresholds for these tests can be found by solving a set of nl non-linear coupled equations. Clearly, the larger the number of agents, the more computationally challenging the rule design.

Consistent with PBPO rules, satisfying the equations is a *necessary* condition for global optimality for a set of thresholds, but not *sufficient*. As the computation of individual rules for each agent is challenging (we provide an example in the sequel to show that the rules are, in general, not the same for all agents), we examine behavior in the case of large n. To this end, we derive the error exponent, and use the computed exponent to show that under the assumption that agent states are independent, it is optimal for all agents to employ the same rule in the limit of large n. While our result has similarity to that of [21] for conditionally independent observations, there are key differences in the analysis highlighted in the sequel and in [20]. Without loss of generality, and for clarity, we focus on binary hypotheses.

A. Person-by-Person Optimal Agent Design for n agents

Let $m = 2$ and $b = 2$. The PBPO decision rule $\gamma_k(y_k, s_k)$ for agent k when $Y_k = y_k$ and $S_k = s_k$ is given as $\gamma_k(y_k, s_k) = 1$ if

$$\frac{p_1(y_k|s_k)}{p_0(y_k|s_k)} > \frac{[g_f(s_k,1) - g_f(s_k,0)]p(h=0|s_k)}{[g_d(s_k,1) - g_d(s_k,0)]p(h=1|s_k)} \tag{4}$$

and $\gamma_k(Y_k, S_k) = 0$ otherwise, where

$$g_f(s_k, u_k) = c_f p_0(U_0 = 1|s_k, u_k) \tag{5a}$$

$$g_d(s_k, u_k) = c_d p_1(U_0 = 1|s_k, u_k) \tag{5b}$$

$$c_f = c_{1,0} - c_{0,0} \tag{6a}$$

$$c_d = c_{0,1} - c_{1,1} \tag{6b}$$

and

$$g_d(s_k, 1) - g_d(s_k, 0) > 0 \tag{7}$$

is satisfied.

For the rest of the paper, we are concerned with minimizing the probability of error, i.e, $c_{1,0} = c_{0,1} = 1$ and $c_{0,0} = c_{1,1} = 0$. Note that the above rule is *person-by-person* optimal, namely that if $\boldsymbol{\psi}^*$ is a PBPO solution, then any other strategy $\boldsymbol{\psi}$ that changes *some*, but not *all*, of the agents rules must satisfy $J_n(\boldsymbol{\psi}^*) \leq J_n(\boldsymbol{\psi})$. That is, in order to achieve a better cost all of the γ_ks must be varied together. As stated before, this means that a PBPO strategy $\boldsymbol{\psi}^*$ need not be globally optimal.

One can show that the PBPO rule for the fusion center to minimize the probability of error when the agents decide $\boldsymbol{u}$ and the state is $\boldsymbol{s}$ is $\psi(\boldsymbol{u}, \boldsymbol{s}) = 1$ when

$$\frac{p(h=1|\boldsymbol{u},\boldsymbol{s})}{p(h=0|\boldsymbol{u},\boldsymbol{s})} > 1, \tag{8}$$

and $\psi(\boldsymbol{u}, \boldsymbol{s}) = 0$ otherwise. Equation (8) is simply just the maximum *a posteriori* (MAP) rule. We observe that the decision regions generated by the fusion center are dependent on the rules used by the agents. To this end, we let $\boldsymbol{\gamma}$ be the collection of the agents decision rules. For the rest of the paper, given a specific collection of the agents decision rules $\boldsymbol{\gamma} \in \Gamma^n$, and fixed n, we denote the probability of error as $J_n(\boldsymbol{\gamma})$.

B. Error Exponent

As stated before, our goal is to minimize $J_n(\boldsymbol{\gamma})$ over all $\boldsymbol{\gamma} \in \Gamma^n$. Letting A_0 and A_1 be the (disjoint) sets of pairs $(\boldsymbol{u}, \boldsymbol{s})$, where the fusion center decides 0 and 1, respectively, we can write $J_n(\boldsymbol{\gamma})$ as,

$$\begin{aligned} J_n(\boldsymbol{\gamma}) = &\sum_{\boldsymbol{s},\boldsymbol{u} \in A_0^C} p_0(\boldsymbol{u}|\boldsymbol{s})p(h=0|\boldsymbol{s})p(\boldsymbol{s}) + \\ &\sum_{\boldsymbol{s},\boldsymbol{u} \in A_1^C} p_1(\boldsymbol{u}|\boldsymbol{s})p(h=1|\boldsymbol{s})p(\boldsymbol{s}). \end{aligned} \tag{9}$$

It should be noted that for both summations in (9) we are only summing over those $(\boldsymbol{u}, \boldsymbol{s})$ in A_h^C such that $p(\boldsymbol{u}, \boldsymbol{s}, h) \neq 0$ for both $h = 0, 1$. That is, we do not condition on events that have measure zero. We define the following key sequence:

$$\alpha_n = \min_{\boldsymbol{s}, h \in \{0,1\}} p(h|\boldsymbol{s}), \tag{10}$$

that is, α_n is the smallest $p(h|\boldsymbol{s})$, $h = 0, 1$, appearing in either summation. Notice that since the fusion center is implementing the *maximum a posteriori* (MAP) rule, we can rewrite (9) as

$$\begin{aligned} &J_n(\boldsymbol{\gamma}) = \\ &\sum_{\boldsymbol{s},\boldsymbol{u}} \min\{p_0(\boldsymbol{u}|\boldsymbol{s})p(h=0|\boldsymbol{s}), p_1(\boldsymbol{u}|\boldsymbol{s})p(h=1|\boldsymbol{s})\}p(\boldsymbol{s}) \end{aligned} \tag{11}$$

Since we are concerned with large n, we focus on the error exponent defined as

$$\lim_{n \to \infty} \frac{1}{n} \log J_n(\boldsymbol{\gamma}). \tag{12}$$

Then, if we let $r_n(\boldsymbol{\gamma}) = \frac{1}{n} \log J_n(\boldsymbol{\gamma})$ and $R_n = \inf_{\boldsymbol{\gamma} \in \Gamma^n} r_n(\boldsymbol{\gamma})$ for all n, we analyze the limiing behavior of R_n. Similar to [22], we would like to derive upper and lower bounds on $r_n(\boldsymbol{\gamma})$. Unfortunately, we cannot use the bounds derived in [22] and [21] since our observations are not conditionally independent. Thus, we need the following which is proved in [20].

Lemma 1. *For any n and $\boldsymbol{\gamma} \in \Gamma^n$, we have*

$$\begin{aligned} \frac{\log \alpha_n}{n} - \frac{\log 2}{n} + \frac{1}{n}\mu(\boldsymbol{\gamma}, \epsilon^*) - \frac{\sqrt{2\mu''(\boldsymbol{\gamma}, \epsilon^*)}}{n} \\ \leq r_n(\boldsymbol{\gamma}) \leq \frac{1}{n}\mu(\boldsymbol{\gamma}, \epsilon^*), \end{aligned} \tag{13}$$

where for $\epsilon \in (0, 1)$ we define [1]

$$\mu(\boldsymbol{\gamma}, \epsilon) = \log \left[\sum_{\boldsymbol{s}} \sum_{\boldsymbol{u}} (p_0(\boldsymbol{u}|\boldsymbol{s}))^{1-\epsilon} (p_1(\boldsymbol{u}|\boldsymbol{s}))^{\epsilon} p(\boldsymbol{s}) \right], \tag{14}$$

[1] We also extend the definition to include the cases $\epsilon = 0$ and $\epsilon = 1$, with $\mu(\boldsymbol{\gamma}, 0) = \lim_{\epsilon \to 0^+} \mu(\boldsymbol{\gamma}, \epsilon)$; $\mu(\boldsymbol{\gamma}, 1) = \lim_{\epsilon \to 1^-} \mu(\boldsymbol{\gamma}, \epsilon)$.

α_n *is defined in* (10), $\mu''(\gamma, \epsilon)$ *is the second derivative of* $\mu(\gamma, \epsilon)$ *with respect to* ϵ, *and* $\epsilon^* = \arg\min_{\epsilon\in[0,1]} \mu(\gamma, \epsilon)$.

The dependence on γ is captured in u, since different rules used by the agents will change the statistics of u. A proof of the upper bound on the exponent is provided in the Appendix. The lower bound can be found by observing that

$$\begin{aligned}\mu'(\gamma, \epsilon) &= \sum_{u,s} Q_\epsilon(u, s) L(u, s)\\ \mu''(\gamma, \epsilon) &= \left\{\sum_{u,s} Q_\epsilon(u, s)(L(u, s))^2\right\} - (\mu'(\gamma, \epsilon))^2,\end{aligned}$$

where all derivatives are with respect to ϵ and $L(u, s)$ is the log likelihood ratio defined as

$$L(u, s) = \log \frac{p_1(u|s)}{p_0(u|s)}.$$

For $\epsilon \in (0, 1)$, define

$$Q_\epsilon(u, s) = \frac{p_0(u|s)^{1-\epsilon} p_1(u|s)^\epsilon p(s)}{\sum_{u',s'} (p_0(u'|s'))^{1-\epsilon} p_1(u'|s')^\epsilon p(s')}.$$

We can interpret $Q_\epsilon(u, s)$ as a probability distribution on $L(u, s)$. Hence, $\mu'(\gamma, \epsilon)$ and $\mu''(\gamma, \epsilon)$ are the mean and variance of $L(u, s)$, respectively, according to $Q_\epsilon(u, s)$. Moreover, one can show

$$\begin{aligned}p_0(u|s)p(s) &= \{\exp[\mu(\gamma, \epsilon) - \epsilon L(u, s)]\} Q_\epsilon(u, s)\\ p_1(u|s)p(s) &= \{\exp[\mu(\gamma, \epsilon) + (1 - \epsilon)L(u, s)]\} Q_\epsilon(u, s).\end{aligned}$$

The lower bound then follows by first noticing that

$$J_n(\gamma) \geq \alpha_n \sum_{s,u} \min\{p_0(u|s), p_1(u|s)\} p(s),$$

and then considering only those (u, s) that are within two standard deviations from the mean of $L(u, s)$ according to $Q_\epsilon(u, s)$. Note that if $\frac{1}{n} \log \alpha_n$ and $\mu''(\gamma, \epsilon)$ are not properly controlled, then the bounds given in (13) could be far apart even for large n. To tighten the bounds for large n, we introduce the following assumptions:

Assumption 1. *For all* n, $\gamma \in \Gamma^n$, *and* $\epsilon \in [0, 1]$*:*

a) $|\mu(\gamma, \epsilon)| < \infty$.
b) There exists a finite constant θ *such that* $|\mu''(\gamma, \epsilon)| \leq n\theta$.
c) $\lim_{n\to\infty}(\log \alpha_n)/n = 0$.

Observe that $\mu(\gamma, \epsilon)$ does not depend on the conditional distribution for the hypotheses, *i.e.* $p(h|s)$. Thus, the term $\frac{1}{n} \log \alpha_n$ can be thought of as the "loss" accrued due to removing the information the network state provides about the true hypothesis. Hence, our assumptions is that this loss goes to zero. In addition to technical reasons, there is an intuitive reason we make this assumption. Recall that the fusion center has perfect knowledge of the network state. If the network's state carries too much information about the hypothesis, the fusion center could drive the probability of error to zero exponentially fast just by looking at the state, regardless of the rules used by the agents.

Under Assumption 1, the bounds given in (13) will be tight for sufficiently large n, and so we define [2]

$$\Lambda_n = \inf_{\gamma\in\Gamma^n} \min_{\epsilon\in[0,1]} \frac{1}{n} \mu(\gamma, \epsilon). \tag{15}$$

Then, under Assumption 1, we have the following.

Theorem 1. *Given the signal model defined by Equations* (2) *and* (3) *and the cost function in Equation* (9)*, the optimal error exponent defined in Equation* (12) *is given by*

$$\Lambda = \lim_{n\to\infty} \inf_{\gamma\in\Gamma^n} \min_{\epsilon\in[0,1]} \frac{1}{n} \mu(\gamma, \epsilon) \tag{16}$$

if the limit exists.

A few remarks are in order. First, the exponent computation makes no assumption on the correlation between the states. Second, even though the state of the system, s, is correlated with the hypotheses, this correlation has no affect on the asymptotics (provided assumption (c) is satisfied).

To provide a further analysis, we consider the following conditions: (1) Agent states are independent *a priori* and (2) Both hypotheses are possible under all s. We can then write,

$$\begin{aligned}\mu(\gamma, \epsilon) &= \log\left[\sum_s \sum_u (p_0(u|s))^{1-\epsilon} (p_1(u|s))^\epsilon p(s)\right]\\ &= \log\left[\sum_s \sum_u \prod_{k=1}^n (p_0(u_k|s_k))^{1-\epsilon} (p_1^(u_k|s_k))^\epsilon p(s_k)\right]\\ &= \log\Bigg[\left\{\sum_{s_1,u_1} (p_0(u_1|s_1))^{1-\epsilon} (p_1(u_1|s_1))^\epsilon p(s_1)\right\}...\\ &\qquad ...\left\{\sum_{s_n,u_n} (p_0(u_n|s_n))^{1-\epsilon} (p_1(u_n|s_n))^\epsilon p(s_n)\right\}\Bigg]\\ &= \sum_{k=1}^n \log\left[\sum_{s_k,u_k} (p_0(u_k|s_k))^{1-\epsilon} (p_1(u_k|s_k))^\epsilon p(s_k)\right]\\ &= \sum_{k=1}^n \mu_k(\gamma_k, \epsilon).\end{aligned} \tag{17}$$

Thus, $\mu(\gamma, \epsilon)$ is decomposable, that is, it is the sum of the $\mu_k(\gamma_k, \epsilon)$s, where agent k is using rule $\gamma_k \in \Gamma$, $k = 1, 2, ..., n$. Notice that the exponent loses this property if one of the previous assumptions is removed. It can be shown that $\mu(\gamma, \epsilon)$ is convex in ϵ and non-positive, and is zero for all ϵ except in the uninteresting case where the fusion center is unable to distinguish between the two hypotheses, *i.e.*, under all states, $p_0(u|s) = p_1(u|s)$ for all u. The same result can be shown for $\mu(\gamma_k, \epsilon)$, $k = 1, 2, .., n$. Thus, $\mu(\gamma, \epsilon)$ is non-increasing in n.

There are several applications (sensor networks, microbial systems) where each agent has common characteristics. We thus characterize this "sameness" and further develop our analysis.

[2]We take the minimum over all ϵ since for any $\gamma \in \Gamma^n$, $\mu(\gamma, \epsilon)$ is continuous in ϵ and defined over a compact set.

Definition 1. Given a collection of n agents, these agents are *identical* if the following conditions hold:

- $p_h(Y_k = y|S_k = s) = p_h(Y_j = y|S_j = s)$ for all $k, j \in \{1, 2, ..., n\}$, $h = 0, 1$.
- agent states are i.i.d *a priori*.

To motivate our asymptotic analysis of the optimizing decision rules, we provide an example for small n which results in distinct rules at each agent even if the agents are identical according to Definition 1. Consider two hypotheses (0, 1), binary messages ($b = 2$), two agents ($n = 2$), and two states (0, 1). We assume that the agent states are i.i.d with $p(S = 1) = .25$. The conditional distribution for the hypothesis is

$$p(h_0|s_1, s_2) = p(h_0|s_1) \begin{cases} .52 & s_1 = 0 \\ .48 & s_1 = 1. \end{cases}$$

The observations y_1 and y_2 are independent, conditioned on the hypothesis and state, take values in $\{0, 1, 2\}$, and have the following common distribution: $(p_0(y = 0|s = 0), p_0(y = 1|s = 0), p_0(y = 2|s = 0)) = (.8, .2, 0)$, $p_1(y = i|s = 0) = \frac{1}{3}$, and

$$p_h(y = i|s = 1) = \begin{cases} 1 & i = h \\ 0 & i \neq h \end{cases}$$

for $i = 0, 1, 2$ and $h = 0, 1$. Notice that even though agent 1 is more correlated with the hypothesis, the two agents are identical by our definition. In a given state, each agent computes a likelihood ratio test. If we enumerate through all of the cases for this discrete observations example, we see that, in a given state each agent can choose from one of two following rules:

A) $u_i = 1$ if and only if $y_i = 0$.
B) $u_i = 1$ if and only if $y_i \in \{0, 1\}$.

An optimal strategy is found by exhaustive enumeration. We find that the optimal strategy is for agent one to use rule B when in state 0 and rule A when in state 1, and for agent two to use rule A regardless of its state. This strategy results in a probability of error of .1185. Clearly, the optimal rule is not the same for each agent.

Since the agents are identical, $\mu_k(\gamma, \epsilon)$ only depends on $\gamma \in \Gamma$ and $\epsilon \in [0, 1]$, since different agents using the same rules will have the same $\mu_k(\gamma, \epsilon)$s. Thus, we drop the k subscript. Let

$$\Lambda_r = \min_{\gamma \in \Gamma} \min_{\epsilon \in [0,1]} \mu(\gamma, \epsilon). \tag{18}$$

That is, Λ_r is the optimal exponent when we restrict ourselves to the set of strategies where all agents use the same rule.

Theorem 2. *Assume the agents are identical according to Definition 1 and that both hypotheses are possible under all states. Then, Λ exists and is equal to Λ_r.*

Proof. See the Appendix. □

Theorem 2 states that there is no loss of asymptotic optimality if all agents use the same rule. Moreover, to find the optimal rule, one only need to solve the optimization problem in (18).

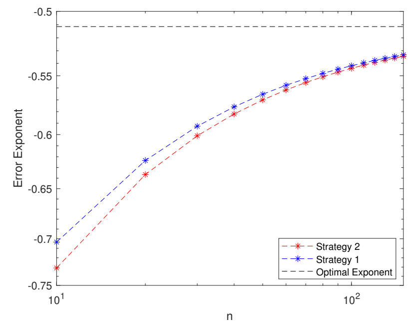

Fig. 2. Error exponents for the two strategies described in Section IV

IV. Numerical Results

In this section, we present numerical results to show the convergence of the error exponent of a sub-optimal identical strategy. Specifically, let the two hypotheses be equally likely and independent of the network state. Assume the network consists of n agents, with $n \in \{10, 20, ..., 150\}$, which can take one of two states, 0 or 1, independently of the other agents with $p(S_k = 1) = .05$ for all $k \in \{1, 2, ..., n\}$. All agents have the following common distribution: $(p_{1,0}, p_{1,0}, p_{2,0}) = (.9, .095, .005)$, $(p_{1,1}, p_{1,1}, p_{2,1}) = (.005, .9, .095)$, $(q_{1,0}, q_{1,0}, q_{2,0}) = (.1, .3, .6)$. and $(q_{1,1}, q_{1,1}, q_{2,1}) = (.875, .03, .095)$, where $p_{y,s} = p_0(y|s)$ and $q_{y,s} = p_1(y|s)$ for $y \in \{0, 1, 2\}$.

Moreover, we consider two strategies.

- *Strategy 1*:All agents use the same rule. When in state 0, agent k sends 0 when $y_k = 0$ and 1 otherwise, and when in state 1 sends 0 when $y_k \in \{0, 1\}$, for all $k \in \{1, 2, ..., n\}$.
- *Strategy 2*: The first $n - 1$ agents use the rule described in Strategy 1. The last agent n sends 0 when in state 0 and $y_n = 1$ and 1 otherwise, and when in state 1 sends 0 when $y_n \in \{0, 1\}$.

In Fig. 2 we plot the error exponent $r_n(\boldsymbol{\gamma}) = \frac{\log J_n(\boldsymbol{\gamma})}{n}$ for both strategies. The dotted black line in Fig. 2 is the optimal error exponent. Notice that even for relatively large n, Strategy 1 never strictly outperforms Strategy 2, but rather *the performance of the two strategies start to converge to each other*. Moreover, as n grows larger, the exponents of the two strategies will converge to the optimal exponent. Hence, it is not necessarily true that for large n the optimal strategy is an identical one, but that for large n *the optimal strategy cannot outperform the asymptotically optimal identical rule*.

V. Conclusions

In this paper, we have formulated the problem of state-dependent decentralized detection. The Bayes optimal tests for each agent are shown to be likelihood ratio tests with state dependent thresholds. Conditioned on this agent decision rule, the optimal fusion center rule is the MAP rule. For the most general case, computing the Bayes optimal tests for each agent in the state-dependent case is NP-hard. The overall probability of error for this system is analyzed. In particular, it is shown that despite the presence of conditionally dependent observations (in contrast to the state-free case typically studied), one can compute the error exponent for the probability of error; furthermore, the resulting asymptotic decision rules for each agent (if they are statistically identical as defined in the body of the paper) are in fact the same. Thus, a significant complexity reduction is incurred for rule computation for the large n case. Numerical results confirm properties for a key special case.

VI. Appendix

A. Proof of Upper Bound

Proof. Assuming the fusion center is implementing the MAP rule, we have

$$\begin{aligned} &J_n(\boldsymbol{\gamma}) = \\ &\sum_{\boldsymbol{s},\boldsymbol{u}} \min\{p_0(\boldsymbol{u}|\boldsymbol{s})p(h=0|\boldsymbol{s}), p_1(\boldsymbol{u}|\boldsymbol{s})p(h=1|\boldsymbol{s})\}p(\boldsymbol{s}) \\ &\overset{(a)}{\leq} \sum_{\boldsymbol{s},\boldsymbol{u}} (p_0(\boldsymbol{u}|\boldsymbol{s})p(h=0|\boldsymbol{s}))^{1-\epsilon}(p_1(\boldsymbol{u}|\boldsymbol{s})p(h=1|\boldsymbol{s}))^{\epsilon} p(\boldsymbol{s}) \\ &\leq \sum_{\boldsymbol{s},\boldsymbol{u}} p_0(\boldsymbol{u}|\boldsymbol{s})^{1-\epsilon} p_1(\boldsymbol{u}|\boldsymbol{s})^{\epsilon} p(\boldsymbol{s}) \end{aligned} \tag{19}$$

where (a) is due to the fact that for any two positive numbers a and b,

$$\min\{a, b\} \leq a^{\epsilon} b^{1-\epsilon} \quad \text{for all } \epsilon \in [0, 1]. \tag{20}$$

Hence,

$$\frac{1}{n}\log J_n(\boldsymbol{\gamma}) \leq \frac{1}{n}\log\left[\sum_{\boldsymbol{s},\boldsymbol{u}} (p_0(\boldsymbol{u}|\boldsymbol{s}))^{1-\epsilon}(p_1(\boldsymbol{u}|\boldsymbol{s}))^{\epsilon} p(\boldsymbol{s})\right]. \tag{21}$$

Since this is true for all ϵ, simply take the minimum over $0 \leq \epsilon \leq 1$. Note that this is also true for all strategies so long as the fusion center implements the MAP rule. Thus, we restrict ourselves to strategies that have the fusion center implement the MAP rule. This proves the upper bound. □

B. Proof of Theorem 2

Proof. Since having each agent use the same rule is a valid strategy, we have $\Lambda_n \leq \Lambda_r$. Then, let $(\gamma^*, \epsilon^*) \in \Gamma \times [0, 1]$ be such that $\mu(\gamma^*, \epsilon^*) \leq \mu(\gamma, \epsilon)$ for all $(\gamma^*, \epsilon^*) \in \Gamma \times [0, 1]$ (assuming such a pair exists). For any set of rules $\boldsymbol{\gamma} \in \Gamma^n$, we have

$$\min_{\epsilon\in[0,1]} \frac{1}{n}\sum_{k=1}^{n} \mu(\gamma_k, \epsilon) \geq \mu(\gamma^*, \epsilon^*) = \Lambda_r. \tag{22}$$

Since this is true for all $\boldsymbol{\gamma}$, taking the infimum over all $\boldsymbol{\gamma} \in \Gamma^n$ gives us $\Lambda_n \geq \Lambda_r$. Together with $\Lambda_n \leq \Lambda_r$, we have that $\Lambda_n = \Lambda_r$ for all n. Hence, $\Lambda = \lim_{n\to\infty} \Lambda_n = \lim_{n\to\infty} \Lambda_r = \Lambda_r$. □

References

[1] J. N. Tsitsiklis, "Decentralized detection," in *Advances in Statistical Signal Processing*. JAI Press, 1993, pp. 297–344.

[2] I. Y. Hoballah and P. K. Varshney, "Distributed bayesian signal detection," *IEEE Transactions on Information Theory*, vol. 35, no. 5, pp. 995–1000, Sep. 1989.

[3] A. Deshmukh, J. Liu, V. V. Veeravalli, and G. Verma, "Information flow optimization in inference networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May. 2020, pp. 8289–8293.

[4] Y. Gu, Y. Jiao, X. Xu, and Q. Yu, "Request–response and censoring-based energy-efficient decentralized change-point detection with iot applications," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6771–6788, Apr. 2021.

[5] G. Gül, "Minimax robust decentralized hypothesis testing for parallel sensor networks," *IEEE Transactions on Information Theory*, vol. 67, no. 1, pp. 538–548, Jan. 2020.

[6] A. Gurevich, K. Cohen, and Q. Zhao, "Sequential anomaly detection under a nonlinear system cost," *IEEE Transactions on Signal Processing*, vol. 67, no. 14, pp. 3689–3703, Jul. 2019.

[7] B. Huang, K. Cohen, and Q. Zhao, "Active anomaly detection in heterogeneous processes," *IEEE Transactions on Information Theory*, vol. 65, no. 4, pp. 2284–2301, Aug. 2018.

[8] M. N. Kurt, Y. Yilmaz, and X. Wang, "Real-time nonparametric anomaly detection in high-dimensional settings," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, Jan. 2020.

[9] C. Zhong, M. C. Gursoy, and S. Velipasalar, "Deep actor-critic reinforcement learning for anomaly detection," in *Proc. IEEE Global Communications Conference (GLOBECOM)*, Dec. 2019, pp. 1–6.

[10] U. S. Shanthamallu, A. Spanias, C. Tepedelenlioglu, and M. Stanley, "A brief survey of machine learning methods and their sensor and iot applications," in *Proc. International Conference on Information, Intelligence, Systems & Applications (IISA)*, Aug. 2017, pp. 1–8.

[11] J. Engebrecht, K. Nealson, and M. Silverman, "Bacterial bioluminescence: isolation and genetic analysis of functions from vibrio fischeri," *Cell*, vol. 32, no. 3, pp. 773–781, Mar. 1983.

[12] J. Boedicker and K. Nealson, "Microbial communication via quorum sensing," *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications*, vol. 1, no. 4, pp. 310–320, Dec. 2015.

[13] M. M. Vasconcelos, U. Mitra, O. Camara, K. P. Silva, and J. Boedicker, "Bacterial quorum sensing as a networked decision system," in *Proc. IEEE International Conference on Communications (ICC)*, May. 2018, pp. 1–6.

[14] M. M. Vasconcelos, U. Mitra, O. Câmara, M. S. Gangan, and J. Boedicker, "A continuous-time decision-making model for bacterial growth via quorum sensing: theory and evidence," in *Proc. ACM International Conference on Nanoscale Computing and Communication*, Sep. 2019, pp. 1–6.

[15] R. R. Tenney and N. R. Sandell, "Detection with distributed sensors," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-17, no. 4, pp. 501–510, Jul. 1981.

[16] D. Teneketzis and Y.-C. Ho, "The decentralized wald problem," *Information and Computation*, vol. 73, no. 1, pp. 23–44, Apr. 1987.

[17] H. Chen, B. Chen, and P. K. Varshney, "A new framework for distributed detection with conditionally dependent observations," *IEEE Transactions on Signal Processing*, vol. 60, no. 3, pp. 1409–1419, Mar. 2012.

[18] J. Tsitsiklis and M. Athans, "On the complexity of decentralized decision making and detection problems," *IEEE Transactions on Automatic Control*, vol. 30, no. 5, pp. 440–446, May. 1985.

[19] E. Drakopoulos and C. C. Lee, "Optimum fusion of correlated local decisions," in *Proc. IEEE Conference on Decision and Control (CDC)*, Dec. 1988, pp. 2489–2494 vol.3.

[20] J. Shaska and U. Mitra, "Decentralized decision-making for multi-agent networks: the state-dependent case supplemental notes," 2021, unpublished, available at https://documentcloud.adobe.com/link/track?uri=urn:aaid:scds:US:fd9f98b1-4514-4668-af37-7f815c4ac559.

[21] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Mathematics of Control, Signals and Systems*, vol. 1, no. 2, pp. 167–182, Jun. 1988.

[22] C. E. Shannon, R. G. Gallager, and E. R. Berlekamp, "Lower bounds to error probability for coding on discrete memoryless channels. I," *Information and Control*, vol. 10, no. 1, pp. 65–103, Jan. 1967.