

Data Fusion Trees for Detection: Does Architecture Matter?

Wee Peng Tay, *Student Member, IEEE*, John N. Tsitsiklis, *Fellow, IEEE*, and Moe Z. Win, *Fellow, IEEE*

Abstract—We consider the problem of decentralized detection in a network consisting of a large number of nodes arranged as a tree of bounded height, under the assumption of conditionally independent and identically distributed (i.i.d.) observations. We characterize the optimal error exponent under a Neyman–Pearson formulation. We show that the Type II error probability decays exponentially fast with the number of nodes, and the optimal error exponent is often the same as that corresponding to a parallel configuration. We provide sufficient, as well as necessary, conditions for this to happen. For those networks satisfying the sufficient conditions, we propose a simple strategy that nearly achieves the optimal error exponent, and in which all non-leaf nodes need only send 1-bit messages.

Index Terms—Decentralized detection, error exponent, sensor networks.

I. INTRODUCTION

MOST of the decentralized detection literature has been concerned with characterizing optimal detection strategies for particular sensor configurations; the comparison of the detection performance of different configurations is a rather unexplored area. We bridge this gap by considering the asymptotic performance of bounded height tree networks. We analyze the dependence of the optimal error exponent on the network architecture, and characterize the optimal error exponent for a large class of tree networks. The problem of optimal decentralized detection has attracted a lot of interest over the last 25 years. Tenney and Sandell [1] are the first to consider a decentralized detection system in which each of several sensors makes an observation and sends a summary (e.g., using a quantizer or other “transmission function”) to a fusion center. Such a system is to be contrasted to a *centralized* one, where the raw observations are transmitted directly to the fusion center. The framework introduced in [1] involves a “star topology” or “parallel configuration”: the fusion center is regarded as the root of a tree, while the sensors are the leaves, directly connected to the root. Several pieces of work follow, e.g., [2]–[12], all of which study the

parallel configuration under a Neyman–Pearson or Bayesian criterion. A common goal of these references is to characterize the optimal transmission function, where optimality usually refers to the minimization of the probability of error or some other cost function at the fusion center. A typical result is that under the assumption of (conditionally) independent sensor observations, likelihood ratio quantizers are optimal; see [6] for a summary of such results.

The study of sensor networks other than the parallel configuration is initiated in [13], which considers a tandem configuration, as well as more general tree configurations, and characterizes optimal transmission strategies under a Bayesian formulation. Tree configurations are also discussed in [14]–[21], under various performance objectives. In all but the simplest cases, the exact form of optimal strategies in tree configurations is difficult to derive. Most of these references focus on person-by-person (PBP) optimality and obtain necessary, but not sufficient, conditions for an optimal strategy. When the transmission functions are assumed to be finite-alphabet quantizers, typical results establish that under a conditional independence assumption, likelihood ratio quantizers are PBP optimal. However, finding the optimal quantizer thresholds requires the solution of a nonlinear system of equations, with as many equations as there are thresholds. As a consequence, computing the optimal thresholds or characterizing the overall performance is hard, even for networks of moderate size.

Because of these difficulties, the analysis and comparison of large sensor networks is apparently tractable only in an asymptotic regime that focuses on the rate of decay of the error probabilities as the number of sensors increases. For example, in the Neyman–Pearson framework, one can focus on minimizing the error exponent¹

$$g = \limsup_{n \rightarrow \infty} \frac{1}{n} \log \beta_n$$

where β_n is the Type II error probability at the fusion center and n is the number of sensors, while keeping the Type I error probability less than some given threshold. Note our convention that error exponents are negative numbers. The magnitude of the error exponent, $|g|$, is commonly referred to as the rate of decay of the Type II error probability. A larger $|g|$ would translate to a faster decay rate, hence a better detection performance. This problem has been studied in [22], for the case of a parallel configuration with a large number of sensors that receive independent and identically distributed (i.i.d.) observations.

The asymptotic performance of another special configuration, involving n sensors arranged in tandem, has been studied in

Manuscript received November 28, 2006; revised November 20, 2007. Published August 27, 2008 (projected). This work was supported in part by the National Science Foundation under Contracts ECS-0426453 and ANI-0335256, the Charles Stark Draper Laboratory Robust Distributed Sensor Networks Program, and an Office of Naval Research Young Investigator Award N00014-03-1-0489. The material in this paper was presented in part at the 44th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, September 2006.

The authors are with the Laboratory for Information and Decision Systems, the Massachusetts Institute of Technology (MIT), Cambridge, MA USA (e-mail: wptay@mit.edu; jnt@mit.edu; moewin@mit.edu).

Communicated by L. Tong, Associate Editor for Detection and Estimation.
Digital Object Identifier 10.1109/TIT.2008.928240

¹Throughout this paper, \log stands for the natural logarithm.




Fig. 1. A tree network of height h , with $n - h$ leaves. Its error probability is no larger than that of a parallel configuration with $n - h$ leaves and a fusion center. If h is bounded while n increases, the optimal error exponent is the same as for a parallel configuration with n leaves.

[23]–[25], under a Bayesian formulation. Necessary and sufficient conditions for the error probability to decrease to zero as n increases have been derived. However, even when the error probability decreases to zero, it apparently does so at a subexponential rate (see [26] for such a result for the Bayesian case). Accordingly, [25] argues that the tandem configuration is inefficient and suggests that as the number of sensors increases, the network “should expand more in a parallel than in [a] tandem” fashion.

Even though the error probabilities in a parallel configuration decrease exponentially, the energy consumption of having each sensor transmit directly to the fusion center can be too high. The energy consumption can be reduced by setting up a directed spanning in-tree, rooted at the fusion center. In a tree configuration, each non-leaf node combines its own observation (if any) with the messages it has received and forms a new message, which it transmits to another node. In this way, information from each node is propagated along a multihop path to the fusion center, but the information is “degraded” along the way. For the case where observations are obtained only at the leaves, it is not hard to see that the detection performance of such a tree cannot be better than that of a parallel configuration with the same number of leaves.

In this paper, we investigate the detection performance of a tree configuration under a Neyman–Pearson criterion. We restrict to trees with bounded height for two reasons. First, without a restriction on the height of the tree, performance can be poor (this is exemplified by tandem networks in which, as remarked above, the error probability seems to decay at a subexponential rate). Second, bounded height translates to a bound on the delay until information reaches the fusion center.

As it is not apparent that the Type II error probability decays exponentially fast with the number of nodes in the network, we first show that under the bounded height assumption, exponential decay is possible. We then obtain the rather counterintuitive result that if leaves dominate (in the sense that asymptotically almost all nodes are leaves), then bounded height trees have the same asymptotic performance as the parallel configuration, even in nontrivial cases. (Such an equality is clear in some trivial cases, e.g., the configuration shown in Fig. 1, but is unexpected in general.) This result has important ramifications: a system designer can reduce the energy consumption in a network (e.g., by employing an h -hop spanning tree that minimizes the overall energy consumption), without losing detection efficiency, under certain conditions.

We also provide a strategy in which each non-leaf node sends only a 1-bit message, and which nearly achieves the same performance as the parallel configuration. These results are coun-

terintuitive for the following reasons: 1) messages are compressed to only 1 bit at each non-leaf node so that “information” is lost along the way, whereas in the parallel configuration, no such compression occurs; 2) even though leaves dominate, there is no reason why the error exponent will be determined solely by the leaves. For example, our discussion in Section V-E indicates that without the bounded height assumption, or if a Bayesian framework is assumed instead of the Neyman–Pearson formulation, then a generic tree network (of height greater than 1) performs strictly worse than a parallel configuration, even if leaves dominate.

Finally, under a mild additional assumption on the allowed transmission functions, we find that the sufficient conditions for achieving the same error exponent as a parallel configuration, are also necessary.

The rest of this paper is organized as follows. In Section II, we present our model in detail. In Section III, we state the Neyman–Pearson problem, provide some motivating examples, and state the main results. In Section IV, we consider “relay trees,” in which observations are only made at the leaves. In Section V, we prove the main results. Finally, in Section VI, we summarize our work and offer some concluding remarks.

II. PROBLEM FORMULATION

In this section, we introduce the model and the required notation. We consider a decentralized binary detection problem involving $n - 1$ sensors and a fusion center; we will be interested in the case where n increases to infinity. We are given two probability spaces $(\Omega, \mathcal{F}, \mathbb{P}_0)$ and $(\Omega, \mathcal{F}, \mathbb{P}_1)$, associated with two hypotheses H_0 and H_1 . We use \mathbb{E}_j to denote the expectation operator with respect to \mathbb{P}_j . Each sensor v observes a random variable X_v taking values in some set \mathcal{X} . Under either hypothesis H_j , $j = 0, 1$, the random variables X_v are i.i.d., with marginal distribution $\mathbb{P}_j^{\mathcal{X}}$.

A. Tree Networks

The configuration of the sensor network is represented by a directed tree $T_n = (V_n, E_n)$. Here, V_n is the set of nodes, of cardinality n , and E_n is the set of directed arcs of the tree. One of the nodes (the “root”) represents the fusion center, and the remaining $n - 1$ nodes represent the remaining sensors. We will always use the special symbol f to denote the root of T_n . We assume that the arcs are oriented so that they all point toward the fusion center. In the sequel, whenever we use the term “tree,” we mean a directed, rooted tree as described above.

We will use the terminology “sensor” and “node” interchangeably. Moreover, the fusion center f will also be called a sensor, even though it plays the special role of fusing; whether the fusion center makes its own observation or not is irrelevant, since we are working in the large n regime, and we will assume it does not.

We say that node u is a *predecessor* of node v if there exists a directed path from u to v . In this case, we also say that v is a *successor* of u . An *immediate predecessor* of node v is a node u such that $(u, v) \in E_n$. An immediate successor is similarly defined. Let the set of immediate predecessors of v be $C_n(v)$. If v is a leaf, $C_n(v)$ is naturally defined to be empty. The *length*




Fig. 2. Both nodes v and u belong to the set A_n , but only node u belongs to the set B_n .

of a path is defined as the number of arcs in the path. The *height* of the tree T_n is the length of the longest path from a leaf to the root, and will be denoted by h_n .

Since we are interested in asymptotically large values of n , we will consider a *sequence* of trees $(T_n)_{n \geq 1}$. While we could think of the sequence as representing the evolution of the network as sensors are added, we do not require the sequence E_n to be an increasing sequence of sets; thus, the addition of a new sensor to T_n may result in some edges being deleted and some new edges being added. We define the height of a sequence of trees to be $h = \sup_{n \geq 1} h_n$. We are interested in tree sequences of bounded height, i.e., $h < \infty$.

Definition 1 (h -Uniform Tree): A tree T_n is said to be h -uniform if the length of every path from a leaf to the root is exactly h . A sequence of trees $(T_n)_{n \geq 1}$ is said to be h -uniform if there exists some $n_0 < \infty$, so that for all $n \geq n_0$, T_n is h -uniform.

For a tree with height h , we say that a node is at *level* k if it is connected to the fusion center via a path of length $h - k$. Hence, the fusion center f is at level h , while in an h -uniform tree, all leaves are at level 0.

Let $l_n(v)$ be the number of leaves of the subtree rooted at the node v . (These are the leaves whose path to f goes through v .) Thus, $l_n(f)$ is the total number of leaves. Let $p_n(v)$ be the total number of predecessors of v , i.e., the total number of nodes in the subtree rooted at v , not counting v itself. Thus, $p_n(f) = n - 1$. We let $A_n \subset V_n$ be the set of nodes whose immediate predecessors include leaves of the tree T_n . Finally, we let $B_n \subset A_n$ be the set of nodes all of whose predecessors are leaves; see Fig. 2.

B. Strategies

Given a tree T_n , consider a node $v \neq f$. Node v receives messages Y_u from every $u \in C_n(v)$ (i.e., from its immediate predecessors). Node v then uses a transmission function γ_v to encode and transmit a summary $Y_v = \gamma_v(X_v, \{Y_u : u \in C_n(v)\})$ of its own observation X_v , and of the received messages $\{Y_u : u \in C_n(v)\}$, to its immediate successor.² We constrain all messages to be symbols in a fixed alphabet \mathcal{T} . Thus, if the in-degree of v is $|C_n(v)| = d$, then the transmission function γ_v maps $\mathcal{X} \times \mathcal{T}^d$ to \mathcal{T} . Let $\Gamma(d)$ be a given set of transmission functions that the node v can choose from. In general, $\Gamma(d)$ is a subset of the set of all possible mappings from $\mathcal{X} \times \mathcal{T}^d$ to \mathcal{T} . For example, $\Gamma(d)$ is often assumed to be the set of quantizers whose outputs are the result of comparing likelihood ratios to some thresholds (cf.

²To simplify the notation, we suppress the dependence of X_v , Y_u , γ_v , etc., on n .

the definition of a log-likelihood ratio quantizer (LLRQ) in Section III-B). For convenience, we denote the set of transmission functions for the leaves, $\Gamma(0)$, by Γ . We assume that all transmissions are perfectly reliable.

Consider now the root f , and suppose that it has d immediate predecessors. It receives messages from its immediate predecessors, and based on this information, it decides between the two hypotheses H_0 and H_1 , using a fusion rule $\gamma_f : \mathcal{T}^d \mapsto \{0, 1\}$.³ Let Y_f be a binary-valued random variable indicating the decision of the fusion center.

We define a *strategy* for a tree T_n , with $n - 1$ nodes and a fusion center, as a collection of transmission functions, one for each node, and a fusion rule. In some cases, we will be considering strategies in which only the leaves make observations; every other node v simply fuses the messages it has received, and forwards a message $Y_v = \gamma_v(\{Y_u : u \in C_n(v)\})$ to its immediate successor. A strategy of this type will be called a *relay strategy*. A tree network in which we restrict to relay strategies will be called a *relay tree*. If in addition, the alphabet \mathcal{T} is binary, we will use the terms 1-bit relay strategy and 1-bit relay tree. Finally, in a relay tree, nodes other than the root and the leaves will be called *relay nodes*.

III. THE NEYMAN–PEARSON PROBLEM

In this section, we formulate the Neyman–Pearson decentralized detection problem in a tree network. We provide some motivating examples, and introduce our assumptions. Then, we give a summary of the main results.

Given a tree T_n , we require that the Type I error probability $\mathbb{P}_0(Y_f = 1)$ be no more than a given $\alpha \in (0, 1)$. A strategy is said to be *admissible* if it meets this constraint. We are interested in minimizing the Type II error probability $\mathbb{P}_1(Y_f = 0)$. Accordingly, we define $\beta^*(T_n)$ as the infimum of $\mathbb{P}_1(Y_f = 0)$, over all admissible strategies. Similarly, we define $\beta_R^*(T_n)$ as the infimum of $\mathbb{P}_1(Y_f = 0)$, over all admissible relay strategies. Typically, $\beta^*(T_n)$ or $\beta_R^*(T_n)$ will converge to zero as $n \rightarrow \infty$. We are interested in the question of whether such convergence takes place exponentially fast, and in the exact value of the Type II error exponent, defined by

$$g^* = \limsup_{n \rightarrow \infty} \frac{1}{n} \log \beta^*(T_n)$$

$$g_R^* = \limsup_{n \rightarrow \infty} \frac{1}{l_n(f)} \log \beta_R^*(T_n).$$

Note that in the relay case, we use the total number of leaves $l_n(f)$ instead of n in the definition of g_R^* . This is because only the leaves make observations and therefore, g_R^* measures the rate of error decay per observation.

We denote the Kullback–Leibler (KL) divergence of two probability measures, \mathbb{P} and \mathbb{Q} , as

$$D(\mathbb{P} \parallel \mathbb{Q}) = \mathbb{E}^{\mathbb{P}} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right]$$

³Recall that in centralized Neyman–Pearson detection, randomization can reduce the Type II error probability. Therefore, in general, the fusion center uses a randomized fusion rule to make its decision. Similarly, the transmission functions γ_v used by each node v , can also be randomized. We avoid any discussion of randomization to simplify the exposition, and because randomization is not required asymptotically, as will become apparent in Section V.

where $\mathbb{E}^{\mathbb{P}}$ is the expectation operator with respect to (w.r.t.) \mathbb{P} . Suppose that X is a sensor observation. For any $\gamma \in \Gamma$, let the distribution of $\gamma(X)$, under hypothesis H_j , be \mathbb{P}_j^γ . Note that $-\mathrm{D}(\mathbb{P}_0^\gamma \|\mathbb{P}_1^\gamma) \leq 0 \leq \mathrm{D}(\mathbb{P}_1^\gamma \|\mathbb{P}_0^\gamma)$, with both inequalities being strict as long as the measures \mathbb{P}_0^γ and \mathbb{P}_1^γ are not indistinguishable.

In the classical case of a parallel configuration, with $n - 1$ leaves directly connected to the fusion center, the optimal error exponent, denoted as g_P^* , is given by [22]

$$g_P^* = \lim_{n \rightarrow \infty} \frac{1}{n} \log \beta^*(T_n) = -\sup_{\gamma \in \Gamma} \mathrm{D}(\mathbb{P}_0^\gamma \|\mathbb{P}_1^\gamma) \quad (1)$$

under Assumptions 1 and 2, stated in Section III-B below.

Our objective is to study g^* and g_R^* for different sequences of trees. In particular, we wish to obtain bounds on these quantities, develop conditions under which they are strictly negative (indicating exponential decay of error probabilities), and develop conditions under which they are equal to g_P^* . At this point, under Assumptions 1 and 2, we can record two relations that are always true

$$g_P^* \leq g_R^*, \quad -\mathrm{D}(\mathbb{P}_0^X \|\mathbb{P}_1^X) \leq g^* \leq z g_R^* \quad (2)$$

where $z = \liminf_{n \rightarrow \infty} l_n(f)/n$. The first inequality is true because all of the combining of messages that takes place in a relay network can be carried out internally, at the fusion center of a parallel network with the same number of leaves. The inequality $-\mathrm{D}(\mathbb{P}_0^X \|\mathbb{P}_1^X) \leq g^*$ follows from the fact that $-\mathrm{D}(\mathbb{P}_0^X \|\mathbb{P}_1^X)$ is the classical error exponent in a centralized system where all raw observations are transmitted directly to the fusion center. Finally, the inequality $g^* \leq z g_R^*$ follows because an optimal strategy is at least as good as an optimal relay strategy; the factor of z arises because we have normalized g_R^* by $l_n(f)$ instead of n .

For a sequence of trees of the form shown in Fig. 1, it is easily seen that $g^* = g_R^* = g_P^*$. In order to develop some insights into the problem, we now consider some less trivial examples.

A. Motivating Examples

In the following examples, we restrict to relay strategies for simplicity, i.e., we are interested in characterizing the error exponent g_R^* . However, most of our subsequent results hold without such a restriction, and similar statements can be made about the error exponent g^* (cf. Theorem 1).

Example 1: Consider a 2-uniform sequence of trees, as shown in Fig. 3, where each node v_i receives messages from $m = (n - 3)/2$ leaves (for simplicity, we assume that n is odd).

Let us restrict to 1-bit relay strategies. Consider the fusion rule that declares H_0 iff both v_1 and v_2 send a 0. In order to keep the Type I error probability bounded by α , we view the message by each v_i as a local decision about the hypothesis, and require that its local Type I error probability be bounded by $\alpha/2$. Furthermore, by viewing the subtree rooted at v_i as a parallel configuration, we can design strategies for each subtree so that

$$\lim_{n \rightarrow \infty} \frac{1}{m} \log \mathbb{P}_1(Y_{v_i} = 0) = g_P^*. \quad (3)$$




Fig. 3. A 2-uniform tree with two relay nodes.




Fig. 4. A 2-uniform tree with a large number of relay nodes.

At the fusion center, the Type II error exponent is then given by

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{n} \log \beta_n \\ &= \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_1(Y_{v_1} = 0, Y_{v_2} = 0) \\ &= \frac{1}{2} \lim_{n \rightarrow \infty} \frac{1}{m} \log \mathbb{P}_1(Y_{v_1} = 0) \\ &\quad + \frac{1}{2} \lim_{n \rightarrow \infty} \frac{1}{m} \log \mathbb{P}_1(Y_{v_2} = 0) \\ &= g_P^* \end{aligned}$$

where the last equality follows from (3). This shows that the Type II error probability falls exponentially and, more surprisingly, that $g_R^* \leq g_P^*$. In view of (2), we have $g_R^* = g_P^*$. It is not difficult to generalize this conclusion to all sequences of trees in which the number $n - l_n(f) - 1$ of relay nodes is bounded. For such sequences, we will also see that $g^* = g_R^*$ (cf. Theorem 1 (iii)). \square

Example 2: We now consider an example in which the number of relay nodes grows with n . In Fig. 4, we let both m and N be increasing functions of n (the total number of nodes), in a manner to be made explicit shortly.

Let us try to apply a similar argument as in Example 1, to see whether the optimal exponent of the parallel configuration can be achieved with a relay strategy, i.e., whether $g_R^* = g_P^*$. We let each node v_i use a local Neyman–Pearson test. We also let the fusion center declare H_0 iff it receives a 0 from all relay sensors. In order to have a hope of achieving the error exponent of the parallel configuration, we need to choose the local Neyman–Pearson test at each relay so that its local Type II error exponent is close to $g_P^* = -\sup_{\gamma \in \Gamma} \mathrm{D}(\mathbb{P}_0^\gamma \|\mathbb{P}_1^\gamma)$. However, the associated local Type I error cannot fall faster than exponentially, so we can assume it is bounded below by $\delta \exp(-m\epsilon)$, for some $\delta, \epsilon > 0$, and for all m large enough. In that case, the overall Type I error probability (at the fusion center) is at least $1 - (1 - \delta e^{-m\epsilon})^N$. We then note that if N increases quickly with m (e.g., $N = m^m$), the Type I error probability approaches 1, and eventually exceeds α . Hence, we no longer have an admissible strategy. Thus, if there is a hope of achieving the optimal exponent g_P^* of the parallel configuration, a more complicated fusion rule will have to be used. \square




Fig. 5. A 2-uniform tree, with $m = l_n(f)/2$.

Our subsequent results will establish that, similar to Example 1, the equalities $g^* = g_R^* = g_P^*$ also hold in Example 2. However, Example 2 shows that in order to achieve this optimal error exponent, we may need to employ nontrivial fusion rules at the fusion center (and for similar reasons at the relay nodes), and various thresholds will have to be properly tuned. The simplicity of the fusion rule in Example 1 is not representative.

In our next example, the optimal error exponent is inferior (strictly larger) than that of a parallel configuration.

Example 3: Consider a sequence of 1-bit relay trees with the structure shown in Fig. 5. Let the observations X_v at the leaves be i.i.d. Bernoulli random variables with parameter $1 - p$ under H_0 , and parameter p under H_1 , where $1/2 < p < 1$. Note that

$$g_P^* = \mathbb{E}_0 \left[\log \frac{d\mathbb{P}_1^X}{d\mathbb{P}_0^X} \right] = p \log \frac{1-p}{p} + (1-p) \log \frac{p}{1-p}.$$

We can identify this relay tree with a parallel configuration involving m nodes, with each node receiving an independent observation distributed as $\gamma(X_1, X_2)$. Note that we can restrict the transmission function γ to be the same for all nodes v_1, \dots, v_m [22], without loss of optimality. We have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{m} \log \beta^*(T_n) \\ &= \min_{\gamma \in \Gamma(2)} \sum_{j=0}^1 \mathbb{P}_0(\gamma(X_1, X_2) = j) \log \left[\frac{\mathbb{P}_1(\gamma(X_1, X_2) = j)}{\mathbb{P}_0(\gamma(X_1, X_2) = j)} \right]. \end{aligned} \tag{4}$$

To minimize the right-hand side (RHS) of (4), we only need to consider a small number of choices for γ . If $\gamma(X_1, X_2) = X_1$, we are effectively removing half of the original $2m$ nodes, and the resulting error exponent is $g_P^*/2$, which is inferior to g_P^* . Suppose now that γ is of the form $\gamma(X_1, X_2) = 0$ iff $X_1 = X_2 = 0$. Then, it is easy to see, after some calculations (omitted), that

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{m} \log \beta^*(T_n) \\ &= p^2 \log \frac{(1-p)^2}{p^2} + (1-p^2) \log \frac{1-(1-p)^2}{1-p^2} \\ &> 2 \left(p \log \frac{1-p}{p} + (1-p) \log \frac{p}{1-p} \right) \end{aligned}$$

and

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{l_n(f)} \log \beta^*(T_n) \\ &> p \log \frac{1-p}{p} + (1-p) \log \frac{p}{1-p} \\ &= g_P^*. \end{aligned}$$




Fig. 6. A 2-uniform tree, with $l_n(v_i) = i + 1$.

Finally, we need to consider γ of the form $\gamma(X_1, X_2) = 1$ iff $X_1 = X_2 = 1$. A similar calculation (omitted) shows that the resulting error exponent is again inferior. We conclude that the relay network is strictly inferior to the parallel configuration, i.e., $g_P^* < g_R^*$. An explanation is provided by noting that this sequence of trees violates a necessary condition, developed in Section V-F for the optimal error exponent to be the same as that of a parallel configuration; see Theorem 1(iv). \square

A comparison of the results for the previous examples suggests that we have $g_P^* = g_R^*$ (respectively, $g_P^* < g_R^*$) whenever the degree of level 1 nodes increases (respectively, stays bounded) as n increases. That would still leave open the case of networks in which different level 1 nodes have different degrees, as in our next example.

Example 4: Consider a sequence of 2-uniform trees of the form shown in Fig. 6. Each node v_i , $i = 1, \dots, m$, has $i + 1$ leaves attached to it. We will see that the optimal error exponent is again the same as for a parallel configuration, i.e., $g_R^* = g^* = g_P^*$. (cf. Theorem 1(ii)). \square

B. Assumptions

In this subsection, we list our assumptions. Assumptions 1 and 2 are similar to the assumptions made in the study of the parallel configuration (see [22]).

Assumption 1: The measures \mathbb{P}_0^X and \mathbb{P}_1^X are equivalent, i.e., they are absolutely continuous w.r.t. each other. Furthermore, there exists some $\gamma \in \Gamma$ such that $-D(\mathbb{P}_0^\gamma || \mathbb{P}_1^\gamma) < 0 < D(\mathbb{P}_1^\gamma || \mathbb{P}_0^\gamma)$.

$$\text{Assumption 2: } \mathbb{E}_0 \left[\log^2 \frac{D\mathbb{P}_1^\gamma}{D\mathbb{P}_0^\gamma} \right].$$

Assumption 2 implies the following lemma; see [22] for a proof.

Lemma 1: There exists some $a \in (0, \infty)$, such that for all $\gamma \in \Gamma$,

$$\begin{aligned} & \mathbb{E}_0 \left[\log^2 \frac{d\mathbb{P}_1^\gamma}{d\mathbb{P}_0^\gamma} \right] \leq \mathbb{E}_0 \left[\log^2 \frac{d\mathbb{P}_1^X}{d\mathbb{P}_0^X} \right] + 1 < a \\ & \mathbb{E}_0 \left[\left| \log \frac{d\mathbb{P}_1^\gamma}{d\mathbb{P}_0^\gamma} \right| \right] < a. \end{aligned}$$

Given an admissible strategy, and for each node $v \in V_n$, we consider the log-likelihood ratio of the distribution of Y_v (the message sent by v) under H_1 , w.r.t. its distribution under H_0

$$\mathcal{L}_{v,n}(y) = \log \frac{d\mathbb{P}_{1,n}^{(v)}}{d\mathbb{P}_{0,n}^{(v)}}(y),$$

where $d\mathbb{P}_{1,n}^{(v)}/d\mathbb{P}_{0,n}^{(v)}$ is the Radon–Nikodym derivative of the distribution of Y_v under H_1 w.r.t. that under H_0 . If Y_v takes values in a discrete set, then this is just the log-likelihood ratio

$\log(\mathbb{P}_1(Y_v = y)/\mathbb{P}_0(Y_v = y))$. For simplicity, we let $L_{v,n} = \mathcal{L}_{v,n}(Y_v)$ and define the log-likelihood ratio of the received messages at node v to be

$$S_n(v) = \sum_{u \in C_n(v)} L_{u,n}.$$

(Recall that $C_n(v)$ is the set of immediate predecessors of v .)

A (1-bit) log-likelihood ratio quantizer (LLRQ) with threshold t for a non-leaf node v , with $|C_n(v)| = d$, is a binary-valued function on \mathcal{T}^d , defined by

$$\text{LLRQ}_{d,t}(\{y_u : u \in C_n(v)\}) = \begin{cases} 0, & \text{if } x \leq t \\ 1, & \text{if } x > t \end{cases}$$

where

$$x = \frac{1}{l_n(v)} \sum_{u \in C_n(v)} \mathcal{L}_{u,n}(y_u). \quad (5)$$

By definition, a node v that uses a LLRQ ignores its own observation X_v and acts as a relay. If all non-leaf nodes use an LLRQ, we have a special case of a relay strategy. We will assume that LLRQs are available choices of transmission functions for all nonleaf nodes.

Assumption 3: For all $t \in \mathbb{R}$ and $d > 0$, $\text{LLRQ}_{d,t} \in \Gamma(d)$.

As already discussed (cf. (2)), the optimal performance of a relay tree is always dominated by that of a parallel configuration with the same number of leaves, i.e., $g_P^* \leq g_R^*$. In Section V, we find sufficient conditions under which the equality $g_R^* = g_P^*$ holds. Then, in Section V-F, we look into necessary conditions for this to be the case. It turns out that nontrivial necessary conditions for the equality $g_R^* = g_P^*$ to hold are, in general, difficult to obtain, because they depend on the nature of the transmission functions available to the sensors. For example, if the sensors are allowed to simply forward undistorted all of the messages that they receive, then the equality $g_R^* = g_P^*$ holds trivially. Hence, we need to impose some restrictions on the set of transmission functions available, as in the assumption that follows.

Assumption 4:

- There exists an $n_0 \geq 1$ such that for all $n \geq n_0$, we have $l_n(v) > 1$ for all v in the set B_n of nodes whose immediate predecessors are all leaves.
- Let X_1, X_2, \dots be i.i.d. random variables under either hypothesis H_j , each with distribution \mathbb{P}_j^X . For $k > 1$, $\gamma_0 \in \Gamma(k)$, and $\gamma_i \in \Gamma$, where $i = 1, \dots, k$. Let $\xi = (\gamma_0, \dots, \gamma_k)$. We also let ν_j^ξ be the distribution of $\gamma_0(\gamma_1(X_1), \dots, \gamma_k(X_k))$ under hypothesis H_j . We assume that

$$g_P^* < \inf_{\xi \in \Gamma(k) \times \Gamma^k} \frac{1}{k} \left[\log \frac{d\nu_1^\xi}{d\nu_0^\xi} \right] \quad (6)$$

for all $k > 1$.

Assumption 4 holds in most cases of interest. Part (a) results in no loss of generality: if in a relay tree we have $l_n(v) = 1$ for some $v \in B_n$, we can remove the predecessor of v , and treat v as a leaf. Regarding part (b), it is easy to see that the

left-hand side (LHS) of (6) is always less than or equal to the RHS, hence we have only excluded those cases where (6) holds with equality. We are essentially assuming that when the messages $\gamma_1(X_1), \dots, \gamma_k(X_k)$ are summarized (or quantized) by γ_0 , there is some loss of information, as measured by the associated KL divergences.

C. Main Results

In this section, we collect and summarize our main results. The asymptotic proportion of nodes that are leaves, defined by

$$z = \liminf_{n \rightarrow \infty} \frac{l_n(f)}{n}$$

plays a critical role.

Theorem 1: Consider a sequence of trees, $(T_n)_{n \geq 1}$, of bounded height. Suppose that Assumptions 1–3 hold. Then

- $g_P^* \leq g_R^* < 0$ and $-\text{D}(\mathbb{P}_0^X | \mathbb{P}_1^X) \leq g^* \leq z g_R^* < 0$;
- if $z = 1$, then $g_P^* = g^* = g_R^*$;
- if the number of non-leaf nodes is bounded, or if $\min_{v \in B_n} l_n(v) \rightarrow \infty$, then $g_P^* = g^* = g_R^*$;
- if Assumption 4 also holds, we have $g_R^* = g_P^*$ iff $z = 1$.

Note that part (i) follows from (2), except for the strict negativity of the error exponents, which is established in Proposition 2. Part (ii) is proved in Proposition 3. Part (iii) is proved in Corollary 1. (Recall that B_n is the set of non-leaf nodes all of whose immediate predecessors are leaves.) Part (iv) is proved in Proposition 5. One might also have expected a result asserting that $g_P^* \leq g^*$. However, this is not true without additional assumptions, as will be discussed in Section V-F.

IV. ERROR BOUNDS FOR h -UNIFORM RELAY TREES

In this section, we consider a 1-bit h -uniform relay tree, in which all relay nodes at level k use a LLRQ with a common threshold t_k . We wish to develop upper bounds for the error probabilities at the various nodes. We do this recursively, by moving along the levels of the tree, starting from the leaves. Given bounds on the error probabilities associated with the messages received by a node, we develop a bound on the log-moment generating function at that node (cf. (8)), and then use the standard Chernoff bound technique to develop a bound on the error probability for the message sent by that node (cf. (7)).

Let $t^{(k)} = (t_1, t_2, \dots, t_k)$, for $k \geq 1$, and $t^{(0)} = \emptyset$. For $j = 0, 1, k \geq 1$, and $\lambda \in \mathbb{R}$, we define recursively

$$\Lambda_{j,0}(\gamma; \lambda) = \Lambda_{j,0}(\gamma, \emptyset; \lambda) = \log \mathbb{E}_j \left[\left(\frac{d\mathbb{P}_1^\gamma}{d\mathbb{P}_0^\gamma} \right)^\lambda \right]$$

$$\Lambda_{j,k}^*(\gamma, t^{(k)}) = \sup_{\lambda \in \mathbb{R}} \left\{ \lambda t_k - \Lambda_{j,k-1}(\gamma, t^{(k-1)}; \lambda) \right\} \quad (7)$$

$$\Lambda_{j,k}(\gamma, t^{(k)}; \lambda) = \max \left\{ -\Lambda_{1,k}^*(\gamma, t^{(k)})(j + \lambda), \Lambda_{0,k}^*(\gamma, t^{(k)})(j - 1 + \lambda) \right\}. \quad (8)$$

The operation in (7) is known as the Fenchel–Legendre transform of $\Lambda_{j,k-1}(\gamma, t^{(k-1)}; \lambda)$ [27]. We will be interested in the case where

$$-\text{D}(\mathbb{P}_0^\gamma | \mathbb{P}_1^\gamma) < 0 < \text{D}(\mathbb{P}_1^\gamma | \mathbb{P}_0^\gamma) \quad (9)$$

$$t_1 \in (-\text{D}(\mathbb{P}_0^\gamma | \mathbb{P}_1^\gamma), \text{D}(\mathbb{P}_1^\gamma | \mathbb{P}_0^\gamma)) \quad (10)$$




Fig. 7. Typical plot of $\Lambda_{0,k-1}(\gamma, t^{(k-1)}; \lambda)$, $k \geq 2$.

$$t_k \in \left(-\Lambda_{1,k-1}^*(\gamma, t^{(k-1)}), \Lambda_{0,k-1}^*(\gamma, t^{(k-1)}) \right),$$

for $1 < k \leq h$. (11)

We now provide an inductive argument to show that the above requirements on the thresholds t_k are feasible. From Assumption 1, there exists a $\gamma \in \Gamma$ that satisfies (9), hence the constraint (10) is feasible. Furthermore, the $\Lambda_{j,1}^*(\gamma, t^{(1)})$ are large deviations rate functions and are therefore positive when t_1 satisfies (10) [27]. Suppose now that $k > 1$ and that $\Lambda_{j,k-1}^*(\gamma, t^{(k-1)}) > 0$. From (8), $\Lambda_{j,k-1}(\gamma, t^{(k-1)}; \lambda)$ is the maximum of two linear functions of λ (see Fig. 7). Taking the Fenchel–Legendre transform, and since t_k satisfies (11), we obtain $\Lambda_{j,k}^*(\gamma, t^{(k)}) > 0$, which completes the induction.

From the definitions of $\Lambda_{j,k}$ and $\Lambda_{j,k}^*$, the following relations can be established. The proof consists of straightforward algebraic manipulations and is omitted.

Lemma 2: Suppose that $\gamma \in \Gamma$ satisfies (9), and $t^{(h)}$ satisfies (10) and (11). For $k \geq 1$, we have

$$\Lambda_{1,k}^*(\gamma, t^{(k)}) = \Lambda_{0,k}^*(\gamma, t^{(k)}) - t_k.$$

Furthermore, the supremum in (7) is achieved at some $\lambda \in (-1, 0)$ for $j = 1$, and $\lambda \in (0, 1)$ for $j = 0$. For $k \geq 2$, we have

$$\Lambda_{1,k}^*(\gamma, t^{(k)}) = \frac{\Lambda_{1,k-1}^*(\gamma, t^{(k-1)})(\Lambda_{0,k-1}^*(\gamma, t^{(k-1)}) - t_k)}{\Lambda_{0,k-1}^*(\gamma, t^{(k-1)}) + \Lambda_{1,k-1}^*(\gamma, t^{(k-1)})}$$

$$\Lambda_{0,k}^*(\gamma, t^{(k)}) = \frac{\Lambda_{0,k-1}^*(\gamma, t^{(k-1)})(\Lambda_{1,k-1}^*(\gamma, t^{(k-1)}) + t_k)}{\Lambda_{0,k-1}^*(\gamma, t^{(k-1)}) + \Lambda_{1,k-1}^*(\gamma, t^{(k-1)})}.$$

Proposition 1 that follows, whose proof is provided in the Appendix, will be our main tool in obtaining upper bounds on error probabilities. It shows that the Type I and II error exponents are essentially upper-bounded by $-\Lambda_{0,h}^*(\gamma, t^{(h)})$ and $-\Lambda_{1,h}^*(\gamma, t^{(h)})$, respectively. Recall that $p_n(v)$ is the total number of predecessors of v , $l_n(v)$ is the number of leaves in the subtree rooted at v , and B_n is the set of nodes all of whose immediate predecessors are leaves.

Proposition 1: Fix some $h \geq 1$, and consider a sequence of trees $(T_n)_{n \geq 1}$ such that for all $n \geq n_0$, T_n is h -uniform. Suppose that Assumptions 1 and 2 hold. Suppose that, for every n , every leaf uses the same transmission function $\gamma \in \Gamma$, which satisfies (9), and that every level k node ($k \geq 1$) uses an LLRQ with threshold t_k , satisfying (10) and (11).

(i) For all nodes v of level $k \geq 1$ and for all $n \geq n_0$, we have

$$\frac{1}{l_n(v)} \log \mathbb{P}_1 \left(\frac{S_n(v)}{l_n(v)} \leq t_k \right) \leq -\Lambda_{1,k}^*(\gamma, t^{(k)}) + \frac{p_n(v)}{l_n(v)} - 1$$

$$\frac{1}{l_n(v)} \log \mathbb{P}_0 \left(\frac{S_n(v)}{l_n(v)} > t_k \right) \leq -\Lambda_{0,k}^*(\gamma, t^{(k)}) + \frac{p_n(v)}{l_n(v)} - 1.$$

(ii) Suppose that for all $n \geq n_0$ and all $v \in B_n$, we have $l_n(v) \geq N$. Then, for all $n \geq n_0$, we have

$$\frac{1}{l_n(f)} \log \mathbb{P}_1 \left(\frac{S_n(f)}{l_n(f)} \leq t_h \right) \leq -\Lambda_{1,h}^*(\gamma, t^{(h)}) + \frac{h}{N}$$

$$\frac{1}{l_n(f)} \log \mathbb{P}_0 \left(\frac{S_n(f)}{l_n(f)} > t_h \right) \leq -\Lambda_{0,h}^*(\gamma, t^{(h)}) + \frac{h}{N}.$$

V. OPTIMAL ERROR EXPONENT

In this section, we show that the Type II error probability in a sequence of bounded height trees falls exponentially fast with the number of nodes. We derive sufficient conditions for the error exponent to be the same as that of a parallel configuration. We show that if almost all of the nodes are leaves, i.e., $z = 1$, then $g_P^* = g^* = g_R^*$. The condition $z = 1$ is also equivalent to another condition that requires that the proportion of leaves attached to bounded degree nodes vanishes asymptotically. We also show that under some additional mild assumptions, this sufficient condition is necessary. We start with some graph-theoretic preliminaries.

A. Properties of Trees

In this section, we define various quantities associated with a tree, and derive a few elementary relations that will be used later.

Recall that B_n is the set of non-leaf nodes all of whose predecessors are leaves. (For an h -uniform tree, B_n is the set of all level 1 nodes.) For $N > 0$, let

$$F_{N,n} = \{v \in B_n : l_n(v) \leq N\}$$

$$F_{N,n}^c = \{v \in B_n : l_n(v) > N\} \quad (12)$$

and

$$q_{N,n} = \frac{1}{l_n(f)} \sum_{v \in F_{N,n}} l_n(v) \quad (13)$$

where the sum is taken to be zero if the set $F_{N,n}$ is empty. Let $q_N = \limsup_{n \rightarrow \infty} q_{N,n}$. For a sequence of h -uniform trees, this is the asymptotic proportion of leaves that belong to “small” subtrees in the network.

It turns out that it is easier to work with h -uniform trees. For this reason, we show how to transform any tree of height h to an h -uniform tree.

Height Uniformization Procedure. Consider a tree $T_n = (V_n, E_n)$ of height h , and a node v that has at least one leaf as an immediate predecessor ($v \in A_n$). Let D_n be the set of leaves that are immediate predecessors of v , and whose paths to the fusion center f are of length $k < h$. Add $h - k$ nodes, $\{u_j : j = 1, \dots, h - k\}$, to V_n ; remove the edges (u, v) , for all $u \in D_n$; add the edges (u_1, v) , and (u_{j+1}, u_j) , for $j = 1, \dots, h - k - 1$; add the edges (u, u_{h-k}) , for all $u \in D_n$. This procedure is repeated for all $v \in A_n$. The resulting tree is h -uniform. \square

The height uniformization procedure essentially adds more nodes to the network, and reattaches some leaves, so that the path from every leaf has exactly h hops. Let $(T'_n = (V'_n, E'_n))_{n \geq 1}$ be the new sequence of h -uniform trees obtained from $(T_n)_{n \geq 1}$, after applying the uniformization procedure. (We are abusing notation here in that T'_n typically does not have n nodes, nor is the sequence $|V'_n|$ increasing.) Regarding notation, we adopt the convention that quantities marked with a prime are defined with respect to T'_n .

Note that $l'_n(f) = l_n(f)$. For the case of a relay network, it is seen that any function of the observations at the leaves that can be computed in T'_n can also be computed in T_n . Thus, the detection performance of T'_n is no better than that of T_n . Hence, we obtain

$$g_R^* \leq \limsup_{n \rightarrow \infty} \frac{1}{l'_n(f)} \log \beta^*(T'_n). \quad (14)$$

Therefore, any upper bound derived for h -uniform trees, readily translates to an upper bound for general trees. On the other hand, the coefficients q_N for the h -uniform trees T'_n (to be denoted by q'_N) are different from the coefficients q_N for the original sequence T_n . They are related as follows. The proof is given in the Appendix.

Lemma 3: For any $N, M > 0$, we have

$$q'_N \leq h(Nq_M + N/M).$$

In particular, if $q_N = 0$ for all $N > 0$, then $q'_N = 0$ for all $N > 0$.

It turns out that the condition $z = 1$ is equivalent to the condition $q_N = 0$ for all $N > 0$. The proof is provided in the Appendix.

Lemma 4: We have $z = 1$ iff $q_N = 0$ for all $N > 0$.

B. An Upper Bound

In this subsection, we develop an upper bound on the Type II error probabilities, which takes into account some qualitative properties of the sequence of trees, as captured by q_N .

Lemma 5: Consider an h -uniform sequence of trees $(T_n)_{n \geq 1}$, and suppose that Assumptions 1–3 hold. For every $\epsilon > 0$, there exists some N such that

$$g_R^* \leq (1 - q_N)(g_P^* + \epsilon).$$

Proof: If $g_P^* + \epsilon \geq 0$, there is nothing to prove, since $q_N \leq 1$ and $g_R^* \leq 0$. Suppose that $g_P^* + \epsilon < 0$. Choose $\gamma \in \Gamma$ such that

$$-D(\mathbb{P}_0^\gamma \| \mathbb{P}_1^\gamma) \leq -\sup_{\gamma' \in \Gamma} D(\mathbb{P}_0^{\gamma'} \| \mathbb{P}_1^{\gamma'}) + \frac{\epsilon}{2} = g_P^* + \frac{\epsilon}{2} < 0.$$

Let $t_k = t = -D(\mathbb{P}_0^\gamma \| \mathbb{P}_1^\gamma) + \epsilon/2 \leq g_P^* + \epsilon$, for $k = 1, \dots, h$, and note that

$$-D(\mathbb{P}_0^\gamma \| \mathbb{P}_1^\gamma) < t < 0. \quad (15)$$

Because of (15), we have $\Lambda_{0,1}^*(\gamma, t^{(1)}) > 0$. Furthermore, using Lemma 2, $\Lambda_{1,1}^*(\gamma, t^{(1)}) = \Lambda_{0,1}^*(\gamma, t^{(1)}) - t > -t$. Now let $k \geq 2$, and suppose that $\Lambda_{1,k-1}^*(\gamma, t^{(k-1)}) > -t$ and $\Lambda_{0,k-1}^*(\gamma, t^{(k-1)}) > 0$. From Lemma 2

$$\Lambda_{0,k}^*(\gamma, t^{(k)}) = \frac{\Lambda_{0,k-1}^*(\gamma, t^{(k-1)})(\Lambda_{1,k-1}^*(\gamma, t^{(k-1)}) + t)}{\Lambda_{0,k-1}^*(\gamma, t^{(k-1)}) + \Lambda_{1,k-1}^*(\gamma, t^{(k-1)})} > 0$$

and

$$\Lambda_{1,k}^*(\gamma, t^{(k)}) = \Lambda_{0,k}^*(\gamma, t^{(k)}) - t_k = \Lambda_{0,k}^*(\gamma, t^{(k)}) - t > -t.$$

Hence, by induction, t_k satisfies (10) and (11), so that Proposition 1 can be applied.

Choose N sufficiently large so that $h/N < \Lambda_{0,h}^*(\gamma, t^{(h)})$. If $q_N = 1$, the claimed result holds trivially. Hence, we assume that $q_N \in [0, 1)$. In this case, for n sufficiently large, there exists at least one node in B_n so that $l_n(v) > N$. We remove all nodes $v \in B_n$ with $l_n(v) \leq N$, and their immediate predecessors. Then, we remove all level-2 nodes v that no longer have any predecessors, and so on. In this way, we obtain an h -uniform subtree of T_n , to be denoted by T''_n . (Quantities marked with double primes are defined w.r.t. T''_n .) We have $l''_n(v) > N$ for all $v \in B''_n$, and $l''_n(f) = \sum_{v \in F_{N,n}^c} l_n(v) = l_n(f)(1 - q_{N,n})$. Consider the following relay strategy on the tree T''_n . (Since this is a subtree of T_n , this is also a relay strategy for the tree T_n , with some nodes remaining idle.) The leaves transmit with transmission function γ , and the other nodes use a 1-bit LLRQ with threshold t . (Note that in the definition (5) of the normalized log-likelihood ratio, the denominator $l_n(v)$ now becomes $l''_n(v)$.)

We first show that the strategy just described is admissible. We apply part (ii) of Proposition 1 to T''_n , to obtain

$$\begin{aligned} & \limsup_{n \rightarrow \infty} \frac{1}{l''_n(f)} \log \mathbb{P}_0(Y_f = 1) \\ &= \limsup_{n \rightarrow \infty} \frac{l''_n(f)}{l_n(f)} \cdot \frac{1}{l''_n(f)} \log \mathbb{P}_0(Y_f = 1) \\ &\leq (1 - q_N) \limsup_{n \rightarrow \infty} \frac{1}{l''_n(f)} \log \mathbb{P}_0 \left(\frac{S_n(f)}{l''_n(f)} > t \right) \\ &\leq (1 - q_N) \left(-\Lambda_{0,h}^*(\gamma, t^{(h)}) + \frac{h}{N} \right) < 0, \end{aligned}$$

hence $\mathbb{P}_0(Y_f = 1) \leq \alpha$, when n is sufficiently large.

To bound the Type II error probability, we use Proposition 1 and Lemma 2, to obtain

$$\begin{aligned} g_R^* &\leq \limsup_{n \rightarrow \infty} \frac{1}{l''_n(f)} \log \beta^*(T''_n) \\ &\leq (1 - q_N) \limsup_{n \rightarrow \infty} \frac{1}{l''_n(f)} \log \mathbb{P}_1 \left(\frac{S_n(f)}{l''_n(f)} \leq t \right) \end{aligned}$$

$$\begin{aligned}
&\leq (1 - q_N) \left(-\Lambda_{1,h}^*(\gamma, t^{(h)}) + \frac{h}{N} \right) \\
&= (1 - q_N) \left(t - \Lambda_{0,h}^*(\gamma, t^{(h)}) + \frac{h}{N} \right) \\
&\leq (1 - q_N)t \\
&\leq (1 - q_N)(g_P^* + \epsilon).
\end{aligned}$$

This proves the lemma. \blacksquare

C. Exponential Decay of Error Probabilities

We now establish that Type II error probabilities decay exponentially. The bounded height assumption is crucial for this result. Indeed, for the case of a tandem configuration, the exponential decay property does not seem to hold.

Proposition 2: Consider a sequence of trees of height h , and let Assumptions 1–3 hold. Then

$$-\infty < g_P^* \leq g_R^* < 0 \text{ and } -\infty < -D(\mathbb{P}_0^X \parallel \mathbb{P}_1^X) \leq g^* < 0.$$

Proof: The lower bounds on g_R^* and g^* follow from (2). Note that g_P^* cannot be equal to $-\infty$ because it cannot be better than the error exponent of a parallel configuration in which all the observations are provided uncompressed to the fusion center. The error exponent in the latter case is $-D(\mathbb{P}_0^X \parallel \mathbb{P}_1^X)$, by Stein's lemma, and is finite as a consequence of Assumption 2.

It remains to show that the optimal error exponents are negative. Every tree of height h satisfies $n \leq l_n(f)h + 1$. From (2), we obtain $g^* \leq g_R^*/h$. Therefore, we only need to show that $g_R^* < 0$. As discussed in connection with (14), we can restrict attention to a sequence of h -uniform trees.

We use induction on h . If $h = 1$, we have a parallel configuration and the result follows from [22]. Suppose that the result is true for all sequences of $(h - 1)$ -uniform trees. Consider now a sequence of h -uniform trees. Let $\epsilon > 0$ be such that $g_P^* + \epsilon < 0$. From Lemma 5, there exists some N such that $g_R^* \leq (1 - q_N)(g_P^* + \epsilon)$. If $q_N < 1$, we readily obtain the inequality $g_R^* < 0$.

Suppose now that $q_N = 1$. We only need to consider a sequence $(n_k)_{k \geq 1}$ such that $\lim_{k \rightarrow \infty} q_{N, n_k} = 1$. Using the inequality (22), we have

$$\frac{|F_{N, n_k}|}{l_{n_k}(f)} \geq \frac{q_{N, n_k}}{N}$$

and

$$\liminf_{k \rightarrow \infty} \frac{|F_{N, n_k}|}{l_{n_k}(f)} \geq \frac{1}{N}. \quad (16)$$

For each node $v \in B_n$, we remove all of its immediate predecessors (leaves) except for one, call it u . The leaf u transmits $\gamma(X_u)$ to its immediate successor v . Since node v receives only a single message, it just forwards it to its immediate successor. The resulting performance is the same as if the nodes v in B_n were making a measurement X_v and transmitting $\gamma(X_v)$ to their successor. This is equivalent to deleting all the leaves of T_n to form a new tree, T_n'' , which is $(h - 1)$ -uniform. The preceding argument shows that $\beta^*(T_{n_k}) \leq \beta^*(T_{n_k}'')$.

We have $l_{n_k}''(f) = |B_{n_k}|$, and from (16)

$$\liminf_{k \rightarrow \infty} \frac{|B_{n_k}|}{l_{n_k}(f)} \geq \liminf_{k \rightarrow \infty} \frac{|F_{N, n_k}|}{l_{n_k}(f)} \geq \frac{1}{N}.$$

Therefore

$$\limsup_{k \rightarrow \infty} \frac{1}{l_{n_k}(f)} \log \beta^*(T_{n_k}) \leq \frac{1}{N} \limsup_{k \rightarrow \infty} \frac{1}{l_{n_k}''(f)} \log \beta^*(T_{n_k}'').$$

By the induction hypothesis, the RHS in the above inequality is negative and the proof is complete. \blacksquare

D. Sufficient Conditions For Matching the Performance of the Parallel Configuration

We are now ready to prove the main result of this section. It shows that when $q_N = 0$ for all $N > 0$, or equivalently, when $z = 1$ (cf. Lemma 4), bounded height tree networks match the performance of the parallel configuration.

Proposition 3: Consider a sequence of trees of height h in which $z = 1$, or equivalently $q_N = 0$ for all $N > 0$. Suppose that Assumptions 1–3 hold. Then

$$g_P^* = g^* = g_R^*.$$

Furthermore, if the sequence of trees is h -uniform, the optimal error exponent does not change even if we restrict to relay strategies in which every leaf uses the same transmission function and all other nodes use a 1-bit LLRQ with the same threshold.

Proof: We have shown $g_P^* \leq g_R^*$ in (2). We now prove that $g_R^* \leq g_P^*$. As already explained, there is no loss in generality in assuming that the sequence of trees is h -uniform (by performing the height uniformization procedure, and using Lemma 3).

For any $\epsilon > 0$, Lemma 5 yields

$$g_R^* \leq g_P^* + \epsilon.$$

Letting $\epsilon \rightarrow 0$, we obtain $g_R^* \leq g_P^*$, hence $g_R^* = g_P^*$. From (2) with $z = 1$, we obtain $g^* \leq g_R^* = g_P^*$.

We now show that $g^* \geq g_P^*$. Consider a tree with n nodes, $l_n(f)$ of which are leaves. We will compare it with another sensor network in which $l_n(f)$ nodes v transmit a message $\gamma_v(X_v)$ to the fusion center and $n - l_n(f) - 1$ nodes transmit their raw observations to the fusion center. The latter network can simulate the original network, and therefore its optimal error exponent is at least as good. By a standard argument (similar to the one in Proposition 4 below), the optimal error exponent in the latter network can be shown to be greater than or equal to

$$\limsup_{n \rightarrow \infty} \frac{l_n(f)}{n} g_P^* - \liminf_{n \rightarrow \infty} \frac{n - l_n(f) - 1}{n} D(\mathbb{P}_0^X \parallel \mathbb{P}_1^X) = g_P^*$$

concluding the proof. \blacksquare

Fix an $\epsilon \in (0, -g_P^*)$. For any tree sequence with $z = 1$, we can perform the height uniformization procedure to obtain an h -uniform sequence of trees. In practice, this height uniformization procedure may be performed virtually at each node, so that the tree sequence simulates an h -uniform tree sequence.

A simple strategy on the height uniformized tree sequence that ϵ -achieves the optimal error exponent is a relay strategy in which

- (i) all leaves transmit with the same transmission function $\gamma \in \Gamma$, where γ is chosen such that $-\text{D}(\mathbb{P}_0^\gamma || \mathbb{P}_1^\gamma) \leq g_P^* + \epsilon/2$;
- (ii) all other nodes use 1-bit LLRQs with the same threshold $t = -\text{D}(\mathbb{P}_0^\gamma || \mathbb{P}_1^\gamma) + \epsilon/2$.

Lemmas 3 and 4, and the proof of Lemma 5 show that this relay strategy ϵ -achieves the optimal error exponent $g_R^* = g^* = g_P^*$. This also shows that there is no loss in optimality even if we restrict the relay nodes to use only 1-bit LLRQs. This may be useful in situations where the nodes are simple, low-cost devices.

Proposition 3 provides sufficient conditions for a sequence of trees to achieve the same error exponent as the parallel configuration. We note a few special cases in which these sufficient conditions are satisfied. The first one is the case where there is a finite bound on the number of nodes that are not leaves. In that case, z is easily seen to be 1. This is consistent with the conclusion of Example 1, where a simpler argument was used. The second is the more general case where nodes in B_n are attached to a growing number of leaves, which implies that $q_N = 0$ for all $N > 0$.

Corollary 1: Suppose that Assumptions 1–3 hold. Suppose further that either of the following conditions holds.

- (i) There is a finite bound on the number of nodes that are not leaves.
 - (ii) We have $\min_{v \in B_n} l_n(v) \rightarrow \infty$.
- Then, $g_P^* = g^* = g_R^*$.

The above corollary can be applied to Example 2. In that example, every level-1 node has m leaves attached to it, with m growing large as n increases. Therefore, the tree network satisfies condition (ii) in Corollary 1, and the optimal error exponent is $g^* = g_R^* = g_P^*$. In this case, even if the number N of level-1 nodes grows much faster than m , we still achieve the same error exponent as the parallel configuration. The above proposed strategy, in which every leaf uses the same transmission function, and every node uses the same LLRQ, will nearly achieve the optimal performance.

We are now in a position to determine the optimal error exponent in Example 4.

Example 4, revisited: Recall that in Example 4, every $v_i \in B_n$ has $i + 1$ of predecessors. It is easy to check that $z = 1$. From Proposition 3, the optimal error exponent is the same as that for the parallel configuration. \square

E. Discussion of the Sufficient Conditions

Proposition 3 is unexpected as it establishes that the performance of a tree possessing certain qualitative properties is comparable to that of the parallel configuration. Furthermore, the optimal performance is obtained even if we restrict the nonleaf nodes to use 1-bit LLRQ's. At first sight, it might appear intuitive that if the leaves dominate in a relay tree ($z = 1$), then the tree should always have the same performance as a parallel configuration. However, this intuition is misleading, as this is not the case for a Bayesian formulation, in which both the Type

I and II error probabilities are required to decay at the same rate. To see this, consider the 2-uniform tree in Fig. 3, where every node is constrained to sending 1-bit messages. Suppose we are given nonzero prior probabilities π_0 and π_1 for the hypotheses H_0 and H_1 . Instead of the Neyman–Pearson criterion, suppose that we are interested in minimizing the error exponent

$$\limsup_{n \rightarrow \infty} \frac{1}{l_n(f)} \log P_e^*$$

where P_e^* is the minimum of the error probability $\pi_0 \mathbb{P}_0(Y_f = 1) + \pi_1 \mathbb{P}_1(Y_f = 0)$, optimized over all strategies. It can be shown that to obtain the optimal error exponent, we only need to consider the following two fusion rules: a) the fusion center declares H_0 iff both v_1 and v_2 send a 0, or b) the fusion center declares H_1 iff both v_1 and v_2 send a 1. Then, using the results in [28], the optimal error exponent for this tree network is strictly worse than that for the parallel configuration. Similarly, if we constrain the Type I error in the Neyman–Pearson criterion to decay faster than a predetermined rate, it can be shown that the optimal Type II error exponent for a tree network can be strictly worse than that of a parallel configuration.

Note that the bounded height assumption is essential in proving $g^* = g_R^* = g_P^*$, when $z = 1$. Although our technique can be extended to include those tree sequences whose height grows very slowly compared to n (on the order of $\log |\log(n/l_n(f) - 1)|$), we have not been able to find the optimal error exponent for the general case of unbounded height tree sequences. As noted before, in a tandem network, the Bayesian error probability decays subexponentially fast [26]. The proof of Proposition 2 in [26] involves the construction of a tree network, with unbounded height, and in which $z = 1$. In that proof, it is also shown that such a network has a subexponential rate of error decay. We conjecture that this is also the case for the Neyman–Pearson formulation.

In summary, for a tree network to achieve the same Type II error exponent as a parallel configuration, we require that the tree sequence have a bounded height, satisfy the condition $z = 1$, and that the error criterion be the Neyman–Pearson criterion. Without any one of these three conditions, our results no longer hold.

F. A Necessary Condition For Matching the Performance of the Parallel Configuration

In this subsection, we establish necessary conditions under which a sequence of relay trees with bounded height performs as well as a parallel configuration. As noted in Section III-B, any necessary conditions generally depend on the type of transmission functions available to the relay nodes. However, under an additional condition (Assumption 4), the sufficient condition for $g_R^* = g_P^*$ in Proposition 3 is also necessary.

Proposition 4: Suppose that Assumptions 1, 2, and 4 hold, and $h \geq 2$. If there exists some $N > 0$ such that $q_N > 0$ (equivalently, $z < 1$), then $g_P^* < g_R^*$.

Proof: Fix some $N > 0$ and suppose that $q_N > 0$. Given a tree T_n , we construct a new tree T_n'' , as follows. We remove all nodes other than the leaves and the nodes in $F_{N,n}$. For all the leaves u that are not immediate predecessors of some $v \in F_{N,n}$,

we let u transmit its message directly to the fusion center. We add new edges (v, f) , for each $v \in F_{N,n}$. This gives us a tree T_n'' of height 2, with $l_n''(f) = l_n(f)$ and $q_N'' = q_N$. The latter tree T_n'' can simulate the tree T_n , hence the optimal error exponent associated with the sequence $(T_n)_{n \geq 1}$ is bounded below by the optimal error exponent associated with the sequence $(T_n'')_{n \geq 1}$. Therefore, without loss of generality, we only need to prove the proposition for a sequence of trees of height 2, and in which $F_{N,n} = B_n$, for some $N > 0$ such that $q_N > 0$; we henceforth assume that this is the case. The rest of the argument is similar to the proof of Stein's lemma in Lemma 3.4.7 of [27]. Suppose that a particular admissible relay strategy has been fixed, and let β_n be the associated Type II error probability. Let $\lambda_n = \mathbb{E}_0[S_n(f)]/l_n(f)$. We show that $S_n(f)/l_n(f)$ is close to λ_n in probability. Let D_n be the set of leaves that transmit directly to the fusion center. The proof of the following lemma is in the Appendix.

Lemma 6: For all $\eta > 0$, $\mathbb{P}_0(|S_n(f)/l_n(f) - \lambda_n| > \eta) \rightarrow 0$, as $n \rightarrow \infty$.

We return to the proof of Proposition 4. Given the transmission functions at all other nodes, the fusion center will optimize performance by using an appropriate likelihood ratio test, with a (possibly randomized) threshold. We can therefore assume, without loss of generality, that this is the case. We let ζ_n be the threshold chosen, and note that it must satisfy

$$\mathbb{P}_0(S_n(f)/l_n(f) \leq \zeta_n) \geq 1 - \alpha. \quad (17)$$

From a change of measure argument (see Lemma 3.4.7 in [27]), we have for $\eta > 0$

$$\begin{aligned} & \frac{1}{l_n(f)} \log \beta^*(T_n) \\ & \geq \lambda_n - \eta + \frac{1}{l_n(f)} \log \mathbb{P}_0 \left(\lambda_n - \eta < \frac{S_n(f)}{l_n(f)} \leq \zeta_n \right). \end{aligned}$$

Using (17) and Lemma 6, we see that the last term goes to 0 as $n \rightarrow \infty$. We also have

$$\begin{aligned} \lambda_n &= \frac{1}{l_n(f)} \left(\sum_{v \in D_n} \mathbb{E}_0 \left[\log \frac{d\mathbb{P}_1^v}{d\mathbb{P}_0^v} \right] + \sum_{v \in F_{N,n}} \mathbb{E}_0[L_{v,n}] \right) \\ & \geq (1 - q_{N,n})g_P^* + q_{N,n}K \end{aligned}$$

where, using the notation in Assumption 4,

$$K = \inf_{\substack{1 < k \leq N \\ \xi \in \Gamma(k) \times \Gamma^k}} \frac{1}{k} \mathbb{E}_0 \left[\log \frac{d\nu_1^\xi}{d\nu_0^\xi} \right] > g_P^*.$$

Then, letting $n \rightarrow \infty$, we have

$$g_R^* \geq (1 - q_N)g_P^* + q_N K - \eta$$

for all $\eta > 0$. Taking $\eta \rightarrow 0$ completes the proof. \blacksquare

The condition that there exists a finite N such that $l_n(v) \leq N$ for a nonvanishing proportion of nodes, in the statement of Proposition 4, can be thought of as corresponding to a situation where relay nodes are of two different types: high-cost relays that can process a large number of received messages ($l_n(v) \rightarrow$




Fig. 8. Every node makes a 3-bit observation. Leaves are constrained to sending 1-bit messages, while level-1 nodes are constrained to sending 4-bit messages.

∞) and low-cost relays that can only process a limited number of received messages ($l_n(v) \leq N$ for some small N). From this perspective, Proposition 4 states that a tree network of height greater than one, with a nontrivial proportion of low-cost relays, will always have a performance worse than that of a parallel configuration.

Together with Proposition 3, we have shown the following.

Proposition 5: Suppose that Assumptions 1–4 hold. Then, $g_R^* = g_P^*$ iff $z = 1$ (or, equivalently, iff $q_N = 0$ for all $N > 0$).

We close with an example in which $z < 1$ and $g^* < g_P^*$. Since there are also easy examples where $z < 1$ and $g_P^* < g^*$, this suggests that one can combine them to construct examples where $z < 1$ and $g^* = g_P^*$. Thus, unlike the case of a relay tree, $z = 1$ is not a necessary condition for $g^* = g_P^*$.

Example 5: Consider the tree network shown in Fig. 8, where every node makes a 3-bit observation. Each leaf then compresses its 3-bit observation to a 1-bit message, while each level 1 node is allowed to send a 4-bit message. (Recall that our framework allows for different transmission function sets $\Gamma(d)$ at the different levels.) We assume Assumptions 1–4 hold.

Consider the following strategy: each level 1 node forwards the two 1-bit messages it receives from its two leaves to the fusion center. It then compresses its own 3-bit observation into a 2-bit message before sending it to the fusion center. Using this strategy, the tree network is equivalent to a parallel configuration with $3m$ nodes, $2m$ of which are constrained to sending 1-bit messages, and m of which are constrained to sending 2-bit messages. Clearly, this parallel configuration performs strictly better than one in which all $3m$ nodes are constrained to sending 1-bit messages, therefore, we have $g^* < g_P^*$. \square

Example 5 shows that, unlike the case of relay trees, a tree can outperform a parallel configuration. On the other hand, Example 5 is an artifact of our assumptions. For example, if we restrict every node in this example to sending only 1 bit, the situation is reversed and we have $g_P^* < g^*$. The question of whether a parallel configuration always performs at least as well as a tree network, i.e., whether $g_P^* \leq g^*$, when every node can send the same number of bits, remains open.

VI. CONCLUSION

We have studied the asymptotic detection performance of tree networks with bounded height, under a Neyman–Pearson criterion. Similar to the parallel configuration, we have shown that the optimal Type II error probability decays exponentially

fast with the number of nodes. In addition, we have shown that if leaves dominate (i.e., $l_n(f)/n \rightarrow 1$), the network can achieve the same performance as if all nodes were transmitting directly to the fusion center. We also provided a simple strategy, in which all leaves use the same transmission function, and all other nodes act as 1-bit relays, which achieves the optimal error exponent to any desired accuracy. The sufficient conditions are easy to achieve in cases of practical interest, hence, a system designer can obtain the optimal performance while ensuring that the network is energy efficient. Once the sufficient conditions are satisfied, the architecture of the network no longer affects its detection error exponent. On the other hand, we also showed that for the practically interesting case, where $z = 1$, the sufficient conditions are also necessary. Thus, in a network where the leaves do not dominate, the error decay rate will be worse than that of a parallel configuration, and will actually depend on the particular network architecture.

Needless to say, our conclusions only hold for the particular setting and criterion we have employed. One issue that has not been touched upon is that, with a relay network, a significantly larger value of n may be required before the asymptotic error exponent yields a good approximation. Moreover, in practice, it would be wasteful to have only the leaves make observations, if n is not large enough. Furthermore, under a Bayesian criterion, the same performance as the parallel configuration can no longer be achieved, although exponential decay is still possible [28]. Finally, the more realistic case where the i.i.d. assumption is violated, remains unexplored, with work mainly limited to the parallel configuration [29]–[34].

Future work includes characterizing the asymptotically optimal performance of tree networks without the bounded height constraint. We would like to understand the rate at which the error probability decays, and its dependence on the rate at which the height of the tree increases. Another intriguing question, which has been left unanswered, is whether the inequality $g_P^* \leq g^*$ is always true under the bounded height assumption, when every node is constrained to sending the same number of bits.

APPENDIX

A. Proof of Proposition 1

We first show part (i). The proof proceeds by induction on k . Suppose that $k = 1$, which is equivalent to the well-studied case where all sensors transmit directly to a fusion center. In this case, $p_n(v) = l_n(v)$. Since $t_1 \in (-D(\mathbb{P}_0^\gamma \parallel \mathbb{P}_1^\gamma), D(\mathbb{P}_1^\gamma \parallel \mathbb{P}_0^\gamma))$, from [27, eq. (2.2.13)], we obtain

$$\frac{1}{l_n(v)} \log \mathbb{P}_1 \left(\frac{S_n(v)}{l_n(v)} \leq t_1 \right) \leq -\Lambda_{1,1}^*(\gamma, t_1).$$

The inequality for the Type I error probability follows from a similar argument.

Consider now the induction hypothesis that the result holds for some k . Given a k -uniform tree rooted at v , the induction hypothesis leads to bounds on the probabilities associated with the log-likelihood ratio $L_{v,n}$ of the message Y_v computed at the node v . We use these bounds to obtain bounds on the log-

moment generating function of $L_{v,n}$. Recall that $L_{v,n}$ equals $\mathcal{L}_{v,n}(0)$ whenever $Y_v = 0$, which is the case if and only if $S_n(v)/l_n(v) \leq t_k$. Fix some $\lambda \in [-1, 0]$. We have

$$\begin{aligned} & \frac{1}{l_n(v)} \log \mathbb{E}_1 [e^{\lambda L_{v,n}}] \\ &= \frac{1}{l_n(v)} \log \left[\mathbb{P}_1(Y_v = 0)e^{\lambda \mathcal{L}_{v,n}(0)} + \mathbb{P}_1(Y_v = 1)e^{\lambda \mathcal{L}_{v,n}(1)} \right] \\ &= \frac{1}{l_n(v)} \log \left[\mathbb{P}_1(Y_v = 0)^{1+\lambda} \mathbb{P}_0(Y_v = 0)^{-\lambda} \right. \\ & \quad \left. + \mathbb{P}_1(Y_v = 1)^{1+\lambda} \mathbb{P}_0(Y_v = 1)^{-\lambda} \right] \\ &\leq \frac{1}{l_n(v)} \log \left[\mathbb{P}_1(Y_v = 0)^{1+\lambda} + \mathbb{P}_0(Y_v = 1)^{-\lambda} \right]. \end{aligned}$$

Using the inequality $\log(a + b) \leq \max\{\log(2a), \log(2b)\}$, we obtain

$$\begin{aligned} & \frac{1}{l_n(v)} \log \mathbb{E}_1 [e^{\lambda L_{v,n}}] \\ &\leq \max \left\{ \frac{1+\lambda}{l_n(v)} \log \mathbb{P}_1(Y_v = 0), -\frac{\lambda}{l_n(v)} \log \mathbb{P}_0(Y_v = 1) \right\} \\ & \quad + \frac{\log 2}{l_n(v)} \\ &\leq \max \left\{ -(1+\lambda)\Lambda_{1,k}^*(\gamma, t^{(k)}), \lambda\Lambda_{0,k}^*(\gamma, t^{(k)}) \right\} \\ & \quad + \frac{p_n(v)}{l_n(v)} - 1 + \frac{\log 2}{l_n(v)} \end{aligned} \quad (18)$$

$$\leq \Lambda_{1,k}(\gamma, t^{(k)}; \lambda) + \frac{p_n(v)}{l_n(v)} + \frac{1}{l_n(v)} - 1 \quad (19)$$

where (18) follows from the induction hypothesis.

Consider now a node u at level $k+1$. The subtree rooted at u is a $(k+1)$ -uniform tree. Each level k node $v \in C_n(u)$ can be viewed as the root of a k -uniform tree and (19) can be applied to $L_{v,n}$. From Markov's inequality, and since $\lambda \in [-1, 0]$, we have

$$\mathbb{P}_1 \left(\frac{S_n(u)}{l_n(u)} \leq t_{k+1} \right) \leq e^{-\lambda l_n(u)t_{k+1}} \mathbb{E}_1 [e^{\lambda S_n(u)}]$$

so that

$$\begin{aligned} & \frac{1}{l_n(u)} \log \mathbb{P}_1 \left(\frac{S_n(u)}{l_n(u)} \leq t_{k+1} \right) \\ &\leq -\lambda t_{k+1} + \frac{1}{l_n(u)} \sum_{v \in C_n(u)} \log \mathbb{E}_1 [e^{\lambda L_{v,n}}] \\ &= -\lambda t_{k+1} + \sum_{v \in C_n(u)} \frac{l_n(v)}{l_n(u)} \cdot \frac{1}{l_n(v)} \log \mathbb{E}_1 [e^{\lambda L_{v,n}}] \\ &\leq -\lambda t_{k+1} + \Lambda_{1,k}(\gamma, t^{(k)}; \lambda) \\ & \quad + \sum_{v \in C_n(u)} \frac{p_n(v)}{l_n(u)} + \frac{|C_n(u)|}{l_n(u)} - 1 \end{aligned} \quad (20)$$

$$= -\lambda t_{k+1} + \Lambda_{1,k}(\gamma, t^{(k)}; \lambda) + \frac{p_n(u)}{l_n(u)} - 1 \quad (21)$$

where (20) follows from the induction hypothesis and (19). Taking the infimum over $\lambda \in [-1, 0]$ (cf. Lemma 2), and using (7), we obtain

$$\begin{aligned} \frac{1}{l_n(u)} \log \mathbb{P}_1 \left(\frac{S_n(u)}{l_n(u)} \leq t_{k+1} \right) \\ \leq -\Lambda_{1,k+1}^*(\gamma, t^{(k+1)}) + \frac{p_n(u)}{l_n(u)} - 1. \end{aligned}$$

A similar argument proves the result for the Type I error probability, and the proof of part (i) is complete.

For part (ii), suppose that for all $n \geq n_0$ and all $v \in B_n$, we have $l_n(v) \geq N$. Note that $l_n(f) \geq N|B_n|$. Furthermore, the number of nodes at each level $k \geq 1$ is bounded by $|B_n|$, which yields

$$\frac{p_n(f)}{l_n(f)} - 1 \leq \frac{n}{l_n(f)} - 1 = \frac{n - l_n(f)}{l_n(f)} \leq \frac{h|B_n|}{N|B_n|} = \frac{h}{N}.$$

Applying the results from part (i), with $k = h$, we obtain part (ii).

B. Proof of Lemma 3

We have $l'_n(f) = l_n(f)$. Furthermore, it can be shown that $|B'_n| \leq h|B_n|$. Therefore

$$\begin{aligned} q'_{N,n} &= \frac{1}{l'_n(f)} \sum_{v \in F'_{N,n}} l'_n(v) \leq \frac{1}{l_n(f)} N|B'_n| \\ &\leq \frac{1}{l_n(f)} Nh (|F_{M,n}| + |F_{M,n}^c|) \\ &\leq hNq_{M,n} + hN/M, \end{aligned}$$

where the last inequality follows from

$$|F_{M,n}| \leq \sum_{v \in F_{M,n}} l_n(v) \quad \text{and} \quad |F_{M,n}^c| \leq l_n(f)/M.$$

Taking the limit superior as $n \rightarrow \infty$, we obtain

$$q'_N \leq h(Nq_M + N/M).$$

Suppose that $q_M = 0$ for all $M > 0$. Then for all $N, M > 0$, we have

$$q'_N \leq hN/M.$$

Taking $M \rightarrow \infty$, we obtain the desired result.

C. Proof of Lemma 4

Suppose that $q_N > 0$ for some $N > 0$. Using the inequality

$$q_{N,n} = \frac{1}{l_n(f)} \sum_{v \in F_{N,n}} l_n(v) \leq \frac{N|F_{N,n}|}{l_n(f)}$$

or

$$|F_{N,n}| \geq \frac{q_{N,n}}{N} l_n(f) \quad (22)$$

we obtain

$$\begin{aligned} \frac{l_n(f)}{n} &\leq \frac{l_n(f)}{|F_{N,n}| + l_n(f)} \\ &\leq \frac{l_n(f)}{q_{N,n} l_n(f)/N + l_n(f)} \\ &= \frac{N}{N + q_{N,n}}. \end{aligned}$$

Letting $n \rightarrow \infty$, we obtain

$$z \leq \frac{N}{N + q_N} < 1.$$

For the converse, suppose that $q_N = 0$ for all $N > 0$. It can be seen that each non-leaf node is on a path that connects some $v \in B_n$ to the fusion center. Therefore, the number of non-leaf nodes $n - l_n(f)$ is bounded by $h|B_n|$. We have

$$\frac{n - l_n(f)}{l_n(f)} \leq \frac{h|B_n|}{l_n(f)} = h \frac{|F_{N,n}| + |F_{N,n}^c|}{l_n(f)} \leq hq_{N,n} + \frac{h}{N}.$$

Therefore

$$\limsup_{n \rightarrow \infty} \frac{n - l_n(f)}{l_n(f)} \leq \frac{h}{N}.$$

This is true for all $N > 0$, which implies that $\lim_{n \rightarrow \infty} l_n(f)/n = 1$.

D. Proof of Lemma 6

For each $v \in B_n$, we have $Y_v = \gamma_v(\{\gamma_u(X_u) : u \in C_n(v)\})$, for some $\gamma_v \in \Gamma(l_n(v))$. Using the first and second part of Lemma 1, there exists some $a_1 \in (0, \infty)$, such that

$$\begin{aligned} \mathbb{E}_0[L_{v,n}^2] &\leq \mathbb{E}_0 \left[\left(\sum_{u \in C_n(v)} \log \frac{d\mathbb{P}_1^{\gamma_u}}{d\mathbb{P}_0^{\gamma_u}} \right)^2 \right] + 1 \\ &\leq l_n(v) \mathbb{E}_0 \left[\sum_{u \in C_n(v)} \log^2 \frac{d\mathbb{P}_1^{\gamma_u}}{d\mathbb{P}_0^{\gamma_u}} \right] + 1 \\ &\leq l_n^2(v) a_1 + 1 \\ &\leq l_n^2(v) a \quad (23) \end{aligned}$$

where $a = a_1 + 1$.

To prove the lemma, we use Chebyshev's inequality, and the inequalities $l_n(v) \leq N$ for $v \in F_{N,n}$, and $|D_n| \leq l_n(f)$, to obtain

$$\begin{aligned} \mathbb{P}_0 \left(\left| \frac{S_n(f)}{l_n(f)} - \lambda_n \right| > \eta \right) \\ \leq \frac{1}{\eta^2 l_n^2(f)} \left(\sum_{v \in D_n} \mathbb{E}_0 \left[\log^2 \frac{d\mathbb{P}_1^{\gamma_v}}{d\mathbb{P}_0^{\gamma_v}} \right] + \sum_{v \in F_{N,n}} \mathbb{E}_0[L_{v,n}^2] \right) \end{aligned}$$

$$\leq \frac{1}{\eta^2 l_n^2(f)} \left(\sum_{v \in D_n} a + \sum_{v \in F_{N,n}} l_n^2(v) a \right) \quad (24)$$

$$\begin{aligned} \leq \frac{a}{\eta^2 l_n(f)} + \frac{a}{\eta^2 l_n(f)} \sum_{v \in F_{N,n}} \frac{l_n(v)}{l_n(f)} N \\ \leq \frac{a(1+N)}{\eta^2 l_n(f)} \quad (25) \end{aligned}$$

where (24) follows from Lemma 1 and (23). The RHS of (25) goes to zero as $n \rightarrow \infty$, and the proof is complete.

ACKNOWLEDGMENT

The authors wish to thank the anonymous reviewers for the careful reading of the manuscript and detailed comments that have improved the presentation.

REFERENCES

- [1] R. R. Tenney and N. R. Sandell, "Detection with distributed sensors," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-17, no. 4, pp. 501–510, Jul. 1981.
- [2] Z. Chair and P. K. Varshney, "Optimal data fusion in multiple sensor detection systems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 22, no. 1, pp. 98–101, Jan. 1986.
- [3] G. Polychronopoulos and J. N. Tsitsiklis, "Explicit solutions for some simple decentralized detection problems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 26, no. 2, pp. 282–292, Mar. 1990.
- [4] P. Willett and D. Warren, "The suboptimality of randomized tests in distributed and quantized detection systems," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 355–361, Mar. 1992.
- [5] J. N. Tsitsiklis, "Extremal properties of likelihood-ratio quantizers," *IEEE Trans. Commun.*, vol. 41, no. 4, pp. 550–558, Apr. 1993.
- [6] J. N. Tsitsiklis, "Decentralized detection," *Advances in Statistical Signal Processing*, vol. 2, pp. 297–344, 1993.
- [7] W. W. Irving and J. N. Tsitsiklis, "Some properties of optimal thresholds in decentralized detection," *IEEE Trans. Autom. Control*, vol. 39, no. 4, pp. 835–838, Apr. 1994.
- [8] R. Viswanathan and P. K. Varshney, "Distributed detection with multiple sensors: Part I—Fundamentals," *Proc. IEEE*, vol. 85, no. 1, pp. 54–63, Jan. 1997.
- [9] B. Chen and P. K. Varshney, "A Bayesian sampling approach to decision fusion using hierarchical models," *IEEE Trans. Signal Process.*, vol. 50, no. 8, pp. 1809–1818, Aug. 2002.
- [10] B. Chen and P. K. Willett, "On the optimality of the likelihood-ratio test for local sensor decision rules in the presence of nonideal channels," *IEEE Trans. Inf. Theory*, vol. 51, no. 2, pp. 693–699, Feb. 2005.
- [11] A. Kashyap, "Comments on the optimality of the likelihood-ratio test for local sensor decision rules in the presence of nonideal channels," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 1274–1275, Mar. 2006.
- [12] B. Liu and B. Chen, "Channel-optimized quantizers for decentralized detection in sensor networks," *IEEE Trans. Inf. Theory*, vol. 52, no. 7, pp. 3349–3358, Jul. 2006.
- [13] L. K. Ekchian and R. R. Tenney, "Detection networks," in *Proc. IEEE Conf. Decision and Control*, Orlando, FL, Dec. 1982, pp. 686–691.
- [14] R. Viswanathan, S. C. A. Thomopoulos, and R. Tumuluri, "Optimal serial distributed decision fusion," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 24, no. 4, pp. 366–376, Jul. 1988.
- [15] A. R. Reibman and L. W. Nolte, "Design and performance comparison of distributed detection networks," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 23, no. 6, pp. 789–797, Nov. 1987.
- [16] Z. B. Tang, K. R. Pattipati, and D. L. Kleinman, "Optimization of detection networks: Part I—Tandem structures," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, no. 5, pp. 1044–1059, Sep. 1991.
- [17] Z. B. Tang, K. R. Pattipati, and D. L. Kleinman, "Optimization of detection networks: Part II—Tree structures," *IEEE Trans. Syst., Man, Cybern.*, vol. 23, no. 1, pp. 211–221, Jan. 1993.
- [18] J. D. Papastavrou and M. Athans, "On optimal distributed decision architectures in a hypothesis testing environment," *IEEE Trans. Autom. Control*, vol. 37, no. 8, pp. 1154–1169, Aug. 1992.
- [19] A. Pete, K. Pattipati, and D. Kleinman, "Optimization of detection networks with multiple event structures," *IEEE Trans. Autom. Control*, vol. 39, no. 8, pp. 1702–1707, Aug. 1994.
- [20] S. Alhakeem and P. K. Varshney, "A unified approach to the design of decentralized detection systems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 31, no. 1, pp. 9–20, Jan. 1995.
- [21] Y. Lin, B. Chen, and P. K. Varshney, "Decision fusion rules in multi-hop wireless sensor networks," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 2, pp. 475–488, Apr. 2005.
- [22] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control, Signals, Syst.*, vol. 1, pp. 167–182, 1988.
- [23] M. E. Hellman and T. M. Cover, "Learning with finite memory," *Ann. Math. Statist.*, vol. 41, no. 3, pp. 765–782, 1970.
- [24] T. M. Cover, "Hypothesis testing with finite statistics," *Ann. Math. Statist.*, vol. 40, no. 3, pp. 828–835, 1969.
- [25] J. D. Papastavrou and M. Athans, "Distributed detection by a large team of sensors in tandem," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 639–653, May 1992.
- [26] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "On the sub-exponential decay of detection probabilities in long tandems," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Honolulu, HI, Apr. 2007, pp. 837–840.
- [27] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. New York: Springer-Verlag, 1998.
- [28] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win, "Bayesian detection in bounded height tree networks," in *Proc. Data Compression Conf.*, Snowbird, UT, Mar. 2007, pp. 243–252.
- [29] E. Drakopoulos and C. C. Lee, "Optimum multisensor fusion of correlated local decisions," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 27, no. 4, pp. 593–606, Jul. 1991.
- [30] M. Kam, Q. Zhu, and W. S. Gray, "Optimal data fusion of correlated local decisions in multiple sensor detection systems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 28, no. 3, pp. 916–920, May 1992.
- [31] R. S. Blum and S. A. Kassam, "Optimum distributed detection of weak signals in dependent sensors," *IEEE Trans. Inf. Theory*, vol. 38, no. 3, pp. 1066–1079, May 1992.
- [32] R. S. Blum, S. A. Kassam, and H. V. Poor, "Distributed detection with multiple sensors: Part II—Advanced topics," *Proc. IEEE*, vol. 85, no. 1, pp. 64–79, Jan. 1997.
- [33] J.-F. Chamberland and V. V. Veeravalli, "How dense should a sensor network be for detection with correlated observations?," *IEEE Trans. Inf. Theory*, vol. 52, no. 11, pp. 5099–5106, Nov. 2006.
- [34] W. Li and H. Dai, "Distributed detection in large-scale sensor networks with correlated sensor observations," in *Proc. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2005.