DECENTRALIZED DETECTION IN UNDIRECTED NETWORK TOPOLOGIES

O. Patrick Kreidl and Alan S. Willsky

Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science
77 Massachusetts Avenue, Cambridge, MA 02139 U.S.A.

ABSTRACT

Consider the well-studied decentralized Bayesian detection problem with the twist of an *undirected* network topology, each edge representing a bidirectional (and perhaps unreliable) finite-rate communication link between two distributed sensor nodes. Every node operates in parallel, processing any particular local measurement in two (discrete) decision stages: the first selects the symbols (if any) transmitted to its immediate neighbors and the second, upon receiving the symbols (or lack thereof) from the same neighbors, decides the value of its local state. We adapt the team solution already known for directed acyclic networks and establish conditions such that the iterative numerical algorithm to collectively optimize the local decision rules admits an efficient message-passing interpretation, featuring an asynchronous distributed implementation in which total computation and communication overhead scales only linearly with the number of nodes. In sharp contrast to the directed case, this message-passing algorithm retains its global correctness and convergence guarantees without restrictions on the network topology.

***Index Terms*—** Bayes procedures, Distributed algorithms, Distributed detection, Markov processes, Message passing

1. INTRODUCTION

Existing research literature on the decentralized detection problem, formally introduced in [1], focuses almost exclusively on the case of unidirectional inter-sensor communication defined on a directed network topology [2, 3]. The analysis is particularly well-developed for a binary hypothesis test in which a set of remote sensors each transmit a finite-alphabet symbol to a common "fusion center" that, in turn, makes the minimum-error team decision. So-called person-by-person optimality conditions analytically reduce to solving a set of nonlinear fixed-point equations, providing a certain conditional independence assumption is upheld [4]. These team fixed-point equations are known to apply for any directed acyclic network topology; moreover, in the special case of a polytree topology, the iterative algorithm for solving these equations offline (i.e., before processing actual measurements) admits an efficient message-passing interpretation [5–7]. Specifically, in each offline iteration, every sensor node adjusts its local decision rule (for subsequent online processing) based on incoming messages from its neighbors and, in turn, sends adjusted outgoing messages to its neighbors. The messages from upstream neighbors define a "likelihood function" for the symbols the node may receive online (e.g., "what does the information from my parents mean to me"), while the messages received from downstream neighbors define a "cost-to-go function" for the symbols the node may transmit online (e.g., "what does the information from me mean to my children, their children and so on"). Each node need only be initialized with local statistics and iterative local computation is invariant to the number of nodes (yet scales exponentially with the number of neighbors, so large networks are taken to be sparsely-connected) [7].

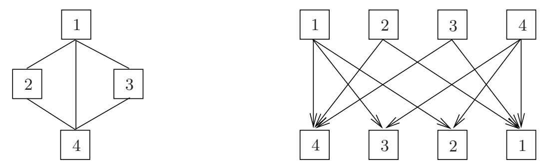

(a) Undirected Topology (b) Equivalent Directed Topology

Fig. 1. Illustration of the key step in our analysis of the decentralized detection problem with bidirectional inter-sensor communication: (a) an undirected topology and (b) its directed counterpart.

We adapt the Bayesian formulation in decentralized detection for the case of bidirectional inter-sensor communication defined on an *undirected* network topology. An undirected architecture is arguably more compatible with emerging concepts in ad-hoc networking [8, 9], since designating a common fusion center or enforcing a directed acyclic topology "on the fly" may require expensive nonlocal coordination among the distributed nodes [10–14]. The key step in our analysis is illustrated in Fig. 1, where we "unravel" the bidirectional communication implied by an undirected topology into an equivalent directed topology in which each node appears both as a transmitter and a receiver. Though the resulting directed network is a polytree, because the node replication violates the critical conditional independence assumption, we cannot readily conclude that the tractable solution known for directed networks is applicable. Our analysis proves it is applicable if the Bayesian cost function is separable across the nodes: specifically, under both the conditional independence and separable cost assumptions, the decision rules at every node reduce to a pair of local likelihood ratio tests. Moreover, the forward-backward message schedule defined on the directed topology translates into a parallel message schedule defined on the undirected topology: in each offline iteration, every node exchanges both types of messages with all of its neighbors, firstly adjusting its stage-one rule and outgoing "likelihood" messages, then

This work was supported by U.S. ARO under MURI grants W911NF-05-1-0207 and W911NF-06-1-0076.

Copyright 2006 IEEE. Published in the IEEE 2007 Statistical Signal Processing Workshop (SSP 2007), scheduled for August 26-29, 2007 in Madison, Wisconsin. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works, must be obtained from the IEEE. Contact: Manager, Copyrights and Permissions / IEEE Service Center / 445 Hoes Lane / P.O. Box 1331 / Piscataway, NJ 08855-1331, USA. Telephone: + Intl. 908-562-3966.

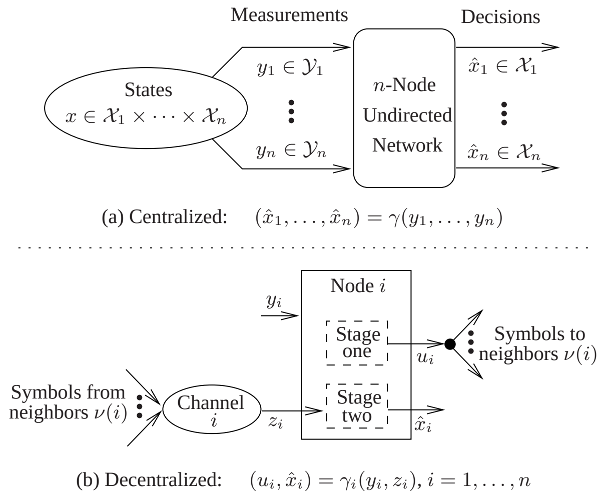

Fig. 2. Our global detection model assuming (a) a centralized strategy and (b) a decentralized strategy for processing measurements.

adjusting its stage-two rule and outgoing "cost-to-go" messages. In sharp contrast to the directed case [7], this message-passing algorithm retains its correctness and convergence guarantees without restrictions on the network topology.

The remainder of this paper is organized as follows. Section 2 and Section 3 adapt the formulation and analysis, respectively, of the Bayesian detection problem starting with an undirected network topology. For brevity, we omit proofs and emphasize the aspects that are in clear contrast with what is known for the more conventional case of directed acyclic networks. We conclude in Section 4 with suggestions for future research, looking towards distributed detection applications in which performance specifications render the particular processing constraints considered here too severe.

2. BAYESIAN FORMULATION

Our centralized and decentralized processing models are illustrated in Fig. 2, both assuming that (i) the hidden state x and observable measurement y take their values in, respectively, a discrete vector space $\mathcal{X} = \mathcal{X}_1 \times \cdots \times X_n$ and Euclidean vector space $\mathcal{Y} = \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_n$; and (ii) the network is to generate a global estimate $\hat{x} \in \mathcal{X}$ based on the spatially-distributed measurement vector $y \in \mathcal{Y}$. The decentralized model in Fig. 2(b) further assumes each component of $\hat{x}$ is determined by an individual sensor in the n-node network defined by an undirected graph $\mathcal{G}$, each edge (i, j) indicating a (perhaps unreliable) low-rate bidirectional communication link between nodes i and j. Every node i, initially observing only the component measurement y_i, operates in two distinct stages: the first stage decides upon the symbols u_i (if any) transmitted to its neighbors $\nu(i) = \{j \mid \text{edge } (i, j) \text{ in } \mathcal{G}\}$ and the second stage, upon receiving the symbols z_i (or lack thereof) from the same neighbors, decides upon the local estimate $\hat{x}_i$. The collections of transmitted symbols u and received symbols z take their values in discrete vector spaces $\mathcal{U} = \mathcal{U}_1 \times \cdots \times \mathcal{U}_n$ and $\mathcal{Z} = \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_n$, respectively.

The probabilistic model starts with a distribution $p(x, y)$ that jointly describes the hidden state process X and noisy measurement process Y. The decentralized decision process implied by undirected graph $\mathcal{G}$ consists of just two stages, all nodes operating in parallel with every node i generating firstly its stage-one decision U_i, observing only the local measurement Y_i, and then its stage-two decision $\hat{X}_i$ upon also receiving the information Z_i described by a conditional distribution $p(z_i|x, y, u_{\nu(i)})$ based on the information $U_{\nu(i)} = \{U_j \mid j \in \nu(i)\}$ collectively transmitted by neighbors $\nu(i)$ in $\mathcal{G}$. Altogether, any particular *strategy* $\gamma : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathcal{U} \times \mathcal{X}$ induces a global decision process $(U, \hat{X}) = \gamma(Y, Z)$. Denote by Γ the set of all such strategies and by $\Gamma^{\mathcal{G}}$ the admissible subset of decentralized strategies implied by $\mathcal{G}$: specifically, denoting by $\mathcal{M}_i$ all stage-one communication rules of the form $\mu_i : \mathcal{Y}_i \rightarrow \mathcal{U}_i$ and by Δ_i all stage-two detection rules of the form $\delta_i : \mathcal{Y}_i \times \mathcal{U}_i \times \mathcal{Z}_i \rightarrow \mathcal{X}_i$, define $\Gamma_i = \mathcal{M}_i \times \Delta_i$ for each node i and let $\Gamma^{\mathcal{G}} = \Gamma_1 \times \cdots \times \Gamma_n$.

The final ingredient to the formulation is the real-valued function

$$c(u, \hat{x}, x) = c(\hat{x}, x) + \lambda c(u, x),$$

expressing a (weighted) sum of the costs incurred over the two stages. Any fixed strategy $\gamma \in \Gamma$ is then penalized by its expected cost,

$$J(\gamma) = E\left[c(U, \hat{X}, X)\right] = E\left[E\left[c(\gamma(Y, Z), X)|\, Y, Z\right]\right]. \quad (1)$$

In turn, the decentralized design problem is expressed by

$$J(\gamma^*) = \min_{\gamma \in \Gamma} J_d(\gamma) + \lambda J_c(\gamma) \text{ subject to } \gamma \in \Gamma^{\mathcal{G}}, \quad (2)$$

where the functions $J_c : \Gamma \rightarrow \mathbb{R}$ and $J_d : \Gamma \rightarrow \mathbb{R}$ quantify the stage-one *communication penalty* and stage-two *detection penalty*, respectively. By construction, fixing the rules $\gamma_i = (\mu_i, \delta_i)$ local to node i is equivalent to specifying the distribution

$$p(u_i, \hat{x}_i|y_i, z_i; \gamma_i) = p(u_i|y_i; \mu_i)p(\hat{x}_i|y_i, u_i, z_i; \delta_i),$$

reflecting the two-stage causal processing implied by $\mathcal{G}$. It follows that fixing a strategy $\gamma \in \Gamma^{\mathcal{G}}$ specifies the distribution

$$p(u, z, \hat{x}|x, y; \gamma) = \prod_{i=1}^{n} p\left(z_i|x, y, u_{\nu(i)}\right) p(u_i, \hat{x}_i|y_i, z_i; \gamma_i)$$

and, in turn, the distribution that determines $J(\gamma)$ in (1) becomes

$$p(u, \hat{x}, x; \gamma) = \int_{y \in \mathcal{Y}} p(x, y) \prod_{i=1}^{n} p(u_i, \hat{x}_i|x, y, u_{\nu(i)}; \gamma_i)\, dy,$$

where the summation over $\mathcal{Z}$ is taken inside the product i.e,

$$p(u_i, \hat{x}_i|x, y, u_{\nu(i)}; \gamma_i) = \sum_{z_i \in \mathcal{Z}_i} p(z_i|x, y, u_{\nu(i)})p(u_i, \hat{x}_i|y_i, z_i; \gamma_i)$$

for each node i.

3. TEAM SOLUTION

This section summarizes the results of applying the team-theoretic analysis already known for directed network topologies (as presented in [7], for example) to the problem formulated in Section 2. As already depicted in Fig. 1, the key idea is to map the set $\Gamma^{\mathcal{G}}$ of all two-stage strategies defined on an undirected topology $\mathcal{G}$ into an equivalent set of strategies defined on a particular two-level directed topology. In contrast to known results for the directed case, our first result is a negative one: specifically, the usual conditional independence assumption does *not* by itself imply that the team-optimal strategy admits a finite parameterization. Our second result establishes that another assumption is needed, namely that the Bayesian cost function is separable across the nodes. When both assumptions hold, the team optimality conditions reduce analytically to a nonlinear fixed-point equation with identical structure to that which arises for the "unraveled" directed counterpart. In turn, the forward-backward message-passing algorithm developed for directed polytrees immediately applies, translating into a parallel message-passing algorithm on the original (and perhaps "loopy") undirected topology.

The basic idea of viewing bidirectional inter-sensor communication as a sequence of unidirectional decision stages has appeared in earlier research literature. A detailed analysis of two sensor nodes performing a global binary hypothesis test appears in [15]. Their model assumes one node is a primary decision-maker and the other acts as a (costly) consultant, the latter only providing input when the former explicitly requests for it. Indeed, their formulation satisfies the assumptions we require for tractability of the two-stage team solution in arbitrary n-node network topologies (and our analysis also accounts for the possibility of unreliable links). More general topologies or more than two decision stages (but still for a global binary hypothesis test and with reliable links) are considered in [16–19], but distinctly assuming that each node processes only a new measurement in every stage, essentially "forgetting" all of its preceding measurements and preserving the critical conditional independence assumption. Our formulation, however, assumes each node processes the same local measurement over successive decision stages, and it is interesting that this subtle difference in the model leads to the dramatic algorithmic ramifications of our negative result.

3.1. Necessary Optimality Conditions

We begin the team-theoretic analysis for the detection model in Fig. 2 by showing that the usual conditional independence assumption is not enough to guarantee that the optimal decentralized strategy γ^* in (2) admits a finite parameterization. In the directed case, however, under this assumption the global minimizer γ^* in (2) reduces to a collection of likelihood-ratio tests, each local rule of the form

$$\gamma_i^*(Y_i, Z_i) = \arg \min_{(u_i,\hat{x}_i)\in\mathcal{U}_i\times\mathcal{X}_i} \sum_{x\in\mathcal{X}} \theta_i^*(u_i, \hat{x}_i, x; Z_i)p(Y_i|x), \quad (3)$$

with rule parameters $\theta_i \in \mathbb{R}^{|\mathcal{U}_i\times\mathcal{X}_i\times\mathcal{X}\times\mathcal{Z}_i|}$ local to each node i coupled to the parameters $\theta_{-i} = \{\theta_j; j \neq i\}$ at all other nodes via a nonlinear fixed-point equation,

$$\theta_i = f_i(\theta_{-i}), \qquad i = 1, \ldots, n. \quad (4)$$

That parameter vector $\theta = (\theta_1, \ldots, \theta_n)$ is finite-dimensional is key to the known theoretical correctness and convergence guarantees for block-cyclic Gauss-Seidel iterative solutions of (4).

Assumption 1 (Conditional Independence) *Conditioned on state $X = x$, the measurement Y_i and received symbol Z_i local to node i are mutually independent as well as independent of all other information observed in the network, namely the measurements Y_{-i} and symbols Z_{-i} local to all other nodes i.e., for every i,*

$$p(y_i, z_i|x, y_{-i}, z_{-i}, u_{\nu(i)}) = p(y_i|x)p(z_i|x, u_{\nu(i)}).$$

Proposition 1 *Let Assumption 1 hold. Consider any particular node i and assume both rules local to all other nodes are fixed at their optimal values in (2), which we denote by $\gamma_{-i}^* = \{\gamma_j^* \in \Gamma_j \mid j \neq i\}$.*

• *Assume the stage-two rule local to node i is fixed at its optimal value $\delta_i^* \in \Delta_i$. The optimal stage-one rule reduces to*

$$\mu_i^*(Y_i) = \arg \min_{u_i\in\mathcal{U}_i} \sum_{x\in\mathcal{X}} a_i^*(u_i, x; Y_i)p(Y_i|x),$$

where the parameter values $a_i^* \in \mathbb{R}^{|\mathcal{U}_i\times\mathcal{X}\times\mathcal{Y}_i|}$ depend on all other fixed rules through a nonlinear operator f_i^1 of the form

$$a_i^* = f_i^1(\delta_i^*, \gamma_{-i}^*). \quad (5)$$

• *Assume the stage-one rule local to node i is fixed at its optimal value $\mu_i^* \in \mathcal{M}_i$. The optimal stage-two rule reduces to*

$$\delta_i^*(Y_i, U_i, Z_i) = \arg \min_{\hat{x}_i\in\mathcal{X}_i} \sum_{x\in\mathcal{X}} b_i^*(\hat{x}_i, x; U_i, Z_i)p(Y_i|x),$$

where parameter values $b_i^* \in \mathbb{R}^{|\mathcal{X}_i\times\mathcal{X}\times\mathcal{U}_i\times\mathcal{Z}_i|}$ depend upon all other fixed rules through a nonlinear operator f_i^2 of the form

$$b_i^* = f_i^2(\mu_i^*, \gamma_{-i}^*). \quad (6)$$

It is instructive to contrast each part of Proposition 1 with (3) in a directed network topology. The only difference in the stage-two rule δ_i^* is that the stage-one communication decision U_i acts as side information (in addition to local channel information Z_i); in particular, parameters b_i^* depend only on local measurement y_i through the discrete symbol $u_i = \mu_i^*(y_i)$, so the likelihood function $p(Y_i|x)$ is the sufficient statistic for measurement Y_i. However, in the stage-one rule μ_i^*, parameters a_i^* are seen to depend explicitly on the local measurement y_i, implying a likelihood-ratio test need not be team-optimal. That is, Proposition 1 tells us that (4) still applies given the undirected model, taking $\theta_i = (a_i, b_i)$ and $f_i = (f_i^1, f_i^2)$, but that the parameter vector θ need not necessarily be finite-dimensional. This theoretical complication is equivalent to that arising when Assumption 1 is violated for even the simplest directed networks (e.g., two nodes in series), in which case the decentralized design problem is known to be NP-complete [4]. Thus, Proposition 1 suggests the same for the problem in (2), even when Assumption 1 is upheld. Indeed, this negative result was already anticipated from the node replication discussed in Fig. 1. It is more interesting that continuing the analysis leads to a positive result: the efficient message-passing interpretation is applicable despite the node replication issue and without restrictions on the undirected network topology.

Assumption 2 (Separable Costs) *The global cost function in both stages of the decision process is additive over nodes of the network,*

$$c(u, \hat{x}, x) = \sum_{i=1}^{n} \left[c(\hat{x}_i, x) + \lambda c(u_i, x)\right].$$

Proposition 2 *If Assumption 2 holds, then Proposition 1 applies with (5) and (6) specialized to the proportionalities*

$$a_i^*(u_i, x; y_i) \propto \alpha_i^*(u_i, x) = p(x)\left[\lambda c(u_i, x) + C_i^*(u_i, x)\right]$$

and

$$b_i^*(\hat{x}_i, x; u_i, z_i) \propto \beta_i^*(\hat{x}_i, x; z_i) = p(x)P_i^*(z_i|x)c(\hat{x}_i, x),$$

respectively, where (i) the likelihood function $P_i(z_i|x)$ for received information Z_i depends upon the fixed stage-one rules in the immediate neighborhood $\nu(i)$ through a nonlinear operator g_i of the form

$$P_i^*(z_i|x) = g_i(\mu_{\nu(i)}^*)$$

and (ii) the cost-to-go function $C_i(u_i, x)$ for transmitted information U_i depends upon the fixed stage-two decision rules in the immediate neighborhood as well as the fixed stage-one decision rules in the two-step neighborhood, denoted by $\nu^2(i) = \bigcup_{j\in\nu(i)} \nu(j) - i$, through a nonlinear operator h_i of the form

$$C_i^*(u_i, x) = h_i(\mu_{\nu^2(i)}^*, \delta_{\nu(i)}^*).$$

3.2. Message-Passing Algorithm

The equations of Proposition 2 are in fact equivalent to those known for the "unraveled" $2n$-node directed (polytree) network in which (parentless) nodes 1 to n employ the rules μ_1^* to μ_n^*, while (childless) nodes $n+1$ to $2n$ employ the rules δ_1^* to δ_n^*. Hence, the efficient message-passing interpretation presented in [7] for directed networks is readily applicable.

Assumption 3 (Measurement/Channel/Cost Locality) *In addition to the conditions of Assumption 1 and Assumption 2, the measurement and channel models as well as both stages of the cost function local to node i do not directly depend on any of the non-local state processes X_{-i} i.e.,*

$$p(y_i, z_i|x, y_{-i}, z_{-i}, u_{\nu(i)}) = p(y_i|x_i)p(z_i|x_i, u_{\nu(i)})$$

and

$$c(u_i, \hat{x}_i, x) = c(\hat{x}_i, x_i) + \lambda c(u_i, x_i).$$

Corollary 1 *If Assumption 3 holds, then Proposition 2 reduces to*

$$\mu_i^*(Y_i) = \arg\min_{u_i \in \mathcal{U}_i} \sum_{x_i \in \mathcal{X}_i} \alpha_i^*(u_i, x_i) p(Y_i|x_i)$$

with stage-one rule parameters $\alpha_i^* \in \mathbb{R}^{|\mathcal{U}_i \times \mathcal{X}_i|}$ given by

$$\alpha_i^*(u_i, x_i) \propto p(x_i) \left[\lambda c(u_i, x_i) + \sum_{j \in \nu(i)} C_{j \to i}^*(u_i, x_i) \right], \qquad (7)$$

and

$$\delta_i^*(Y_i, Z_i) = \arg\min_{\hat{x}_i \in \mathcal{X}_i} \sum_{x_i \in \mathcal{X}_i} \beta_i^*(\hat{x}_i, x_i; Z_i) p(Y_i|x_i)$$

with stage-two rule parameters $\beta_i^* \in \mathbb{R}^{|\mathcal{X}_i \times \mathcal{X}_i \times \mathcal{Z}_i|}$ given by

$$\begin{aligned} \beta_i^*(\hat{x}_i, x_i; z_i) \propto c(\hat{x}_i, x_i) \sum_{x_{\nu(i)}} p(x_i, x_{\nu(i)}) \times \\ \sum_{u_{\nu(i)}} p(z_i|x_i, u_{\nu(i)}) \prod_{j \in \nu(i)} P_{j \to i}^*(u_j|x_j) \end{aligned}, \qquad (8)$$

where each node i produces both a likelihood message for every neighbor $j \in \nu(i)$ given by

$$P_{i \to j}^*(u_i|x_i) = \int_{y_i} p(y_i|x_i) p(u_i|y_i; \mu_i^*) dy_i \qquad (9)$$

and a cost-to-go message for each neighbor $j \in \nu(i)$ given by

$$\begin{aligned} C_{i \to j}^*(u_j, x_j) = \sum_{x_i} \sum_{\hat{x}_i} c(\hat{x}_i, x_i) \sum_{x_{\nu(i)-j}} p(x_i, x_{\nu(i)}|x_j) \times \\ \sum_{u_{\nu(i)-j}} p(\hat{x}_i|x_i, u_{\nu(i)}; \delta_i^*) \prod_{m \in \nu(i)-j} P_{m \to i}^*(u_m|x_m), \end{aligned} \qquad (10)$$

$$\begin{aligned} p(\hat{x}_i|x_i, u_{\nu(i)}; \delta_i^*) = \\ \sum_{z_i} p(z_i|x_i, u_{\nu(i)}) \int_{y_i} p(y_i|x_i) p(\hat{x}_i|y_i, z_i; \delta_i^*) dy_i. \end{aligned}$$

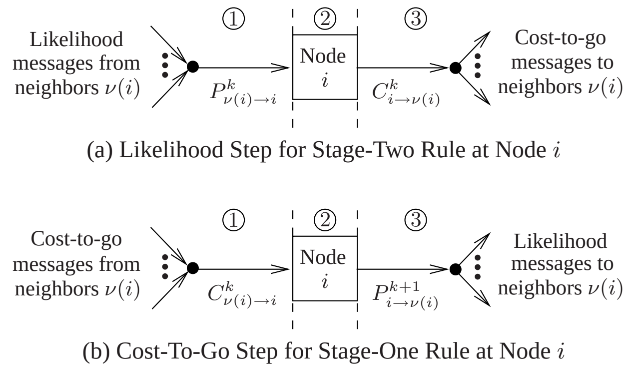

Fig. 3. The kth parallel message-passing iteration as discussed in Corollary 2, each node i interleaving its purely-local computations with only nearest-neighbor communications.

Corollary 1 implies that the rule parameters (α_i^*, β_i^*) local to node i are completely determined by the incoming messages from neighbors $\nu(i)$ on the original undirected network topology $\mathcal{G}$. Specifically, we see in (8) that the stage-two parameters β_i^* depend upon the incoming likelihood messages $P_{\nu(i) \to i}^* = \{P_{j \to i}^*; j \in \nu(i)\}$, the right-hand-side summarized by operator $f_i^2(P_{\nu(i) \to i}^*)$. Meanwhile, we see in (7) that the stage-one parameters α_i^* depend upon the incoming cost-to-go messages $C_{\nu(i) \to i}^* = \{C_{j \to i}^*; j \in \nu(i)\}$, the right-hand-side summarized by operator $f_i^1(C_{\nu(i) \to i}^*)$. Similarly, the satisfaction of Corollary 1 at all nodes other than i depends upon the outgoing messages from node i to its neighbors $\nu(i)$. The outgoing likelihood messages $P_{i \to \nu(i)}^* = \{P_{i \to j}^*; j \in \nu(i)\}$ expressed in (9) are summarized by operator $g_i(\alpha_i^*)$, while the outgoing cost-to-go messages $C_{i \to \nu(i)}^* = \{C_{i \to j}^*; j \in \nu(i)\}$ expressed in (10) are summarized by operator $h_i(\beta_i^*, P_{\nu(i) \to i}^*)$. Altogether, we see that Corollary 1 specializes the nonlinear fixed-point equations in (4) to

$$\left. \begin{array}{rcl} \beta_i &=& f_i^2\left(P_{\nu(i) \to i}\right) \\ \alpha_i &=& f_i^1\left(C_{\nu(i) \to i}\right) \\ P_{i \to \nu(i)} &=& g_i\left(\alpha_i\right) \\ C_{i \to \nu(i)} &=& h_i\left(\beta_i, P_{\nu(i) \to i}\right) \end{array} \right., \quad i = 1, \ldots, n. \qquad (11)$$

Corollary 2 *Initialize stage-one parameters $\alpha^0 = (\alpha_1^0, \ldots, \alpha_n^0)$ and stage-two parameters $\beta^0 = (\beta_1^0, \ldots, \beta_n^0)$, then generate the sequence $\{(\alpha^k, \beta^k); k = 1, 2, \ldots\}$ by iterating (11) in a parallel message schedule defined on the undirected graph $\mathcal{G}$, each node interleaving local updates of stage-one and stage-two decision rules with nearest-neighbor exchanges of likelihood and cost-to-go messages e.g., as illustrated in Fig. 3,*

$$\begin{array}{rcll} P_{i \to \nu(i)}^k &:=& g_i\left(\alpha_i^{k-1}\right) & \textit{from } i = 1, \ldots, n, \\ \beta_i^k &:=& f_i^2\left(P_{\nu(i) \to i}^k\right) & \textit{from } i = 1, \ldots, n, \\ C_{i \to \nu(i)}^k &:=& h_i\left(\beta_i^k, P_{\nu(i) \to i}^k\right) & \textit{from } i = 1, \ldots, n \textit{ and} \\ \alpha_i^k &:=& f_i^1\left(C_{\nu(i) \to i}^k\right) & \textit{from } i = 1, \ldots, n. \end{array}$$

If Assumption 3 holds, the associated sequence $\{J(\gamma^k)\}$ converges.

Note that Corollary 2 applies without restrictions on the undirected topology, whereas a directed topology must be a polytree for

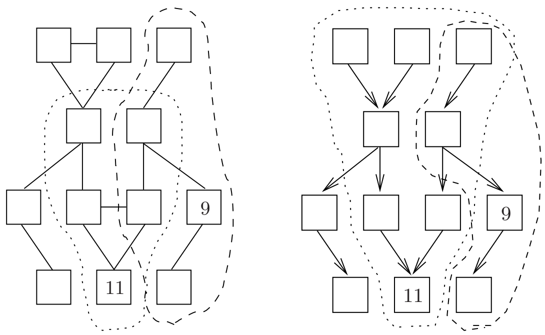

(a) Undirected Topology (b) Directed Topology

Fig. 4. Comparison of the team coupling captured by the offline message-passing algorithm in (a) an undirected network or (b) a directed network. In (a), the incoming messages for each rule depend directly only on the rules of nodes within a two-step neighborhood (i.e., its immediate neighbors and the immediate neighbors of each such neighbor); in (b), the incoming messages depend directly upon the rules of all ancestors (a node's parents, the parents of each such parent, and so on), all descendants (i.e., a node's children, the children of each such child, and so on) as well as the ancestors of each such descendant. The dashed and dotted curves show these subsets for nodes 9 and 11, respectively, each of which similarly intersects with such subsets (not shown) of other nodes—the team coupling in each topology is the extent that the respective n subsets intersect.

the analogous efficient message-passing interpretation to retain its correctness and convergence guarantees [7]. Also note that each type of network implies different explicit online constraints: in the directed case, each node's online data is related only to the measurements local to itself and its ancestors (i.e., its parents, the parents of each such parent, and so on); in the undirected case, each node's online data is related only to the measurements local to itself and its immediate neighbors. The different online processing constraints manifest themselves in different team couplings, in the sense discussed in Fig. 4, sought by the respective offline message-passing algorithms. These architectural considerations suggest directed networks are preferable when only a few nodes are to provide state estimates (and assuming the network is such that these few nodes are at the end of its forward partial order), and undirected networks are preferable when all nodes are to provide state estimates. In general, such comparisons will also depend upon the particular pair of topologies as well as the prior, measurement, channel or cost models.

4. CONCLUSION

We have adapted the Bayesian formulation and team-theoretic analysis in decentralized detection for a simplest architecture with bidirectional inter-sensor communication defined on an *undirected* network topology. Our results show that, relative to what is known for the simplest directed architecture, the model requires more restrictive assumptions to avoid worst-case offline complexity, yet less restrictive assumptions to attain maximal offline efficiency. Future work includes an empirical performance analysis of the strategies found by the undirected message-passing algorithm and, in anticipation of performance specifications that render the two-stage processing constraints considered here too severe, an extension to more elaborate decision architectures (e.g., multiple communication stages, a hybrid or hierarchical network topology).

5. REFERENCES

[1] R.R. Tenney and N.R. Sandell, Jr., "Detection with distributed sensors," *IEEE Trans. on AES*, 17(4), pp. 501–510, Sep 1981.

[2] J.N. Tsitsiklis, "Decentralized detection," in *Advances in Statistical Signal Processing*, 2, pp. 297–344, 1993.

[3] R. Viswanathan and P.K. Varshney, "Distributed detection with multiple sensors: Part I–Fundamentals," *Proc. of the IEEE*, 85(1), pp. 54–63, Jan 1997.

[4] J.N. Tsitsiklis and M. Athans, "On the complexity of decentralized decision making and detection problems," *IEEE Trans. on AC*, 30(5), pp. 440–446, May 1985.

[5] Z.B. Tang, K.R. Pattipati, and D.L. Kleinman, "Optimization of detection networks: Part II–Tree structures," *IEEE Trans. on SMC*, 23(1), pp. 211–221, Jan/Feb 1993.

[6] A. Pete, K.R. Pattipati, and D.L. Kleinman, "Optimization of decision networks in structured task environments," *IEEE Trans. on SMC*, A-26(6), pp. 739–748, Nov 1996.

[7] O.P. Kreidl and A.S. Willsky, "An efficient message passing algorithm for optimizing decentralized detection networks," Tech Report LIDS-P-2726, MIT Laboratory for Information and Decision Systems, Cambridge, MA, Dec 2006.

[8] C.-Y. Chong and S.P. Kumar, "Sensor networks: Evolution, opportunities, and challenges," *Proc. of the IEEE*, 91(8), pp. 1247–1256, Aug 2003.

[9] A.M. Sayeed, D. Estrin, G.G. Pottie, and K. Ramchandran, "Guest editorial: Self-organizing distributed collaborative sensor networks," *IEEE J. on SAC*, 23(4), pp. 689–692, Apr 2005.

[10] J.-F. Chamberland and V.V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. on SP*, 51(2), pp. 407–416, Feb 2003.

[11] S.A. Aldosari and J.M.F. Moura, "Detection in decentralized sensor networks," in *Proc. of the IEEE ICASSP*, 2, pp. 277–280, May 2004.

[12] B. Chen and P.K. Willett, "On the optimality of the likelihood-ratio test for local sensor decision rules in the presence of non-ideal channels," *IEEE Trans. on IT*, 51(2), pp. 693–699, Feb 2005.

[13] J.-J. Xiao and Z.-Q. Luo, "Universal decentralized detection in a bandwidth constrained sensor network," *IEEE Trans. on SP*, 53(8), pp. 2617–2624, Aug 2005.

[14] V. Saligrama, M. Alanyali, and O. Savas, "Distributed detection with packet losses and finite capacity links," *IEEE Trans. on SP*, 54(11), pp. 4118–4132, Nov 2006.

[15] J.D. Papastavrou and M. Athans, "A distributed hypothesis-testing team decision problem with communications cost," in *Proc. 25th IEEE CDC*, Dec 1986, pp. 219–225.

[16] D. Teneketzis and Y.C. Ho, "The decentralized wald problem," *Information and Computation*, 73, pp. 23–44, Apr 1987.

[17] H.R. Hashemi and I.B. Rhodes, "Decentralized sequential detection," *IEEE Trans. on IT*, 35(3), pp. 509–520, May 1989.

[18] S.Alhakeem and P.K. Varshney, "A unified approach to the design of decentralized detection systems," *IEEE Trans. on AES*, 31(1), pp. 9–20, Jan 1995.

[19] S.Alhakeem and P.K. Varshney, "Decentralized Bayesian detection with feedback," *IEEE Trans. on SMC*, A-26(4), pp. 503–513, Jul 1996.