Decentralized Detection with Long-Distance Communication

(Invited Paper)

O. Patrick Kreidl and Alan S. Willsky
MIT Department of Electrical Engineering and Computer Science
77 Massachusetts Avenue, Cambridge, Massachusetts 02139 U.S.A.
Email: {opk,willsky}@mit.edu

***Abstract*— We consider the well-studied decentralized Bayesian detection problem with the twist that a small subset of nodes (all arranged in a given directed network topology) may also communicate to "long-distance" neighbors via an overlay undirected network topology. Providing a certain condition on the interface between the two networks is upheld, the natural combination of the efficient team solutions already known for each type of network alone remains team-optimal in the combined network.**

I. Introduction

Existing research literature on the decentralized detection problem, formally introduced in [1] as a special class of (static) team decision problems [2], typically assumes all measurements relate to a common hidden state. As such, the usual formulation consists of a set of remote sensors, each receiving its own measurement and transmitting information to a lone "fusion center" that, in turn, makes the final state-related decision [3], [4]. In many distributed sensing applications of contemporary interest [5], [6], each local measurement may relate to only part of the global state; as such, there may be multiple fusion centers, each making the final decision about different (but generally correlated) state variables. Some generalizations along these lines have been examined before [7], but the restriction to unidirectional communication links limits the extent to which final decisions can be coordinated and, moreover, the supporting team-optimal design algorithms appear to be difficult to implement in large networks.

This paper extends the decentralized Bayesian detection formalism to problems in which a relatively small subset of sensors, hereafter called the "leader" nodes, are endowed with the opportunity to communicate to one another via bidirectional links. Fig. 1 illustrates the hybrid networks encompassed by our problem formulation, combining the usual directed (acyclic) network $\mathcal{G}^D$ with an overlay undirected network $\mathcal{G}^U$. Online decision-making in the directed non-leader network proceeds from parentless to childless nodes in the forward partial order implied by $\mathcal{G}^D$. The novelty of our model is that each leader node in the undirected network, after receiving information from its (non-leader) parents but before transmitting information to its (non-leader) children in $\mathcal{G}^D$, may also exchange information with its neighboring leaders in $\mathcal{G}^U$. The fact that the leader network can connect nodes that are not spatial neighbors in the non-leader network captures the opportunity for "long-distance" communication.

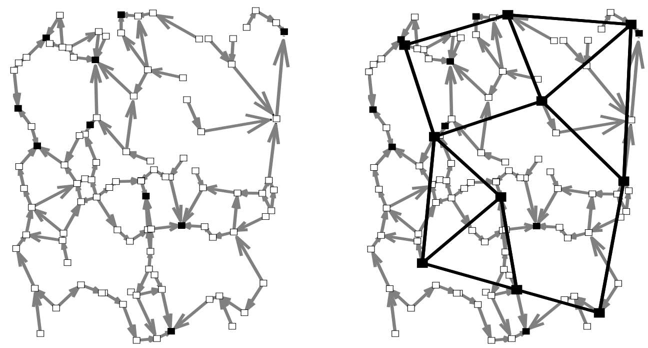

(a) A "Non-Leader" Network (b) An Overlay "Leader" Network

Fig. 1. Illustration of the hybrid networks $\mathcal{H} = \mathcal{G}^D \cup \mathcal{G}^U$ considered in this paper. The (a) "non-leader" network is defined by any spanning directed acyclic graph $\mathcal{G}^D$, where the fusion centers (chosen to be the childless nodes here) are designated by the filled markers. The (b) "leader" network is defined by any undirected graph $\mathcal{G}^U$, connecting via "long-distance" links an arbitrary yet relatively small subset of the nodes in (a). Note that the leader nodes in (b) need not necessarily coincide with the designated fusion centers in (a).

Our main result is that, providing a certain condition on the interface between the two sub-networks is upheld, the natural combination of the efficient team solutions already known for each type of network alone remains team-optimal in the hybrid network $\mathcal{H} = \mathcal{G}^D \cup \mathcal{G}^U$. Specifically, we identify a class of hybrid network models in which (i) every node's optimal rule for online measurement processing simplifies to a local likelihood-ratio test and (ii) the offline iterative algorithm for computing the globally-coupled thresholds [3], [8] admits a message-passing interpretation analogous to those presented in [9] and [10] for directed and undirected networks, respectively. The latter conclusion is arguably surprising, considering the extent to which a hybrid network topology appears to be non-tree structured (and offline efficiency for directed networks alone requires a tree-structured topology [9], [11]). As will be discussed, this curious result manifests itself here via essentially the same probabilistic structure exploited in the analysis of undirected networks alone (where offline efficiency also does *not* require a tree topology [10]).

The remainder of this paper is organized as follows. Section II generalizes the canonical Bayesian detection formulation to the case of a hybrid network topology. Section III summarizes the efficient team solution for these hybrid networks, leveraging those already known for either type of network alone. (For brevity, we must omit proofs and full details of the message-passing equations, which can both be found in [12]). To conclude, Section IV applies the algorithm to a specific hybrid network of binary detectors, quantifying the performance gained by permitting long-distance communication.

II. Bayesian Formulation

This section generalizes the Bayesian decentralized formulation with costly, unreliable communication presented in [9] and [10] (for directed and undirected networks, respectively) to encompass the hybrid networks illustrated in Fig. 1. Fig. 2 depicts the associated centralized and decentralized processing models, both assuming that (i) the hidden state x and observable measurement y take their values in, respectively, a discrete product space $\mathcal{X} = \mathcal{X}_1 \times \cdots \times X_n$ and Euclidean product space $\mathcal{Y} = \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_n$; (ii) a given distribution $p(x, y)$ jointly describes the hidden state process X and noisy measurement process Y; and (iii) the network is to generate a global estimate $\hat{x} \in \mathcal{X}$ based on the realized measurement vector $Y = y$. The decentralized model in Fig. 2(b) further assumes each component of $\hat{x}$ is determined by an individual sensor in the n-node hybrid network $\mathcal{H} = \mathcal{G}^D \cup \mathcal{G}^U$. Each edge (i, j) in either graph indicates a (perhaps unreliable) low-rate communication link between the respective pair of nodes: the distinction, however, is that in the (directed) non-leader topology $\mathcal{G}^D$ the links are only unidirectional while in the undirected leader topology $\mathcal{G}^U$ the links are bidirectional. In turn, the details of each local processing rule depend on whether or not the node is part of the leader network.

A. Formulation with No Leader Network

Let us first consider the decentralized processing model in the canonical case of no leader network. Decision(s) are generated in the forward partial order implied by $\mathcal{H} = \mathcal{G}^D$; that is, each node i, observing only the component measurement y_i and the symbol(s) z_i received on incoming links with all parents $pa(i) = \{j \mid \text{edge } (j, i) \text{ in } \mathcal{G}^D\}$ (if any), decides upon both its component estimate $\hat{x}_i$ and the symbol(s) u_i transmitted on outgoing links with all children $ch(i) = \{j \mid \text{edge } (i, j) \text{ in } \mathcal{G}^D\}$ (if any). The symbols received and transmitted at every node i take their values in discrete sets $\mathcal{Z}_i$ and $\mathcal{U}_i$, respectively. Let the multipoint-to-point channel into each node i be modeled by conditional distribution $p(z_i|x, y, u_{pa(i)})$, describing the side information Z_i at node i based on its parents' transmitted symbols $u_{pa(i)} = \{u_j \in \mathcal{U}_j \mid j \in pa(i)\}$ (and, in general, the processes (X, Y) of the environment external to the network).

Altogether, the collections of transmitted symbols u and received symbols z take their values in discrete product spaces $\mathcal{Z} = \mathcal{U}_1 \times \cdots \times \mathcal{U}_n$ and $\mathcal{Z} = \mathcal{Z}_1 \times \cdots \times \mathcal{Z}_n$, respectively. By construction, the global decision process $\hat{X}$ is generated component-wise in the forward partial order of $\mathcal{G}^D$, each node i individually generating both U_i and $\hat{X}_i$ upon observing both Y_i and Z_i. It follows that any particular *strategy* $\gamma : \mathcal{Y} \times \mathcal{Z} \to \mathcal{U} \times \mathcal{X}$ induces a global decision process $(U, \hat{X}) = \gamma(Y, Z)$.

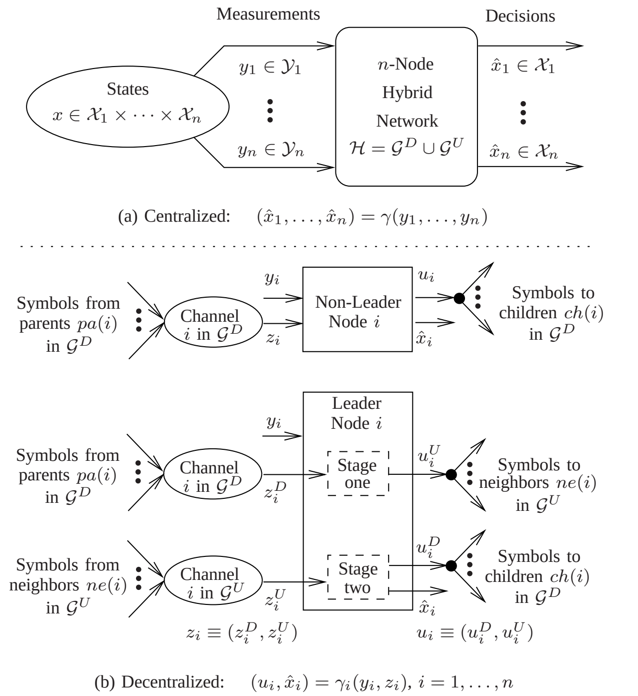

Fig. 2. Our global detection model assuming (a) a centralized strategy and (b) a decentralized strategy for processing measurements. In the latter, the details of each local processing rule depend on whether or not the node is in the leader network.

Denote by $\bar{\Gamma}$ the set of all such strategies and by $\Gamma \subset \bar{\Gamma}$ the *admissible* subset of these strategies in the decentralized network; specifically, denoting by Γ_i the subset of all rules $\gamma_i : \mathcal{Y}_i \times \mathcal{Z}_i \to \mathcal{U}_i \times \mathcal{X}_i$ local to node i, we let $\Gamma = \Gamma_1 \times \cdots \times \Gamma_n$.

The Bayesian criterion is essentially the same as for the centralized processing model, but also accounting for the communication-related decision process U. That is, every possible realization of the joint process $(U, \hat{X}, X)$ is assigned a real-valued cost of the form

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

where non-negative constant λ specifies the unit conversion between detection costs $c(\hat{x}, x)$ and communication costs $c(u, x)$. Any fixed strategy $\gamma \in \bar{\Gamma}$ is then penalized by its expected cost,

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

In turn, the decentralized design problem is expressed by

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

where the functions $J_c : \bar{\Gamma} \to \mathbb{R}$ and $J_d : \bar{\Gamma} \to \mathbb{R}$ quantify the *communication penalty* and the *detection penalty*, respectively.

B. Formulation with a Leader Network

Starting with a given n-node directed acyclic graph $\mathcal{G}^D$, a hybrid network $\mathcal{H}$ is formed by also introducing a particular undirected graph $\mathcal{G}^U$. As discussed in Section I, the former is the non-leader network and the latter is the leader network.

Moreover, we typically assume that the vertex set $\mathcal{V}^U$ of $\mathcal{G}^U$ is of cardinality much smaller than n. The processing model local to each non-leader node $i \notin \mathcal{V}^U$ remains unchanged from that discussed above. However, each leader node $\ell \in \mathcal{V}^U$ can, after receiving symbols from its parents $pa(\ell)$ but before transmitting symbols to its children $ch(\ell)$ in the non-leader network $\mathcal{G}^D$, exchange symbols with its neighboring leaders $ne(\ell) = \{j \mid \text{edge } \{\ell, j\} \text{ in } \mathcal{G}^U\}$.

Specifically, each leader node $\ell \in \mathcal{V}^U$ operates in two distinct stages: the first decides the symbols u_ℓ^U transmitted to its neighboring leaders based on component measurement y_ℓ and symbols z_ℓ^D from its non-leader parents; the second, upon also receiving symbols z_ℓ^U on the links with its neighboring leaders, decides the state estimate $\hat{x}_\ell$ and the symbols u_ℓ^D for its non-leader children. Thus, the symbols received and transmitted at each leader node ℓ take their values in composite sets $\mathcal{Z}_\ell = \mathcal{Z}_\ell^D \times \mathcal{Z}_\ell^U$ and $\mathcal{U}_\ell = \mathcal{U}_\ell^D \times \mathcal{U}_\ell^U$, respectively. Similarly, each leader node is taken to have two independent channel models, $p(z_\ell^D|x, y, u_{pa(\ell)})$ describing information Z_ℓ^D received from its parents in $\mathcal{G}^D$ and $p(z_\ell^U|x, y, u_{ne(\ell)}^U)$ describing information Z_ℓ^U received from its neighbors in $\mathcal{G}^U$. Altogether, the formulation in (2) still applies providing that every leader's rule set is augmented to $\Gamma_\ell = \mathcal{M}_\ell \times \Delta_\ell$, where $\mathcal{M}_\ell$ and Δ_ℓ consist of all functions $\mu_\ell : \mathcal{Y}_\ell \times \mathcal{Z}_\ell^D \to \mathcal{U}_\ell^U$ and $\delta_\ell : \mathcal{Y}_\ell \times \mathcal{Z}_\ell^D \times \mathcal{U}_\ell^U \times \mathcal{Z}_\ell^U \to \mathcal{U}_\ell^D \times \mathcal{X}_\ell$, respectively.

Observe that a leader node will not communicate within the undirected network (and, in turn, with any of its children in the directed network) until it has received symbols from all of its parents in the directed network. As illustrated in Fig. 3, this opens up the possibility for *gridlock*, in which online processing can stall because information required before a leader node may proceed cannot be realized until after this same leader node transmits information. The following assumption on hybrid network $\mathcal{H}$ ensures the absence of any such gridlock. For any given directed network, the *ancestors* of any node i consist of its parents $pa(i)$, the parents $pa(j)$ of each such parent $j \in pa(i)$ and so on; similarly, the *descendants* consist of its children $ch(i)$, the children $ch(j)$ of each such child $j \in ch(i)$ and so on.

Assumption 1 (Absence of Gridlock): In hybrid network $\mathcal{H} = \mathcal{G}^D \cup \mathcal{G}^U$, for every adjacent pair of leader nodes in the undirected network $\mathcal{G}^U$, there exists no non-leader node that, in the directed network $\mathcal{G}^D$, is both an ancestor of one and a descendant of the other.

With Assumption 1 in place, the strategy-dependent distribution that determines $J(\gamma)$ in (1) becomes well-defined. In particular, for each non-leader node $i \notin \mathcal{V}^U$, fixing the rule $\gamma_i \in \Gamma_i$ is equivalent to specifying the distribution

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

Here, we have introduced the superscript-D notation on both u_i and $u_{pa(i)}$ for compatibility with the leader node notation, recognizing that (i) $u_i^D \equiv u_i \in \mathcal{U}_i$ for every non-leader node i and (ii) $u_j^D \equiv u_j \in \mathcal{U}_j$ for every parent $j \in pa(i)$ unless node j is also a leader node in which case $u_j^D \in \mathcal{U}_j^D$. For each leader node $\ell \in \mathcal{V}^U$, we use u_ℓ and z_ℓ to denote (u_ℓ^U, u_ℓ^D) and (z_ℓ^U, z_ℓ^D), respectively, so that similarly fixing the rule $\gamma_\ell = (\mu_\ell, \delta_\ell)$ is equivalent to specifying the distribution

$$p(u_\ell, \hat{x}_\ell|y_\ell, z_\ell; \gamma_\ell) = p(u_\ell^U|y_\ell, z_\ell^D; \mu_\ell) p(u_\ell^D, \hat{x}_\ell|y_\ell, z_\ell, u_\ell^U; \delta_\ell)$$

and, in turn,

$$p(u_\ell, \hat{x}_\ell|x, y, u_{pa(\ell)}^D, u_{ne(\ell)}^U; \gamma_\ell) = \sum_{z_\ell^D \in \mathcal{Z}_\ell^D} p(z_\ell^D|x, y, u_{pa(\ell)}^D) p(u_\ell, \hat{x}_\ell|x, y, z_\ell^D, u_{ne(\ell)}^U; \gamma_\ell)$$

with

$$p(u_\ell, \hat{x}_\ell|x, y, z_\ell^D, u_{ne(\ell)}^U; \gamma_\ell) = \sum_{z_\ell^U \in \mathcal{Z}_\ell^U} p(z_\ell^U|x, y, u_{ne(\ell)}^U) p(u_\ell, \hat{x}_\ell|y_\ell, z_\ell; \gamma_\ell).$$

It follows that fixing the entire strategy $\gamma \in \Gamma_1 \times \cdots \times \Gamma_n$ determines the conditional distribution

$$p(u, \hat{x}|x, y; \gamma) = \prod_{i \notin \mathcal{V}^U} p(u_i^D, \hat{x}_i|x, y, u_{pa(i)}^D; \gamma_i) \times \prod_{\ell \in \mathcal{V}^U} p(u_\ell, \hat{x}_\ell|x, y, u_{pa(\ell)}^D, u_{ne(\ell)}^U; \gamma_\ell)$$

and, in turn, the distribution underlying (1) via

$$p(u, \hat{x}, x; \gamma) = \int_{y \in \mathcal{Y}} p(x, y) p(u, \hat{x}|x, y; \gamma)\, dy.$$

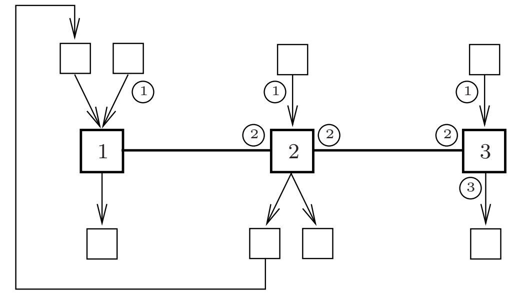

(a) Gridlock Between Leader Nodes 1 & 2

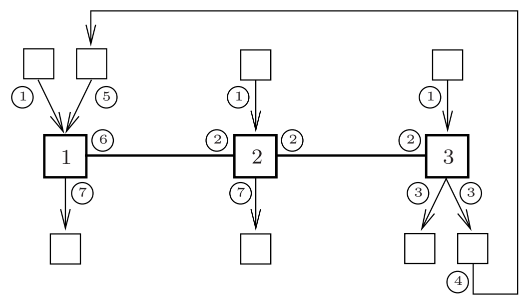

(b) Absence of Gridlock

Fig. 3. Two simple hybrid networks, (a) one with gridlock and (b) one without gridlock. In each case, the leader (non-leader) nodes are the large (small) squares, while the circled numbers beside the links indicate the sequential partial-ordering of the nodes' communication decisions. Note that (a) violates Assumption 1 while (b) does not.

III. Main Result for Hybrid Networks

As highlighted in Section I, when considering a directed network or an undirected network alone, previous work identifies the minimal model assumptions by which team-optimality conditions associated with (2) give rise to an offline iterative algorithm that admits an efficient message-passing interpretation [9], [10]. Firstly, the optimal decentralized strategy γ^* is guaranteed to have a finite parameterization (i.e., reduce to a collection of likelihood-ratio tests) only under the conditional independence assumption; in the undirected case, the separable cost assumption is also required. Secondly, total offline computation/communication overhead scales linearly with the number of nodes n only under the measurement/channel/cost locality assumption; in the directed case, a polytree topology is also required. Our analysis of hybrid networks starts with these assumptions in place.

Assumption 2: (Conditional Independence & Measurement/Channel Locality): In hybrid network $\mathcal{H}$, the global probabilistic model satisfies

$$p(y_i, z_i|x, y_{-i}, z_{-i}, u^D_{pa(i)}) = p(y_i|x_i)p(z_i|x_i, u^D_{pa(i)})$$

for every non-leader node $i \notin \mathcal{V}^U$ and

$$p(y_\ell, z_\ell|x, y_{-\ell}, z_{-\ell}, u^D_{pa(\ell)}, u^U_{ne(\ell)}) = p(y_\ell|x_\ell)p(z^D_\ell|x_\ell, u^D_{pa(i)})p(z^U_\ell|x_\ell, u^U_{ne(\ell)})$$

for every leader node $\ell \in \mathcal{V}^U$.

Assumption 3 (Separable Costs & Cost Locality): In hybrid network $\mathcal{H}$, the global cost function satisfies

$$c(u, \hat{x}, x) = \sum_{i=1}^{n} c(u_i, \hat{x}_i, x_i)$$

with

$$c(u_i, \hat{x}_i, x_i) = \begin{cases} c(u^D_i, \hat{x}_i, x_i) & , \quad i \notin \mathcal{V}^U \\ c(u^U_i, x_i) + c(u^D_i, \hat{x}_i, x_i) & , \quad i \in \mathcal{V}^U \end{cases} .$$

Assumption 4 (Polytree Non-Leader Network): In hybrid network $\mathcal{H}$, the directed sub-network $\mathcal{G}^D$ is a polytree.

Interestingly, it turns out that Assumptions 1 to 4 are *not* sufficient for team-optimality conditions to admit an efficient message-passing interpretation. Additional restrictions on hybrid network $\mathcal{H}$, or more specifically on the interface between the two types of sub-networks, are required. The following assumption encapsulates the class of hybrid networks for which both efficiency and correctness of the offline message-passing equations is retained. The *lineage* of each leader node $\ell \in \mathcal{V}^U$ refers to the union of its ancestors and descendants in $\mathcal{G}^D$, which we note may consist of both leader and non-leader nodes.

Assumption 5 (Hybrid Interface Restrictions): In a hybrid network $\mathcal{H} = \mathcal{G}^D \cup \mathcal{G}^U$, for every pair of leader nodes within a distance two of each other (in $\mathcal{G}^U$), the respective lineages (in $\mathcal{G}^D$) have no node in common.

Notice that the conditions in Assumption 5 subsume those of Assumption 1, as the latter is satisfied given every pair of adjacent leader nodes (i.e., leader nodes within a distance of one on $\mathcal{G}^U$) have disjoint lineages. To convey some intuition behind this assumption, first recall from Section II that the flow of online measurement processing proceeds from parentless nodes to childless nodes, where every leader node along the way, upon receiving symbols from all of its parents in $\mathcal{G}^D$, exchanges symbols with all of its neighbors in $\mathcal{G}^U$ before it transmits symbols to its children in $\mathcal{G}^D$. The key analysis step is to "unravel" the hybrid network to an equivalent directed network whose partial-order corresponds to the proscribed flow of online processing. Fig. 4 illustrates a simple hybrid topology and its "unraveled" counterpart, replicating all leader nodes to reflect their two-stage decision processes.

Intuitively-speaking, Assumption 5 (together with Assumption 4) basically ensures that this unraveled network retains an overall polytree topology, established to be essential for efficiency in pure directed networks [9]. That the replicated leader nodes (i.e., two nodes always having equal measurements), which appear to violate the conditional independence

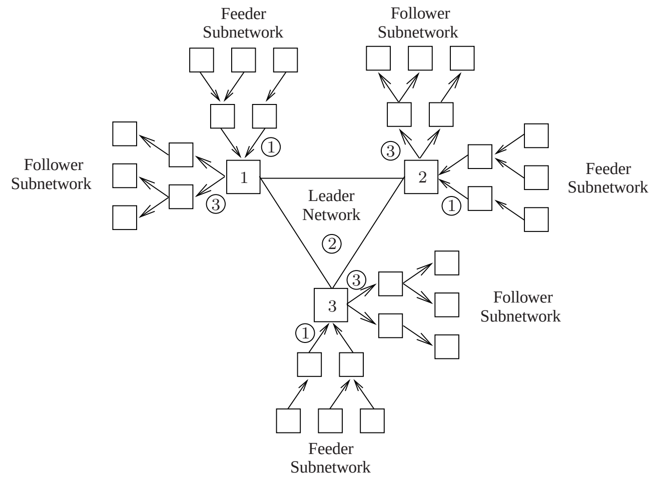

(a) A Simple Hybrid Network

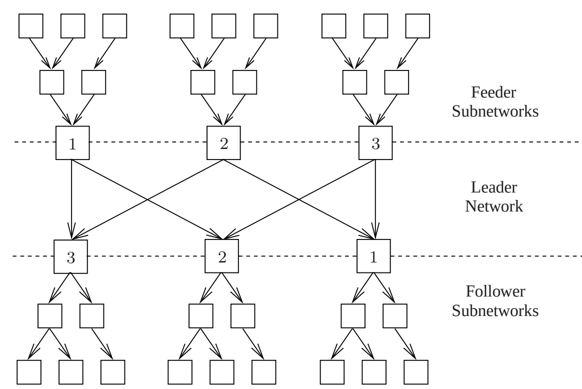

(b) The "Unraveled" Directed Counterpart to (a)

Fig. 4. A (a) simple hybrid network and (b) its equivalent directed counterpart, the latter's forward partial order in correspondence with the flow of online processing in (a). This example lies in the class of hybrid networks for which team-optimality conditions reduce to efficient message-passing equations, where the associated offline iterative algorithm corresponds to repeated forward-backward sweeps in (b).

assumption, cause no additional complexity here is established in exactly the same manner as for pure undirected networks [10]. The arguably peculiar fact that Assumption 4 places restrictions only between leader nodes within distance two of one another (in $\mathcal{G}^U$) similarly follows from known results for pure undirected networks. Altogether, the offline message-passing algorithm is also seen to respect this "unraveled" topology, iterating over a set of fixed-point equations (obtained from the natural combination of those already known for each type of sub-network alone) in correspondence with repeated forward-backward sweeps in the equivalent directed network.

Proposition 1 (Hybrid Offline Efficiency): Consider a hybrid network $\mathcal{H} = \mathcal{G}^D \cup \mathcal{G}^U$ and let Assumptions 2 to 5 hold. Team-optimality conditions for (2) are satisfied by the collection of non-leader rules

$$\gamma_i^*(Y_i, Z_i) = \arg\min_{(u_i,\hat{x}_i)} \sum_{x_i} \phi_i^*(u_i, \hat{x}_i, x_i; Z_i) p(Y_i|x_i)$$

for $i \notin \mathcal{V}^U$ and the collection of leader rules

$$\mu_\ell^*(Y_\ell, Z_\ell^D) = \arg\min_{u_\ell^U} \sum_{x_\ell} \alpha_\ell^*(u_\ell^U, x_\ell; Z_\ell^D) p(Y_\ell|x_\ell)$$

$$\delta_\ell^*(Y_\ell, Z_\ell) = \arg\min_{(u_\ell^D,\hat{x}_\ell)} \sum_{x_\ell} \beta_\ell^*(u_\ell^D, \hat{x}_\ell, x_\ell; Z_\ell) p(Y_\ell|x_\ell)$$

for $\ell \in \mathcal{V}^U$, where the real-valued parameters

$$\phi^* = \{\phi_i^*; i \notin \mathcal{V}^U\} \bigcup \{(\alpha_\ell^*, \beta_\ell^*); \ell \in \mathcal{V}^U\}$$

denote any solution to nonlinear fixed-point equations having a block structure in correspondence with edges in the network $\mathcal{H}$. (Proofs and equations can be found in Chapter 4 of [12].)

IV. Example

We close with an application of the efficient team solution to a particular 25-node network, comparing performance with and without the presence of a leader network of binary detectors, holding all other aspects of the problem constant. Costs are chosen to specialize detection penalty J_d to the sum node-error-rate across the six gateway nodes (filled markers), each making its own final state-related decision, and communication penalty J_c to the sum link-use-rate across the entire network, each node having the (cost-free) option to remain silent on its outgoing unit-rate links. Each sensor is modeled as the same scalar linear-Gaussian binary detector, while each multipoint-to-point link is modeled as the same interference channel with loss probability $q \in [0, 1]$. (See [12] for more details of the underlying probabilistic model and transmission scheme.)

Fig. 5 displays both the specific network topologies (with and without the leader network) and the optimized tradeoff curves obtained via the offline message-passing algorithm. Notice that the random construction of the hybrid network has left one leader node without any supporting non-leader subnetwork. Nonetheless, with the leader network in place, up to roughly 80% of the centralized detection performance lost by the purely myopic strategy (i.e., the strategy when J_c is zero) can be recovered; without the leader network, only up to roughly 40% of this performance gap is recovered.

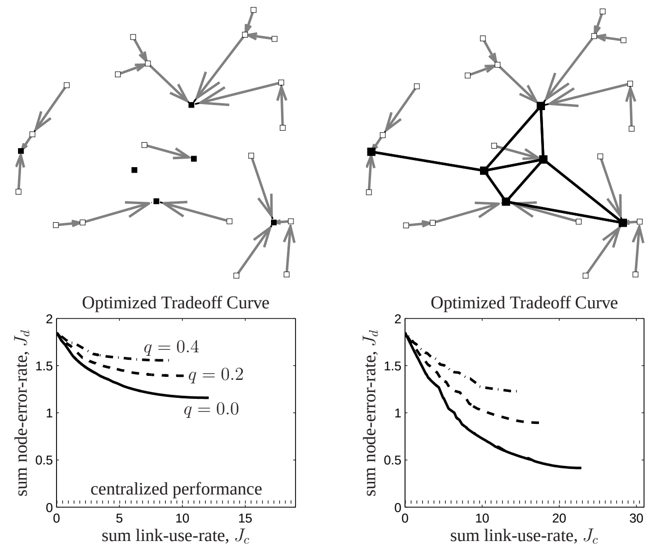

(a) Without Leader Network (b) With Leader Network

Fig. 5. Comparison of the optimized tradeoff curves achieved by the offline message-passing algorithms for a 25-node network of binary detectors with and without the undirected leader network: the filled (unfilled) markers denote the gateway (communication-only) nodes responsible for making (supporting) the final state-related decisions. Only in (b) are these gateway nodes provided the opportunity for long-distance communication. Three curves are shown in each case, corresponding to channel loss probabilities of $q = 0$ (solid line), 0.2 (dashed line) and 0.4 (dash-dotted line). Each such curve is obtained by increasing λ in increments of 10^{-4}, starting with $\lambda = 0$, and for each such λ declaring convergence in iteration k when $J(\gamma^{k-1}) - J(\gamma^k) < 10^{-3}$. Also shown is a simulation-based Monte-Carlo approximation of the centralized detection penalty (dotted line), computed using 1000 samples from $p(x, y)$.

V. Acknowledgment

This work supported by the U.S. Army Research Office under MURI grant W911NF-06-1-0076.

References

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

[2] J. Marschak, R. Radner, et al., *The Economic Theory of Teams*, 1972.

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

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

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

[6] A.M. Sayeed, D. Estrin, et al., "Self-organizing distributed collaborative sensor networks," *IEEE J. on SAC*, 23(4), pp. 689-692, 2005.

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

[8] Z. Tang, K.R. Pattipati, and D.L. Kleinman, "An algorithm for determining the decision thresholds in a distributed detection problem," *IEEE Trans. on SMC*, 21(1), pp. 231-237, 1991.

[9] O.P. Kreidl and A.S. Willsky, "An efficient message passing algorithm for optimizing decentralized detection networks," in *Proc. of IEEE CDC*, pp. 6776-6783, 2006.

[10] O.P. Kreidl and A.S. Willsky, "Decentralized detection in undirected network topologies," in *IEEE SSP Workshop*, 2007.

[11] Z. 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, 1993.

[12] O.P. Kreidl, *Graphical Models and Message-Passing Algorithms for Network-Constrained Decision Problems*, MIT Ph.D. thesis, 2008.