

Collaborative Multiagent Gaussian Inference in a Dynamic Environment Using Belief Propagation

(Extended Abstract)

Stefano Ermon
Department of Computer
Science
Cornell University
Ithaca, New York
ermonste@cs.cornell.edu

Carla Gomes
Department of Computer
Science
Cornell University
Ithaca, New York
gomes@cs.cornell.edu

Bart Selman
Department of Computer
Science
Cornell University
Ithaca, New York
selman@cs.cornell.edu

ABSTRACT

The problem of multiagent Gaussian inference in a dynamic environment, also known as distributed Kalman filtering, is formulated into the framework of message passing algorithms. Upon generalizing the derivation of the standard Kalman filter to the distributed case, we propose novel solutions that outperform current state of the art techniques.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Algorithms

Keywords

Distributed problem solving, Reasoning

1. INTRODUCTION

Many distributed inference problems can be modeled as a network $G = (V, E)$ of sensing devices that can perform local computations and communicate with other nodes, collaborating to produce global information from individual local data. In particular, the focus of this paper is on Bayesian estimation, where a probability model is assumed to be known and one is interested in computing the posterior distribution of a collection of hidden variables (“the state” x), given the evidence collected in the network.

In many real world problems it is critical to introduce the dynamics of the system into the model, for example in the case of tracking a moving object or monitoring an environment over time. Here we consider the case of a state of the world $x \in \mathbb{R}^N$ evolving in time according to a discrete time linear dynamical system:

$$x_{k+1} = A_k x_k + w_k, \quad (1)$$

Cite as: Collaborative Multiagent Gaussian Inference in a Dynamic Environment Using Belief Propagation (Extended Abstract), S. Ermon, C. Gomes and B. Selman, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX.

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

where the noise w_k is Gaussian. Furthermore we assume that each sensor $i \in V$ in the network measures an output $y_k(i)$ at each time step k that is a linear combination of the state variables corrupted by additive Gaussian noise.

Ideally one would like to compute at each node $i \in V$ and for each time step k the minimum mean square error estimate (that is also the maximum likelihood estimate) of the state x_k given by

$$\hat{x}(k|k) \triangleq E[x_k | \cup_{i \in V} \{y_\ell(i) | \ell \leq k\}]. \quad (2)$$

Notice that this is a statistical *causal* inference problem, because only information coming from the past can be used to infer properties about the present (or future) states.

When there is only one agent, or one central node is assumed to know all the information available in the network, this formulation is analogous to a Kalman filtering problem. However in many settings a centralized solution, in which a single computational node receives and elaborates all the information available, is either not feasible due to their communication and energy restrictions or not desirable because it introduces a single point of failure. Therefore there is a need for distributed solutions where inference is performed locally at each node on the basis of information that is retrieved both locally and by communication with neighboring nodes.

In this paper, we show how to generalize the classical result ([1]) that derives the Kalman filter in terms of Belief Propagation (BP) to the distributed case by choosing a suitable graphical model. The model defines an inference problem that is equivalent to the system theoretic definition given by equations (1) and (2), but embeds the structure of the communication network G . This will naturally lead to the definition of a message passing algorithm that at the same time also defines a communication protocol, because it is sufficient to interpret messages as real communications that take place between nodes.

A key feature is that the graphical model proposed relies heavily on the notion of *spanning tree* of the network, that imposes an hierarchy among the nodes and therefore enforces an ordered flow of information in the corresponding corresponding communication protocol. In a practical application, the algorithm proposed in [3] for a static inference problem can be used to organize the nodes of the network into a spanning tree with high-quality communication links and without need of central coordination.

The resulting solution enables all the nodes in the net-

work to compute an estimate that is as good as a centralized one while at the same time minimizing the total use of communication resources. In fact the messages computed according to the BP updates allow the nodes to locally elaborate and fuse the information they receive before transmitting it again, being able to summarize it and thus reducing the number of messages needed. Therefore the proposed architecture distributes the computational burden and also reduces the communication resources used.

Nevertheless, even if the algorithm is fully distributed and *causal*, it might not be possible to use it in *real-time* if there are latencies in the communication links. When the time needed for the information to travel around the network is greater than the time step of the dynamical system, the previous approach is not viable and it is not generally possible to get an estimate as good as the centralized one anymore.

If that is the case, then the message passing approach is particularly useful because it becomes fundamental to maximize the flow of information in order to be able to take decisions readily, even on the base of the partial information received. However dealing with delays is not an easy task, mainly because an intrinsic feature of dynamic estimation problems is that the temporal order in which information is used becomes important. In particular, the graphical model interpretation defines an exact solution that is expensive both in terms of space and time complexity (KF-delayed), yet also leads to a much more efficient approximation technique (BP-approx) based on message passing.


The performance of the approximate solution is compared in simulation against a popular *consensus* based solution ([2]) developed specifically for this setting and that implicitly assumes communication latencies. The benchmark is a tracking application, with an experimental setup very similar to the one described in [2].

We consider a target moving on noisy circular trajectories in a plane. A network of 50 sensing devices is randomly generated according to the *random geometric model*, and the nodes make noisy measurements of the position of the target either along the x -axis, or along the y -axis. Clearly, no individual sensor can estimate the position of the target by itself, but it becomes possible using information coming from neighboring nodes.


We compare several approaches by measuring the average error of the estimates of the nodes and their *disagreement*, defined as the variance with respect to the average estimate. Having a good level of agreement is helpful in many applications where nodes are also taking decisions based on their estimated state, mainly because it improves the level of cooperation. The approximate solution (BP-approx) and the *consensus* based solution (DKF) are compared against the centralized Kalman filter (CKF) that computes (2) using all the information available at time k (not implementable with communication latencies). We also consider the exact estimation algorithm (KF-delayed), that represents a lower bound for the average error in the presence of latencies. The results are shown in figure 1(a) and 1(b) respectively.

The improvement of BP-approx over DKF on the average error is of about 8%, while on average the improvement on the disagreement is of almost 50%. Empirically we have also seen that the performance gap tends to increase with higher noise levels, measured by a larger trace of Q . Moreover we can see that the approximation given by BP-approx is almost as good as the theoretical optimum in presence of

delays given by KF-delayed.



(a) Comparison of the performance of the algorithms in terms of average error.



(b) Comparison of the performance of the algorithms in terms of disagreement.

Figure 1: Simulative comparison between the algorithms.

2. ACKNOWLEDGMENTS

This research is funded by NSF grant 0832782.

3. REFERENCES

- [1] F. Kschischang, B. Frey, and H. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on information theory*, 47(2):498–519, 2001.
- [2] R. Olfati-Saber. Distributed Kalman filtering for sensor networks. In *Proc. of the 46th IEEE Conference on Decision and Control*, 2007.
- [3] M. Paskin and C. Guestrin. Robust probabilistic inference in distributed systems. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 436–445. AUAI Press Arlington, Virginia, United States, 2004.