

Distributed data-aggregation consensus for sensor networks

Relaxation of consensus concept and convergence property

Michael Rabbat
McGill University

Prepared By:
McGill University
3480 University Street
Montréal, Québec H3A 0E9

PWGSC Contract Number: W7707-135656

Contract Scientific Authority: Stephane Blouin, 902-463-3100 ext. 216

The scientific or technical validity of this Contract Report is entirely the responsibility of the Contractor and the contents do not necessarily have the approval or endorsement of the Department of National Defence of Canada.

Contract Report
DRDC-RDDC-2014-C152
August 2014

- © Her Majesty the Queen in Right of Canada, as represented by the Minister of National Defence, 2014
- © Sa Majesté la Reine (en droit du Canada), telle que représentée par le ministre de la Défense nationale, 2014

Abstract

This project considers the problem of localization and tracking an underwater target from bearings-only measurements obtained by static underwater sensors. Nodes communicate their local observations over volatile, unreliable underwater channels with the aim being to improve the overall accuracy of each sensor's estimate, while also aligning the estimates so that all nodes agree on the target location. This report summarizes our initial investigation of having each node locally run a particle filter for state estimation.

This page intentionally left blank.

1 Introduction

Note: *At the time the present contract was awarded, the outcome of the “Adaptive Multi-sensor Biomimetics for Unsupervised Submarine Hunt” (AMBUSH) Technology Investment Fund (TIF) competition was not known. Since then, the TIF AMBUSH project has been awarded. The present contract is thus viewed as an initial effort which will lead into and benefit the extended TIF project. In light of this, the Project Authority has authorized the contractor to write a shorter-than-normal contract report so as to focus on research and staffing efforts which will ultimately benefit the TIF project.*

We consider a network of underwater sensors which cooperate to localize and track the position of target emitting an acoustic signal. The sensors measure the direction to the target, relative to true North. Although rough distance information can also be obtained from the data, it appears to be highly unreliable and so for the time being we focus on the problem of cooperative bearings-only localization and tracking.

The underwater sensors communicate over acoustic channels, and this is a very challenging and unreliable medium [1]. The quality of the communication channel changes dramatically as thermoclines and other environmental factors vary. Consequently, communication links are directed and time-varying: a node i may receive messages transmitted by j while j does not receive messages transmitted by i , and the times when i does receive messages from j may be sporadic.

The aim of the project is to develop a method for collaborative target localization and tracking in this volatile environment. Consensus algorithms [2] are very attractive for use in such environments because they do not involve routing messages over multiple hops. Instead, nodes exchange messages with their immediate neighbors and fuse data with the aim of reaching a state where all nodes agree on the same value. The existing literature on consensus algorithm provides conditions under which convergence to a consensus is guaranteed in spite of time-varying and directed network connectivity [2, 3, 4, 5].

2 Literature Review

Consensus dynamics have been studied intensively in the systems and control community during the past decade [2, 6]. The vast majority of analyses consider continuous time dynamics which imply that nodes are in constant communication with their neighbors. There are also discrete-time versions of the consensus dynamics considered, however these also assume synchronous communication and updating.

Given the dynamic nature of the communication medium and the network in this project, a different class of consensus algorithm called *randomized gossip* is more suitable [7, 8]. In asynchronous randomized gossip algorithms, pairs of neighboring nodes exchange messages and perform updates in an asynchronous and unattended manner, and they also

The class of broadcast gossip algorithms [9, 10, 11, 12] are particularly relevant to this project since underwater communications should be modeled as a broadcast medium; when one node transmits, all other nodes potentially receive the message. Of course, in practice, only a (random) subset will actually receive the message. In contrast to synchronous consensus dynamics [2] and asynchronous pairwise randomized gossip [7, 8], broadcast gossip algorithms do not require that nodes know the identities of their neighbors. Instead, in broadcast gossip, nodes periodically and asynchronously broadcast their current estimate, and those other nodes who receive the message perform an update and aggregate the received information with their own local estimate. Consequently, nodes do not need to know any characteristics of the network (size, or other structure-related parameters), and the algorithm is highly robust to changing network size, topology, and other conditions.

The majority of work on consensus and gossip has focused on the distributed averaging problem, where each node has an initial value and the objective is to reach a state where all nodes have the average of the initial values. A smaller subset of the literature has considered the estimation and tracking of time-varying signals, typically in a linear-quadratic-Gaussian framework, with the aim being to minimize mean-squared error or mean-squared deviation [13, 14, 15, 16]. Bearings measurements are highly nonlinear, and the noise in underwater acoustic measurements is likely not Gaussian. Thus, these methods are not likely to be well-suited to the problem at hand.

Recently, there has been interest in developing consensus-based approaches to distributed particle filtering [17, 18, 19, 20]. Particle filters are a sequential Monte Carlo method for recursive state estimation [21]. Instead of assuming that the target dynamics and measurements follow a particular parametric distribution (e.g., Gaussian measurements or linear dynamics), particle filters represent uncertainty in the target state using a collection of weighted point masses (the particles). Although particle filters can be more computationally intensive than the Kalman family of filtering approaches, they can also provide significantly better tracking performance, especially when the target dynamics are non-linear and/or the measurement noise is non-Gaussian. In distributed particle filtering approaches, each node runs a local particle filter. In order to update the state estimates, the measurements from different sensors need to be fused and incorporated, and different approaches have been proposed to accomplishing this using consensus algorithms [17, 18, 19, 20], with the aim being to aggregate and diffuse information over the network so that all nodes agree on the target state and track.

3 Proposed Solution

We propose to adopt a distributed particle filtering approach in this work. This choice is motivated by the observations drawn from the dataset provided by the Scientific Authority at the start of the project. In this dataset, the targets dynamics are linear for a significant portion of the time, but they become reasonably non-linear when the target turns or maneuvers. Moreover, the measurements contain a very significant amount of noise and many measurements should be treated as coming from clutter. For this reason, approaches which

attempt to sequentially localize the target without accounting for its dynamics are unlikely to succeed, and some smoothing or filtering will significantly improve the performance.

In contrast to existing approaches to distributed particle filtering [17, 18, 19, 20], which require a considerable amount of communication since nodes send information about the particle cloud, we propose to have each node run a local particle filter but nodes will instead communicate their measurements directly. When a node receives the message with measurements from any neighbor, it will incorporate the information into its own particle filter. Another attractive feature of the particle filter is that it is straightforward to incorporate a varying number of measurements at each step, so nodes can easily make use of whatever information is received from their neighbors.

4 Results

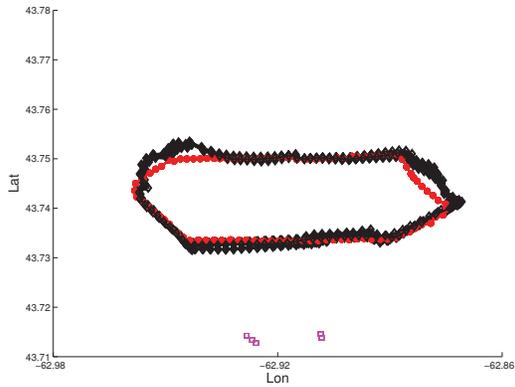
4.1 Modeling

A significant portion of the team’s effort so far has focused on investigating and processing the dataset provided. Because of the significant noise levels in the data and the limited number of measurements at each time step, it is necessary to exploit as much side information as possible. This includes accounting for the fact that the target must obey the laws of physics and cannot change locations arbitrarily in a fixed time interval. We have developed an initial motion model. The motion model assumes the target moves at a constant velocity and switches between heading straight ahead, executing a gradual turn, and executing a sharp turn. The model parameters have been estimated from the data. We have also begun to characterize the distribution of noise in the data. Each sensor produces roughly between 0 and 5 bearings measurements at each time step, and the offset between the measurements recorded and the “true” bearings (calculated based on the GPS coordinates provided for the ship towing the target) has somewhat heavier tails than a Gaussian distribution. These models are used in the initial particle filter implementation.

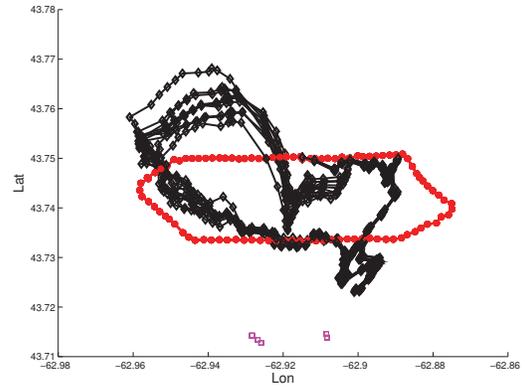
4.2 Tracking Results

As a baseline we consider the performance of a centralized particle filter which receives all of the measurements from all sensors. For the application scenario provided, if every sensor always received the messages sent by every other node (i.e., the network topology corresponded to a complete graph at every step), then each node would effectively be implementing a replica of this centralized filter. Thus, we can compare the performance of the distributed filters with that of this centralized filter.

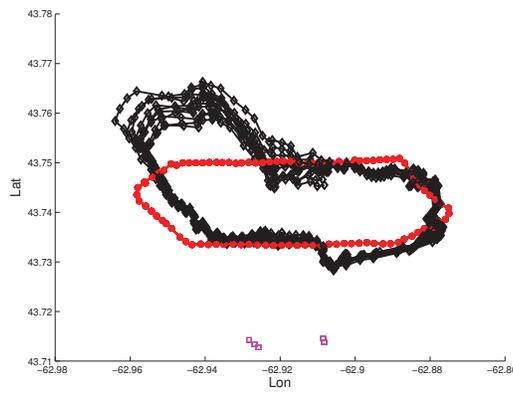
Preliminary results are shown in Figure 1. Fig. 1(a) shows the results of tracking if the sensors were to measure the true bearings at each time step (corresponding to one minute), with a small amount of additive white Gaussian noise. This “idealized” setting shows the performance limits arising from the configuration of the sensors and the current motion model. Fig. 1(b) shows the performance of a particle filter which uses all of the measurements available at every time step, and Fig. 1(c) shows the performance of a particle filter



(a) Tracking using the true bearings



(b) Tracking using all the measurements



(c) Tracking using the best measurements

Figure 1: Centralized particle filter using 10,000 particles. The red curve shows the true target trajectory and the black curves correspond to 10 different runs of the particle filter. The magenta squares indicate the locations of each sensor.

which clairvoyantly selects only the most accurate measurement from each sensor at each time step (disregarding the others). There is a distinct improvement from (b) to (c), but clearly there is still plenty of room for improvement.

4.3 Discussion

Now that we have a clear understanding of the performance of the centralized particle filter, we are conducting experiments to understand how distributed particle filters which only receive a subset of the observations will perform in comparison. Preliminary results indicate that performance depends significantly on the percentage of transmissions received from neighboring nodes, and on the position of the particular sensor. More detailed results will be provided in the final report.

5 Ongoing and Future Work

As is evident from the initial results shown in Figure 1, there is plenty of room for improvement, even in the centralized particle filter. To address the significant amount of spurious measurements due to clutter, we are working on incorporating *probabilistic data association* methods [22] into the particle filter. These methods assess each incoming message with respect to the current state estimate to distinguish (via a hypothesis test) whether or not the measurement is due to clutter or the target. This is done before fusing the measurement into the estimate, and should lead to improvements in the accuracy of the estimator. The improvement in the centralized estimator will be directly applicable to the distributed estimators as well.

On the theoretical side there is plenty of work to be done. Existing analyses of broadcast gossip algorithms assume the network topology is fixed, although not necessarily known. Little work has considered time-varying networks or unreliable links. We will begin to pursue an analysis along these lines. In terms of relaxation of the consensus concept, we are working on developing bounds on the ϵ -averaging time [7, 8], the worst-case number of iterations that must be performed in order to reach within ϵ of a consensus with high probability (i.e., with probability $1 - \delta$ for some $\delta > 0$). The ϵ -averaging time depends in general on the network size and structure (through the second largest eigenvalue of the graph Laplacian matrix), as well as the desired level of accuracy, ϵ , and the desired confidence δ . In relation to the goals of this project, such results will allow us to guarantee approximate consensus in finite amounts of time. The extension to time-varying topologies requires additional work to understand exactly how the time-varying nature of the topology affects the ϵ -averaging time.

These aspects and others will be the focus of our attention moving forward into the TIF AMBUSH project.

References

- [1] Sozer, E., Stojanovic, M., and Proakis, J. (2000), Underwater acoustic networks, *IEEE Journal of Oceanic Engineering*, 25(1), 72–83.
- [2] Olfati-Saber, R., Fax, J., and Murray, R. (2007), Consensus and Cooperation in Networked Multi-Agent Systems, *Proceedings of the IEEE*, 95(1), 215–233.
- [3] Tahbaz-Salehi, A. and Jadbabaie, A. (2007), Necessary and Sufficient Conditions for Consensus Over Random Independent and Identically Distributed Switching Graphs, In *Proceedings of the 46th IEEE Conference on Decision and Control*.
- [4] Olfati-Saber, R. and Murray, R. M. (2004), Consensus Problems in Networks of Agents with Switching Topology and Time-Delays, *IEEE Transactions on Automatic Control*, 49(9), 1520–1533.

- [5] Cortés, J. (2008), Distributed algorithms for reachign consensus on general functions, *Automatica*, 44(3), 726–737.
- [6] Cao, M., Morse, S. A., and Anderson, B. D. O. (2008), Reaching a consensus in a dynamically changing environment: Convergence rates, measurement delays, and asynchronous events, *SIAM Journal on Control and Optimization*, 47, 601–623.
- [7] Boyd, S., Ghosh, A., Prabhakar, B., and Shah, D. (2006), Randomized Gossip Algorithms, *IEEE Transactions on Information Theory*, 52, 2508–2530.
- [8] Dimakis, A. G., Kar, S., Moura, J. M., Rabbat, M. G., and Scaglione, A. (2010), Gossip Algorithms for Distributed Signal Processing, *Proceedings of the IEEE*, 98(11), 1847 – 1864.
- [9] Aysal, T., Yildiz, M., Sarwate, A., and Scaglione, A. (2009), Broadcast Gossip Algorithms for Consensus, *IEEE Transactions on Signal Processing*, 57(7), 2748–2761.
- [10] Fagnani, F. and Frasca, P. (2011), Broadcast gossip averaging: Interference and unbiasedness in large Abelian Cayley networks, *IEEE J. Selected Topics in Signal Processing*, 5(4), 866–875.
- [11] Franceschelli, M., Giua, A., and Seatzu, C. (2011), Distributed averaging in sensor networks based on broadcast gossip algorithms, *IEEE Sensors Journal*, 11(3), 808–817.
- [12] Shaochuan, W. and Rabbat, M. (2012), Broadcast gossip algorithms for consensus on stongly connected digraphs (extended version). arXiv:1208.4895.
- [13] Olfati-Saber, R. and Shamma, J. (2005), Consensus filters for sensor networks and distributed sensor fusion, In *Proc. IEEE Conf. on Decision and Control*, Seville, Spain.
- [14] Xiao, L., Boyd, S., and Lall, S. (2005), A scheme for robust distributed sensor fusion based on average consensus, In *Proceedings of the ACM/IEEE Conference on Information Processing in Sensor Networks*, Los Angeles, CA, USA.
- [15] Khan, U. and Moura, J. (2008), Distributing the Kalman Filter for Large-Scale Systems, *IEEE Trans. Signal Processing*, 56(10), 4919–4935.
- [16] Lopes, C. and Sayed, A. (2008), Diffusion Least-Mean Squares Over Adaptive Networks: Formulation and Performance Analysis, *IEEE Trans. Signal Processing*, 56(7), 3122–3136.
- [17] Farahmand, S., Roumeliotis, S., and Giannakis, G. (2011), Set-Membership Constrained Particle Filter: Distributed Adaptation for Sensor Networks, *IEEE Trans. on Signal Processing*, 59(9), 4122 –4138.

- [18] Hlinka, O., Slučiak, O., Hlawatsch, F., Djurić and, P., and Rupp, M. (2010), Likelihood consensus: Principles and application to distributed particle filtering, In *the Forty Fourth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, pp. 349–353.
- [19] Oreshkin, B. and Coates, M. (2010), Asynchronous distributed particle filter via decentralized evaluation of Gaussian products, In *Proc. ISIF Int. Conf. Information Fusion*, Edinburgh, UK.
- [20] Ü, D., Coates, M. J., and Rabbat, M. G. (2011), Distributed auxiliary particle filters using selective gossip, In *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Proc. (ICASSP)*, pp. 3296–3299, Prague, Czech Republic.
- [21] Ristic, B., Arulampalam, S., and Gordon, N. (2004), *Beyond the Kalman filter: particle filters for tracking applications*, Artech House.
- [22] Kirubarajan, T. and Bar-Shalom, Y. (2004), Probabilistic data association techniques for target tracking in clutter, *Proceedings of the IEEE*, 92(3), 536–557.