Graph-based Compression of Particle Clouds for Gossip-Based Decentralized Particle Filters

Michael G. Rabbat[1], Mark J. Coates[1], Jun Ye Yu[1], Deniz Üstebay[1], and Stephane Blouin[2]

[1]Department of Electrical and Computer Engineering, McGill University, Email: michael.rabbat@mcgill.ca, jun.y.yu@mail.mcgill.ca, deniz.ustebay@mail.mcgill.ca, mark.coates@mcgill.ca

[2]DRDC Atlantic Research Centre, Email: stephane.blouin@drdc-rddc.gc.ca

May 30, 2014

Abstract

In gossip-based decentralized particle filtering, a network of sensors collaborate to track one or more targets. Each node in the network maintains and updates a copy of the particle filter, and gossip algorithms are used to simultaneously synchronize the filters and fuse measurements after each measurement round. Particle filtering methods are attractive for tracking when the state dynamics and/or the measurement model are nonlinear and/or non-Gaussian. The accuracy of particle filters is generally related to the number of particles used to approximate the posterior density. A significant challenge in the decentralized setting is to accurately fuse and synchronize the weighted particle approximations while reducing the communication overhead. This is especially important when nodes communicate over unreliable wireless channels, in which case bandwidth-limitations restrict the amount of information that can be communicated in each measurement round. Also, when nodes are battery-powered, the energy consumed for each wireless transmission is significantly higher than that consumed for computation, and so reducing the amount of information communication can extend the lifetime of the system.

We propose a graph-based method to synchronize the particle weights to values which are approximately equal to those that would be computed if every node had access to the measurements from all other sensors. Our approach leverages the fact that, in the decentralized particle filtering setup, it is possible to keep the particle locations at all nodes synchronized through the use of shared randomness, and only the particle weights need to be evaluated when new measurements are obtained. The weight update requires evaluating the joint likelihood of the measurements from all sensors evaluated at the particle locations. Under the assumption that measurements at each node are conditionally independent given the state, the joint log-likelihood factorizes into the sum of node-specific log-likelihoods, so linear methods can be used for synchronization. Moreover, the joint log-likelihood values at nearby particles tend to be similar. In order obtain a parsimonious representation of the particle weights, we construct a graph over the particles by connecting nearby particles with an edge. The Laplacian eigenvectors of this graph are then used as a basis for transform coding of the weight values, allowing us to efficiently compute approximate weight values by only agreeing on the values of a few most significant transform coefficients over the network. By varying the number of transform coefficients used we can trade off the accuracy of the approximation with the amount of communication overhead. The efficacy of the proposed approach is illustrated using simulated examples as well as data from at-sea trials of bearings-only tracking problems.

DRDC-RDDC-2015-N037