Localizing Large Numbers of Targets Without Data Association Using Teams of Mobile Robots

Philip Dames and Vijay Kumar*

I. Introduction

This work considers the problem of localizing large, but unknown, numbers of targets with low cost, low resolution sensors on multiple mobile robot platforms. We formulate the multi-target localization problem using the probability hypothesis density (PHD) filter, which allows us to simultaneously estimate the number of targets and their positions within the environment without the need for data association. In some situations, the robot team must search in communication-limited environments, either due to large amounts of clutter, wireless interference, or large distance scales. In these instances, a centralized solution that relies upon consistent inter-robot communication would perform poorly. To circumvent this, we propose the use of a network architecture, outlined in Figure 1, that allows robots to communicate with each other on a local peer-to-peer basis, and with a server or the cloud via access points, exchanging measurements and poses to update their beliefs about the targets and plan future actions. The server provides a mechanism to collect and synthesize information from all robots and to share the global, albeit time-delayed, belief state to robots near access points. Such communication concerns only arise due to the presence of multiple agents in the team.

We design a decentralized control scheme that exploits this communication architecture and the PHD representation of the belief state. Specifically, robots move to maximize mutual information between the target set and measurements, both those they collect themselves and those available by accessing the server, balancing local exploration with knowledge sharing across the team. The expected information gain of self-collected measurements is computed, over a finite time horizon, as the mutual information between the target locations and the binary event of receiving no target detections, effectively hedging against non-informative actions in a computationally tractable manner. Additionally, robots coordinate their actions with other robots exploring the same local region of the environment. We present simulation results with small teams of robots equipped with range-only sensors seeking tens to hundreds of targets in indoor environments, such as those seen in Figure 2.

A. Related Work

One common decentralized architecture is Decentralized Data Fusion (DDF), first described by Grime and Durrant-Whyte [3], in which each robot manages its own copy of the joint belief and aggregates data from the other robots through channel filters which only admit information that is distinct from their current belief. The DDF framework is particularly amenable to Gaussian beliefs, as the information form of the Kalman filter allows for efficient, low-bandwidth updates. However, more complicated belief representations, such as the PHD, often require overly conservative approaches to data fusion. A centralized approach reduces redundant computations and ensures that all robots have an identical belief state, but is not always possible to implement in practice.

Our solution takes the best of each of these approaches, allowing robots to communicate on a peer-to-peer basis in a decentralized fashion while also including communication with a centralized server or a cloud, which robots may access via the existing network infrastructure in the environment. This idea of robots relying on information from a server has been called cloud robotics and has recently generated quite

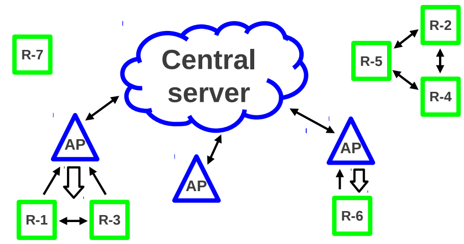

Fig. 1. Diagram of the network structure. Robots (green squares) are able to opportunistically communicate on a peer-to-peer basis with nearby robots as well as exchange information with the central server through access points (blue triangles). The communication links originating from robots are all relatively low-bandwidth while the downlink from the server may be higher bandwidth.

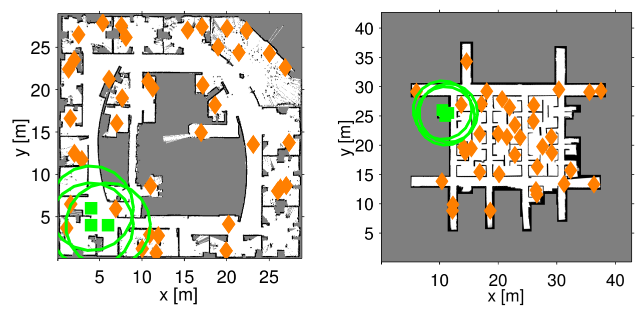

(a) Environment 1 (b) Environment 2

Fig. 2. Example maps used in simulation runs. Robots are indicated by green squares, sensor footprints by green circles, and targets by orange diamonds.

*Philip Dames and Vijay Kumar are with the Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA {pdames, kumar}@seas.upenn.edu

a bit of excitement [4]. A similar idea was also used for estimation and control of groups of robots by Michael et al. [6], where an Asymmetric Broadcast Control (ABC) was used to synthesize locally derived information and provide low-resolution global information to the group.

Situations with multiple targets, and the problem of simultaneously estimating the number of targets and their locations, is not a trivial extension of the single target case. Our approach relies on a filtering technique that is based on the mathematical construct of a random finite set [5] which, as we will see, completely avoids the need to solve data association, *i.e.,* matching measurements to target tracks. This offers a significant computational advantage, especially when the number of targets is large, and allows us to deal with unreliable sensors which can experience false negatives (*i.e.,* missing targets that are actually there), false positives (*i.e.,* seeing targets that are not actually there), and noisy measurements.

There is a relatively limited body of work on active control based on the RFS framework, with attempts to maximize information using Rényi's definition by Ristic and Vo [7] and Ristic et al. [8]. In this work, the measurement model involves a summation over all possible data associations, and the authors present simulation results of a single robot seeking three targets in an open environment. Dames et al. [2] use Shannon's information to track a small number of targets, but do not assume that the target positions are independent. In all of these works, the resulting control calculations quickly become computationally intractable for large numbers of targets and measurements. This current work builds upon the framework presented by Dames and Kumar [1] where, unlike most other work on information-based control, we explicitly consider the information benefits of communication.

II. Preliminary Results

The results of a series of simulations in environments 1 and 2, which contain 36 and 50 targets respectively, are shown in Figure 3. Note that we see that the average expected target cardinality approaches the true number of targets and that the standard deviation across the trials tends to decrease over time, showing that the estimates are consistent. Figure 3 also shows typical final PHDs in each environment, and we see that the targets are generally well-localized. The estimation is not perfect: there are occasional false positive detections, and densely packed targets are occasionally represented by a single cluster of high mass rather than individual clusters of unit mass. This is to be expected when the filter lacks target labels because a noisy measurement from one target looks equivalent to a good measurement of a nearby target.

When looking for tens of targets with three robots, the computational load is relatively small. Assuming that robots move at a constant speed of 0.5 m/s, robots spend an average of 1–3% of the time on filtering updates, 1–5% of the time on control computations, and the remaining time driving, with variations due to the complexity of the environment.

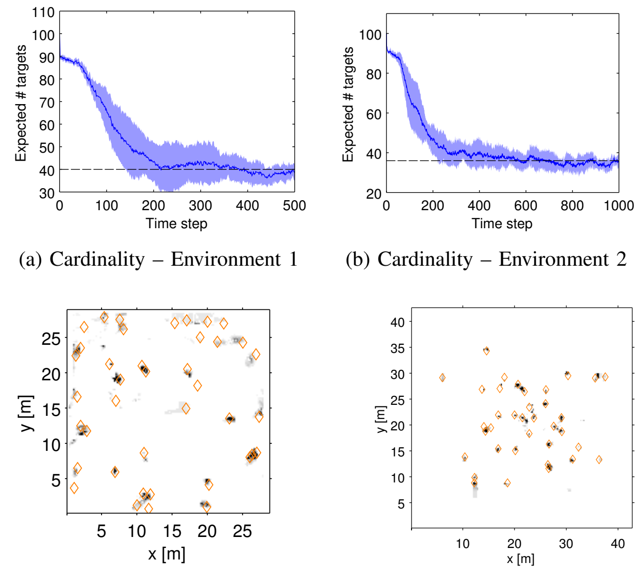

(a) Cardinality – Environment 1 (b) Cardinality – Environment 2

(c) PHD – Environment 1 (d) PHD – Environment 2

Fig. 3. Data showing the expected number of targets over time for environments 1 and 2, where (a) and (b) show the average expected number of targets over time. The solid blue line is the average across 10 trials and the shaded blue area is ± 1 standard deviation. Typical final PHD estimates are shown in (c) and (d) for the same environments, with the true target locations denoted by the orange diamonds and the estimated target density field in the background (grayscale, with black corresponding to a higher expected number of targets).

References

[1] P. Dames and V. Kumar. Cooperative Multi-Target Localization with Noisy Sensors. In *IEEE Intl. Conf. on Robotics and Automation*, pages 1877–1883, May 2013.

[2] P. Dames, M. Schwager, V. Kumar, and D. Rus. A Decentralized Control Policy for Adaptive Information Gathering in Hazardous Environments. In *IEEE Conf. on Decision and Control (CDC)*, pages 2807–2813, December 2012.

[3] S. Grime and H. Durrant-Whyte. Data Fusion in Decentralized Sensor Networks. *Control engineering practice*, 2(5):849–863, 1994.

[4] E. Guizzo. Cloud Robotics: Connected to the Cloud, Robots Get Smarter. *IEEE Spectrum. Available via http://spectrum. ieee. org/automaton/robotics/robotics-software/cloud-robotics*, 2011.

[5] R. Mahler. *Statistical Multisource-Multitarget Information Fusion*, volume 685. Artech House Boston, 2007.

[6] N. Michael, J. Fink, S. Loizou, and V. Kumar. Architecture, Abstractions, and Algorithms for Controlling Large Teams of Robots: Experimental Testbed and Results. *Int. J. Robotics Research*, pages 409–419, 2011.

[7] B. Ristic and B.N. Vo. Sensor Control for Multi-Object State-Space Estimation Using Random Finite Sets. *Automatica*, 46(11):1812–1818, 2010.

[8] B. Ristic, B.N. Vo, and D. Clark. A Note on the Reward Function for PHD Filters with Sensor Control. *IEEE Trans. Aerosp. Electron. Syst.*, 47(2):1521–1529, 2011.