Consistent Decentralized Graphical SLAM with Anti-Factor Down-Dating

Alex Cunningham, Vadim Indelman and Frank Dellaert

***Abstract*—This report presents our recent and ongoing work developing a consistent decentralized data fusion approach for robust multi-robot SLAM in dangerous, unknown environments. The DDF-SAM 2.0 approach extends our previous work by combining local and neighborhood information in a single, consistent augmented local map, without the overly conservative to avoiding information double-counting in the previous DDF-SAM approach. We introduce the anti-factor as a means to subtract information in graphical SLAM systems, and illustrate its use to both replace information in an incremental solver and to cancel out neighborhood information from shared summarized maps. Evaluations in a synthetic example environment demonstrate that we avoid double-counting information.**

I. INTRODUCTION

A key enabling technology for robust multi-robot teams capable of operation in challenging environments is developing a Simultaneous Localization and Mapping (SLAM) system capable of fusing sensory information in a decentralized manner. In our previous work on decentralized robot perception [1], [2], we introduced DDF-SAM (which we will denote as DDF-SAM 1.0 in this paper), a SLAM paradigm extending Smoothing and Mapping (SAM) [3] techniques to the Decentralized Data Fusion (DDF) problem by robots performing local SLAM and sharing a subset of map variables with neighbors in the form of *summarized maps*. The key component of the approach to send summarized densities between robots allows enables application-specific approaches choices of information to transmit, which allows for throttling to match the available communication bandwidth between platforms.

However, DDF-SAM 1.0 has two key shortcomings: 1) an overly conservative approach for avoiding information double-counting, and 2) reliance on a batch marginalization approach for map summarization. The method employed for avoiding information double-counting enforced a strict separation between a robot's local map and the neighborhood map comprised of data from neighboring robots - an approach that avoids double-counting, but results in each robot maintaining two separate, but incomplete maps of its environment. Furthermore, the batch summarization technique used to marginalize non-shared variables from the local map performs a computationally expensive and separate factorization of the entire local system, which increases in complexity over time.

Our recent and ongoing work, DDF-SAM 2.0, addresses these shortcomings with the following contributions towards effective decentralized perception: 1) the *augmented local system* as a single, consistent incremental solver on each robot blending local and neighborhood map information, 2) the *anti-factor* as a tool to avoid double-counting by down-dating summarized maps, and 3) a discussion of a summarization approach designed to exploit the Bayes Tree [4] factorization.

The *double-counting* problem in distributed systems is overconfidence due correlated information propagating through the network. Consider an example with three robots A, B, and C: robot A shares local information with robot B, then robot B incorporates this information into its own map and shares information with robot C. Robot C incorporates the information from robot B (which contains information from robot A), and finally shares its new estimate with robot A. By completing this cycle in the network, robot A has now received information derived from its own local estimates in addition to measurements from robots B and C. If robot A treats the message from robot C as independent of its own measurements, it will double-count its local measurements, leading to overconfidence.

II. DDF-SAM 2.0

This report outlines DDF-SAM 2.0, an approach to preventing the double-counting of information, even while combining both local and neighborhood information into a single incremental Bayes tree solver, resulting in an *augmented* local graph. This representation and corresponding algorithm enables each robot to maintain a consistent SLAM solution with an extended sensor horizon provided by neighboring robots. The following sections refer to the core representations of probabilistic inference for SLAM, primarily factor graphs as a sparse factorization of measurements, the variable elimination algorithm for performing inference and marginalization, and the Bayes tree as a partially solved form of a factor graph after elimination. For a detailed treatment, see [3], [4].

It is a simple extension of DDF-SAM 1.0 to assemble a single factor graph containing both the local density and summarized maps from neighboring robots, and incrementally maintain single solver, but the difficulty comes when summarizing this augmented local system to share with other robots. The key insight enabling the augmented local system comes from additive nature of information when incorporating additional measurements in linear systems: not only is it possible to add information to a system, but to *subtract* information as well. To that end we introduce the *anti-factor* to down-date factorized estimates.

Updating the solver with local factors remains essentially the same as in a single-robot iSAM case, in which we simply add new factors to the existing graph and update the Bayes tree solution. While adding a new neighborhood system is an operation that consists of adding a new set of factors to the underlying linear solver on a robot, because summarized maps actually *replace* information from previous timesteps, we use anti-factors to cancel the previously added information. In DDF-SAM 1.0, we solved this problem by reconstructing the entire graph from only the latest cached summarized maps, but this results in a large batch optimization procedure to integrate a new summarized map.

With a down-dating procedure, however, we can subtract out the information from the old summarized map from time k and then add the new summarized map from $k+1$. In order to add a new summarized graph while removing double-counted information, we instead add the anti-factor negation of the summarized map from time k and the new map from time $k+1$ simultaneously. From the perspective of updating an incremental solver, this is no different an operation than adding new local factors.

While it is possible to use the existing Bayes tree to compute a summarization, the measurements from the neighboring robots are still present in the system, we can cancel out this neighborhood information through the use of anti-factors. The summarized graph sent by robot r becomes a graph combining a summarized density corresponding to the local information on r, with anti-factor negations of cached summarized maps from other robots. While this results in adding additional factors to a summarized graph before transmission, it should be noted that linear factors in Hessian form can be combined additively, ensuring the size of summarized maps does not necessarily increase with neighborhood size. Because we have exactly canceled out the information from other robots from before transmission, no information will be double-counted as these summarized maps propagate between robots.

While DDF-SAM 1.0 used batch elimination approach to create an exact joint density on the shared variables, we can also use the Bayes tree to perform summarization as a part of periodic variable reordering. Rather than maintaining a separate full system as in DDF-SAM 1.0, we can choose a variable elimination ordering to simultaneously minimize fill-in as well as place all of the wanted variables at the top of the Bayes tree. With the orderings in the incremental solver created to facilitate which allows us to extract the summarized map directly from the Bayes tree. This approach, like batch summarization, is an exact marginalization operation, so no information is lost in the process.

III. Evaluation

In our simulated example, we model robot three platforms, each equipped with a downwards facing camera, an inertial measurement unit and GPS. The camera runs an image-based feature detector, which produces landmarks with globally unique labels; for a treatments of data associations between platforms, see [2]. To evaluate the exact cancellation of information during inference without the effect of linearization error, the robots share linearization points upon sighting new landmarks.

Example plots comparing the augmented local system and neighborhood-only solution appear in Fig. 1. These plots

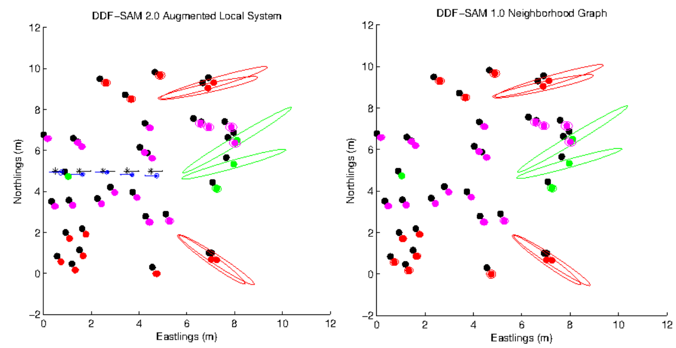

Figure 1. The DDF-SAM 2.0 augmented local system (left) solution and the DDF-SAM 1.0 neighborhood-only map 2D plots with covariances tracking the center robot in the simulated scenario, with comparison to ground truth. Black denotes ground truth for both poses (present only in augmented local system) and landmarks, blue is the final trajectory with covariances, green landmarks are purely local landmarks, magenta landmarks overlap with neighbors, and red landmarks are only observed through neighbors.

were extracted after a small number of poses to ensure covariance ellipses are readily visible, and are projected into a 2D plane for clarity of presentation. In these cases, we compare the effect of approximate summarization with exact batch summarization and the pure local solution on variables shared between robots. Note that the marginal covariance ellipses for variables in the neighborhood-only system, which is guaranteed consistent by construction, are the same as the local augmented system.

We computed a numerical score for the information present, defined as the trace of the information matrix of each summarized map. We computed both Bayes tree-based summarization and the full batch partial elimination, and computed the difference in information scores for each technique, and found that the traces of the information matrices differed on the order of $1e{-}7$ at most, across all robots over the full trajectory. This deviation is within numerical error tolerances.

IV. Future Work

In future work, we will relax the core assumptions made by the evaluation presented in this paper by incorporating relinearization into the the approach, as well as handling separate linearization points between robot platforms. We also are actively investigating the use of approximate summarization techniques to allow for faster computation and communication at the cost of less precise maps.

References

[1] A. Cunningham, M. Paluri, and F. Dellaert, "DDF-SAM: Fully distributed slam using constrained factor graphs," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2010.

[2] A. Cunningham, K. Wurm, W. Burgard, and F. Dellaert, "Fully distributed scalable smoothing and mapping with robust multi-robot data association," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2012.

[3] F. Dellaert and M. Kaess, "Square Root SAM: Simultaneous localization and mapping via square root information smoothing," *Intl. J. of Robotics Research*, vol. 25, pp. 1181–1203, Dec 2006.

[4] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Intl. J. of Robotics Research*, vol. 31, pp. 217–236, Feb 2012.