Distributed Perception and Estimation: a Short Survey

Vadim Indelman*

***Abstract*— We present a short survey on the state of the art in distributed perception and estimation, a central problem in multi-robot systems. Autonomously perceiving the world by a group of robots requires a number of key capabilities, including data association, mapping and robot localization, and as such is tightly related with distributed multi-robot simultaneous localization and mapping (SLAM). In this article we first review dominant approaches in multi-robot perception, SLAM, and collaborative localization (CL), and then focus on two of the key challenges in these approaches: consistent estimation and robust perception.**

I. Introduction and Problem Formulation

Distributed perception and estimation is a central problem in multi-robot systems, and as such has been investigated in recent years by different research communities, including the control, robotics, artificial intelligence and computer vision communities. This fundamental problem is encountered in numerous application domains, such as multi-robot tracking, cooperative localization (CL) and navigation (CN), multi-robot simultaneous localization and mapping (SLAM), distributed multi-view 3D reconstruction and mapping, and multi-robot monitoring. Autonomously perceiving the world by a group of robots requires a number of key capabilities, including map merging, data association, and robot localization. As such, distributed perception is tightly related with distributed multi-robot SLAM.

The general problem can be formulated within the decentralized data fusion (DDF) paradigm [19], where robots (or mobile sensors) are to infer random variables of interest based on local sensor measurements and information communicated by nearby robots. A distributed architecture has no central computational unit; instead, each robot performs inference on its own using the currently available information. Moreover, it naturally supports communication topology that may be changing over time as the robots move. Such a framework has therefore a number of desirable characteristics [19], such as scalability to large number of robots, and robustness to failure since, as opposed to a centralized framework, a decentralized architecture does not have a single point of failure.

The identity of the random variables to be inferred changes according to the problem at hand and in many cases differs for each robot. For example, in tracking applications one is typically interested in estimating the target position and possibly additional target states (e.g. velocity). In other problems, however, each robot thrives to infer a different set of latent variables. Thus, in collaborative localization and navigation, each robot aims to estimate its own state and possibly the states of other robots. A similar situation arises also in distributed perception and multi-robot SLAM, where robots need to infer the model of the surrounding environment and also localize themselves.

This article aims to provide an overview of the recent developments and accomplishments in distributed perception and estimation, with a focus on multi-robot SLAM and cooperative localization. We also discuss the state of the art addressing two of the key challenges often encountered in these problems: consistent distributed inference, and robust distributed perception. Both aspects are crucial for safe and reliable multi-robot systems. The former addresses the issue of double counting information that can easily happen in robot sensor networks and could result in failure due to overconfident estimation. In the latter aspect, robust distributed perception, we review distributed approaches that aim to identify and reject outliers in data association.

Throughout this short article we aim to provide intuitive explanations of the addressed problems and mostly refrain from mathematical exposition of thereof, while the interested reader can easily follow the cited publications for further details. Moreover, we refer the reader to recent surveys on important aspects in distributed perception and estimation that are not covered herein, such as distributed computer vision algorithms [47] and multi-sensor distributed target tracking [49].

II. Literature Review

A. Multi-Robot Perception, SLAM and CL

A key capability in multi-robot systems is collaborative localization and mapping in unknown or partially unknown environments, problems that are instantiations of distributed estimation. By sharing information between the robots, the performance of individuals in the group can be significantly improved. Multi-robot localization, mapping and SLAM have been extensively investigated by the robotics community in the last two decades.

The developed approaches vary in the considered observation models, inference engines and undertaken assumptions. Below we review some of these approaches; due to space limitation we primarily focus on distributed multi-robot inference approaches that are based on iterative nonlinear optimization techniques and extended Kalman and information filters, while noting there is a rich branch of literature that manages nonlinearity with nonparametric techniques such as particle filters (see e.g. [23], [25], [10]).

In a seminal work [34], Kurazume et al. proposed one of the first approaches for cooperative positioning with multiple

*Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel.

robots: robots are divided into two groups, each time one group remains stationary while the other moves and takes measurements of the former group. Another groundbreaking method for multi-robot localization has been developed by Roumeliotis and Bekey [48], where a fully distributed estimation algorithm was developed, employing an (extended) Kalman filter and assuming the robots are capable of taking relative pose measurements with respect to each other and share a common global frame.

Nerurkar, et al. [43] presented a distributed maximum a posteriori (MAP) estimator for multi-robot collaborative localization, proposing a distributed data-allocation scheme that enables robots to simultaneously process and update their local data. Their approach incorporates multi-robot relative observations (e.g. bearing, orientation, range) to attain higher levels of estimation accuracy, as done in [48]. Bailey et al. [6] also consider joint localization with relative observations (range-bearing measurements), developing a graph-based approach that fuses local information in a decentralized fashion and uses a central server to fuse local information from different robots and multi-robot relative observations. The authors state the latter can be decentralized to improve robustness.

So far, the mentioned approaches considered multi-robot observations involving robot states from the same time instant. In the context of multi-robot perception and SLAM, where robots operate in and make observations of unknown environments, the corresponding multi-robot constraints describe different robots observing a mutual scene, *not necessarily* at the same time. The resulting measurement equations (or measurement likelihood terms) no longer can be represented in terms of robot states at a given time, and instead either involve additional random variables (e.g. landmarks) or robot states from different time instances, thereby posing additional challenges for distributed inference. The research community has been investigating in the last years approaches for distributed inference in these scenarios.

These research efforts include approaches that extend the smoothing and mapping (SAM) paradigm, originally introduced by Dellaert and Kaess [17], to the multi-robot case (e.g. [1], [15], [32], [16], [14]). Thus, Anderson et al. [1] developed a collaborative SAM (C-SAM) approach within a centralized multi-robot framework, while Cunnigham et al. [15], [14] presented an extension of SAM considering a distributed architecture and formulating the problem within a DDF framework [19]. In both cases, to preserve sparsity [20] and support efficient computations, robots infer their past and current states and map the environment in which they operate. As such, information fusion in these approaches also involves *map merging*, a fundamental problem in multi-robot systems that has been extensively investigated in recent years (see, e.g. [33], [11]). A related approach was also developed by Indelman et al. [27], considering distributed cooperative localization and navigation using multiple view geometry to formulate constraints representing image observations of mutual scenes acquired by different robots, possibly at different time instances. The approach was later generalized to arbitrary multi-robot observation models [28]. More recently, Walls et al. [52], [53] develop an approach for underwater cooperative localization considering *faulty* low-bandwidth communication channels. Their approach is based on the decentralized extended information filter (DEIF) algorithm [54] and is capable of optimally fusing information transmitted by different robots over an extremely faulty communication channel and, remarkably, exactly reproduce the estimate of a centralized filter [52].

While many of the above approaches typically assume the initial relative pose between the robots is known, i.e. the robots share a common reference frame, significant research endeavors have also been recently devoted to relax this assumption, enabling distributed inference and perception also when the robots do not initially share a common reference frame. See, e.g., [26], [58], [1], [10], [16], [30] and the references therein.

Another approach for decentralized estimation uses consensus algorithms, which were developed initially to address multi-robot control problems (see, e.g. [38]). These approaches solve the rendezvous problem for multi-agent control, and have been recently applied also to estimation problems. Thus, Yang et al. [56] develop a distributed approach for target tracking using a consensus estimator. Another impressive example is the work by Aragues et al. [2], where the authors combine consensus with information filters to perform distributed map merging. The latter is an important problem in distributed perception, with one of the challenges being the possibility that the maps to be merged do not actually overlap. Additional research on distributed map merging include, e.g., [33], [11].

Distributed and decentralized inference via belief propagation (BP) [45] has also received attention in literature. The BP algorithm produces posterior probability distributions equivalent to centralized algorithms when run on networks without loops. When the communication topology includes cycles, one may run multiple iterations of the BP algorithm, an algorithm known as loopy BP (LBP). Unfortunately, LBP is not guaranteed to converge when run on loop networks, and if it does, it might not converge to the correct posterior distribution (see, e.g., [57]). Nevertheless, the algorithm has been empirically shown to converge to approximately correct posteriori belief in many cases [46]. Recently, the BP algorithm has been applied to decentralized multi-robot SLAM [18], where instead of using an LBP over loopy graphs (representing communication topology with cycles), the authors use BP over a spanning tree of the graph. Interestingly, it is noted [18] that using LBP over loopy graphs may double count information, and therefore produce inconsistent belief (see Section II-B).

Having discussed some of the key approaches in distributed inference and perception, we now review literature addressing two particular challenges in these problems: consistent distributed inference (prevent information double counting) and robust distributed perception.

B. Consistent Decentralized Inference

When considering decentralized data fusion one has to be careful not to double count information, i.e. use the same information more than once, as otherwise estimation may become inconsistent, overconfident. Clearly, this is particularly undesired in multi-robot systems where it could lead to mission failure and compromise safety.

To better understand the notion of information double counting, we recall that the Bayes rule can be used to fuse either different independent probability distribution functions (pdf) or dependent pdfs if explicitly accounting for this dependency (see, e.g., [7]). The problem occurs when this dependency exists and is neglected. Unfortunately, in decentralized inference double counting can easily happen. To see that, consider a cyclic communication topology between robots A, B and C: A talks with B, B talks with C, and C communicates with A. Assume A transmits a message, that is conditioned on its local measurements (e.g. $p\left(x|Z_A\right)$), to B, which calculates and then passes the posterior $p\left(x|Z_B, Z_A\right)$ to C. Robot C does the same and passes $p\left(x|Z_C, Z_B, Z_A\right)$ to A. At this point, if A treats the received message, $p\left(x|Z_C, Z_B, Z_A\right)$, as independent with respect to its local belief $p\left(x|Z_A\right)$, it will double count information - the measurement Z_A will be effectively used twice.

The problem of fusing information from different sources while properly keeping track of common information (Z_A in the example above), or calculating the appropriate cross-covariance terms, has been addressed by several researchers [24], [44], [48], [5], [28], [14].

An often considered scenario is of multi-robot observations involving only current time instances, e.g. inter-robot relative measurements. For this case Roumeliotis et al. [48] maintain an augmented covariance matrix for all robots, with the required cross-covariances, and show that appropriate Kalman filter calculations can be distributed among different robots. Grime et al. [24] and Nettleton et al. [44] address a similar problem using joint information. Their approach models correlation between different states, which can be subtracted since update step in information form is simple addition. Bahr et al. [5] introduced a method for consistent cooperative localization with a bank of filters for tracking the origins of measurements to prevent double-counting. A key advantage of their approach is that broadcasts do not need to be received by all participating robots, making it in particular attractive to distributed multi-robot framework. Fallon et al. [21] developed a distributed bookkeeping strategy to ensure that information is incorporated in a consistent manner.

The problem becomes more complicated if robots can share with each other marginal distributions of also other latent variables than their current states (e.g. past pose, landmarks), as common in multi-robot SLAM (see Section II-A). For example, robot r may observe an area that was observed by another robot r' some time ago, say t_i. In order for robot r to incorporate this information, robot r' has to transmit its marginal distribution either over the mutually observed landmarks or over the appropriate past state $x_i^{r'}$. As earlier, a robot that receives such a message has to track its common information (or to calculate the appropriate cross-covariance term). The difficulty is due to the fact that the identity of these latent variables is typically unknown ahead of time while maintaining all the possible cross-covariance terms is impractical (as opposed to [48]).

This challenging problem has received considerably less attention, while the above mentioned approaches typically cannot be directly applied. In particular, Indelman et al. [28] used a graph-based approach to calculate the correlation terms for consistent information fusion in the EKF framework. Cunningham et al. [14] develop a consistent DDF framework for smoothing and mapping, where information double counting is avoided by down-dating information, similarly to [24] [44]. Consistent distributed inference has been recently investigated also by Walls et al. [52] in the context of cooperative underwater navigation with *faulty* communication (as mentioned in Section II-A).

A general framework for fusing correlated information without the correlation being known has been proposed by Julier and Uhlmann [31]. Their approach, known as covariance intersection (CI), produces a consistent posterior (but suboptimal) covariance even though the correlation is unknown, and as such is appealing for distributed inference. Arambel et al. [4] present an application of CI for a group of space vehicles, where relative position measurements are communicated in a ring topology. Recently, Carrillo-Arce et al. [12] suggested to use CI for decentralized multi-robot cooperative localization, showing their approximate approach has linear processing and communication complexity in the number of robots.

C. Distributed Robust Perception

A key requirement for reliable and robust multi-robot operation is the ability to autonomously and consistently perceive the world. To do so, distributed data association approaches aim to determine the correct association between local measurements of the world (e.g. images) and measurements communicated by other robots. Assuming data association is solved, one can proceed to distributed inference, as discussed in Section II-A.

While data association aspects have been extensively investigated by the computer vision community, existing approaches (e.g. image matching using the RANSAC algorithm [22]) typically assume a centralized framework, and as such are not directly applicable to distributed multi-robot systems.

A distributed multi-robot framework, where each robot has access to only partial information and, furthermore, when this information is obtained only incrementally as the robots move and explore the environment, makes the data association problem even more challenging. A particularly challenging aspect is perceptual aliasing, where different environments that are similar in appearance (e.g. two similar buildings or corridors) should not be mistakenly considered as the same environment.

Distributed data association approaches only recently began to be investigated. Aragues et al. [3] presents a distributed approach for consistently matching several sets of features observed by a team of robots, while detecting and resolving conflicting associations. Li et al. [37] develops a distributed optimization framework for dealing with outliers using consensus algorithms and duality theory [8]. Also Montijano et al. [39], [40] developed distributed consensus approaches for outlier rejection. In particular, their approach integrates a distributed RANSAC [22] with distributed averaging [55]. The latter facilitates distributed calculation of maximum likelihood estimation, but is not resilient to outliers. To filter the outliers, the authors developed a decentralized variation of RANSAC, where a set of hypotheses are computed from random subsets of robots and then voted by all of them. Similar to the original RANSAC algorithm, the hypothesis with the best support (in this case - the larger number of votes) is considered as the correct hypothesis. A related idea has been also developed in the context of distributed multi-robot SAM [16].

Another related challenge in robust perception is generating correct loop closure constraints which are essential for high-accuracy inference over time. Unfortunately, current state of the art methods in place recognition (e.g. FAB-MAP [13]) are not error-free and do occasionally produce incorrect results, especially in the presence of perceptual aliasing. As earlier, introducing outlier correspondences (i.e. spurious measurements) into the inference layer can lead to catastrophic results. To be resilient to outliers overlooked by data association approaches, the robotics community has been therefore recently focusing on robust graph optimization techniques (e.g. [50], [35], [51], [36], [9]). However, these approaches are typically developed for the single robot case, and do not consider multi-robot systems.

The latter has been recently explored in [30], [42], considering the problem of robust multi-robot inference while the initial relative poses of the robots are unknown. The developed approach incorporates data association aspects into Bayesian inference and then, similarly to [36], resorts to expectation-maximization (EM) [41] to efficiently calculate the maximum a posteriori (MAP) solution over camera (robot) poses while marginalizing out latent variables that represent for each correspondence if it is inlier or outlier. In a subsequent work [29], the authors started considering incremental aspects of robust data association in the presence of perceptual aliasing, addressing the question whether sufficient amount of information has been obtained to make decisions regarding the unknown relative poses and multi-robot data association.

III. Conclusions

In this short survey we reviewed the state of the art in distributed perception and estimation. The primary focus was given to distributed multi-robot SLAM and cooperative localization, and to two particular challenges that are often encountered in these and other related problems: consistent decentralized inference, and distributed robust perception.

While significant progress has been made in recent years in each of the mentioned aspects, research on distributed perception and estimation is far from being mature. In particular, research directions that could be of interest for future research include consistent high-accuracy distributed inference for large group of robots (where extensive book-keeping is impractical), robust *online* distributed perception in the presence of perceptual aliasing, and active aspects of distributed perception and inference. The latter could be particularly beneficial to further improve performance and robustness of multi-robot systems.

References

[1] L. Andersson and J. Nygards. C-SAM : Multi-robot SLAM using square root information smoothing. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2008.

[2] R. Aragues, J.Cortes, and C. Sagues. Distributed consensus on robot networks for dynamically merging feature-based maps. *IEEE Transactions on Robotics*, 2012.

[3] R. Aragues, E. Montijano, and C. Sagues. Consistent data association in multi-robot systems with limited communications. In *Robotics: Science and Systems (RSS)*, Zaragoza, Spain, June 2010.

[4] P.O. Arambel, C. Rago, and R.K. Mehra. Covariance intersection algorithm for distributed spacecraft state estimation. In *American Control Conference*, volume 6, pages 4398–4403, 2001.

[5] A. Bahr, M.R. Walter, and J.J. Leonard. Consistent cooperative localization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3415–3422, May 2009.

[6] T. Bailey, M. Bryson, H. Mu, J. Vial, L. McCalman, and H. Durrant-Whyte. Decentralised cooperative localisation for heterogeneous teams of mobile robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2011.

[7] T. Bailey, S. Julier, and G. Agamennoni. On conservative fusion of information with unknown non-gaussian dependence. In *Intl. Conf. on Information Fusion, FUSION*, 2012.

[8] Dimitri P Bertsekas and John N Tsitsiklis. *Parallel and distributed computation: numerical methods*, volume 23. Prentice hall Englewood Cliffs, NJ, 1989.

[9] L. Carlone, A. Censi, and F. Dellaert. Selecting good measurements via l1 relaxation: A convex approach for robust estimation over graphs. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2667–2674, 2014.

[10] L. Carlone, M. Kaouk Ng, J. Du, B. Bona, and M. Indri. Rao-Blackwellized particle filters multi robot SLAM with unknown initial correspondences and limited communication. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 243–249, 2010.

[11] Stefano Carpin. Fast and accurate map merging for multi-robot systems. *Autonomous Robots*, 25(3):305–316, 2008.

[12] L. C. Carrillo-Arce, E. D. Nerurkar, J. L. Gordillo, and S. I. Roumeliotis. Decentralized multi-robot cooperative localization using covariance intersection. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 1412–1417, 2013.

[13] M. Cummins and P. Newman. FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance. *Intl. J. of Robotics Research*, 27(6):647–665, June 2008.

[14] A. Cunningham, V. Indelman, and F. Dellaert. DDF-SAM 2.0: Consistent distributed smoothing and mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Karlsruhe, Germany, May 2013.

[15] 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.

[16] 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)*, St. Paul, MN, 2012.

[17] F. Dellaert and M. Kaess. Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *Intl. J. of Robotics Research*, 25(12):1181–1203, Dec 2006.

[18] Joseph Djugash and Sanjiv Singh. Motion-aided network slam with range. *Intl. J. of Robotics Research*, 31(5):603 – 625, April 2012.

[19] H. Durrant-Whyte and M. Stevens. Data fusion in decentralized sensing networks. In *4th Intl. Conf. on Information Fusion*, 2001.

[20] R.M. Eustice, H. Singh, and J.J. Leonard. Exactly sparse delayed-state filters for view-based SLAM. *IEEE Trans. Robotics*, 22(6):1100–1114, Dec 2006.

[21] Maurice F Fallon, Georgios Papadopoulos, and John J Leonard. A measurement distribution framework for cooperative navigation using multiple auvs. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 4256–4263, 2010.

[22] M. Fischler and R. Bolles. Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography. *Commun. ACM*, 24:381–395, 1981.

[23] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3):325–344, 2000.

[24] S. Grime and H.F. Durrant-Whyte. Data fusion in decentralized sensor networks. *Control Engineering Practice,*, 2(1):849–863, 1994.

[25] A. Howard. Multi-robot simultaneous localization and mapping using particle filters. *Intl. J. of Robotics Research*, 25(12):1243–1256, 2006.

[26] A. Howard, G. S. Sukhatme, and M. J. Matarić. Multi-robot mapping using manifold representations. *Proceedings of the IEEE - Special Issue on Multi-robot Systems*, 94(9):1360 – 1369, Jul 2006.

[27] V. Indelman, P. Gurfil, E. Rivlin, and H. Rotstein. Distributed vision-aided cooperative localization and navigation based on three-view geometry. *Robotics and Autonomous Systems*, 60(6):822–840, June 2012.

[28] V. Indelman, P. Gurfil, E. Rivlin, and H. Rotstein. Graph-based distributed cooperative navigation for a general multi-robot measurement model. *Intl. J. of Robotics Research*, 31(9), August 2012.

[29] V. Indelman, N. Michael, and F. Dellaert. Incremental distributed robust inference from arbitrary robot poses via em and model selection. In *RSS Workshop on Distributed Control and Estimation for Robotic Vehicle Networks*, July 2014.

[30] V. Indelman, E. Nelson, N. Michael, and F. Dellaert. Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.

[31] S.J. Julier and J.K. Uhlmann. A non-divergent estimation algorithm in the presence of unknown correlations. In *American Control Conference*, pages 2369–73, 1997.

[32] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller. Multiple relative pose graphs for robust cooperative mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3185–3192, Anchorage, Alaska, May 2010.

[33] Kurt Konolige, Dieter Fox, Benson Limketkai, Jonathan Ko, and Benjamin Stewart. Map merging for distributed robot navigation. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 212–217, 2003.

[34] R. Kurazume, S. Nagata, and S. Hirose. Cooperative positioning with multiple robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, volume 2, pages 1250–1257, 1994.

[35] Y. Latif, C. D. C. Lerma, and J. Neira. Robust loop closing over time. In *Robotics: Science and Systems (RSS)*, 2012.

[36] Gim Hee Lee, Friedrich Fraundorfer, and Marc Pollefeys. Robust pose-graph loop-closures with expectation-maximization. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 556–563, 2013.

[37] Jixin Li, Ehsan Elhamifar, I-Jeng Wang, and René Vidal. Consensus with robustness to outliers via distributed optimization. In *IEEE Conference on Decision and Control*, pages 2111–2117, 2010.

[38] M. Mesbahi and M. Egerstedt. *Graph Theoretic Methods for Multiagent Networks*. Princeton University Press, 2010.

[39] E. Montijano, S. Martinez, and C. Sagues. Distributed robust data fusion based on dynamic voting. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 5893–5898. IEEE, 2011.

[40] Eduardo Montijano, Sonia Martínez, and Carlos Sagues. Distributed robust consensus using ransac and dynamic opinions. *IEEE Transactions on Control Systems Technology*, 23(1), January 2015.

[41] T. K. Moon. The expectation-maximization algorithm. *Signal processing magazine, IEEE*, 13(6):47–60, 1996.

[42] Erik Nelson, Vadim Indelman, Nathan Michael, and Frank Dellaert. An experimental study of robust distributed multi-robot data association from arbitrary poses. In *Intl. Sym. on Experimental Robotics (ISER)*, 2014.

[43] E.D. Nerurkar, S.I. Roumeliotis, and A. Martinelli. Distributed maximum a posteriori estimation for multi-robot cooperative localization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1402–1409, May 2009.

[44] Eric W Nettleton and Hugh F Durrant-Whyte. Delayed and asequent data in decentralized sensing networks. In *Intelligent Systems and Advanced Manufacturing*, pages 1–9, 2001.

[45] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

[46] Avi Pfeffer and Terry Tai. Asynchronous dynamic bayesian networks. 2012.

[47] Richard J Radke. A survey of distributed computer vision algorithms. In *Handbook of Ambient Intelligence and Smart Environments*, pages 35–55. Springer, 2010.

[48] S.I. Roumeliotis and G.A. Bekey. Distributed multi-robot localization. *IEEE Trans. Robot. Automat.*, August 2002.

[49] D. Smith and S. Singh. Approaches to multisensor data fusion in target tracking: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 18(12):1696, Dec 2006.

[50] N. Sunderhauf and P. Protzel. Towards a robust back-end for pose graph slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1254–1261. IEEE, 2012.

[51] Niko Sünderhauf and Peter Protzel. Switchable constraints vs. max-mixture models vs. rrr–a comparison of three approaches to robust pose graph slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013.

[52] J. M. Walls and R. M. Eustice. An exact decentralized cooperative navigation algorithm for acoustically networked underwater vehicles with robustness to faulty communication: Theory and experiment. 2013.

[53] Jeffrey M Walls, Alexander G Cunningham, and Ryan M Eustice. Cooperative localization by factor composition over a faulty low-bandwidth communication channel. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2015.

[54] Sarah E Webster, Jeffrey M Walls, Louis L Whitcomb, and Ryan M Eustice. Decentralized extended information filter for single-beacon cooperative acoustic navigation: Theory and experiments. *Robotics, IEEE Transactions on*, 29(4):957–974, 2013.

[55] Lin Xiao, Stephen Boyd, and Sanjay Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on*, pages 63–70, 2005.

[56] Peng Yang, Randy A Freeman, and Kevin M Lynch. Distributed cooperative active sensing using consensus filters. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 405–410, 2007.

[57] Jonathan S Yedidia, William T Freeman, and Yair Weiss. Understanding belief propagation and its generalizations. *Exploring artificial intelligence in the new millennium*, 8:236–239, 2003.

[58] X. Zhou and S.I. Roumeliotis. Multi-robot SLAM with unknown initial correspondence: The robot rendezvous case. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 1785–1792. IEEE, 2006.