IEEE CSM Special Issue

Title: “Distributed Control and Estimation of Robotic Vehicle Networks”

Special Issue Editors: Nisar Ahmed (University of Colorado, Boulder), Jorge Cortes (UC San Diego), Sonia Martinez (UC San Diego)

Many applications, such as environmental monitoring, security and surveillance, scientific exploration, and intelligent transportation, share a fundamental need to accomplish multiple tasks across space and time that are beyond the capabilities of a single autonomous platform. The growing demand for autonomous multivehicle networks has stimulated broad interest in distributed control and estimation (DCE) strategies that support cooperative and coordinated vehicle autonomy. Ideally, distributed approaches would not only perform as well as centralized methods, but also lead to better scalability, naturally parallelized computation, and resilience to communication loss and hardware failures.

However, as evidenced by the ongoing work in both the robotics and control systems communities, many significant theoretical and technological hurdles must still be overcome to meet the demanding challenges facing autonomous multivehicle networks in real applications. Many key concepts for DCE have been known for some time, but experience has shown that successfully transitioning that theory to robust practical implementation can be quite difficult. Examples include hurdles in the implementation of distributed planning and computation, which can translate into slowly reactive platforms and inefficient practical executions. Another equally and sometimes dismissed important problem is that new advances in one area (for instance, computing power, communication) enable new algorithms that display previously unfeasible capabilities, which in turn need adequate integration with existing capabilities in other areas (for instance, sensing and perception). As new emerging technologies for sensing, actuation, computation, and communication continue to enable novel algorithms and platform designs, this becomes ever more challenging as problem domains as well as end-user expectations can be radically transformed. To keep pace with a fast-changing technological landscape, researchers must constantly assess to what extent standing problems have been solved by research and technological advances, which relevant problems have received insufficient attention, which assumptions need to be reassessed, and what new challenges and solutions to reach for.

The genesis of this special issue was a Robotics: Science and Systems Conference Workshop on the topic of distributed control and estimation for robotic vehicle networks held in Berkeley in July 2014. The high number of workshop participants led to a large number of contributions to the special issue. Consequently, this is the first of two issues that will be published in the IEEE Control Systems Magazine. This issue presents diverse perspectives on estimation, perception, planning, and control from specialists within the robotics and control communities.

A key motivation for organizing this special issue is the observation that specialized research efforts on DCE within the control and robotics communities have not been well coordinated with each other in recent years, which is unfortunate considering the deep-rooted connections of these communities. This has subsequently led to a plethora of theoretical frameworks and design techniques for DCE, leading to an overwhelming number of practical design possibilities and exciting new pathways for innovative research, but also increasing the risk of unnecessarily duplicating effort across different communities. There are indeed many strong similarities among favored approaches in both communities, and many of their differences can be traced either to the influence of operational requirements for different applications (for example, aerospace, maritime, and supply management) or inspiration by different disciplines (such as artificial intelligence and machine learning in robotics versus communications and information theory in controls).

This special issue represents one step in a larger effort to close the knowledge gaps between specialists in the robotics and control communities, and aims to help readers (researchers and practitioners alike) better appreciate the fundamental connections between distributed control, planning, perception, and decision-making for multivehicle networks.

Key technical questions

When closing the loop for multivehicle systems, it is usually convenient to assume that distributed control/planning and distributed estimation/perception problems can be treated separately. Indeed, state-of-the-art techniques for distributed planning (such as graph-based trajectory generation, and consensus-/graph-based task allocation) and perception (such as multi-robot graph-based simultaneous localization and mapping (SLAM), and Bayesian/consensus sensor fusion for cooperative tracking) can often lead to good results when used together. However, the assumed "separation principle" is heuristic and leaves open many questions regarding how off-the-shelf solutions for different parts of the same problem should be jointly selected or modified to work best together, and the guarantees (if any) on the optimal and robust behavior of the resulting design. Alternative integrated approaches have also emerged for multivehicle systems (for example, distributed optimization, model predictive control, and reinforcement learning), which formally capture and exploit subtle, yet important, dynamic linkages between the control and estimation problems. However, there are also many questions surrounding these approaches including how reasonable for general applications are the assumptions/approximations required for analytical and computational tractability and how to leverage state-of-the-art planning/perception methods for individual mobile robots.

This special issue's goal is to contribute to the following key technical challenges:

1. Coordination tasks leading to complex optimization problems need to be solved in a scalable manner by multiple agents and in a timely fashion. Proposed approaches must be robust to communication failures, lack of global knowledge (including mismatches across agents), and uncertainty. Such algorithmic solutions may intelligently exploit the rich and heterogeneous nature of multirobot systems.
2. Communication may play different roles in different problems, ranging from being an integral part of the solution to being unnecessary. In particular, tradeoffs between communication, computation, and memory need to be carefully established in order to obtain algorithmic solutions that minimally sacrifice the performance of their centralized counterparts. Depending on the specific context, different kind of errors in communication may be more tolerable than others. In this regard, not all applications are created equal, for instance, underwater robots must overcome different communication

challenges than those faced by ground or air robots. Also, what is an acceptable scalable solution in some settings may not be in others, such as in problems that involve variables that scale poorly with the network size. Strategies or technology that make communication both robust and scalable are particularly needed.

3. In tasks such as cooperative localization in GPS-denied environments, the ability of agents to share relevant partial information in a decentralized way is key. Agents need methods to account for changing network topologies, the presence of correlated, common information from the past in loopy networks, and the possibility that not all agents are receiving or observing useful data. Mismatches across information sets, such as different frames of reference or data association ambiguities, inevitably arise. Methods that enable distributed vehicles to efficiently reconcile these inconsistencies are necessary.
4. The combination of real-time distributed control and estimation with computationally intensive state-of-the-art perception systems, such as computer vision or dense lidar mapping, presents formidable challenges. This is especially true for small agile platforms like micro unmanned aerial vehicles (UAVs). The advent of ubiquitous computing technology such as smartphones presents untapped opportunities for integration into larger distributed perception networks and can be exploited for improved decision-making.

Broader impact: distributed multi-vehicle systems and society

Beyond the technical issues raised above, there are also many concerns about the human and socioeconomic impact autonomous multivehicle networks will have, once applications such as the ones mentioned above are established. Indeed, it is important to keep in mind that multivehicle systems are fundamentally a human tool, and that their deployment into “human environments” naturally raises questions about human-vehicle interaction, safety, and security. In turn, these issues point to yet another a strong connection between DCE research and other noteworthy research areas, such as information networks and cyber-physical systems.

For instance, the issues associated with introducing increased autonomy into air traffic management are typical of the concerns now being echoed in other multivehicle application domains such as autonomous highway management, multi-UAV/single operator team design, and persistent littoral surveillance by autonomous ship fleets. As an example, enabling human-vehicle interaction in multirobot teams is emerging as an important problem in need of further development to enhance performance and adaptability. Questions on this theme include how a single human operator, or a group of them, can manage autonomous vehicle networks, possibly composed by very large number of agents; what type of information needs to be made available to human operators to ensure safe and secure performance, and vice versa, what information the human operators should provide; how the roles of human operators may change once distributed autonomy takes over, and how autonomy can/should be implemented in distributed multi-vehicle systems. Our hope is that the contributions of this special issue pave the way towards further work in DCE that addresses these important questions.

Contributions of the special issue

The contributions of the special issue touch upon distributed estimation, perception, planning and control. The articles all share in common the emphasis on developing fundamental understanding and algorithms that allows agents with partial information to achieve global objectives, and the issue of information consistency across the network. The articles are particularly fitting to the theme of the special issue in that they do not focus solely on one aspect, but instead touch upon many of the challenges laid out above: the connection between distributed estimation, control, and decision-making; robustness in communication against noise, delays, message drops and under dynamic interaction topologies; scalability of the proposed strategies; and information sharing.

The article "Incremental Distributed Inference from Arbitrary Poses and Unknown Data Association" by Vadim Indelman, Erik Nelson, Jing Dong, Nathan Michael, and Frank Dellaert develops a distributed, online multirobot approach to enable a group of robots to concurrently establish a common reference frame and resolve multirobot data association on-the-fly. This research addresses the challenge 3 mentioned above. While both of these problems have been previously addressed assuming the other problem is solved, few attempts have been made to solve the two problems simultaneously. To enable cooperative inference about variables of interest (for instance, observed objects, tracked targets), it is essential to establish a common reference frame and environment map between the robots, so that they can correctly interpret the information transmitted by other team members. Solving the two coupled problems concurrently enables the robot team, scattered in a complex, initially unknown environment, to establish collaboration without requiring any prior knowledge. The article contributions are threefold: first, the development of an expectation-maximization approach to determine initial relative poses and multirobot data association within an incremental, online, distributed framework. Second, the development of a model-selection approach for selecting the most probable cluster among several possible candidates. This approach uses the Chinese Restaurant Process to cope with potential perceptual aliasing and to account for the possibility there is no correct solution given the information obtained thus far. The third contribution is an extensive evaluation in real-world multirobot experiments, both in indoor and outdoor environments.

In “Scalable Multi-Robot Localization and Mapping with MOARSLAM”, John Morrison, Dorian Galvan-Lopez, and Gabe Sibley provide insight into distributed SLAM with multiple coordinating devices and robots that do not have a priori knowledge of the environment structure or global “eyes in the sky” to help localize individual agents. After reviewing the fundamentals of multirobot SLAM, the authors describe the Multiple Operator Augmented Relative SLAM (MOARSLAM) approach and provide experimental demonstrations of the technique on augmented reality and outdoor mapping applications with real data sets. This work thus addresses the challenge 4 in what concerns real data assimilation. MOARSLAM enables a distributed network of robots or mobile devices equipped with monocular or stereo cameras and inertial measurement units to achieve large-scale cooperative localization and mapping through a client-server architecture, where clients may request place-matching queries, upload local map updates or download updated maps. This architecture allows the agents to process measurements locally and independently using state-of-the-art robotic SLAM techniques. Agents can therefore operate asynchronously even in the event of network disconnections, while still providing the benefits of a fully coordinated multi-robot SLAM system that resolves challenging issues such as alignment of overlapping multi-agent maps, loop-closure detection, and memory management for long-term operation. Techniques such as MOARSLAM represent the first step on the road toward more general long-term SLAM distributed methods that address key scalability issues, which at their core are all related to realization of, and recovery from, estimation inconsistencies in large-scale mobile networks.

In contrast to the partially centralized mapping and localization strategy taken by MOARSLAM, the "Cooperative Localization for Mobile Agents" article by Solmaz Kia, Stephen Rounds, and Sonia Martinez focuses on a fully decentralized view of the multi-agent localization problem. In the cooperative localization paradigm presented here, agents use relative measurements to each another as a source of position information in unknown environments, rather than relying solely on environmental features for map building. To achieve optimum localization performance across all agents in the network in the context of recursive Kalman filtering, cooperative localization strategies theoretically require each agent in the network to track every other agent in the network as if they were a set of "moving landmarks". Such joint tracking perfectly accounts for cross correlations between different agent states that are introduced by relative measurements between platforms, but is computationally intractable and infeasible for even modestly sized networks. The authors show how a scalable synchronized interim master message-passing approach can overcome this limitation, by enabling each agent to update only the relevant correlation statistics, and thus drastically improve communication and computation requirements for practical cooperative localization in mobile robotic networks. The research of this paper is aligned with the challenge 2 described above.

Switching focus from distributed estimation to distributed control in robotic vehicle networks, "Distributed Optimal Control of Multiscale Dynamical Systems: A Tutorial," by Silvia Ferrari, Greg Foderaro, Pingping Zhu, and Thomas Wettergren considers the problem of optimal control for multiscale dynamical systems of interacting agents, in which optimal state and control laws must be determined for each individual agent, but the cost function depends on the macroscopic state of all agents over time. This paper addresses research issues that pertain to challenge 1 above. The authors present a class of distributed optimal control (DOC) methods to render a tractable solution to this challenging problem. In particular, a technique based on restriction operators, such as time-varying density functions, is presented to reduce the computation required by optimizing system dynamics over large temporal and/or spatial scales. The authors discuss direct DOC implementation methods based on discretization, which can be shown to not compromise closed-loop agent performance, as well as indirect implementation techniques, which relax certain assumptions on the restriction operator for improved computational performance and robustness in the presence of constraints. The broad applicability of the DOC approach is shown through several examples, including a multi-agent formation and path-planning application, an image reconstruction problem, and a robust-control application for collaborative networks.

Overall, with this special issue and the selected papers, we have aimed to not only highlight recent progress on distributed control and estimation for multivehicle networks, but also bring together different perspectives from the controls and robotics communities that can serve to promote a much-needed integrative research agenda.