IEEE CSM Special Issue - part II
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)

This edition of IEEE CSM presents the second part of the special issue on “Distributed Control and Estimation of Robotic Vehicle Networks”. The topic of distributed control and estimation (DCE) is of increasing relevance to robotic vehicle systems that share a fundamental need to accomplish multiple tasks across space and time that are beyond the capabilities of a single platform. The growing demand for autonomous mobile multi-robot networks has stimulated broad interest in DCE strategies that naturally support cooperative, collaborative, and coordinated vehicle autonomy. When designed and implemented properly, distributed approaches 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.

The special issues arose with the intent of bringing together recent advances made by the robotics and controls research communities towards realizing the full potential of DCE methods. The goal was to identify fundamental questions and hurdles on the intertwined problems of estimation, perception, planning, and control that still remain and must be addressed to facilitate more widespread deployment of multi-vehicle technologies in practice.

Key technical challenges

We provide here for completeness a summary description of the key technical challenges on DCE that motivated the organization of both special issues. This description provides a conceptual framework to understand how the various contributed articles fit together and address important hurdles that must be overcome to transition multivehicle networks into practical widespread use. The special issues highlight contributions to the following key technical challenges:

1) Coordination tasks leading to complex optimization problems need to be solved in a scalable manner. Basic understanding and guidelines are necessary to determine what type of coordination approach (centralized, hierarchical, distributed, or some combination thereof) is more suitable for specific classes of tasks. Whatever the approach employed, algorithmic solutions must be robust to communication failures, lack of global knowledge (including mismatches across agents), and uncertainty. In addition to verification and validation issues, researchers must still grapple with better understanding of the interplay between model representations/abstractions and real world physical constraints, and understanding how to beat the curse of dimensionality to achieve practical scalability for multi-agent systems.

2) Communication may play different roles depending on the problem scenario, ranging from being an integral part of the solution to being unnecessary. Trade-offs between communication, computation, and memory need to be carefully weighed to obtain algorithmic solutions that minimally sacrifice the performance of their centralized counterparts and efficiently use the

resources available to the agents. The accurate characterization of how information accuracy affects the level of achievement of the task at hand is particularly important in this regard. The development of realistic wireless communication and signal models is also critical. As robotic vehicle networks grow more sophisticated, they will increasingly rely on signals of opportunity in both outdoor and indoor settings (for instance, GPS, WiFi, cell networks, or RF). Hence, the challenge of identifying effective models of complex signal propagation patterns must be addressed for mobile robotic networks to operate in cluttered and contested environments.

3) In cooperative tasks such as localization or target tracking, the ability of agents to share relevant (possibly partial) information in a decentralized way is key. Agents need methods to account for changing network topologies, the fact that information from different agents might be correlated in a way that the agents do not know exactly about, and the possibility that not all agents are receiving or observing useful data. Methods that enable distributed vehicles to efficiently reconcile inconsistencies across information sets and deal with reliable estimates across the network are necessary.

4) The combination of real-time DCE with computationally intensive state-of-the-art perception and learning presents formidable challenges. This issue becomes more relevant with decreasing robot size (for instance, small agile platforms). Recent years have seen amazing gains in the capability of artificial intelligence and machine learning algorithms, though there is still a long way to go for these to be fully translatable and reliable “off-the-shelf” components in networks of autonomous robots. The advent of ubiquitous computing technology such as smartphones presents untapped opportunities for integration into larger distributed networks and can be exploited for improved decision making.

The introduction of autonomous robotic vehicle networks in human environments (whether on the road, in the air, in the home or workplace, or elsewhere) naturally raises concerns about human-vehicle interaction and safety/security. Even though the contributions of the special issue do not deal directly with these topics, we believe it is important to emphasize the importance of studying the human and socioeconomic impact autonomous multi-vehicle networks will have.

Contributions of the special issue

The contributions of the special issue squarely address the four technical challenges outlined above by focusing on the interplay between coordination and communication, and the issue of common information across multiple agents. Regarding the first aspect, emphasis is put on understanding to what extent centralized coordination, distributed coordination, or a hybrid combination of both is adequate to the task at hand in multi-agent robotic networks. Regarding the second aspect, special care is taken in ensuring the consistency of the estimates across the network given that the correlation of the information being fused is unknown by the individual agents. A distinctive feature of the contributions is that none of them is restricted to any particular challenge, but instead they all touch upon several of them, as we describe below.

The article “Distributed Data Fusion: Neighbors, Rumors, and the Art of Collective Knowledge” by Mark E. Campbell and Nisar Ahmed discusses solutions to the distributed data fusion problem from a Bayesian perspective. This research addresses the challenges 2 and 3 mentioned earlier. In distributed data fusion problems, a group of agents sense their local environment, communicate with other agents as needed, and collectively try to infer knowledge about a particular process. The focus of the article is on developing Bayesian distributed data fusion solutions which are either equivalent to the centralized solution, or as close as possible to the centralized solution. A centralized estimate of the process uses all of the sensor data to create an estimate and has the benefits of being the ‘best’ solution given the information available. However, the centralized solution has its own challenges, including single point failures and large amounts of data being communicated. These challenges can be easily overcome by distributed approaches, which instead face the challenge of creating accurate estimates across multiple agents. Given the complexities of the communication network, the sensors, or the process itself, identical estimates cannot always be obtained. The article argues for Bayesian distributed data fusion techniques as a highly practical and robust framework for tackling distributed estimation problems. The technical developments are illustrated in the example of kids playing hide and seek, for a variety of increasingly complex variants, starting from a baseline case of linear models and Gaussian distributions, and adding various forms of complexity, including nonlinear dynamic models, sensors with dynamics, non-Gaussian distributions, network delays, and network topology. The article concludes by discussing open research questions on the complexity of the representations, connections between distributed planning and estimation methods, and various applications of interest.

In “A Distributed Optimization Framework for Localization and Formation Control with Applications to Vision-Based Measurements”, Roberto Tron, Justin Thomas, Giuseppe Loianno, Kostas Daniilidis, and Vijay Kumar introduce the reader to a set of theoretical and practical tools that have proved useful in obtaining distributed algorithmic solutions to multi-agent mutual localization and formation control problems. The mutual localization problem for a multi-agent system consists of estimating the pose of each agent with respect to a common reference frame. The multi-agent formation control problem consists of maneuvering the agents to achieve a specified set of relative positions or directions. Both problems involve a geometric aspect, given by the geometry of the poses (rotation and/or translation) of the agents, and a graph-theoretic aspect, where vertices in a graph represent agents and edges are associated to measurements or other pairwise quantities. This research fits into challenge 1 and also touches aspects of challenge 4. The article focuses on vision-based settings where bearing (i.e., relative direction without distance) measurements have special importance. Obtaining such measurements on aerial vehicles from vision and Inertial Measurements Units (IMU) sensors requires strategies for identifying and tracking neighboring agents in the images obtained from the onboard camera. The discussion shows how the notions of shape decomposition and rigidity characterize the well-posedness of both the mutual localization and formation control problems, and help encode the desired solutions as global minimizers of network-wide objective functions. The article also describes the design of distributed algorithms to minimize these network-wide costs and how to obtain estimates of their region of attraction that, in some cases, lead to global convergence results under the assumption of ideal measurements.

The article “Tutorial on the Role of Information Assumptions in Decentralized Task Allocation” by Luke B. Johnson, Han-Lim Choi and Jonathan P. How explores the benefits and costs associated with coordinating multi-agent teams. The paper establishes a distinction between *centralized* task allocation, where a single control center receives all relevant information necessary to define the motion and actions of every agent in the team, and *decentralized* task allocation, where individual agents make their own decisions about motion and actuation with the use of local communication and *a priori* information about the planning environment. The authors argue that the use of one type of coordination versus the other is a design choice and should depend on the specific task and communication environment. After providing example tasks to validate this point, the article focuses on understanding two different approaches to decentralized task allocation and how choices about the approach affect the overall mission performance. In the first algorithmic approach, termed implicit coordination, consensus among the agents is achieved on the agents’ local information spaces. The second algorithmic approach, termed plan consensus, relies on achieving consensus in the assignment domain. The key differences between these approaches is the type of information that must be communicated between the agents. This leads to algorithms that perform well in environments with vastly different information assumptions. The final section of the article introduces hybrid approaches that use consensus in both the information domain and the assignment domain simultaneously. The research of this article is aligned with challenges 1 and 2.

The last article of the issue also deals with research issues that pertain to challenges 1 and 2. The article “Safe Coordinated Maneuvering of Teams of Multi-rotor UAVs” by Venanzio Cichella, Ronald Choe, Syed Bilal Mehdi, Enric Xargay, Naira Hovakimyan, Isaac Kaminer, Vladimir Dobrokhodov, Antonio Manuel Pascoal, and A. Pedro Aguiar addresses the problem of safely coordinating a fleet of multi-rotor UAVs over a time-varying communication network and in the presence of static and dynamic obstacles. Their capability for hovering still in the air of multirotor UAVs, together with their small size, low cost, and high agility make this type of aircraft especially interesting as a research and development platform. The aim of the article is to present an integrated framework to solve the cooperative mission planning, coordinated motion control, and collision avoidance problems. The discussion assumes that the UAVs have access to their states and can perceive external objects within a given detection range. The proposed approach is a *hybrid* strategy combining centralized and decentralized coordination. A central unit is responsible for the mission planning, generating a solution to the trajectory-generation problem free of conflicts, satisfying the dynamic constraints of the vehicles, and the mission specific temporal constraints. This solution is communicated to the vehicles before the beginning of the mission. Decentralized controllers embedded onboard the UAVs ensure that the vehicles meet the desired temporal assignments of the mission, while flying along the trajectories, even in the presence of faulty communication networks, temporary link losses, and switching topologies. The hybrid controller is more robust than purely decentralized solutions, as the centralized planning phase is not affected by faults in the communication network. At the same time, the mission execution is safer with respect to purely centralized solutions, since decentralized control allows each vehicle to react in a timely fashion to other vehicles’ failures and potentially hazardous maneuvers without having to communicate with a central station. The discussion pays particular attention to the crucial role played by the quality of the service provided by the supporting communication network during mission execution and to safety, understood as the capability of each UAV to be able to avoid unpredicted static and dynamic obstacles in the airspace, and at the same time guarantee a minimum separation distance to the other vehicles involved in the cooperative mission.

When viewed in conjunction, this second special issue together with the first one published in April provide a fresh look on distributed control and estimation strategies of robotic vehicle networks and make substantial inroads into several of the challenges currently faced by robotics and control researchers.