Networked Robotics

Sarah Tang[1] and Vijay Kumar[1]

1 Definition

Networked robotics studies teams of robots that utilize a communication network to coordinate with each other, sensors, computers, or humans to accomplish complex goals. Robots can be terrestrial, aerial, or underwater and can communicate implicitly — detecting each other using sensors, such as cameras or LIDAR — or explicitly — sending messages in the form of light, sound, or radio signals. Research in this area aims to enable teams of robots to self-organize to complete complex tasks. The availability of multiple robots allows for greater efficiency and redundancy such that tasks can still be completed even if some robots fail. The communication network allows robots to leverage data collected by other agents, for example, sensor data about a remote location or feedback data from a previous attempt of the same task, to adapt their own actions. These capabilities give networked robots the potential to impact many industries, including manufacturing, construction, field robotics, environmental monitoring, and entertainment. The remainder of this article will discuss algorithms for autonomy or networked autonomy, present applications in industry, and propose avenues of future research for networked robots.

2 Research Challenges

The key capability of networked robotic systems is their ability to accomplish tasks or achieve effectiveness beyond the capabilities of a single agent. The most straightforward tasks involve *coordination*, where robots in the network simultaneously perform component sub-tasks in parallel, and the existence of multiple vehicles allows for tasks to be completed with greater efficiency. Robots are typically identical

[1]GRASP Lab, University of Pennsylvania, 3330 Walnut Street, Philadelphia, PA 19104 e-mail: {sytang, kumar}@seas.upenn.edu

This is a post-peer-review, pre-copyedit version of an article published in Encyclopedia of Robotics. The final authenticated version is available online at: http://dx.doi.org/10.1007/978-3-642-41610-1_21-1.

and sub-tasks have structured variations, such as different locations for package delivery or different steps of a manufacturing assembly line. In scenarios like mapping, robots can receive and utilize information about areas they have never visited. New robots can also be deployed as substitutes for those that are reaching the end of their battery lives, allowing for longer-duration tasks, such as persistent monitoring.

For more complex tasks, robots must exhibit a more sophisticated level of *cooperation*. In these scenarios, agents collectively complete tasks that would be completely infeasible for a single agent, such as forming an image for a visual display. In more constraining scenarios, robots could be directly coupled with each other, as in the case of cooperative manipulation.

Broadly speaking, there are two types of network architectures. In *centralized* networks, a significant portion of data aggregation and planning is carried out by a “base station”, which can communicate with all robots. Most often, however, a *decentralized* network, where each agent determines its own actions based on local information, is more desirable. Though planning becomes more challenging, as global information is no longer available, the network becomes more resilient, as it no longer has a single point-of-failure, and can span larger spatial distances, as robots are no longer constrained to the base station’s communication range. This paradigm reflects phenomena found in nature, where animals, such as termites, ants, bees, or wasps, can complete sophisticated tasks (for example, lifting objects weighing many times their individual body weights or constructing large, expansive habitats) by communicating only with their nearest neighbors.

Central to the concept of networked robotics is the Perception - Action - Communication (PAC) loop, which allows robots to use their sensor inputs and communications with their neighbors to inform future actions. A particular challenge is constructing individual PAC loops that give rise to collective progress towards a specified task. As robots take actions, their relationships with their neighbors and environment change and, consequently, the information they receive also changes. These dynamics are difficult to characterize, yet are essential to understanding the behavior of networked systems. Section 3 will provide more details.

Another important aspect is the interface for command and control for Human-Swarm Interaction (HSI), commonly used in robotic testbeds for research or education. Here, a human operator typically enters team-level commands to a central processor, which then broadcasts appropriate commands to each robot. The base station might also compile information about robots’ statuses and relay feedback to the operator. We discuss this in more detail in Section 4.

Recent years have further seen the development of decentralized, heterogenous multi-robot teams composed of agents that are diverse in type, mobility modes (i.e. underwater vs. land vs. aerial robots), and communication and sensing capabilities. The ultimate vision is one of inter-robot *collaboration*, where utilization of agents’ complementary capabilities allows for completion of new types of tasks, such as mapping environments that span multiple terrains. Heterogenous teams pose novel research challenges, like formulating abstractions and metrics of diversity and fusing data from disparate sensors. These will be discussed in Section 5.

Networked systems can also often benefit from learning algorithms that allow robots to refine their models of their environment and tasks or their control policies as new data is received; these types of learning and adaptation algorithms will be discussed in Section 6. As robotic networks grow in size, spatial distribution, and diversity, new challenges have opened up in designing robust and resilient networks, which will be discussed in Section 7. Finally, Section 8 will discuss successful experimental realizations of networked robots and Section 9 will conclude with a discussion of future research frontiers.

3 Perception-Action-Communication

A fundamental component of an autonomous networked robot is its *Perception-Action-Communication (PAC)* loop, pictured in Fig. 1. In this paradigm, each vehicle can have discrete, continuous, or hybrid representations of their environment at different spatial-temporal resolutions. For example, robots' dynamics are often modeled as continuous differential equations, while topology of the communication network often take the form of a graph or adjacency matrix. Within each agent, these models are built and refined by its perception module, which informs its planner and controller of the next best action to take. These actions are subsequently executed and communicated to (or sensed by) nearby agents. The robot in turn incorporates new information about its neighbors' actions into the next iteration of the PAC loop.

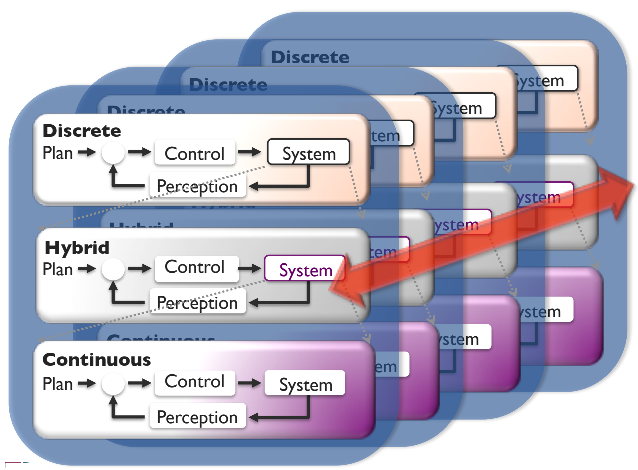

Fig. 1: A hierarchy of PAC loops in networked robotic systems.

What differentiates networked robots from static sensor networks is their ability to move and respond on-the-fly to new information. As a result, many research efforts revolve around designing modules of PAC loops and analyzing them in the presence of a dynamic sensing and communication network topologies. The simplest mathematical model to represent these dynamics is described as follows.

Let the network of interest contain N robots. The variables representing the parameters of each robot will be denoted with an index $i \in [1,N]$. Each robot has an associated state, $\mathbf{x}_i \in X_i \subset \mathbb{R}^n$, and control input, $\mathbf{u}_i \in U_i \subset \mathbb{R}^m$. A function $f_i : X_i \times U_i \to TX_i$ specifies a dynamic model for each robot:

$$\dot{\mathbf{x}}_i = f_i(\mathbf{x}_i, \mathbf{u}_i). \tag{1}$$

The communication network, that is, who can talk to whom, is typically modeled as a *communication graph*:

$$\mathscr{G}(t) = (\mathscr{V}(t), \mathscr{E}(t)), \tag{2}$$

where $\mathscr{V}(t)$ is a set of N vertices with locations given by the set $\{\mathbf{x}_i(t) \mid i \in [1,N]\}$. Robots have limited communication ranges, and an edge (i,j) is part of the set $\mathscr{E}(t)$ if robots i and j can communicate with each other. Note that because robots are mobile, $\mathscr{G}(t)$ is a time-varying graph.

A commonly used representation of $\mathscr{G}(t)$ is the *adjacency matrix* $\mathscr{A}^c(t)$, with elements defined as:

$$\mathscr{A}^c_{ij}(t) = \begin{cases} 1 & \text{if } (i,j) \in \mathscr{E}(t), \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

It is often more informative to represent the communication network with a positive-semidefinite weighing function, $w(\mathbf{x}_i(t), \mathbf{x}_j(t)) : X_i \times X_j \to \mathbb{R}$, that models not only the ability of two robots to send signals to each other, but also the strength of that signal. The set of edges can then be defined as:

$$\mathscr{E}(t) = \{(i,j) \mid w(\mathbf{x}_i(t), \mathbf{x}_j(t)) > 0\}. \tag{4}$$

Elements of the adjacency matrix can be defined as:

$$\mathscr{A}^c_{ij}(t) = w(\mathbf{x}_i(t), \mathbf{x}_j(t)). \tag{5}$$

An analogous *sensing graph*, $\mathscr{A}^s(t)$, can be defined to represent robots that can observe each other, but not necessary exchange information. A network's sensing and communication graphs are often different.

With sensing and communication inputs, robots can maintain estimates of their neighbors' states, modeled with the equation:

$$\hat{\mathbf{x}}^{(i)}_j(t) = h_i(\mathbf{x}_i(t), \mathbf{z}_{ij}(t)). \tag{6}$$

$\hat{\mathbf{x}}^{(i)}_j(t)$ represents the estimate of robot j maintained by robot i, which depends on robot i's current state and $\mathbf{z}_{ij}(t)$, its measurement of robot j's state. As can be seen from this model, robots' perception, control, and communication protocols are closely intertwined. Control algorithms must often simultaneously navigate robots towards a desired location and guarantee maintenance of communication graph properties, such as k-hop or algebraic connectivity. This can be accomplished with a range of approaches such as gradient-based or sampling-based methods; a survey of methods can be found in (Zavlanos et al, 2011). Stability analysis of these controllers have also been studied extensively; (Oh et al, 2015) presents a survey of these results.

Controllers rely on both information about other agents in the network, such as relative positions and orientations (Cornejo and Nagpal, 2015), and the communication or sensing networks themselves, such as the connectivity of the communication graph (Yang et al, 2010). However, these states typically must be estimated from noisy information. The interplay between the reliability of robots' communications and the quality of estimates is explicitly explored in (Schwager et al, 2011). Notably, for systems consisting of robots with first-order dynamics, the network's closed-loop dynamics are stable when using a straightforward flooding algorithm for estimation and proportional controller, even in the presence of arbitrary delays in the network update time. This relationship between stability, estimation, control, and communication, however, becomes more complicated for higher order vehicles. A survey of Kalman filtering methods for networked systems can be found in (Ribeiro et al, 2010). A further survey of control and estimation methods can be found in (Cao et al, 2013).

An equally important perception challenge is estimating the state of the environment. For instance, in data collection tasks, robots work together to map a scalar field quantity across a desired area. Each time a robot takes a measurement, the data is propagated through the communication network to inform future navigation. Robots navigate with strategies such as maximizing the information gain (Julian et al, 2012) or searching for peaks in the field (Atanasov et al, 2014). In persistent monitoring tasks, robots must arrange themselves to provide optimal coverage of a specified area; a survey of coverage methods is given by (Choset, 2001). In multi-robot Simultaneous Localization and Mapping (SLAM) tasks, robots collectively construct a map of a desired area. Bringing sensing data from multiple agents together to achieve consensus on a single map is a challenging task. While one option is to have all robots upload their local maps onto a central server, where the data is then fused, it is often more advantageous to achieve consensus on each robots' local map by propagating each robot's data through the network. Limited communication bandwidth again poses a major challenge, as robots often cannot communicate entire map portions to each other, and information-rich lower complexity representations, such as object-based models, must be found (Choudhary et al, 2017).

In the mapping scenarios previously described, the challenge of allocating target locations to specific robots in the team can be modeled as an *unlabeled multi-robot planning problem*. In this paradigm, all goal locations must be visited by some robot in the network, but it does not matter which robot visits which goal. Algorithms for this problem have been proposed in both obstacle-free (Turpin et al, 2014a) and cluttered environments (Turpin et al, 2014b). A related problem is the *labeled multi-robot planning problem*, where each robot is assigned a non-interchangeable goal location to visit. This situation arises, for example, in a delivery scenario, when robots are each carrying a unique package. Solutions for this problem have also been proposed for discrete (Luna and Bekris, 2011) and continuous (Yu and Rus, 2015; Tang and Kumar, 2018) spaces. In other scenarios, robots might need to maintain formations while navigating through an unknown environment. While controllers are responsible for maintaining a given formation, higher-level planners are needed to determine which formation to take and when (Alonso-Mora et al, 2016). The common challenge of these problems is again in achieving coordination at the global level using only local exchanges of information.

4 Human-Swarm Interaction

An important aspect of networked robotic systems is designing interfaces for *Human-Swarm Interaction (HSI)*. It is generally impractical to maintain a one-to-one correspondence between robots and operators, especially for large teams. As a result, research efforts have largely focused on designing scalable interfaces that allow a single operator to control tens to hundreds or even thousands of robots. One prominent design decision is how to derive commands for individual agents from a high-level operator command. For example, a chosen "leader" can follow the operator command exactly while the remaining agents follow in a set formation. An equally important challenge is communicating feedback to the operator. This is particularly crucial when the operator is spatially separated from the team and depends on this feedback to perceive the robots. Bandwidth limitations often preclude relaying the complete state of each vehicle while maintaining a meaningful update rate. Furthermore, the large amount of information becomes difficult to interpret quickly. As a result, different ways of summarizing the state of the team, such as robots' centroid. Feedback in the form of audio and light signals have also been demonstrated (McLurkin et al, 2006). Other research challenges include designing controllers for robots that guarantee safety, even in the presence of unsafe operator commands (Pickem et al, 2016) and accounting for errors in the robots' state information. A survey of these efforts can be found in (Kolling et al, 2016).

5 Heterogeneity and Diversity

While a large portion of networked robotics research has utilized identical agents, enabling true collaboration requires the utilization of a heterogenous team. Heterogeneity in the network also raises fundamental research questions in modeling and analysis. Given a team of diverse robots, how should their *capabilities*, that is, what sub-tasks they can complete, be abstracted? Similarly, how should complex tasks be modeled and how can we determine whether the collective capabilities of robots will be sufficient to complete them? In the presence of multiple different robots that can accomplish a given sub-task, which robot should be chosen to do so? Answering these questions requires the development of new mathematical models that capture the diversity of the team and the disparate characteristics of the agents. One possible model represents types of robots as *species* in a community with distinctive *traits*. This allows for the formulation of an optimization problem to distribute robots amongst sub-tasks (Prorok et al, 2015). It is further possible to define metrics for team diversity and quantitatively assess its relationship with team performance (Prorok et al, 2017).

In the planning community, the *k-color problem* is one model that abstracts planning and task allocation in a heterogenous team. Robots are partitioned into different groups or colors. Each goal must be visited by a robot of the correct color, but robots within a group are interchangeable. Traditional sampling-based methods for single-robot systems have been successfully extended to this setting (Solovey and Halperin, 2014). The diversity of heterogenous teams has commonly been leveraged for mapping tasks. In particular, air-ground robot teams have been shown to be capable of mapping multi-floor buildings (Charrow et al, 2015) and using their differing viewpoints to accurately localize targets (Chaimowicz et al, 2005).

6 Learning and Adaptation

The model-based methods for planning, estimation, and control described thus far can be further enhanced by allowing agents to learn and adapt to new information. This is especially effective in a multi-agent setting, as the communication network allows robots to also leverage information gained by their neighbors. However, this also present novel challenges in data fusion and analysis.

Learning algorithms allow robots to refine policies or parameters. For example, robots can use a distributed iterative learning algorithm to learn controller parameters (Hock and Schoellig, 2016) or policies without explicitly modeling the system dynamics (Levine and Abbeel, 2014). Similarly, a distributed reinforcement learning algorithm can be used to allow robots to update their control policies with feedback about the rewards gained from their own and their neighbors' actions (Varshavskaya et al, 2008). Deep reinforcement learning methods have also been extended to account for the distributed nature and limited sensing and communication capabilities of networked robots (Omidshafiei et al, 2017).

Adaptation algorithms further allow robots to alter their existing policies when their environments or networks change. For example, (Mather and Hsieh, 2011) presents a method for synthesizing the team's feedback control strategies while (Schwager et al, 2017) describes an adaptive coverage algorithm allows robots to continuously refine their estimate of the most promising subsequent position at which to collect data.

7 Resilience, Robustness, and Privacy

As networked robots operate in dynamic, unstructured environments, it is important for them to be *resilient*, that is, able to continue performing their tasks even in the presence of communication or robot failures. Resilience at the network level can be achieved by appropriate design of robots' PAC loops. For example, past research has proposed controllers that guarantee various measures of network integrity, such as a certain communication rate (Zavlanos et al, 2013) or existence of end-to-end communication between two nodes even when some intermediate communication links fail (Fink et al, 2012). Resilience can also be guaranteed at the individual level by equipping robots with redundant sensors and actuators, however, this comes at a cost of increased payload and computation demands.

A related consideration is network *robustness*, or the ability to reject unexpected external changes. From a controls perspective, individual robots should accurately execute their desired motions even in the presence of unmodeled disturbances. Like resilience, robustness is also a macro-scale consideration. For example, in consensus problems, proposed control policies guarantee robots will converge on the correct value despite the presence of malicious agents that deliberately broadcast misleading values to their neighbors (Saulnier et al, 2015). Finally, networks should ensure *privacy*, that is, outside agents should not be able to conclude information about agents in the network from observable actions. Unfortunately, this is a difficult task, as past work has shown that even observations of communication flow in the network, without any knowledge of the communications' content, can expose information. The problem of designing for privacy in heterogenous networks has been explored (Prorok and Kumar, 2016), however, much work still needs to be done towards achieving provably private networked systems.

8 Applications

Recent years has seen a dramatic increase in the size of networked robotic testbeds, a survey of which is given in (Jimnez-Gonzlez et al, 2013). Notably, (Pickem et al, 2016) presents the Robotarium, a 20-robot system designed to provide researchers anywhere in the world with a remote, robotic testbed. In addition to communication interfaces, the Robotarium also provides safe control protocols that prevent the execution of remote commands that would result in inter-robot collisions. (McLurkin et al, 2006) describes a 112-robot testbed, where robots provide human operators with audio and light feedback. The largest robotic network to date is a thousand-robot swarm (Rubenstein et al, 2014), which has demonstrated the ability to self-organize into complex shapes. These platforms have allowed for research experiments on an unprecedented scale.

Networked robotic systems have been successfully applied to coordination, cooperation, and collaboration tasks in both industry and academia. Amazon Robotics has successfully deployed a network of robots, pictured in Fig. 3, for autonomous

(a) Vehicles from Amazon Robotics used for warehouse management (Wurman et al, 2008).

(b) 49-Crazyflie testbed at the University of Southern California (Preiss et al, 2017).

Fig. 2: Examples of multi-agent systems.

warehouse inventory management. In the area of aerial robotics, quadrotors have emerged as a popular platform for first response (Mohta et al, 2014) and construction (Lindsey et al, 2011). As a result, a number of testbeds have emerged to facilitate research in coordinated quadrotor behaviors, most notably the 20-robot testbed at the University of Pennsylvania (Kushleyev et al, 2012) and the 49-robot testbed at the University of Southern California (Preiss et al, 2017), pictured in Fig. 2b. These vehicles have been used for applications such as cooperative manipulation (Sreenath and Kumar, 2013), pictured in 3b and target tracking (Hausman et al, 2015). There have also been a number of academic (Augugliaro et al, 2013) and industry (Intel, 2017) realizations of quadrotor teams for visual performances. Fig. 3a illustrates 300 Intel Shooting Star drones in choreographed flight.

Another type of cooperative robot team occurs in modular robots — vehicles that interconnect and rearrange themselves to form larger structures that subsequently act as a single mobile agent. Two prominent realizations of this concept are the 3D M-Block (Romanishin et al, 2015) and the SMORES modules (Jing et al, 2016), pictured in Fig. 3c. These systems can autonomously separate and rearrange into new configurations as necessary; this flexibility gives them extraordinary potential to be usable for a vast array of tasks.

Finally, systems of ground (Ramaithitima et al, 2016), aerial (Schwager et al, 2011), and underwater (Detweiler et al, 2014) robots have been used as sensor networks. These types of networks can be particularly useful for environmental monitoring, both for continuous, long-term data collection for periods of time that would

(a) 300 Intel Shooting Star drones performing choreography (Intel, 2017).

(b) A team of quadrotors performing cooperative manipulation (Sreenath and Kumar, 2013).

(c) The SMORES modular robot platform (Jing et al, 2016).

Fig. 3: Networked robots for tasks requiring cooperation.

be impractical for humans to sustain or for information gathering in environments, such as chemical spills, that are too hazardous for humans.

Significant milestones have also been reached in the deployment of heterogenous multi-robot networks. (Nikolaidis et al, 2015) present experiments quantifying the efficacy of human-robot cross-training. (Charrow et al, 2015) demonstrates a ground and aerial robot carrying out mapping tasks under a common planning paradigm. (Chaimowicz et al, 2005) demonstrates a team of cooperative ground and aerial robots for mapping tasks in urban environments.

9 Future Research Directions

While many exciting research results have been achieved in the field of networked robotics, many open problems still remain. Much of the existing literature has focused on analyzing systems with simplistic robot models with first- or second-order dynamics. Further work needs to be done in developing control, estimation, and communication algorithms for robots with more realistic and perhaps higher-order models. Current approaches to the synthesis and analysis of networks are typically designed as completely centralized (ie. the base station aggregates all data from all robots) or decentralized (ie. they broadcast to and receive information from all neighbors within range). Future networks could make deliberate decisions about whom to communicate with and when for better control over their resources. This requires research in building more advanced communication protocols and algorithms for reaching consensus across the individual models held by agents in the network.

Most methods thus far have also been model-based, and there is also vast potential for applying recent advances in machine learning, in particular deep learning, to robot teams. Rather than relying on hand-coded models of their environment and tasks, modules of PAC loops could be learned from sensed or communicated data and adapted on-the-fly depending on the task at hand. Research challenges towards this goal include finding efficient representations of data that can be transferred at low bandwidth and combining disparate types of information from different robots.

In heterogenous teams, information-rich abstractions can be used in planning and task allocation to allow robots to autonomously manage their computation and power resources and autonomously distribute sub-tasks to the most capable robot-types in a complex mission. Methods for fusing information from varying sensing and communication mechanisms can further allow robots to share data amongst each other. In the areas of robustness and resilience, robots could be equipped with more informative and nuanced metrics for information or sensor uncertainty and safeguards against privacy breaches. These current and future advances in autonomous, cooperative networked robotic systems poses them as powerful tools for socio-economic impact.

References

Alonso-Mora J., Montijano E., Schwager M., Rus D. (2016) Distributed multi-robot formation control among obstacles: A geometric and optimization approach with consensus. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 5356–5363

Atanasov N., Ny J. L., Pappas G. J. (2014) Distributed algorithms for stochastic source seekign with mobile robot networks. ASME Journal of Dynamic Systems, Measurement, and Control 137(3), DOI 10.1115/1.4027892

Augugliaro F., Schoellig A. P., D'Andrea R. (2013) Dance of the flying machines: Methods for designing and executing an aerial dance choreography. IEEE Robotics Automation Magazine 20(4):96–104, DOI 10.1109/MRA.2013.2275693

Cao Y., Yu W., Ren W., Chen G. (2013) An overview of recent progress in the study of distributed multi-agent coordination. IEEE Transactions on Industrial Informatics 9(1):427–438, DOI 10.1109/TII.2012.2219061

Chaimowicz L., Cowley A., Gomez-Ibanez D., Grocholsky B., Hsieh M. A., Hsu H., Keller J. F., Kumar V., Swaminathan R., Taylor C. J. (2005) Deploying Air-Ground Multi-Robot Teams in Urban Environments, pp. 223–234

Charrow B., Kahn G., Patil S., Liu S., Goldberg K. Y., Abbeel P., Michael N., Kumar V. (2015) Information-theoretic planning with trajectory optimization for dense 3d mapping. In: Robotics: Science and Systems

Choset H. (2001) Coverage for robotics – a survey of recent results. Annals of Mathematics and Artificial Intelligence 31(1):113–126, DOI 10.1023/A:1016639210559

Choudhary S., Carlone L., Nieto C., Rogers J., Christensen H. I., Dellaert F. (2017) Distributed mapping with privacy and communication constraints: Lightweight algorithms and object-based models. The International Journal of Robotics Research (IJRR)

Cornejo A., Nagpal R. (2015) Distributed range-based relative localization of robot swarms. In: Workshop on the Algorithmic Foundations of Robotics (WAFR), pp. 91–107

Detweiler C., Banerjee S., Doniec M., Jiang M., Peri F., Chen R. F., Rus D. (2014) Adaptive decentralized control of mobile underwater sensor networks and robots for modeling underwater phenomena. Journal of Sensor and Actuator Networks 3(2):113–149, DOI 10.3390/jsan3020113

Fink J., Ribeiro A., V.Kumar (2012) Robust control for mobility and wireless communication in cyber?physical systems with application to robot teams. Proceedings of the IEEE 100(1):164–178

Hausman K., Müller J., Hariharan A., Ayanian N., Sukhatme G. S. (2015) Cooperative multi-robot control for target tracking with onboard sensing. The International Journal of Robotics Research 34(13):1660–1677

Hock A., Schoellig A. P. (2016) Distributed iterative learning control for a team of quadrotors. IEEE Conference on Decision and Control (CDC)

Intel (2017) Drones at intel. URL https://newsroom.intel.com/press-kits/drones-at-intel/

Jimnez-Gonzlez A., de Dios J. R. M., Ollero A. (2013) Testbeds for ubiquitous robotics: A survey. Robotics and Autonomous Systems 61(12):1487–1501

Jing G., Tosun T., Yim M., Kress-Gazit H. (2016) An end-to-end system for accomplishing tasks with modular robots. In: Proceedings of Robotics: Science and Systems

Julian B. J., Angermann M., Schwager M., Rus D. (2012) Distributed robotic sensor networks: An information-theoretic approach. The International Journal of Robotics Research 31(10):1134–1154, DOI 10.1177/0278364912452675

Kolling A., Walker P., Chakraborty N., Sycara K., Lewis M. (2016) Human interaction with robot swarms: A survey. IEEE Transactions on Human-Machine Systems 46(1):9–26, DOI 10.1109/THMS.2015.2480801

Kushleyev A., Mellinger D., Kumar V. (2012) Towards a swarm of agile micro quadrotors. In: Robotics: Science and Systems

Levine S., Abbeel P. (2014) Learning neural network policies with guided policy search under unknown dynamics. In: Neural Information Processing Systems (NIPS)

Lindsey Q., Mellinger D., Kumar V. (2011) Construction of cubic structures with quadrotor teams. In: Robotics: Science and Systems (RSS), Los Angeles, CA, USA

Luna R., Bekris K. E. (2011) Push and swap: Fast cooperative path-finding with completeness guarantees. In: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI), pp. 294–300

Mather T., Hsieh M. A. (2011) Distributed robot ensemble control for deployment to multiple sites. In: Robotics: Science and Systems

McLurkin J., Smith J., Frankel J., Sotkowitz D., Blau D., Schmidt B. (2006) Speaking swarmish: Human-Robot interface design for large swarms of autonomous mobile robots. In: AAAI spring symposium, pp. 72–75

Mohta K., Turpin M., Kushleyev A., Mellinger D., Michael N., Kumar V. (2014) Quadcloud: A rapid response force with quadrotor teams. In: The 14th International Symposium on Experimental Robotics (ISER), Springer International Publishing, pp. 577–590

Nikolaidis S., Lasota P., Ramakrishnan R., Shah J. (2015) Improved human-robot team performance through cross-training, an approach inspired by human team training practices. The International Journal of Robotics Research (IJRR) 34(14):1711–1730

Oh K.-K., Park M.-C., Ahn H.-S. (2015) A survey of multi-agent formation control. Automatica 53:424–440, DOI https://doi.org/10.1016/j.automatica.2014.10.022

Omidshafiei S., Pazis J., Amato C., How J. P., Vian J. (2017) Deep decentralized multi-task multi-agent reinforcement learning under partial observability. URL https://arxiv.org/abs/1703.06182

Pickem D., Glotfelter P., Wang L., Mote M., Ames A., Feron E., Egerstedt M. (2016) The robotarium: A remotely accessible swarm robotics research testbed. https://arxiv.org/abs/1609.04730, accessed: 30 June 2017

Preiss J. A., Honig W., Sukhatme G. S., Ayanian N. (2017) Crazyswarm: A large nano-quadcopter swarm. In: IEEE International Conference on Robotics and Automation (ICRA)

Prorok A., Kumar V. (2016) A macroscopic privacy model for heterogeneous robot swarms. In: International Conference on Swarm Intelligence

Prorok A., Hsieh M. A., Kumar V. (2015) Fast redistribution of a swarm of heterogeneous robots. In: International Conference on Bio-inspired Information and Communications Technologies

Prorok A., Hsieh M. A., Kumar V. (2017) The impact of diversity on optimal control policies for heterogeneous robot swarms. IEEE Transactions on Robotics (T-RO) 33(2):346–358

Ramaithitima R., Whitzer M., Bhattacharya S., Kumar V. (2016) Automated creation of topological maps in unknown environments using a swarm of resource-constrained robots. IEEE Robotics and Automation Letters (RA-L) 1(2):746–753

Ribeiro A., Schizas I., Roumeliotis S., Giannakis G. (2010) Kalman filtering in wireless sensor networks – incorporating communication cost in state estimation problems. IEEE Control Systems Magazine 30(2):66–86

Romanishin J. W., Gilpin K., Claici S., Rus D. (2015) 3d m-blocks: Self-reconfiguring robots capable of locomotion via pivoting in three dimensions. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 1925–1932

Rubenstein M., Cornejo A., Nagpal R. (2014) Programmable self-assembly in a thousand-robot swarm. Science 345(6198):795–799, DOI 10.1126/science.1254295

Saulnier K., Saldana D., Prorok A., Pappas G. J., Kumar V. (2015) Resilient flocking for mobile robot teams. IEEE Robotics and Automation Letters (R-AL) 2(2):1039–1046, DOI 10.1109/LRA.2017.2655142

Schwager M., Michael N., Kumar V., Rus D. (2011) Time scales and stability in networked multi-robot systems. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 3855–3862

Schwager M., Vitus M. P., Rus D., Tomlin C. J. (2017) Robust adaptive coverage for robotic sensor networks. In: International Symposium of Robotics Research (ISRR), pp. 437–454

Solovey K., Halperin D. (2014) k-color multi-robot motion planning. The International Journal of Robotics Research (IJRR) 33(1):82–97

Sreenath K., Kumar V. (2013) Dynamics, control and planning for cooperative manipulation of payloads suspended by cables from multiple quadrotor robots. In: Robotics: Science and Systems (RSS)

Tang S., Kumar V. (2018) A complete algorithm for generating safe trajectories for multi-robot teams. In: Bicchi A., Burgard W. (eds) Robotics Research: Volume 2, Springer International Publishing, Cham, pp. 599–616

Turpin M., Michael N., Kumar V. (2014a) CAPT: Concurrent assignment and planning of trajectories for multiple robots. The International Journal of Robotics Research 33(1):98–112

Turpin M., Mohta K., Michael N., Kumar V. (2014b) Goal assignment and trajectory planning for large teams of interchangeable robots. Autonomous Robots 37(4):401–415

Varshavskaya P., Kaelbling L. P., Rus D. (2008) Efficient distributed reinforcement learning through agreement. In: International Symposium on Distributed Autonomous Robotic Systems (DARS)

Wurman P. R., D'Andrea R., Mountz M. (2008) Coordinating hundreds of cooperative, autonomous vehicles in warehouses. AI Magazine 29(1):9–20

Yang P., Freeman R. A., Gordon G. J., Lynch K. M., Srinivasa S. S., Sukthankar R. (2010) Decentralized estimation and control of graph connectivity for mobile sensor networks. Automatica 46(2):390–396

Yu J., Rus D. (2015) An effective algorithmic framework for near optimal multi-robot path planning. In: The International Symposium on Robotics Research (ISRR), Sestri Levante, Italy

Zavlanos M. M., Egerstedt M. B., Pappas G. J. (2011) Graph-theoretic connectivity control of mobile robot networks. Proceedings of the IEEE 99(9):1525–1540, DOI 10.1109/JPROC.2011.2157884

Zavlanos M. M., Ribeiro A., Pappas G. J. (2013) Network integrity in mobile robotic networks. IEEE Transactions on Automatic Control 58(1):3–18, DOI 10.1109/TAC.2012.2203215