41. Networked Robots

Vijay Kumar, Daniela Rus, Gaurav S. Sukhatme

This chapter discusses *networked robots*, multiple robots operating together coordinating and cooperating by *networked communication* to accomplish a specified task. This chapter presents an overview of the field with an emphasis on recent results and research challenges. Multiple robots enable new capabilities and the communication network enables new approaches and solutions that are difficult with just perception and control. Communication enables new control and perception capabilities in the system (e.g., access to information outside the perception range of the robot system). Conversely, control enables solutions for problems that are difficult without mobility (e.g., localization). Section 41.1 defines the field, examines the benefits of networking in robot coordination, and discusses applications. Section 41.2 highlights a few projects focused on networked robotics and discusses the application potential of the field. Section 41.3 discusses the research challenges at the intersection of control, communication, and perception. Section 41.4 defines a model for the control of a networked system which is used in Sects. 41.5–41.8 to examine specific research issues and opportunities facilitated by the interplay between communication, control, and perception.

41.1 Overview

The term *networked robots* refers to multiple robots operating together coordinating and cooperating by *networked communication* to accomplish a specified task. Communication between entities is fundamental to cooperation (and coordination), hence there is a central role for the communication network in networked robots. Networked robots may also involve coordination and cooperation with stationary sensors, embedded computers, and human users. The central feature of networked robots is the ability of the system to perform tasks that are well beyond the abilities of a single robot or multiple uncoordinated robots.

The *IEEE technical committee on networked robots* has adopted the following definition of a networked robot:

A networked robot is a robotic device connected to a communications network such as the Internet or local-area network (LAN). The network could be wired or wireless, and based on any of of a variety of protocols such as the transmission control protocol (TCP), the user datagram protocol (UDP), or 802.11. Many new applications are now being developed ranging from automation to exploration. There are two subclasses of networked robots:

1. Teleoperated, where human supervisors send commands and receive feedback via the network. Such systems support research, education, and public awareness by making valuable resources accessible to broad audiences.

2. Autonomous, where robots and sensors exchange data via the network. In such systems, the sensor network extends the effective sensing range of the robots, allowing them to communicate with each other over long distances to coordinate their activity. Sensing, actuation, and computation need no longer be collocated. A broad challenge is to develop a science base that couples communication, perception, and control to enable such new capabilities.

This definition of autonomous networked robots also includes a third class of distributed systems, mobile sensor networks, which is a natural evolution of sensor networks. Robot networks allow robots to measure spatially and temporally distributed phenomena more efficiently. The robots in turn can deploy, repair, and maintain the sensor network to increase its longevity, and utility. The focus of this chapter is *autonomous networked robots*.

Embedded computers and sensors are becoming ubiquitous in homes and factories, and increasingly wireless ad hoc networks or plug-and-play wired networks are becoming commonplace. Human users interact with embedded computers and sensors to perform tasks ranging from monitoring (e.g., supervising the operation of a factor and surveillance in a building) to control (e.g., running an assembly line consisting of sensors, actuators, and material-handling equipment). In most of these cases, the human users, embedded computers, and sensors are not collocated and the coordination and communication happens through a network. *Networked robots* extends this vision to multiple robots functioning in a wide range of environments performing tasks that require them to coordinate with other robots, cooperate with humans, and act on information derived from multiple sensors.

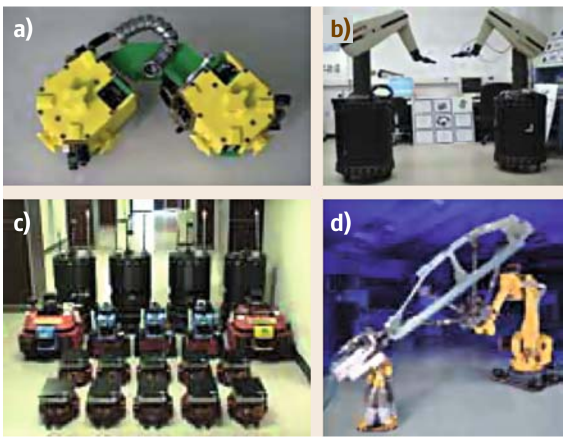

Fig. 41.1a–d Small modules [41.1] can automatically connect and communicate information to perform locomotion tasks **(a)**; robot arms [41.2] on mobile bases can cooperate to perform household chores **(b)**; swarms of robots [41.3] can be used to explore an unknown environment **(c)**; and industrial robots can cooperate in welding operations **(d)**

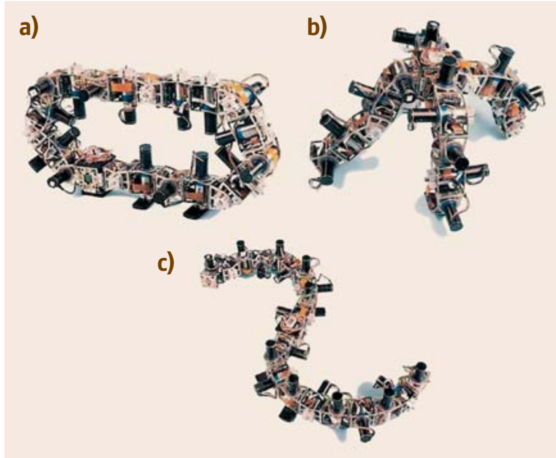

Fig. 41.2a–c Robotic modules [41.4] can be reconfigured to morph into different locomotion systems including a wheel-like rolling system **(a)**, a snake-like undulatory locomotion system **(b)**, a four-legged walking system **(c)**

Figure 41.1 shows prototype concepts derived from academic laboratories and industry. In all these examples, independent robot or robotic modules can cooperate to perform tasks that a single robot (or module) cannot perform. Robots can automatically couple to perform locomotion tasks (also see Fig. 41.2) and manipulation tasks that either a single robot cannot perform, or that would require a special-purpose larger robot to perform. They can also coordinate to perform search and reconnaissance tasks exploiting the efficiency that is inherent in parallelism. They can also perform independent tasks that need to be coordinated. Examples in the manufacturing industry include, for example, fixturing and welding.

Besides being able to perform tasks that individual robots cannot perform, networked robots also result in improved efficiency. Networking gives each robot access to information outside its perception range. Tasks such as searching or mapping can, in principle, be performed faster with an increase in the number of robots. A speed up in manufacturing operations can be achieved by deploying multiple robots performing operations in parallel but in a coordinated fashion.

Another advantage of using the network to connect robots is the ability to connect and harness physically removed assets. Mobile robots can react to information sensed by other mobile robots at a remote location. Industrial robots can adapt their end-effectors to new parts being manufactured upstream in the assembly line. Human users can use machines that are remotely located via the network.

The ability to network robots also enables fault tolerance in design. If robots can dynamically reconfigure themselves using the network, they are more tolerant to robot failures. This is seen in the Internet where multiple gateways, routers, and computers provide for a fault-tolerant system (although the Internet is not robust in other ways). Similarly, robots that can *plug* and *play* can be swapped in and out automatically to provide for a robust operating environment.

Finally, networked robots have the potential to provide great synergy by bringing together components with complementary benefits and making the whole greater than the sum of the parts.

Applications for networked robots abound. The US military routinely deploys unmanned vehicles that are reprogrammed remotely based on intelligence gathered by other unmanned vehicles, sometimes automatically. The deployment of satellites in space, often by astronauts in a shuttle with the shuttle robot arm requires the coordination of complex instrumentation onboard the space shuttle, human operators on a ground station, the shuttle arm, and a human user on the shuttle. Home appliances now contain sensors and are becoming networked. As domestic and personal robots become more commonplace, it is natural to see these robots working with sensors and appliances in the house while cooperating with one or more human users. Networked robots will likely be used as critical ingredients in the environmental observatories of the future. Large-scale ecological monitoring precludes the use of monolithic infrastructure, and is envisioned to be built as a distributed, networked robotic system.

41.2 State of the Art and Potential

The Network Robot Forum established in Japan in 2003 [41.5] estimates that the size of the networked robot industry will be over $200B by 2013, much larger than the industrial robot market for manufacturing applications. This growth is broad-based, across many industries. There is a strong connection between this industry and the industry connected to sensor networks. Sensor networks have been projected to grow dramatically in terms of commercialization and market value [41.6]. Robot networks are analogous to sensor networks except that they allow sensors to have mobility and allow the geographical distribution of the sensors to be adapted based on the information acquired.

A system of robots, embedded computers, actuators, and sensors has tremendous potential in civilian, defense, and manufacturing applications. Nature provides the proof of concept of what is possible [41.7]. Group behaviors in nature can be found in organisms that are only microns to those that are several meters in length. There are numerous examples of simple animals that execute simple behaviors with modest sensors and actuators but communicate with and sense nearest neighbors to enable complex emergent behaviors that are fundamental to navigation, foraging, hunting, constructing nests, survival, and eventually growth. As seen in Fig. 41.3, relatively small agents are able to manipulate objects that are significantly larger in terms of size and payload by cooperating with fairly simple individual behaviors. The coordination between agents is completely decentralized, allowing scaling up to large numbers of robots and large objects [41.8]. Individuals do not recognize each other. In other words, there is no labeling or identification of robots. The number

Fig. 41.3 Ants are able to cooperatively manipulate and transport objects often in large groups, without identified or labeled neighbors, and without centralized coordination

of agents in the team is not explicitly encoded. Agents are identical, enabling robustness to failures and modularity. There is minimal communication, and even that which is present is only between neighbors. Furthermore, the optimal mode of group coordination may be scale dependent. Studies of wasps show strong evidence of centralized coordination among species with small colony sizes, but a distributed, decentralized coordination in larger colonies [41.9]. All these attributes are relevant to networked robots.

Biology has shown how simple decentralized behaviors in unidentified individuals (e.g., insects and birds exhibiting swarming behaviors) can exhibit a wide array of seemingly intelligent group behaviors. Similarly networked robots can potentially communicate and cooperate with each other, and even though individual robots may not be sophisticated, it is possible for networked robots to provide a range of intelligent behaviors that are beyond the scope of intelligent robots.

The significance and potential impact of networked robots is apparent from the following examples.

The manufacturing industry has always relied on integration between sensors, actuators, material-handling equipment, and robots. Today companies are finding it easier to reconfigure existing infrastructure by networking new robots and sensors with existing robots via wireless networks. There is also an increasing trend toward robots interacting with each other in operations like welding and machining, and robots cooperating with humans in assembly and material-handling tasks. Workcells consist of multiple robots, numerous sensors and controllers, automated guided vehicles, and one or two human operators working in a supervisory role. However, in most of these cells, the networked robots operate in a structured environment with very little variation in configuration and/or operating conditions.

There is a growing emphasis on networking robots in applications of field robotics, for example, in the mining industry. Like the manufacturing industry, operating conditions are often unpleasant and the tasks are repetitive. However, these applications are less structured and human operators play a more important role.

In the health care industry, networks allow health care professionals to interact with their patients, other professionals, expensive diagnostic instruments, and in the future surgical robots. Telemedicine is expected to provide a major growth impetus for remote networked robotic devices that will take the place of today's standalone medical devices.

There are already many commercial products, notably in Japan, where robots can be programmed via and communicate with cellular phones. For example, the MARON robot developed by Fujitsu lets a human user dial up their robot and instruct it to conduct simple tasks including sending pictures back to the user via a cellular phone. Indeed these robots will interact with other sensors and actuators in the home – door openers equipped with Bluetooth cards and actuators and computer-controlled lighting, microwaves, and dishwashers. Indeed the Network Robot Forum [41.5] is already setting standards for how stationary sensors and actuators can interact with other robots in domestic and commercial settings.

Environmental monitoring is a key application for networked robots. By exploiting mobility and communication, robotic infrastructure enables observation and data collection at unprecedented scales in various aspects of ecological monitoring. This is significant for environmental regulatory policies (e.g., clean air and water legislation), as well as an enabler of new scientific discovery, for example, it is possible to obtain maps of salinity gradients in oceans, temperature and humidity variations in forests, the and chemical composition of air and water in different ecological systems [41.10]. In addition to mobile sensor networks, it is also possible to use robots to deploy sensors and to retrieve information from the sensors. Mobile platforms allow the same sensor to collect data from multiple locations while communication allows the coordinated control and aggregation of information. Examples include systems built for aquatic [41.11], terrestrial [41.12], and subsoil monitoring [41.13]. There are many efforts to developed networked underwater platforms [41.14–16]. Networks of static and robotic devices have been developed for aquatic monitoring [41.11] and to obtain high-resolution information on the spatial and temporal distributions of plankton assemblages and concomitant environmental parameters. The RiverNet project [41.17] at Rensselaer Polytechnic Institute (RPI) has focused on the development of robotic sensor networks for monitoring a river ecosystem. Recent work at University of California, Los Angeles (UCLA), University of Southern California (USC), University of California, Riverside, and University of California, Merced on the networked infomechanical system project [41.12] has focused on the development of robotic networks for monitoring the forest canopy, with a view to providing data for modeling canopy and undercover growth. Networked robotic mini-rhizotrons [41.13] are being deployed in the forest to monitor root growth in the soil.

In the defense industry, countries like the USA have invested heavily in the concept of networked, geographically distributed assets. Unmanned aerial vehicles like the Predators are operated remotely. Information from sensors on the Predators triggers the deployment of other vehicles and weapon systems at a different remote location and allows commanders in a third location to control and command all these assets. The US military is engaged in the large Future Combat Systems initiative to develop network-centric approaches to deploying autonomous vehicles. The network-centric tactical paradigms for modern warfare have created networked robots for defense and homeland security. While networked robots are already in operation, current approaches are limited to human users commanding a single vehicle or sensor system. However, it takes many human operators (between 2–10 depending on the complexity of the system) to deploy complex systems like unmanned aerial vehicles. A Predator unmanned aerial vehicle (UAV) is operated from a tactical control station, which may be on an aircraft carrier, with a basic crew of 3–10 operators.

The eventual goal, however, is to enable a single human user to deploy networks of unmanned aerial, ground, surface, and underwater vehicles. There have been several recent demonstrations of multirobot systems exploring urban environments [41.20, 21] and the interiors of buildings [41.19, 22] to detect and track intruders, and transmit all of the above information to a remote operator. These examples show that it is possible to deploy networked robots using an off-the-shelf 802.11b wireless network and have the team be remotely tasked and monitored by a single operator. An example of a project with heterogeneous vehicles in an urban setting is shown in Fig. 41.4. An example of a project with heterogeneous vehicles in an indoor setting is shown in Fig. 41.5 wherein robots map an environment and deploy themselves to form a sensor network to detect intruders.

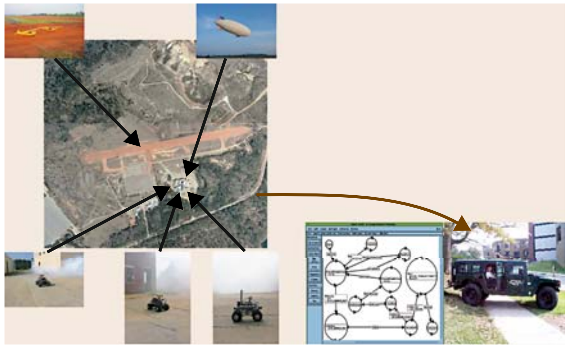

Fig. 41.4 A single operator commanding a network of aerial and ground vehicles from a command and control vehicle in an urban environment for scouting and reconnaissance in a recent demonstration by the University of Pennsylvania, Georgia Tech. and University of Southern California [41.18]

Fig. 41.5 Under the DARPA SDR program, a team from the University of Southern California, the University of Tennessee, and Science Applications and International Corporation (SAIC) demonstrated mapping, and intruder detection by a team of networked robots [41.19]

Many research projects are addressing group behaviors or collective intelligence by realizing swarming behaviors observed in nature. For example, the European Union (EU) has several EU-wide coordinated projects on collective intelligence or swarm intelligence. The I-Swarm project in Karlsruhe [41.23] and the Swarm-Bot project at Ecole Polytechnique Fédérale de Lausanne (EPFL) [41.24] are examples of swarm intelligence. The Laboratory for Analysis and Architecture of Systems (LAAS) has a strong group in robotics and artificial intelligence. This group has had a long history of basic and applied research in multirobot systems. The integration of multiple unmanned vehicles for applications such as terrain mapping and fire-fighting is addressed in [41.25]. A recent multi-university US project addresses the development of networked vehicles for swarming behaviors [41.26]. Projects such as these are exploring the scalability of the basic concepts to large numbers of robots, sensors, and actuators.

41.3 Research Challenges

While there are many successful embodiments of networked robots with applications to manufacturing industry, the defense industry, space exploration, domestic assistance, and civilian infrastructure, there are significant challenges that have to be overcome.

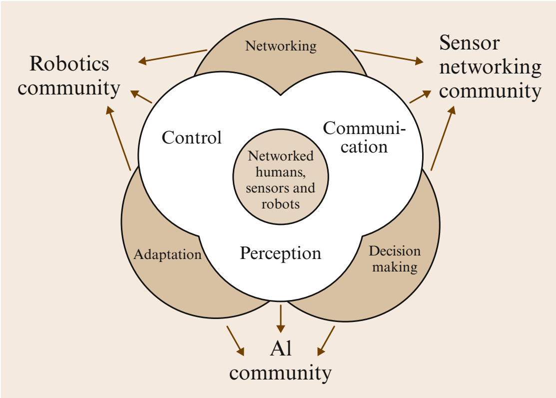

Fig. 41.6 The paradigm of *networked robots* introduces fundamental challenges at the intersection of control, perception, and communication that is of interest to the robotics, sensor networks, and artificial intelligence communities

The problem of coordinating multiple autonomous units and making them cooperate creates problems at the intersection of communication, control, and perception. Who should talk to who and what information should be conveyed, and how? How does each unit move in order to accomplish the task? How should the team members acquire information? How should the team aggregate information? These are all basic questions that need basic advances in control theory, perception, and networking. In addition, because humans are part of the network (as in the case of the Internet), we have to device an effective way for multiple humans to be embedded in the network and command/control/monitor the network without worrying about the specificity of individual robots in the network. Thus the underlying research challenges lie at the intersection of control theory, perception, and communication/networks, as shown in Fig. 41.6.

It is also worth noting that robot networks are dynamic unlike networks of sensors, computers or machines which might be networked together in a fixed topology. When a robot moves, its neighbors change and its relationship to the environment changes. As a consequence, the information it acquires and the actions it executes must change. Not only is the network topology dynamic, but the robot's behavior also changes as the topology changes. It is very difficult to predict the performance of such dynamic robot networks, yet it is this analysis problem that designers of robot networks must solve before deploying the network.

This notion of a changing topology inevitably leads us to complicated mathematical models. Traditionally, models of group behavior have been built on continuous models of dynamics of individuals, including local interactions with neighbors, and models of control and sensing with a fixed set of neighbors. While dynamics at the level of individual units may be adequately described by differential equations, the interactions with neighbors are best described by edges on a graph. Modeling, analysis, and control of such systems will require a comprehensive theoretical framework and new representational tools. New mathematical tools that marry dynamical system theory, switched systems, discrete mathematics, graph theory, and computational geometry are needed to solve the underlying problems. We need a design methodology for solving the *inverse problem in navigation* – behaviors for controlling individuals to achieve a specified aggregate *motion* and *shape* of the group, and the application to active perception and coverage. An overview of some of these methods is provided in Sect. 41.4.

Problems of perception have been studied extensively in the robotics community. However, the perception problems in a system of networked, mobile sensor platforms bring a new set of challenges; for example, consider the problem of estimating the state of the network. State estimation requires the estimation of the state of robots and the environment based on local, limited-range sensory information. Localization of n vehicles in an m-dimensional configuration space requires $O((nm)^k)$ computations, where k is somewhere between 3 and 6, depending on the algorithm and domain-specific assumptions. The estimation problem is further exacerbated by the fact that not all robots in the network may be able to get the necessary information in a time-critical fashion. There are deep issues of representation and algorithmic development, which are discussed in Sect. 41.6.

The paradigm of active perception [41.27] links the control of sensor platforms to perception, bringing control theory and perception together in a common framework. Extending this paradigm to networked robots requires approaches of distributed control to be merged with decentralized estimation. Robots can move in order to localize themselves with respect to their neighbors, to localize their neighbors, and also to identify, localize, and track features in the environment. These problems are discussed in Sect. 41.7.

As discussed earlier, the communication network is central to the functioning of a network of robots. However, if the network consists of mobile agents with transmitters and receivers with finite power, there is no guarantee that all agents can talk to teach other. Unlike a static sensor network, robots in a network can move toward each other to facilitate communication and adaptively maintain a communication network. Some basic algorithmic problems and several pertinent results are provided in Sect. 41.8.

41.4 Control

The control of individual robots is critical to the performance and scope of robot networks. Indeed motion coordination algorithms have been proposed for the purpose of improving communication performance [41.28, 29], localization [41.30, 31], information integration, deployment [41.32], and coverage [41.33–35], among other tasks. Mobility allows the group of robots to self-deploy, and self-organize by relocating themselves in support of communication, sensing, or task needs; for example, they can reconfigure to guarantee a desired communication bandwidth, k-hop connectivity, or algebraic connectivity, enabling message delivery from one robot to another. The group can also self-organize to position sensors so as to cover a desired area and adapt to shifts in the focus of monitoring activities. Controlling sensor position also supports map making, tracking of objects and events, and goal-directed navigation for users of the network. Finally, mobility allows robots to accomplish tasks such as navigation, reconnaissance, transportation, and search and rescue.

Given a group of mobile sensors, we would like to have distributed control capabilities that realize desirable global specifications. Thus, it is necessary to be able to automatically determine the necessary position and orientation of the group members and/or the distribution of group members, and their motion to achieve the desired task. At a lower level, the robots must be able to use information from the communication network and from their own sensors to derive local estimates, reason about the spatial network (their neighbors and their relationship to the environment), and then use the appropriate control policies to achieve the desired group specifications. We briefly outline the simplest mathematical model that is necessary to formulate such problems in order to provide a better sense of the underlying challenges.

In a robot network, we have multiple agents or nodes in which each agent is a physical entity that can be a robot, a vehicle with actuators and sensors, a sensor platform (possibly static) or even a communication relay node. Each agent A_i is characterized by an identifier, $i \in I \subset Z$, a state $x_i \in X_i \subset \mathbb{R}^n$, and control inputs $u_i \in U_i \subset \mathbb{R}$, with $f_i : X_i \times U_i \to TX_i$ specifying the dynamics:

$$\dot{x}_i = f_i(x_i, u_i) . \tag{41.1}$$

The state x_i will consist of the position (and orientation), r_i in some d-dimensional space, and its velocity, $\dot{r}_i : x_i = \left(r_i^T, \dot{r}_i^T\right)^T$, with $n = 2d$. $\mathcal{N}^{\mathrm{c}}(r_i)$ and $\mathcal{N}^{\mathrm{s}}(r_i)$ are neighborhoods of r that define the range and field of view of the communication hardware and sensors, respectively.

A network of robots S consists of N agents with a *sensing graph* and a *communications graph* that is defined by the physical distribution of the agents. The sensing graph (and similarly the communications graph) is defined by a map $E^{\mathrm{s}} : X^1 \times X^2 \dots X^N \to I \times I$, where the edges of the graph are formed dynamically depending on the physical proximity of pairs of agents. Specifically, the $N \times N$ adjacency matrix, $\mathcal{A}^{\mathrm{s}}$ (and similarly $\mathcal{A}^{\mathrm{c}}$) has entries

$$\mathcal{A}^{\mathrm{s}}{}_{ij} = \begin{cases} 1, & \text{if } r_j \in \mathcal{N}^{\mathrm{s}}(r_i) , \\ 0, & \text{otherwise} . \end{cases} \tag{41.2}$$

Agent A_i has estimates of its own state and the states of neighbors (e.g., A_j), and these estimates are derived from information associated with edges in the sensing and communication graph

$$\hat{x}_j^{(i)} = h(x_i, z_{ij}) , \tag{41.3}$$

where z_{ij} represents measurements of the state of agent A_j available to A_i by sensing or communication channels and h is the estimator used by A_i. Note that z_{ij} may have dimension less than n and may therefore not contain complete information about $x_{ij} = x_i - x_j$. Clearly the relative position vector denoted by $r_{ij} = r_i - r_j$ and its magnitude are important quantities that may need to be estimated for biological and artificial agents.

Finally, A_i can encode n_{b_i} behaviors, which we will denote by $\mathcal{B}_i = B_1, B_2, \dots, B_{n_{b_i}}$. Each behavior B_j is a controller, a function $k_j : \mathbb{R} \times X_i \to U_i$. All agents can be assigned identical or different behaviors. Each behavior represents a set of unsynchronized, locally executed computations (for control or estimation) being carried out for some collective purpose, with each processor using in its computations only data from its *neighboring* processors. Furthermore, even for a fixed assignment of behaviors, each processor's neighbors typically *change* with time because the processors are moving in and out of the sets $\mathcal{N}^{c}$ and $\mathcal{N}^{s}$. Thus the methodology for modeling and analyzing such systems will require the merging of graph theory and dynamical system theory at a fundamental level.

The reader is directed to the many survey articles on this subject for further information. An overview of challenges for the controls community is presented in [41.36]. The underlying theory for networked mobile systems has been explored in the context of automated highway systems [41.37], cooperative robot reconnaissance [41.19] and manipulation [41.38], formation flight control [41.39], and the control of groups of unmanned vehicles [41.21]. Our goal in the following sections is to explore the connections between communication, perception, and control.

41.5 Communication for Control

Communication networks allow physically disconnected entities to exchange information. At the lowest level, when groups of vehicles coordinate their actions, communication allows vehicles to exchange state information [41.40–42]. At a higher level, robots can plan navigation and exploration tasks based on an integrated map of the world derived from information acquired from different robots [41.43].

The use of communication for control in the multivehicle context has been addressed in the PATH project where formations of inline vehicles were studied [41.37]. Problems of the stability of the formation [41.44], the convergence of the formation to shapes [41.45], and the overall performance of the system [41.46] are of great interest. The performance of the system is directly influenced by the interconnections between agents. In addition to impacting on stability [41.37], feedback of states from different agents and feedforward information from the plans of different agents affects the rates at which the system of agents can respond to external stimuli [41.46] or to commands from human operators [41.47].

In addition, communication can be used for high-level control and planning of robots. There is great interest in using static sensor nodes as beacons to guide robot navigation. In [41.48], the problem of coverage and exploration of an unknown dynamic environment using a mobile robot is considered. An algorithm is presented which assumes that global information is not available [neither a map, nor global positioning system (GPS) information]. The algorithm deploys a network of radio beacons that assists the robot in coverage. The network is also used by the robot for navigation. The deployed network can also be used for applications other than coverage (such as multirobot task allocation). A similar idea was presented using potential-field-based navigation in [41.43]. In this work the notion of no-go or danger areas was incorporated into the navigation cost function. Recent work along these lines with experimental data from sensor nodes is reported in [41.49].

In such communication-enabled cooperative control and planning (see also [41.50]), the communication network plays an important role in the creation of a shared representation of information. This notion of a shared representation is important to the scaling of coordinated control algorithms to large numbers of devices. For example, in [41.41], the information form of the Kalman filter is used to derive a framework for decentralizing estimation and fusion algorithms. This approach was shown to be applicable to multiple heterogenous ground and aerial platforms [41.30]. Such methodologies are transparent to the specificity and identity of the cooperating vehicles. This is because vehicles share a common representation, which consists of a certainty grid that contains information about the probability of detection of targets, and an information vector–matrix pair that is used in the information form of the Kalman filter [41.21]. Observations are propagated through the network by changing both the certainty grid and the information vector/matrix. This allows each vehicle to choose the action that maximizes a utility function, which is the combined mutual information gain from onboard sensors towards the detection and localization of features in the environment.

Thus, in summary, at the lowest level, communication enables either partial or complete state feedback of the network and allows agents to exchange information for feedforward control. At the higher levels, agents can share information for planning and for control. This is also discussed in Sect. 41.6 where the communication network is shown to enable a network-centric approach to perception.

41.6 Communication for Perception

While individual robots have sensors and the ability to build maps and models by integrating sensory information, networked robots can exchange information and leverage sensory data, maps, and models from other robots. The challenge is to exploit communication for perception in tasks such as distributed mapping in the presence of the delays, limited bandwidth, and disruption that are typical of communication networks.

Distributed localization is the term used to describe the merging of communication and perception for state estimation. Localization is an essential tool for the development of low-cost robot networks for use in location-aware applications and ubiquitous networking [41.51]. Location information is needed to track the placement of the nodes and to correlate the values measured by the node with their physical location. Distributed computation and robustness in the presence of measurement noise are key ingredients for a practical localization algorithm that will give reliable results over a large-scale network.

The methods for distributed localization can be classified into two broad classes: algorithms that rely on anchor nodes for localization and algorithms that use no beacons. Localization may be computed using range information between nodes, bearing information, or both.

In [41.28] a theoretical foundation for network localization in terms of graph rigidity theory is provided. The problem is solved when nodes have perfect range information and it is showm that a network has a unique localization if and only if its underlying graph is *generically globally rigid*. In [41.52] the Cramér–Rao lower bound (CRLB) for network localization is derived. This work computes the expected error characteristics for an ideal algorithm, and compares this to the actual error in an algorithm based on multilateration, drawing the important conclusion that the error introduced by the algorithm is just as important as the measurement error in assessing end-to-end localization accuracy. In [41.53] a distributed algorithm that uses no beacons and is guaranteed to compute correct location information under measurement noise for nodes that can range to neighbors is presented. This algorithm relies on the notion of robust quadrilaterals to compute robustly a global system of coordinates among the nodes. The computation supports moving nodes. Extensions of this work to passive tracking have been discussed in [41.54]. Localization based on the propagation of location information from known reference nodes based on connectivity includes [41.55,56]. Mobility-assisted localization is introduced in [41.57]. Other techniques use distributed propagation of location information using multilateration [41.52, 58].

Two approaches for cooperative relative localization of mobile robot teams are given in [41.59, 60]. Neither method uses GPS, landmarks, or maps of any kind; instead, robots make direct measurements of the relative pose of nearby robots and broadcast this information to the team as a whole. In [41.59], each robot processes this information independently to generate an egocentric estimate for the pose of other robots using a Bayesian formalism with a particle filter implementation. In [41.60], maximum-likelihood estimation (MLE) and numerical optimization is used to achieve a similar result.

A key issue is to be able to scale these computations for building a shared representation to large numbers of robots and sensors. This problem was studied in recent experiments under the US Defense Advanced Research Projects Agency (DARPA-funded software for distributed robotics (SDR) program. The goal of these experiments was to develop and demonstrate a multi-robot system capable of carrying out a specific mission. This required the ability to deploy a large number of robots into an unexplored building, map the building interior, detect and track intruders, and transmit all of the above information to a remote operator. A report on one set of experiments is presented in [41.19]. A tiered strategy for deploying the robots is described, where highly capable robots formed the first wave to enter and map a building, followed by a second wave which used the resulting map to self-deploy and monitor the environment for intruders. Both approaches relied extensively on networking the robots using commercial 802.11b wireless technology. This task involved both communication for building a shared representation as well control for perception.

Another important set of problems arises when robot networks are used for identifying, localizing, and then tracking targets in a dynamic setting. An embedded stationary wireless sensor network is like a virtual sensor spread over a large geographical area. Such a network can provide information to mobile robots about remote locations. Robot networks allow this virtual sensor to move in response to external stimuli and to track moving targets. Indeed, it is possible to cast this scenario as a pursuit-evasion game with robotic sensor networks [41.61]. For example, the Tenet project at USC addresses the design of network primitives and abstractions for tiered network architectures, with robotic pursuit evasion as one of the target applications. Algorithms for guiding the sampling strategy of a robotic boat to model and locate phenomena of interest (e.g., hotspots) in aquatic environments are discussed in [41.11]. The networked infomechanical systems (NIMS) project has focused on sensor-assisted techniques for mobile robot-based adaptive sampling for event response [41.62] and field reconstruction [41.63].

The information collected by the nodes in a sensor network can be processed at a central location or in a decentralized fashion. Such in-network data processing techniques make better use of network communication and computation resources than centralized processing. This also enables the network to compute accurate and up-to-date global pictures of the global perception landscape that are available to all the robots in the system. Methods for in-network data processing with static nodes include artificial potential-field computation, gradient computations, particle filters, Bayesian inference, and signal processing. Algorithms have been developed for computing maps, paths, and predictors [41.43, 48, 64].

A recent DARPA demonstration showed how communication networks can be used effectively in perception tasks involving heterogenous robots [41.20]. In cooperative search, identification, and localization unmanned aerial vehicles (UAVs) can be used to cover large areas, searching for targets. However, sensors on UAVs are typically limited in their accuracy of localization of targets on the ground. On the other hand, ground robots can be deployed to accurately locate ground targets but have the disadvantage of not being able to move rapidly and see through obstacles such as buildings or fences. In [41.30], the synergy between these two devices is exploited by creating a seamless network of UAVs and unmanned ground vehicles (UGVs). As discussed in Sect. 41.5, the key to such network-centric approaches for search and localization is a shared representation of state information, which in this case is easily scalable to large numbers of UAVs and UGVs and is transparent to the specificity of individual platforms. However, how to do this more generally and for more unstructured information remains an issue for future research.

41.7 Control for Perception

Networked mobile robots enable the exploration of dynamic environments and the recovery of three-dimensional information via distributed active perception [41.27]. Since the nodes are mobile, a natural question is: where should the nodes be placed in order to ensure successful integration of information from multiple nodes, and to maximize the quality of the estimates returned by the team? Since there is a cost associated with transmitting and processing data, it is important to consider which sensor readings should be used in the state estimation and what information should be communicated to the rest of the system. The quality of the information computed by the network depends on the locations of the sensor platforms both in an absolute and relative sense. The quality also depends on the noise characteristics of each sensor, and the communication network.

A robot network goes well beyond a fixed sensor network, which can only collect data at fixed positions in space; for example, when an event is detected at a specific location it is possible to direct more sensors toward the location of observation of the event for more information (for example, higher-resolution data or higher sampling frequency). Reconfiguring the node locations for adaptive resolution sampling relies on distributed control strategies.

Various strategies have been introduced for controlling mobile sensor network coverage. Mobile sensing agents are controlled using gradients of information-based objective functions [41.65]. Stability results are derived without concerns for the optimality of the network configuration, but local guarantees are provided. A body of results reported in [41.66] and [41.67] describes decentralized control laws for positioning mobile sensor networks optimally with respect to a known event distribution density function. This approach is advantageous because it guarantees that the network (locally) minimizes a cost function relevant to the coverage problem. However, the control strategy requires that each agent have a complete knowledge of the event distribution density, thus it is not reactive to the sensed environment. The work by [41.68, 69] generalizes these results to situations in which the nodes estimate rather than know ahead of time the event distribution density function. A local (decentralized) control law requires that each agent can measure the value and gradient of the distribution density function at its own position. This results in a sensor network that is reactive to its sensed environment while maintaining or seeking a near-optimal sensing configuration. In addition, the distribution density function approximation yields a closed-form expression for the control law in terms of the vertices of an agent's Voronoi region. This eliminates the need for the numerical integration of a function over a polynomial domain at every time step, thereby providing a significant reduction in computational overhead for each agent. Other work in event monitoring for unknown distributions includes [41.33]. *Krause* et al. [41.70] have recently proposed an approach for sensor placement that considers both the sensing quality and communication cost of imperfect sensing and communication components. They use a parametric model for link reception rate that assumes no acknowledgement and no temporal correlation of lossy links.

Beginning with the art gallery problem, there have been multiple efforts to determine an optimal configuration of sensors to cover a given region [41.71–73]. A variant which allows the use of mobile sensors is known as the *watchmen tours* problem. In these approaches the sensor model is abstract and not well suited to real environments and cameras. Distributed geometric optimization methods [41.67] have also been used for mobile sensor network reconfiguration. A related class of methods is the use of estimation-theoretic optimization metrics and the application of information filters to coordinate network-wide motion [41.30]. There are other distributed optimization methods which use a distributed control law and show that it optimizes a global metric of interest, such as using a potential field or other linear control law based only on local neighbor interactions [41.74]. Research focusing on the control of cameras with pan, tilt, and zoom capabilities is due to [41.34, 75, 76]. In [41.75] an approach is developed to calibrate a pan–tilt–zoom camera automatically over its full zoom range and to build very high-resolution panoramas. In [41.34], the cameras are constantly moved to track observed targets, using a factor graph. A recent algorithm due to [41.77] significantly improves on this by positioning cameras to make the network better suited to detect and classify targets as they emerge. Pan–tilt–zoom cameras allow the construction of far more flexible vision systems than static cameras.

41.8 Control for Communication

In Sect. 41.5, we briefly discussed the benefits of using the communication network to synthesize and improve controller design. Conversely, the movement of robots affects the network and data transmission in the network. This gives rise to many challenges. If the controllers for individual robots are known, can we provide guarantees about communication in the network and can we develop robust information routing and networking algorithms in the presence of robot motion? Another challenge concerns how information propagates and diffuses in these networks. If the robots move under a given control model, how does a piece of information propagate through the network and what can we say about when and where that piece of information will be heard? If we know the answers to such questions, it may be possible to design controllers to realize desired communication network characteristics.

One simple control strategy that can affect network performance is to control the robot motion to ensure messages are transmitted between designated nodes. The movement of robots in a network of robots and sensors may cause network partitioning when nodes go out of range. However, the ability of the robots to move in a controlled way also leads to an opportunity to address the information routing problem in disconnected networks by turning the robots into relay nodes. The key idea here is to enable the robot holding a current message to an unavailable destination to modify their trajectory in order to relay a message. This problem has been formulated as an optimization problem. The goal is to minimize the trajectory modifications necessary to send a message to its destination. Several solutions have been proposed depending on the information that is available to the robots. If the robots' trajectories are known, path planning techniques can be used to compute who moves where to relay what. If the robots' trajectories are not known, a distributed spanning tree can be created to enable robots to keep track of each other. Each robot is assigned a region of movement and a parent in the spanning tree. When the robot leaves its region, the parent is informed. When the robot moves too far away, the spanning tree is modified.

Mobile robots can be used to create desired network topologies under suitable models of network communication. If a robot is used to emplace nodes in an environment (or if sensor nodes robotically self-deploy) to build a network, the problem is referred to as deployment. It is possible to control the motion of individual nodes to guarantee that a specified topology is maintained [41.29]. It is also possible to repositioning nodes with the explicit aim of changing the network topology – the so-called mobility-based topology control problem.

A distributed algorithm for the deployment of mobile robot teams has been described by [41.78] using the concept of virtual pheromones: localized messages from one robot to another. These messages are used to generate either a gas expansion or a guided growth deployment model. Similar algorithms based on artificial potential fields are described in [41.79, 80], where the latter incorporates a connectivity constraint. An incremental deployment algorithm for mobile sensor networks is given in [41.32]; nodes are deployed one at a time into an unknown environment, with each node making use of information gathered by previously deployed nodes to determine its deployment location. The algorithm is designed to maximize network coverage while ensuring that nodes retain line of sight with one another.

Most work on network topology control has dealt with uncontrolled deployments, where there is no explicit control of the positions of individual nodes. The primary mechanisms proposed are power control and sleep scheduling. These methods involve pruning an already existing, well-connected communication graph in order to save power while ensuring that the resultant subgraph preserves connectivity. Given a network that is connected when all nodes are operating at maximum power, the aim of power control is to use the minimum power level at each node for which the network remains connected [41.81]. Given an overdeployed network, sleep scheduling seeks to activate a minimal subset of nodes to maintain connectivity and achieve other desired metrics [41.82]. In contrast, controlled deployments are feasible when the positions of individual nodes can be altered. Such deployments are interesting for two reasons. First, network topology with wireless communication relates directly to proximity relations and hence the position of the nodes. Second, there is increasing evidence that a large number of deployments are likely to involve careful, nonrandom placement of nodes. The positioning of nodes is controlled either by the nodes themselves or by external agents. Such networks present a different and interesting scenario for topology control since it is possible to exploit control of the motion and placement of the nodes to build efficient topologies. A local, completely decentralized technique for topology control using mobility is given in [41.83].

An important application for networked robots is in monitoring and surveillance, where it is important that the robots cover the space while remaining within communication range [41.84]. Probing environment and adaptive sleeping protocol (PEAS) was one of the first attempts to address communication connectivity and sensing coverage simultaneously using heuristic algorithms [41.85]. *Wang* et al. [41.86] proposed a new coverage configuration protocol (CCP) to produce an approach that simultaneously optimizes coverage and connectivity while maximizing the number of nodes that are placed into sleep mode. Furthermore, they also identified three different classes of coverage–connectivity problems with respect to the ratio of radio and sensing ranges and recognized the critical ratio where the former range is twice as long as the latter. *Zhang* and *Hou* proved that, if the communication range is at least twice the sensing range, complete coverage of a convex area guarantees network communication connectivity, and then used this theorem as a basis for a localized density control algorithm [41.82]. This was subsequently generalized to show that the condition that the communication range is twice the sensing range is sufficient and is the tight lower bound to guarantee that complete coverage preservation implies communication connectivity among nodes if the original network topology is connected [41.87].

In summary, if the state of the communication network and the desired state of the communication network is known to each agent, it should be possible to synthesize distributed controllers to move agents to attain desired network characteristics. However, the assumptions on the global state are clearly not justified. Also, the desired motion to optimize network characteristics will conflict with the motion that is required to perform the desired task. However, as the brief discussion above illustrates, there are many interesting studies that point to promising directions for future work in this very fertile research field.

41.9 Conclusions and Further Reading

The paradigm of networked robots offers significant potential for accomplishing tasks that cannot be accomplished by individual robots. Indeed this paradigm is critical to such tasks as environmental monitoring, surveillance and reconnaissance, and security for civilian or defense purposes. However, there are many scientific challenges to realizing this vision for networked robots. The main overarching challenges are summarized here.

Technical challenges to scalability: We do not yet have a methodology for creating self-organizing robot networks that are robust to labeling (or numbering), with completely decentralized controllers and estimators, and with provable emergent response. This requires basic research at the intersection of control, perception, and communication.

Performing physical tasks in the real world: Most of our present applications emphasize going from static sensor networks to mobile sensor networks and, as such, are able to acquire and process information. We are a long way from creating robust robot networks that can perform physical tasks in the real world.

Human interaction for network-centric control and monitoring: Advances over the last decade have provided human users with the ability to interact with hundreds or thousands of computers on the Internet. It is necessary to develop similar network-centric approaches to interfacing, both for control and for monitoring.

Finally, a major challenge is to create robot networks that are proactive and anticipate our needs and commands rather than reacting (with delays) to human commands.

While there are major challenges ahead, there is no denying the tremendous potential of networked robots. This chapter and other related chapters in Part D (Chap. 31) and Part E (Chap. 40) show that the research community is making steady progress and confronting these challenges head on.

The references for this chapter includes many pointers to the specific research issues covered here. For some excellent further background on networked robotics we direct the reader to [41.5, 6, 19, 26, 29, 50, 64, 77].

References

41.1 Z. Butler, K. Kotay, D.L. Rus, M. Vona: Self-reconfiguring Robots, Commun. ACM **45**(3), 39–45 (2002)

41.2 O. Khatib, K. Yokoi, K. Chang, D. Ruspini, R. Holmberg, A. Casal: Coordination and decentralized cooperation of multiple mobile manipulators, J. Robot. Syst. **13**(11), 755–764 (1996)

41.3 L.E. Parker: The effect of heterogeneity in teams of 100+ mobile robots. In: *Multi-Robot Systems Volume II: From Swarms to Intelligent Automata*, ed. by A. Schultz, L.E. Parker, F. Schneider (Kluwer, Dordrecht 2003)

41.4 M. Yim, Y. Zhang, D. Duff: Modular robots, IEEE Spect. **39**(22), 30–34 (2002)

41.5 N. Hagita: Introduction to network robots, Introduction to Network Robots (Barcelona, 2005)

41.6 D. Estrin: *Embedded, Everywhere* (National Academies Press, Washington 2001)

41.7 J. Parrish, S. Viscido, D. Grunbaum: Self-organized fish schools: An examination of emergent properties, Biol. Bull. **202**, 296–305 (2002)

41.8 N. Franks, S. Pratt, E. Mallon, N. Britton, D. Sumpter: Information flow, opinion polling and collective intelligence in house-hunting social insects, Philos. Trans. R. Soc. London. B **357**, 1567–1584 (2002)

41.9 R.L. Jeanne: Group size, productivity, and information flow in social wasps. In: *Information Processing in Social Insects* (Birkhauser, Basel 1999)

41.10 Argo Floats: *A global array of 3,000 free-drifting profiling floats for environmental monitoring* (Argo Information Center, Ramonville 2007), http://www.argo.ucsd.edu/index.html

41.11 G.S. Sukhatme, A. Dhariwal, B. Zhang, C. Oberg, B. Stauffer, D.A. Caron: The design and development of a wireless robotic networked aquatic microbial observing system, Environmen. Eng. Sci. **24**(2), 205–215 (2006)

41.12 W. Kaiser, G. Pottie, M. Srivastava, G.S. Sukhatme, J. Villasenor, D. Estrin: Networked infomechanical systems (nims) for ambient intelligence. In: *Ambient Intelligence* (Springer, Berlin, Heidelberg 2005) pp. 83–114

41.13 Amarss: *Networked minirhizotron planning and initial deployment* (Center for Embedded Networkedsensing, Los Angeles 2007), http://research.cens.ucla.edu/

41.14 H. Singh, J. Catipovic, R. Eastwood, L. Freitag, H. Henriksen, F.F. Hover, D. Yoerger, J.G. Bellingham, B.A. Moran: An integrated approach to multiple AUV communications, navigation and docking, IEEE Oceans (1996) pp. 59–64

41.15 I. Vasilescu, M. Dunbabin, P. Corke, K. Kotay, D. Rus: Data collection, storage, and retrieval with an underwater sensor network, Proceedings of the ACM Sensys 2005 (San Diego, 2005)
41.16 N. Leonard, D. Paley, F. Lekien, R. Sepulchre, D.M. Fratantoni, R. Davis: Collective motion, sensor networks and ocean sampling, Proc. IEEE (2006)
41.17 D.O. Popa, A.S. Sanderson, R.J. Komerska, S.S. Mupparapu, D.R. Blidberg, S.G. Chappell: Adaptive sampling algorithms for multiple autonomous underwater vehicles, IEEE/OES AUV2004: A Workshop on Multiple Autonomous Underwater Vehicle Operations (2004)
41.18 M.A. Hsieh, A. Cowley, J.F. Keller, L. Chaimowicz, B. Grocholsky, V. Kumar, C.J. Talyor, Y. Endo, R. Arkin, B. Jung, D. Wolf, G. Sukhatme, D.C. MacKenzie: Adaptive teams of autonomous aerial and ground robots for situational awareness, J. Field Robot. **24**(11–12), 991–1014 (2007)
41.19 A. Howard, L.E. Parker, G.S. Sukhatme: Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection, Int. J. Robot. Res. **25**(5–6), 431–447 (2006)
41.20 L. Chaimowicz, A. Cowley, B. Grocholsky, M.A. Hsieh, J.F. Keller, V. Kumar, C.J. Taylor: Deploying air-ground multi-robot teams in urban environments, Third Multi-Robot Systems Workshop (Washington, 2005)
41.21 B. Grocholsky, R. Swaminathan, J. Keller, V. Kumar, G. Pappas: Information driven coordinated air-ground proactive sensing, IEEE International Conference on Robotics and Automation (IEEE, Barcelona 2005)
41.22 D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, B. Stewart: Distributed multi-robot exploration and mapping, Proc. IEEE (2006)
41.23 J. Seyfied, M. Szymanski, N. Bender, R. Estana, M. Theil, H. Worn: The I-Swarm Project: Intelligent Small World Autonomous Robots for Micro-manipulation, SAB 2004 International Workshop (Santa Monica 2004)
41.24 F. Mondada, G.C. Pettinaro, A. Guignard, I.W. Kwee, D. Floreano, J.-L. Deneubourg, S. Nofli, L.M. Gambardella, M. Dorigo: Swarm-Bot: A new distributed robotic concept, Auton. Robot. **17**, 193–221 (2004)
41.25 A. Ollero, S. Lacroix, L. Merino, J. Gancet, J. Wiklund, V. Remuss, I. Veiga, L.G. Gutierrez, D.X. Viegas, M.A. Gonzalez, A. Mallet, R. Alami, R. Chatila, G. Hommel, F.J. Colmenero, B. Arrue, J. Ferruz, R. Martinez de Dios, F. Caballero: Architecture and perception issues in the COMETS multi-UAV project, IEEE Robot. Autom. Mag. (2005)
41.26 Q. Li, D. Rus: Navigation protocols in sensor networks, ACM Trans. Sensor Netw. **1**(1), 3–35 (2005)
41.27 R. Bajcsy: Active perception, Proc. IEEE **76**, 996–1005 (1988)
41.28 T. Eren, D. Goldenberg, W. Whitley, Y.R. Yang, S. Morse, B.D.O. Anderson, P.N. Belhumeur: Rigidity, computation, and randomization of network localization, Proceedings of IEEE INFOCOM '04 (Hong Kong 2004)
41.29 A. Hsieh, A. Cowley, V. Kumar, C.J. Taylor: Towards the deployment of a mobile robot network with end-to-end performance guarantees, IEEE Int'l Conf. on Robotics and Automation (Orlando, 2006)
41.30 B. Grocholsky, S. Bayraktar, V. Kumar, C.J. Taylor, G. Pappas: Synergies in feature localization by air-ground teams, Proc. 9th Int'l Symposium of Experimental Robotics (Singapore 2004)
41.31 B. Grocholsky, E. Stump, V. Kumar: An extensive representation for range-only SLAM, International Symposium on Experimental Robotics (Rio de Janeiro, 2006)
41.32 A. Howard, M.J. Matarić, G.S. Sukhatme: An incremental self-deployment algorithm for mobile sensor networks, Auton. Robot. **13**(2), 113–126 (2002), Special Issue on Intelligent Embedded Systems
41.33 Z. Butler, D. Rus: Controlling mobile networks for monitoring events with coverage constraints, Proceedings of the IEEE International Conference on Robotics and Automation (2003)
41.34 M. Chu, J. Reich, F. Zhao: Distributed attention for large video sensor networks, Intelligent Distributed Surveillance Systems 2004 seminar (2004)
41.35 B. Jung, G.S. Sukhatme: Tracking targets using multiple robots: The effect of environment occlusion, Auton. Robot. **13**(3), 191–205 (2002)
41.36 R.M. Murray, K.J. Åström, S.P. Boyd, R.W. Brockett, G. Stein: Future directions in control in an information-rich world, IEEE Contr. Syst. Mag. (2003)
41.37 A. Pant, P. Seiler, K. Hedrick: Mesh stability of look-ahead interconnected systems, IEEE Trans. Autom. Control. **47**, 403–407 (2002)
41.38 T. Sugar, J. Desai, V. Kumar, J.P. Ostrowski: Coordination of multiple mobile manipulators, Proc. IEEE Int. Conf. Robot. Autom. **3**, 2022–2027 (2001)
41.39 R.W. Beard, J. Lawton, F.Y. Hadaegh: A coordination architecture for spacecraft formation control, IEEE Trans. Control Syst. Technol. **9**, 777–790 (2001)
41.40 A. Das, J. Spletzer, V. Kumar, C. Taylor: Ad hoc networks for localization and control, Proc. of the IEEE Conf. on Decision and Control (2002) pp. 2978–2983
41.41 J. Manyika, H. Durrant-Whyte: *Data Fusion and Sensor Management: An Information-Theoretic Approach* (Prentice Hall, Upper Saddle River 1994)
41.42 H.G. Tanner, A. Jadbabaie, G.J. Pappas: Stable flocking of mobile agents, part i: Fixed topology, Proceedings of the 42nd IEEE Conference on Decision and Control (2003) pp. 2010–2015
41.43 Q. Li, D. Rus: Navigation Protocols in Sensor Networks, ACM Trans. Sensor Netw. **1**(1), 3–35 (2005)

41.44 J.M. Fowler, R. D'Andrea: Distributed control of close formation flight, Proc. of the IEEE Conf. on Decision and Control (2002) pp. 2972–2977

41.45 J.P. Desai, J.P. Ostrowski, V. Kumar: Modeling and control of formations of nonholonomic mobile robots, IEEE Trans. Robot. Autom. **17**(6), 905–908 (2001)

41.46 H.G. Tanner, V. Kumar, G.J. Pappas: Leader-to-formation stability, IEEE Trans. Robot. Autom. **20**(3) (2004)

41.47 S. Loizou, V. Kumar: Relaxed input to state stability properties for navigation function based systems, Proceedings of the 45th IEEE Conference on Decision and Control (San Diego, 2006)

41.48 M. Batalin, G.S. Sukhatme: Coverage, exploration and deployment by a mobile robot and communication network, Telecommun. Syst. J. **26**(2), 181–196 (2004), Special Issue on Wireless Sensor Networks

41.49 K. J. O'hara, V. Bigio, S. Whitt, D. Walker, T. R. Balch: *Evaluation of a Large Scale Pervasive Embedded Network for Robot Path Planning* (ICRA 2006) pp. 2072–2077

41.50 V. Kumar, N. Leonard, A.S. Morse (eds.): *Cooperative Control*, *Vol. 309 of Lecture Notes in Control and Information Sciences* (Springer, Berlin, Heidelberg 2004)

41.51 J. Chen, S. Teller, H. Balakrishnan: Pervasive pose-aware applications and infrastructure, IEEE Comput. Graph. Appl. **23**(4), 14–18 (2003)

41.52 A. Savvides, C.-C. Han, M. Srivastava: Dynamic Fine-Grained localization in Ad-Hoc networks of sensors, Proceedings of the Seventh Annual International Conference on Mobile Computing and Networking (MOBICOM-01) (ACM, New York 2001) pp. 166–179

41.53 D. Moore, J. Leonard, D. Rus, S.J. Teller: Robust distributed network localization with noisy range measurements, SenSys '04: Proceedings of the 2nd International conference on Embedded networked sensor systems (2004) pp. 50–61

41.54 C. Detweiler, J. Leonard, D. Rus, S. Teller: Passive Mobile Robot Localization within a Fixed Beacon Field, Proceedings of the 2006 International Workshop on Algorithmic Foundations of Robotics (2006)

41.55 N. Bulusu, J. Heidemann, D. Estrin: Adaptive beacon placement, Proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS-01) (Los Alamitos 2001) pp. 489–498

41.56 S. N. Simic, S. Sastry: *Distributed localization in wireless ad hoc networks*, Technical Report UCB/ERL M02/26, EECS Department (University of California, Berkeley 2001), http://www.eecs.berkeley.edu/Pubs/TechRpts/2002/4010.html

41.57 P. Corke, R. Peterson, D. Rus: Communication-assisted localization and navigation for networked robots, Int. J. Robot. Res. **4**(9), 116 (2005)

41.58 R. Nagpal, H.E. Shrobe, J. Bachrach: Organizing a global coordinate system from local information on an ad hoc sensor network, Lect. Notes Comput. Sci. **2634**, 333–348 (2003)

41.59 A. Howard, M.J. Matarić, G.S. Sukhatme: Putting the 'i' in 'team': An ego-centric approach to cooperative localization, IEEE International Conference on Robotics and Automation (Taipei, 2003) pp. 868–892

41.60 A. Howard, M.J. Matarić, G.S. Sukhatme: Localization for mobile robot teams using maximum likelihood estimation, IEEE/RSJ International Conference on Intelligent Robots and Systems (Lausanne, 2002) pp. 434–459

41.61 R. Vidal, O. Shakernia, H.J. Kim, D.H. Shim, S. Sastry: Probabilistic pursuit-evasion games: Theory, implementation and experimental evaluation, IEEE Trans. Robot. Autom. **18**(5), 662–669 (2002)

41.62 M. Batalin, M.H. Rahimi, Y. Yu, D. Liu, A. Kansal, G.S. Sukhatme, W. Kaiser, M. Hansen, G. Pottie, M. Srivastava, D. Estrin: Call and response: Experiments in sampling the environment, ACM SenSys (Baltimore, 2004) pp. 25–38

41.63 M.H. Rahimi, W. Kaiser, G.S. Sukhatme, D. Estrin: Adaptive sampling for environmental field estimation using robotic sensors, IEEE/RSJ International Conference on Intelligent Robots and Systems (2005) pp. 747–753

41.64 F. Zhao, L. Guibas: *Wireless Sensor Networks: An Information Processing Approach* (Morgan Kaufmann, New York 2004)

41.65 F. Zhang, B. Grocholsky, V. Kumar: Formations for localization of robot networks, IEEE Int'l Conf. on Robotics and Automation (New Orleans, 2004)

41.66 J. Cortes, S. Martinez, T. Karatas, F. Bullo: Coverage control for mobile sensing networks, IEEE Trans. Robot. Autom. **20**(2), 243–255 (2004)

41.67 J. Cortes, S. Martinez, F. Bullo: Spatially-distributed coverage optimization and control with limited-range interactions, ESAIM Contr. Optim. Calc. Variat. **11**, 691–719 (2005)

41.68 M. Schwager, J. McLurkin, D. Rus: Distributed coverage control with sensory feedback for networked robots, Proc. Robot. Sci. Syst. (RSS) (2006)

41.69 M. Schwager, J.-J. Slotine, D. Rus: Decentralized Adaptive Control for Coverage for Networked Robots, Proceedings of the 2007 International Conference on Robotics and Automation (2007)

41.70 A. Krause, C. Guestrin, A. Gupta, J. Kleinberg: Near-optimal sensor placements: Maximizing information while minimizing communication cost, Fifth International Conference on Information Processing in Sensor Networks (IPSN'06) (2006)

41.71 V. Chvatal: A combinatorial theorem in plane geometry, J. Combin. Theory Ser. **18**, 39–41 (1975)

41.72 J. O'Rourke: *Art Gallery Theorems and Algorithms* (Oxford Univ. Press, New York 1987)

41.73 S. Fisk: A short proof of Chvatal's watchmen theorem, J. Combin. Theory Ser. **24**, 374 (1978)

41.74 A. Jadbabaie, J. Lin, A.S. Morse: Coordination of groups of mobile autonomous agents using nearest neighbor rules, IEEE Trans. Autom. Control **48**(6), 988–1001 (2003)

41.75 S. Sinha, M. Pollefeys: Camera network calibration from dynamic silhouettes, Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (2004)

41.76 R. Collins, A. Lipton, H. Fujiyoshi, T. Kanade: Algorithms for cooperative multisensor surveillance, Proc. IEEE (2001)

41.77 A. Kansal, W. Kaiser, G. Pottie, M. Srivastava, G.S. Sukhatme: Reconfiguration methods for mobile sensor networks. Submitted to ACM Transactions on Sensor Networks (2006)

41.78 D. Payton, M. Daily, R. Estkowski, M. Howard, C. Lee: Pheromone robotics, Auton. Robot. **11**, 319–324 (2001)

41.79 A. Howard, M.J. Matarić, G.S. Sukhatme: Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem, Proceedings of the International Symposium on Distributed Autonomous Robotic Systems (2002) pp. 299–308

41.80 S. Poduri, G.S. Sukhatme: Constrained coverage for mobile sensor networks, IEEE International Conference on Robotics and Automation (New Orleans 2004) pp. 165–172

41.81 R. Wattenhofer, L. Li, P. Bahl, Y.M. Wang: A cone-based distributed topology-control algorithm for wireless multi-hop networks, IEEE/ACM Trans. Netw. (2005)

41.82 H. Zhang, J.C. Hou: On deriving the upper bound of alphalifetime for large sensor networks, ACM Trans. Sensor Netw. **1**(6) (2005)

41.83 S. Poduri, S. Pattem, B. Krishnamachari, G.S. Sukhatme: A unifying framework for tunable topology control in sensor networks. Technical Report CRES-05-004 (2005)

41.84 M.A. Hsieh, V. Kumar: Pattern generation with multiple robots, IEEE Int'l Conf. on Robotics and Automation (Orlando, 2006)

41.85 F. Ye, G. Zhong, J. Cheng, L. Zhang, S. Lu: Peas: A robust energy conserving protocol for long-lived sensor networks, International Conference on Distributed Computing Systems (2003)

41.86 X. Wang, G. Xing, Y. Zhang, C. Lu, R. Pless, C. Gill: Integrated coverage and connectivity configuration in wireless sensor networks, First ACM Conf. on Embedded Networked Sensor Systems, SenSys (2003)

41.87 D. Tian, N.D. Georganas: Connectivity maintenance and coverage preservation in wireless sensor networks. *AdHoc Networks* (2005)