

We begin by providing a brief overview of single-robot control philosophies and architectures, including behavior-based control, in Section 1. In Section 2 we move from single robots to multi-robot systems (MRS) and discuss the additional challenges this transition entails. In Section 3 we use empirical case studies to discuss and demonstrate three important ways in which robots can interact, and thus coordinate, their behavior. In Section 4 we discuss formal approaches to the design and analysis of MRS that are of fundamental importance if the full potential of MRS is to be achieved. Finally, in Section 5 we briefly discuss the future of coordinated behavior-based MRS and conclude the chapter.

1. Overview of Robot Control Architectures

In this section, we briefly discuss the most popular approaches and techniques for the control of a single robot. In the next section we proceed with the fundamental principle of this chapter, the control of multiple robots, and how it is related to, and different from, the control of a single robot.

1.1 Single Robot Control

We define *robot control* as the process of mapping a robot's sensory information into actions in the real world. We do not consider entities that make no use of sensory information in control decisions as robots, nor do we consider entities that do not perform actions as robots, because neither category is truly interacting in the real world. Any robot must, in one manner or another, use incoming sensory information to make decisions about what actions to execute. There are a number of control philosophies dictating how this mapping from sensory information to actions should occur, each with its advantages and disadvantages. A continuum of approaches to robot control can be described as a spectrum spanning from deliberative to reactive control.

The *deliberative* approach to robot control is usually computationally intensive due to the use of explicit reasoning or planning using symbolic representations and world models (3). For the reasoning processes to be effective, complete and accurate models of the world are required. In domains where such models are difficult to obtain, such as in dynamic and fast-changing environments or situations with significant uncertainty in the robot's sensing and action, it may be impossible for the robot to act in an appropriate or timely manner using deliberative control (3,4).

In contrast to deliberative control, the *reactive* approach to robot control is characterized by a tight coupling of sensing to action, typically involving no intervening reasoning (5,6). Reactive control does not require the acquisition or maintenance of world models, as it does not rely on the types of complex reasoning processes utilized in deliberative control. Rather, simple rule-based methods involving a minimal amount of computation,

internal representations, or knowledge of the world are typically used. This makes reactive control especially well suited to dynamic and unstructured worlds where having access to a world model is not a realistic option. Furthermore, the minimal amount of computation involved means reactive systems are able to respond in a timely manner to rapidly changing dynamics.

A middle ground between deliberative and reactive philosophies is found in *hybrid* control, exemplified by three-layered architectures (7,8). In this approach, a single controller includes both reactive and deliberative components. The reactive part of the controller handles low-level control issues requiring fast response time, such as local obstacle avoidance. The deliberative part of the controller handles high-level issues on a longer time-scale, such as global path planning. A necessary third component of hybrid controllers is a middle layer that interfaces the reactive and deliberative components.

Three-layered architectures aim to harness the best of reactive controllers in the form of dynamic and time-responsive control and the best of deliberative controllers in the form of globally efficient actions over a long time-scale. However, there are complex issues involved in interfacing these fundamentally differing components and the manner in which their functionality should be partitioned is not yet well understood.

Behavior-based control, described in detail in the next section, offers an alternative to hybrid control. It can also include both deliberative and reactive components, but unlike hybrid control, it is composed of a set of independent modular components that are executed in parallel (1,2).

The presented spectrum of control approaches is continuous and a precise classification of a specific controller on the continuum may be difficult. The distinction between deliberative and reactive control, and hybrid and behavior-based control is often a matter of degree, based on the amount of computation performed and the response time of the system to relevant changes in the world. In a specific domain, the choice of controller is dependent on many factors, including how responsive the robot must be to changes in the world, how accessible a world model is, and what level of efficiency or optimality is required.

1.2 Behavior-Based Control

The control methodology we focus on in this chapter is behavior-based (BB) control. The BB approach to robot control must not be classified as strictly deliberative or reactive, as it can, and in many cases is, both. However, BB control is most closely identified (often incorrectly so) with the reactive side of the control spectrum, because primary importance is placed on maintaining a tight, real-time coupling between sensing and action (7,8).

Fundamentally, a behavior-based controller is composed of a set of modular components, called behaviors, which are executed in parallel. A *behavior* is a control law that clusters a set of constraints in order to achieve and maintain a goal (1,2). Each behavior receives

inputs from sensors and/or other behaviors and provides outputs to the robot's actuators or to other behaviors. For example, an obstacle avoidance behavior might send a command to the robot's wheels to turn left or right if the robot sensors detect the robot is moving directly toward an obstacle. There is no centralized world representation or state in a BB system. Instead, individual behaviors and networks of behaviors maintain any models or state information.

Many different behaviors may independently receive input from the same sensors and output action commands to the same actuators. The issue of choosing a particular action given inputs from potentially multiple sensors and behaviors is called *action selection* (12). One of well-known mechanism for action selection is the use of a predefined behavior hierarchy, as in the Subsumption Architecture (10), in which commands from the highest-ranking active behavior are sent to the actuator and all others are ignored. (Note, however, that the Subsumption Architecture has most commonly been used in the context of reactive and not BB systems.) Numerous principled as well as ad hoc methods for addressing the action selection problem have been developed and demonstrated on robotic systems. These include varieties of command fusion (13) and spreading of activation (14), among many others. For a comprehensive survey on action selection mechanisms, see (15).

BB systems are varied, but there are two fundamental tenets all BB systems inherently adhere to: 1) the robot is embodied and 2) the robot is situated. A robot is *embodied* in the sense that it has a physical body and its behavior is limited by physical realities,

uncertainties, and consequences of its actions, all of which may be hard to predict or simulate. A robot is *situated* in the sense that it is immersed in the real world and acts directly on the sensory information received from that world, not on abstract or processed representations of the world.

BB control makes no assumptions on the availability of a complete world model; therefore, it is uncommon for a BB controller to perform extensive computation or reasoning relying on such a model. Instead, BB controllers maintain a tight coupling of sensing and action, allowing them to act in a timely manner in response to dynamic and fast-changing worlds. However, BB systems have also demonstrated elegant use of distributed representations enabling robot mapping and task learning (16,17,18).

This section has discussed approaches and philosophies to the control of a single robot, with a focus on the BB approach. In the next section, the scope is expanded to consider the control of a coordinated group of multiple robots.

2. From Single Robot Control to Multi-Robot Control

In this section we discuss the advantages and additional issues involved in the control of multi-robot systems (MRS) as compared to the single-robot systems (SRS) discussed in the previous section. An *MRS* is a system composed of multiple, interacting robots. The

study of MRS has received increased attention in recent years. This is not surprising, as continually improving robustness, availability, and cost-effectiveness of robotics technology has made the deployment of MRS consisting of increasingly larger numbers of robots possible. With the growing interest in MRS comes the expectation that, at least in some important respects, multiple robots will be superior to a single robot in achieving a given task. In this section we outline the benefits of a MRS over a SRS and introduce issues involved in MRS control and how they are similar and different to those of SRS control.

This chapter is focused on distributed MRS in which each robot operates independently under local sensing and control. *Distributed MRS* stand in contrast to *centralized MRS*, in which each robot's actions are not completely determined locally, as they may be determined by an outside entity, such as another robot or by any type of external command. In distributed MRS, each robot must make its own control decisions based only on limited, local, and noisy sensor information. We limit our consideration in this chapter to distributed MRS because they are the most appropriate for study with regard to systems that are scalable and capable of performing in uncertain and unstructured real-world environments where uncertainties are inherent in the sensing and action of each robot. Furthermore, this chapter is centered on achieving system-level coordination in a distributed BB MRS. Strictly speaking, the issues in a centralized MRS are more akin to a scheduling or optimal assignment and less of a problem of coordination in a distributed system.

2.1 Advantages and Challenges of Multi-Robot Systems

Potential advantages of MRS over SRS include a reduction in total system cost by utilizing multiple simple and cheap robots as opposed to a single complex and expensive robot. Also, multiple robots can increase system flexibility and robustness by taking advantage of inherent parallelism and redundancy. Furthermore, the inherent complexity of some task environments may require the use of multiple robots, as the necessary capabilities or resource requirements are too substantial to be met by a single robot.

However, the utilization of MRS poses potential disadvantages and additional challenges that must be addressed if MRS are to present a viable and effective alternative to SRS. A poorly designed MRS, with individual robots working toward opposing goals, can be less effective than a carefully designed SRS. A paramount challenge in the design of effective MRS is managing the complexity introduced by multiple, interacting robots. As such, in most cases just taking a suitable SRS solution and scaling it up to multiple robots is not adequate.

2.2 Necessity of Coordination in Multi-Robot Systems

In order to maximize the effectiveness of a MRS, the robots' actions must be spatio-temporally coordinated and directed towards the achievement of a given system-level task or goal. Just having robots interact is not sufficient in itself to produce interesting or practical system-level coordinated behavior. The design of MRS can be quite challenging because unexpected system-level behaviors may emerge due to unanticipated ramifications of the robots' local interactions. In order for the interacting robots to produce coherent task-directed behavior, there must be some overarching coordination mechanism that spatio-temporally organizes the interactions in a manner appropriate for the task.

The design of such coordination mechanisms can be difficult; nonetheless, many elegant handcrafted distributed MRS have been demonstrated, both in simulation and on physical robots (19,20,21). The methods by which these systems have achieved task-directed coordination are diverse and the possibilities are seemingly limited only by the ingenuity of the designer. From a few robots performing a manipulation task (22,23), to tens of robots exploring a large indoor area (24,25), to potentially thousands of ecosystem monitoring nanorobots (26,27), as the number of robots in the system increases, so does the necessity and importance of coordination. The next section examines mechanisms by which system-level coordination can be successfully achieved in a MRS.

3. From Local Interactions to Global Coordination

Given the importance of coordination in a MRS, we now address the issue of how to organize the robots' local interactions in a coherent manner in order to achieve system-level coordination. There are many mechanisms by which the interactions can be organized. We classify them into three broad and often overlapping classes: interaction through the environment, interaction through sensing, and interaction through communication. These classes are not mutually exclusive because MRS can, and often do, simultaneously utilize mechanisms from any or all of these classes to achieve system-level coordinated behavior.

In the following sections we describe each of these interaction classes in detail. Through the discussion of empirical case studies we demonstrate how each type of interaction can be used to achieve system-level coordination in a MRS.

3.1 Interaction Through the Environment

The first mechanism for interaction is through the robots' shared environment. This form of interaction is *indirect* in that it consists of no explicit communication or physical interaction between robots. Instead, the environment itself is used as a medium of indirect communication. This is a powerful approach that can be utilized by very simple robots with no capability for complex reasoning or direct communication.

An example of interaction through the environment is demonstrated in *stigmergy*, a form of interaction employed by a variety of insect societies. Originally introduced in the biological sciences to explain some aspects of social insect nest-building behavior, stigmergy is defined as the process by which the coordination of tasks and the regulation of construction do not depend directly on the workers, but on the constructions themselves (28). This concept was first used to describe the nest-building behavior of termites and ants (29). It was shown that coordination of building activity in a termite colony was not inherent in the termites themselves. Instead, the coordination mechanisms were found to be regulated by the task environment, in that case the growing nest structure. A location on the growing nest stimulates a termite's building behavior, thereby transforming the local nest structure, which in turn stimulates additional building behavior of the same or another termite.

Through the careful design of robot sensing, actuation, and control features, it is possible to utilize the concept of stigmergy in task-directed MRS. This powerful mechanism of coordination is attractive as it typically requires minimal capabilities of the individual robots. The robots do not require direct communication, unique recognition of other robots or even distinguishing other robots from miscellaneous objects in the environment, or the performance of computationally intensive reasoning or planning.

Stigmergy, and more generally interaction through the environment, has been successfully demonstrated as a mechanism to coordinate robot actions in a number of

MRS. It has been demonstrated in an object manipulation domain (30) in which a large box was transported to a goal location through the coordinated pushing actions of a group of robots. There was no globally agreed upon plan as to how or over what trajectory the box should be moved; however, each robot could indirectly sense the pushing actions of other robots through the motions of the box itself. Through simple rules, each robot decided whether to push the box or move to another location based on the motions of the box itself. As a large enough number of robots pushed in compatible directions, the box moved, which in turn encouraged other robots to push in the same direction.

Other examples of the use of stigmergy in MRS include distributed construction in which a given structure was built in a specified construction sequence (31). The individual robots were not capable of explicit communication and executed simple rule-based controllers in which local sensory information was directly mapped to construction actions. The construction actions of one robot altered the environment, and therefore the subsequent sensory information available for it and all other robots. This new sensory information then activated future construction actions. In the following subsection we discuss in detail how the concept of stigmergy was utilized in a MRS object clustering task domain (28).

3.2 Interaction Through the Environment Case Study: Object Clustering

We now describe an empirical case study in an object clustering task domain for which interaction through the environment was used to achieve system-level coordination. The clustering task domain requires a group of objects, originally uniformly positioned in an enclosed environment, to be re-positioned by a group of robots into a single dense cluster of objects. There is no *a priori* target location for the cluster in the environment. Rather, the position of the cluster is to determined dynamically at the time of task execution.

The particular approach to the object clustering task we describe here is from work presented in (28). There, the robots performing the task were extremely simple, capable only of picking up and transporting and dropping a single object at a time. The robots had very limited local sensing and no explicitly communication, memory of past actions, or recognition of other robots. Even with these highly limited capabilities, a homogeneous MRS composed of such robots was shown to be capable of successfully and robustly performing the object clustering task.

The robots in this task domain were able to coherently achieve system-level coordination in the formation of a single cluster. The mechanism by which they achieved coordination was an example of interaction through the environment. The robots communicated through their individual placement of objects over time, thus modifying the task environment, and thereby indirectly influencing the future object-placement behaviors of other robots and themselves. The location of the final cluster was not determined through explicit communication, negotiation, or planning on the part of the robots. Rather, it was determined through a symmetry break in the initially uniform distribution of objects.

Once a small cluster began to form, it was likely to grow larger. During the early stages of task execution, several clusters were likely to be formed. However, over time, a single large cluster resulted.

The robots in this work were designed in a manner that carefully exploited the physical dynamics of interaction between the robots and their environment. Their hardware and rules were tuned so as to be probabilistically more likely to pick up an object that is not physically proximate to other objects (thus conserving clusters), to not drop objects near boundaries (thus avoiding hard-to-find objects), and to be probabilistically more likely to deposit an object near other objects (thereby building up clusters). Together, these properties resulted in a form of positive feedback in which the larger a cluster of objects became, the more likely it was to grow even larger.

Similar approaches employing stigmergy were also demonstrated in the physical segregation and sorting of a collection of object classes. Additional studies with physical robots have been conducted and, by making various changes in the robots and the task environment, it has been demonstrated that one can influence the location of the final cluster by initializing the initial distribution of objects in a non-uniform manner (28).

Given this specific example of system-level coordination achieved through the use of interaction through the environment, in the following subsection we move on to the next method of organizing the robot's interactions: interaction through sensing.

3.3 Interaction Through Sensing

The second mechanism for interaction among robots is through sensing. As described in (19), interaction through sensing ‘refers to local interactions that occur between robots as a result of sensing one another, but without explicit communication.’ As with interaction through the environment, interaction through sensing is also *indirect* as there is no explicit communication between robots; however, it requires each robot to be able to distinguish other robots from miscellaneous objects in the environment. In some instances, each robot may be required to uniquely identify all other robots, or classes of other robots. In other instances, it may only be necessary to simply distinguish robots from other objects in the environment.

Interaction through sensing can be used by a robot to model the behavior of other robots or to determine what another robot is doing in order to make decisions and respond appropriately. For example, flocking birds use sensing to monitor the actions of other birds in their vicinity to make local corrections to their own motion. It has been shown that effective flocking results from quite simple local rules followed by each bird responding to the direction and speed of the local neighbors (32).

In the follow subsection we describe a case study in a formation marching domain in which interaction through sensing is used to achieve coordinated group behavior. Other

domains in which interaction through sensing has been utilized in MRS include flocking (33), in which each robot adjusts its motions according to the motions of locally observed robots. Through this process, the robots can be made to move as a coherent flock through an obstacle-laden and dynamic environment. Interaction through sensing has also been demonstrated in an adaptive division of labor domain (34). In that domain, each robot dynamically changes the task it is executing based on the observed actions of other robots and the observed availability of tasks in the environment. Through this process, the group of robots coherently divides the labor of the robots appropriately across a set of available tasks.

3.4 Interaction Through Sensing Case Study: Formation Marching

In this section we describe an empirical case study of a formation marching task domain for which interaction through sensing was used to achieve system-level coordination. The formation marching task domain requires a group of robots to achieve and maintain relative positions to one another as the group moves through the environment in a global formation. Each robot in the MRS operates under local sensing and control and is not aware of global information such as all other robot's positions and headings. In some environments, the formation may need to be perturbed in order for the group to move through a constrained passage or around obstacles. In such cases, the formation needs to correctly re-align after the perturbation.

The approach to formation marching described here was presented in (35). The general idea of the approach is that every robot in the MRS positioned itself relative to a designated neighbor robot. This neighbor robot, in turn, positioned itself relative to its own designated neighbor robot. As all robots are only concerned with their relative positions with respect to their neighbor robot, no robot is aware of, or needs to be aware of, the global positions and headings of all robots in the formation. Each robot only needs to be capable of determining the distance and heading to its neighbor. The global geometry of the formation was then determined through the defined chain of neighbors.

A “leader” robot has no neighbors and independently determines the speed and heading of the entire formation. Therefore, as the leader robot moves forward, the robot(s) that had the leader as their neighbor also move forward. This forward motion propagates down the chain of designated neighbors, causing the entire formation to move.

The formation could be dynamically changed by altering the structure of the local neighbor relationships. For example, if the desired formation is a line, each robot may be designated a neighbor robot to its left or right for which it desires to stay next to in order to maintain the line formation. If a cue is given to all robots to change to a diamond formation, each robot may follow a new neighbor at a different relative position and the line formation would then be dynamically changed to a diamond.

In the following subsection we move to the next method of organizing the robot's interactions: interaction through communication.

3.5 Interaction Through Communication

The third mechanism for interaction among robots is through explicit communication. Unlike the first two forms of interaction, described above, which were indirect, in interaction through communication robots may communicate with others directly. Such *robot-directed* communication can be used to request information or action from other robots or to respond to received requests.

Communication in physical robotics is not free or reliable and can be constrained by limited bandwidth and range, and unpredictable interference. When utilizing it, one must consider how and toward what end it is used. In some domains, such as the Internet, communication is reliable and of unlimited range; however, in physical robot systems, communication range and reliability are important factors in system design (2,36).

There are many types of communication. Communication could be direct from one robot to another, direct from one robot to a class of other robots, or broadcast from one robot to all others. Furthermore, the communication protocol can range from simple protocol-less schemes to a complex negotiation-based and communication-intensive schemes. The

information encoded in a communication may be state information contained by the communicating robot, a command to one or more other robots, or a request for additional information from other robots, etc.

Communications may be task-related rather than robot-directed, in which case it is made available to all (or a subset) of the robots in the MRS. A common task-related communication scheme is *publish/subscribe messaging*. In publish/subscribe messaging, subscribing robots request to receive certain categories of messages, and publishing robots supply messages to all appropriate subscribers.

In the next subsection, we describe a case study of the effective use of interaction through communication.

3.6 Interaction Through Communication Case Study: Multiple Target Tracking

The case study on interaction through communication in this section is focused on the use of explicit communication in a multi-target tracking task as discussed in (37). In multi-target tracking, the goal is to have a set of robots with limited sensing ranges position and orient themselves such they are able to acquire and track multiple objects moving through their environment. The locations, trajectories, and number of targets are not known *a priori*. These difficulties are compounded in a distributed MRS, where the system must determine which robot(s) should monitor which target(s). Robots redundantly tracking

the same target may be wasting resources and letting another target remain untracked. In this domain, explicit communication between the robots has been shown to be capable of effectively achieving system-level coordination.

In the implementation described in (37), each robot had a limited sensing and communication range. Communication was used by each robot to transmit the position and velocities of all targets within its sensing range to all other robots within its communication range. This simple communication scheme involved no handshaking or negotiation.

Each robot was constantly evaluating the importance of its current tracking activities and possible changes in position that could increase the importance of its tracking activities. Communication was used to allow each robot to keep a local map of target movements within the robot's communication range but outside its sensing range. As a result, the group as a whole effectively tracked a maximum number of targets with a minimum number of available robots.

This demonstration of the use of interaction through communication concludes the discussion of MRS coordination mechanisms. As was mentioned above, any given MRS is likely to use any or all of the three mechanisms in varying degrees to achieve system-level coordination. Through an improved understanding of each of these mechanisms of coordination, one is better positioned to design a MRS utilizing the most appropriate combination of mechanisms for achieving a given task. In the next section we provide a

