72. Collective Manipulation and Construction

Lynne Parker

Many practical applications can make use of robot collectives that can manipulate objects and construct structures. Examples include applications in warehousing, truck loading and unloading, transporting large objects in industrial environments, and assembly of large-scale structures. Creating such systems, however, can be challenging. When collective robots work together to manipulate physical objects in the environment, their interactions necessarily become more tightly coupled. This need for tight coupling can lead to important control challenges, since actions by some robots can directly interfere with those of other robots. This chapter explores techniques that have been developed to enable robot swarms to effectively manipulate and construct objects in the environment. The focus in this chapter is on decentralized manipulation and construction techniques that would likely scale to large robot swarms (at least 10 robots), rather than approaches aimed primarily at smaller teams that attempt the same objectives. This chapter first discusses the swarm task of object transportation; in this domain, the objective is for robots to collectively move objects through the environment to a goal destination. The chapter then discusses object clustering and sorting, which requires objects in the environment to be aggregated at one or more locations in the environment. The final task discussed is that of collective construction and wall building, in which robots work together to build a prespecified structure. While these different tasks vary in their specific objectives for collective manipulation, they also have several commonalities. This chapter explores the state of the art in this area.

72.1 **Object Transportation** 1395
72.1.1 Transport by Pushing 1396
72.1.2 Transport by Grasping 1397
72.1.3 Transport by Caging 1400
72.2 **Object Sorting and Clustering** 1401
72.3 **Collective Construction and Wall Building** 1402
72.4 **Conclusions** 1404
References 1404

72.1 Object Transportation

Some of the earliest work in swarm robotics was aimed at the object transportation task [72.1–6], which requires a swarm of robots to move an object from its current position in the environment to some goal destination. The primary benefit of using collective robots for this task is that the individual robots can combine forces to move objects that are too heavy for individual robots working alone or in small teams. However, the task is not without its challenges; it is nontrivial to design decentralized robot control algorithms that can effectively coordinate robot team members during object transportation. A further complication is that the interaction dynamics of the robots with the object can be sensitive to certain object geometries [72.7, 8] and object rotations during transportation [72.8], thus exacerbating the control problem.

There are many ways to compare and contrast alternative distributed techniques to collective object transport. The most common distinctions are:

- Local knowledge only versus some required global knowledge (e.g., of team size, state, position).
- Homogeneous swarms versus heterogeneous swarms (e.g., teams with leaders and followers).
- Manual controller design versus autonomously learned control.

- 2-D (two-dimensional) vs. 3-D (three-dimensional) environments.
- Obstacle-free environments versus cluttered environments.
- Static environments versus dynamic environments.
- Dependent on fully functioning robots versus systems robust to error.

Alternatively, we can compare transportation techniques by focusing on the specific manipulation technique employed. The manipulation techniques used for collective object transportation can be grouped into three primary methods [72.9]: pushing, grasping, and caging. The pushing approach requires contact between each robot and the object, in order to impart force in the goal direction; however, the robots are not physically connected with the object. In the grasping approach, each robot in the swarm is physically attached to the object being transported. Finally, the caging approach involves robots encircling the object so that the object moves in the desired direction, even without the constant contact of all the robots with the object.

This section outlines some of the key techniques developed to address this object transportation task, organized according to these three main techniques.

72.1.1 Transport by Pushing

A canonical task often used as a testbed in distributed robotics is the box pushing task. The number, size, or weight of the boxes can be varied to explore different types of multirobot cooperation. This task typically involves robots first locating a box, positioning themselves at the box, and then moving the box cooperatively toward a goal position. Typically, this task is explored in 2-D. The domain of box pushing is also popular because it has relevance to several real-world applications [72.10], including warehouse stocking, truck loading and unloading, transporting large objects in industrial environments, and assembling of large-scale structures.

The pushing technique was first demonstrated in the early work of *Kube* and *Zhang* [72.1], inspired by the cooperative transport behavior in ants [72.7]. They proposed a behavior-based approach that combined behaviors for seeking out the object (illuminated by a light), avoiding collisions, following other robots, and motion control. An additional behavior to detect stagnation was used to ensure that the collective did not work consistently against each other. In this approach, all robots acted similarly; there was no concept of a leader and followers. While some of the robots in the swarm might not contribute to the pushing task due to poor alignment or positioning along the nondominant pushing direction, Kube and Zhang showed that careful design of these behaviors enabled the robot swarm to distribute along the boundary of the object and push it. Figure 72.1 shows five robots cooperatively pushing a lighted box.

Other researchers have explored different aspects of box pushing in multirobot systems. While much of this early work involved demonstrations of smaller robot teams, many of these techniques could theoretically scale to larger numbers of robots. Task allocation and action selection are often demonstrated using collective box pushing experiments; examples of this work include that of *Parker* [72.11, 12], who illustrated aspects of adaptive task allocation and learning; *Gerkey* and *Mataric* [72.13], who present a publish/subscribe dynamic task allocation method; and *Yamada* and *Saito* [72.14], who develop a behavior-based action selection technique that does not require any communication.

Other work using box pushing as an implementation domain for multirobot studies includes *Donald* et al. [72.15], who illustrates concepts of information invariance and the interchangeability of sensing, communication, and control; *Simmons* et al. [72.16], who demonstrate the feasibility of cooperative control for building planetary habitats, *Brown* and *Jennings* [72.17], and *Böhringer* et al. [72.18], who explored notions of strong cooperation without communication in pusher/steerer models, *Rus* et al. [72.19], who

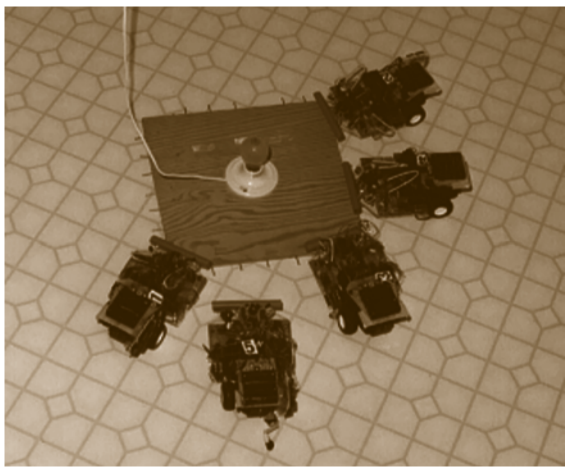

Fig. 72.1 Demonstration of five robots collectively pushing a lighted box (after [72.7])

studied different cooperative manipulation protocols in robot teams that make use of different combinations of state, sensing, and communication, and *Jones* and *Mataric* [72.20], who developed general methods for automatically synthesizing controllers for multirobot systems.

Most of this existing work in box pushing has focused, not on box pushing as the end objective, but rather on using box pushing for demonstrating various techniques for multirobot control. However, for studies whose primary objective is to generate robust cooperative transport techniques, work has more commonly focused on manipulation techniques involving grasping and caging, rather than pushing, since grasping and caging provide more controllability by the robot team.

72.1.2 Transport by Grasping

Grasping approaches for object transportation in swarm robotics typically make use of form closure and force closure properties [72.21]. In *form closure*, the object motion is constrained via frictionless contact constraints; in *force closure*, frictional contact forces exerted by the robots prevent unwanted motions of the manipulated object. The earliest work representing the grasping technique is that of *Wang* et al. [72.4]. This approach uses form closure, along with a behavior-based control approach that is similar to the early swarm robot pushing technique of *Kube* and *Zhang* [72.1]. The technique of Wang et al. called BeRoSH (for Behavior-based Multiple Robot System with Host for Object Manipulation), incorporates behaviors for pushing, maintaining contact, moving, and avoiding objects. In this approach, the goal pose of the object is provided directly to each robot from an external source (i. e., the Host); otherwise, the robots work independently according to their designed behaviors. As a collective, the swarm exhibits form closure. Wang et al. showed that this form closure technique can successfully transport an object to its desired goal pose from a variety of different starting locations.

Another early work using the force closure grasping technique is that of *Stilwell* and *Bay* [72.2] and *Johnson* and *Bay* [72.3]. They developed distributed leader–follower techniques that enable swarms of tank-like robots to transport pallets collectively while maintaining a level height of the pallet during transportation (Fig. 72.2). In their approaches, a pallet sits atop several tank-like robots; the weight of the pallet creates a coupling with the robots that could be viewed similar to a grasp. To transport the pallet, one vehicle is designated as the leader. This leader then perturbs the dynamics of the system to move the swarm in the desired direction, and with the desired pallet height. The remaining robots in the swarm react to the perturbations to stabilize the forces in the system. The system is fully distributed, and requires robots to only use local force information to achieve the collective motion. The individual robots do not require knowledge of the pallet mass or inertia, the size of the collective, or the robot positions relative to the pallet's center of gravity. They showed the control stability of their approach for this application, even in the presence of inaccurate sensor data.

A related approach is that of *Kosuge* and *Oosumi* [72.5], who also used a decentralized leader–follower approach for multiple holonomic robots grasping and moving an object, in a manner similar to that of [72.2]. Their approach defines a compliant motion control algorithm for each velocity-controlled robot. The main difference of this work compared to [72.2] is that the control algorithm specifies the desired internal force as part of the coordination algorithm. This approach was validated in simulation for robots carrying an aluminum steel pipe.

Another related approach is that of *Miyata* et al. [72.6], who addressed the need for nonholonomic vehicles to regrasp the object during transport. Their approach includes a hybrid system that makes use of both centralized and decentralized planners. The centralized planner develops an approximate motion plan for the object, along with a regrasping plan at low resolution; the decentralized planner precisely estimates object motion and robot control at a much higher resolution.

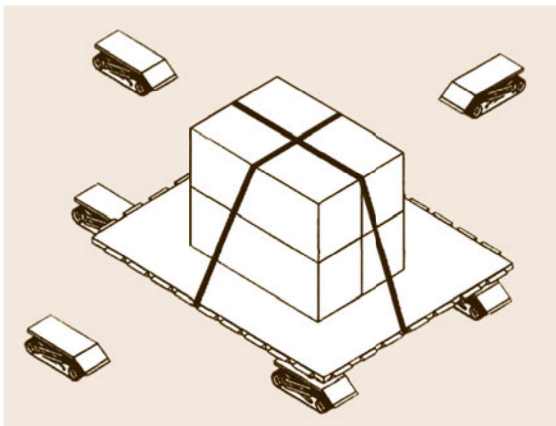

Fig. 72.2 Cooperative transport of a pallet using tank-like robots (after [72.2])

They demonstrated the effectiveness of this approach in simulation.

Sugar and *Kumar* [72.22] developed distributed control algorithms enabling robots with manipulators to grasp and cooperatively transport a box. In this work, a novel manipulator design enables the locomotion control to be decoupled from the manipulation control. Only a small number of the team members need to be equipped with actively controlled end effectors. This approach was shown to be robust to positioning errors related to the misalignment between the two platforms and errors in the measurement of the box size.

Cooperative stick pulling [72.23, 24] was explored by *Ijspeert* et al.; this task requires robots to pull sticks out of the ground (Fig. 72.3). The robot controllers are behavior-based, and include actions such as looking for sticks, detecting sticks, gripping sticks, obstacle avoidance, and stick release. Experiments show that the dynamics are dependent on the ratio between the number of robots and sticks; that collaboration can increase superlinearly with certain team sizes; that heterogeneity in the robots can increase the collaboration rate in certain circumstances; and that a simple signalling scheme can increase the effectiveness of the collaboration for certain team sizes. A main objective of this research was to explore the effectiveness of various modeling techniques for group behavior. These modeling techniques are discussed in more detail in a separate chapter.

The SWARM-BOTS project is a more recent example of the use of grasping for collective transport; it also makes use of self-assembly as a novel approach for achieving distributed transport. In this work [72.25], *s-bot* robots are developed that have grippers enabling the robots to create physical links with other s-bots or objects, thus creating assemblies of robots. These assemblies can then work together for navigation across rough terrain, or to collectively transport objects. The s-bots are cylindrical, with a flexible arm and toothed gripper that can connect one s-bot to another (Fig. 72.4).

The decentralized control of the SWARM-BOT robots is learned using evolutionary techniques in simulation, then ported to the physical robots. The learned s-bot control [72.26] consists of an assembly module, which is a neural network that controls the robot prior to connection, and a transport module, which is a neural network that enables the s-bot to move the object toward the goal after a grasp connection is made. The self-assembly process involves the use of a red-colored seed object, to which other s-bots are attracted. S-bots initially light themselves with a blue ring, and then are attracted to the red color, while being repulsed by the blue color. Once robots make a connection, they color themselves red.

The interaction of these attractive and repulsive forces across the s-bots enables the robots to self-assemble into various connection patterns. Once the s-bots have self-assembled, they use the transport module to align toward a light source, which indicates the target position. The s-bots then apply pushing and pulling motions to transport the object to the destination. Similar to the approach of *Kube* and *Zhang* [72.1], the s-bots also check for stagnation and execute a recovery move when needed. The authors demonstrate [72.8] how the evolutionary learning approach allows the collective

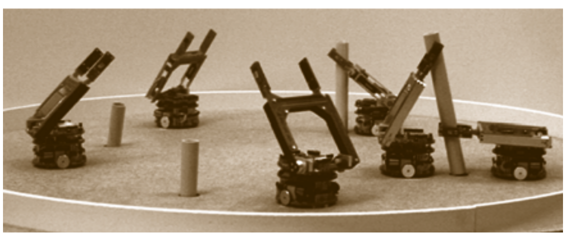

Fig. 72.3 Stick pulling experiment using robot collectives (after [72.23])

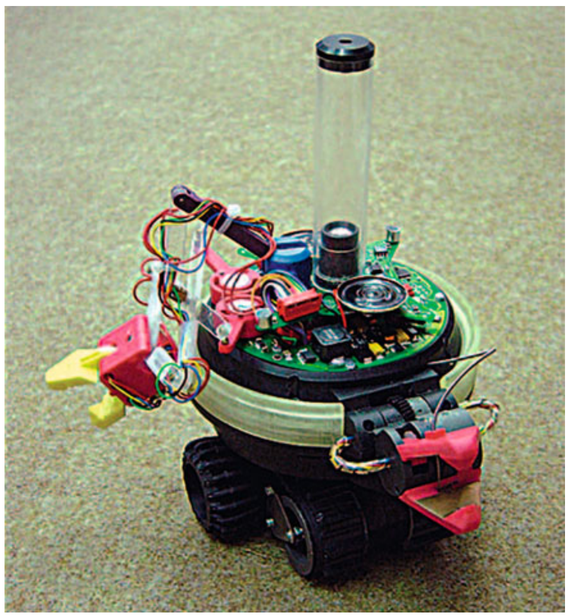

Fig. 72.4 An s-bot, developed as part of the SWARM-BOTS project (after [72.25])

to successfully deal with different object geometries, adapt to changes in target location, and scale to larger team sizes.

This technique for collective transport using self-assembly was demonstrated [72.25] in an interesting application of object transport, in which 20 s-bots self-assembled into four chains in order to pull a child across the floor (Fig. 72.5). In this experiment, the user specifies the number of assembled chains, the distribution of the s-bots into the chains, the global localization of the child, and the global action timing. The s-bots then autonomously form the chains using self-assembly and execute the pull.

Several additional interesting phenomena regarding collective transport were discovered in related studies with the SWARM-BOTS. *Nouyan* et al. [72.27] showed that the different collective tasks of path formation, self-assembly, and group transport can be solved in a single system using a homogeneous robot team. They further introduce the notion of *chains with cycle directional patterns*, which facilitate swarm exploration in unknown environments, and assist in establishing paths between the object and goal. The paths established by the robot-generated chains mimic pheromone trails in ants. In [72.28], *Groß* and *Dorigo* determined that, while robots that behave as if they are solitary robots can still collectively move objects, robots that learn transport behaviors in a group can achieve a better performance. In [72.29], *Campo* et al. showed that the SWARM-BOTS robots could effectively transport objects even with only partial knowledge of the direction of the goal. They investigated four alternative control strategies, which vary in the degree to which the robots negotiate regarding the goal position during transport. Their results showed that negotiating throughout object transport can improve motion coordination. All of these works are based on inspiration from biological systems.

Fig. 72.5 SWARM-BOTS experiment in which s-bots self-assemble to pull a child across the floor (after [72.25])

The work of *Berman* et al. [72.31] is not only bio-inspired, but also seeks to directly model the group retrieval behavior in ants. Their studies examined the ants' roles during transport in order to define rules that govern the ants' actions. They further explored measurements of individual forces used by the ants to guide food to their nest. They found that the distributed ant transport behavior exhibits an initial disordered phase, which then transitions to a more highly coordinated phase of increased load speed. From these studies, a computational dynamic model of the ant behavior was designed and implemented in simulations, showing that the derived model matches ant behavior. Ultimately, this approach could be adapted for use on physical robot teams.

Once a robot collective has begun transporting an object, the question arises as to how new robots can join the group to help with the transport task. *Esposito* [72.30] addresses this challenge by adapting a grasp quality function from the multifingered hand literature. This approach assumes that robots know the object geometry, the total number of robots in the swarm, and the actuator limitation. Individual robot contact configurations are defined relative to the object center and object boundary. The objective is to find an optimal position for a new robot by optimizing across the grasping wrench space. A numerical algorithm was developed to address this problem, which incorporates the force closure criteria. This ap-

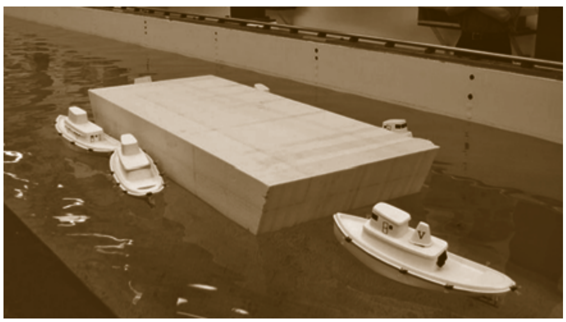

Fig. 72.6 Illustration of unmanned tugboats autonomously transporting a barge (after [72.30])

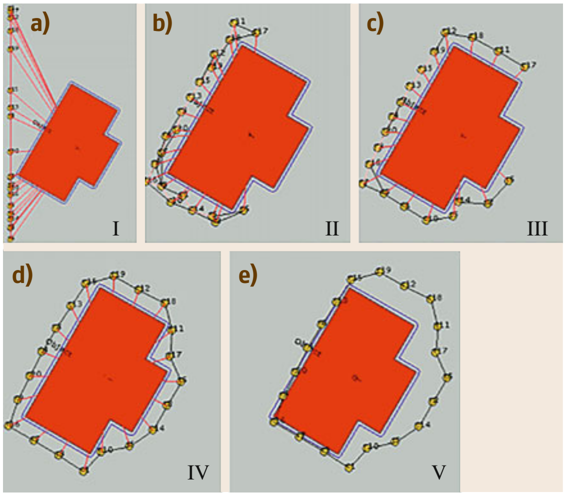

Fig. 72.7a–e Illustration in simulation of object closure by 20 robots (after [72.32])

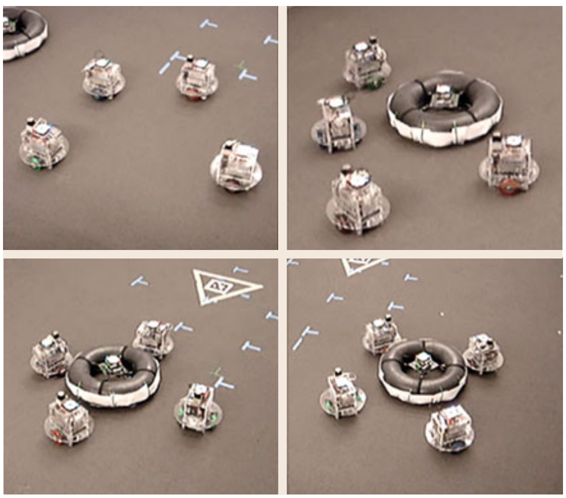

Fig. 72.8 Demonstration of the use of vector fields for collective transport via caging (after [72.33])

proach was demonstrated on unmanned tugboats collectively moving a barge, as illustrated in Fig. 72.6. In this demonstration, the robots are equipped with articulated magnetic attachments that allow them to grasp the barge. This approach is scalable to larger numbers of robots, with constant best case runtime, and median runtimes polynomial in the number of robots.

72.1.3 Transport by Caging

The caging approach simplifies the object manipulation task, compared to the grasping approach, by making use of the concept of *object closure* [72.34]. In object closure, a bounded movable area is defined for the object by the robots surrounding it. The benefit of this approach is that continuous contact between the object and the robots is not needed, thus making for simpler motion planning and control techniques, compared to grasping techniques based on the form or force closure. *Wang* and *Kumar* [72.32] developed this object-closure technique under the assumptions that the robots are circular and holonomic, the object is star-shaped, the robots know the number of robots in the collective, and can estimate the geometric properties of the object, along with the distance and orientation to other robots and the object. Their approach causes the robots to first approach the object independently, and then search for an *inescapable formation*, which is a configuration of the robots from which the object cannot escape. Finally, the robots execute a formation control strategy to guide the object to the goal destination. The object approach technique is based on potential fields, in which force vectors attract the robot toward the object and generally away from other robots. *Song* and *Kumar* [72.35] proved the stability of this potential field approach for collective transport. Robots search for proper configurations around the object by representing the problem as a path finding problem in configuration space. This work describes a necessary and sufficient condition for testing for object closure. Later work [72.36] presents a fast algorithm to test for object closure. Experiments with 20 robots validate the proposed approach (Fig. 72.7).

A further enhancement of this vector-based control strategy was developed in [72.33], which can account for inter-robot collisions. This latter strategy implements three primary behaviors – approach, surround, and transport. In this variant of the work, robots converge to a smooth boundary using control-theoretic techniques. This work was implemented on a collective of physical robots, as illustrated in Fig. 72.8.

72.2 Object Sorting and Clustering

Collective object sorting and clustering requires robot teams to sort objects from multiple classes, typically into separate physical clusters. There are different types of related tasks in this domain [72.37], including clustering, segregation, patch sorting, and annular sorting. Early discussions of this task in robot swarms were given by *Deneubourg* et al. [72.38], with the ideas inspired by similar behaviors in ant colonies. The objective is to achieve clustering and sorting behaviors without any need for hierarchical decision making, inter-robot communication, or global representations of the environment. Deneubourg et al. showed that stigmergy could be used to cluster scattered objects of a single type, or to sort objects of two different types. To achieve the sorting behavior, the robots sensed the local densities of the objects, as well as the type of object they were carrying. Clustering resulted from a similar mechanism operating on a single type of object. *Beckers* et al. [72.39] achieved clustering from even simpler robots and behaviors, via stigmergic threshold mechanisms.

Holland and *Melhuish* [72.37] explored the effect of stigmergy and self-organization in swarms of homogeneous physical robots. The robots are programmed with simple rule sets with no ability for spatial orientation or memory. The experiments show the ability of the robots to achieve effective sorting and clustering, as illustrated in Fig. 72.9. In this work, a variety of influences were explored, including boundary effects and the distance between objects when deposited. The authors concluded that the effectiveness of the developed sorting behaviors is critically dependent on the exploitation of real-world physics. An implication of this finding is that simulators must be used with care when exploring these behaviors.

Wang and *Zhang* [72.40, 41] explored similar aims, but focused on discovering a general approach to the sorting problem. They conjecture that the outcome of the sorting task is dependent primarily on the capabilities of the robots, rather than the initial configuration. This conjecture is validated in simulation experiments, as illustrated in Fig. 72.10.

Other work on this topic includes that of *Yang* and *Kamel* [72.42], who present research using three colonies of ants having different speed models. The approach is a two-step process. The first step is for clusterings to be visually formed on the plane by agents walking, picking up, or setting down objects according to a probabilistic model, which is based on *Deneubourg* et al. [72.38]. The second step is for clusters to be combined using a hypergraph model. Experiments were conducted in simulation to show the viability of the approach. The authors also discovered that having too many agents can lead to a deterioration in the swarm performance.

Martinoli and *Mondada* [72.43] implemented another bio-inspired approach to clustering, in which the robot behavior is similar to that of a Braitenberg vehicle. They also discovered that large numbers of robots can cause interference in this task, concluding that noncooperative task cannot always be improved with more robots.

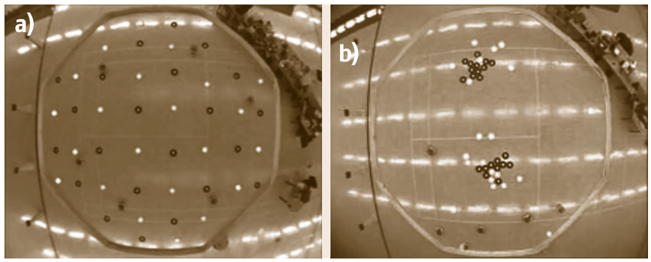

Fig. 72.9 Results of physical robot experiments in sorting. Panel **(a)** shows the starting configuration, while **(b)** shows the sorting results after 1.75 h (after [72.37])

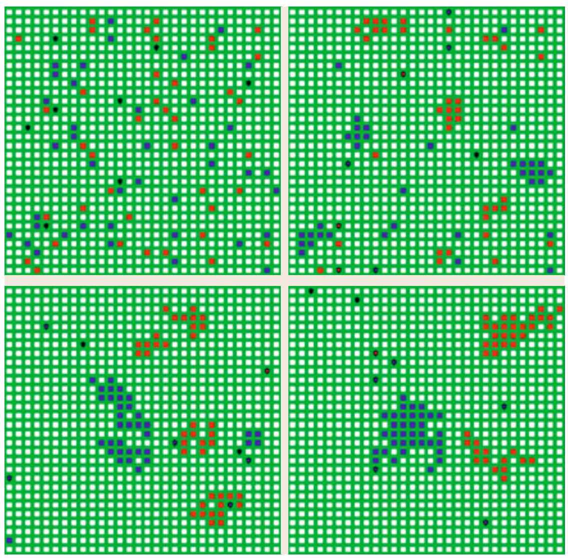

Fig. 72.10 Results of simulations of sorting tasks, with 8 robots and 40 objects of two types (after [72.40])

72.3 Collective Construction and Wall Building

The objective of the collective construction and wall building task is for robots to build structures of a specified form, in either 2-D or 3-D. This task is distinguished from self-reconfigurable robots, whose bodies themselves serve as the dynamic structure. This section is focused on the former situation, in which manipulation is required to create the desired structure. The argument in favor of this separation of mobility and structure is that, once formed, the structure does not need to move again, and thus the ability to move could serve as a liability [72.44]. Furthermore, robotic units that serve both as mobility and structure might not be effective as passive structural elements.

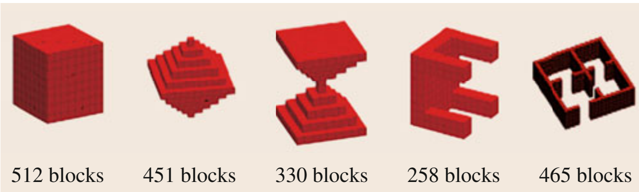

Fig. 72.11 Experiments for a variety of 3-D structures, built autonomously by a system of simple robots and blocks (after [72.44])

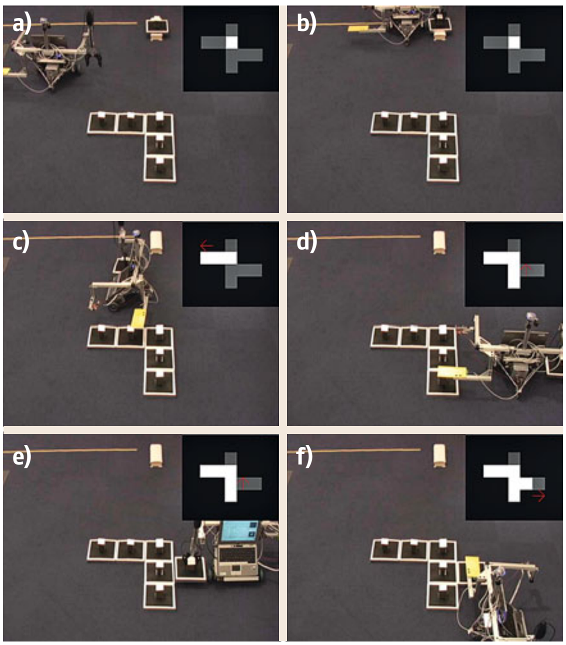

Fig. 72.12a–f Proof-of-principle experiments for 2-D construction, using a single robot (after [72.45])

Werfel et al. have extensively explored this topic, developing distributed algorithms that enable simplified robots to build structures based on provided blueprints, both in 2-D [72.45–47] and in 3-D [72.44]. In their 3-D approach, the system consists of idealized mobile robots that perform the construction, and smart blocks that serve as the passive structure. The robots' job is to provide the mobility, while the blocks' role is to identify places in the growing structure at which an additional block can be placed that is on the path toward obtaining the desired final structure. The goal of their work is to be able to deploy some number of robots and free blocks into a construction zone, along with a single block that serves as a seed for the structure, and then have the construction to proceed autonomously according to the provided blueprint of the desired structure.

Several simplifying assumptions are made in this work [72.44], such as the environment being weightless and the robots being free to move in any direction in 3-D, including along the surface of the structure under construction. This work does not address physical robot navigation and locomotion challenges, grasping challenges, etc.

In this approach, blocks are smart cubes; they can communicate with attached neighbors, they share a global coordinate system, and they can communicate with passing robots regarding the validity of block attachments to exposed faces. Once robots have transported a free block to the structure, they locate attachment points in one of three ways: random movement, systematic search, or gradient following. A significant contribution of this work is the development of the block algorithm that enables the blocks to specify

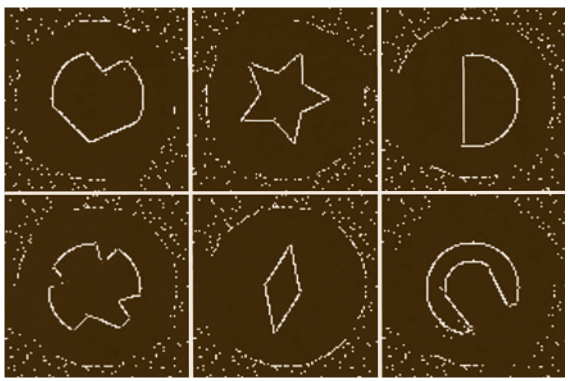

Fig. 72.13 Geometric structures built by a team of 30 robots, in simulation (after [72.48])

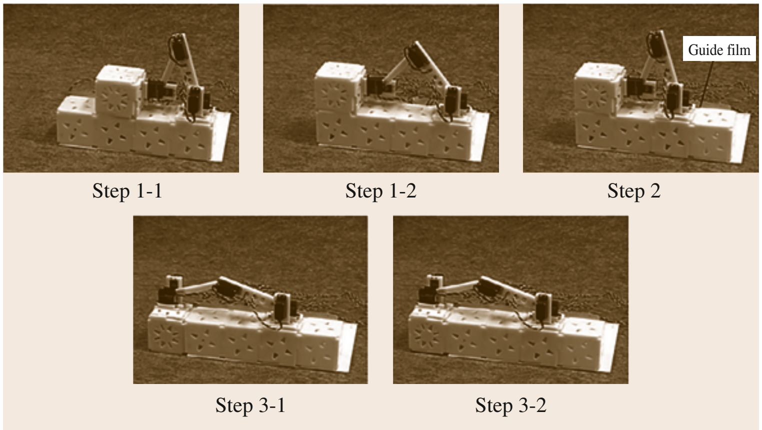

Fig. 72.14 Experiments with prototype hardware designed for multirobot construction tasks (after [72.49])

how to grow the developing structure with guarantees, and with only limited required communication. More specifically, the communication requirement between blocks scales linearly in the size of the structure, while no explicit communication between the mobile robots is needed.

Experiments using this approach have shown the ability of the system to build a variety of structures in simulation, as illustrated in Fig. 72.11. A proof-of-principle physical robot experiment using a single robot in the 2-D case [72.45] is illustrated in Fig. 72.12.

Werfel [72.48] also describes a system for arranging inert blocks into arbitrary shapes. The input to the robot system is a high-level geometric program, which is then translated by the robots into an appropriate arrangement of blocks using their programmed behaviors. The desired structure is communicated to the robots as a list of corners, the angles between corners, and whether the connection between corners is to be straight or curved. Robots are provided with behaviors such as *clearing*, *doneClearing*, *beCorner*, *collect*, *seal*, and *off*. Figure 72.13 shows some example structures built using this system in simulation.

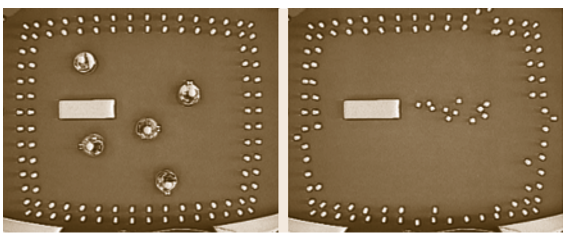

Fig. 72.15 Experimental trial demonstrating a swarm building a loose wall via a spatiotemporal varying template (after [72.50])

Hardware challenges of collective robot construction are addressed by *Terada* and *Murata* [72.49]. In this work, a hardware design is proposed that defines passive building blocks, along with an assembler robot that constructs structures with the robots. Figure 72.14 shows the prototype hardware completing an assembly task. In principle, multiple assembler robots could work together to create larger construction teams more closely aligned with the concept of swarm construction.

Other related work on the topic of collective construction includes the work of *Wawerla* et al. [72.51], in which robots use a behavior-based approach to build a linear wall using blocks equipped with either positive or negative Velcro, distinguished by block color. Their results show that adding 1 bit of state information to communicate the color of the last attached block provides a significant improvement in the collective performance. The work by *Stewart* and *Russell* [72.50, 52] proposes a distributed approach to building a loose wall structure with a robot swarm. The approach makes use of a spatiotemporal varying light-field template, which is generated by an organizer robot to help direct the actions of the builder robots. Builder robots deposit objects in locations indicated by the template.

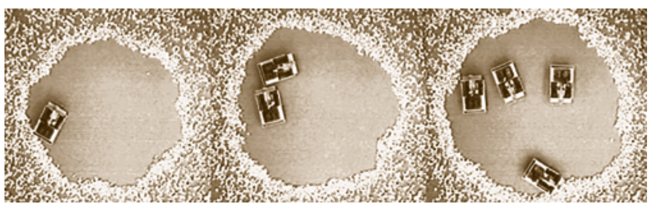

Fig. 72.16 Experiments of blind bulldozing for site clearing using physical robots (after [72.53])

Figure 72.15 shows the results from one of the experiments using this approach on physical robots.

Another type of construction is called *blind bulldozing*, which is inspired by a behavior observed in certain ant colonies. Rather than constructing by accumulating materials, this approach achieves construction by removing materials. This task has practical application in site clearing, such as would be needed for planetary exploration [72.54]. Early ideas of this concept were discussed by *Brooks* et al. [72.55], which argues for large numbers of small robots to be delivered to the lunar surface for site preparation. *Parker* et al. [72.53], further developed this idea by proposing robots using force sensors to clear an area by pushing material to the edges of the work site. In this approach, the robot system collective behavior is modeled in terms of how the nest grows over time. Stigmergy is used to control the construction process, in that the work achieved by each robot affects the other robots' behaviors through the environment. Figure 72.16 shows some experiments using this approach on physical robots. The authors argue that blind bulldozing is appropriate in applications where the cost, complexity, and reliability of the robots is a concern.

72.4 Conclusions

This chapter has surveyed some of the important techniques that have been developed for collective object transport and manipulation. While many advances have been made, there are still many open challenges that remain. Some open problems include: How to deal with faults in the robot team members during task execution; how to address construction in dynamic and cluttered environments; how to enable humans to interact with the robot swarms; how to extend more of the existing techniques to 3-D applications; how to design formal techniques for predicting and guaranteeing swarm behavior; how to realize larger scale systems on physical robots; and how to apply swarm techniques for manipulation and construction to practical applications.

References

72.1 C.R. Kube, H. Zhang: Collective robotics: From social insects to robots, Adapt. Behav. **2**(2), 189–218 (1993)

72.2 D. Stilwell, J.S. Bay: Toward the development of a material transport system using swarms of ant-like robots, IEEE Int. Conf. Robot. Autom. (1993) pp. 766–771

72.3 P.J. Johnson, J.S. Bay: Distributed control of simulated autonomous mobile robot collectives in payload transportation, Auton. Robot. **2**(1), 43–63 (1995)

72.4 Z. Wang, E. Nakano, T. Matsukawa: Realizing cooperative object manipulation using multiple behaviour-based robots, Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (1996) pp. 310–317

72.5 K. Kosuge, T. Oosumi: Decentralized control of multiple robots handling an object, Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (1996) pp. 318–323

72.6 N. Miyata, J. Ota, Y. Aiyama, J. Sasaki, T. Arai: Cooperative transport system with regrasping car-like mobile robots, Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (1997) pp. 1754–1761

72.7 C.R. Kube, E. Bonabeau: Cooperative transport by ants and robots, Robot. Auton. Syst. **30**(1), 85–101 (2000)

72.8 R. Groß, M. Dorigo: Towards group transport by swarms of robots, Int. J. Bio-Inspir. Comput. **1**(1), 1–13 (2009)

72.9 Y. Mohan, S.G. Ponnambalam: An extensive review of research in swarm robotics, World Congr. Nat. Biol. Inspir. Comput. 2009 (2009) pp. 140–145

72.10 D. Nardi, A. Farinelli, L. Iocchi: Multirobot systems: A classification focused on coordination, IEEE Trans. Syst. Man Cybern. B **34**(5), 2015–2028 (2004)

72.11 L.E. Parker: ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots, Proc. IEEE/RSJ/GI Int. Conf. Intell. Robot. Syst. (1994) pp. 776–783

72.12 L.E. Parker: Lifelong adaptation in heterogeneous teams: Response to continual variation in individual robot performance, Auton. Robot. **8**(3), 239–269 (2000)

72.13 B.P. Gerkey, M.J. Mataric: Sold! Auction methods for multi-robot coordination, IEEE Trans. Robot. Autom. **18**(5), 758–768 (2002)

72.14 S. Yamada, J. Saito: Adaptive action selection without explicit communication for multirobot box-pushing, IEEE Trans. Syst. Man Cybern. C **31**(3), 398–404 (2001)

72.15 B. Donald, J. Jennings, D. Rus: Analyzing teams of cooperating mobile robots, IEEE Int. Conf. Robot. Autom. (1994) pp. 1896–1903

72.16 R. Simmons, S. Singh, D. Hershberger, J. Ramos, T. Smith: First results in the coordination of hetero-

geneous robots for large-scale assembly, ISER 7th Int. Symp. Exp. Robot. (2000)
72.17 R.G. Brown, J.S. Jennings: A pusher/steerer model for strongly cooperative mobile robot manipulation, Proc. 1995 IEEE Int. Conf. Intell. Robot. Syst. (1995) pp. 562–568
72.18 K. Böhringer, R. Brown, B. Donald, J. Jennings, D. Rus: Distributed robotic manipulation: Experiments in minimalism, Lect. Notes Comput. Sci. **223**, 11–25 (1997)
72.19 D. Rus, B. Donald, J. Jennings: Moving furniture with teams of autonomous robots, Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (1995) pp. 235–242
72.20 C. Jones, M.J. Mataric: Automatic synthesis of communication-based coordinated multi-robot systems, Proc. IEEE/RJS Int. Conf. Intell. Robot. Syst. (2004) pp. 381–387
72.21 A. Bicchi: On the closure properties of robotic grasping, Int. J. Robot. Res. **14**(4), 319–334 (1995)
72.22 T.G. Sugar, V. Kumar: Control of cooperating mobile manipulators, IEEE Trans. Robot. Autom. **18**(1), 94–103 (2002)
72.23 A.J. Ijspeert, A. Martinoli, A. Billard, L.M. Gambardella: Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment, Auton. Robot. **11**(2), 149–171 (2001)
72.24 A. Martinoli, K. Easton, W. Agassounon: Modeling swarm robotic systems: A case study in collaborative distributed manipulation, Int. J. Robot. Res. **23**(4/5), 415–436 (2004)
72.25 F. Mondada, L.M. Gambardella, D. Floreano, S. Nolfi, J.L. Deneuborg, M. Dorigo: The cooperation of swarm-bots: Physical interactions in collective robotics, IEEE Robot. Autom. Mag. **12**(2), 21–28 (2005)
72.26 R. Groß, E. Tuci, M. Dorigo, M. Bonani, F. Mondada: Object transport by modular robots that self-assemble, IEEE Int. Conf. Robot. Autom. (2006) pp. 2558–2564
72.27 S. Nouyan, R. Groß, M. Bonani, F. Mondada, M. Dorigo: Group transport along a robot chain in a self-organised robot colony, Proc. 9th Int. Conf. Intell. Auton. Syst. (2006) pp. 433–442
72.28 R. Groß, M. Dorigo: Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling, Adapt. Behav. **16**(5), 285–305 (2008)
72.29 A. Campo, S. Nouyan, M. Birattari, R. Groß, M. Dorigo: Negotiation of goal direction for cooperative transport, Lect. Notes Comput. Sci. **4150**, 191–202 (2006)
72.30 J.M. Esposito: Distributed grasp synthesis for swarm manipulation with applications to autonomous tugboats, IEEE Int. Conf. Robot. Autom. (2008) pp. 1489–1494
72.31 S. Berman, Q. Lindsey, M.S. Sakar, V. Kumar, S.C. Pratt: Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems, Proc. IEEE **99**(9), 1470–1481 (2011)
72.32 Z. Wang, V. Kumar: Object closure and manipulation by multiple cooperating mobile robots, IEEE Int. Conf. Robot. Autom. (2002) pp. 394–399
72.33 J. Fink, N. Michael, V. Kumar: Composition of vector fields for multi-robot manipulation via caging, Robot. Sci. Syst. (2007) pp. 25–32
72.34 Z.D. Wang, V. Kumar: Object closure and manipulation by multiple cooperating mobile robots, IEEE Int. Conf. Robot. Autom. (2002) pp. 394–399
72.35 P. Song, V. Kumar: A potential field based approach to multi-robot manipulation, IEEE Int. Conf. Robot. Autom. (2002) pp. 1217–1222
72.36 Z. Wang, Y. Hirata, K. Kosuge: Control a rigid caging formation for cooperative object transportation by multiple mobile robots, Proc. IEEE Int. Conf. Robot. Autom. (2004) pp. 1580–1585
72.37 O. Holland, C. Melhuish: Stigmergy, self-organization, and sorting in collective robotics, Artif. Life **5**(2), 173–202 (1999)
72.38 J.L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, L. Chretien: The dynamics of collective sorting robot-like ants and ant-like robots, Proc. 1st Int. Conf. Simul. Adapt. Behav. Anim. Anim. (1990)
72.39 R. Beckers, O. Holland, J. Deneubourg: From local actions to global tasks: Stigmergy and collective robotics, Proc. 14th Int. Workshop Synth. Simul. Living Syst. (1994) pp. 181–189
72.40 T. Wang, H. Zhang: Multi-robot collective sorting with local sensing, IEEE Intell. Autom. Conf. (2003)
72.41 T. Wang, H. Zhang: Collective Sorting with Multiple Robots, IEEE Int. Conf. Robot. Biomim. (2004) pp. 716–720
72.42 Y. Yang, M. Kamel: Clustering ensemble using swarm intelligence, IEEE, Swarm Intell. Symp. (2003) pp. 65–71
72.43 A. Martinoli, F. Mondada: Collective and cooperative group behaviours: Biologically inspired experiments in robotics, Lect. Notes Comput. Sci. **223**, 1–10 (1997)
72.44 J. Werfel, R. Nagpal: Three-dimensional construction with mobile robots and modular blocks, Int. J. Robot. Res. **27**(3/4), 463–479 (2008)
72.45 J. Werfel, Y. Bar-Yam, D. Rus, R. Nagpal: Distributed construction by mobile robots with enhanced building blocks, IEEE Int. Conf. Robot. Autom. (2006) pp. 2787–2794
72.46 J. Werfel: Building patterned structures with robot swarms, Proc. 19th Int. Joint Conf. Artif. Intell. (2005) pp. 1495–1502
72.47 J. Werfel, R. Nagpal: Extended stigmergy in collective construction, IEEE Intell. Syst. **21**(2), 20–28 (2006)
72.48 J. Werfel: Building blocks for multi-robot construction, Distrib. Auton. Robot. Syst. **6**, 285–294 (2007)

72.49 Y. Terada, S. Murata: Automatic modular assembly system and its distributed control, Int. J. Robot. Res. **27**, 445–462 (2008)
72.50 R.L. Stewart, R.A. Russell: A distributed feedback mechanism to regulate wall construction by a robotic swarm, Adapt. Behav. **14**(1), 21–51 (2006)
72.51 J. Wawerla, G.S. Sukhatme, M.J. Mataric: Collective construction with multiple robots, IEEE/RSJ Int. Conf. Intell. Robot. Syst. (2002) pp. 2696–2701
72.52 R.L. Stewart, R.A. Russell: Building a loose wall structure with a robotic swarm using a spatio-temporal varying template, IEEE/RSJ Int. Conf. Intell. Robot. Syst. (2004) pp. 712–716
72.53 C.A.C. Parker, H. Zhang, C.R. Kube: Blind bulldozing: Multiple robot nest construction, IEEE/RSJ Int. Conf. Intell. Robot. Syst. (2003) pp. 2010–2015
72.54 T. Huntsberger, G. Rodriguez, P.S. Schenker: Robotics challenges for robotic and human mars exploration, Proc. Robot. (2000) pp. 340–346
72.55 R.A. Brooks, P. Maes, M.J. Mataric, G. More: Lunar based construction robots, Proc. Int. Conf. Intell. Robot. Syst. (1990)