

Occlusion-Based Cooperative Transport with a Swarm of Miniature Mobile Robots

Jianing Chen, *Student Member, IEEE*, Melvin Gauci, *Student Member, IEEE*, Wei Li, *Student Member, IEEE*, Andreas Kolling, *Member, IEEE*, and Roderich Groß, *Senior Member, IEEE*

Abstract—This paper proposes a strategy for transporting a large object to a goal using a large number of mobile robots that are significantly smaller than the object. The robots only push the object at positions where the direct line of sight to the goal is occluded by the object. This strategy is fully decentralized and requires neither explicit communication nor specific manipulation mechanisms. We prove that it can transport any convex object in a planar environment. We implement this strategy on the e-puck robotic platform and present systematic experiments with a group of 20 e-pucks transporting three objects of different shapes. The objects were successfully transported to the goal in 43 out of 45 trials. When using a mobile goal, teleoperated by a human, the object could be navigated through an environment with obstacles. We also tested the strategy in a 3-D environment using physics-based computer simulation. Due to its simplicity, the transport strategy is particularly suited for implementation on microscale robotic systems.

Index Terms—Cooperative transport, cooperation without communication, occlusion, e-puck, swarm robotics.

I. INTRODUCTION

THE transport of large and heavy objects toward specific goal locations is a task that lends itself to the use of multiple robots. However, a survey of the literature reveals that multirobot systems that are capable of solving this task are often sophisticated even in proof-of-concept studies. One of the problems is visual occlusion. To move the object in the correct direction, robots must interact not only with the object, but also with the goal. As the object is often larger than the robots, it

Manuscript received September 11, 2014; revised December 25, 2014; accepted February 3, 2015. Date of publication March 5, 2015; date of current version April 2, 2015. This paper was recommended for publication by Associate Editor L. Pallottino and Editor T. Murphey upon evaluation of the reviewers' comments. This research was funded by the Engineering and Physical Sciences Research Council (grant no. EP/K033948/1) and a Marie Curie European Reintegration Grant within the 7th European Community Framework Programme (grant no. PERG07-GA-2010-267354). The work of M. Gauci was supported by a Strategic Educational Pathways Scholarship (Malta). The scholarship is part-financed by the European Union—European Social Fund (ESF) under Operational Programme II—Cohesion Policy 2007–2013, “Empowering People for More Jobs and a Better Quality of Life.” A preliminary version of this work was presented at the 2013 International Conference on Robotics and Automation [1].

The authors are with Sheffield Robotics and the Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, S1 3JD, U.K. (e-mail: j.n.chen@sheffield.ac.uk; m.gauci@sheffield.ac.uk; wei.li1@sheffield.ac.uk; a.kolling@sheffield.ac.uk; r.gross@sheffield.ac.uk).

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors. The material consists of a video that shows a swarm of robots transporting objects of various shapes and sizes to a goal.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TRO.2015.2400731

may occlude their view of the goal. The problem of how robots should perceive the goal and potentially inform each other about its position is not simple to solve [2], [3], and often imposes limitations on the system. For instance:

- 1) In a 2-D environment, the robots could perceive the goal using sensors that are positioned higher than the object [4], [5]. However, this imposes a limitation on the maximum possible height of the object.
- 2) A centralized system could be used, whereby an infrastructure is in place to handle the localization of and communication with robots [6]. The applicability of such systems is restricted to environments where these infrastructures are available.
- 3) A decentralized system could be used that relies on inter-robot communication. For example, some of the robots could perceive the goal and inform other robots that are not able to perceive the goal [7]–[11]. This solution usually requires a reliable communication technology, which may limit the system's scalability with respect to the number of robots.
- 4) The object itself could be considered as part of the solution, whereby it is explicitly designed or modified in such a way to assist robots in transporting it to the goal [3]. This, however, results in a system with limited generalizability to other objects.

The novelty of the transport strategy presented in this paper is that rather than treating occlusion as a problem to be overcome, occlusion is used to organize a swarm of robots to push a large object to a goal. The basic idea is to push the object across the portion of its surface, where it occludes the direct line of sight to the goal. This results in the transportation of the object along a path that may not be optimal, but always arrives at the goal. As shown in this paper, the strategy can be implemented in a fully decentralized manner. The robots use on-board cameras to perceive the object and goal. They do not need to communicate explicitly with each other. The performance of the group scales well with the number of robots, making it possible to transport objects of various shapes and sizes.

The simplicity of the strategy makes it particularly suited for the implementation on mobile robots that have limited capabilities [12], [13]. In the long term, such simple multi-robot strategies could be implemented at very small scales. Potential applications for swarms of such minimalist robots could be the delivery of drugs through the vascular network of humans or the removal of debris within narrow fluid pipelines.

This paper extends preliminary work that was presented in [1]. It presents for the first time a mathematical analysis of

the transport strategy, proving its correctness for objects of arbitrary convex shapes. Moreover, it presents the results from a new set of experiments that assess 1) the effectiveness of the strategy in transporting objects of different shapes and sizes, and 2) the ability of the strategy to transport an object toward a dynamic target. Finally, results obtained from simulation suggest that the strategy can also be implemented in 3-D environments.

This paper is organized as follows. Section II discusses related work. Section III describes the problem formulation and introduces the transport strategy in a platform-independent manner. Section IV provides a proof of the strategy's correctness for objects of arbitrary convex shapes moving in a planar environment. Section V presents a set of experiments using the e-puck robotic platform. Section VI studies the strategy when the goal is a mobile robot controlled by a human. Section VII presents a conceptual implementation of the strategy in a 3-D environment using physics-based computer simulation. Section VIII concludes the paper.

II. RELATED WORK

Over the past 20 years, multirobot object transportation has become a canonical task for studying cooperation in groups of robots. The three most common types of strategies are pulling, pushing, and caging.

Transport by pulling involves making a number of robots connect themselves to the object, for example, through grasping [14]–[16] and/or lifting [17]–[20]. In nature, such behavior seems to require relatively little intelligence on behalf of the individuals [21]. However, the pulling strategy is still difficult to be applied on robotic systems because of the complexity of the physical mechanisms.

Transport by pushing is a simple way of manipulating an object when the object is relatively large compared with the robots. The problem of stabilizing the moving direction of the object, while being pushed by a single robot, is similar to the inverted pendulum problem; the controller design is difficult compared with the simple physical mechanism it requires. In a multi-robot pushing system, increasing the number of pushing robots not only increases the overall pushing force but also simplifies the stabilization problem because the pushing forces distributed over multiple contact points on the object can be used to reach equilibrium [22]. For example, in [23], a physical system that uses two six-legged robots to push a large rectangular object was presented. In the experiment, the object is movable by one robot, but the performance was improved significantly when the object was pushed by two robots that cooperated through wired communication. Nevertheless, it is still a problem for robots in cooperative transport to choose good pushing positions and speeds.

Cooperative transport by caging is a special case of pushing. It requires a group of robots to organize themselves into a formation around the object in a way that the object is caged inside the formation [5], [24]. As long as the formation of the robots is maintained, while they are moving, the object will follow the group of robots. Depending on the shape of the ob-

ject, caging can be a complex problem [25], [26]. As the caging solutions often require a certain number of robots and a considerable amount of information about the object, it is challenging to design a single caging system that is scalable in terms of the number of robots and flexible in terms of object types. In [6], a caging system that copes with a variable number of robots is presented. A group of robots orbit around an object that has corners. The object is however only moved by a few robots at a time, which imposes a limit on the object's weight.

It is desirable for a cooperative transport system to be scalable with respect to the number of robots [2]. One common point of pushing/caging-based systems (including all systems referred to before except for [6]) is that the number of robots is not large, typically fewer than five. One important factor that limits the number of robots is the use of interrobot communication to achieve highly cohesive behavior. There are a few works that have studied how a relatively large group of robots can be used in a cooperative transport task when the controller only requires local information. For example, in [4] and [27], a system that took inspiration from ants is studied. The robots simply map the perceptual cues obtained from a small number of sensors onto nine motion primitives. Due to the simplicity of the control method, the number of robots working simultaneously in the cooperative transport task is flexible and a physical system containing 3 to 6 robots was used in experiments. In [3], a physical system that includes up to 100 Kilobots was used to study a decentralized strategy for collective transport. The strategy was evaluated in situations where the robots resided within the object being transported. In [28], a large swarm of Kilobots was controlled using a global input signal issued by a human operator to transport objects toward a goal.

It is also very common that the dimensions of the object need to be limited so that pushing robots can directly perceive other robots or the goal. For example, many systems (including all pushing systems referred to before) require the object to be lower than some of the sensors on the robots. An alternative decentralized approach is through role differentiation using explicit interrobot communication. For example, [8] presents a box pushing system where robots assume different roles. In the case of cooperative transport, the roles are "pusher" and "watcher." The watcher is in front of the object and observes the goal, while the pushers are behind the object. The robots communicate through WiFi. In [9], an underwater box-pushing system is presented with three robotic fish; two of them work as pushers while the other works as an observer. The fish can share sensing information through explicit communication to work out the approximate pose of the box, the two pushers can push on appropriate positions without seeing the goal directly.

One important property of the method we propose is that neither consistent perception of the goal nor explicit communication are required for robots that are pushing the object. This removes most of the limitations discussed above. Assuming the robots are significantly smaller than the object, they can organize themselves into positions where the direct line of sight to the goal is occluded.

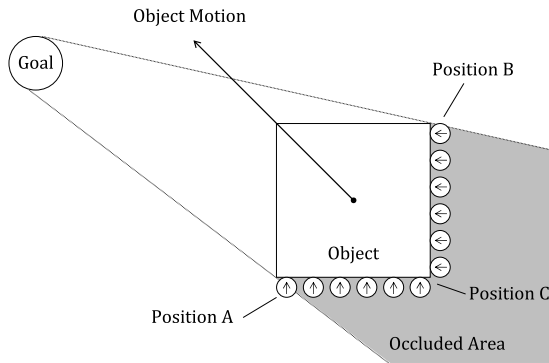


Fig. 1. Illustration of how a swarm of robots can push a large object in a 2-D planar environment (adapted from [1]). The robots keep pushing only along the section of the object’s perimeter that occludes their views of the goal. As a consequence, the motion of the object will be approximately toward the goal.

III. METHODOLOGY

A. Problem Formulation

The task we consider is as follows. A bounded environment contains a convex-shaped object, a goal, and a number of robots. The environment is otherwise free of obstacles. The aim is that the robots, which are initially placed in arbitrary locations, push the object to the goal. Note that the goal specified in the problem may not be the final destination of the transportation. In a broader scenario, the goal could be moving, or it could be one of a series of way points (see Section VI).

We make the following assumptions. The object and the goal can each be recognized by the robots. The dimension of the object is large enough to occlude the robots’ perception of the goal when they are behind it (see Fig. 1). The robots can perceive the goal from any point within the environment, unless it is occluded by the object.

B. Occlusion-Based Cooperative Transport Strategy

Consider a number of robots that can distribute themselves uniformly around the section of the object’s surface that occludes their view of the goal (the “back side” of the object), as shown in Fig. 1. Then, if all the robots push the object by moving in a direction perpendicular to the object’s surface at their points of contact, the motion of the object will be approximately toward the goal. As the object moves, its occluded surface changes over time, thus changing the direction of motion. If the robots keep pushing *only* against the occluded surface, the object will eventually reach the goal.¹

The occlusion-based cooperative transport strategy can be realized using a fully decentralized behavior and without explicit communication among the robots. In Fig. 2, the behavior of the individual robots is given in the form of a state machine. A robot first searches the object using an algorithm that is suitable for the environment (“*Search Object*”). For bounded environments, as considered in this paper, the robot performs a random walk.

¹The strategy could in principle be also used for transporting objects that are not tall enough to occlude the robots’ view of the goal. If a robot reached the object, but the goal was visible “behind” it, the robot would then still push.

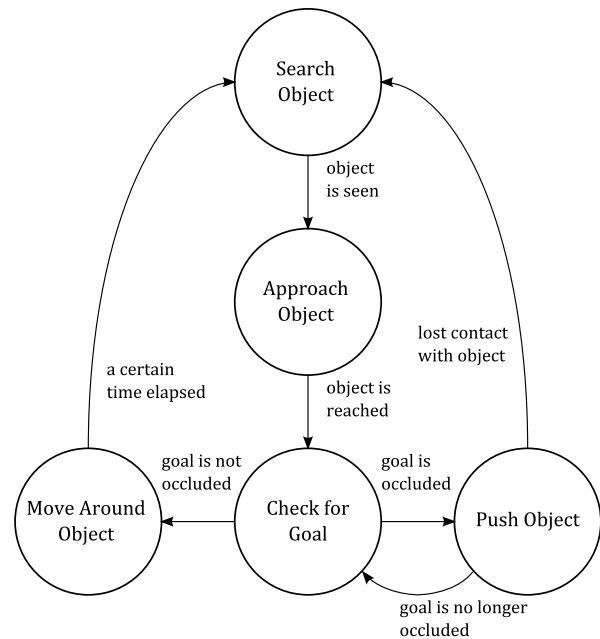


Fig. 2. State machine representation of the individual robot behavior realizing the occlusion-based cooperative transport strategy. The start state is “*Search Object*.” If the object is lost at any stage, the robot restarts from “*Search Object*.” The behavior is fully decentralized and does not require explicit interrobot communication.

More sophisticated search algorithms could help our strategy to also cope with unbounded environments. Once the object is seen the robot moves toward it (“*Approach Object*”). When the robot has reached the object, it enters state “*Check for Goal*” to work out whether the goal can be seen from its position. If the goal cannot be seen, the robot will push the object simply by moving against it (“*Push Object*”). If the goal can be seen, the robot will attempt to find another position around the object (“*Move Around Object*”), for example, executing a left-hand-wall-following behavior.

Although not strictly necessary, a behavior realizing the above strategy should also prevent robots from colliding with each other and the boundaries of the environment. This can greatly improve performance because robots move with fewer collisions (if any). Hence, in our implementations, robots and the boundary are treated as obstacles to avoid. The goal, if embodied, is also treated as an obstacle, while it still serves as the target of transportation.

When a group of robots execute the overall behavior, they eventually end up at different positions along the occluded section of the object due to the stochastic nature of the system. However, the more robots that are used, the more likely it is that they approximate a uniform distribution (as shown in Fig. 1).

IV. MATHEMATICAL ANALYSIS

In this section, we analyze the occlusion-based cooperative transport strategy for the case of arbitrary convex objects in planar environments. We prove that, under some idealized assumptions, the strategy always succeeds in moving the object

to the goal. Note that the transport strategy is not suited for objects of *arbitrary* concave shapes (for a counter example, see Appendix A).

A. Modeling of the Occlusion Problem

We assume that each of the goals and robots are points (without embodiment). Let $\mathbf{c} \in \mathbb{R}^2$ be the center of mass of a rigid convex object with respect to a coordinate frame in which $\mathbf{g} = [0, 0]^T$ is the goal point. Let the perimeter of the object be described by a closed, convex, and differentiable curve given by

$$\mathbf{p}(\theta) = \begin{bmatrix} r(\theta) \cos \theta \\ r(\theta) \sin \theta \end{bmatrix} + \mathbf{c} \quad (1)$$

with $\theta \in [0, 2\pi]$ and $r : [0, 2\pi] \rightarrow \mathbb{R}$ differentiable and satisfying $r(2\pi) = r(0)$. By specifying $r(\theta)$, any convex shape can be approximated by \mathbf{p} . Initially, \mathbf{g} is outside \mathbf{p} .

The inward pointing normal vector on $\mathbf{p}(\theta)$, named $\mathbf{N}(\theta)$, is the derivative of $\mathbf{p}(\theta)$ rotated by $\frac{\pi}{2}$

$$\mathbf{N}(\theta) = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \mathbf{p}'(\theta). \quad (2)$$

Points along \mathbf{p} where the direct line of sight to \mathbf{g} is occluded are between the two tangent points of \mathbf{p} from point \mathbf{g} . We write the two tangent points as $\mathbf{p}(\alpha)$ and $\mathbf{p}(\beta)$, $\alpha, \beta \in [0, 2\pi]$. As tangent points, they satisfy

$$\begin{aligned} \mathbf{p}(\alpha) \cdot \mathbf{N}(\alpha) &= \mathbf{p}(\beta) \cdot \mathbf{N}(\beta) = 0 \\ \mathbf{p}(\theta) \cdot \mathbf{N}(\theta) &> 0 \quad \forall \theta \in (\alpha, \beta). \end{aligned} \quad (3)$$

Since \mathbf{p} is convex and \mathbf{g} is outside \mathbf{p} , α and β are well defined. For convenience, write $\mathbf{a} = \mathbf{p}(\alpha)$ and $\mathbf{b} = \mathbf{p}(\beta)$. Additionally, they are named; therefore, \mathbf{a} is the tangent point on the right side of vector $(\mathbf{c} - \mathbf{g})$, while \mathbf{b} is the one on the left side. Strictly speaking, \mathbf{a} and \mathbf{b} satisfy

$$\begin{aligned} a_x c_y - a_y c_x &> 0 \\ b_x c_y - b_y c_x &< 0 \end{aligned} \quad (4)$$

with x and y subscripts denoting the x and y coordinates. These properties of \mathbf{a} and \mathbf{b} will play an important role in the proof of the transport strategy later.

Fig. 3 illustrates the above definitions. In colloquial terms, all points $\mathbf{p}(\theta)$ with $\theta \in (\alpha, \beta)$ are on the occluded perimeter of the object, while all other points on \mathbf{p} are visible from \mathbf{g} .

B. Resultant Force Applied on the Object

Lemma 1: Assume that $n \rightarrow \infty$ robots are uniformly distributed along the occluded perimeter of the object and they are the only robots asserting a force on the object. The direction of the resultant force asserted on the object by the robots is equal to the direction of the vector $(\mathbf{b} - \mathbf{a})$ rotated by $\frac{\pi}{2}$ and its magnitude is proportional to $\|\mathbf{b} - \mathbf{a}\|$.

Proof: According to the strategy, all robots along the occluded perimeter assert normal forces on \mathbf{p} . Without loss of generality, let the magnitude of the force be one unit force per




Fig. 3. If normal forces are uniformly applied on the blue section of the convex-shaped object's perimeter (major arc \mathbf{ab} in this diagram), the combined force vector \mathbf{F} is the vector $(\mathbf{b} - \mathbf{a})$ rotated by $\frac{\pi}{2}$ and its magnitude is proportional to the length $\mathbf{b} - \mathbf{a}$ (chord \mathbf{ab} in this diagram). Point \mathbf{q} is an affecting point of \mathbf{F} . \mathbf{c} denotes the center of mass of the object. \mathbf{g} denotes the goal.

unit length. The combined force is the definite integral given by

$$\mathbf{F} = \int_{\alpha}^{\beta} \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \mathbf{p}'(\theta) d\theta. \quad (5)$$

The solution of the definite integral in (5) is

$$\mathbf{F} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} (\mathbf{p}(\beta) - \mathbf{p}(\alpha)) \quad (6)$$

which is

$$\mathbf{F} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} (\mathbf{b} - \mathbf{a}). \quad (7)$$

□

We can also derive the torque around the z -axis caused by the robots. For this, with slight abuse of notation, we interpret all previous points as embedded in the x, y plane in \mathbb{R}^3 . Again, we assume that the magnitude of the force is one unit force per unit length. Then, the magnitude of the torque around z -axis contributed by all robots with respect to point \mathbf{c} is

$$Q = \int_{\alpha}^{\beta} [(\mathbf{p}(\theta) - \mathbf{c}) \times \mathbf{N}(\theta)] \cdot \hat{\mathbf{z}} d\theta \quad (8)$$

where $\hat{\mathbf{z}}$ represents a unit vector pointing along the z -axis. The part within the integral is equal to

$$\begin{bmatrix} r(\theta) \cos(\theta) \\ r(\theta) \sin(\theta) \\ 0 \end{bmatrix} \times \begin{bmatrix} -r'(\theta) \sin(\theta) - r(\theta) \cos(\theta) \\ r'(\theta) \cos(\theta) - r(\theta) \sin(\theta) \\ 0 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad (9)$$

which can be simplified to $r'(\theta)r(\theta)$. Then, (8) can be written as

$$Q = \int_{\alpha}^{\beta} r'(\theta)r(\theta)d\theta. \quad (10)$$

Its solution is

$$Q = \frac{r^2(\beta) - r^2(\alpha)}{2}. \quad (11)$$

Lemma 2: If the combined force contributed by the robots \mathbf{F} is considered as a single force, while Q is the torque induced by \mathbf{F} , the mid point of segment \mathbf{ab} is an affecting point of \mathbf{F} .

Proof: Naming the affecting point of \mathbf{F} as \mathbf{q} , \mathbf{F} , \mathbf{q} and Q must satisfy

$$Q = [(\mathbf{q} - \mathbf{c}) \times \mathbf{F}] \cdot \hat{\mathbf{z}}. \quad (12)$$

The above equation can be transformed into

$$\mathbf{q} \cdot (\mathbf{b} - \mathbf{a}) = \frac{r^2(\beta) - r^2(\alpha)}{2} + \mathbf{c} \cdot (\mathbf{b} - \mathbf{a}) \quad (13)$$

which can be viewed as the vector equation of a line.

While \mathbf{q} can be any point on (13), we make \mathbf{q} a convenient point on (13), which is

$$\mathbf{q} = \frac{\mathbf{a} + \mathbf{b}}{2}. \quad (14)$$

□

C. Motion Dynamics of the Object

As the object is moved, \mathbf{a} and \mathbf{b} can change over time. We assume that the robots react instantly to such changes so that the occluded perimeter is always uniformly filled up with pushing robots. Thus, (7) is valid at any point in time as long as \mathbf{g} is outside \mathbf{p} . In other words

$$\mathbf{F}(t) = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} (\mathbf{b}(t) - \mathbf{a}(t)). \quad (15)$$

From (15), it follows that the rotation of the object does not affect the relationship between \mathbf{a} , \mathbf{b} , and \mathbf{F} . According to Newton's laws, the translation dynamics of the center of mass of the object are

$$\mathbf{v} = \dot{\mathbf{c}}, \dot{\mathbf{v}} = \frac{\mathbf{F}}{M} \quad (16)$$

where $\dot{\mathbf{v}}$ (respectively $\dot{\mathbf{c}}$) is the derivative of \mathbf{v} (respectively \mathbf{c}) with respect to time t , and M is the object's mass.

We can apply a quasi-static analysis to the case here in which some robots are pushing a rigid object slowly [29]. Then, the translation dynamics of the object is

$$\dot{\mathbf{c}} = k\mathbf{F} \quad (17)$$

where $k \in \mathbb{R}^+$ is a positive constant that transfers \mathbf{F} proportionally to the velocity of the object.

D. Convergence of the Object's Distance to the Goal

Theorem 1: The distance between the object's center of mass (\mathbf{c}) and the goal (\mathbf{g}) is strictly decreasing over time if the velocity of the object is governed by (17). As $t \rightarrow \infty$, \mathbf{g} will be on the object perimeter \mathbf{p} .

Proof: Let $l(t) = \mathbf{c}(t) \cdot \mathbf{c}(t)$ be the squared distance of the center of mass \mathbf{c} to goal \mathbf{g} , then its derivative with regard to time is

$$\dot{l} = 2k\mathbf{c} \cdot \mathbf{F}. \quad (18)$$

Substituting \mathbf{F} with (7), we get

$$\mathbf{c} \cdot \mathbf{F} = (b_x c_y - b_y c_x) - (a_x c_y - a_y c_x). \quad (19)$$

According to (4), $\mathbf{c} \cdot \mathbf{F} < 0$. Hence, $l(t)$ is strictly decreasing. Since $l(t) \geq 0$ for all $t > 0$ (as long as \mathbf{g} is outside \mathbf{p}), we get $\lim_{t \rightarrow \infty} l(t) = L \in \mathbb{R}$. Therefore,

$$\lim_{t \rightarrow \infty} \mathbf{c} \cdot \mathbf{F} = \lim_{t \rightarrow \infty} b_x c_y - b_y c_x + a_y c_x - a_x c_y = 0 \quad (20)$$

which together with (4) implies that

$$\begin{aligned} \lim_{t \rightarrow \infty} b_x c_y - b_y c_x &= 0 \\ \lim_{t \rightarrow \infty} a_y c_x - a_x c_y &= 0. \end{aligned} \quad (21)$$

In other words, the areas of the triangles \mathbf{gca} and \mathbf{gcb} approach zero as $t \rightarrow \infty$. Since \mathbf{c} is always inside \mathbf{p} , the triangles \mathbf{gca} and \mathbf{gcb} can never have 0 area unless $\mathbf{a} = \mathbf{g}$ and $\mathbf{b} = \mathbf{g}$ (see Fig. 3). Hence, as $t \rightarrow \infty$, \mathbf{g} will be on \mathbf{p} . In other words, the object will ultimately coincide with the goal and stop moving. □

V. EXPERIMENTS WITH OBJECTS OF DIFFERENT SHAPES

To assess the occlusion-based cooperative transport strategy in a 2-D planar environment, a decentralized controller is implemented on a centimeter-scale mobile robot platform.

In our previous work [1], a preliminary version of the controller was validated by experiment using a rectangular box of dimensions 42 cm \times 39 cm as the object. Using this relatively regular object, we demonstrated the feasibility of the transport strategy; the object was transported successfully to the goal in all 30 trials that were conducted.

After analyzing the transport strategy mathematically, we obtained an indication of objects with not-unusual shapes that are nevertheless challenging for the strategy to handle. In this section, a new set of experiments is introduced to evaluate the strategy using objects of these shapes. The results obtained are compared with predictions from the mathematical model. The section also describes the robotic system, as well as the controller, which is an improved version over [1].

A. Robot Platform and Sensing

For the physical implementation, we use the e-puck, which is an off-the-shelf differential-wheeled robot [30]. The e-puck is around 7.0 cm in diameter, around 5.5 cm high, and weighs approximately 150 g. Its maximum speed is 12.8 cm/s. Fig. 4(a) shows a photograph of an e-puck. In this study, each e-puck was fitted with a black "skirt" to give it a uniform color. In addition, it was fitted with a green top marker to facilitate the post-analysis of videos taken by an overhead camera.

Fig. 4(b) shows a schematic of the e-puck including the locations of the sensors used in this study. The e-puck has eight infrared proximity sensors distributed around its body; they are 3.1 cm above the ground. It also has a directional color camera in the front of its hull that is 2.8 cm above the ground.

The infrared proximity sensors measure, at a rate of 50 times per second, the proximity to embodied items: the object, the goal, the environment boundary, and other e-pucks. The proximity to the first three items (passive items) is estimated by




Fig. 4. e-puck robot. (a) e-puck fitted with a black skirt and a green top marker. (b) Top-view schematic of an e-puck (adapted from [1]), indicating the locations of its wheels, camera, and proximity sensors.

sending pulses of infrared light and measuring their reflections (discarding possible contributions from ambient infrared light). We found that this method does not provide reliable estimates for the proximity of e-pucks—neither the black skirts nor the plain e-pucks would be suitable reflectors. To mitigate this problem, we use a customized sampling routine, whereby the e-puck emits infrared light almost continuously (see [1] for details).²

The directional color camera is used to recognize both object and goal. The object is the only blue item in the environment, the goal is the only red item in the environment.³ The camera provides images of resolution up to 640×480 at around 18 frames per second. The image is however subsampled to 40×15 pixels.⁴ Each captured image is processed to provide four scalar values: 1) the number of pixels that are considered blue and red, and 2) the horizontal distribution biases of the blue and red pixels. For details, see [1].

B. Controller

The e-puck controller is a state machine implementing the individual behaviors of the transport strategy (see Fig. 2).

²Note that the sensors are *not* used for explicit interrobot communication, which in principle would be possible [3], [31].

³The e-puck's wheels, which are partly red too, are hidden behind the skirt.

⁴To achieve this, a customized library was used.

The robot performs a random walk and approaches any blue object seen by its camera. If the robot loses sight of the object, it resumes the search. When it reaches the object, it does a full rotation to look for the red goal. If the goal is not seen, the robot starts pushing the object. If the goal is seen, the robot executes a left-hand-wall-following behavior, which relocates the robot to a position where the goal may be occluded by the object.

When in the pushing formation, a robot's perception of the goal may not only be occluded by the object but by its neighboring robots as well. However, the robots at the two ends of the formation (i.e., at Positions A and B in Fig. 1) can effectively monitor the visibility of the goal. These robots can be considered as observers. When an observer perceives the goal, it leaves the formation. Consequently, its neighbor becomes an observer. Thus, pushing robots that are no longer in the occluded perimeter happen to leave in a recursive manner. For e-pucks, this behavior is utilized so that only observers are required to scan the environment for the goal, while the other pushing robots can be devoted exclusively to pushing the object.

During transport, a pushing robot moves perpendicularly toward the object's surface in front of it. If the object has a curved perimeter (e.g., a circle), this means the distance between two pushing robots will become smaller when the object starts moving. Thus, collisions between the robots in the pushing formation will occur. This problem is magnified by the e-puck's design: two e-pucks will easily get stuck when they collide. In our previous work [1], the e-pucks avoided collisions by leaving the pushing formation. In the version used in this experiment, an improved implementation was used to let the pushing robots adjust their moving direction to avoid collisions and/or leave the pushing formation.

In order to make the controller work in a real environment, basic behaviors like collision avoidance and error handling are added into the state machine. For most of the state transition conditions, certain sensory inputs are compared with a preset threshold. In each of the states, specific low-level motion controllers are activated to achieve the required motion. Each of these controllers calculates the left and right wheel speed by summing the weighted input of the proximity sensors and of values extracted from the camera.

The implementation of the motion controllers is detailed in [1]. The full state machine used on the e-puck and the input weights are found in [32].

C. Experimental Setup

1) *Objects*: We conducted experiments with three objects of different shapes and sizes:

a) *Circular object*: Theoretically, this is an ideal case as the resulting force points directly to the goal. However, in practice, the curved perimeter could make the robots more prone to collide with each other, as the object is being moved. Therefore, it is essential that the collision avoidance mechanism in the pushing state is effective. As the pushing force of e-pucks is rather limited, at least three robots are required to push this object.⁵

⁵Depending on the floor condition and robot power, occasionally this object may also be pushed by just two robots.

TABLE I
SUMMARY OF THE EXPERIMENTAL SETUP AND DATA

Object Characteristics				Experimental Results						
Shape and Size	Height	Mass	Pushing Force	Successful	Completion Time (s)		Path Efficiency		AE (°)	
			Required		Trials	mean	σ	mean	σ	mean
Circular, 40 cm diameter	10 cm	222 g	$\approx 0.75 N$	15 out of 15	220.0	26.3	0.914	0.029	26.7	16.8
Triangular, 45 – 60 – 75 cm	14 cm	432 g	$\approx 1.5 N$	14 out of 15	255.1	63.0	0.793	0.099	90.1	36.2
Rectangular, 58.5 × 13.5 cm	6.5 cm	160 g	$\approx 0.5 N$	14 out of 15	295.4	183.1	0.766	0.192	204.6	79.2




Fig. 5. Experimental setup. The robots were placed approximately in such a formation because the self-calibration of the proximity sensors on the e-puck requires a certain amount of space around the robot.

b) Scalene Triangular Object: This is a simple example of an asymmetrical object. In this case, the ratio of the lengths of the triangle's sides is 3:4:5. According to Lemma 2, the robots cannot push this object along a straight line, because the resultant force vector will never pass through the object's centroid (i.e., the resultant torque can never be zero). As a result, depending on which side(s) the robots are pushing from, the object will rotate clockwise or anticlockwise. Two robots pushing on the same side near the sharpest corner are enough to rotate this object. On the other hand, it takes at least four robots pushing on the same side in order to induce a translational motion.

c) Elongated Rectangular Object: This shape is problematic for the occlusion-based transport strategy, because the resultant force can deviate by almost 90° degrees from the ideal direction of transport. The object easily rotates if the pushing formation is not uniform; in fact, one robot pushing at one end is sufficient to induce a rotation. It takes at least two robots pushing on the same side to give this object a translational motion.




Fig. 6. Completion time of the circular object, scalene triangular object, and elongated rectangular object.

The physical details of the three objects are given in Table I. The mass of each of the objects was chosen so that it is theoretically possible for the e-pucks to push the object from all directions. The side of the objects are painted blue. Two orange markers of different size are attached on top of each object, so that its position and orientation can be tracked in an offline analysis.

2) Environment: The environment of the experiment is a rectangular arena of size 400 cm × 225 cm that is bounded by 50-cm-high walls. The floor of the arena has light gray color, and its walls are painted in white. The goal is a red cylinder of 25 cm diameter and 42 cm height.

3) Trial Procedure: For each of the objects, we conducted 15 trials, that is, we conducted 45 trials in total. The number of robots used in each trial was 20. This was much larger than the least number of robots required for pushing the objects. The strategy benefits from the use of more robots when dealing with objects of various sizes and shapes.

The initial configuration of a trial is illustrated in Fig. 5. The object's centroid was positioned as indicated. The orientation of the object was generated using a random number generator. The robots were placed in a zone between the object and the goal. The actual positions of the robots were loosely snapped to a grid to ensure a minimum gap between robots, which is required by our self-calibration routine for the e-puck. Before starting a trial, each robot was rotated by a random proportion of a full rotation to obtain its initial orientation. The trials were started by issuing

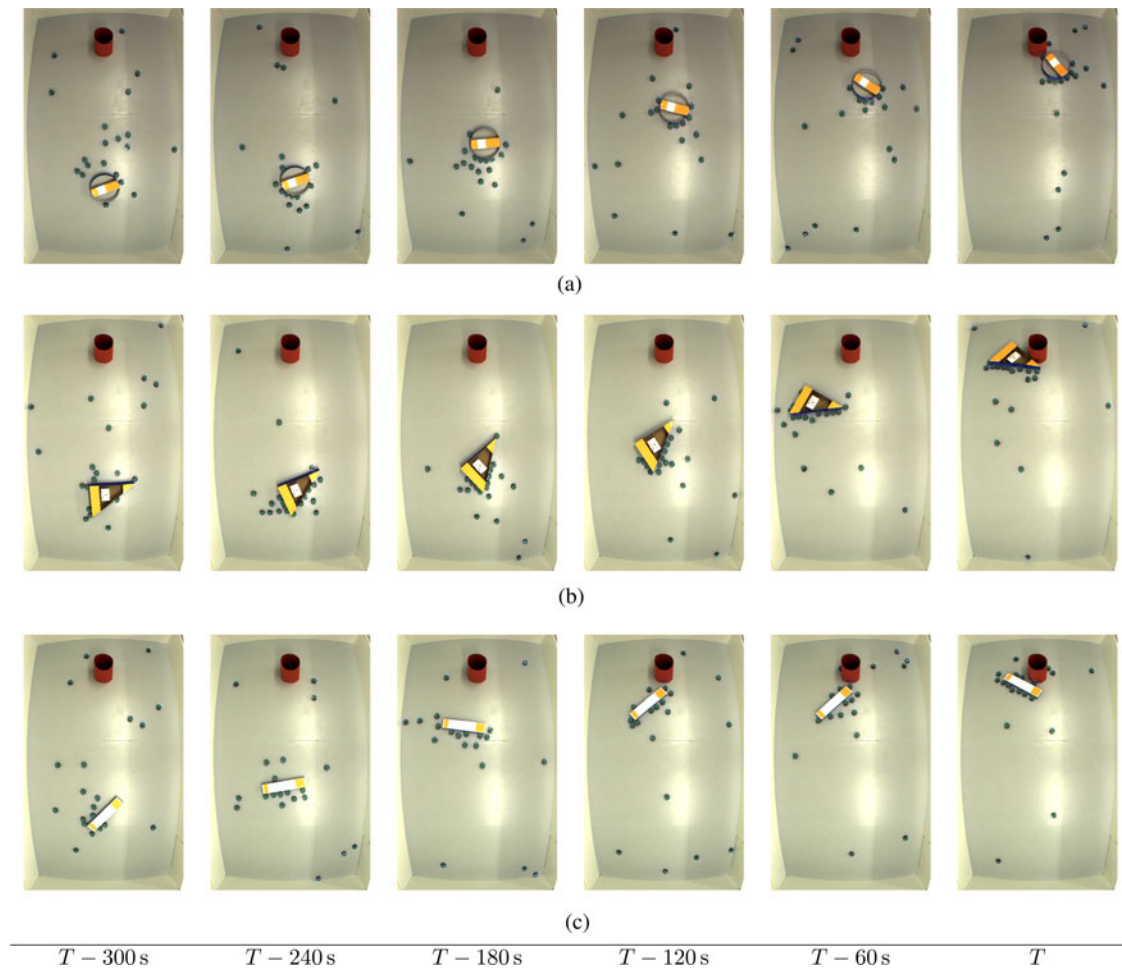


Fig. 7. Snapshots showing three trials with similar durations in the systematic experiments with a circular, triangular, and rectangular object, respectively. T is the total duration of the videos (in s), which ends at the moment when the object collides with the goal. Videos of all the 45 trials are available in [32]. (a) Circular object. (b) Triangular object. (c) Rectangular object.

a signal via an infrared remote control that is received by all robots simultaneously.⁶ The robots were programmed to stop automatically after 15 min.

A trial was stopped if either of the following situations happened:

- 1) The object collided with the goal object. The trial was then considered successful.
- 2) All of the robots stopped automatically due to the 15-min time limit. This means the trial was unsuccessful.
- 3) The object was too close to the wall and, thus, cannot be transported via pushing anymore. For example, either side of the triangular object fully touched the wall. This means the trial was unsuccessful.

The trials were recorded with an overhead camera. The videos were used for the offline tracking of the object. The accompanying video shows three experimental trials, one for each type of object. Videos of all 45 trials are available in [32].

⁶The e-puck's top features an infrared receiver, which can decode the modulated infrared signal from a TV remote.

D. Results

a) Successful Trials: Overall, 43 out of the 45 trials were successful. The object reached the goal within 15 min. One trial with the triangular object failed. The other failed trial was with the rectangular object. In both cases, one side of the object became very close to the boundaries of the arena. This was due to the limited width of the arena and a relatively large error in the transport direction. Fig. 7 shows snapshots from three successful trials.

b) Completion Time: The completion time T_k is defined as the time elapsed from the start of a trial until the centroid of the object is less than 62.5 cm away from the center of the goal (i.e., when the centroid of the object is within the goal region in Fig. 5).

A box-and-whisker plot⁷ of the completion time is given in Fig. 6. The deviations of completion times for the triangular

⁷The line inside the box represents the median of the data. The edges of the box represent the lower and the upper quartiles (25th and 75th percentiles) of the data, while the whiskers represent the lowest and the highest data points that are within 1.5 times the interquartile range from the lower and the upper quartiles, respectively. Crosses represent outliers.

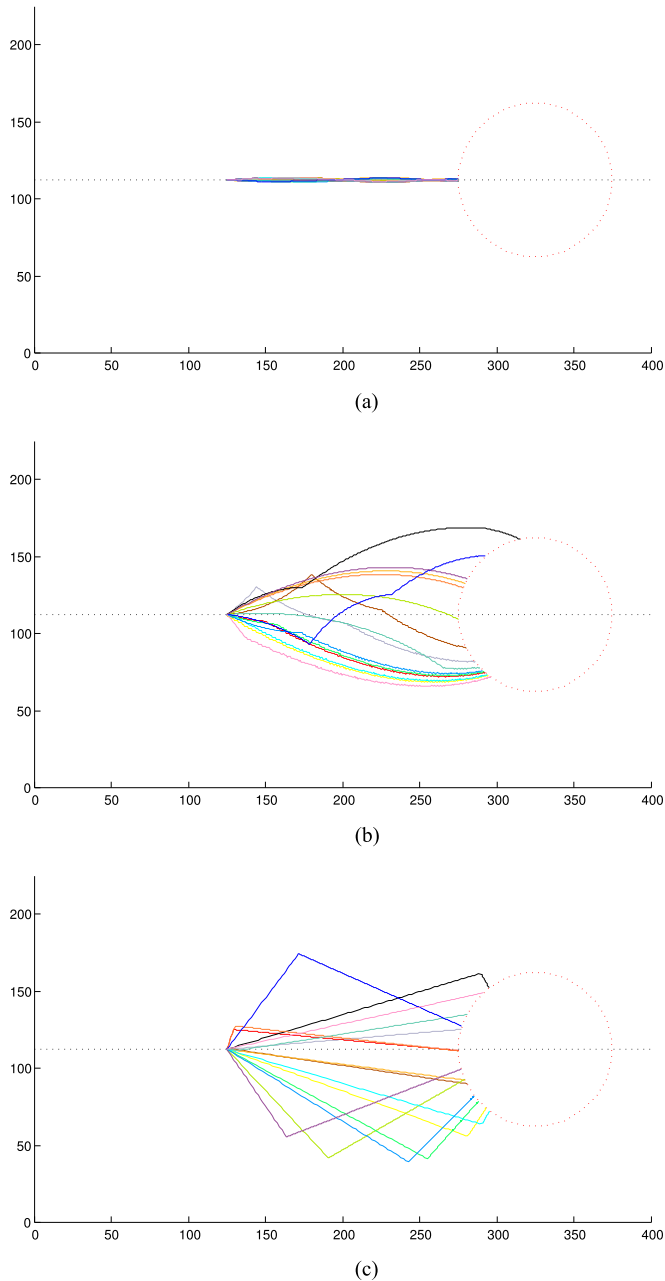


Fig. 8. Predicted paths of the centroid of the objects based on Lemmas 1 and 2. These paths are plotted using the same ratio on both of the axes; therefore, they can be compared with Fig. 9. (a) Circular object. (b) Triangular object. (c) Rectangular object.

and rectangular objects are larger than for the circular object, which shows that the shape of the object will affect the transport. During the trials, it was observed that if the elongated rectangular object reaches an orientation with either of its two small sides pointing toward the goal, it cannot be pushed effectively anymore. In Fig. 7(c), it can be observed from the last three snapshots that such a situation stalled the transport for at least 60 s. While the robots manage to rotate the object out of such situations, it is uncertain how long this takes.

c) Object Paths: According to Lemmas 1 and 2, the resultant force and torque applied on the object can be calculated given

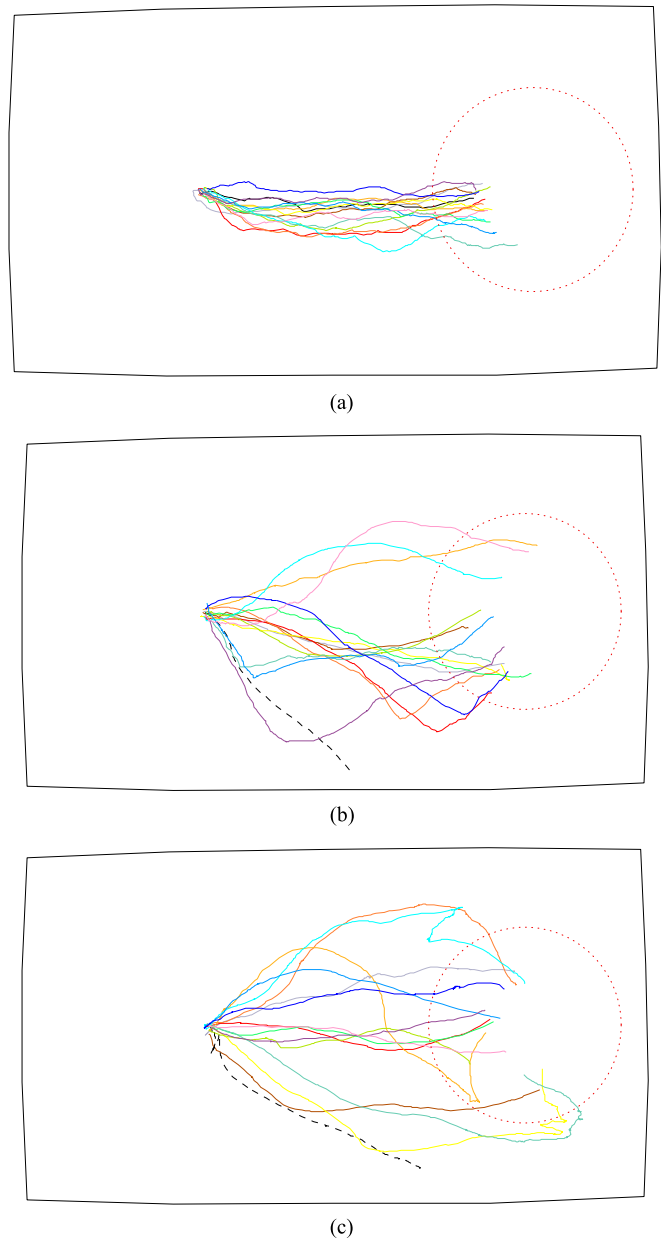


Fig. 9. Actual paths of the centroid of the objects. The dashed black lines are the paths of the two failed trials. The dotted red line is the goal region. It can be observed that the strategy has an effect to correct the direction in which the object is moved. Sometimes, this correction resulted in a significant change in the transport direction. (a) Circular object. (b) Triangular object. (c) Rectangular object.

the initial position and orientation of the object and goal position (assuming an infinite number of point robots are equally dispersed around the occluded perimeter of the object). When the force and torque are directly transferred to the linear and angular velocities of the object, it is possible to predict the objects' paths for the trials. The predicted paths are given in Fig. 8. In addition, the actual paths of the objects were traced from the videos recorded by the overhead camera. These paths are given in Fig. 9.

The differences between each pair of individual paths in Figs. 8 and 9 are obvious; in only some trials, the prediction

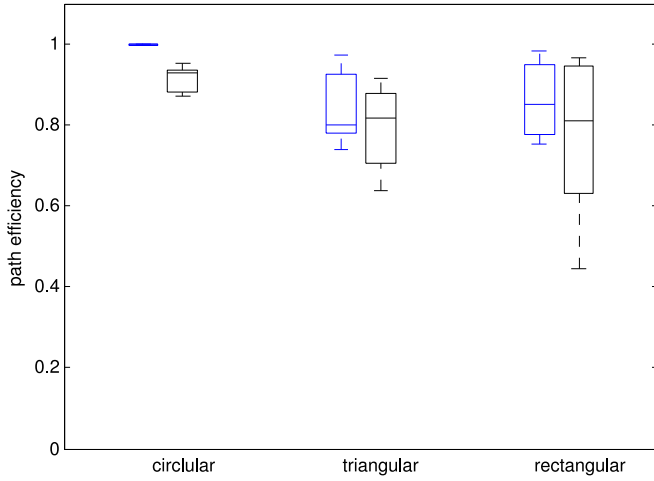


Fig. 10. Path efficiency in the successful trials. This metric compares the length of the path that the object moved with the length of the ideal straight path to reach the goal. For each of the objects, the predicted PE and actual PE are shown in blue (left) and black (right), respectively.

is close to the actual paths. This result was however expected, as many of the idealized assumptions made in Section IV are violated in a physical environment. For example, the robots will not be able to react instantaneously to changes in the object's occluded perimeter. Moreover, the robot's embodiment raises the issue of physical interferences. However, the overall distributions of the paths show a good correspondence.

- 1) The circular object tends to move directly to the goal.
 - 2) The paths of the triangular object are typically curved.
 - 3) The paths of the rectangular object have a more widespread uniform distribution.
- d) *Path Efficiency*: We define the path efficiency of a trial as

$$PE = \frac{s_{\min}}{s} \quad (22)$$

where s_{\min} is the distance between the start position and the goal region, and s is the length of the path of the object when its centroid enters the goal region. An ideal transport path would have a PE of 1.

For all successful trials, both the actual PE values and the PE values corresponding to the predicted paths shown in Fig. 8 are calculated. Fig. 10 shows a box-and-whisker plot of predicted PE versus actual PE for each of the objects. The predicted and actual PE s both reveal the difference in the efficiency when transporting objects of different shapes.

e) *Accumulated Angular Error*: The efficiency of a pushing-based transport strategy may also be affected if a substantial amount of unnecessary object rotation occurs in the process.

We define the accumulated angular error (AE) as the difference between the relative difference in the orientations at the beginning and the end of a trial and the total amount of changes of orientation. Let $\mathbf{p}(t)$ and $\mathbf{q}(t)$ be the centroids of the two tracking markers on top of the object in the video of a trial at time step t . Then, the orientation vector of the object at time step t is

$$\mathbf{a}(t) = \mathbf{p}(t) - \mathbf{q}(t). \quad (23)$$

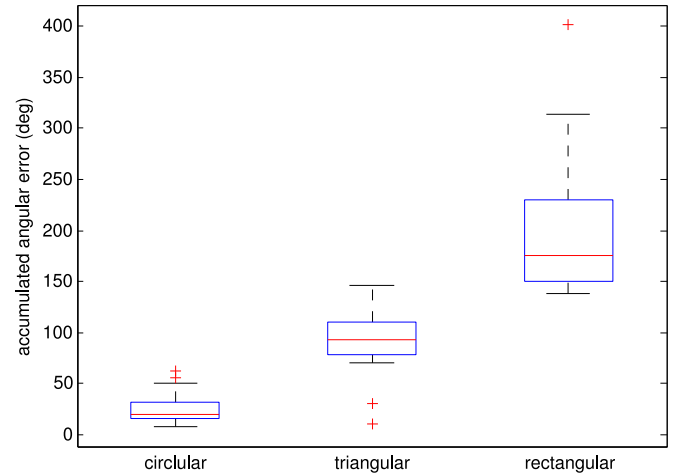


Fig. 11. Accumulated AE when the object enters the goal region. This metric reflects how much unnecessary rotation appeared in the transportation.

The step interval used in the offline video tracking is 1 s. The change of the orientation between two time steps t_0 and t_1 is defined as

$$D(t_0, t_1) = \left| \arccos \frac{\mathbf{a}(t_0) \cdot \mathbf{a}(t_1)}{\|\mathbf{a}(t_0)\| \|\mathbf{a}(t_1)\|} \right|. \quad (24)$$

The accumulated AE is calculated as

$$AE = \left| D(T_k, 0) - \sum_{t=1}^{T_k} D(t, t-1) \right|. \quad (25)$$

Note that the relative difference between the object's initial orientation and its orientation when it reaches the goal ($D(T_k, 0)$) is excluded, because we focus on quantifying the unnecessary effort on rotation (e.g., two continuous rotations that cancel out each other).

This metric will be zero if the transport process is ideal. Fig. 11 shows the box-and-whisker plot of the accumulated AE of the successful trials. Due to the length of the elongated rectangular object, randomness in the distribution of the pushing robots can cause a torque that is big enough to rotate the object rapidly. However, it is also due to the randomness in such rotations that this object will not always point with one of its ends toward the goal, which would cause the occluded surface for pushing to be very small.

VI. EXPERIMENTS WITH A MOVING GOAL

In a more complex environment, the goal may not be perceived from any position around the object. For example, there could be obstacles between the object and goal, or the distance between the two could be bigger than the range of sensors of the robots. The transport strategy as it stands cannot deal with such an environment. However, it is possible to adapt the goal in the strategy to expand the capability of the transport system.

If the goal is a mobile robot, it can change its position, while the object is being transported. It could navigate along a complex route and, thereby, lead the object to its final destination. How to implement such an intelligent goal robot is a research topic in itself [33]. In this section, we present an experiment in which




Fig. 13. Snapshots of one of the trials in the systematic experiments where the transport group pushes the object toward a teleoperated goal robot and thereby through an environment with obstacles. In the first snapshot ($t = 1$ min), the way points for the goal robot are highlighted.




Fig. 14. Traces of the object's centroid through an environment with obstacles (20 experimental trials).

- 2) J : Short range object sensor. This sensor can detect whether there are objects along the line of sight of it. Its normal vector is $(1.0, 0.57, -0.57)$ in the robot's local coordinate system. Its range is 40 cm.
- 3) K : Ambient light sensor. This omnidirectional sensor can detect whether the robot is directly illuminated by the goal light source. It simply checks the line of sight between the robot and the goal light.
- 4) W : Obstacle sensor. This sensor can detect whether there are obstacles along its line of sight. The environment boundary, other robots, and the embodiment of the goal light are considered as obstacles in the environment. The sensor's normal vector is $(1.0, -0.57, -0.57)$ in the robot's local coordinate system. Its range is 40 cm.

Note that these sensors are designed to directly meet the requirements of the behavior described in Section III-B, which simplifies the controller implementation.




Fig. 15. In this 3-D physics-based simulation, a swarm of robots are pushing an object (the blue capsule) toward a light source (the white sphere). The robots only push across the shadow side of the object, where the direct line of sight to the goal light is occluded by the object.

TABLE II
MAPPING FROM INPUTS TO MOTION OUTPUTS. STATES CORRESPOND TO THOSE IN FIG. 2 AND ARE GIVEN FOR INFORMATION ONLY

Inputs				Motion Outputs			
W	I	J	K	State	$speed$	yaw	$pitch$
0	0	0	—	Search Object	0.6	$[-0.03, 0.07]$	$[-0.1, 0.1]$
0	1	0	—	Approach Object	0.8	0.0	0.0
0	0	1	1	Move Around Object	0.3	$[0.02, 0.12]$	$[-0.3, 0.3]$
0	1	1	1	Push Object	0.0	$[-0.13, -0.03]$	$[-0.1, 0.1]$
0	0	1	0		0.2	$[-0.03, 0.17]$	$[-0.1, 0.1]$
0	1	1	0	Avoid Obstacles	0.7	$[-0.2, 0.2]$	$[-0.2, 0.2]$
1	—	—	—		—0.8	$[-0.3, 0.3]$	$[-0.3, 0.3]$

B. Robot Controller

The overall behavior of the robot follows the state machine description shown in Fig. 2. Due to the specific design of this




Fig. 16. Conceptual robot used in the simulations. In this image, the three thrusters of the robot (p_0 , p_1 and p_2) and the beams of three line-of-sight sensors (I , J and W ; all truncated) are shown. The robot also has an omnidirectional ambient light sensor (K ; not shown).

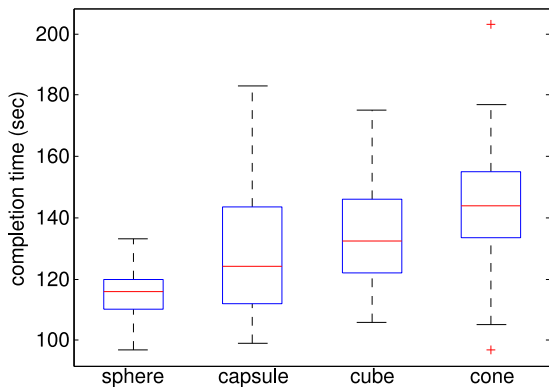


Fig. 17. Completion time of the simulation trials in a 3-D environment.

robot, both the state machine and low-level motion controller can be implemented using a single reactive controller.

Table II shows how the controller input (from the four binary sensors) is directly mapped to the motion output (speed, yaw, and pitch values). The parameters in column Motion Outputs were manually derived based on the overall behavior described in Section III-B. Where a range is provided, the motion output is randomly chosen at every control cycle following a uniform distribution in this range. We do not claim optimality of these parameters. They were chosen to give a working configuration for these proof-of-concept simulations.

Table II also indicates the equivalent states (see Fig. 2). Note that state “Check for Goal” is no longer required: the robot can check whether the goal is visible in an instant using its omnidirectional ambient light sensor (K).

C. Simulation Setup

One hundred robots were randomly placed in a bounded space of dimension 800 cm \times 500 cm \times 500 cm.

Consider the environment as a box of which the two diagonal vertices are positioned at (0, 0, 0) and (800, 500, 500) in the global coordinate system. The goal light was fixed at position (650, 250, 250). The object was initialized at (280, 250, 250), while its initial orientation was randomized using uniform spherical distribution.

Four types of objects were used:

- 1) a sphere with a radius of 41 cm;
- 2) a capsule with side length 60 cm and a radius of 30 cm;
- 3) a cube with side length 66 cm;

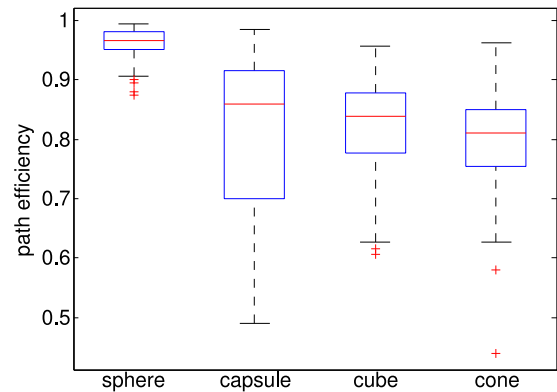


Fig. 18. Path efficiency of the simulation trials in a 3-D environment.

4) a cone with a height of 100 cm and a radius of 52 cm. The mass of these objects were all approximately 280 kg (calculated from their volumes using the density of water).

For each type of object, 100 simulation trials were run. When the centroid distance between the object and goal light was less than 90 cm, a trial was stopped, and considered successful. A trial was also stopped when 900 s elapsed.

D. Simulation Results

In all 400 trials, the object reached the goal within the time limit. The box plot of the completion times (in simulated seconds) for each of the objects is shown in Fig. 17. The path efficiency of the trials is shown in Fig. 18.

Typical situations of the four objects are shown in the accompanying video and the online supplementary material [32].

According to both of the numeric results and the direct observation, the transport task was successfully completed by the robots. Similar to the 2-D case, the shape of the object affected the performance of the strategy.

VIII. CONCLUSION

This paper introduced a cooperative transport strategy that uses a large number of relatively simple and small mobile robots to transport a large object that can occlude the robots perception of the goal. The strategy makes robots push along the surface of the object, where the robots’ line of sight to the goal is occluded by the object itself. By ensuring that the robots only push the object over the occluded surface, the object will eventually reach the goal (but the orientation of the object cannot be controlled). This paper focused on studying the strategy in a 2-D work space. A mathematical formulation of the strategy was provided. We proved that any convex-shaped object will always be successfully transported to the goal point and that the same is not necessarily true for objects of concave shape.

The main advantage of the occlusion-based cooperative transport strategy is that it is suitable for a decentralized system using a large number of relatively simple robots. The robots do not need to communicate (explicitly) with each other. The system is also fully scalable and not sensitive to the exact number of robots that are deployed; in fact, more robots make the strategy work better.




Fig. 19. On the perimeter of this concave object, both red and blue segments are occluded from the goal. The coordinates in the drawing are the forces brought by each of the segments measured in grid units assuming infinite number of point robots are uniformly distributed over these segments. The combined force brought by the robots at the blue segments is zero, whereas the combined force brought by the robots at the two red segments pushes the object away from the goal.

The strategy was implemented on a system of 20 physical e-puck robots. A systematic experiment was performed to verify the implementation using three particularly challenging types of objects. In 43 out of 45 trials in total, the objects were successfully transported to the goal. The self-correction effect introduced by the occlusion-based strategy can be clearly observed in these trials. Depending on the shape of the objects to be transported, the paths traced by them on average were 9.5% to 32.6% longer than the shortest possible path. The paths were compared with predictions from the mathematical model. While most individual paths differed substantially, their overall distribution showed a good correspondence. In an extended experiment, an extra e-puck was used as the goal. This goal robot was remotely controlled by a human operator. Following the path of the goal robot, the transport robots pushed the object in all 20 trials through an environment with obstacles.

A physics-based simulation was used to show an implementation of the transport strategy in a 3-D environment using a swarm of conceptual robots that have only four binary sensors. The simulation shows that the transport strategy has potential to be implemented in a 3-D environment using a large swarm of simple robots. For example, nanorobot swarms could transport materials such as drugs within the human body.

To the best of our knowledge, this is the first successful attempt of using a large number of autonomous robots to push a large nonspecific object. Moreover, the experiment using a mobile goal is a successful instance of human-robot interaction in which a human remotely controls a swarm of robots through a single agent robot. In future work, the goal robot could also be one of a series of way points formed by a group of robots (e.g., mimicking a trail of virtual pheromones [33], [34]). Such a system may accomplish a more complex cooperative transport task autonomously. The strategy itself may also be improved. For example, multiple layers of robots could push objects that are heavy but small in surface area.

APPENDIX A

ANALYSIS FOR CONCAVE OBJECTS

In Section IV, it has been proven that the combined force introduced by the transport strategy always reduces the distance between an arbitrarily convex-shaped object and the goal. This property may not hold for some extreme concave objects (depending on their relative distance and orientation to the goal). For instance, Fig. 19 shows a counter example with $c \cdot \mathbf{F} > 0$. In other words, the resultant force asserted by all robots will move the object away from the goal.

REFERENCES

- [1] J. Chen, M. Gauci, and R. Groß, "A strategy for transporting tall objects with a swarm of miniature mobile robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2013, pp. 863–869.
- [2] R. Groß and M. Dorigo, "Towards group transport by swarms of robots," *Int. J. Bio-Inspired Comput.*, vol. 1, no. 1/2, pp. 1–13, 2009.
- [3] M. Rubenstein, A. Cabrera, J. Werfel, G. Habibi, J. McLurkin, and R. Nagpal, "Collective transport of complex objects by simple robots: Theory and experiments," in *Proc. Int. Conf. Auton. Agents Multi-Agent Syst.*, 2013, pp. 47–54.
- [4] C. Kube and E. Bonabeau, "Cooperative transport by ants and robots," *Robot. Auton. Syst.*, vol. 30, no. 1–2, pp. 85–101, 2000.
- [5] G. A. S. Pereira, V. Kumar, J. Spletzer, C. J. Taylor, and M. F. M. Campos, "Cooperative transport of planar objects by multiple mobile robots using object closure," in *Experimental Robotics VIII*. New York, NY, USA: Springer, 2002, pp. 275–284.
- [6] J. Fink, M. Hsieh, and V. Kumar, "Multi-robot manipulation via caging in environments with obstacles," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2008, pp. 1471–1476.
- [7] K. Kosuge and T. Oosumi, "Decentralized control of multiple robots handling an object," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 1996, vol. 1, pp. 318–323.
- [8] B. Gerkey and M. Mataric, "Sold!: Auction methods for multirobot coordination," *IEEE Trans. Robot. Autom.*, vol. 18, no. 5, pp. 758–768, Oct. 2002.
- [9] Y. Hu, L. Wang, J. Liang, and T. Wang, "Cooperative box-pushing with multiple autonomous robotic fish in underwater environment," *IET Control Theory Appl.*, vol. 5, no. 17, pp. 2015–2022, 2011.
- [10] R. Groß, F. Mondada, and M. Dorigo, "Transport of an object by six pre-attached robots interacting via physical links," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2006, pp. 1317–1323.
- [11] K. Kashiwazaki, N. Yonezawa, M. Endo, K. Kosuge, Y. Sugahara, Y. Hirata, T. Kanbayahi, K. Suzuki, K. Murakami, and K. Nakamura, "A car transportation system using multiple mobile robots: iCART II," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2011, pp. 4593–4600.
- [12] S. Kernbach, D. Häbe, O. Kernbach, R. Thenius, G. Radspieler, T. Kimura, and T. Schmickl, "Adaptive collective decision-making in limited robot swarms without communication," *Int. J. Robot. Res.*, vol. 32, no. 1, pp. 35–55, 2013.
- [13] M. Gauci, J. Chen, W. Li, T. Dodd, and R. Groß, "Self-organized aggregation without computation," *Int. J. Robot. Res.*, vol. 33, no. 8, pp. 1145–1161, 2014.
- [14] E. Tuci, R. Groß, V. Trianni, F. Mondada, M. Bonani, and M. Dorigo, "Cooperation through self-assembly in multi-robot systems," *ACM Trans. Auton. Adapt. Syst.*, vol. 1, no. 2, pp. 115–150, 2006.
- [15] N. Miyata, J. Ota, Y. Aiyama, J. Sasaki, and T. Arai, "Cooperative transport system with regrasping car-like mobile robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 1997, vol. 3, pp. 1754–1761.
- [16] T. Sugar and V. Kumar, "Control of cooperating mobile manipulators," *IEEE Trans. Robot.*, vol. 18, no. 1, pp. 94–103, Feb. 2002.
- [17] Z.-D. Wang, Y. Hirata, Y. Takano, and K. Kosuge, "From human to pushing leader robot: Leading a decentralized multirobot system for object handling," in *Proc. IEEE Int. Conf. Robot. Biomimetics*, 2004, pp. 441–446.
- [18] Y. Aiyama, M. Hara, T. Yabuki, J. Ota, and T. Arai, "Cooperative transportation by two four-legged robots with implicit communication," *Robot. Auton. Syst.*, vol. 29, no. 1, pp. 13–19, 1999.

- [19] A. Trebi-Ollennu, H. Das Nayar, H. Aghazarian, A. Ganino, P. Pirjanian, B. Kennedy, T. Huntsberger, and P. Schenker, "Mars rover pair cooperatively transporting a long payload," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2002, vol. 3, pp. 3136–3141.
- [20] D. Stilwell and J. Bay, "Toward the development of a material transport system using swarms of ant-like robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1993, vol. 1, pp. 766–771.
- [21] S. Berman, Q. Lindsey, M. S. Sakar, V. Kumar, and S. Pratt, "Experimental study and modeling of group retrieval in ants as an approach to collective transport in swarm robotic systems," *Proc. IEEE*, vol. 99, no. 9, pp. 1470–1481, 2011.
- [22] K. M. Lynch and M. T. Mason, "Stable pushing: Mechanics, controllability, and planning," *Int. J. Robot. Res.*, vol. 15, no. 6, pp. 533–556, 1996.
- [23] M. J. Matarić, M. Nilsson, and K. T. Simsarian, "Cooperative multi-robot box-pushing," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 1995, vol. 3, pp. 556–561.
- [24] A. Sudsang and J. Ponce, "A new approach to motion planning for disc-shaped robots manipulating a polygonal object in the plane," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2000, vol. 2, pp. 1068–1075.
- [25] G. A. S. Pereira, M. F. M. Campos, and V. Kumar, "Decentralized algorithms for multi-robot manipulation via caging," *Int. J. Robot. Res.*, vol. 23, no. 7–8, pp. 783–795, 2004.
- [26] Z. Wang, Y. Hirata, and K. Kosuge, "Control multiple mobile robots for object caging and manipulation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2003, vol. 2, pp. 1751–1756.
- [27] C. R. Kube and H. Zhang, "Task modelling in collective robotics," *Auton. Robot.*, vol. 4, no. 1, pp. 53–72, 1997.
- [28] A. Becker, G. Habibi, J. Werfel, M. Rubenstein, and J. McLurkin, "Massive uniform manipulation: Controlling large populations of simple robots with a common input signal," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 520–527.
- [29] K. Lynch, "The mechanics of fine manipulation by pushing," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1992, vol. 3, pp. 2269–2276.
- [30] F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klapotcz, S. Magnenat, J.-C. Zufferey, D. Floreano, and A. Martinoli, "The e-puck, a robot designed for education in engineering," in *Proc. 9th Conf. Auton. Robot Syst. Compet.*, 2009, vol. 1, pp. 59–65.
- [31] S. Suzuki, H. Asama, A. Uegaki, S. Kotosaka, T. Fujita, A. Matsumoto, H. Kaetsu, and I. Endo, "An infra-red sensory system with local communication for cooperative multiple mobile robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 1995, vol. 1, pp. 220–225.
- [32] J. Chen, M. Gauci, W. Li, A. Kolling, and R. Groß. (2014). Online supplementary material. [Online]. Available: <http://naturalrobotics.group.shef.ac.uk/supp/2014-002/>
- [33] S. Nouyan, R. Groß, M. Bonani, F. Mondada, and M. Dorigo, "Teamwork in self-organized robot colonies," *IEEE Trans. Evol. Comput.*, vol. 13, no. 4, pp. 695–711, Aug. 2009.
- [34] K. Sugawara, T. Kazama, and T. Watanabe, "Foraging behavior of interacting robots with virtual pheromone," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, vol. 3, pp. 3074–3079.




Jianing Chen (S'11) received the B.Eng.(Hons) degree in automatic control and systems engineering from University of Sheffield, Sheffield, U.K., in 2010, and is currently working toward the Ph.D. degree at the same institution.

His research interests include self-organizing systems and computational intelligence.



Melvin Gauci (S'11) received the B.Eng.(Hons) degree (*summa cum laude*) in electrical engineering from University of Malta, Malta, in 2009 and the M.Sc.(Eng.) degree (with distinction) in automatic control and systems engineering, and also the Ph.D. degree, from University of Sheffield, Sheffield, U.K., in 2010 and 2014, respectively. His Ph.D. work investigated swarm robotic systems with minimal information processing.

He is a Postdoctoral Fellow with the Wyss Institute for Biologically Inspired Engineering, Harvard University, MA, USA. His research interests are broadly in robotics and computational intelligence, and specifically in embodied self-organizing systems and co/evolutionary machine learning.



Wei Li (S'13) received the B.Eng. degree in automation and the M.Eng. degree in control science and engineering from Harbin Institute of Technology, Harbin, China, in 2009 and 2011, respectively. He is currently working toward the Ph.D. degree with University of Sheffield, Sheffield, U.K.

His research interests include robotics and computational intelligence, particularly as applied to agent behavior learning and system identification.



Andreas Kolling (S'04–M'10) received the M.S. degree in computer science and the B.S. degree in mathematics from Jacobs University Bremen, Bremen, Germany, in 2006 and 2004, respectively. He received the Ph.D. degree in electrical engineering and computer science from University of California, Merced, CA, USA, in 2009.

From 2010 to 2012 he was a Postdoctoral Fellow with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA, and University of Pittsburgh, PA. From 2012 to 2013 he was a Postdoctoral Fellow at Linköping University, Sweden. Since October 2013 he has been a Lecturer with the Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, U.K. His research interests include planning for multirobot systems, pursuit evasion, graph theory, cooperative robotics, and machine learning. He has served as an Associate Editor for ICRA 2014–2015 and a PC member for AAAI 2013–2014 and AAMAS 2012–2014.



Roderich Groß (S'06–M'08–SM'12) received the Diploma degree (*summa cum laude*) in computer science from TU Dortmund, Germany, in 2001 and the Ph.D. degree in engineering sciences from Université libre de Bruxelles, Brussel, Belgium, in 2007.

From 2005 to 2009 he was a JSPS Fellow with Tokyo Institute of Technology, Japan; a Research Associate with University of Bristol, U.K.; a Marie Curie Fellow with Unilever, U.K.; and a Marie Curie Fellow with Ecole Polytechnique Fédérale de Lausanne, Switzerland. Since 2010 he has been with the Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, U.K., where he is currently a Senior Lecturer. His research interests include autonomous robots, swarm intelligence, and evolutionary computation. He has served as an Editor of IROS 2015; an Associate Editor of *Swarm Intelligence*, IEEE COMPUTATIONAL INTELLIGENCE MAGAZINE, IROS 2012–2014, and ICRA 2012–2015; and a Part Editor of the Springer Handbook of Computational Intelligence.