

Collective Transport of Complex Objects by Simple Robots: Theory and Experiments

Michael Rubenstein
Harvard University
mrubens@seas.harvard.edu

Adrian Cabrera
Harvard/EPFL
adricabdelu@gmail.com

Justin Werfel
Harvard University
justin.werfel@wyss.harvard.edu

Golnaz Habibi
Rice University
gh4@rice.edu

James McLurkin
Rice University
jmlurkin@rice.edu

Radhika Nagpal
Harvard University
rad@eecs.harvard.edu

ABSTRACT

Ants show an incredible ability to collectively transport complex irregular-shaped objects with seemingly simple coordination. Achieving similarly effective collective transport with robots has potential applications in many settings, from agriculture to construction to disaster relief. In this paper we investigate a simple decentralized strategy for collective transport in which each agent acts independently without explicit coordination. Using a physics-based model, we prove that this strategy is guaranteed to successfully transport a complex object to a target location, even though each agent only knows the target direction and does not know the object shape, weight, its own position, or the position and number of other agents. Using two robot hardware platforms, and a wide variety of complex objects, we validate the strategy through extensive experiments. Finally, we present a set of experiments to demonstrate the versatility of the simple strategy, including transport by 100 robots, transport of an actively moving object, adaptation to change in goal location, and dealing with partially observable goals.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Experimentation, Theory

Keywords

Transport, Multi-Robot Systems, Collective Intelligence

1. INTRODUCTION

Ants show an incredible ability to collectively transport complex irregularly shaped objects, even man-made ones, without *a priori* knowledge of object shape and mass [1, 2, 3]. This form of transport allows the ant colony to collectively achieve a task that individuals cannot achieve, as well as other benefits such as faster transport speeds, robustness to individual failure, and better adaptability to varying object sizes. Collective transport with multi-robot systems has the potential for many of the same benefits in a number

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.



Figure 1: (Left) 100 robots transporting an “H” object. (Right) Kilobots transporting an actively moving object.

of settings, such as agriculture, disaster relief, warehouses, mining, etc. Effective algorithms for collective transport would allow robots with fixed capabilities to cooperatively transport large, heavy objects, and adapt team size to transport a variety of object sizes efficiently and simultaneously.

In this paper, we focus on a problem of multi-agent coordination in collective transport: can a large group of agents successfully transport an arbitrarily shaped heavy object to a target destination, without *a priori* knowledge of the details of object shape and mass or even the number of other agents taking part in the goal? Studies suggest that ants achieve a high level of success with complex objects, not through coordinated deliberation and communication, but rather through decentralized actions of many individuals applying forces to a single object [1, 2]. In swarm robotics, researchers have developed collective transport strategies inspired by social ants [5, 6]; these algorithms have shown great promise, but have not been analyzed and generalized to complex shapes. Other more analytical approaches have also been developed; however, these tend to depend on sophisticated robots that maintain full knowledge of object shape, position, and orientation at all times [8, 9, 10].

In this paper we investigate a simple decentralized strategy for collective transport, in which each agent independently grips the object and applies a force in the direction of a goal, without any explicit coordination between agents. Using a theoretical model, we show that this strategy is provably guaranteed to successfully transport a complex object to a target location – even though each agent has no knowl-

edge of the object shape, weight, its own position on the object perimeter, or the position and number of other agents – so long as there are enough agents to overcome static friction. Furthermore, we can theoretically predict many important properties, such as speed of transport, minimum number of agents required, and scalability of performance as a function of number of robots, and show that the object trajectory and rotation are near optimal. We implement and evaluate the strategy through extensive experiments using two hardware testbeds, the Kilobot robot swarm [11] and the r-one robot swarm [12], and a wide variety of complex object choices. Our results show that the behavior of real robots is well captured by the theoretical model, including predicted performance and scalability, in spite of the many unmodeled sensing and motion errors that real robots exhibit. We also demonstrate an extension to the simple decentralized strategy for the case where the target location is only observed by a few robots; the extension makes use of an algorithmic analogy between collective transport and flocking with informed agents [15]. Finally, we present a set of demonstrations to show that, while the decentralized approach is seemingly very simple, it is still quite versatile. We show several complex collective transport scenarios, including transport by 100 robots (Fig. 1), transporting an actively moving object, and adapting to changes in goal location.

2. RELATED WORK

There has been a considerable amount of research on collective transport in social insects [1, 2, 3] as well as groups of robots [1, 4, 6, 9]. In this section we primarily focus on transport strategies for autonomous mobile robots that use grippers to move objects, and where collective behavior is required to move objects too heavy for a single robot.

In the control theory community, many approaches have been developed using centralized planners, leader-follower schemes, or distributed coordination between robots [8, 9, 10]. Much of this work focuses on a slightly different task called *distributed planar manipulation*, where robots transport a lightweight object of known shape along a pre-specified global trajectory with control of both object orientation and total force, e.g., in a factory setting. An advantage of these methods is that they are analytically tractable and generalize to complex object shapes. A disadvantage is that the task and methods rely on robots knowing global object position and orientation at all times, and many methods also rely on knowing position of all other team members. This requires sophisticated robots that can estimate object geometry and can track absolute position and orientation of the object and other robots. This limits the applicability and scalability of the strategies.

A complementary approach has been pursued by swarm robotics research, where decentralized transport strategies have been developed using inspiration from social insects or through evolutionary computation [1, 5, 6]. These algorithms rely on simple local sensing by individual robots, with no explicit knowledge of object shape or explicit coordination between robots. An advantage of these methods is that they rely on simple robot capabilities, and many have been experimentally tested on platforms such as the Swarmbots [5]. This research has also pursued interesting extensions, such as robots forming pulling chains. A disadvantage of the decentralized methods is that they have not been theoretically analyzed, and evaluation is done only in

simulation or through limited robot experiments; in one interesting exception [4], the cooperative behavior is analyzed but requires a leader-follower strategy. The lack of analysis makes it unclear whether the ant-inspired decentralized approaches generalize to more complex object geometries or larger numbers of robots. In general, for both types of prior research, experimental verification of the transport strategy has been limited to small numbers of robots (2-6) and simple object shapes (e.g., circles, squares, “L” shape).

Our work builds on the swarm robotics approach and we present an analysis of a simple decentralized strategy that generalizes to complex shapes; furthermore, we present experimental results with large numbers of robots and complex objects to show that the method is robust to noise and variation present in real robots.

3. PHYSICAL MODEL FOR COLLECTIVE TRANSPORT

The act of collective transport can be broken down into two phases. In the first phase, independent agents scattered in the environment seek out the object to be transported, and move themselves to a location around the object and grip it. The second phase involves the agents applying forces to the object to transport it to the goal. In this paper we explore a model for this second phase of collective transport. Because the agents are all attached to a single object, this object forms an implicit means of coordinating the effects of their actions. As we show in section 4.2, this model also provides insights that aid in the first phase of transport.

3.1 Physics-Based Model

We model the object to be transported as an arbitrary connected 2D shape, with N agents applying forces to it at arbitrary points. The object has mass m , and can slide across a level planar environment with coefficients of static and kinetic friction μ_s, μ_k respectively. We assume a quasi-static system, meaning the friction forces dominate inertial forces. With this assumption, we treat the model as a first-order system, allowing us to neglect inertial terms.

Each agent i applies a force \vec{f}_i to the object in the plane. We assume each agent can apply \vec{f}_i in any direction, independent of how or where it is attached. This can be achieved with an appropriate choice of robotic platform, e.g., using omnidirectional grippers [5] or holonomic movement [13]. We consider the force applied by each agent \vec{f}_i to be dependent on the agent’s velocity \vec{v}_i as

$$\vec{f}_i = C(V_{max} - \vec{v}_i \cdot \hat{d})\hat{d}, \quad (1)$$

where V_{max} is an agent’s maximum speed, $C = f_0/V_{max}$ where f_0 is the magnitude of the maximum force an agent can apply, and \hat{d} is a unit vector in the direction the agent is applying force. Intuitively, this means an agent applies its maximum force f_0 to the object when it is stationary, and can apply no force when moving at its maximum velocity in the direction \hat{d} . This is a simplified model of force application by mobile robots. For this paper we will assume that f_0 and V_{max} are identical for all agents, and all agents have speed in the range $0 \leq \|\vec{v}_i\| \leq V_{max}$.

Using this model, we investigate the following question: if all agents $\{1, 2, \dots, N\}$ apply forces in the direction of the goal, what is the object’s resulting translation and rotation? We assume the goal is defined and sensed by agents in such




Figure 2: Example of object with 2 agents (red circles) applying force in the goal direction. All r vectors originate at the marked object center-of-mass.

a way that all agent forces are parallel, e.g., agents may aim for a given compass direction or a distant point source. First, we show that the object's center of mass will travel in a straight line towards the goal and we compute the object's steady-state speed when its translational velocity is much greater than its rotational velocity. Second, we show that during transport an object's rotational velocity will become negligible compared to its translational velocity, with total rotation no more than 180° .

3.2 Translation

Multiple forces applied by the agents cause both translation and rotation of the object. Rigid body dynamics states that the translation of an object's center-of-mass due to multiple forces on the object is the same as if those forces were all applied at the object center-of-mass. The agent forces in the direction of the goal are opposed by friction \vec{f}_{fric} , so the total force on the object is

$$\vec{F}_{tot} = \sum_{i=1}^N \vec{f}_i - \vec{f}_{fric}. \quad (2)$$

If we assume the object is not yet moving or it is translating with negligible rotation, then its velocity is the same at all points: \vec{v}_i for all agents i equals the center-of-mass velocity \vec{v}_{obj} . This allows the following simplification of the force from all agents

$$\sum_{i=1}^N \vec{f}_i = NC(V_{max} - \vec{v}_{obj} \cdot \hat{d}) \hat{d}. \quad (3)$$

Because of the quasi-static assumption, if the object is moving, \vec{v}_{obj} is pointed in the direction of the total agent force on the object; therefore $\vec{v}_{obj} \cdot \hat{d} = \|\vec{v}_{obj}\|$. This allows us to rewrite eq. 2 as

$$\sum_{i=1}^N \vec{f}_i = NC(V_{max} - \|\vec{v}_{obj}\|) \hat{d}. \quad (4)$$

The condition for the robots to be able to overcome static friction and move an object initially at rest is

$$NCV_{max} > \mu_s mg. \quad (5)$$

The steady-state velocity \vec{v}_{ss} for a moving object is that at which kinetic friction $\vec{f}_{fric} = -\mu_k mg \hat{d}$ balances the force exerted by agents. Setting this equal to the sum of the agent forces and solving for $\|\vec{v}_{obj}\|$ gives

$$\|\vec{v}_{obj}\| = \|\vec{v}_{ss}\| = \left(V_{max} - \frac{\mu_k mg}{NC} \right). \quad (6)$$

Thus the steady-state speed $\|\vec{v}_{ss}\|$ at which N agents can transport an object approaches the maximum speed of an unencumbered agent V_{max} asymptotically as N increases.

3.3 Rotation

The translation speed calculation above assumes the object is translating much faster than rotating, allowing us to neglect the effects of rotation on an agent's velocity. Here we justify that assumption by showing that for any initial configuration, the object undergoes a maximum rotation of 180° in a single direction. As inertial terms are neglected due to the quasi-static assumption, there will be no overshoot or oscillation of the object rotation. This stable bounding of rotation guarantees that rotational speed will ultimately become negligible, justifying the approximation in sec. 3.2.

We define \vec{r}_{ref} as a vector from the object's center of mass to an arbitrary but fixed reference edge location on the object (see fig. 2). This vector \vec{r}_{ref} forms an angle θ with the goal direction. We also define a vector \vec{r}_i from the object's center-of-mass to the location of each agent i ; the angle between \vec{r}_{ref} and \vec{r}_i is a constant θ_i (fig. 2).

The total torque on the object due to agent forces is

$$\sum_i \vec{\tau}_i = \sum_i \left(\|\vec{r}_i\| \|\vec{f}_i\| \sin(\theta + \theta_i) \right) (-\hat{k}), \quad (7)$$

where \hat{k} is a unit vector pointing out of the plane. If the object is moving, then agent i 's speed in the direction of the goal is

$$\vec{v}_i \cdot \hat{d} = (\vec{v}_{obj} + \vec{v}_{r,i}) \cdot \hat{d} = v_{obj} + \vec{v}_{r,i} \cdot \hat{d}, \quad (8)$$

where v_{obj} is the object's center-of-mass speed (as per section 3.2, its direction is toward the goal), and $\vec{v}_{r,i}$ is agent i 's velocity due purely to object rotation; $\vec{v}_{r,i}$ has magnitude $\|\vec{r}_i\| \dot{\theta}$ and angle $(\theta + \theta_i + \pi/2)$, so that $\vec{v}_{r,i} \cdot \hat{d} = \|\vec{r}_i\| \dot{\theta} \cos(\theta + \theta_i + \pi/2) = -\|\vec{r}_i\| \dot{\theta} \sin(\theta + \theta_i)$. Therefore the force applied by each agent is

$$\vec{f}_i = C(V_{max} - v_{obj} + \|\vec{r}_i\| \dot{\theta} \sin(\theta + \theta_i)) \hat{d}. \quad (9)$$

From eqs. 7 and 9 the total torque due to the agents is

$$\sum_i \vec{\tau}_i = - \sum_i \left(\|\vec{r}_i\| C(V_{max} - v_{obj} + \|\vec{r}_i\| \dot{\theta} \sin(\theta + \theta_i)) \cdot \sin(\theta + \theta_i) \right) \hat{k}, \quad (10)$$

which using harmonic addition can be written as

$$\sum_i \vec{\tau}_i = (K_1 \sin(\theta + K_2) - K_3 \dot{\theta}) \hat{k}, \quad (11)$$

with $K_3 = C \sum_i \|\vec{r}_i\|^2 \sin^2(\theta + \theta_i)$ and constants

$$K_1 = -C(V_{max} - v_{obj}) \cdot \sqrt{\left(\sum_i \|\vec{r}_i\| \cos \theta_i \right)^2 + \left(\sum_i \|\vec{r}_i\| \sin \theta_i \right)^2} \quad (12)$$

$$K_2 = \tan^{-1} \left(\frac{\sum_i \|\vec{r}_i\| \sin \theta_i}{\sum_i \|\vec{r}_i\| \cos \theta_i} \right). \quad (13)$$

From the quasi-static assumption, without inertial terms the sign of the object's rotational velocity matches that of the total agent torque on the object; in particular, $\sum \vec{\tau}_i = 0$ iff $\dot{\theta} = 0$. We can express this as $\sum \vec{\tau}_i = \dot{\theta} \mathcal{F}$, where \mathcal{F} is a strictly positive-valued function.

We can then rewrite eq. 11 as an expression for $\dot{\theta}$:

$$\dot{\theta} = \frac{K_1 \sin(\theta + K_2)}{\mathcal{F} + K_3}. \quad (14)$$

To characterize the object’s rotation, we are concerned only with the sign of $\dot{\theta}$ as a function of θ . The denominator of eq. 14 is strictly positive, so $\dot{\theta}$ takes the sign of the numerator. Friction torque opposes rotation, so it will affect the speed of rotation but not its sign.

There are two cases to consider:

(A) $K_1 = 0$. In this case $\sum \vec{\tau}_i = 0$ for any θ ; there is no preferred direction, and no rotation will occur regardless of the initial angle.

(B) $K_1 < 0$. In this case, there are two equilibrium points (i.e., $\sum \vec{\tau}_i = \dot{\theta} = 0$), where $\sin(\theta + K_2) = 0$. Around $\theta = -K_2$, rotation acts to oppose small angular displacements ($\dot{\theta}(-K_2 + \epsilon) < 0, \dot{\theta}(-K_2 - \epsilon) > 0$ for $\epsilon > 0$); around $\theta = \pi - K_2$, small initial angular displacements increase.

Thus, eq. 10 has either (A) a neutral equilibrium for all angles, or (B) one stable equilibrium at $\theta = -K_2$ and one unstable equilibrium at $\theta = \pi - K_2$, with K_2 given by eq. 13. In the former case, there is no rotation; in the latter, for any initial orientation, the object will rotate away from the unstable equilibrium until it reaches the stable equilibrium, a rotation of no more than 180° . This bound implies that rotation must stop or become negligible.

3.4 Key Properties of the Model

Using our physics-based model for collective transport, we can make the following claims about the strategy in which all agents apply force in the goal direction.

1. **Transport Success:** Given enough agents to overcome static friction (eq. 5), we can guarantee the object will eventually reach the goal following an optimal straight path, with at most 180° of initial rotation.
2. **Agnostic:** The strategy is agnostic to the object shape, location of object center-of-mass, attachment location of agents, and number of agents (assuming they can overcome static friction). By agnostic, we mean that the agent’s behavior and transport success do not change based on any of these factors.
3. **Scalability:** From eq. 6, we can compute the steady state transport speed as a function of object mass, μ_k , number of agents, and agent parameters C, V_{max} . This allows us to predict how the performance of this strategy scales with the number of agents and the object mass. In particular, the transport speed for a given object will approach the maximum agent velocity V_{max} asymptotically as $1/N$ where N agents are participating. Also, for a given μ_k and agent force model, the transport steady-state speed is determined solely by the ratio of object mass to the number of agents, and does not decline due to conflicts as more agents participate.

These properties suggest that simple agents — without any knowledge of object shape, their attachment point to the object, or complex inter-agent communication — can be very effective at collective transport. Furthermore we can strongly reason about the collective behavior, such as path and speed traveled, even for complex objects.

However, the model also makes some simplifying assumptions that may not be true for real robots. For example, we

assumed all agents are identical, massless, frictionless, and can apply force instantaneously without error in the desired direction. Due to sensing and locomotion errors, real robots may take time to apply force in the goal direction and may not always apply a force directly towards the goal. Another simplifying assumption is that all agents know the direction of the goal at all times. While there are many examples of all agents knowing the goal location, such as in social insects [2] or in some robot systems [9], there are also many scenarios where this may not be true. For example, visual occlusion of the goal or the use of low-cost short-range sensors may prevent some robots from directly observing the goal.

In the next section we present a series of robot experiments using real robot hardware to demonstrate that the key properties of the physical model still hold, even in systems that do not perfectly match the simplifying assumptions made for the model.

4. COLLECTIVE TRANSPORT: EXPERIMENTAL VALIDATION

In this section we present a series of robot experiments using two hardware platforms — the open-source, commercially available Kilobot [11] and the open-source r-one [12] — to both verify and extend the collective transport model.

First, with the Kilobot platform, we demonstrate that the simple physics-based model is in fact a good predictor of the second phase of collective transport. The *transport success*, *agnostic*, and *scalability* properties still hold, even when implemented with noisy and very limited sensing and locomotion. Then, using the more capable r-one platform, we use a gripper design compatible with the physics-based model, and demonstrate both the first and second phase of collective transport. Additionally, the r-one platform allows us to present an extension where only a few robots can directly sense the goal location. We show that by an analogy between heading alignment in flocking and collective transport, the r-ones can achieve the same properties with only a few informed agents by exploiting inter-agent communication. Finally, the generality of the approach allows us to demonstrate some unique collective transport tasks, such as changing goal location mid-way through transport and transporting an actively moving or “squirming” object.

4.1 Model Validation

To show that the *transport success*, *agnostic*, and *scalability* properties hold in a real robot system, we used the Kilobot robot platform (fig. 3). The Kilobot is a low-cost robot that allows one to easily test collective algorithms with large numbers of robots. The Kilobot moves on a flat surface using differential drive with two vibration motors, and has a sensor to detect local ambient light levels; this allows us to implement phototaxis (i.e., moving towards a light source). With only a single light sensor, the implementation of phototaxis on Kilobots requires a slight rotation back and forth across the direction of travel in order to determine the direction to the light. This slight side-to-side rotation will cause the force applied by the robot to not always be perfectly aligned to the light direction, and on rare occasions sensor or locomotion noise may cause a robot to apply force in a completely different direction to the light source, taking tens of seconds to re-align to the light direction. By marking the target location with a light source (fig. 3), this phototaxis behavior allows Kilobot to apply forces in the goal direction.



Figure 3: (Upper left) A single Kilobot robot, with motors (A), light sensor (B), and forward direction (C) marked. (Upper right) A Kilobot robot in an object gripping ring. (Bottom) Experimental setup showing a test object with robots (D), the goal light source (E), and camera tracking system (F).

The objects to be transported are laser-cut from foam core poster-boards, and elevated to the same height as a Kilobot using metal legs. The object’s center-of-mass path and orientation are tracked using an overhead camera system. Since Kilobot does not have any way to grip the object, we emulate a gripper by creating circular rings along the edge of the object to be transported, and place the robots within these rings (see fig. 3). A Kilobot can rotate within this ring, allowing it to apply forces to the object in any direction, matching the model.

The first set of Kilobot experiments is designed to evaluate the *transport success* and *agnostic* properties of the collective transport model. This means that the object’s center-of-mass should always move in a straight line to the goal independent of the object shape, center-of-mass location, number of robots, and robot attachment points. Also, the object should rotate no more than 180° .

To test this we designed four different experiments that varied different factors, shown in fig. 4: (A) *Shape of objects*: four different object shapes with the same mass, (B) *Object center-of-mass*: for a single shape, placing a metal weight in one of four locations, (C) *Robot positions*: for a single shape, using different locations of the six robots, (D) *Number of robots*: for a single shape, using 4, 8, or 12 robots. All objects are approximately 20 cm in size, and were started in the orientation shown in fig. 4. The objects were then transported until they reached a goal approximately 80 cm away. The object’s center-of-mass x,y position and orientation was recorded ten times a second for every trial. Each initial condition (4 different shapes, 4 center-of-mass locations, 4 different starting locations, and 3 different numbers of robots) was tested three times for a total of 45 trials.

Fig. 5 shows the x,y paths of some representative transport runs. As can be seen, the paths of the objects form a relatively straight line towards the goal. The divergence of the paths as they approach the goal is because the light




Figure 4: Objects and starting orientations for experiments. (A) Four test shapes with the same mass. (B) Same shape with four locations of center-of-mass, using metal washers to shift center-of-mass. (C) Same shape with different robot attachment positions. (D) Same shape moved by increasing numbers of robots. (E) The order in which robots are added to measure change of speed vs. number of robots. (F) Objects that keep weight proportional to number of robots for 32,16,8,4,2 robots. Note: objects without robots shown had robots in every circular ring.

source marking the goal is a light bulb, not a point light source, and the spread is approximately the same as the diameter of the bulb (6 cm). To quantitatively show how straight the paths are, we computed a *straightness-metric* for each path which is $dist_{min}/dist_{traveled}$ where $dist_{min}$ is the distance between object starting point and object ending point, and $dist_{traveled}$ is the length of the actual path traveled by the object. The average *straightness-metric* for all 45 runs was .997 and the minimum was .987.

For experiment set (B), the objects undergo rotation that can be predicted from the location of the center of mass and the locations of the robots relative to the center of mass. Fig. 5 shows with each starting orientation, the object rotates to approximately the same equilibrium orientation, and does so by rotating less than 180° , as predicted by the model.

Another key property of the physical model is the *scalability* of transport. This property predicts that if all other factors are kept the same, the steady state speed of transport will asymptotically approach the maximum agent velocity V_{max} as more agents participate. To demonstrate this, we created experiment set (E) where a single object shape is transported by increasing numbers of robots. The steady state speed of the object was measured. At least three robots were needed to overcome static friction for this object. As fig. 6 shows, the steady state object speed asymptotically approaches .55 cm/s, which is near the average measured speed of .58 cm/s for an unencumbered Kilobot moving towards a light source.

The *scalability* property also predicts that the steady state object speed v_{ss} will stay constant with varying object masses m and number of robots N , so long as the ratio m/N is maintained. To demonstrate this, we created a series of objects that allowed 2,4,8,16, and 32 robots to transport the




Figure 5: (Left) The recorded path of object transport under varying conditions and objects. Path colors for experiment (A) are red, (B) are green, (C) are black, (D) are blue. The object starts at (0,0) and the goal is on the right. (Right) The recorded orientation of object transport under varying starting orientations for experiment (B). Illustrations of three object orientations are shown; the object is a black box and center-of-mass location is the dot.




Figure 6: (Left) Screen captures of an example experiment. (Center) The speed of transport as number of robots increases if all other factors are kept constant, with 4 experiments per robot number. The blue line shows the least squares best fit of the form $a_1 - a_2/N$ where a_1 and a_2 are constants, the same form as eq. 6. (Right) The speed of transport as the number of robots increases proportionally to the object mass, with at least 5 experiments per robot number. The black line shows the average speed for all experiments in this graph. In both graphs, the result from each experiment is marked with a blue circle, and the average for each N is marked in red.

object while keeping the mass/robot constant (experiment F in fig. 4). The object was constructed so that its mass maintained a ratio of 2.2 grams per robot. Fig. 6 shows that each of these objects is transported at approximately the same steady state speed, further verifying the scalability of the transport model.

4.2 Attachment and Transport

For the second part of the experimental validation, we use a more capable robot platform, the r-one (fig. 7), which is a fully autonomous open-source robot platform [12]. The r-one can move on the floor using wheeled differential drive, detect the direction of a light source using four light sensors, and determine if it has collided with objects in the environment using a bump sensor. The r-one also has an infra-red (IR) communication system that can detect heading difference between itself and neighbors, allowing each robot to determine the relative orientation between itself and its neighbors. We use an overhead camera system with AprilTag fiducial markers [14] placed on each robot and the object to automatically log robot and object position and orientation during experiments. For object transport, we designed an omni-directional gripper (fig. 7) inspired by the

SwarmBot gripper design [6]. This gripper surrounds the r-one, and has a bearing-like design that allows the gripper to freely rotate around the robot. The gripper can attach to the object using a gendered Velcro lining. The object to be transported is a flat piece of wood with Velcro along its perimeter. As the r-one can freely rotate within the gripper to apply forces in any direction, it is compatible with the physical model.

This r-one platform allows us to evaluate both phases of collective transport, where the individual robots approach and grip (phase one) and then transport the object to the goal (phase two) using the same decentralized strategy as before. The individual robots use a simple finite state automaton program to switch between two states: finding and attaching to the object, and moving towards the goal marked by the light source. In the starting state, each robot spirals outward from its starting location until it bumps the object and attaches to it. Once attached, the robot detects the direction of the light source using its four light sensors, and maintains its heading towards the light source while moving forward, thus applying a force on the object in the goal direction.

Note that the robot does not plan or control exactly where it attaches along the object perimeter; however, we know from the model in section 3 that robot positions should not affect transport success. Also note that the robots are not synchronized, so they may attach and start attempting to move the object towards the goal at different times.

Fig. 8 shows an experimental run of the full transport algorithm, showing the paths of the three robots before and after attachment as well as the path of the object center-of-mass. Five such tests were run, where three robots transport a triangular-shaped object (side length 30 cm and weight 650 grams) a distance of approximately 2 meters starting from a detached state. For the five tests, the paths of the object center-of-mass are shown in fig. 8 and the average *straightness-metric* was .934. The results suggest that the r-one second phase transport behavior closely matches the Kilobots, as expected by the model. The *straightness-metric* value is lower with the r-one than with Kilobots, because the r-one uses multiple uncalibrated light sensors to determine the direction to the goal, and variances between these sensors result in larger errors in moving towards a light source.

4.3 Consensus-Based Collective Transport

In many scenarios, not all agents may be capable of detecting the goal location, e.g., due to visual occlusion of the goal by the object, or due to only some robots having GPS capabilities. A fundamental assumption behind the decentralized strategy presented so far is that all agents align their forces by orienting their motion towards the common goal direction. This reliance on alignment of forces and movement allows us to make an analogy between collective transport and another well-known collective behavior: flocking[15, 16]. In flocking, agents observe the headings of their neighbors and use that to modify their own heading; this iterative process results in the whole group probably achieving consensus and aligning their heading even in the absence of a target destination. Couzin et al. [15] demonstrated an extension of this process where a few informed individuals, who know the goal location, can influence the whole flock to align in the desired direction.

Using this flocking consensus behavior, we can extend the basic decentralized strategy from section 3 to the scenario where only a few agents observe the goal, if all agents are allowed to communicate relative heading with their neighbors. This strategy was first introduced and explored experimentally in [7] using 4 Swarmbots and circular objects; our main contribution is the theoretical guarantee that the approach will work correctly for more complex shapes and scenarios. We tested this strategy using the r-one robots, that can communicate heading and implement flocking motion [12].




Figure 7: (Left) Picture of an r-one with an omnidirectional velcro gripper (A), light sensors (B), AprilTag (C) and IR communication sensors (D). (Right) Picture of 4 robots pushing the object to a light source.

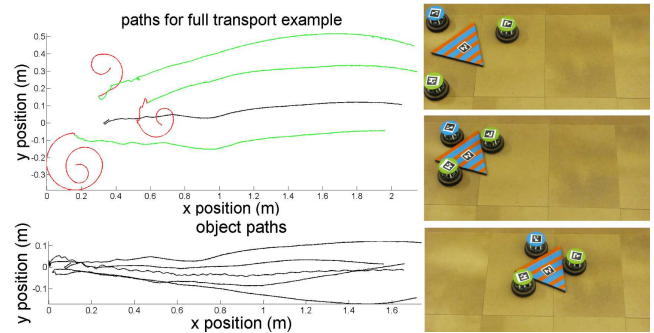


Figure 8: (Upper left) The paths of three robots (red during phase 1 and green during phase 2) and object (black). (Lower left) The path of object center-of-mass over five trials. (Right) Image of robots in starting position, after attachment, and after transport. For all cases, the goal is on the right.

Algorithmically, each robot uses one of two behaviors. A robot that can see the goal uses the strategy from section 4.2, which is to align towards the goal while moving forward. A robot that cannot see the goal uses its IR sensors to determine the heading difference between itself and its neighbors. It rotates in order to minimize its total heading or consensus error between itself and all its N neighbors, including neighbors that can see the goal. Consensus error is defined as $\sum_{i=1}^N \|\theta_i\|$, where θ_i is the heading difference between a

robot and its neighbor i , and $-\pi \leq \theta_i \leq \pi$. As shown in [15, 16], if the robots form a connected network and at least one can see the goal, this results in the robots achieving heading consensus in the direction of the goal. If the robots also move forward as they adjust their heading, then once heading consensus is reached, all robots will be applying forces towards the goal, matching the model in section 3.

To test this consensus-based collective transport, we ran five trials where four r-ones were attached to the object in random orientations and positions and only one of the robots was allowed to sense the light direction. For these five trials, the object was successfully transported to the goal, and the path of the object had an average *straightness-metric* of .921 which is similar to the case where all robots can sense the goal. Since robots are aligning and transporting at the same time, not all robots will be applying force to the object in the correct direction. We expect that transport speed should be higher when more robots are aligned. Fig. 9 shows how the average consensus error relates to overall transport speed. As expected, the scatterplot shows an obvious inverse relationship between object speed and consensus error over the five runs. The second plot shows the speed and error over time for a single experiment, including the initial startup phase when robots are not yet aligned as well as a manual disturbance phase where the robots were purposefully rotated to break alignment consensus. In both cases, we see that high alignment error results in low speed, due to many opposing forces on the object, but that as the headings align, the robots are able to successfully recover and transport the object.

4.4 Complex Transport Tasks

As the described method for collective transport can succeed under a wide range of conditions, it allows for demon-

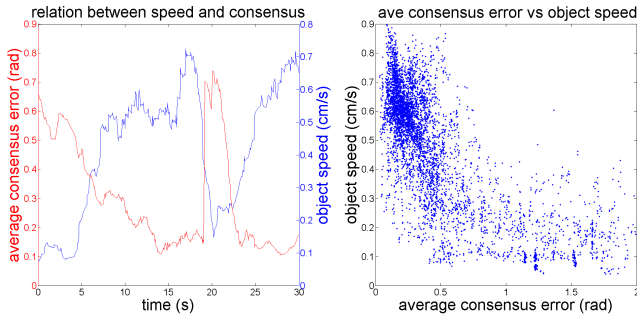


Figure 9: (Left) The average consensus error and object speed during a test of consensus-based transport during startup and during a manual disturbance (at time=18 s) to robot headings. (Right) A scatterplot of object speed vs. average consensus error for 1/10th of a second intervals during 5 tests of consensus-based transport.




Figure 10: The path of an object during four experiments where the goals changes mid-experiment. Object starts on left, moves towards first goal to the right, then goal moves to lower left.

strations that are unique to this method. For example, the simplicity and scalability of this method allowed for a demonstration of 100 Kilobots transporting an object (fig. 1). Since this method for transport is agnostic to object shape, it can even transport an object whose shape is slowly changing or an actively moving shape. In nature, ants have been observed carrying wriggling worms, but this has never been shown in artificial systems. To demonstrate this, an experiment was designed where Kilobots transported a shape that had an actively articulated joint, which caused the object to “squirm” while it was transported, see fig. 1. Additionally, as the algorithm is memoryless, it can allow changing or moving goal location. Fig. 10 shows the path of transport during an experiment where the goal location is changed mid-experiment.

5. CONCLUSION

In this paper, we investigate a simple decentralized method for collective transport and present both theory and experiments to evaluate this method. Using a physics-based model, we are able to show that, given enough agents to overcome static friction, the method is guaranteed to move the object in an optimal straight-line path to the goal with no more than 180° of rotation. Additionally, this method does not require agents to know object shape, center-of-mass, precise attachment points, or number of other agents. This method allows for agents with limited capabilities to complete complex transport tasks, provided an agent can determine goal direction. We also show how simple inter-agent communication of heading, similar to flocking, can allow agents to collectively transport objects even when some or most lack knowledge of goal location.

For future work, this physical model may be useful for analyzing other collective transport strategies, for example strategies where robots can only apply forces in a limited direction (e.g., only push) and therefore may need to detach and reattach in different locations to move the object [1]. Another area of future research is to compare our model to transport by ants, such as *Pheidologeton diversus* [2], where ants appear to apply force in the nest direction by surrounding the object such that ants in front are walking backwards and pulling, ants in the back push, and ants on the sides shuffle sideways [2]. By designing specific objects and scenarios [3], it may be possible to compare the behavior of ant transport with the results from our models and experiments.

6. ACKNOWLEDGMENTS

Thanks to Amanda Boone and Kevin Koch for the design and assembly of the r-one gripper rings. This research was funded by the Wyss Institute at Harvard, and NSF Grants CCF-0926148, CCF-0643898, and CPS-0931871.

7. REFERENCES

- [1] E. Bonabeau *et al*, *Swarm intelligence: From Natural to Artificial Systems*, Chapter 7, 1999.
- [2] M. Moffett, *Cooperative Food Transport By an Asiatic Ant*, *National Geographic Research*, vol. 4 (386-394), 1988.
- [3] S. Berman *et al*, *Study of Group Food Retrieval by Ants as a Model for Multi-Robot Collective Transport Strategies*, *RSS 2010*.
- [4] D. Stilwell, J. Bay, *Toward the development of a material transport system using swarms of ant-like robots*, *ICRA 1993*.
- [5] R. Gross, M. Dorigo, *Towards Group Transport by Swarms of Robots*, *International Journal of Bio-Inspired Computation*, 1(1-2):1-13, 2009.
- [6] R. Gross *et al*, *Object Transport by Modular Robots that Self-Assemble*, *ICRA 2006*.
- [7] A. Campo *et al*, *Negotiation of goal direction for cooperative transport*, *ANTS 2006*.
- [8] G. Montemayor, J. Wen, *Decentralized Collaborative Load Transport by Multiple Robots*, *ICRA 2005*.
- [9] Q. Lindsey, M. Shomin, V. Kumar, *Cooperative Quasi-Static Planar Manipulation with Multiple Robots*, *IDETC/CIE 2010*.
- [10] H. Su, V. Krovi, *Decentralized Dynamic Control of a Nonholonomic Mobile Manipulator Collective: a Simulation Study*, *ASME-DSCC 2008*.
- [11] M. Rubenstein *et al*, *Kilobot: A Low Cost Scalable Robot System for Collective Behaviors*, *ICRA 2012*.
- [12] J. McLurkin *et al*, *A Low-Cost Multi-Robot System for Research Teaching and Outreach*, *DARS 2010*.
- [13] P. Muir, C. Neuman, *Kinematic Modeling for Feedback Control of an Omnidirectional Wheeled Mobile Robot*, *ICRA 1987*.
- [14] E. Olson, *AprilTag: A Robust and Flexible Visual Fiducial System*, *ICRA 2011*.
- [15] I. Couzin, J. Krause, N. Franks, S. Levin, *Effective Leadership and Decision Making in Animal Groups on the Move*, *Nature*, 433 (513-516) 2005.
- [16] H. Tanner, A. Jadbabaie, G. Pappas, *Stable flocking of mobile agents, Part I: Fixed topology*, *CDC 2003*.