Altruistic Relationships for Optimizing Task Fulfillment in Robot Communities

Christopher M. Clark¹, Ryan Morton², and George A. Bekey³

Computer Science, California Polytechnic State Univ., San Luis Obispo, CA, 93407 USA, ¹cmclark@calpoly.edu, ²spacemort@gmail.com, ³gbekey@calpoly.edu

1 Abstract

This paper introduces the concept of a multi-robot community in which multiple robots must fulfill their individual tasks while operating in a shared environment. Unlike typical multi-robot systems in which global cost functions are minimized while accomplishing a set of global tasks, the robots in this work have individual tasks to accomplish and individual cost functions to optimize (e.g. path length or number of objects to gather).

A strategy is presented in which a robot may choose to aid in the completion of another robot's task. This type of "altruistic" action leads to evolving altruistic relationships between robots, and can ultimately result in a decrease in the individual cost functions of each robot. However, altruism with respect to another robot must be controlled such that it allows a relationship where both robots are altruistic, but protects an altruistic robot against a selfish robot that does not help others.

A quantitative description of this altruism is presented, along with a law for controlling an individual's altruism. With a linear model of the altruism dynamics, altruistic relationships are proven to grow when robots are altruistic, but protect an altruistic robot from a selfish robot. Results of task planning simulations are presented that highlight the decrease in individual robot cost functions, as well as evolutionary trends of altruism between robots.

2 Introduction

As the number of mobile robots increases within our homes, industry, and for scientific exploration, there is an increase in the number of situations where many robots will have to work together within a common workspace. Such situations will promote the necessity for a peaceful and constructive Multi-Robot Community (MRC).

In this work, we depart from the situation in which a group of robots are designed, owned, or operated by a *single* individual or organization with the purpose of achieving a common goal. Here, an MRC consists of robots that may be designed, owned, or operated by *several* individuals or organizations that may have different goals. Robots may be designed to act in their own best interest and accomplish their own goals without concern of the goals of other robots in the community. That is, robots in an MRC may be selfish.

This work proposes that robots in an MRC can achieve their goals more efficiently through the use of altruism. We define an altruistic robot as one which assists others in the attainment of their goals even if such actions may be harmful to itself. Specifically, if robots are willing to perform the tasks of other robots while decreasing their own efficiency in the short run, large gains in individual and global efficiency can be made over long time horizons.

As an example, consider two autonomous robots commissioned to pick up courier mail from around the city. While these robots might be owned and operated by different and possibly competing organizations, it may be in the best interest of both robots for them to accomplish each other's tasks when appropriate. That is, if robot A's pickup location is closer to robot B, then robot B should consider performing this pickup for robot A. This would be an altruistic action for robot B because it reduces efficiency without accomplishing any of its own pickups. However, robot A might later reciprocate this altruistic action, thereby building an altruistic relationship as a result of which both A and B may complete their tasks more rapidly or with less total distance traveled.

In the next section, related work is briefly covered. Section 3 defines a framework for an MRC as well as the problem being addressed in this paper. Section 4 presents a method for robots to control their altruism within the MRC framework, highlighting the ability for robots to create altruistic relationships while protecting them against purely selfish robots. In section 5, simulation results are presented. Finally, conclusions and future work are provided in section 6.

3 Background

Multi-Robot Systems (MRS) have been an active area of robotics research [2], due to the several potential advantages over single robot systems. Namely, they offer the possibility for greater spacio-temporal sampling, force multiplication, and robustness to failure. Advancements in the areas of MRS mission planning, MRS motion planning, MRS localization, MRS mapping, and the most related subject of task allocation have occurred over the last two decades.

3.1 Task Allocation

Of particular relevance to this work is the MRS task allocation problem, in which the MRS must accomplish a set of tasks characterized by their geo-

graphic location. The problem is to determine the optimal assignment of task points to robots, along with the optimal sequence for robots to visit these task points that minimizes the time to visit all task points. This is a variation of the Multi-Traveling Salesperson Problem (MTSP), a problem with many applications but no polynomial time solution (e.g. [12]). Regardless, many good heuristic driven methods have been developed that yield sub-optimal solutions.

One popular method of assigning tasks to robots in an MRS is to use a Market-Based Auction approach [6]. In this method, tasks are auctioned off to robots with the highest bids. Bids are typically based on the ability of the robot to accomplish the task, while considering the additional cost of traveling to the task site. While this method is not guaranteed to find optimal solutions, it is efficient and can lead to near optimal solutions.

3.2 Altruism

In the literature on robotics there are extensive treatments of cooperation among robots, but little mention of altruistic behaviors. Cooperation may in fact be altruistic, but it is generally not described in those terms in the literature. Examples of work involving cooperation are [4], [3], [10], and [8].

Work directly involving altruism includes that of [11] and [7] who describe robot behaviors in terms of a satisfaction index and transmission/reception of signals from other robots. A robots progress in a given task can be measured by its satisfaction in the task, which corresponds to the fitness or performance index indicated above. Thus, a robot needing help with a task may emit an attraction “please help me” signal. Lucidarme et. al. [7] propose an altruism vector based on the satisfaction index of a robot and the signals emitted by other robots; a given robot decides on altruistic behavior based on the magnitude of this vector. Similarly, [1] describes a software architecture for robot colonies based on robot tropisms, defined as their likes and dislikes. Reinforcement of particular behaviors strengthens them in future scenarios. Here too a robot can call for help to other robots when it needs assistance in moving heavy objects beyond its capability.

The emergence of cooperative behaviors has been studied extensively in game theory, under the name Prisoners Dilemma, e.g.[9]. However, while the winning strategies in this situation call for cooperation, altruism is not in the discussions known to us.

3.3 Reputation Management

Reputation Management (RM) occurs when an agent evaluates the actions of other agents, forms opinions about those agents, and then uses these opinions to adjust its own actions. The field of RM involves analysis of such processes with applications ranging from interpersonal relationships to economics. A survey of RM with an emphasis on its application to the online marketplace

is presented in [5]. A related example can be found in [13], where RM is applied to the general area of "Electronic Communities". This work demonstrates the positive development of altruistic relationships in which the *trust* of other agents can be built up over time. This has close similarities to the application within an MRC, but uses the trust to assess the quality of information from other agents. Here, we use this trust to determine if robots should be altruistic to one another, thereby improving individual performance.

4 Multi-Robot Communities

A Multi-Robot Community $C = \{r_1, r_2, \dots, r_N\}$ is a set of N robots that can interact through some shared workspace W . In this community, each robot r_i will have a set of L_i individual tasks to accomplish described by the possibly dynamic set $T_i = \{t_{i1}, t_{i2}, \dots, t_{iL_i}\}$. Such tasks may include taking measurements, picking up materials, placing objects, etc.

A task t_{ij} is considered to be completed if any of the robots within C visits the task location. Therefore, once tasks are assigned, each robot r_i plans a sequence S_i of task locations to visit that minimizes path length.

$$S_i = \{t_{kn}, t_{lo}, \dots, t_{mp}\} \quad (1)$$

where indices k, l, m, n, o, p are arbitrary at this point to reflect the possibility that robot r_i 's task sequence S_i may include any of the n^{th} , o^{th} , or p^{th} tasks belonging to any of the k^{th} , l^{th} , or m^{th} robots within the community.

While different task sequencing algorithms may be used, it should be clear that the effectiveness of the task allocation is related to how close the planned sequences are to the optimal task sequence S_i^* .

The cost incurred to accomplish a task sequence is calculated based on the 2 dimensional euclidean distance $d_{kn,lo}$ between two task locations t_{kn} and t_{lo} . Hence, a candidate for each robot's personal cost function is

$$J_i = J_i(S_i) = w_i \sum_{t_{kn}, t_{lo} \in S_i} d_{kn,lo} \quad (2)$$

Where w_i has units $s \cdot m^{-1}$ and for this paper has value 1. Note that, instead of eq. (2), several other cost functions could be used within the altruistic MRC framework that follows. The global cost function typically used to characterize performance for task allocation in an MRS would be

$$J_{global} = \frac{1}{N} \sum_i^N J_i \quad (3)$$

This cost function can be used to measure performance of an MRC. However, individual robots participating within an MRC will most likely use only their individual cost function described in eq. (2). Hence, controllers should be designed with such cost functions in consideration.

4.1 Task Fulfillment in MRC

A common approach taken for optimizing task allocation is to implement a market-based auction system in which robots place bids for tasks (e.g. [6]). Through inter-robot communication, each robot r_i can communicate its bid on each task as it is auctioned. The task will be awarded to the robot with the lowest bid for that task.

A robot's bid on a particular task will be calculated based on the ability to complete the task in minimal time. For example, if τ_{ij} is the time for robot r_i to accomplish task j and all other tasks awarded in previous bidding wars, then the bid b_{ij} can be set to equal τ_{ij} .

In this work, each robot r_i can auction any of its own tasks from T_i . Other robots can choose to bid on robot r_i 's tasks, thus acting in an altruistic manner if they win the bid and complete the task. This choice is based on the level of altruism one robot may have towards another robot.

To be precise, we define the variable $\alpha_{ij} \in [0, \infty]$ as the level of altruism robot r_i has towards robot r_j . A robot r_i will bid on other robot r_j 's task t_{jk} according to:

$$b_{i,jk} = \begin{cases} J_i(S_i \cup t_{jk}) & \text{if } \alpha_{ij} > J_i(S_i \cup t_{jk}) - J_i(S_i) \\ \infty & \text{else} \end{cases} \quad (4)$$

The goal is to control this value of α such that robots will behave altruistically towards one another. This will lead to robots doing each other's tasks when more efficient, thereby decreasing cost functions J_i .

However, since robots are trying to minimize their own cost, they have incentive to act in a selfish manner by maximizing the number of their tasks completed by other robots and minimizing the number of tasks they themselves complete. Hence it is not always beneficial to simply set α_{ij} to some constant that ensures altruistic behavior (e.g. setting $\alpha_{ij} = \infty$). Instead, robots can adapt α_{ij} on-the-fly in response to the actions of other robots, or more specifically in response to the complimentary altruism, α_{ji} .

When such dynamic altruistic behavior is allowed, analysis of the altruistic dynamics within an MRC is required. To do so, a standard discretized linear time-invariant state space model is proposed:

$$\alpha_{t+1} = A\alpha_t + B\mathbf{u}_t \quad (5)$$

In this equation, $\alpha = [\alpha_{12} \ \alpha_{21} \ \alpha_{13} \ \alpha_{31} \ \dots \ \alpha_{ij}]^T$ is the state vector, t is the time step, $A = I^{N \times N}$ is the state matrix, $B = 1$ is the input matrix, and $\mathbf{u} = [u_{12} \ u_{21} \ u_{13} \ u_{31} \ \dots \ u_{ij}]^T$ is the control input. While this implies a simple model, different values for A and B are possible and may require more complicated control laws.

5 Altruism Controller Design

This section provides a controller for setting $\mathbf{u}$ in equation (5), that allows for mutually altruistic relationships to form. Before proceeding to the design of control laws for α_{ij} , it should be noted that the system assumes: 1) all robots can reliably communicate with one another to conduct auctions and bidding, and 2) all robots can either directly measure the altruistic nature α_{ji} of other robots. This second assumption implies that individual cost functions have the same units and can be compared.

In the control strategy proposed, each robot r_i will try to set its altruistic nature α_{ij} towards another robot to be that which the other robot r_j has toward it. This is accomplished through the following proportional control law:

$$u_{ij} = K_{ij}(\alpha_{ji} + \epsilon_{ij} - \alpha_{ij}) \quad (6)$$

The control gain $K_i > 0$ determines the rate at which α_{ij} approaches the desired value of $\alpha_{ji} + \epsilon_{ij}$. The first term in this desired value is the altruism that robot r_j has towards robot r_i . The second term, $\epsilon_{ij} > 0$, indicates how much more altruistically robot r_i will act towards robot r_j .

It is important to note that ϵ_{ij} is used to allow altruism to grow between two robots in that if both robots use such a control law, their complimentary values α_{ij} and α_{ji} will grow throughout time. Consider the resulting state model of the altruistic relationship between 2 robots that can occur, regardless of the other robots in C . Note that without losing generality, it is assumed $K_{12} = K_{21} = K$ and $\epsilon_{12} = \epsilon_{21} = \epsilon$.

$$\begin{pmatrix} \alpha_{12} \\ \alpha_{21} \\ \epsilon \end{pmatrix}_{t+1} = \begin{pmatrix} 1-K & K & K \\ K & 1-K & K \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha_{12} \\ \alpha_{21} \\ \epsilon \end{pmatrix}_t \quad (7)$$

The stability of the system can be evaluated through a coordinate transformation $e_1 = \alpha_{12} - \alpha_{21}$. Given this transformation, the system can be restated as error dynamics:

$$e_{1,t+1} = (1 - 2K)e_{1,t} \quad (8)$$

Hence if $|1 - 2K| < 1$ the error dynamics will be stable and it follows that the error $(\alpha_{12} - \alpha_{21}) \rightarrow 0$ as $t \rightarrow \infty$.

Now, consider a desired rate of change of altruism to be $K\epsilon$, then the error in the rate of change of altruism is:

$$e_{2,t+2} = (\alpha_{12,t+1} - \alpha_{12,t}) - K\epsilon \quad (9)$$

Substituting the top row of eq. (7) into eq. (9) yields:

$$e_{2,t+2} = -Ke_{1,t} \quad (10)$$

Fig. 1. Task fulfillment in an MRC when for all i, j , the altruistic nature in (a) of robots $\alpha_{ij} = 0\%$ and individual robot paths must span large portions of the workspace to visit individual task points. In (b), $\alpha_{ij} = 5\%$ and when $\alpha_{ij} = 100\%$ (c). With altruism, robots tend to create more localized paths that demonstrate increased efficiency. Altruism are percentages of greatest distance a task can add to a robots path in the given workspace.

Again if $|1 - 2K| < 1$, $e_1 \rightarrow 0$ as $t \rightarrow +\infty$. It then follows that $e_2 \rightarrow 0$ as $t \rightarrow +\infty$. More explicitly, the rate of change of altruism ($\alpha_{12_{t+1}} - \alpha_{12_t}$) stabilizes to $K\epsilon$.

Thus, for gain conditions $0 < K < 1$, the mutual altruisms α_{12} and α_{21} will both match each other and grow over time.

If, on the other hand, robot r_i attempts to be altruistic towards a selfish robot r_j , (i.e. $\alpha_{ji} = 0$), then the state transition from eq. (7) reduces to:

$$\begin{pmatrix} \alpha_{12} \\ \epsilon \end{pmatrix}_{t+1} = \begin{pmatrix} 1 - K & K \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha_{12} \\ \epsilon \end{pmatrix}_t \quad (11)$$

In this case we can define the error to minimize with the transformation $e_3 = \alpha_{12} - \epsilon$. Substituting top row of eq. (11) into this transformation yields error dynamics:

$$e_{3_{t+1}} = (1 - K)e_{3_t} \quad (12)$$

Hence if $0 < K < 1$, then $e_3 \rightarrow 0$ as $t \rightarrow +\infty$, and it follows that the corresponding state α_{12} is stable. In fact, this behaves like a Proportional control system where ϵ is the desired state.

To note, the α_{ij} controller presented above is dependent on the reciprocal altruism α_{ji} . If robot r_j can not be trusted, then its altruism must be estimated based directly on its actions and any task point t_{ik} on which it bids. In this work, a conservative estimate is used:

$$\hat{\alpha}_{ij,t+1} = \begin{cases} \max(\hat{\alpha}_{ij,t}, J_j(S_j \cup t_{ik}) - J_j(S_j)) & \text{if robot } j \text{ bids} \\ \min(\hat{\alpha}_{ij,t}, J_j(S_j \cup t_{ik}) - J_j(S_j)) & \text{else} \end{cases} \quad (13)$$

6 Results

To demonstrate the general effect on an MRC when robots create altruistic relationships, a simulated task fulfillment experiment was conducted in which 500 tasks were randomly created within a 6.4x6.4 m 2D workspace as shown in Fig. (1). Each task was randomly assigned to be within the task set of one of 8 robots operating in the space. After each of the 500 tasks are assigned, robots auction them off to all robots using an assumed wireless communication system. Robots will bid if they have sufficiently large altruism towards the auctioneer. The order of auctioning is random.

Fig. 2. Potential effects of increased altruism within an MRC

For each run of the experiment, robots had fixed but equal values of altruism. Between experiments, the altruism was incremented by 1% up to 20%, then incremented by 10%. In Fig. (2), the average robot path cost in meters is plotted for these values of altruism. Note that the values for altruism are normalized with respect to the percentage of the maximum value defined as the greatest distance between two points in the workspace (i.e. the length of the diagonal connecting opposite corners). The path costs show a dramatic decrease when compared with the case with no altruism. The cost plateaus where additional tasks typically won't cost a robot more than 20% of the largest distance a task point can occur on path cost for the given workspace. This can be observed in Fig. (1) as well, where in (a) no altruism results in paths that cross one another and the entire environment. However, in (c) where 100% altruism is used, paths are more localized and the personal cost functions decrease to the effect that the global cost function is also minimized.

To demonstrate the effects of the control law, consider the case where two robots start with random values for altruism with respect to each other (e.g. 0%, 18%). As shown in Fig. 3(a), the altruism between the two robots first converges and then grows with time (i.e. as new tasks are auctioned). It can be noted that the rising slope calculated from Fig. 3(a) is 1/16 which matches the expected value of $K\epsilon = (0.25)(0.25)$.

When robots are acting selfishly, the controller invoked by the altruistic robot is stable, (see Fig.3(b)). The altruism towards such selfish robots reaches equilibrium at the point ϵ greater than 0. Hence, the altruistic robot is still protected against selfish robots. Note that increasing the gain K increases the speed of convergence. In (c) and (d), cases where the altruism of other robots are estimated using equation (13) are presented. In (c), robot 1 has altruism fixed at 100%. Robot 2's altruism towards 1 starts at 100% and converges to the maximum bid used by robot 1. In (d), robot 1 has altruism fixed at 0, it can be seen that robot 2's altruism converges towards ϵ .

7 Conclusions and Future Work

This paper presented the idea of a Multi-Robot Community in which several robots sharing a common workspace are attempting to complete their individual tasks. It was shown that altruistic actions, where robots assist with each other's tasks, can lead to decreased path costs for individual robots.

While a controller was presented that can set altruism such that altruistic relationships can evolve between two altruistic robots and still protect against selfish robots, the real contribution comes from the idea of analyzing the stability of altruism as a linear time-invariant system. Future controllers can be designed and analyzed in a similar fashion.

Considerable work is still required, especially to resolve the assumptions listed above. In the future it is hoped that better estimation of another robot's altruism can be achieved, selfish robots that have variable values for alpha are addressed (e.g. gaming robots), situations with uneven task distribution are considered (5 tasks for one robot 100 for another), and practical implementation and experimentation may occur.

References

1. Agah A, Bekey GA (1997) Tropism-based cognition: a novel software architecture for agents in colonies. *Journal of Experimental and Theoretical Artificial Intelligence* 9: 393-404
2. Arkin RC, Bekey GA(1997) *Robot Colonies*. Springer Publishing.
3. Arkin RC, Balch T (1998). Cooperative multiagent robotic systems, in *Artificial Intelligence and Mobile Robots*. D. Kortenkamp, R. P. Bonasso and R. Murphy, AAAI Press/MIT Press: 277-295
4. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* RA-2: 14-23
5. Dellarocas C (2006) Reputation Mechanisms. *Handbook on Information Systems and Economics*. Hendershott T (ed). Elsevier Publishing. 154-165
6. Goldberg D, Cicirello V, Dias MB, Simmons R, Smith S, Stentz A (2003) Task Allocation Using a Distributed Market-Based Planning Mechanism. *International Conference on Intelligent Agents*

Fig. 3. Stability of altruism within an MRC. In (a), both robots use the altruistic controller but with different initial values (0% and 18%). In (b), only one robot uses the controller while the other robot maintains a fixed value of 0. In (c) and (d), the altruism must be estimated by robots.

7. Lucidarme P, Simonin O, Liegeois A (2002) Implementation and evaluation of a satisfaction/altruism based architecture for multi-robot systems. Proc. IEEE Int. Conf. on Robotics and Automation 1007–1012
8. Mataric M (1997) Behavior based control: Examples from navigation, learning and group behavior. Jour. of Experimental and Theoretical Artificial Intelligence 9(2-3): 67–83
9. Nowak MA, May RM, Sigmund K (1995) The arithmetics of mutual help. Scientific American 272(6):76–81
10. Parker LE (1998) ALLIANCE: An architecture for fault tolerant multirobot cooperation. IEEE Trans. on Robotics and Automation 14(No. 2): 220–240
11. Simonin O, Ferber J (2000) Modeling self-satisfaction and altruism to handle action and reactive cooperation. SAB'00 Proceedings: 314–323
12. Sofge D, Schultz A, De Jong K (2002) Evolutionary Computational Approaches to Solving the Multiple Traveling Salesman Problem Using a Neighborhood Attractor Schema. Applications of Evolutionary Computing - EvoWorkshops
13. Yu B, Singh MP (2000) A Social Mechanism of Reputation Management in Electronic Communities. Cooperative Information Agents (CIA) 2000. Springer-Verlag, Berlin Heidelberg 154–165.