Cost, Precision, and Task Structure in Aggression-based Arbitration for Minimalist Robot Cooperation

Tanushree Mitra and Dylan A. Shell

Department of Computer Science and Engineering
Texas A&M University

Abstract. This paper reexamines a multi-robot transportation task, introduced and studied by Vaughan and his collaborators, in which constrained space induces inter-agent interference. Previous research demonstrated the effectiveness of an arbitration mechanism inspired by biological signaling where the level of aggression displayed by each agent effectively prioritizes the limited resources. This paper shows that apart from determining the correct fitness of an individual several other factors, such as signaling cost, precision of the outcome and properties of the resource and task are key to determine an effective arbitration technique. Based on these factors we present a taxonomy of the arbitration mechanisms. The large signalling costs incurred by our simple robots using minimal set of sensors permit us to identify scenarios in which a dominance hierarchy outperforms, not only to no arbitration, but also aggression-based mechanisms. We identify how memory of past interactions can be used to the advantage of an agent, albeit with a trade-off between cost and outcome accuracy. We also show that the importance of a particular aggressive interaction to long-term task performance is not trivial to determine and depends on the task structure. Results help us identify instances where agents may manipulate interactions to alter the frequency and duration of aggressive encounters, affecting the overall task performance.

1 Introduction

Spatial interference is a common and important phenomenon in navigation tasks involving multiple robots; it is a particular instance of the general problem of resource competition amongst agents attempting to achieve their own ends while interacting with others. When autonomous agents have an incentive to cooperate, a worthwhile question is how best to mitigate the negative effects of resource contention. Motivated by the methods which animals employ to contest resources, Vaughan and his collaborators (*cf.* [5], [2], and [6]) have shown how displays of stylized aggression can effectively resolve resource conflicts in a multi-robot transportation task. That line of work has produced increasingly effective methods for assessing the level of aggression that an individual agent should exhibit in order to prioritize the limited resource effectively.

This paper shows that determining the correct fitness of an individual at a particular time is only one of several aspects of effective conflict resolution. An important consideration is the cost of the aggressive signaling. In fact, the analogous biological mechanism is directly concerned with the interplay of signal precision and cost: aggressive displays allow animal to assess the strength of their resource competitors before they decide to engage in a costly fight [1]. Animals, after all, organize themselves into a dominance hierarchy which they can use to resolve future resource competition [4] in an inexpensive albeit static way.

In this paper, we examine the multi-robot transportation task domain that Vaughan and his collaborators have studied. Specifically, we study a two robot interference scenario, the goal of which is to cooperatively perform the maximum number of collective transportation tasks in a given time. We present a taxonomy of arbitration mechanisms for two-agent spatial interference, including a characterization of conflict resolution models, introduce a notion of outcome accuracy and explicitly consider interaction cost. Results from physical robot experiments and data-driven simulations led to following contributions:–

1. We show that there exist similar problem instances in which either dynamic aggression or a static dominance hierarchy are advantageous.
2. We also demonstrate how memory of past interactions with respect to the task structure and *properties of the resource* can result in improved future task performance.
3. The paper shows that varying the *properties of assigned task*, the frequency of spatial interference and the cost incurred in its resolution varies significantly.
4. A new "minimalist" resource arbitration method is introduced which produces dynamic outcomes—albeit with comparatively high costs—suitable for simpler robots (with fewer sensors) than heretofore known.
5. We identify how agents may manipulate interactions to alter the frequency and duration of aggressive encounters, affecting the overall task performance in repetitive tasks.

We begin by giving a brief overview of previous research of interference in multi-robot systems. Next, in Section 3 we propose a taxonomy of arbitration mechanisms for two-agent spatial interference. In Section 4, we provide a comparative study of the various arbitration models. Results of the study are based on physical robot experiments. Finally in Section 5, we designed a custom simulator based on our physical robot data from previous section and anlyzed interference under systemic variation of environment.

2 Related Work

Goldberg and Matarić [3] have suggested using interference as a tool for evaluating multi-robot controllers, *viz.* identifying trade-off between performance time and interference. Vaughan *et al.* [5] compared a dominance hierarchy to the aggression-based strategy in a multi-robot transportation task in a simulated environment. Our data below show that one strategy can be preferable to the other, but exactly which depends on the shared resource and also on the individual task dynamics. Brown *et al.* [2] introduced the concept of rational aggression where the level of aggression is determined by the investment made by Zuluaga and Vaughan [6] further improved on Brown *et al.*'s performance by basing the level of aggression on the investment in the shared resource. Details on the relevant strategies compared in this work are below. Brown *et al.* [2] introduced the concept of rational aggression where the level of aggression is determined by the investment made by Zuluaga and Vaughan [6] further improved on Brown *et al.*'s performance by basing the level of aggression on the investment in the shared resource. Details on the relevant strategies compared in this work are below.

This paper uses the same controlled scenario (depicted in Fig. 1) as this previous line of research. Two robots perform repetitive transportation tasks by moving around cycles in the environment (shown as blue and green walls they must follow). Competition for space occurs at a narrow shared region, the girdle, where only one robot traverse at a time; different arbitration models determine who gives way.

Additionally, in [3] the authors suggested that changing the environment could play a role in altering interference properties. This manifest itself in [2], where changes in the (simulated) environment produced large standard deviations in their results. Part of our work is an attempt to analyze and mitigate interference under systematic variation of the environment.

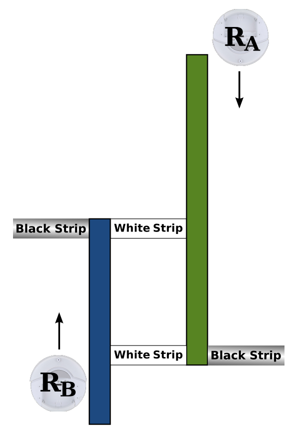

Fig. 1: **The task involves robots navigating their respective regions (around blue and green walls, respectively) but also having to resolve a conflict in the shared space (between the "white strips").**

2.1 Arbitration Models

A resource arbitration model determines which of the two robots should have access to the shared resource. In aggression based models each robot determines a quantity, their level of aggression, and engages in a dynamic behavior involving giving or taking way in the shared space.

Vaughan's random aggression: Each agent picks a random aggression at each encounter, resulting in a random outcome; *i.e.,* the resource is gained by a random agent.

Vaughan's personal space method: The level of aggression is determined by the amount of free space visible behind the robot in the event of interference.

Rational aggression based on local task investment: This is modeled after [6], aggressive interaction based on local task investment, where the robot "displays its aggression" by moving backward a distance inversely proportional to the distance traveled so far in the constrained region, and then moving forward until it bumps again. The robot's controller is shown in Figure 2(a).

Linear dominance hierarchy: A fixed dominance is assigned to each robot before they start their spatial navigation and each one follows this at every encounter to determine who gets right-of-way. Figure 2(b) shows this.

Cutting your Losses: Some memory of local task performance is added to the rational aggression method. When a robot meets an opponent inside the girdle, it displays its level of aggression for at most ϕ attempts. At the same time it measures and remembers the cost it incurs in this display, by measuring the lose or gain in the shared space distance from the time it starts its aggressive display. Figure 2(c) gives this mechanism.

Random walk: As soon as a robot encounters an opponent, it backs a random distance. It then moves forward and if the opponent is still in the girdle, it again moves back by a random distance. The opponent also follows the same protocol. Eventually one of the robots is pushed out of the girdle. Figure 2(d) illustrates this mechanism.

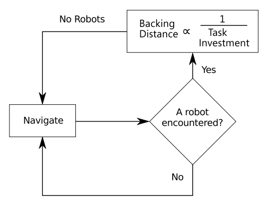

(a) Aggression-based arbitration.

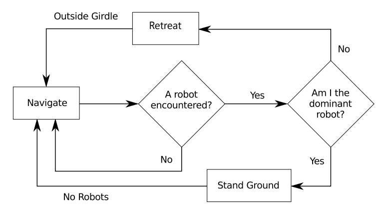

(b) Fixed dominance hierarchy.

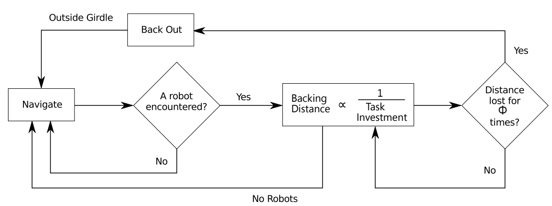

(c) Cutting one's losses.

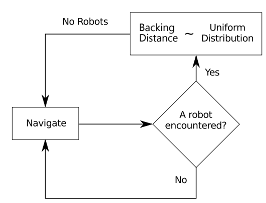

(d) Random walk arbitration.

Fig. 2: **Details for four different arbitration mechanisms.**

3 Taxonomy of spatial conflict resolution models

We propose a taxonomy of conflict resolution models with the following axes:–

- Dynamic vs. static: An arbitration method is static if and only if it does not employ information about a particular encounter to resolve that conflict.
- Deterministic (DET) vs. Probabilistic (PROB): A method is deterministic if and only if, given the same scenario, the resource is always awarded to the same agent.
- High (HOA) vs. Low outcome accuracy (LOA): The former has higher probability of selecting the rational winner. The robot with the greatest local investment should gain the resource in rational interactions.
- Costly (HIGHCOST) vs. cheap (LOWCOST): Time, energy and other resources may be involved in an arbitration mechanism. Their utility depends on the comparative saving and/or trade-off of these costs.

Fig. 3: **An example of a robot giving way. After bumping each other, both robots move back and one exits the girdle to make room.**

Different arbitration models can be denoted by a quadruplet:

Model	Classification			
Linear dominance hierarchy	STATIC	DET	LOA	LOWCOST
Vaughan's random aggression	STATIC	PROB	LOA	LOWCOST
Random walk	DYNAMIC	PROB	LOA	LOWCOST
Vaughan's Personal space	DYNAMIC	DET	HOA	HIGHCOST
Rational aggression	DYNAMIC	PROB	HOA	HIGHCOST
Cutting your losses	DYNAMIC	PROB	HOA	LOWCOST

4 Comparative Study

4.1 Implementation Details

We imposed spatial interference on physical robots by making them navigate through an environment as shown in Figure 1. Interference occurs when two iRobot *Creates*'s R_A and R_B, 33cm in diameter attempt to cross a girdle ∼53cm wide from opposite directions as shown by the arrows in Figure 1. R_B's transportation task length is almost half that of R_A. R_A does 10 traversals, while R_B covers 20. We assigned these numbers so as to avoid situations where the robot performing the shorter task finishes all its trips while the other one keeps traversing an encounter-free region.

4.2 Aggressive Interaction and Linear Dominance

Both these models were executed in environments with different girdle lengths. The aim is to assess the role the shared resource plays on arbitration outcomes.

Varying Girdle Length: The utility of aggressive interaction is reduced when both robots have large, almost equal aggression levels, a phenomenon which happens when encounters are at the center of a large enough girdle. This can be observed in Figure 4, as the aggression strategy performs increasingly poorly with increasing girdle length. A dominance hierarchy, despite it not necessarily resolving the conflict toward the agent with the greater investment, proves to be a better arbitration method in such cases.

However for encounters at the ends of the girdle, the short aggressive interactions coupled with the ability to produce a rational winner makes aggression based arbitration beneficial over dominance. Vaughan *et al.* [5] provide an instance where choice of aggression level proves no better than random selection of aggression (Vaughan's random aggression). In fact, the outcome of such a random mechanism is an average drawn from the outcomes of following either extremes of the linear dominance hierarchy. In the data above the advantage of such a mechanism can be seen.

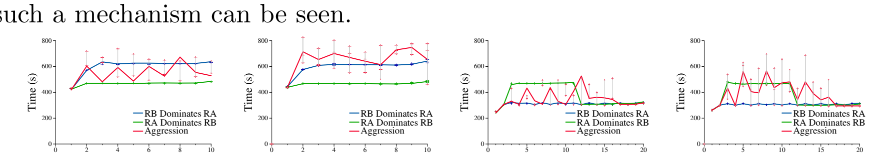

(a) R_A,GL=3. 35 (b) R_A,GL=4.7 (c) R_B,GL=3.35 (d) R_B,GL=4.7

Fig. 4: **Average task times of R_A and R_B with varying girdle lengths (GL) in meters, fixed task ratio 25:38. (Results are from three experiments averaged for each of the three strategies, for each of the three cases. We show 2 cases for space constraints.)**

An important question is "how precisely can the outcome of the arbitration be predicted given the initial position of encounter?" When the robots have approximately equal local task investment, as in the setup above, the noise in the robot's interactions breaks the symmetry. In Figure 5 the mixed red and blue region near the center of the girdle (girdle proportion ∼0.5) shows that there are situations when R_A's local investment is less than R_B, but R_A wins or vice-versa. These are the few instances when the outcome of the aggressive encounter is neither particularly accurate nor rational and there is high cost involved, decreasing their utility in such situations. Moreover, in instances where encounters occur at girdle ends, the inaccuracy of dominance hurts only when the less dominant robot has higher local investment and still retreats. If the task ratios are appropriate, then such scenarios may occur only rarely, making dominance arbitration the superior arbitration model for such environments.

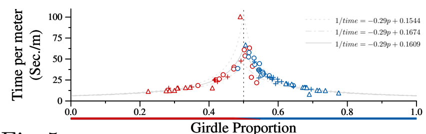

Fig. 5: **Time to resolve an aggressive interaction of two physically grounded robots. The horizontal bar shows the robot which gets right of way when the point of encounter inside the girdle (length normalized) varies, red being robot $\mathbf{R}_A$ and blue $\mathbf{R}_B$**

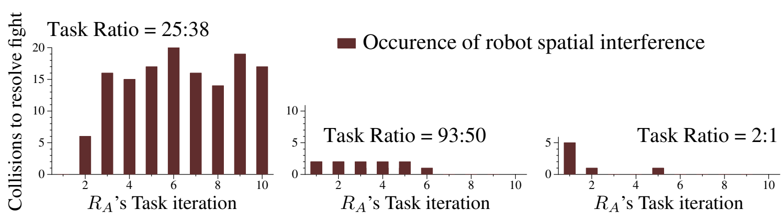

Fig. 6: **The duration of aggressive interaction of two physically grounded robots changes as we vary task ratio, suggesting the importance of task structure in spatial interference. These experiments were performed on physical robots. We reexamine effects of task structure in spatial interference through exhaustive simulation in the next section.**

Varying Task Ratio: We examined how the properties of the task assigned to each agent influence aggressive encounters. This factor dictates the time when a robot starts its journey inside the girdle relative to the other and, thus, the initial position where they end up meeting. There is also the chance that they do not meet at all. The variation of the duration of aggressive interactions as shown in Figure 6 indicates the importance of task structure in aggression based arbitration. Task structure is reexamined through exhaustive simulation below.

Cutting your Losses The utility of aggressive encounters can be improved by adding memory of recent performance. The robot measures and remembers the loss or gain in the shared space distance from the time it starts its aggressive display. If it repetitively loses distance, then it is unlikely to win the whole interaction. In such a situation it is beneficial to retreat. The greater the number of confirmations about the gain/loss in distance, the more accurate its decision. Figure 7 shows this decrease in error with an increase in the number of confirmations. The tradeoff here is whether to take an early

Random Walk The attraction of random walk arbitration is its minimalism compared to other arbitration methods: robots do not need to sense or estimate

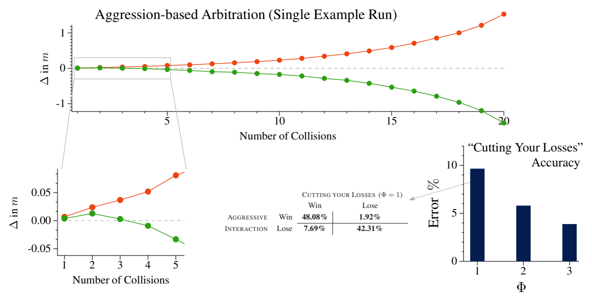

Fig. 7: **The intuition behind the "cutting your losses" strategy is illustrated via an example of the aggression-based interaction. The sign of the single-time gain (denoted Δ in the graphs) indicates a likely win or loss. Waiting longer before measuring $sgn(\Delta)$ reduces the estimate error due to the "escalation" dynamics. The results are from physical robot experiments.**

their positions, since the position the bump takes place implicitly encodes the dynamic variable. It is perhaps somewhat surprising that a dynamic arbitration mechanism is possible despite neither of the robots having the means to record their investment in the task.

5 Task Ratio—A Macroscopic Study

Designed a custom simulator that uses physical robot data to model dominance and aggressive robot interactions for different girdle lengths across a range of task ratios. We ran the simulator, varying girdle lengths (GL) from 10m to 150m and, for each girdle length, task lengths varied from 15m to 150m. The results represent complementary foci: either minimizing absolute signal cost via a static arbitration, or incurring whatever cost to ensure a dynamic arbitration.

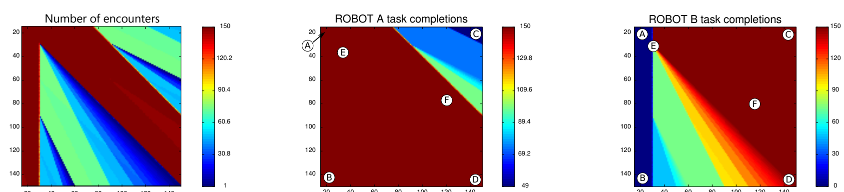

(a) Number of encounters (b) Laps completed by R_A (c) Laps completed by R_B

Fig. 8: **Dominance for girdle length = 30m, R_A is dominator. The x-axis shows R_A's task length, y-axis that of R_B. (a) Color bars shows the relative number of encounters,(b), (c) Color bar shows the relative number of laps finished when at least one of the robots completes 150 laps. Points A to F are detailed in Figure 9.**

Dominance Model Figure 8 shows the performance in a girdle length of 30m when R_A is the dominator. Interesting regions from these plots were selected to investigate the interaction dynamics for the first 20,000 seconds (long enough to show long-term behavior). These are marked with A—F in Figure 9, and described in detail in that caption.

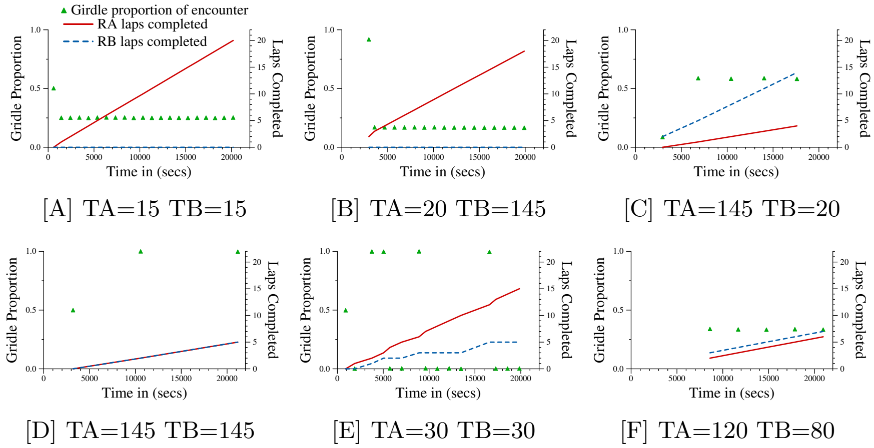

Fig. 9: **Dominance for girdle length = 30m, R_A is dominator. The girdle proportions are with respect to the position of robot R_A. TA corresponds to the task length of R_A, and TB to that of R_B. Results presented are based on simulation experiments. We observe the following at each point:**

A & B— Every time R_B makes an attempt to cross the girdle, it meets the dominator R_A and is pushed out of the girdle, making no progress whereas R_A finishes more than 15 laps during the alloted time. This is clearly a model of resource starvation.

E— In this case, R_A and R_B meet frequently and with R_A being the dominator, R_B is able to complete fewer trips than R_A with such frequent spatial interference.

C— The encounter position is close to R_B's girdle entry point, so even if R_B retreats (R_A being the dominator) the local investment made by R_B is less. R_A is the rational winner with aggression but at the cost of aggressive interaction time. Also they do not encounter one another every time R_B enters the girdle. But R_B's shorter task allows it to complete more trips than R_A.

F & D— The number of task iterations completed by both robots are almost equal. D belongs to the region where the number of encounters are less frequent (Figure 8(a)) and, if they occur at all, they are at the girdle end when R_A is about to exit (see Figure 9[D]). In such a situation R_A is the rational winner and the dominance hierarchy (with dominator R_A) is the best model to follow.

Aggression Model The model was run to compare with the dominance method.

Collective best performance across varying girdle length: Certain combinations of task lengths takes significantly longer to complete 100 tasks (Figure omitted due to space restrictions). One might think that this is due to severe interference for these task ratios. However, plotting the interference count we find that this is not always true. The following considers one fixed girdle length.

Collective vs. individual best performance for fixed girdle length: There are regions of high interference corresponding to regions of low task completion time (Figure omitted due to space restrictions). These are the instances where the robots met often but engaged in less costly aggressive interactions. On the other hand, there also exist regions of low interference but with high task times. These regions involve high cost aggressive interactions.

We further investigate as to what happens for certain combinations of task lengths, similar to what we did for dominance model. Positions of first encounter for every task iteration completed in the first 20,000s are shown (Figure 10).

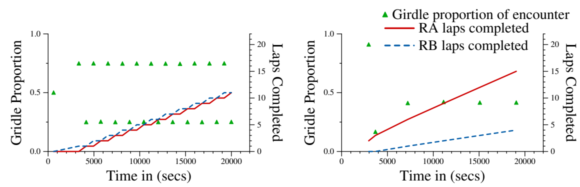

[A] TA=15 TB=15 [B] TA=20 TB=145 [C] TA=145 TB=20

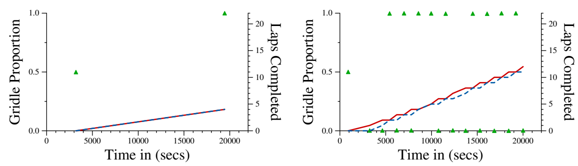

[D] TA=145 TB=145 [E] TA=30 TB=30 [F] TA=120 TB=80

Fig. 10: **Aggression GL30. The girdle proportion are w.r.t. the position of robot R_A.**

A— R_A and R_B meet frequently. The position of encounter inside the girdle results in costly interactionsi.The number of task iterations completed by R_A and R_B is almost equal for the first 20,000 seconds of simulation and every time a rational winner is chosen. Compare this with how R_B performs when R_A is the dominator (Fig. 9[A] and 9[B]). R_B did not make any progress during the time allotted. Here the trade-off is to, either to engage in costly aggressive interactions obtaining a rational winner, giving a fair chance of winning to each robot or to resort to (cheap) dominance and bias towards one agent.

B— R_A and R_B do not meet as frequently here, but whenever they do, long aggressive interactions result: R_B is the rational winner in all cases. Compare this with Figure 9[B] where R_B hardly made any progress. The best resolution mechanism would be with R_B as dominator so that no aggression cost need be paid.

C— This is the reciprocal of case B and our earlier conclusion (but for R_A) holds true.

D— The number of task iterations completed by R_A and R_B are almost equal. We notice that D belongs to the region where the number of encounters are less frequent and they occur at the ends of the girdle when R_A is about to exit (Figure 9[D]). We had earlier concluded that, in such a situation R_A is the rational winner and the dominance hierarchy (with R_A being the dominator) is the best interference model to follow. We see that aggressive arbitration is also a reasonable interference resolution mechanism. The reason being that these regions have cheap interactions.

E— Compare Figure 10[E] with Figure 9[E]. The number of laps which R_B completes with aggressive signaling doubles compared with when it is dominated. However, the decrease in the number of laps of R_A is not that significant in both these modes of arbitration. We do see more frequent encounters in case of aggressive interactions, but all of these take place at the very ends of the girdle resulting in cheap arbitration, making it beneficial.

6 Conclusion

Several factors contribute to conflict resolution and its effectiveness:

Cost vs. precision of the arbitration mechanism—Time and energy cost are incurred in resolving resource conflicts. This influences the utility of aggressive displays in the first place.

Properties of the shared resource for which the agents are competing—This affects, among other things, the cost of communicating its aggression and what constitutes a worthwhile investment.

The task which each agent is assigned to perform—This can be coupled through the shared resource. This coupling, effects individual and collective dynamics.

The inherent noise in the "communication" channel—Noise plays a role in dynamic arbitration mechanisms: it can be beneficial in breaking symmetry, a situation which occurs when agents have identical aggression.

From all these facts we can conclude that there cannot be just one single best arbitration mechanism catering to all situations. We have also shown instances where a small variation of task ratio may cause a significant change in the task dynamics. With a prior knowledge of this entire task performance space, a single unfavorable interaction can be predicted beforehand, and by adding a wait to its task navigation, the robot can shift its performance to a more favorable region.

References

1. Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behav.*, 33(4):1152, 1985.
2. Sarah Brown, Mauricio Zuluaga, Yinan Zhang, and Richard T. Vaughan. Rational aggressive behaviour reduced interferrence in a mobile robot team. In *Proc. International Conference on Advanced Robotics*, Seattle, WA, U.S.A., July 2005.
3. Dani Goldberg and Maja J. Matarić. Interference as a Tool for Designing and Evaluating Multi-Robot Controllers. In *Proc. AAAI National Conference on Artificial Intelligence*, pages 637–642, Providence, RI, U.S.A., July 1997.
4. Devi Stuart-Fox. Testing game theory models: fighting ability and decision rules in chameleon contests. *Proc. Ro.l Soc. B: Biological Sciences*, 273(1593):1555, 2006.
5. Richard T. Vaughan, Kasper Støy, Gaurav S. Sukhatme, and Maja J. Matarić. Go ahead, make my day: Robot conflict resolution by aggressive competition. In *From Animals to Animats: Proc. Simulation of Adaptive Behavior*, pages 491–500, Paris, France, August 2000.
6. Mauricio Zuluaga and Richard T. Vaughan. Reducing spatial interference in robot teams by local-investment aggression. In *Proc. IEEE International Conference on Intelligent Robots and Systems*, Edmonton, Alberta, August 2005.