Adaptive Robot Coordination using Interference Metrics*

Avi Rosenfeld, Gal A Kaminka, and Sarit Kraus
Bar Ilan University
Department of Computer Science
Ramat Gan, Israel
{rosenfa, galk, sarit}@cs.biu.ac.il

Abstract

One key issue facing robotic teams is effective coordination mechanisms. Many robotic groups operate within domains where restrictions such as limiting areas of operation are liable to cause spatial conflicts between robots. Our previous work proposed a measure of coordination, interference, *that measured the total time robots dealt with resolving such conflicts. We found that a robotic group's productivity was negatively correlated with interference: Effective coordination techniques minimized interference and thus achieved higher productivity. This paper uses this result to create adaptive coordination techniques that are able to dynamically adjust the efforts spent on coordination to match the number of perceived coordination conflicts in a group. Our robots independently calculate a projected level of interference they will encounter. By using this metric as a guide, we are able to create adaptive coordination methods that can quickly and effectively adjust to a given domain's spatial limitations. We present two adaptation heuristics that are completely distributed and require no communication between robots. Using thousands of simulated trials, we found that groups using these approaches achieved a statistically significant improvement in productivity over non-adaptive coordination methods.*

1. Introduction

Groups of robots are likely to accomplish certain tasks more quickly and robustly than single robots [3, 5]. However, the physical environment where such teams operate often pose a challenge for the robots to perform properly. For example, domains such as robotic search and rescue, vacuuming, and waste cleanup are all characterized by limited operating spaces where the robots are likely to collide. Improved coordination methods in such domains result in more productive groups.

Our previous work [9] defined a measure called *interference* to facilitate comparison between various coordination methods. Interference is defined as the total time each robot spends in resolving conflicts with other robots. This not only includes the time robots collide, but also the time robots spend preventing such collisions and the time they engage in resolution behaviors after such an event. It was found that a strong negative correlation exists between interference in a group and its productivity. However, this does not mean that robots should avoid the coordination activities which constitute interference as such behaviors are often critical for achieving cohesive team behavior. Rather, it was suggested that the coordination method of choice needs to appropriately match the needs of the domain. As such, interference should be kept to a minimum, while still sufficiently high to meet the coordination requirements of the environment.

This paper builds on this hypothesis by presenting two applications whereby robots dynamically adapt their coordination techniques based on the amount of interference they project will be encountered. Our first method works by tweaking the strength within one coordination method to adapt to its environment. Our second approach proceeds to dynamically self select between a range of mutually exclusive coordination methods. In order to quickly adapt to a changing environment, we use weight-based heuristics by which every robot in the group is capable of quickly modifying its resolution methods to match its estimates of resource conflicts. Our approach is completely distributed, and requires no communications between robots. We found that our adaptive methods result in statistically significant higher average productivity than those of non-adaptive methods regardless of the group size.

The remainder of this paper is organized as follows. The next section discusses the correlation between interference and any group's productivity. We present the problem of matching the best coordination method to a given domain.

* This material is based upon work supported in part by the NSF under grant #0222914 and by the BSF under grant #2002401. Sarit Kraus is also affiliated with UMIACS.

Section 3 presents our adaptive coordination algorithm, and presents our hypothesis that such an approach can effective adapt to the dynamic nature of many robotic domains. Such a method will be able to overcome the shortcoming in static methods. In section 4 we present and evaluate our experiments with dynamic groups to confirm this hypothesis. We discuss related work in section 5. Section 6 concludes and describes possible future directions.

2. Interference versus Productivity

A strong inverse relationship seems to exist between a robotic group's productivity and the length of time these robots engage in coordination behaviors. We previously found [9] a strong negative correlation between the total amount of time robots spend in resolution behaviors, a concept referred to as *interference*, and the productivity of the group. While adding robots may speed up the time to complete certain tasks, and can even be necessary for completing other tasks, these robots can trigger collisions which detract from the group's performance.

Our previous work [9], contrasted various coordination algorithms within the foraging domain. The foraging domain has been extensively studied, and is formally defined as locating target items from a search region S, and delivering them to a goal region G [4]. Various coordination methods have been developed that could work within this domain [10, 3, 14, 8]. We compared algorithms including the concepts of *Aggression* [14], a *dynamic Bucket Brigade* [8], and the use of a repulsion schema mechanism (*Noise* group) [1]. Among others, we compared three additional groups called *Gothru*, *Repel Fix* and *Timeout*. Gothru represents idealized group behavior without any possibility for interference and can only exist in simulation. These robots were never affected by obstacles, and were allowed to simply pass through teammates. Repel Fix resolved collisions by moving away from a teammate for a fixed period of time, here set to 500 cycles. The Timeout method only reacted once a robot detected it had not sufficiently moved for 100 cycles. After this point, it attempted to become unstuck by moving in a random walk for 150 cycles.

Figure 1 graphically presents the results that motivate this work. The X-axis represents the number of robots in the group, and the Y-axis corresponds to the number of foraging pucks that the group brought to the goal region. Notice how Gothru is the only group to achieve positive gains in productivity over all group sizes. The levels of interference that existed in all other groups eventually caused the group's productivity to decrease with the addition of robots. We found a very high negative correlation between the total time groups spent reasoning about and reacting to collisions, and the corresponding productivity. However, no one group was successful in minimizing this level of interference across all group sizes. Our conclusion was that static groups are often not suited for minimizing interference over all conditions. This paper describes how to create adaptive coordination methods that are able to react to the dynamic conditions within robotic domains and thus achieve better productivity across all group sizes.

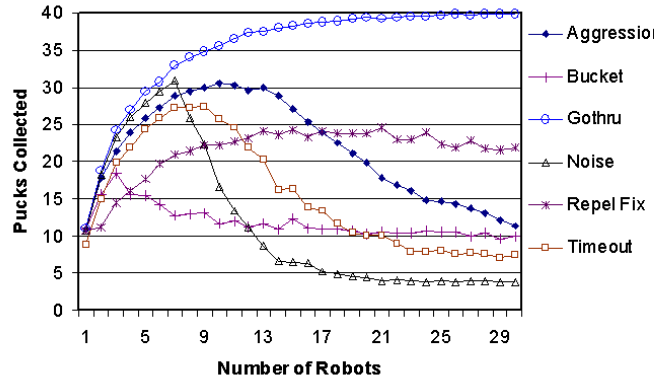

Figure 1. Comparing Coordination Methods

3. Adaptive Coordination

We focus on adaptive methods which use weight-based heuristics to dynamically modify team coordination algorithms to match perceived environmental conditions. By observing the triggers for episodes of interference, we believe it is possible to create coordination methods that move between simple and complex techniques as needed. We present two variations of this approach and their advantages over static methods. In the first technique we have the robots self adjust within one coordination method to match the perceived environmental conditions. Our second technique involves adaptation between distinct coordination algorithms.

In order to demonstrate the shortcomings within static methods, we studied 5 variations of the Repel group. We chose values of 10, 50, 100, 200, and 500 cycles for the length of time these robots would repel after a projected collision. As was the case in our previous work, we used the robot simulator, Teambots [2], to collect data for these groups. We left other details of our setup identical to the implementation previously used. As such, Teambots [2] simulated the activity of groups of Nomad N150 robots in a foraging area that measured approximately 5 by 5 meters. We used a total of 40 target pucks, 20 of which where stationary within the search area, and 20 moved randomly. For each group, we measured how many pucks were delivered to the goal region by groups of 1 – 30 robots within 9 minutes. For statistical significance, we averaged the results of 50 trials with the robots being placed at random initial positions for each run. Thus, this experiment consisted of a total of 15,000 trials of 9 minutes of simulated robotic activity.

The best variation of the Repel coordination method depended on the size of the group. As the group size grew, robots required increasingly more aggressive coordination methods to fight collisions. Among these groups, Repel50 had the highest productivity in groups up to 10 robots. Between 10 and 15 robots the Repel100 group did best. The Repel200 group fared better over the next 5 robots, and the Repel500 group had the highest productivity between 20 – 30 robots. The goal of our first adaptive approach is to create an algorithm that can select the best repel amount for the environment given the group size.

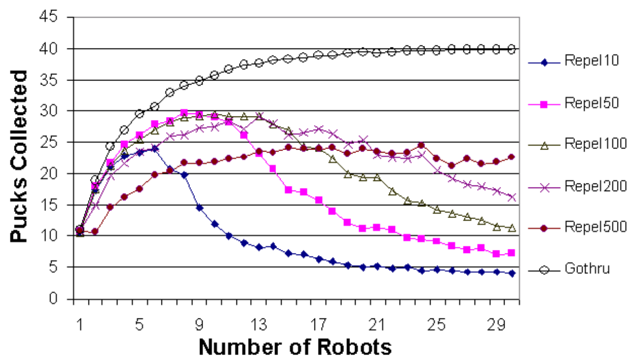

Figure 2. Static Repel Group Productivity

Adaptation could also help to switch between distinct coordination methods. As figure 1 demonstrates, there is no one best coordination method for all sized groups within our domain. Our Noise group fared best in groups up until 7 robots. The Aggression method fared best in groups 8 – 17 robots. Repel500 had the highest average group productivity after this group size. Our second adaptive approach aims to select between various coordination approaches based on the needs of the group.

Both of our adaption techniques are based on heuristics that are sensitive to the triggers of interference the robots face within their domain. As the robot senses the probability of collisions is high, it uses more robust interference resolution mechanisms. If the threat of collisions is low, more simple coordination methods are used. Specifically, our algorithm works as follows: We first initialize a base value, V_{init}, representing the supposed interference level the robot will encounter. For each cycle that passes where a robot detects no impending collisions, it decreases its value of V by a certain amount, W_{down}. For each cycle where the robots detects it is approaching another teammate or obstacle, it increases its value V by a certain amount, W_{up}. Thus, the robots' value V which is constantly in flux and aims to model itself based on the projected amount of interference that particular robot independently encounters.

The coordination method used is based on the distinct values of V within each robot. In our first adaptation method, we translate the various values for V into the number of cycles the *Repel* method uses to repel once it detects a collision is imminent. In our second adaptation method the values for V are used to switch between a set of coordination techniques ranging from those with little computational overhead, but are not likely to scale, to more robust methods with higher overheads. In both methods, when collisions are unlikely, the value V is low and thus coordination methods with little overhead will be used. Thus, the interference will be low, and the level of the group's productivity will be high. When the robot contains a high value of V, it will select more aggressive coordination behaviors to more effectively resolve the projected collision. These computationally high behaviors are needed to prevent the robot from re-triggering its resolution behavior for the same event. While these repulsion activities themselves constitute interference, these behaviors are a necessary evil as more simple behaviors would not suffice.

4. Creating and Evaluating Adaptive Coordination Methods

We began by experimenting with various values for these weights within our first approach of adapting within one coordination method. We found many nearly optimal combinations for the values of V_{init}, W_{up}, and W_{down}. Our adaptive approach was flexible in that a value of V_{init} being originally set too high was soon corrected by the weights in W_{down}. Ultimately, we found that a value of V_{init} = 350 seemed to work best. We used values for W_{down} ranging from 0 to 200 based on how quickly the repel mechanism was previously triggered. Our values for W_{up} ranged from 0 to 550 based on how quickly the robot neared its next collision. This led to a heuristic that took a graduated approach — it would adjust the amount it would repel in the case of a collision fairly quickly up or down based on how frequently collisions occurred within the domain.

We found that our adaptive Repel group produced statistically significant higher levels of performance than that of even the best static method we studied. For statistical significance we ran our adaptive group for 50 trials over a range of 1 – 30 robots. The adaptive Repel team on average collected 1.5 pucks more than the best of the static groups. Figure 3 graphically depicts the success of this group. We conducted the two-tailed paired t-test on our data to confirm the statistical significance of our findings. We compared the averaged productivity values of our adaptive Repel group to all of the non-adaptive methods over the range of 30 robots. All scores were far below the needed 0.05 for significance with the highest p-value for the Null hypothesis being only 0.00013 (between our dynamic group and the Repel100 group).

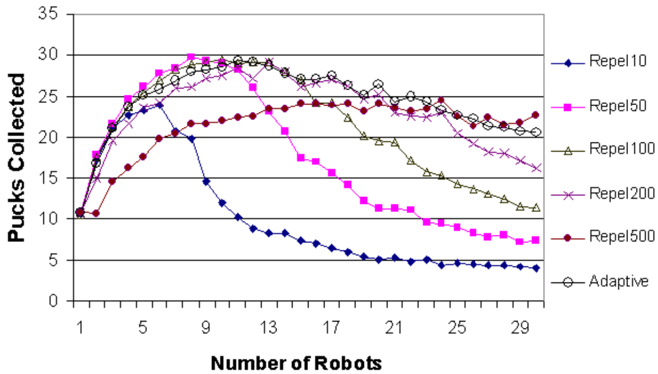

Figure 3. Adapting within Repel Groups

Our second approach used the values of V to switch between 3 distinct coordination methods — Noise, Aggression, and Repel500. The Noise group had the least costly coordination method, and was most effective in small groups up until 7 robots. At the other extreme, the Repel500 fared poorly in small groups but had the best productivity in groups larger than 17 robots. In our implementation we set the values of both W_{down} and W_{up} to be one. We set threshold values of V for each of the three states at 100, 200 and 300 accordingly. Thus, if V increased by a total of 100, the robot would assume a more robust coordination method was required and would transition to use the next most expensive coordination method, say from Noise to Aggression. If this method was still insufficient to resolve projected collisions, W_{up} would increase the value of V until the next threshold was reached. Conversely, if that method was sufficient to resolve that incident of a projected collision, the value of W_{down} would begin to decrease the value of V and the robot could eventually move down to the next lower method of coordination.

Our adaptive coordination heuristic was even more effective within this method. As figure 4 demonstrates, our adaptive group averaged significantly higher productivity than the 3 coordination methods it was based on. Once again, we performed the two-tailed paired t-test on our data and found a p-value below 0.0001 between all groups, demonstrating the strong statistical improvement. In fact, in all groups sized over 2 robots, the adaptive group always scored better than all static methods it was based on, often by more than 20 percent. This result was unexpected. We had assumed adaptation would only be capable of achieving results in line with the best levels of productivity for the methods it was based on, not significantly higher.

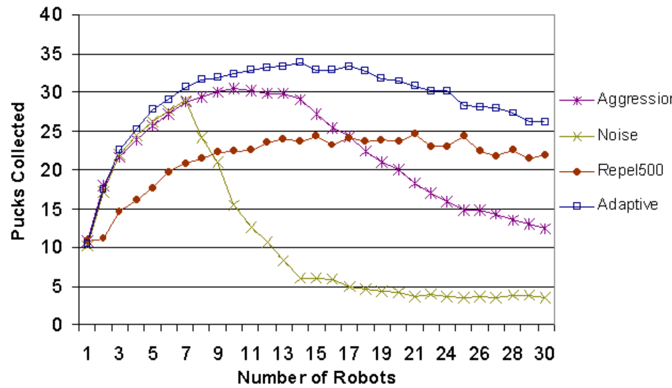

Figure 4. Adapting Between Methods

In order to understand this phenomenon, a closer analysis of our trials is necessary. As was the case in our first adaptive application, values of V were used to determine the strength of the coordination method to be used. Values of V below 100 translated to using the original method, values between 100 – 200 to the Aggression method, and values above 200 in the Repel500 method. Throughout our experiments, we found that values for V on average only ranged between 0 and 200. This implies that the adapting groups on average never used the extreme coordination method Repel500. This is slightly surprising as we found that best performance in groups over 17 robots were achieved by using this method. However, we hypothesized that the best coordination method often changes in the course of a trial. Robots operating within dynamic environments will encounter periods when collisions are more or less likely. As such, ideal performance is likely to require different methods throughout the course of even one trial, even within one group size. Thus, at times where collisions are frequent, and only then, should more robust coordination methods be used. For example, in one trial of 25 robots, the entire team spent 56 percent of their time in the original behavior, 11 percent in Aggression behavior, and 33 percent in the Repel500 behavior. Thus, the average value of V was never above 200.

Our working hypothesis is that fluctuations in the level of collisions even within one trial allow for this adaptive method to outperform the static ones it is based on. Our adaptive method is capable of capitalizing on these fluctuations by changing its fundament coordination method as needed. We believe this led towards the marked improvement in our adaptive method over the static one. For ex-

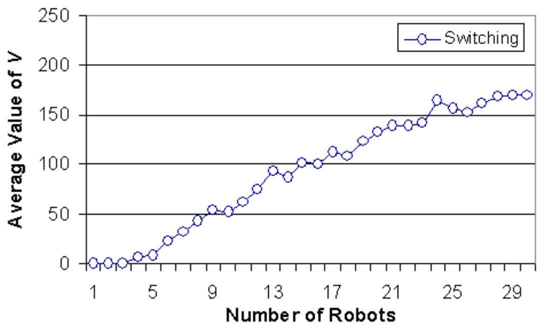

Figure 5. Threshold Values (V) in Switching Method

ample, figure 6 represents the percentage of robots that are colliding throughout the course of three trials (540000 cycles) in groups of 25 robots. The X-axis in this graph represents the number of cycles elapsed in the trial (measured in hundreds of cycles), while the Y-axis measures the percentage of robots colliding at that time. We found that these values do in fact fluctuate, at times sharply, throughout our trials. This illustrates the danger in attempting to converge on one ideal coordination method, even within one trial. As traditional learning methods require many iterations to converge, they would likely fail to be flexible in light of quickly changing environments. We believe that our adaptive heuristic based approach was able to yield significantly improved group results by quickly changing between coordination methods based on the triggers for interference. This allowed these group to reduce their overall interference and increase performance within the dynamic domain the robots operated within.

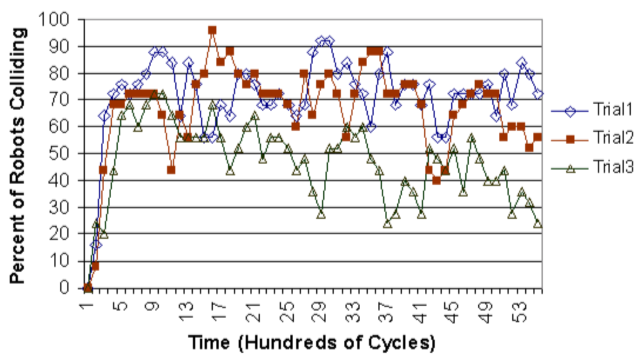

Figure 6. Fluctuations in Collisions over Time

5. Related Work

Our previous work in interference metrics [9] proved quite successful in contrasting various robotic coordination methods. In this work, we use this metric to create adaptive methods based on coordination algorithms for homogeneous robots that use no communication and are not preprogrammed to operate only within certain portions of the domain. The methods of Arkin and Balch [1], Vaughan et al. [14], and Ostergaard et al. [8] all similarity use heuristics to create group activity without communication or prior knowledge of the operating environment. Other algorithms such as those within the work of Fontan and Matarić [10] and the territorial arbitration scheme in Goldberg and Matarić [3] prevent collisions by limiting robots to specific areas within foraging domains. Jäger and Nebel [5] present an algorithm that can dynamically create these areas in a vacuuming domain, but require the robots to communicate locally. Another group of algorithms preassign values so that certain robots inherently have a greater priority to resources than others. Vaughan et al.'s fixed hierarchy system [14] and Goldberg and Matarić's caste arbitration algorithm [3] present variations of this idea for foraging robots.

We found that sharp spatial fluctuations can exist even over the course of one trial for one group size. These fluctuations complicate the problem of having robots learn from their environment. Previously, work in Mahadevan and Connell [6] found reinforcement learning based on Q Learning to be effective for a box pushing robot. While they concede that behavior based learning is especially slow to converge within robotic domains, using a behavior based approach did speed up the process. Matarić [7] studied various reinforcement learning approaches on foraging robots and stressed that the time to learn can be quite long if certain events, such as the inter-robot collisions in our domain, occur sporadically. However, the time to learn certain tasks could be diminished by using behaviors that use implicit knowledge of their domain. Both of these approaches highlight the difficulty in exclusively using traditional learning methods within robotic domains. In order to speed the time our robots adapted to spatial changes, we used methods based on heuristics, and not traditional learning methods. In tasks such as interference resolution where robots must react quickly and near-optimal results are sufficient, our method is likely to be of an advantage over those that exclusively use reinforcement learning. Additionally, fluctuations in spatial constrictions over the course of a trial may make forcing a convergence of an ideal robotic coordination method undesirable.

Several methods have also been proposed to use learning to create better coordination within robotic groups. One possibility is to have team members explicitly communicate learned information learned about their environment with others. Work by Tan [12] represents one proponent of this approach. Other approaches, such as work by Sen et al. [11] attempt to create better coordination through implicitly learning about the domain. Our robots follow this second approach as they never communicate, and create their interference gauges by observing their environment.

Our system is completely distributed as each robot contains a distinct and possibly diverse estimate of interference from that of its teammates. In contrast, work by Tews [13]used a centralized mechanism for coordinating the activities of all members of the group. While having one mechanism to oversee coordination simplifies the process of gathering inputs from various team members, it creates a single possible point of failure. Another key difference is that our work intentionally allowed robots to create estimates for interference without input from teammates. We believe this allowed robots to quickly fit a coordination to its ever changing environment as it sees fit, regardless of its teammates' conditions.

6. Conclusion and Future Work

In this paper we presented a method for dynamically adjusting coordination methods based on the conditions robots sense in its operating domain. Our use of interference metrics allowed us to create powerful adaptive heuristics based on the projected amount of interference a robot will face. We implemented two adaptive methods, one based on dynamically adjusting the strength within one coordination method, and a second based on adapting between the fundamental coordination method used by the robots in the group. In both cases we found a statistically significant improvement in performance by using the adaptive methods, with the second method strongly outperforming the basic static coordination techniques it was based on. The spatial constrictions which cause interference in the foraging domain are common to many areas such as waste cleanup, area coverage in vacuuming, search and rescue domains, and creating collision-free trajectories in restricted spaces. We believe our approach of dynamic coordination methods will greatly benefit designers of robotic teams in these domains as well.

For future work, several directions are possible. The addition of explicit communication may speed the adaption process within coordination groups. Our work also has intentionally limited itself to studying homogeneous groups of robots. An interesting study would be to apply our system to groups of heterogeneous robots. We leave for future work contrasting our approach versus those with communication and traditional reinforcement learning. One main disadvantage of our current approach lies in the manual initial work in setting the weights within our heuristic. Before our adaptive coordination methods could begin, work was needed to set the weights in our robots. It may be possible to simplify the process of initially setting these weights. One possibility would be to pass information based on previous trials and use a combination of classical reinforcement learning in addition to our heuristic based weight system. It is possible that combining these approaches might lead towards creating robot groups with even more effective adaptation.

References

[1] R. Arkin and T. Balch. Cooperative multiagent robotic systems. In *Artificial Intelligence and Mobile Robots*. MIT Press, 1998.

[2] T. Balch. Teambots, www.teambots.org.

[3] D. Goldberg and M. Matarić. Interference as a tool for designing and evaluating multi-robot controllers. In *AAAI/IAAI*, pages 637–642, 1997.

[4] D. Goldberg and M. Matarić. Design and evaluation of robust behavior-based controllers for distributed multi-robot collection tasks. In *Robot Teams: From Diversity to Polymorphism*, 2001.

[5] M. Jager and B. Nebel. Dynamic decentralized area partitioning for cooperating cleaning robots. In *ICRA 2002*, pages 3577–3582.

[6] S. Mahadevan and J. Connell. Automatic programming of behavior-based robots using reinforcement learning. In *National Conference on Artificial Intelligence*, pages 768–773, 1991.

[7] M. Mataric. Reinforcement learning in the multi-robot domain. In *Autonomous Robots*, pages 73–83, 1997.

[8] E. Ostergaard, G. Sukhatme, and M. Matarić. Emergent bucket brigading. In *Autonomous Agents*, 2001.

[9] A. Rosenfeld, G. Kaminka, and S. Kraus. A study of marginal performance properties of robotic groups, Poster to appear in AAMAS '04.

[10] M. Schneider-Fontan and M. Matarić. A study of territoriality: The role of critical mass in adaptive task division. In *From Animals to Animats IV*. MIT Press, 1996.

[11] S. Sen, M. Sekaran, and J. Hale. Learning to coordinate without sharing information. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 426–431, Seattle, WA, 1994.

[12] M. Tan. Multi-agent reinforcement learning: Independent vs. cooperative learning. In M. N. Huhns and M. P. Singh, editors, *Readings in Agents*, pages 487–494. Morgan Kaufmann, San Francisco, CA, USA, 1997.

[13] A. Tews. Adaptive multi-robot coordination for highly dynamic environments. In *CIMCA 2001*, July 2001.

[14] R. Vaughan, K. Støy, G. Sukhatme, and M. Matarić. Go ahead, make my day: robot conflict resolution by aggressive competition. In *Proceedings of the 6th int. conf. on the Simulation of Adaptive Behavior*, 2000.