A Study of Marginal Performance Properties in Robotic Teams*

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

In this paper we describe how the productivity of homogeneous robots scales with group size. Economists found that the addition of workers into a group results in their contributing progressively less productivity; a concept called the Law of Marginal Returns. We study groups that differ in their coordination algorithms, and note that they display marginal returns only until a certain group size. After this point the groups' productivity drops with the addition of robots. However, the group size where this phenomenon occurs varies between groups. To determine the cause for the differences between coordination algorithms, we define a measure of interference that enables comparison, and find a high negative correlation between interference and productivity. Effective coordination algorithms maintain marginal productivity over larger groups by reducing the team's interference levels. Using this result we are able to examine the productivity of robotic groups in several simulated domains in thousands of trials. We find that groups in theory always produce marginally, but that spatial limitations within domains cause robots to deviate from this ideal.

1. Introduction

Teams of robots are likely to accomplish certain tasks more quickly and effectively than single robots [6, 11, 9]. To date, only limited work has been performed on studying how performance scales with the addition of robots to such groups. Should one expect linear, exponential, or decreasing changes in productivity as the group size grows? Previous work by Rybski et al. [11] demonstrated that groups of identical robots do at times demonstrate marginal returns. As such, their productivity curves resembled logarithmic functions; the first several robots within their group added the most productivity per robot and each additional robot added successively less. In contrast, Fontan and Matarić [13] found that robotic groups operating within a similar domain contained a certain group size, a point they call "critical mass", after which the net productivity of the group dropped. Similarly, Vaughan et al. [15] wrote that the rule of "too many cooks" applies to their groups and adding robots decreases performance after a certain group size.

Economists have studied the gains in productivity within groups. According to their Law of Marginal Returns, if one factor of production is increased while the others remain constant, the overall returns will relatively decrease after a certain point [4]. As the size of the group becomes larger, the added productivity by each successive worker is likely to become negligible, but never negative. This classical model contains no reference to a concept similar to a "critical mass" group size after which the added worker decreases the total productivity of the group.

Our research goal is to understand when the marginal returns predicted by the economic model would be consistently realized as work by [11] found they were, and when adding robots would decrease performance as [13] and [15] described. Towards this goal, we first analyze several existing group coordination algorithms and empirically observe the different groups' productivity with the addition of robots. We observe that the different coordination techniques affect the productivity graphs of these groups during scale up. To determine the cause for the differences between coordination algorithms, we define a measure of interference that facilitates comparison, and find a high negative correlation between group interference and productivity. Effective coordination algorithms maintain marginal productivity over larger groups by reducing interference levels. Using this result we are able to examine robotic group productivity in several simulated domains in thousands of trials. We find that groups in theory always produce marginally, but that competition over space causes robots to deviate from this ideal.

* 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.

2. Comparing Group Coordination Methods

The Law of Marginal Returns, also often called the Law of Diminishing Returns, is well entrenched as a central theory within economics. Most economic domains have spatial limitations and other finite production resources. These limiting factors cause the groups' performance to typically increase marginally with the addition of labor. Brue [4] demonstrated that economists from the Enlightenment Period until modern times often did not provide empirical evidence for their theories. He concluded, ”more empirical investigation is needed on whether this law is operational” within new domains, and ”conjectures by 19th century economists about input and outputs ... simply won't do!” The first goal of this paper was to provide this robust study for robotic groups. We wished to ascertain when adding homogeneous robots will hurt group performance as [13] and [15] predict they will after a certain team size, and when adding robots continuously adds to the team's performance.

2.1. Studying Group Coordination Techniques

We began with studying the task of robotic foraging. The previously mentioned works by [11, 13, 15] all performed their research within this domain. The foraging domain has been extensively studied, and is formally defined in [7] as consisting of locating target items from a search region S, and delivering them to a goal region G. There are many applications for this domain such as waste cleanup and planetary exploration.

Previous work by Fontan and Matarić [13] noted that proper coordination lies at the root of effective group behavior. As a result, we implemented several existing group coordination algorithms on homogeneous robots without communication. We then proceeded to study how groups of these robots behave during size scale up. Our *Noise* group uses a method of mutual repelling away from obstacles such as opposing robots. This method is described by [1] and uses a repulsion schema any time a robot projects it is in danger of colliding. It additionally adds a noise element into its direction vector to prevent becoming stuck at a local minima. The *Aggression* group developed by [15] is meant to resolve possible collisions by pushing its teammate(s) out of the way. They posit that possible collisions can best be resolved by having the robots compete and having only one robot gain access to the resource in question. A third approach, given by [10], is to have the robots spread out over their operating domain by using a dynamic *Bucket Brigade* mechanism. In this method, a robot drops the item it is carrying when it detects another robot nearby. In theory, the next closest robot should retrieve the recently dropped object and carry it closer to the goal.

We added three more coordination methods. Our *Gothru* group was allowed to ignore all obstacles, and as such spent no time engaged in coordination behaviors. This ”robot” could only exist in simulation as it simply passes through obstacles such as other robots. However, this group was still not allowed to exit the boundaries of the field. We used this group to benchmark ideal performance without productivity lost because of teammates. At the other extreme, our *Stuck* group also contained no coordination behaviors but simulated a real robot. As such, this group was likely to become stuck when another robot blocked its path. Like the Stuck group, our *Timeout* group contained no repulsion vector to prevent collisions. However, these robots did add noise to the direction vector after a certain threshold had been exceeded where their position did not significantly move.

2.2. Initial Experiment Setup

We used a well tested robotic simulator, Teambots [3], to collect data. We strongly preferred using a simulator as it allowed us the ability to perform thousands of trials of various team sizes and compositions. The sheer volume of this data allowed us to make statistical conclusions that would be hard to duplicate with manually setup trials of physical robots. Using a simulator even allows us to research behaviors, such as Gothru's, that cannot exist with physical robots.

In this experiment, Teambots [3] simulated the activity of groups of Nomad N150 robots. The field measured approximately 5 by 5 meters. Our implementation of foraging followed Balch's [2] multi-foraging task in which the robots attempt to retrieve two or more types of objects. There were a total of 40 such target pucks, 20 of which where stationary within the search area, and 20 moved randomly. Each trial measured how many pucks were delivered by groups of 1 – 30 robots within 9 minutes. For statistical significance, we averaged the results of 100 trials with the robots being placed at random initial positions for each run. Thus, this experiment simulated a total of 24,000 trials of 9 minute intervals.

The simulated robots we studied were based on the same behaviors. The only software differences between the robots lay within their implementation of the previously described teamwork coordination behaviors. Each robot had three common behaviors: wander, acquire, and deliver. In the *wander* phase, the robots originated from a random initial position, and proceeded in a random walk until they detected a resource targeted for collection. This triggered the *acquire* behavior. While performing this second behavior, the robots prepared to collect the puck by slowing down and opening up their grippers to take the item. Assuming they successfully took hold of the object, the deliver behavior was triggered. At times the puck moved, or was moved by another robot, before the robot was able to take it. Once this target resource moved out of sensor range, the robot reverted once again to the wander behavior. The *deliver* behavior consisted of taking the target resource to the goal location which was in the center of the field.

We implemented a total of eight coordination methods. Our previously described *Noise* team was the default team included in the Teambots distribution [3]. Our *Stuck* and *Gothru* groups both removed all coordination behaviors but one group was blocked by teammates, while the other passed through. The *Bucket Brigade* coordination behavior was initiated once a robot detected a teammate within 2 radii. Then, these robots would drop the target being carried, move backwards for 25 cycles, and finally revert to the random walk behavior. The *Aggression* group was based on the random function of aggressive behaviors described in Vaughan et al. [15]. For every cycle a robot found themselves within 2 radii of a teammate, it selected either an aggressive or timid behavior. In order to decide, we had each robot choose a random number between 1 and 100. If the random number was lower than fifty, it became timid and back away for 100 cycles. Otherwise it proceeded forward, mimicking the aggressive behavior. As all robots within two radii choose whether to continue being aggressive every cycle, one or both of the robots assumed the timid behavior before a collision occurred. Similar to the Aggression group, our *Repel Fix* group backtracked for 100 cycles but mutually repelled like the Noise group. Our *Repel Rand* group moved backwards for a random interval uniform over 1 – 200 and also mutually repelled. Our *Timeout* group moved with a random walk for 150 cycles once these robots did not significantly move for 100 cycles. If the timeout threshold was set too low, the robot may consider itself inactive while approaching a target or its home base. However, if this value was set too high, it did not successfully resolve possible collisions in a timely fashion. We experimented with various values until we found that this combination seemed to work well.

2.3. Initial Results

Figure 1 graphically represents the results from this experiment. Our X-axis represents the various group sizes ranging from 1 to 30 robots. The Y-axis depicts the corresponding average number of pucks the group collected.

According to the economic Law of Marginal Returns, marginal returns will be achieved when one or more items of production are held in fixed supply while the quantity of homogeneous labor increases. In this domain, the fixed number of pucks acted as this limiting factor of production. Consequently, one would expect to find production graphs consistent with marginal returns. However, only the

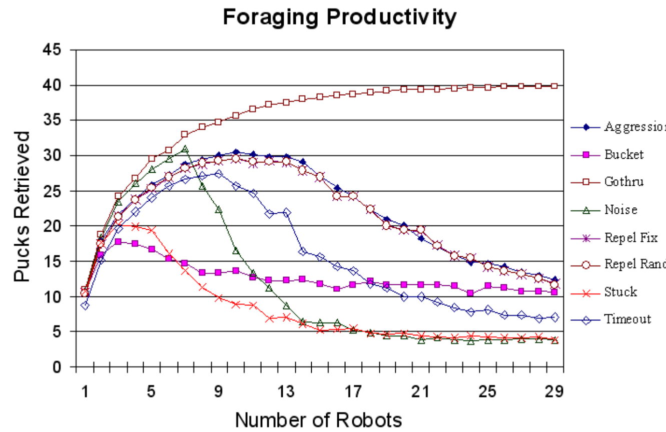

Figure 1. Initial Foraging Results

Gothru group demonstrated this quality over the full range of group sizes. All other groups contained a critical point (CP1) where maximal productivity was reached. After the group size exceeded this point, productivity often dropped precipitously. Eventually, the groups reached a level (CP2) where the addition of more robots ceased to significantly negatively effect the groups' performance.

With the exception of the Aggression, Repel Fix, and Repel Rand groups, all groups' productivity graphs differed significantly. For example, the Stuck group reached its CP1 point with an average of only 20.94 pucks collected with groups of 3 robots. The Aggression group reached a maximum of 30.84 pucks collected in groups of 13 robots. Even among equally sized groups, the differences were large. When comparing foraging groups of 10 robots, the Stuck group gathered only 8.58 pucks - far fewer than Gothru's 35.62 pucks, while the Aggression group collected 30.52 pucks, only 5.2 fewer than Gothru.

Our resulting research was motivated by these results. The Gothru group was the only group capable of realizing marginal gains throughout the entire range of 30 robots. However, many groups demonstrated the positive quality of maintaining increasing productivity over a larger range of robots. For example, the Noise group only kept marginal gains until groups of seven robots, while the aggressive group kept this quality until groups of 13 robots. We also found that the positive qualities of improved performance and maintaining marginal performance over larger groups are not always synonymous. The Noise group kept positive marginal performance over a smaller range than the Aggression group, yet performed better in groups sized seven or less. A closer look at the various coordination models was needed to draw lessons about how to create groups with both properties.

3. Why does Performance Drop?

We needed a mechanism for understanding why certain coordination mechanisms were more effective than others. We posited that differences among robotic groups were often sparked from clashes in spatial constraints. Specific to foraging, conflicts arose over which robot in the group had the right to go to the home base first. As the group size grew, this problem became more common. This caused the groups to deviate from the ideal marginal productivity, depicted by the Gothru group, by greater amounts. The length of time robots clashed with their teammates because of joint resources, such as the home base location, served as the basis in comparing coordination models within any domain.

Previous work by Goldberg and Matarić [6] found a connection between the level of interference a group demonstrated and its corresponding performance. They defined interference as the length of time robots collide, and we began by using this definition to equate between our coordination algorithms. This measure sufficed for some robots, such as those simulated by the Stuck group, because they did not engage in any other coordination behaviors. However, this metric of interference could not explain the differences between all groups. Many robots, such as those simulated by the Aggression group, never collided. If one takes the position that only collisions constitute interference within robotic groups, these robots do not interfere. Yet we clearly observed how the addition of robots detracted from the groups' productivity after its maximal productivity point.

In this section we present our measure of interference. We describe scale up experiments in foraging and search domains that are characterized by resources that lend themselves to group conflicts. We find that interference and productivity are strongly negatively correlated in such domains, and use this metric to explain differences in productivity between all teams. We posit that in the absence of spatial conflicts, all teams should consistently demonstrate marginal gains during scale up. We confirm this idea by easing the "space crunch" in our domains and notice how all groups consistently demonstrate marginal returns. We conclude that any domain with group spatial conflicts will suffer from deviations in marginal performance once interference cannot be resolved.

3.1. Interference: Measure of Coordination

We define interference as the length of time an agent is involved with, either physically or computationally, projected collisions, real or imaginary, from other robots and obstacles. This period of involvement often extends well beyond the actual collision between two robots. Any time spent before a supposed collision in replanning and avoidance activities must also be recorded. Similarly, all post-collision resolution activity must be included as well. Thus, according to our definition, the Gothru group has zero interference because it never engages in any interference resolution behaviors and represents idealized group performance. The Aggression group engages in interference resolution behaviors before a collision ever happens. Its various timid and aggressive behaviors to avoid collisions all constitute interference by our definition. The Bucket Brigade group demonstrates that interference can exist after a collision is prevented. For this group, one needs to measure the productivity lost by handing off the resource from one robot to the next. Many times this group lost productivity during this process because the second robot never properly took the dropped target. Only this measure takes into the account the total interference resolution process.

According to our hypothesis, we expected to see a negative correlation between levels of interference and productivity in three respects. We reasoned that the degree to which a group deviates from the idealized marginal gains is proportional to the amount of average interference within the group. This can impact where the group hits maximal performance. Those groups which reached CP1 with a small number of robots spiked high levels of interference much faster than those where this point was delayed. Second, even before groups hit their maximum productivity point, we hypothesized that the more productive groups have lower levels of interference than their peers. Finally, we expected that differences in where the productivity of the groups eventually plateau can be attributed to the group's saturation level of interference. Those robots that more effectively deal with interference even in large groups will have CP2 values at higher levels.

In order to confirm this hypothesis, we reran our eight foraging groups and logged their interference levels according to our definition. The Gothru group never registered any interference. For all remaining groups, we used the simulator to measure the number of cycles the robots in the groups collided. For all groups other than the Stuck and Gothru groups, we additionally measured the number of cycles the robots triggered interference resolution behaviors when they were not colliding. In the Noise and repulsion groups, this represented the number of cycles spent in repelling activities. In the Aggression group, it was the number of cycles spent in timid and aggressive behaviors. In the Timeout group, this was the cycles spent trying to resolve a collision once the robot timed out. In the Bucket Brigade group, internal behaviors alone did not suffice to measure interference by our definition. We only recorded cycles spent when the robots came close to another and consequently dropped the resource they were carrying. However, we could not measure the time lost when the second robot did not effectively take that resource as we did not have omnipotent knowledge of such events. As a result, our measurement for interference for this group did not necessarily represent an exact measurement, but an underestimate.

Figure 2 represents the result from this trial. The X-axis once again represents the group size, and the Y-axis represents the average number of interference cycles that each robot within the group registered over the 100 trials.

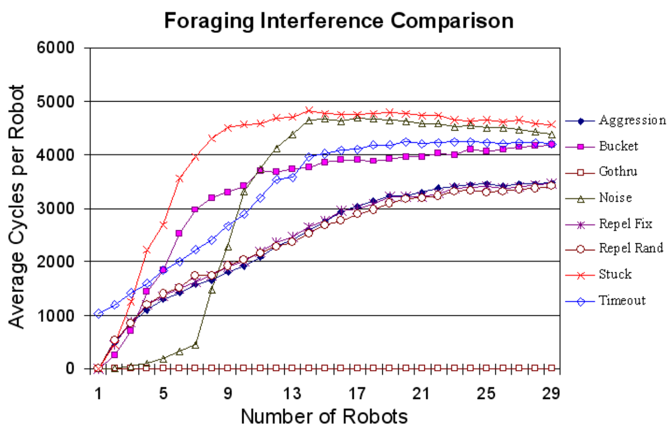

Figure 2. Foraging Interference

We found that CP1 typically occurred for all groups when the average interference level within each robot of the group reached a level between 2500 and 3000 cycles. The longer the group was able to maintain classically diminishing returns, the more cycles of interference were needed to cause the critical point. This is because CP1 will only be reached once the productivity lost due to interference is larger than the total marginal productivity of the group. Before this point, the total production of the group increases, albeit marginally. For example, the Stuck group, which reached its critical point with only four robots, needed closer to only 2500 cycles to cause this critical point. The Aggression group hit CP1 with 13 robots, and consequently needed closer to 3000 cycles to counter the productivity of more robots.

Even when viewing the differences between productivity among equally sized groups, interference differences were significant. We found a very strong average negative correlation of -0.94 between the differences in groups' performance and their interference level over the entire range of 1 to 30 robots. For example, the Noise group most closely followed the idealized Gothru productivity graph for groups up until 7 robots, and registered significantly less interference than the other groups. This interference resolution mechanism had little overhead, and needed fewer cycles to resolve a possible collision. However, this method didn't scale well beyond this point. When the group size became larger than seven, its interference levels grew exponentially and the group's performance quickly decayed. In contrast, the Aggression and other repelling groups had significant levels of interference from the onset, but interference levels only grew linearly with respect to the group size. As a result, this group proved more effective with larger group sizes.

We also found that the eventual performance plateau (CP2) was strongly correlated with interference. Some groups levelled off at significantly smaller interference levels than other groups. For example, even in group sizes above 20 robots, the Bucket Brigade group registered an average interference level of 400 fewer cycles less than the Stuck group. Consequently, it collected on average over 5 pucks more than this group at this level.

As one would expect, most groups performed equally well with one robot, as coordination behaviors should only be triggered in groups of two robots or more. The one exception was the Timeout group which collected on average 8.7 pucks with one robot, or about 2 pucks fewer than the other groups. As we defined interference as the time spend on resolving collisions, or even perceived collisions, such a result is quite plausible. At times these robots timed out while slowing down to pick up a puck or avoid an obstacle even by themselves. As we defined such internal reasoning as interference, these robots interfered with themselves in the amount of about 1000 average cycles per trial.

Two of our groups have slight underestimates for interference; however, this did not change our overall results. As previously mentioned, the Bucket Brigade group interfered if a second robot did not successfully receive the resource handed off to it. We found that this did occur at times when there were relatively small groups of these robots. Thus, the correlation between their productivity and that of other groups' among groups of 2–6 robots dropped to -0.80. By discounting this range, the average overall correlation reached almost -0.97. However, after 6 robots we found that there were enough robots in the area to ensure a second robot would quickly take the resource, and the amount of this underestimate was less significant. The Noise group also registered an underestimate for interference. These robots actually used two repulsion fields for collision resolution. They triggered a strong repulsion field when they sensed another robot or obstacle 0.1 meters away. We only measured the number of times this repulsion field was triggered. However, a second, much weaker repulsion field was triggered from 1.5 meters away. In this instance, our underestimate did not seem to significantly statistically detract from our results. With or without the data from this group, the average correlation between groups was -0.94.

3.2. Competing over Spatial Resources

We proceeded to study if our results were limited to foraging or were a general phenomenon seen when robotic groups are faced with restriction production resources. We created a new spatially limited search domain where the task goal was to find the exit out of the room as quickly as possible. We placed groups of robots within a room of 1.5 by 1.5 meters with one exit 0.6 meters wide. We reasoned a critical productivity point would once again form in this domain. Too few robots would result in a long search time until the exit was found. However, too many robots would cause interference as the exit was only physically wide enough for one robot.

We ran simulated trials of seven of our eight foraging groups ranging in sizes from 1 - 23 robots (the room holds 23 robots) and averaged the results from 100 trials for statistical significance. We omitted the Bucket Brigade group as this coordination method was not relevant to this domain. We then measured the length of time it took the first robot from each group to completely exit the room. We ended the trial at that point and recorded the time elapsed. Thus, this experiment constitutes over 16,000 trials of variable length.

Figure 3 presents our productivity graphs and corresponding interference levels from this experiment. The X-axis in both graphs depict the size of our groups. In the upper section, we flipped the Y-axis to represent the search time of zero as the highest point. As in our foraging graphs, we represent better performance as higher values in this graph. In the lower graph the Y-axis represents our average measurement of interference per robot in the group.

We found that the time to complete the search task was strongly negatively correlated in our new domain as well. We observed that with the exception of the Gothru group, all groups ceased to demonstrate marginal returns at some point. In the Repel Fix group this point occurred with only 5 robots, while the Timeout group reached this point with 14. The Noise group had the lowest level of interference through groups of 16 robots, and was able to most closely approximate Gothru's performance until this group size. After this point the Timeout group fared the best. We found that certain interference resolution mechanisms work best in specific domains. While the repulsion methods were quite effective in foraging, the interference levels in these groups grew exponentially in this domain. Overall, the average statistical correlation for groups of 1-23 robots between the time elapsed to exit the room and their corresponding interference level was -0.94.

3.3. Easing Spatial Restrictions

According to our hypothesis, deviations of productivity in robot groups are strongly correlated with interference. Once our foraging and search groups ceased to successfully resolve inference they realized their critical group sizes. Adding further robots only hurt the groups' performance. We posit that the physical space limitations existent within many robotic groups often cause this interference. The one home base area within the foraging domain and the one exit within the search domain create a competition over space between robots that cannot always be properly resolved.

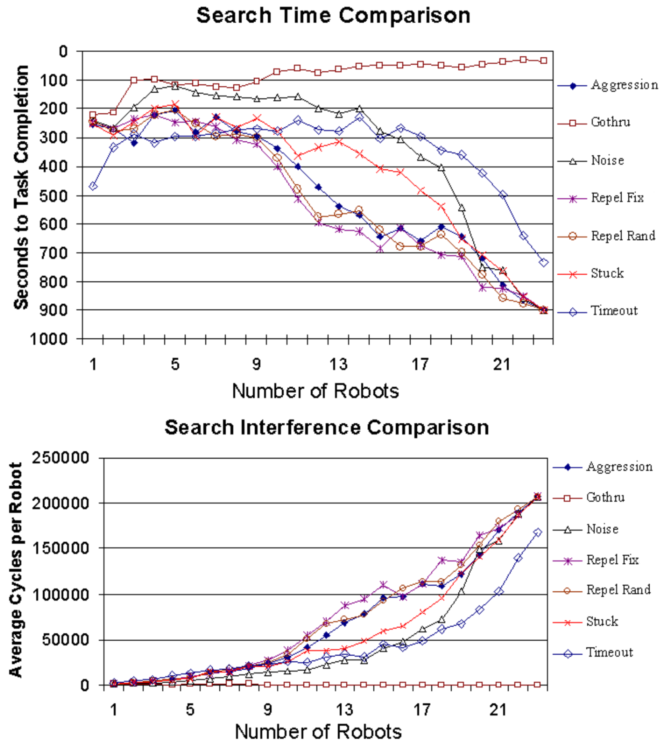

Figure 3. Search Time and Interference

We were able to confirm that our robotic groups always demonstrated marginal returns once restrictions over physical space were eased. We changed the foraging group requirement of returning the pucks to one centralized home base location. Instead, they were allowed to consider the pucks to be in the home base immediately. With the exception of the Bucket Brigade group, we reused all 8 previously studied foraging groups. Once again, we omitted this method because it was not applicable to our new domain. We left all other environmental factors such as the number of trials, the size and shape of the field and the targets to be delivered identical. Thus, Teambots [3] simulated 21,000 trials of 9 minute intervals in this experiment.

As figure 4 shows, all groups did indeed always achieve marginal returns in the modified foraging domain. While Gothru still performed the best, the differences between it and other groups' coordination methods were not as pronounced. The level of interference all groups demonstrated was also minimal, and thus not displayed. We concluded that not every foraging domain needed to have a critical point for productivity where marginal gains during scale up ceased.

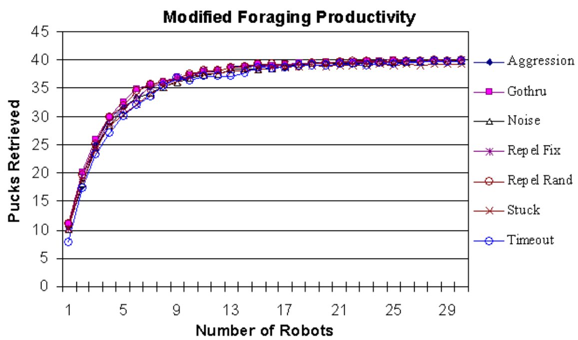

Figure 4. Modified Foraging Domain

Within the search domain, we hypothesized that limitations in the room size and width of the exits created the large amounts of interference during scale up. In order to ease this restriction, we doubled the size of the room to become approximately 3 by 3 meters, and widened the exit to allow free passage out of the room by more than one robot. Once again, we measured the time elapsed (in seconds) until the first robot left the room and averaged 100 trials for each point. This experiment also constituted over 16,000 trials of varying lengths. Figure 5 graphically shows that our modified domain consistently realized marginal increases in faster search times with respect to group size. Once again, interference levels were also negligible in our new domain. Thus, we concluded that achieving marginal productivity gains was always possible once competition over spatial resources was removed.

4. Related Work

Many algorithms have been developed to assist in robotic group coordination. We studied the methods of Arkin and Balch [1], Vaughan et al. [15], and Ostergaard et al. [10] that can be implemented on homogeneous robots. All of these methods resolve spatial conflicts without foreknowledge of the operating domain and do not have any need for communication. Other algorithms exist that require advance knowledge of the physical details of the operating domain. Examples of these algorithms include the territorial allocation method developed by Fontan and Matarić [13] and the territorial arbitration scheme in Goldberg and Matarić [6]. Both methods limit each foraging robot to a specific area or zone and thus prevent collisions. Jäger and Nebel [9] presented 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. This group of coordination methods is similar to the Aggression method we studied, but it preassigns robots to be aggressive or meek. The fixed hierarchy system within Vaughan et al. [15] and the caste arbitration algorithm within Goldberg and Matarić [6] implemented variations of this idea on foraging robots.

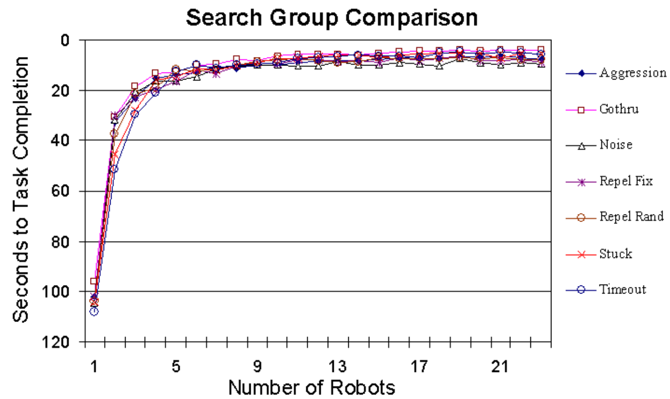

Figure 5. Modified Search Domain

To date, only limited work exists on improving robot scalability. The work by Fontan and Matarić [13] found that groups of 3 robots performed best within their foraging domain. Adding more robots only hurt performance when using their territorial coordination method. Jäger and Nebel [8] presented a collision avoidance technique for use in trajectory planning among robot groups that requires local communication. They noted that their coordination method will not scale beyond groups of 4 robots. Rybski et al. [11] found marginally increasing productivity up to groups of 5 foraging robots, but did not study larger sizes. Our use of interference to contrast coordination methods during scale up is based on Goldberg and Matarić [6], but their definition of interference as collisions did not suffice in our study.

Within the general agent community, Shehory et al. [14] presented a scalable algorithm for a package delivery domain suitable for groups of thousands of agents. He based his algorithm on concepts borrowed from physics. Later work by Sander et al. [12] studied how computational geometry techniques could be applied to groups in the same domain. Both found that group productivity did scale marginally with the addition of agents and that a point existed where adding agents did not significantly improve the productivity of their system. Their agents did not compete over physical space, and they never found that adding agents hurt group performance. Specific to the search domain, work by [5] studied the scalability qualities of their PHA* algorithm, and found that their algorithm yields marginally better results with the addition of agents. Their agents also never impeded the proper functioning of their teammates, and thus search performance times always improved during scale up.

We demonstrated in our paper that the spatial restrictions within robotic domains often prevented marginal gains from being realized as group sizes grew. The corollary of this hypothesis is that marginal returns will always be achieved in domains that do not clash over resources. It is not surprising that groups of agents should therefore always realize marginal returns during scale up once group interference issues have be resolved or are not applicable.

5. Conclusion and Future Work

In this paper we presented a comprehensive study on the productivity of robotic groups during scale up. As the size of robotic groups grew, effective coordination methods were critical towards achieving optimal or near optimal team productivity. The limited space inherent in many environments, such as the foraging and search domains we studied, make this task difficult. Using our novel, non-domain specific definition of interference, we were able to equate between the effectiveness of various existing coordination algorithms. Our interference metric measured the total time these robots dealt with resolving team conflicts and found a strong negative correlation between this metric and the corresponding productivity of that group. Groups demonstrated marginal gains only when their interference level was low. If the new robot added too much interference into the system, it detracted from the group's productivity and marginal productivity gains would cease.

Many robotic domains also contain the limited space and production resources that our foraging and search domains exemplify. We predict robotic groups involved with planetary exploration, waste cleanup, area coverage in vacuuming, and planning collision-free trajectories in restricted spaces will all benefit from use of our interference metric. We plan to implement teams of real robots in these and other domains in the future.

This paper limited its study to homogeneous robots without communication. Additionally, we did not study coordination methods which require pre-knowledge of their domain or algorithms that use other forms of preprocessing. We leave the study of how to widen our metric to allow contrasting robots with differing capabilities such as communication, foreknowledge of domains, and preprocessing requirements for future work. This paper demonstrates how the best coordination method can be observed by comparing different interference levels. We also plan to expand our work to dynamically select the best coordination method in any high interference domain. By creating robotic teams that adjust to the triggers of interference, we believe it will be possible to improve the performance of these groups.

References

[1] R. Arkin and T. Balch. *Cooperative Multiagent Robotic Systems*. In Artificial Intelligence and Mobile Robots, MIT Press, 1998.

[2] T. Balch. Reward and diversity in multirobot foraging, In IJCAI-99 Workshop, 1999.

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

[4] S. L. Brue. Retrospectives: The law of diminishing returns. *The Journal of Economic Perspectives*, 7(3):185–192, 1993.

[5] A. Felner, R. Stern, and S. Kraus. PHA*: Performing A* in unknown physical environments. In *AAMAS 2002*, pages 240–247.

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

[7] 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.

[8] M. Jager and B. Nebel. Decentralized collision avoidance, deadlock detection, and deadlock resolution for multiple mobile robots. In *IROS 2001*.

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

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

[11] P. Rybski, A. Larson, M. Lindahl, and M. Gini. Performance evaluation of multiple robots in a search and retrieval task, In Workshop on Artificial Intelligence and Manufacturing, 1998.

[12] P. Sander, D. Peleshcuk, and B. Grosz. A scalable, distributed algorithm for efficient task allocation. In *AAMAS 2002*, pages 1191–1198.

[13] M. Schneider-Fontan and M. Matarić. From animals to animats iv, pages 553–561. MIT Press. 1996.

[14] O. Shehory, S. Kraus, and O. Yadgar. Goal satisfaction in large scale agent-systems: A transportation example. In *ATAL-98, pages 277-292*.

[15] 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*.