COMBINING MULTIPLE AUTONOMOUS MOBILE SENSOR BEHAVIORS USING LOCAL CLUSTERING

Rustam Stolkin
Jeffrey V. Nickerson
Stevens Institute of Technology, NJ
{rstolkin, jnickerson}@stevens.edu

ABSTRACT

In crisis situations it may be necessary to re-establish communications via ad-hoc networks of communicating way-stations. Crisis, defence, or surveillance scenarios may require the distribution of sensor units over some region of interest. In both cases the use of communicating and sensing, autonomous, mobile robots will become prevalent in the near future.

Existing literature describes various control rules for forming teams of robots into optimally distributed communicating or sensing grids. We show how these robot behaviors exhibit advantages and disadvantages related to the range between robots. Further we describe a new approach which enables several different behaviors to be combined, utilizing the best behavior for any particular inter-robot range. This method is demonstrated using computer model simulations and the improved performance of the combined behaviors, relative to that for individual behaviors, is graphed over important situational parameters.

INTRODUCTION

This paper investigates distributed control rules for coordinating the movement of a collection of communicating robots, autonomous mobile sensors or other kinds of mobile unit, so as to create useful formations. In particular, we address the problem of enabling a number of robots, starting from random positions, to form an equi-spaced lattice over some designated region.

The work relates to the theme of this workshop, situation management, in the following way. It has been observed that emergency, military and other situations can often be characterized as involving the motion of many objects through time and space (Kokar 2004). This paper considers situations in which many mobile sensing units are being deployed into an area. In many such cases, the optimal final configuration of these objects is an equi-spaced lattice. However, in other cases, it may be useful for the sensor clustering density to increase towards the periphery of the group, perhaps to protect against external objects or to quarantine internal objects. As we will see, distributed control rules can be adjusted to achieve these different configurations, and thus, might play a role in emergency situation management.

This work has many possible applications. When monitoring, measuring or surveilling an area of territory, it will often be desirable to disperse a large number of sensors over that area. Often an equi-spaced lattice of sensor placements will be optimal, but other kinds of lattice spacing may also be useful, depending on the situation. In many crisis situations, existing communications networks are observed to break down (e.g. Maloney and Shays 2004). Alternatively, when first responders arrive in a new location, no fixed communications network may exist. In these cases it may be desirable to rapidly create an ad-hoc, wireless communications network using a number of mobile relaying nodes. The optimal distribution for these nodes will also often be an equi-spaced lattice.

It is desirable that such coordinated movement be achieved by means of distributed control rules, i.e. each mobile unit should be autonomous and determine its own movement independently. This strategy eliminates the need for large computational burden on a centralized controller, makes the system robust against communications breakdown to such a controller and allows flexible adaption to local conditions, obstacles or the breakdown of individual units.

This paper makes the simplifying assumption that each robot is aware of the positions of all other robots. This information might be derived from a combination of onboard sensors which enable a robot to determine its own position, onboard sensors which detect the positions of other robots, and communication of this positional information from one robot to another. Future work may explore scenarios in which robots use limited, incomplete or erroneous positional data.

This paper examines several popular control rules from the literature and some new variants. We show how each control rule has advantages and disadvantages which relate to inter-robot range. We then propose a new method, which enables several different behaviors to be combined, utilising the best behavior for any particular inter-robot range. The control rules and the new combination method are demonstrated using computer model simulations. Finally we present results, graphing the improved performance of the combined behavior as compared to the individual behaviors.

SPRING MODELS

Many existing control rules have been inspired by the behavior of physical or biological systems. Some models (e.g. Martinson and Payton, 2004) make the robots behave as if each were connected to its nearest neighbors by springs (figure 1).

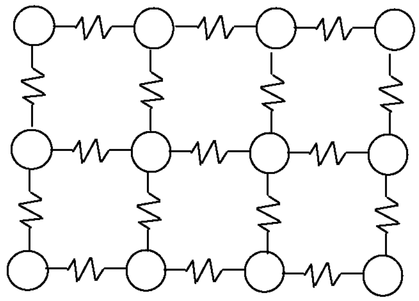

Figure 1. Spring lattice model. Circles represent robots.

Each robot controls its acceleration as if acted upon by a force:

$$\mathbf{F} = k\sum_{i=1}^{4}\left[(\mathbf{x}_i - \mathbf{x}_r) - \frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|}R\right] \qquad \textbf{(1)}$$

where $\mathbf{x}_r$ is the robot's own position, $\mathbf{x}_i$ are the positions of the four adjacent neighbors in the desired configuration, k represents the stiffness of the springs and R is the unstretched length of the spring.

This model is useful in that it should cause robots to attract to equilibrium positions at points in an equi-spaced lattice. It also results in the efficient behavior of high accelerations when robots are far from their desired positions. It is limited in that each robot must have a pre-designated position in the final lattice. This will not result in optimal position assignments given most random starting configurations, nor will it be robust against breakdown of individual robots. It is preferable that any robot be able to assume any position in the final lattice. Automatic re-allocation of lattice position is not a trivial problem.

Additionally, a spring model will often lead to indefinite oscillation of robots. This can be prevented by incorporating damping into the model. Spears and Spears (2004) suggests the incorporation of "friction". Balch and Arkin (1998) rely on the use of a "dead-zone" of zero motion incentive around destination positions.

We have experimented with an alternative spring model, in which each robot behaves as if it were connected to every other robot by a spring (figure 2). This will also result in the desirable behavior of rapid accelerations at large distances, but allows each robot to flexibly choose its final position in the lattice, depending on the starting positions.

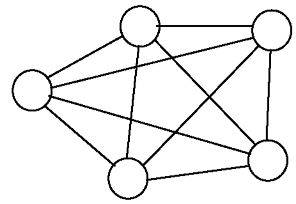

Figure 2. Alternative spring model. Each robot attracts to every other.

Robots can be prevented from oscillating by connecting each robot to each other robot with a viscous damper or "dash-pot" in parallel with each spring (figure 3).

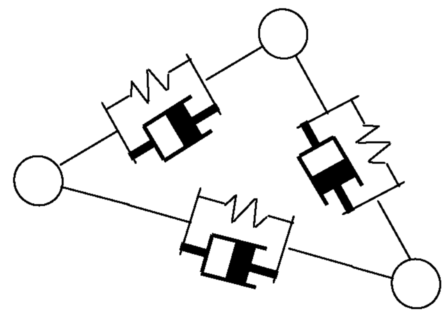

Figure 3. Spring and damper model

Now each robot r controls its acceleration as if acted upon by a force:

$$\mathbf{F} = k\sum_{i=1}^{n}\left[(\mathbf{x}_i - \mathbf{x}_r) - \frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|}R\right] - c\sum_{i=1}^{n}\frac{d}{dt}(\mathbf{x}_r - \mathbf{x}_i) \qquad \textbf{(2)}$$

where $\mathbf{x}_r$ is the robot's own position and $\mathbf{x}_i$ are the positions of all n robots of the entire group, k is the stiffness of the springs and c is the viscosity of the damping system. This is equivalent to a Proportional Derivative (PD) controller. A simpler damping method simulates motion through a viscous fluid:

$$\mathbf{F} = k\sum_{i=1}^{n}\left[(\mathbf{x}_i - \mathbf{x}_r) - \frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|}R\right] - c\frac{d\mathbf{x}_r}{dt} \qquad \textbf{(3)}$$

Control rules based on this damped spring model do result in an ordered distribution of robots over an area (figure 4). However, for larger numbers of robots, the inter-robot spacing is observed to vary with distance from the center of the group (figure 5). Robots become more tightly packed with radius from the group center. This is due to robots being more strongly attracted to distant individuals than to close neighbors.

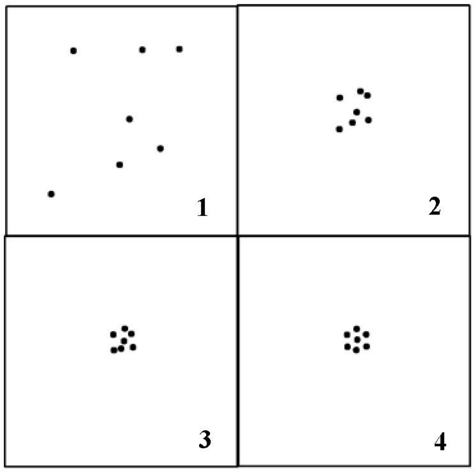

Figure 4. Four stages in chronological order, from a simulation of the motion of 7 robots obeying a spring and damper control rule.

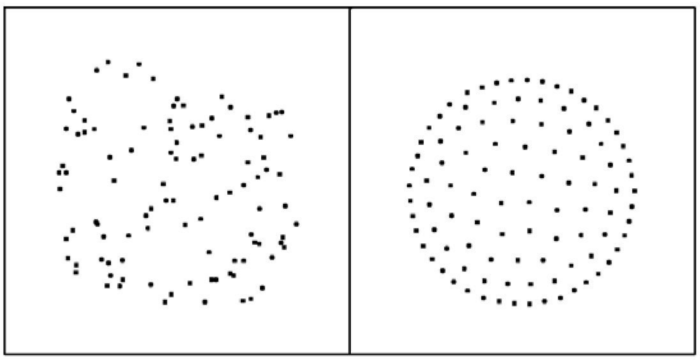

Figure 5. Initial and final configurations for a group of 100 robots obeying a spring and damper control rule. Notice that the inter-robot spacing in the final configuration *decreases* with distance from the centre of the group. This is most easily observed by looking at the circumferential spacing along each concentric ring.

VELOCITY CONTROL METHODS

Similar behavior can be produced much more simply, merely by controlling the velocity of each robot to be inversely proportional to the sum of the displacements to all other robots. Robots can be prevented from coalescing by measuring this displacement from a ring of pre-specified radius around each robot (figure 6).

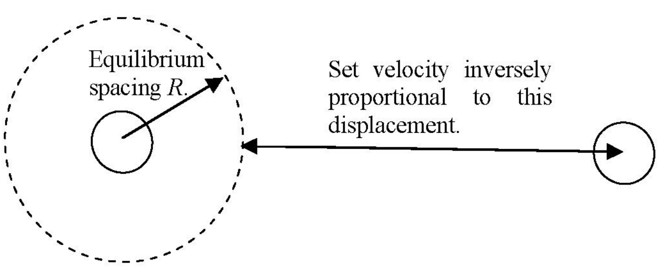

Figure 6. A simple velocity control rule.

Now, each robot directly controls its own velocity according to:

$$\mathbf{V} \propto \sum_{i=1}^{n} \left[(\mathbf{x}_i - \mathbf{x}_r) - \frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|} R \right] \qquad \textbf{(4)}$$

where $\mathbf{V}$ is the desired robot velocity, R is the desired final inter-robot spacing, $\mathbf{x}_r$ is the robot's own position and $\mathbf{x}_i$ are the positions of all other n robots in the group.

It can easily be shown that this velocity control rule is equivalent to the spring/damper rule for the case of critical damping, with the inter-robot distance decaying exponentially in each case. This rule therefore also results in distributions in which the inter-robot distance decreases with radius from the group center (figure 7). As we have mentioned before, such configurations might be useful in focusing attention on the periphery of a region.

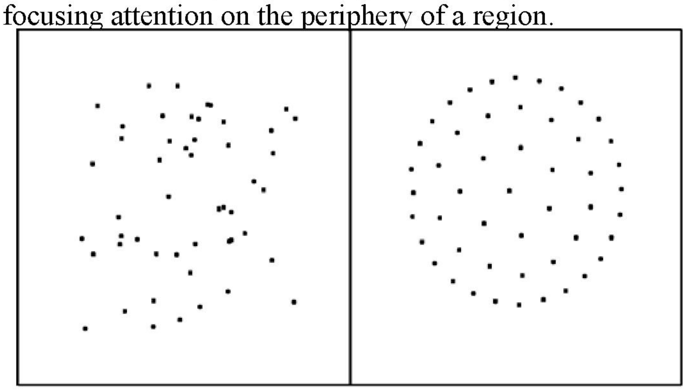

Figure 7. Initial and final configurations for a group of 50 robots obeying a velocity control rule. Notice the increased crowding towards the extremities of the group. This is most easily observed by looking at the circumferential spacing along each concentric ring of robots.

ARTIFICIAL GRAVITY MODELS

Spears and Gordon (1999) use a model inspired by gravitation which they refer to as *artificial physics*. In such models, each robot is attracted to every other robot by a gravity-like inverse square attraction law, centered on a point at a pre-specified radius from that robot:

$$\mathbf{F} \propto \sum_{i=1}^{n} \left\{ \frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|} \cdot \left[(\mathbf{x}_i - \mathbf{x}_r) - \frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|} R \right]^{-2} \right\} \qquad \textbf{(5)}$$

One disadvantage of this method is that, if robots are separated by large distances, the initial response will be extremely slow, with small accelerations, poor exploitation of the robots' maximum speeds and long convergence times. In contrast, spring methods produce the desirable behavior of increased accelerations with increased robot separations. Another disadvantage of gravitation models is that attraction to a ring around a robot implies an equilibrium position at the center of that ring. Thus if two robots overshoot while converging (thus crossing the desired distance of inter-robot spacing and becoming too close to each other), there may be little incentive for the overshot robots to correct their positions. In several reported experiments (e.g. Spears and Spears, 1999) one observes a grid of robots which forms a predominantly correct distributed lattice, but where an extra individual has become trapped interstitially for this reason (e.g. figure 8).

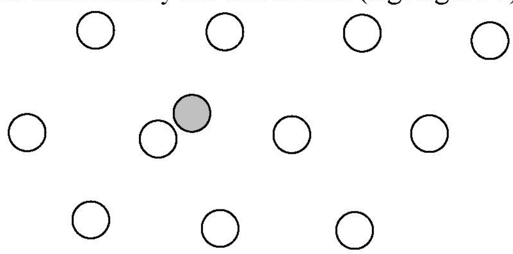

Figure 8. One robot becomes interstitially trapped.

POTENTIAL FIELD MODELS

Several authors (e.g. Dudenhoeffer, 2000) make use of the attraction/repulsion combination:

$$|\mathbf{F}| = \left(\frac{A}{r^{\alpha}} - \frac{B}{r^{\beta}} \right) \quad \textbf{(6)}$$

where $\mathbf{F}$ is a force corresponding to the desired acceleration of the robot, r is the distance between the robot being controlled and another robot which attracts it, and other values are pre-specified constants. The direction of the force $\hat{\mathbf{F}}$ is in the direction of the attracting robot.

This model is an example of a class of control rules which model potential fields of attraction and/or repulsion (see Vail and Veloso, 2003, and Khatib, 1985). These models are sometimes referred to as "social potential fields" (Reif and Wang 1998), and are often said to be inspired by flocking, herding or schooling behavior in the animal kingdom (Gazi and Passino 2002).

These models do result in approximately equi-spaced final distributions of robots. For two robots, the inter-robot equilibrium distance is given by:

$$r = \left(\frac{A}{B} \right)^{\frac{1}{(\alpha - \beta)}} \quad \textbf{(7)}$$

For larger numbers of robots, the relationship is more complicated since robots are attracted to distant individuals in addition to their closest neighbors. In fact, to produce truly equi-distant spacings, each robot must attract only to its closest neighbors and not to other distant individuals. The inverse square attraction law approximates this well since the inter-robot forces drop off rapidly outside of each robot's local neighborhood (in contrast, spring-like forces increase with distance and decrease for close neighbors). The higher the value of α, the better will be the approximation to equi-distant robot spacing in the final configuration.

High α values produce desirable equi-distant robot spacings because they result in attractive forces which decrease rapidly with distance. For the same reason, these forces have the disadvantage of causing slow responses, especially when the initial inter-robot spacing is large. For these reasons, an α value of 2 is often a good compromise between speed of convergence and evenness of final robot spacing.

ATOMIC ATTRACTION/REPULSION MODEL

In our opinion a better, and seemingly obvious analogy is that of attraction/repulsion between atoms in a crystal lattice. The atoms of many elements form regular, distributed patterns, with a constant inter-atomic spacing. If one aims to cause a group of robots to distribute equidistantly over a surface according to a distributed control law, then it seems sensible that each robot should mimic the behavior of such atoms, since this is the way that such atoms behave and their attraction/repulsion laws are well known. The attraction/repulsion behavior of atoms in many materials, is in fact a subset of the behavior described above, with an α value of 2 :

$$\mathbf{F} = \left(\frac{A}{r^{2}} - \frac{B}{r^{\beta}} \right) \quad \textbf{(8)}$$

Where β is dependent on the element (e.g. Sodium atoms have a value of around 7).

We also need to include a damping term. A physical analogy for this term might be heat loss from a vibrating crystal lattice of atoms. We provide damping by causing the robots to behave as though moving through a viscous fluid. The robot control law now becomes:

$$\mathbf{F} = \sum_{i=1}^{n} \left[\frac{(\mathbf{x}_i - \mathbf{x}_r)}{|\mathbf{x}_i - \mathbf{x}_r|} \cdot \left(\frac{A}{|\mathbf{x}_i - \mathbf{x}_r|^2} - \frac{B}{|\mathbf{x}_i - \mathbf{x}_r|^{\beta}} \right) \right] - c\frac{d\mathbf{x}_r}{dt} \quad \textbf{(9)}$$

The atomic attraction/repulsion model is a useful control rule in that it results in approximately equi-spaced robot configurations (figure 9).

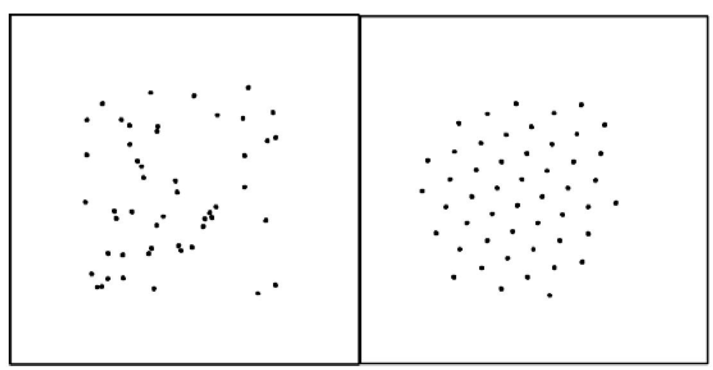

Figure 9. Initial and final configurations for a group of 50 robots obeying an atomic attraction/repulsion control rule. Notice that, in contrast to the spring and velocity models, this rule results in equal inter-robot spacings throughout the lattice.

Unfortunately, as with the gravitation model, the atomic attraction/repulsion model suffers from slow convergence because of the inverse-square attraction rule.

LOCAL CLUSTERING

We have seen that distributed robot control rules based on spring-like attraction models offer the efficiency advantage of rapid accelerations at large inter-robot ranges. However, spring-like attraction is not able to produce robot formations with equal inter-robot spacings even though these formations are highly desirable in many sensor or communications scenarios. In contrast, control rules based on atomic attraction/repulsion models do produce the desired formations but suffer from very slow convergence due to weak inter-robot attraction at large ranges.

We now describe a new method, which exploits the advantages of each of these models, enabling high speed spring-like attraction at long ranges combined with equi-spaced formation producing atomic behavior at close ranges. This method, which we refer to as "local clustering" behavior consists of a simple heuristic control rule. Each robot, along with its close neighbors, can be described as belonging to a "local cluster". Membership of the cluster extends to any new robot which moves to within a specified radius of any existing cluster member. Thus, if two clusters should meet, they will merge into a single larger cluster.

The process of establishing cluster memberships is useful in that it enables the simultaneous adoption of two different kinds of robot behavior. Clusters can move as a whole unit, relative to other clusters, according to one kind of behavior model whilst, simultaneously, within each cluster, cluster members move relative to other members according to a different behavior model.

The overall behavior of each robot is thus the superposition of the group behavior of the cluster to which the robot belongs and the behavior of the robot relative to other members of its cluster.

By setting cluster-to-cluster behavior to follow a spring-like attraction rule and member-to-member behavior to follow an atomic attraction/repulsion rule, the best aspects of each kind of behavior are combined. Now distant robots rapidly accelerate towards the center of the group, using spring-like attraction, but arrange themselves in an orderly atomic lattice once they get there (figure 10).

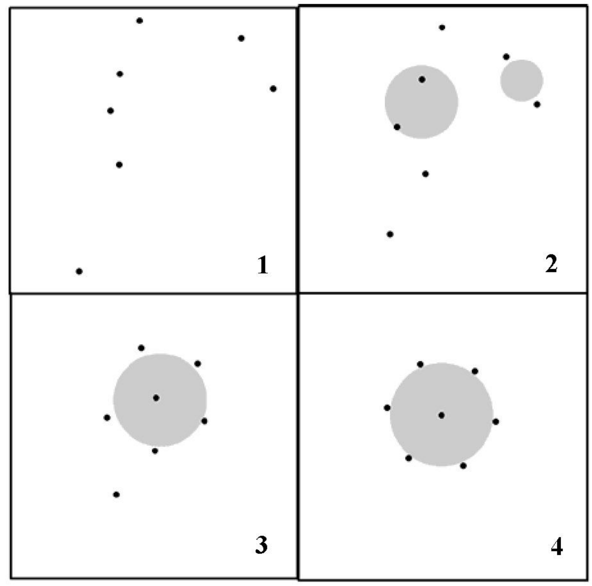

Figure 10. Four stages in chronological order, from a simulation of the motion of 7 robots using the local clustering heuristic. The large, shaded circles are intended to illustrate the position and approximate membership radius of clusters and are centered at cluster centroids.

Additionally, if several robots meet during their journey towards the main group, they form a small traveling cluster which begins to arrange and order itself during that journey (figure 11).

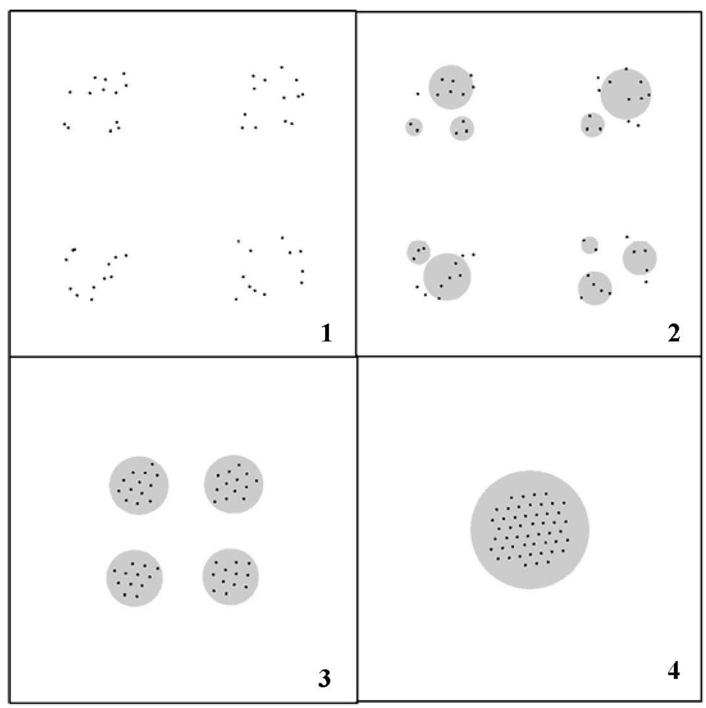

Figure 11. Random initial positions, but with local population concentrations, for 50 robots. Note the efficient behavior of performing the lattice arranging function during the journey towards the final formation.

RESULTS

The local clustering method, for combining multiple robot behaviors, allows groups of robots to form equi-spaced lattice arrangements far more quickly than with single behaviors. The following results compare the performance of pure atomic attraction/repulsion behavior (a relatively fast inverse power law attraction rule) against a combination of the atomic behavior and spring-like behavior, achieved using the local clustering method. Both behaviors result in the same end condition of an equi-spaced lattice distribution. In all cases, robots are modeled with the limitations of maximum attainable speeds and accelerations. If the control rules demand greater speeds or accelerations than the maximum capabilities of the robot, then these are hard-limited at the maximum attainable values.

The most efficient control rules (in terms of speed of lattice formation) will exploit the capabilities of the robot by demanding maximum speeds during as much of the motion as possible. Figure 12 explores the variation of convergence time with variation in the top speed capabilities of the robots. Every experiment was performed with the same number of robots, starting from the same initial random positions. As the speed capabilities of the robots increase, the improved efficiency of the clustering method becomes more dramatic, since the spring component of the combined behavior allows better use to be made of each robot's maximum speed.

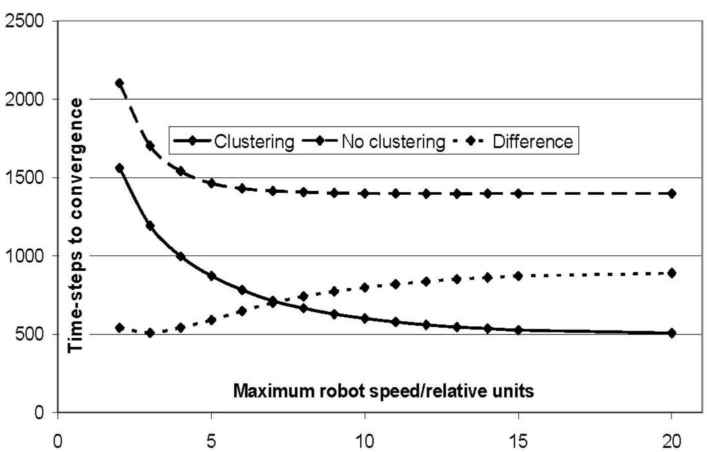

Figure 12. Decrease in convergence time with increase in robot maximum speed capability.

A disadvantage of potential field, artificial gravitation or atomic attraction/repulsion models is that, although they enable robots to form regular lattices, they are extremely slow, particularly in scenarios with large initial inter-robot distances. Figure 13 explores the variation of convergence time with variation in the distance that has to be traveled by the robots in order to meet and form a lattice. In each experiment, the same number of robots was used, with robots randomly distributed about a specified mean initial range from the center of the final formation.

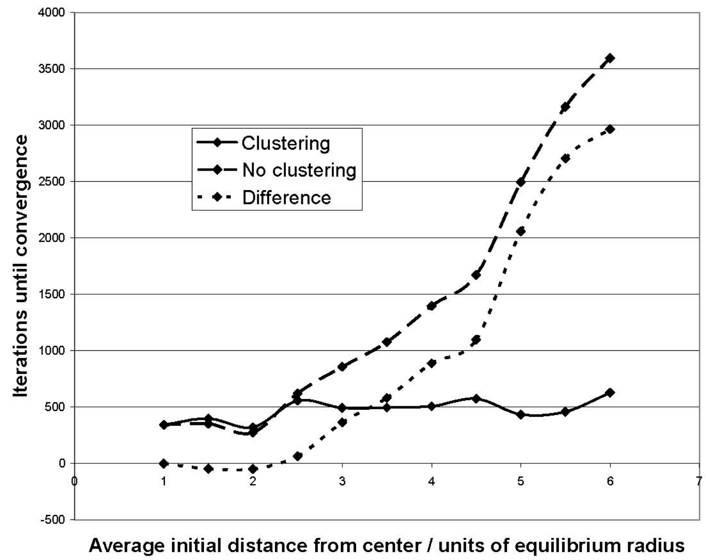

Figure 13. Variations of convergence time with average initial robot distance from final distribution center.

The improved speed of the local clustering method, as compared with an individual inverse square attraction law method (in this case the atomic model), becomes increasingly more dramatic with increased inter-robot range. At long ranges, we expect the convergence time for the atomic model to increase exponentially with distance, whereas we expect only a linear increase with distance for the local clustering method. With greater values of α (equation 5), the performance of the individual behavior will be even worse.

It is not obvious how to independently investigate the variation of performance with varying numbers of robots, since it is not possible to change the number of robots without also changing the initial distribution of robot positions. Figure 14 depicts an experiment which attempts to vary the overall number of robots, and thus density of initial distribution, independent of average robot journey length. For an initial distribution, robots are spaced evenly around a circular locus of specified radius. The robots are also perturbed radially by small random amounts. In all experiments, the initial circular locus radius is kept constant, but different numbers of robots are introduced, thus increasing the density of robots without varying the average robot journey length.

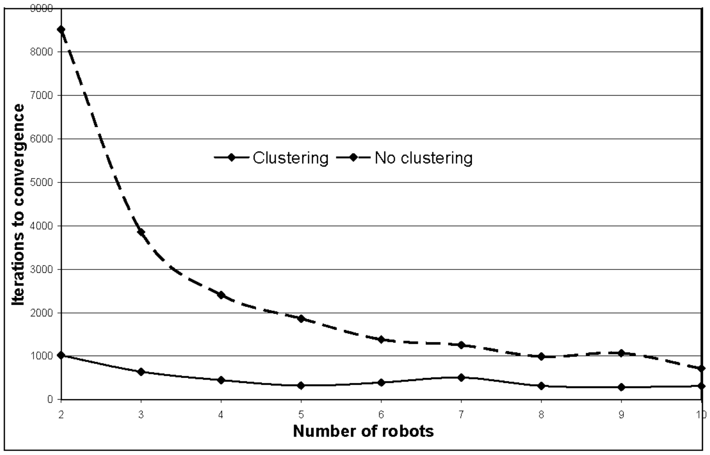

Figure 14. Variation in convergence time with increased density of initial robot positions.

Since the atomic attraction rule suffers less from slowness at close inter-robot ranges, its performance improves significantly with increased robot density. At high initial densities, the performance of pure atomic behavior will tend towards that of the combined behavior, local clustering method. Thus the superior performance of the local clustering method is most significant for scenarios in which robots are initially distributed over a wide area.

CONCLUSION

An effective method, for forming groups of robots into equi-spaced lattices, is for each robot to mimic the behavior of atomic attraction and repulsion. Unfortunately this results in very slow convergence for distant robots. A faster approach involves a spring-like attraction model, but this fails to produce a lattice of constant inter-robot spacings.

We have described a new method which enables the combination of both kinds of behavior, enabling equi-spaced lattice formation but with high accelerations and rapid convergence of distant robots.

The improvements, due to the combined behavior approach, are most significant when robots begin their motion with large inter-robot distances. The combined behavior also improves faster than individual behaviors in response to increased robot top speed capability. The improvements are less significant when the initial distribution of robots is highly concentrated.

The local clustering concept can also be applied more generally. It is a very flexible system which enables the simultaneous use of many different kinds of behaviors. It is particularly useful when different behaviors are desirable at different ranges.

The concept readily extends to larger numbers of behaviors than two. For example, clusters of clusters could be used in order to utilize three different behaviors at three different ranges.

Future work may explore more detailed and realistic models of both robot sensory capabilities and inter-robot communications capabilities. Both uncertainty in sensory measurements and intermittency in communications can be modeled, as well as the variation of these capabilities with range. Control rules should be modified to be robust against such difficulties. We also plan to explore control rules for different kinds of robot arrangements and formations. More specifically, we have shown that different control rules can produce patterns that may be advantageous according to the situation. By having such a repertoire of rules, and by reasoning about the situation, we might be better able to deploy large numbers of robots in emergency situations.

ACKNOWLEDGEMENT

This research was supported in part by the Office of Naval Research, Swampworks, grant #N00014-04-1-0244.

REFERENCES

T. Balch, R. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*. Vol. 14 (6). 1998.

D. Dudenhoeffer, M. Jones. A formation behavior for large-scale micro-robot force deployment. *Proceedings of the 2000 Winter Simulation Conference*. 2000.

V. Gazi, K. Passino. A class of attraction/repulsion functions for stable swarm aggregations. *Proceedings of the IEEE Conference on Decision and Control*. Vol. 3, pages 2842-2847. 2002.

M. M. Kokar, Situation Awareness, Issues and Challenges, *Proceedings of the Seventh International Conference on Information Fusion*, pages 533-534, 2004.

O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *Proceedings of the IEEE Conference on Robotics and Automations*. 1985.

C. Maloney, C. Shays. 9/11 *Can You Hear Me Now Act*, Radio Bill Q&A. http://www.house.gov/maloney/-issues/Homeland/051604RadioBillQandA.pdf

E. Martinson, D. Payton. Lattice formation in mobile autonomous sensor arrays. *SAB'04 Workshop on swarm robotic*s. 2004.

J. Reif, H. Wang, "Social potential fields: A distributed behavioral control for autonomous robots," *Robotics and Autonomous Systems*, vol. 27, pages 171–194. 1998.

W. Spears, D. Spears. An overview of physicomimetics for swarm robotics. *SAB'04 Workshop on Swarm Robotics*. 2004.

W. Spears, D. Gordon. Using artificial physics to control agents. International conference on information intelligence and systems. *IEEE International Conference on Information, Intelligence, and Systems*. 1999.

D. Vail, M. Veloso. Dynamic multi-robot coordination. In *Multi-Robot Systems: From Swarms to Intelligent Automata*, Volume II, A. Schultz, L. Parker, and F. Schneider (eds). pages 87-100, Kluwer Academic Publishers, 2003.