Artificial Physics, Swarm Engineering, and the Hamiltonian Method

Sanza Kazadi*, John R. Lee†, Julie Lee‡

Abstract— **This paper describes the application of a swarm engineering methodology known as the *Hamiltonian method of swarm design* to the artificial physics problem. We demonstrate how to use this methodology to create swarms of predefined global properties by applying it to the basic artificial physics problem which creates locally hexagonal grids of agents, but fails to generate global hexagons. A condition for global hexagonal structure is derived, and two methods are described which accomplish this goal. Neither method requires global information.**

Keywords: ***swarm engineering, artificial physics, Hamiltonian Method of Swarm Design***

1 Introduction

Since the mid 1980's, scientists have been investigating swarms of agents [1, 3, 5, 6]. Swarms, or bidirectionally communicating groups of agents, have been of interest for a number of reasons . People have studied *emergence*, or the ability of the agents to affect a property of their system despite being unaware of that property, extensively. This is thought to be one of the most intriguing qualities of swarms, since the effect tends to occur without careful planning or analysis, but can be fascinating and quite useful. Some swarms have other advantageous properties including fault tolerance, the ability to concentrate resources dynamically, and the ability to change the group behavior despite being unaware of that behavior.

One of the approaches to generating these emergent behaviors is *artificial physics* [1, 7, 8, 9]. Artificial physics (AP) is a control mechanism for a group of agents introduced by Spears [8]. The basic idea behind the methodology is that simple force laws "borrowed" from physics can be used to control the movement of simple agents in a decentralized group. Since physical force laws are simple vector equations of distances and "charge", the implementation of such laws is remarkably simple, as is the gathering of the information required to implement these laws on real, physical platforms. Agents need only be able to determine the distance and direction from any of the other agents, and the appropriate behavior can be calculated as though the agent itself were a "particle" reacting according to these physical force laws. Self-repair, self-organization, fault-tolerance are expected to result from the application of this framework. Moreover, small changes to a basic AP system can result in complex global behavior using only simple controls and primitive sensors. Since the method is not built with a specific group size in mind, it can be applied to groups of any size without change to any control algorithms.

The main problem with this control mechanism is that there is an apparent disconnect between the global variables governing the desired global outcome and the local behaviors of the agents. In any complex system, the goal of the system is to generate a predetermined global outcome; this requirement forms the starting point for any complex system engineering. However, it has been shown elsewhere that systems built from very simple subsystems can generate unpredictable global outcomes with wild changes in global variable as a result of small changes to the interactions between the individual pieces [12]. The difficulties in generating pre-specified global properties in swarms of agents governed by the machinery of artificial physics derives from just this property of complex systems: it is relatively simple to design the local structures, but difficult to generate global structures from these local structures. Without some machinery built in which ensures the global properties are designed properly, it is unlikely that any local design will correctly generate the desired global properties the first time it is created.

Some researchers have tackled the problem of generating global structures[11, 8, 4, 10]. Winfield et al. seem to focus on generating an understanding of what some of the emergent properties of a constructed system will be. However, it is not clear that this generalizes to generating behaviors once the desired global behavior has been decided on. While Spears' intent was to use completely autonomous agents, it was necessary for him to use an ordering to generate the global structure. This ordering, however, might be viewed as global information.

In this paper, we apply a methodology termed a *middle-meeting methodology*[5] to an artificial physics system consisting of simple agents that locally arrange themselves into a hexagonal lattice. The middle meeting methodology has aspects of top-down design as well as aspects of bottom-up design and requires both to be applied in order to generate a global design. While this method is not optimal in the sense that there is a readily applicable algorithm for generating local behaviors once

*Jisan Research Institute, skazadi@jisan.org, To whom correspondences should be sent.
†Jisan Research Institute, jrlee@jisan.org
‡Jisan Research Institute, julie@jisan.org

the global goal has been decided upon, it is an important first step in generating provable global properties. Our theoretical results are applied to a swarm of simulated agents. We observe that the simulated agents perform the desired task of arranging themselves in a large hexagon, a result previously unattainable using the basic artificial physics system.

The remainder of the paper is organized as follows. The middle meeting methodology will be applied to an AP problem in section 2. This application will generate local behaviors of the individual agents. In section 3, we employ these behaviors on a simulated swarm of two-dimensional autonomous agents. Section 4 concludes the paper with a discussion of the relevant issues.

2 Swarm engineering artificial physics systems

In this section, we use the swarm engineering approach to explore the hexagonal configuration problem. This problem centers around the correct placement of agents in a hexagonal grid. In order to create a hexagonal grid, each agent must take its place in a larger hexagonally shaped structure. The task has been accomplished previously [7] by Spears and colleagues. However, to date, no method has yet emerged that allows the global task to be accomplised without explicit global information.

This section begins with an investigation of the way in which microscopic hexagonal lattices may be built up from an originally randomly organized group of agents. These microscopic hexagonal lattices must then be combined in such a way as to produce a macroscopic hexagonal lattice. We assume that the agents are capable of local sensing and carrying out moderate computation. The theory behind the system demonstrates that the final solution is the desired macro-hexagonal state.

Application of this theory in real, physical systems should be direct. In order to make that happen, the agents should be restricted to actions, computations, etc. that could realistically occur on a mobile platform and with real local sensors. We restrict ourselves to sensors and actions that are realistic in terms of the actuators and sensors commonly available or likely to be available in the near future. Where reasonable, if precedents exist in the literature, we will omit the details of some group behaviors that we can reasonably assume can be done in tandem with the actions we are developing here.

In accordance with the middle-meeting methodology described in [5], we begin by first creating a global property using local agent-level properties. Once this has been built so that the desired global structure occurs when a specific unique value of the global property occurs, we continue by generating local behaviors that will yield the desired global property.

2.1 Micro-theory

The first step is to list the various sensory capabilities of the agents. We assume that each agent is capable of measuring the distance from itself to any of its neighbors. That is, each agent i has access to a measurement we label as d_{ij} which indicates the distance between agents i and j. Note that each d_{ij} can be calculated as

$$d_{ij} = |\overrightarrow{x_{ij}}| = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}. \qquad (1)$$

It is interesting to note that in the following derivations, no other individual measurement is required to generate the micro and macro behaviors of the swarm. As a result, the agents need have no other sensory capabilities. We also note that the distances are two-dimensional distances, indicating the tacit assumption that our agents are constrained to a two-dimensional surface. Such a limitation might be reasonable for a swarm constrained to act on the ground, or airborne swarms limited to a particular altitude by design or behavior.

In order to generate the hexagon, agents that are too close to one another must move away from each other, and those that are far away from one another must move closer to one another. As with the original artificial physics studies, we adopt the convention that the actions of the agents will be calculated and then implemented *as though external forces are at work on the agents.* I.e., we assume that each agent senses the distances to each of the other agents within sensory range. Then, using this data, the agent calculates what the next move will be. The implementation is carried out by activating the agent's actuators, which simulate the effect one might expect if the motion came from a real physical system.

Using this micro-measurement, we now proceed to construct a global property. The global property is a property which has a well-defined and unique value when the desired global structure is achieved. Moreover, that value is not capable of being achieved in any other configuration. Generating that value, then, is a matter of undertaking behaviors that produce this value. The need for local properties derives from the requirement that the value be attainable using behaviors that are achievable by the individual agents.

We begin by defining a function known as the *force function, G.* This function determines the virtual force cased by the proximity of two agents. Too close, and the force will be strongly repulsive; too far, and it will be attractive[9]. G is defined as

$$G(t) = \left\{ \begin{array}{cc} -G_1 & t > D_0 \\ G_2 & t < D_0 \end{array} \right\}. \qquad (2)$$

When many agents are involved, the effEct of all the agents is cumulative. A single agent behaves as though

it is reacting to the net force

$$\overrightarrow{F_j} = \sum_{i \neq j}^{N} \left(\frac{\overrightarrow{x_i} - \overrightarrow{x_j}}{|\overrightarrow{x_i} - \overrightarrow{x_j}|} \right) \left(G\left(|\overrightarrow{x_i} - \overrightarrow{x_j}| \right) \right). \quad (3)$$

The force equation defines the way in which the individual agent will behave at the microscopic level. However, we don't yet know whether or not this will lead to the desired global state – that of a single macroscopic hexagonal structure. We use force to derive the global property *energy*, as

$$\begin{array}{l} E = \sum_{i \neq j}^{N} \int_{0}^{(D_0 - D_{ij})} \overrightarrow{F_{ij}}\left(\overrightarrow{x} \right) \cdot d\overrightarrow{x} \\ = \sum_{i \neq j}^{N} \left(G\left(|\overrightarrow{x_i} - \overrightarrow{x_j}| \right) \right) \left(D_0 - D_{ij} \right) \end{array} . \quad (4)$$

Energy measures the total energy required to move all individuals from their current positions to their equilibrium positions. It decreases as the system gets more perfect (stable) and increases as the system gets less perfect (unstable). Our goal is to achieve the smallest value of E possible, because E at its lowest value produces the most stable hexagonal lattice possible. We defer the proof of this fact for a later section of the paper.

In the following discussions, we approximate the force function as

$$G(t) = \lim_{k \to \infty} \left\{ G_1 + (G_2 - G_1) \left(\arctan\left[k\left(t - D_0 \right) \right] + \frac{\pi}{2} \right) \right\}. \quad (5)$$

Note that when $k \to \infty$, the function has the same behavior as that given above in (2).

In order for the system, which is likely to be in a random initial state, to continually evolve to a minimal energy, the time-derivative of the energy must be negative. Therefore, we begin by computing the time derivative and work backward to the condition that each member of the swarm must adhere to to make the global task occur. Accordingly,

$$\frac{dE}{dt} = \sum_{i \neq j}^{N} \left(-G\left(|\overrightarrow{x_i} - \overrightarrow{x_j}| \right) \left(\frac{\triangle x \frac{d \triangle x}{dt} + \triangle y \frac{d \triangle y}{dt}}{D_{ij}} \right) \right)$$

$$+ \sum_{i \neq j}^{N} (\left(\lim_{k \to \infty} \left(\frac{2k\left(D_0 - D_{ij} \right)\left(G_2 - G_1 \right)}{\pi \left(1 + k^2 \left(|\overrightarrow{x_i} - \overrightarrow{x_j}| - D_0 \right)^2 \right)} \right) \right) < 0. \quad (6)$$

We can ignore the second term, as it is always zero except when $|\overrightarrow{x_i} - \overrightarrow{x_j}| = D_0$. This corresponds to a physical situation that will not generally occur and cannot be stable. As a result, we focus exclusively on the first term.

We have noted that the energy is positive definite and that the minimal energy is the desired state. Thus, we are searching for all cases in which the first term of (6) is negative. This indicates that the system energy is decreasing which is what we want it to do.

There are two cases leading to a negative value for the first term. First, if $D_{ij} < D_0$, G will be negative if $\Delta x \frac{\Delta x}{dt} + \Delta y \frac{\Delta y}{dt} > 0$. This will happen if the two agents are moving further away from one another. As a result, two agents that are too close to one another must move away from one another. On the other hand, if $D_{ij} > D_0$ then $\Delta x \frac{\Delta x}{dt} + \Delta y \frac{\Delta y}{dt} < 0$. This means that agents that are too far apart must start coming closer to one another. Combining these two behaviors necessarily produces a behavior that makes the energy in the system non-increasing.

Despite the fact that the energy is non-decreasing, it is still possible for the agents to be in a stable, low energy lattice that is not a perfect hexagon. This is because the energy of the local hexagon is very low and the agents are very stable in their current positions. Although the global energy is smallest when perfect hexagon is achieved, the stable lattices do not change because the energy required to move the misplaced agents out of their equilibrium positions is greater than the immediate difference in the global energy. In order to achieve the perfect hexagon, the lattice needs to be shaken to give the lattice the edge to move the misplaced agents out of their equilibrium. Once those misplaced agents are out of their equilibrium positions, they have to go back to a stable position. They can either go back to the original equilibrium position or they can find a new equilibrium position that is more stable than the original equilibrium position.

This analysis has begun with the creation of a property called energy. This property has the requirement that the specific numerical value we are trying to achieve may only be attained with a unique system configuration. The fact that the system is non-degenerate in this sense means that once the specific value is achieved, the desired system configuration will be achieved.

The first step in generating behaviors has also been achieved in this subsection. In order to generate a hexagonal lattice, we must first generate hexagonal sublattices. This can be accomplished by placing each individual agent at a specific distance from any of its neighbors. We've generated a method of achieving this utilizing an analysis of our energy property. This has generated a very general method which may be applied in a variety of ways to individual agents.

The next two subsections will deal with the question of the minimal energy configuration of the larger lattice and with the methods of moving the entire lattice into an energy configuration which is closer to this global energy configuration.

2.2 Lowest Energy Configuration

In previous sections, we have assumed that the lowest energy configuration of the lattice was a perfect hexagon.

We prove this fact using a simple geometrical argument.

In the lattice, the minimum distances between each pair will always be greater than the equilibrium distance due to the high repulsion coefficient:

$$G(t) = \left\{ \begin{array}{cc} -G & t > D_0 \\ 300G & t < D_0 \end{array} \right\}. \quad (7)$$

Such high repulsion coefficient of G makes it energetically impossible for any agent to get closer than the equilibrium distance other than transiently. Distance error, defined as $|(D_0 - D_{ij})|$, is zero only when the pair are at their equilibrium distance. As a result, the only instances that the distance error from a pair of agents is nonzero are when the pair are further away from one another than the equilibrium distance. Energy of a pair is given by

$$E = (G(x_i - x_j))(D_0 - D_{ij}), \quad (8)$$

Accordingly, the energy is always equal to the sum of all the energy that comes from two agents that are more than the equilibrium distance apart:

$$E_{excess} = \sum_{i \neq j}^{N} G_{Attraction}(D_0 - D_{ij}) \quad (9)$$

The main factors in calculating energy then, are the number of pairs that are more than the equilibrium distance apart and the magnitude of the distance. In hexagonal lattices, all six sides are completely symmetric. This symmetry minimizes the number of pairs of agents whose distances apart are greater than the equilibrium distance. Moreover, any agent moved out of this configuration will increase the sum of the distances between it and all other agents, thereby increasing the overall energy. Therefore, the lowest energy configuration will be the completely symmetric, perfect hexagonal lattice, as any other configuration will have a larger total distance error.

We have proved the following theorem.

Theorem 2.1 *The macroscopic state of the lattice of agents which minimizes the energy is the large hexagonal state.*

2.3 Transitioning between minima

What we've demonstrated in the preceding subsections is that we can create a property and use this property to motivate a class of behaviors, all of which are guaranteed to cause the desired global property. Moreover, we've demonstrated that this system has a unique minimal value to this property, and that minimal value occurs at the system configuration we're after in this system. As a result, minimizing the property generates a system with the desired configuration.

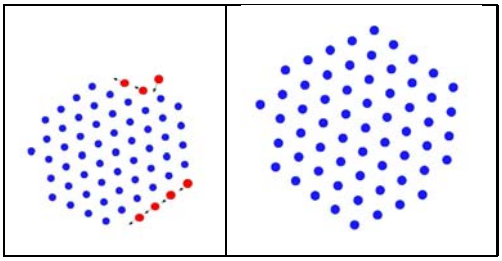

Figure 2.1: This figure illustrates a single translation of the agents in our system during which an outlier inserts itself between two agents, forcing two other agents to move left, and reducing the overall system energy. At the same time, four of the agents at the bottom right of the group translate down in an energy-neutral shift.

Now, we turn to the difficult task of correcting the initial configuration. As depicted in Figure 2.1, the actual configuration of the lattice of agents can be different from the desired configuration when the energy is minimized.

We implemented this theory to move the misplaced agents out of their temporary equilibrium positions and into their true equilibrium positions. The lattice as a whole moves in a hexagonal course, providing shocks at each turn. This gives the agents enough difference in energy to move out and possibly find a equilibrium position with even lower energy. This idea finds its roots in simulated annealing, which finds the incorrect element of the set and changes it, instead of replacing the whole set. With decreasing energy function and implemented movements, the lattice will always achieve a perfect hexagon.

3 Simulation

In this section, we apply two swarm control methods derived from our understanding of the problem outlined in section 2 to a virtual swarm using a two dimensional simulation. We begin by describing the simulation, focusing on the elements of the simulation which should make the control approaches easily useful to real robot platforms. Next we describe, in turn, two different solutions to the control problem given in section 2.

3.1 Description of the simulation

Our simulation is based on the artificial physics simulations of Spears et al. [8]. Briefly, that simulation focused on generating agent control using analogs of radial physical laws such as the law of gravity or electrostatics. Individual agents are represented as points and are capable of moving within the two-dimensional plane that makes up their environment. Each agent is assumed to have sensors capable of determining the distances and directions to its nearest neighbors. These distances and directions

are then used to generate motion vectors. By utilizing a force function similar to that given in section 2, Spears and his collegues were able to have the individual agents arrange themselves in a locally hexagonal lattice.

Our simulation is essentially identical to theirs. All agents are represented as points and are capable of movement within a two-dimensional plane, which is their environment. Agents are assumed to be equipped with directional sensors which provide them with range and bearing to other agents. Each agent is assumed to have actuators which allow them to move, stop on a dime, and turn in place; inertia is not considered to be much of a consideration.

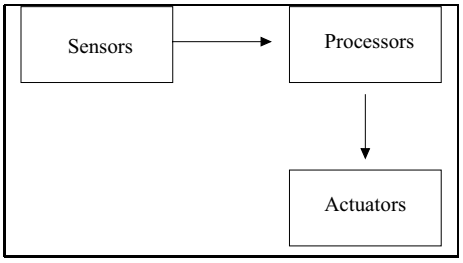

Figure 3.1: This figure illustrates the data pathway of the agents in our system. All agents are reactive, and so data enters through the sensors, is processed directly in the processors, and is translated to actuator commands. No memory is utilized in the control of these agents.

Figure 3.1 illustrates the data pathway of the agents. In our system, all agents are reactive; memory is not required or utilized in control algorithms. Thus, all sensory data is directly processed and translated to actuator commands in the same way that the original artificial physics simulations were. As a result, reaction times can be exceedingly high. However, no special hardware is assumed here, and so moderate timescales easily implemented on real agents may be used for real robotic implementations.

As indicated in section 2, the basic behavior is nearly identical to the original behavior. This quickly changes a randomly organized group of agents to a relatively well ordered group of agents whose energy value is significantly lower than the original energy level. However, it alone is incapable of generating a perfect hexagonal lattice, and so the perturbations described in section 2 must be applied. The basic behavior and energy minimization is illustrated in Figure 3.2.

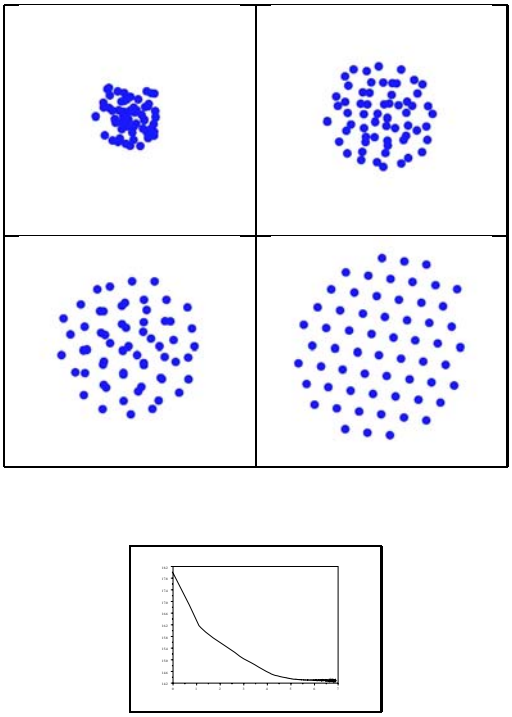

Figure 3.2: Basic artificial physics behavior produces a semi-ordered lattice from a random lattice quickly. As expected, the energy decreases quickly. However, when the energy levels off, the system is caught in a higher energy state than the absolute minimum state.

We already noted that the basic behavior could be augmented by adding behaviors that encouraged energy-requiring transitions between minimal energy states. We shall see two different methods of achieving this in the next two subsections. Note, that in both cases, the requirement of energy increase is a guiding factor.

3.2 Behavioral addendum 1

It is clear from the form of the G function that increasing the energy of the system entails utilizing a behavior that either moves the agents toward one another or moves them away from one another. As a result, we now look at methods of moving agents in such a way that they will increase the system's energy, but also somehow accomplish the translations discussed in section 2.

An important assumption is that the agents have the capability to synchronize their actions. Others have shown [2] that relatively simple behaviors may be utilized to obtain decentralized synchronization of agents. We assume that our agents have this added ability for the first solution.

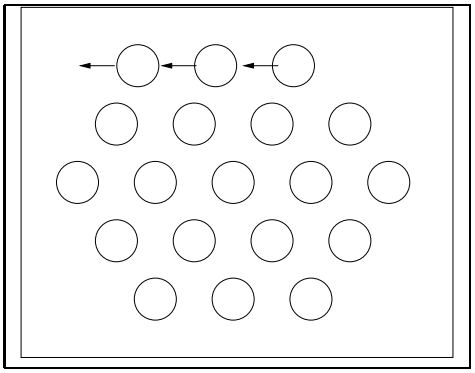

Figure 3.3: Moving a group of agents left, as shown, creates a "drag" on the outer agents who have no surrounding agents to "push" them forward if they are left behind during an imperfectly synchronized turn. As a result, the outer agents' energy levels may be increased enough to shift.

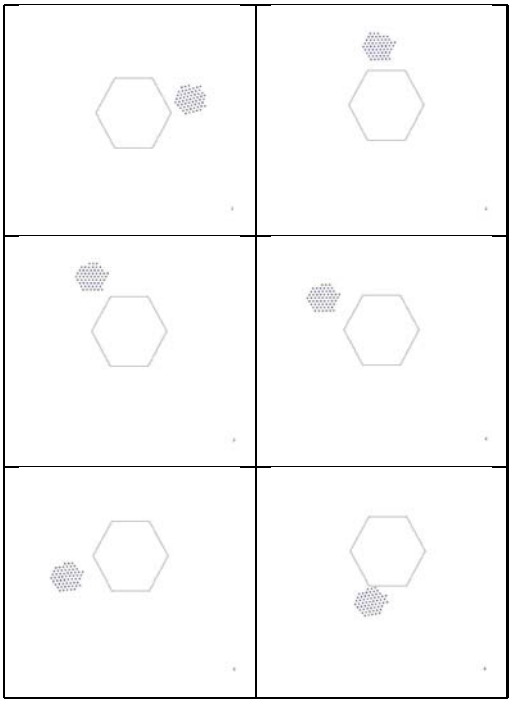

Figure 3.4: This set of images illustrates the "drifting" behavior which accomplishes the energy increase of the group of agents and the shifting of the outermost agents. As can be seen, this results, after several cycles, if necessary, in the completion of the hexagonal structure.

We also note, as depicted in Figure 3.3, that when agents move in tandem and nearly perfect synchrony, the initial movement tends to create temporary movements of outer agents. This increases the outer agents' energies sufficiently to create a shift in agents as the entire group moves in a single direction.

As we've noted, it is possible to generate a hexagon through energy reducing shifts occurring in the group of agents, bringing agents to more perfect positions and lowering the overall energy. However, in order to do this, one must move the group in such a way that the shifts happen along the outer edges of the large hexagon that is to be formed. Thus, the group movements must be aligned with the edges of the large hexagon. As a result, the group of agents must execute movements that are hexagonal, and this will both maintain the interior of the group's organization and effect movements of the outlying agents, since these shifts require a small amount of energy.

We implement this behavior, as illustrated in Figure 3.4. As can be seen, the motions of the group move the group over a rather wide area. During these motions, the outermost agents shift, while the innermost agents stay in formation. At the completion of sometimes several cycles, the agents take on a very stable formation which does not support further shifts. This formation is the hexagonal formation, and is the only one which does not allow shifts.

It is interesting to examine the graph of energy versus time as the system settles into its lowest energy state. As the system shifts into a new state, though momentary variation continues, the base energy tends to step down until it reaches a minimal value. This is illustrated in Figure 3.5.

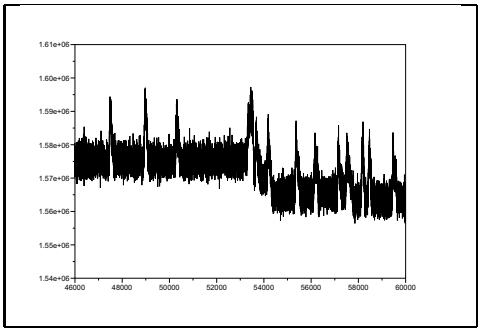

Figure 3.5: The energy of a configuration steps down as the configuration becomes more perfect. The energy configuration completes this step-down pattern when the desired state is reached.

The main drawback to this method is that the entire group must have an added layer of complexity as it executes the drifting behavior in tandem. Moreover, this behavior requires space, and the group may not be deployed until the behavior is complete. We have not explicitly built a behavior into the system which allows the agents to detect when they are in a large hexagon, as this added ability is beyond the scope of this study.

The next subsection examines a behavior that does not require the space of this method.

3.3 Behavioral addendum 2

As we saw in the section 2.2, it is possible to create group behaviors which accomplish the shifting of the agents using behaviors that increase the energy of the system. However, the first method had a few drawbacks that are difficult to envision allowing for deployable groups. In this section, we examine a more direct method of increasing the energy of a system in such a way that it goes into a shifting action but which does not require synchrony among the agents or motions of the agents.

As a result, we now look at a second method of accomplishing the same thing which requires neither the coordination nor the space required by the first method. We start once again by considering what we can do to increase the energy of the system. As before, we realize that the energy can be increased by an individual agent by moving the agent out of equilibrium. A movement away from another agent will increase the energy much less than a movement towards the group. We therefore consider movements of agents towards one another, carefully attenuated so as to produce small jumps required for the sliding action, but not for the rearrangement of the entire array.

The agents to which this behavior will be added must also be carefully controlled, as agents in the interior will generate larger energy changes with small movements than those on the exterior. Therefore the behavior requires the agents to consider their surroundings to qualify them for the random triggering. Thus, the behavior's trigger is limited to those agents whose neighbors fall below 3, which happens only to agents in external positions in the lattice, or to those at corners as a result of the agents' jostling around.

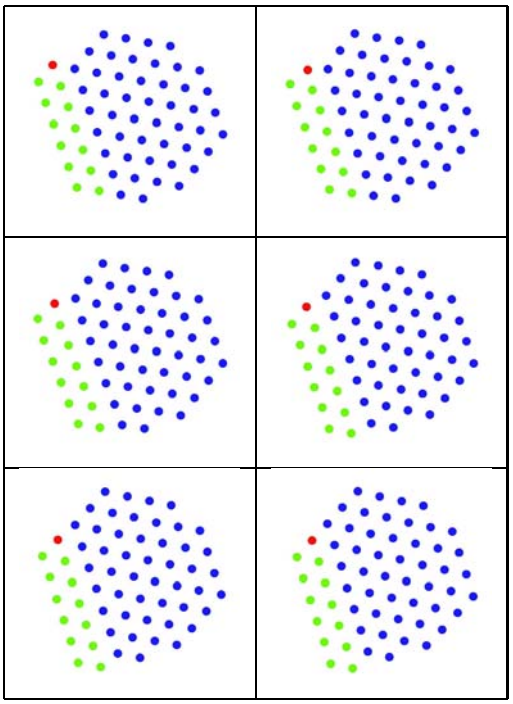

Figure 3.6: The approach of a single agent can cause a set of other agents to slide into another position, with little change in the overall energy of the group.

Figure 3.6 illustrates this behavior of an agent outside of a group of other agents. The exterior agent momentarily approaches the group. The approach increases the energy of the lattice, and the energy reduction behavior moves the exterior agents in a translation motion. These actions can be initiated by single agents even though the entire group is in motion or static.

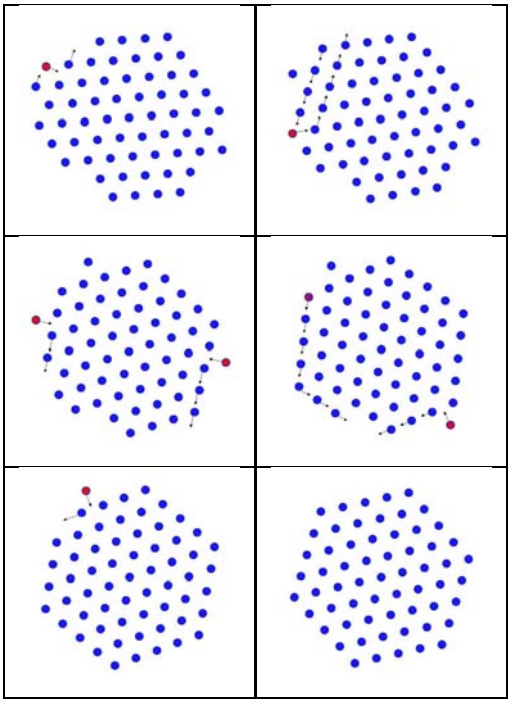

Figure 3.7: This set of images illustrates the "closing" behavior. The behavior initiates a number of different sliding behaviors that eventually causes the generation of a hexagonal structure. As before, the energy inches down to a minimum under the action of the individual agent perturbations.

Figure 3.7 illustrates the behavior of a group of agents working under this behavior. The group begins with a state very far from the desired hexagonal state. However, as the translation movements continue, the system falls into a lower energy state, which eventually is that of the hexagonal organization. In all simulated runs[1], the system will settle into the hexagonal state eventually.

4 Summary and Conclusion

In this paper, we examined the swarm control problem in the context of the hexagonal lattice construction as a result of the basic artificial physics system. The original artificial physics system was capable of producing imperfect hexagonal lattices with little or no symmetry. Our purpose, then, was to enhance this basic behavior so as to produce perfect symmetric hexagonal lattices.

Our method consisted of generating a *property* which is a measurable constructed from other properties (or measurables) that the agents could realistically be expected to be able to measure. In our case, the only thing we needed was the distance and bearing to construct a property that had the proper requirements:

[1]The simulation was run more than 10000 times.

1. The property was constructed from existing properties that the individual agents could measure.
2. The property has a single well-defined value at the desired system configuration that cannot be obtained in any other way.

Using these requirements, we were able to generate simple behaviors that moved the system from any initial state to a state more like the desired state. Not only were we able to generate one method, but we were able to generate two methods, which were modified from the basic method that itself was identical to Spears' original algorithm.

The power of this method lies in the ability of the engineer to create one or more properties whose numerical values are unique to the state that the system is in. The engineer, then, needs only chart a path through the allowed phase space of the system to the final desired value, hopefully utilizing behaviors which individual agents can accomplish on their own, with or without guidance from a central controller. The method can be applied to single properties or to vectors of properties, provided that the desired vector is well-defined in the same way a single property might be. We believe that the method is so powerful, in fact, that we now coin a term for this method: ***The Hamiltonian Method of Swarm Design***.

In the future, we intend to apply this method to swarms of greater complexity than this one. We expect that this method of not only swarm design, but complex system design, may be applied to a large number of different systems including, but not limited to, systems of autonomous mechanical agents, computing systems, economic systems, and social systems. While some of this research is currently under way, we expect that the exploration of all fields to which this methodology might be applied will reveal an extraordinarily vast scope. Moreover, we expect that an extension to this work will be able to solve the problem originally posed ten years ago which led us to these results: "Is it possible that the global specification of a problem is enough to yield the basic requirements of the solution including all actuators, sensors, processing, and other capabilities of agents in the solution?" We believe the answer is yes.

References

[1] Hettiarachchi S., and W. Spears. *DAEDALUS for Agents with Obstructed Perception*. **Proceedings of the 2006 IEEE Mountain Workshop on Adaptive and Learning Systems**.

[2] Holland, O, Melhuish, C., Hoddell, S. *Chorusing and controlled clustering for minimal mobile agents*. **ECAL97, Brighton, 1997**.

[3] Jones, C., Mataric, M. *Adaptive division of labor in large-scale multi-robot systems*. In **Proceedings of**

the IEEE International Conference on Intelligent Robots and Systems (IROS'03), Las Vegas, NV, pages 1969-1974, Oct 2003.

[4] Jones, C., Mataric, M.*From local to global behavior in intelligent self-assembly.* In **Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'03), Taipei, Taiwan, Sep 2003,** pages 721-726, 2003.

[5] Kazadi, S. *On the Development of a Swarm Engineering Methodology.* **Proceedings of IEEE Conference on Systems, Man, and Cybernetics**, Waikoloa, Hawaii, USA, pp.1423-1428, October 2005.

[6] Reynolds, Craig. Boids *Flocks, Herds, and Schools: a Distributed Behavioral Model*, **¡http://www.red3d.com/cwr/boids/¿** Dec. 2006. (Accessed 08 Jan. 2007).

[7] Diana F., William M. Spears, Oleg Sokolsky, and Insup Lee (1999). *Distributed Spatial Control, Global Monitoring and Steering of Mobile Physical Agents.* **IEEE International Conference on Information,Intelligence, and Systems**, November, 1999.

[8] Spears, William M. and Diana F. Gordon (1999). *Using Artificial Physics to Control Agents.***IEEE International Conference on Information, Intelligence, and Systems**, November, 1999.

[9] Spears, W., D. Spears, J. Hamann, and R. Heil. *Distributed, Physics-Based Control of Swarms of Vehicles.* **Autonomous Robots**, Volume 17(2-3), August 2004.

[10] Wiegand, P., M. Potter, D. Sofge, and W. Spears. *A Generalized Graph-Based Method for Engineering Swarm Solutions to Multiagent Problems.* **To appear in PPSN 2006**.

[11] Winfield, A., Sa, J., Gago, M.C., Dixon, C., Fisher, M. *On Formal Specification of Emergent Behaviours in Swarm Robotics Systems.* **Int. J. Advanced Robotic Systems**, 2(4),363-370, 2005

[12] Wolfram, S. A **New Kind of Science**. Champagne, Il: Wolfram Media, Inc., May 2002.