

Analysis of a Phase Transition in a Physics-Based Multiagent System

Diana F. Gordon-Spears and William M. Spears

Computer Science Department
College of Engineering
University of Wyoming
Laramie, WY 82071
{dspears, wspears}@cs.uwyo.edu

In *Proceedings of the FAABS'02 Workshop*

Abstract. This paper uses physics-based techniques to analyze a physics-based multiagent system. Both qualitative and quantitative analyses are provided to better understand and predict a system phase transition. These analyses yield deep insights into the system behavior. Furthermore, they have been tested in a practical context on actual robots and proven to be quite effective for setting system parameters.

1 Motivation

The objective of this research is to design rapidly deployable, scalable, adaptive, inexpensive, and robust networks (swarms) of autonomous distributed sensing agents (e.g., robots). The agents are to monitor, detect, track, report, and respond to environmental conditions within a specified physical region. This is done in a distributed manner by deploying numerous physical mobile robotic agents (each carrying one or more sensors), and then aggregating and fusing the distributed data collected from the agents into a single, coherent tactical picture. Potential applications for multiagent networks include search and rescue, virtual space telescopes consisting of micro-satellites, surveillance and perimeter defense, and distributed sensors for mapping biological hazards.

One of the most significant challenges when dealing with swarms of mobile agents is how to design and analyze the desirable collective behavior. It is assumed that there is no global controller for the multiagent system. The desired aggregate behavior must emerge from purely local interactions, despite restricted situational awareness and a limited communication range. In other words, the agents need to self-assemble (self-organize). Not only do we want desirable global behavior to emerge from local interactions between agents, but we also would like there to be some measure of fault-tolerance i.e., the global behavior degrades gradually if individual agents are damaged. Self-repair is also desirable, where the system repairs itself after being damaged. Finally, *formal (mathematical) behavioral assurances, i.e., predictions of how and when a system will behave in a certain manner, are crucial. Such assurances are the focus of this paper.*

2 Relation to Alternative Approaches

System analysis enables behavioral assurance. Here, we adopt a physics-based approach to analysis. We consider this approach to fit under the category of “formal methods,” not in the traditional sense of the term but rather in the broader sense, i.e., a formal method is a mathematical technique for designing and/or analyzing a system. The two main *traditional* formal methods used for this purpose are theorem proving and model checking. Why do we use a physics-based method instead of these more traditional methods? The gist of theorem proving (model checking) is to begin with a theorem (property) and prove (show) that it holds for the target system. But what if you don’t know how to express the theorem or property in the first place? For example, suppose you visually observe a system behavior that you want to control, but you have no idea what causes it or how to express your property in concrete, logic-based or system-based terms? In particular, there may be a property/law relating various system parameters that enables you to predict or control the observed phenomenon, but you do not understand the system well enough to write down this law.

For such a situation, the traditional, logic-based formal methods are not directly applicable. One potentially applicable approach is empirical. Certainly machine discovery [5] and a vast array of statistical techniques could be applied to solve this problem of deriving laws empirically. We have instead chosen a theoretical (formal) physics-based approach for the following reasons:

- Empirical techniques can tell you *what* happens, but not *why* it happens. By gaining a deep, scientific understanding of a system, it is possible to generalize and extend analogies to a wide variety of aspects of the system in a way that is inconceivable to do with empirically derived laws. We have demonstrated this advantage by developing a deep, highly predictive, and practical framework using a physics potential energy analysis [10] in addition to the phase transition analysis presented here. The practicality of our framework has been validated on a set of actual robots. Minor modifications to the theory are trivial to implement once a deep understanding of the physics has been achieved.
- Other scientists (in other fields) tend to find theoretical results with a solid scientific rationale to be easier to understand and to apply than empirically derived results. Furthermore, a scientifically sound theory is easier for other scientists to build upon.
- Empirical methods can be combinatorially explosive.
- If a physics-based analysis technique is predictive of a system built on physics-based principles, then this analysis provides formal verification of the correctness of the system implementation. No such claims can be made for empirical results.

Note that we will not necessarily abandon empirical methods altogether. Empirical methods form a nice backup approach to apply when physics-based approaches become too difficult. So far, we have made a good deal of progress

in analyzing our multiagent system using physics alone, and have not yet seen the need for an alternative. However, we may see such a need in the future.

3 The Physicomimetics Framework

We have created a framework for the design and analysis of multiagent systems that is grounded in natural laws of physics [9]. We call this framework “physicomimetics” or “artificial physics” (AP). In the AP framework, artificial (virtual) forces control agents. We use the term “artificial” because although we will be motivated by natural physical forces, we are not restricted to only natural physical forces. Although AP forces are virtual, agents *act* as if they are real. Thus the agent’s sensors must see enough to allow it to compute the forces to which it is reacting. The agent’s effectors must allow it to respond to this perceived force.

There are potential advantages to this approach. First, in the real physical world, collections of small entities yield surprisingly complex behavior from very simple interactions between the entities. Thus there is a precedent for believing that complex behavior can be achieved through simple local interactions. This is required for very small agents (such as nanobots), since their sensors and effectors will necessarily be primitive. Two, since the approach is largely independent of the size and number of agents, the results should scale well to larger agents and larger sets of agents. Finally, because this approach is physics-based, it is amenable to well-understood, physics-based analysis techniques.

In the AP framework, agents are treated as physical particles, which could range in size from nanobots to satellites. Particles move in response to the virtual forces that are exerted upon them by their neighbors – in essence the particles act as if they were part of a molecular dynamics simulation. Particles have a position, mass, velocity, and momentum. Friction is included, for self-stabilization. The net action of the system of particles is to reduce potential energy in a continuously changing virtual potential field.

Each particle i has position $\mathbf{p} = (x, y)$ and velocity $\mathbf{v} = (v_x, v_y)$. We use a discrete-time approximation to the continuous behavior of the particles, with time-step Δt . At each time step, the position of each particle undergoes a perturbation $\Delta \mathbf{p}$. The perturbation depends on the current velocity $\Delta \mathbf{p} = \mathbf{v} \Delta t$. The velocity of each particle at each time step also changes by $\Delta \mathbf{v}$. The change in velocity is controlled by the force on the particle $\Delta \mathbf{v} = \mathbf{F} \Delta t / m$, where m is the mass of that particle and \mathbf{F} is the force on that particle. Our simulation is written in Java and is Web-accessible at <http://www.cs.uwyo.edu/~wspears/ap.html>. For a detailed description of how AP relates to prior research, see [9].

Given a problem, by defining it within the AP framework, we map the problem to one of minimizing potential energy (PE). As constraints are violated or performance degrades, PE increases, thus triggering a reactive response. *However, the PE field is never computed by the agents – it is an intellectual framework. The particles respond only to forces. Global patterns are automatically self-assembled via local force interactions between agents.* Given a set of initial

conditions and some desired global behavior, it is necessary to define what sensors, effectors, and force laws are required for the desired global behavior to emerge. We call this the “design problem.”

4 Hexagonal Lattice Sensing Grids

Let us consider an example of design. In this example, AP is applied to a swarm of micro-air vehicles (MAVs) whose mission is to form a hexagonal lattice, which acts as a distributed sensing grid [2]. Such a lattice creates a virtual antenna or synthetic aperture radar to improve the resolution of radar images.

Since MAVs have simple sensors and primitive CPUs, our goal is to provide the simplest possible control rules that require minimal sensors and effectors. At first blush, creating hexagons would appear to be somewhat complicated, requiring sensors that can calculate range, the number of neighbors, their angles, etc. However, it turns out that only range and bearing information are required. To understand this, recall an old high-school geometry lesson in which six circles of radius R can be drawn on the perimeter of a central circle of radius R (the fact that this can be done with only a compass and straight-edge can be proven with Galois theory). Figure 1 illustrates this construction. If the particles (shown as small circular spots) are deposited at the intersections of the circles, they form a hexagon with a particle in the middle.




Fig. 1. How circles can create hexagons.

The construction indicates that hexagons can be created via overlapping circles of radius R . To map this into a force law, imagine that each particle repels other particles that are closer than R , while attracting particles that are further than R in distance. Thus each particle can be considered to have a circular “potential well” around itself at radius R – neighboring particles will want to be at distance R from each other. The intersection of these potential wells is a form of constructive interference that creates “nodes” of very low potential energy where the particles will be likely to reside (again these are the small circular




Fig. 2. The force law, when $R = 50$, $G = 1200$, $p = 2$ and $F_{max} = 1$.

spots in the figure). Thus the particles serve to create the very potential energy surface they are responding to!¹

With this in mind, we define a force law $F = Gm_i m_j / r^p$, where F is the magnitude of the force between two particles i and j , r is the range between the two particles, and p is some power.² The “gravitational constant” G is set at initialization. The force is repulsive if $r < R$ and attractive if $r > R$. Each particle has a sensor that detects the range and bearing to nearby particles. The only effector is to be able to move with velocity \mathbf{v} . To ensure that the force laws are local in nature, particles have a visual range of only $1.5R$. Also, due to the discrete-time nature of the model, it is important to define a maximum force F_{max} that can be obtained.

Figure 2 shows the magnitude of the force, when $R = 50$, $G = 1200$, $p = 2$, and $F_{max} = 1$ (the system defaults). There are three discontinuities in the force law. The first occurs where the force law transitions from F_{max} to $F = Gm_i m_j / r^p$. The second occurs when the force law switches from repulsive to attractive at R . The third occurs at $1.5R (= 75)$, when the force goes to 0.

Figure 3 shows how an initial universe of $N = 200$ particles that began in a single small, random cluster has evolved over 1000 time steps into a hexagonal lattice, using this very simple force law. For a radius $R = 50$ and $p = 2$ we have found that a gravitational constant of $G = 1200$ provides good results.

Note that in Figure 3 we observe a clustering effect, i.e., each node in the lattice may contain multiple particles. Clustering was an emergent property that we had not expected, and it provides increased robust behavior, because the disappearance (failure) of individual particles (agents) from a cluster will have minimal effect. This form of fault-tolerance is a result of the setting of G , which we explore in Section 6. The pattern of particles shown in Figure 3 is quite stable, and does not change to any significant degree as t increases past 1000.

¹ Again, potential energy is never actually computed by agents. Agents only compute local force vectors for their current location. PE is computed for visualization/analysis purposes only.

² The default setting has $m_i = 1.0$ for all particles.




Fig. 3. A hexagonal lattice ($t = 1000$).

5 Square Lattices

Given the success in creating hexagonal lattices, we were inspired to investigate square lattices also. The success of the hexagonal lattice hinges upon the fact that nearest neighbors are R in distance. Clearly this is not true for squares, since if the distance between particles along an edge is R , the distance along the diagonal is $\sqrt{2}R$. The problem is that the particles have no way of knowing whether their relationship to neighbors is along an edge or along a diagonal.

Once again we use a very simple approach. Suppose that at creation each particle is given another attribute, called “spin”. Half of the particles are initialized to be spin “up”, whereas the other half are initialized to be spin “down”. Spins do not change during the evolution of the system.³



Fig. 4. Forming a square using two spins.

Consider the square depicted in Figure 4. Particles that are spin up are open circles, while particles that are spin down are filled circles. Particles of unlike spin are distance R from each other, whereas particles of like (same) spin are distance $\sqrt{2}R$ from each other. This “coloring” of the particles extends to square lattices, with alternating spins along the edges of squares, and same spins along the diagonals.

The construction in Figure 4 indicates that square lattices can be created if particles can sense not only range and bearing to neighbors, but also the spins of their neighbors. Thus the sensors need to be able to detect one more bit of information, spin. We use the same force law as before: $F = Gm_i m_j / r^p$. In this case, however, the range r is renormalized to be $r/\sqrt{2}$ if the two particles have the same spin. Then once again the force is repulsive if $r < R$ and attractive if $r > R$. The only effector is to be able to move with velocity v . To ensure that

³ Spin is merely a particle label, like a color.




Fig. 5. The size of clusters as t increases.




Fig. 6. Cluster size drops suddenly as G is decreased linearly after $t = 1000$.

the force laws are local in nature, particles cannot even see or respond to other particles that are further than $1.7R$.⁴ Clustering also occurs in square lattices.

6 Phase Transition

We decided to study the evolution of cluster size as a lattice emerges. In this section, we focus on hexagonal lattices, although the results with square lattices are a straightforward analogy. For each particle i , we count the number of particles that are close to i ($0 < r < 0.2R$). The particle i is always included, so minimum cluster size is 1.0. This is averaged over all particles and displayed for every time step. Results are shown in Figure 5. At $t = 0$, all particles are very close to one another, yielding a high clustering. Immediately, the particles fly apart, due to the repulsive force, so that by $t = 6$ the particles are all effectively separated. However, after $t = 6$ clusters re-emerge, with the final cluster size being around 2.5. Clearly the re-emergence of clusters serves to lower the total potential energy of the system, and the size of the re-emerged clusters depends on factors such as G , R , and the geometry of the system.

Next, we summarize an interesting follow-on experiment. The previous experiment was repeated, evolving the system until $t = 2500$. However, after $t = 1000$

⁴ The constant is 1.3 if particles have like spin and 1.7 otherwise.

we lowered G by 0.5 for every time step. The results are shown in Figure 6. Our expectation was that the average cluster size would decrease linearly with G , but we were surprised. The average cluster size remained quite constant until about $t = 2000$, which is where G is 700. At this point (which is a phase transition), the cluster size dramatically dropped until roughly $t = 2200$ (where $G = 600$), where the particles are separated again.⁵ This appears very similar to a phase transition in natural physics, demonstrating that AP can yield behavior very similar to that demonstrated in natural systems.

7 Qualitative Analysis

Recall that our primary objective is to provide behavioral assurance. Our goal is to model and analyze multiagent system behavior, thus enabling predictions and, if needed, corrections. Complex, evolving multiagent systems are notoriously difficult to predict. It is quite disconcerting when they exhibit anomalous behaviors for which there is no explanation. This is especially a concern when there are abrupt, radical shifts in their performance. Avoidance of abrupt, undesirable shifts in the performance of complex systems is one of the motivations for formally verifying invariance (safety) properties [3]. If a property can be expressed in temporal logic, model checking is applicable. However, the clustering property is spatial, rather than temporal, and we do not know how to express it. Therefore, we use physics-based methods for property derivation.

Consider the observed phase transition just described. It appears similar to real-world phase transitions. To gain initial insights, we began with a *qualitative* modeling and analysis of this transition, using physics techniques. In particular, we have explored mathematical visualizations of the PE fields, which are computed by performing a line integral over the forces. A line (path) integral is a measure of the work that would need to be done by a virtual particle to get to some position in the force field. By definition:

$$V = - \int_s \mathbf{F} \bullet d\mathbf{s}$$

where $\mathbf{s} = x\mathbf{i} + y\mathbf{j}$ is the path. The actual computation is done via the method of finite differences [4].

A line integral may be used to calculate PE if the force is (or is approximately) conservative because in that case the work to get from point a to point b is independent of the path taken (which is an assumption of the line integral). This can be verified by determining the *curl* of the force field [1]. More formally, V , the potential energy field, is a scalar. The gradient of V , i.e., $\nabla V = \mathbf{i} \partial V / \partial x + \mathbf{j} \partial V / \partial y$. ∇ is a vector operator, which in two dimensions is $\mathbf{i} \partial / \partial x + \mathbf{j} \partial / \partial y$. The curl of \mathbf{F} , $\nabla \times \mathbf{F}$, equals $\partial F_y / \partial x - \partial F_x / \partial y$. The interpretation is that the curl is a measure of rotation of the vector force field near a point. If the curl is

⁵ We define the precise time of the phase transition to be when the average cluster size falls below 1.5.

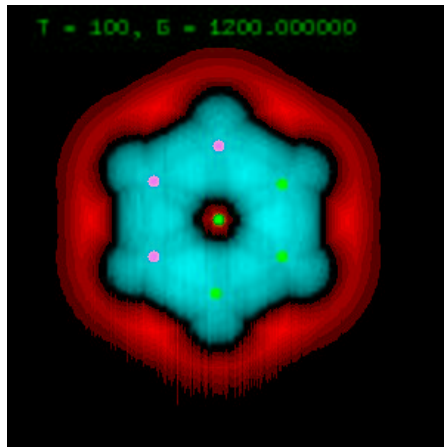


Fig. 7. The PE field when $G = 1200$.




Fig. 9. The PE field when $G = 600$.




Fig. 8. The PE field when $G = 800$.

0 then the force field is conservative and the computation of the potential field is meaningful. As it turns out, due to the radial symmetry of our force law, the curl is 0 everywhere, except for small deviations from 0 at the three force law discontinuities.

Figure 7 illustrates the PE field, for a system of seven particles (there are three on the right that are hard to see) that have stabilized into a hexagonal formation, where $G = 1200$. Light grey represents positive potential energy (the lightest shade represents the maximum), while dark grey represents negative potential energy. Black represents zero potential energy. A virtual particle placed in this field wants to move from regions of high potential energy to low potential energy. For example, consider the central particle, which is surrounded by a region of negative PE, followed by a region of zero PE, and then by a region of positive PE. A virtual particle that is close to the center will tend to move towards that center. Thus the central particle is in a PE well that can attract another particle. This is not surprising, since we showed earlier that a G of 1200 results in clustering.

Now consider lowering G to 800. Figure 8 illustrates the PE field. The central PE well is not as deep as it was previously, but it is still a PE well and can attract another particle. There has not yet been a phase transition. Finally, consider lowering G to 600. Figure 9 illustrates the PE field. Note that the central PE well no longer exists! The central particle is now surrounded by regions of lower PE – thus a particle near the center will tend to move away from that center. A phase transition has occurred and clustering ceases.

8 Quantitative Analysis

We have now set the stage for a more *quantitative* theoretical analysis of the cluster size phase transition. Specifically, we would like to be able to predict *when* the phase transition will occur. In particular, we want to predict the value of G in terms of other parameter settings for which the transition will occur. Before we do this, note that there are, in fact, multiple phase transitions occurring on the microscopic level. In other words, clusters break up in sub-phases, from the largest cluster size to the smallest (of size one). These mini-phase transitions occur at distinct times. Thus, if we want to make a quantitatively precise prediction, then we need to choose which mini-phase transition we wish to predict. The one of interest in this paper is the final mini-phase transition, i.e., the final cluster fragmentation, which occurs when the average cluster size prior to the phase transition is above 1.5, and the average cluster size after the phase transition is 1.0. This final phase transition is qualitatively different from the others, in that the formation changes behavior from a solid to a liquid.

Based on the qualitative analysis (above), we hypothesize that a standard physics *balance of forces* law will be predictive. In the context of our phase transition, this quantitative law states that the phase transition will occur at precisely those parameter values when the *cohesion force*, which keeps a particle within a cluster, exactly equals the *fragmentation force*, which repels the particle

out of the cluster. To quantitatively specify this law, it is necessary to derive the expressions for these forces.

Figure 9 indicates that a particle placed near the central particle can escape along trajectories that avoid the outer perimeter particles. This has been confirmed via observation of the simulation. We depict these most likely escape paths in Figure 10. In this figure, which depicts a canonical scenario, there are two particles in the central cluster, and one particle in all other clusters (and all particles have a mass $m = 1$, although we can easily generalize our analysis). Let us focus on one of the two particles in the center, and call this particle “A.” Due to the symmetry, without loss of generality we can focus on any of the escape paths for particle A. Let us arbitrarily focus on the escape path along the horizontal axis. Particle A can be expelled from its cluster along this axis by the other central particle, which exerts a repulsive force of F_{max} , because the range between particles, r , is very small. Therefore, the fragmentation force upon particle A is equal to F_{max} .




Fig. 10. How particles escape.

Next, we derive an expression for the cohesion force on A. Particle A is held near the center by the outer perimeter clusters/particles. Due to the geometry of the situation, it is simple to show that the cohesion force is $2\sqrt{3}G/R^p$. As above, without loss of generality we focus on the horizontal axis as the escape path of particle A. Consider the force exerted by one of the outer perimeter clusters just above this horizontal axis, on particle A. Because clusters are R apart, the magnitude of this force is G/R^p . The projection of this force on the horizontal axis escape path is $\sqrt{3}/2$ times the magnitude of this force – because the angle between the chosen outer perimeter cluster and the horizontal axis is 30 degrees. Since there are four outer perimeter clusters exerting this force (the remaining two have a force of 0 after projection), we multiply this amount by four to get a total cohesion force of $2\sqrt{3}G/R^p$.

When the cohesion force is greater than the fragmentation force, the central cluster will remain intact. When the fragmentation force is greater, the central cluster will separate. Thus, our law states that the phase transition will occur roughly when the two forces are in balance: $F_{max} = 2\sqrt{3}G/R^p$. We can now state that the phase transition will occur when $G = F_{max}R^p/2\sqrt{3}$. We denote this value of G as G_t . For the system described in the previous section $G_t = 722$,

which is consistent with our visualizations. Our phase transition law is:

$$G_t = \frac{F_{max} R^p}{2\sqrt{3}}$$

We tested this law for varying values of R , F_{max} , and p . The results are shown in Tables 1 and 2, averaged over 10 independent runs, with $N = 200$. The system evolved until equilibrium with a high value of G . Then G was gradually lowered. Cluster size was monitored, and we noted the value of G when the average cluster size dropped below 1.5. The observed values are very close to those that are predicted (within 6%), despite the enormous range in the magnitude of predicted values (approximately four orders). Plus, the variance is low.

Table 1. Predicted/observed values of G_t when $p = 2$.

R	F_{max}		
	0.5	1.0	2.0
25	90/87	180/173	361/342
50	361/355	722/687	1,440/1,430
100	1,440/1,410	2,890/2,840	5,780/5,630

Table 2. Predicted/observed values of G_t when $F_{max} = 1$.

R	p		
	1.5	2.5	3.0
25	36/35	902/874	4,510/4,480
50	102/96	5,100/5,010	36,100/35,700
100	289/277	28,900/28,800	289,000/291,000

These results are quite promising. We have provided a very good predictor of the phase transition G_t , which incorporates the most important system parameters p , R , and F_{max} . It is important to notice that N (the number of particles) does not appear in our equations. This is because the phase transition behavior is largely unaffected by N , which is a nice feature.

As with hexagonal lattices, square lattices also display a similar phase transition as G decreases. The derivation of a predictive, quantitative law for square lattices is a straightforward analogue of the analysis for hexagonal lattices. The one difference between the two analyses is that in a square lattice, one of the two particles in the central cluster is expelled along a path to one of the outer

perimeter particles, rather than between them (see Figure 11). As with hexagonal lattices, square lattices also have mini-phase transitions. Again we focus on the final mini-phase transition, i.e., the final cluster fragmentation, which occurs when the average cluster size prior to the phase transition is above 1.5, and the average cluster size after the phase transition is 1.0.




Fig. 11. How particles escape.

In Figure 11, there are two particles in the cluster at the center of the formation, and one particle in all other clusters. Let us focus on one of the two particles in the center, and call this particle “A.” Using the same reasoning as with hexagons, the fragmentation force upon particle A is equal to F_{max} . Particle A is held near the center by the outer perimeter clusters/particles. Using the geometry of the situation as we did with hexagons, it is simple to show that the total cohesion force on A is $(2\sqrt{2} + 2)G/R^p$. The phase transition will occur when $G = F_{max}R^p/(2\sqrt{2} + 2)$. Our phase transition law for square lattices is:

$$G_t = \frac{F_{max}R^p}{2\sqrt{2} + 2}$$

Table 3. Predicted/observed values of G_t when $p = 2$ for square lattices.

R	F_{max}		
	0.5	1.0	2.0
25	65/69	130/136	259/278
50	259/272	519/530	1,036/1,066
100	1,036/1,112	2,071/2,138	4,143/4,405

We tested this law for varying values of R , F_{max} , and p . The results are shown in Tables 3 and 4, averaged over 10 independent runs, with $N = 200$. The observed values are very close to those that are predicted (within 7%), and the variance among runs is low.

There are several uses for these equations. Not only can we predict the value of G_t at which the phase transition will occur, but we have observed that we can

Table 4. Predicted/observed values of G_t when $F_{max} = 1$ for square lattices.

R	p		
	1.5	2.5	3.0
25	26/26	647/651	3,236/3,312
50	73/74	3,662/3,730	25,891/26,850
100	207/206	20,713/21,375	207,125/211,350

use G_t to help design our systems. For example, running the system at $G \approx 0.9G_t$ will yield the best unclustered formations, while a value of $G \approx 1.8G_t$ will yield the best clustered formations. In addition to empirical testing of our laws, we also tested them on a group of six actual robots. The goal of the robots was to get into geometric lattice formations. Using our laws, we were able to optimize robot behavior with or without clustering, whichever was desired. For details, see [10]. *Finally, note that the predictiveness of the equations is a form of verification of the correctness (in a physics sense) of the AP code!*

9 Summary and Related Work

This paper has presented both qualitative *and* quantitative analyses of a phase transition that occurs in a physics-based multiagent system. The qualitative analysis is based on a formal model of the potential energy and related fields, visualized using numerical analysis techniques. Observations of the field dynamics provided key insights into the physics phenomena precipitating the phase transition. We also present a novel quantitative analysis, that is inspired by traditional physics-style mathematical analyses of balance of forces. This analysis has resulted in *new laws*, which predict precisely (for a wide variety of parameter settings) *when* the phase transition will occur.

The significance of this work is that predictive analyses of swarm behavior are scarce. Multiagent swarms with emergent behavior are notoriously difficult to predict, and only a few researchers have tackled such an endeavor. Our advantage in this respect is that AP is physics-based, and thus we are able to employ traditional physics analysis techniques.

The most related analysis of a swarm under distributed control was presented by Lerman at FAABS'00 [6], and further developed in a subsequent technical report [7]. Lerman applied differential equations to mathematically analyze a multiagent system that uses behavior-based methods for control. Behavior-based multiagent systems assume that a subgroup of the agents will jointly and simultaneously perform a single behavior (such as flocking or homing) for a period of time, then switch to another behavior. Lerman formulated differential equations about state-to-state transitions that can be used to predict the fraction of agents that will be performing a given behavior after a specified time. The similarity to the work presented here is that both analyses predict a behavioral transition.

The difference is that behavior-based systems are heuristic, rather than physics-based. Therefore, the full gamut of deep, knowledge-based physics techniques is not applicable to behavior-based systems. Because these systems are heuristic, it is unclear to what extent they will be amenable to further rigorous analyses – beyond predicting state transitions.

Numaoka has also done a related analysis on emergent collective behavior of swarms [8]. Like Lerman, Numaoka developed a model to predict when a heuristically controlled swarm would switch between strategies (which are like behaviors). The focus of Numaoka’s analysis, unlike ours and Lerman’s, however, was on how many “instigators” would be required to coerce the other agents into switching strategies.

In the future, we plan to formally analyze *all* important aspects of AP systems – for the purpose of developing behaviorally assured swarms.

Acknowledgements

We are grateful to David Book for much advice and guidance on selecting and representing the fields for visualization.

References

1. W. Fulks. *Advanced Calculus*. Wiley & Sons, 1978.
2. J. Kellogg, C. Bovais, R. Foch, H. McFarlane, C. Sullivan, J. Dahlburg, J. Gardner, R. Ramamurti, D. Gordon-Spears, R. Hartley, B. Kamgar-Parsi, F. Pipitone, W. Spears, A. Sciambi, and D. Srull. The NRL Micro Tactical Expendable (MITE) air vehicle. *The Aeronautical Journal*, 106(1062):431–441, 2002.
3. R. Kurshan. *Computer-Aided Verification of Coordinating Processes*. Princeton University Press, 1994.
4. R. Landau and M. Paez. *Computational Physics*. Wiley & Sons, 1997.
5. P. Langley. *Data-driven discovery of physical laws*. *Cognitive Science* 5: 31–54, 1981.
6. K. Lerman. Design and mathematical analysis of agent-based systems. In Rash, Rouff, Truszkowski, Gordon, and Hinchey, editors. *Lecture Notes in AI, volume 1871*: 222–234. Springer-Verlag, 2001.
7. K. Lerman and A. Galstyan. A general methodology for mathematical analysis of multi-agent systems. Tech. Report ISI-TR-529, USC Information Sciences, 2001.
8. C. Numaoka. Phase transitions in instigated collective decision making. *Adaptive Behavior*, 3(2):185–222, 1995.
9. W. Spears and D. Gordon. Using artificial physics to control agents. In *IEEE International Conference on Information, Intelligence, and Systems*: 281–288, 1999.
10. W. Spears and D. Gordon-Spears and R. Heil. Distributed, physics-based control of swarms of vehicles. In review.