

Fluid-Like Swarms with Predictable Macroscopic Behavior

Diana Spears, Wesley Kerr, and William Spears

Department of Computer Science
University of Wyoming, Laramie, WY 82071
dspears@cs.uwyo.edu

Abstract. This paper is concerned with assuring the safety of a swarm of agents (simulated robots). Such behavioral assurance is provided with the physics method called *kinetic theory*. Kinetic theory formulas are used to predict the macroscopic behavior of a simulated swarm of individually controlled agents. Kinetic theory is also the method for controlling the agents. In particular, the agents behave like particles in a moving gas.

The coverage task addressed here involves a dynamic search through a bounded region, while avoiding multiple large obstacles, such as buildings. In the case of limited sensors and communication, maintaining spatial coverage – especially after passing the obstacles – is a challenging problem. Our kinetic theory solution simulates a gas-like swarm motion, which provides excellent coverage. Finally, experimental results are presented that determine how well the macroscopic-level theory, mentioned above, predicts simulated swarm behavior on this task.

1 Safe Swarms

The research in this paper is designed with two objectives in mind: to effectively accomplish a difficult surveillance task, and to accomplish it in a manner that is “safe.” By “safe” we mean that the multi-agent collective that accomplishes the task is, in the aggregate, both *predictable* (behaviorally assured) and *controllable*.

The traditional approach to achieving safe agents is to engineer safety into the individual agents (e.g., [1] [2] [3]) and, sometimes, also into the particular interactions between these individual agents. This is typically accomplished with formal methods, such as model checking [4] or theorem proving [5], control theory [1], or other formalisms [3].

This paper explores an alternative view of safety. Our alternative view is motivated by a desire to model swarms (i.e., very large numbers) of agents cooperatively performing a task. The modern swarm philosophy is one of *emergent behavior*, which is defined as producing intelligent macroscopic behavior in the aggregate from lots of simple, unintelligent agents. The key to swarm agent/robotics emergent behavior is that even though the individual agent behaviors are easy to understand and may be expressed as simple rules, describing the behavior of the swarm as a whole requires a paradigm shift. In other words,

the description of the macroscopic behavior is not a straightforward function of descriptions of the microscopic behaviors – because safety of a swarm is often too computationally difficult to achieve by engineering safety into all the individual agents, which may have complex interactions.

The alternative view of safety that we propose for swarms is founded upon physics. Our choice is motivated by the fact that physics is the most accurate of all the scientific disciplines at predicting the macroscopic behavior of huge numbers of interacting particles using very simple mathematical formulas. Furthermore, physics disciplines, such as fluid dynamics, utilize these formulas for the design of systems that have desirable properties. Rather than engineer the properties into the individual particles (which of course cannot be done in most real-world situations), these disciplines promote engineering using principles at the macroscopic level. For example, in fluid dynamics, there is a field called the “control and management of turbulence dynamics,” in which questions are addressed such as how to ensure that desirable macroscopic fluid properties are preserved when the flow is controlled in a specified manner. Solutions to these questions are stated as control theoretic equations, which are expressed in terms of macroscopic properties of the fluid, such as velocity and pressure [6].

In summary, physicists have developed succinct formulas that are highly predictive of complex, multi-particle behavior. Such formulas can be used at an abstract level to design multi-agent systems with desirable properties, where agents are modeled as particles. The end result is behavioral assurance by engineering safety into the collective, rather than into the individual agents.

2 Physics of Large, Multi-particle Systems

The study of many-particle systems is one of the most active research areas in modern physics [7]. Today it is well known that although one could write down the equations of motion of individual atoms and their interactions, the complexity of doing this for a large number of particles is too daunting. The problem is not just quantitative (which would lead one to suspect that it could be solved with improved computational power), but it is also *qualitative* [7]. For example, how could one hope to understand the human abilities of natural language and planning by studying the properties of individual neurons? Likewise, physicists often predict macroscopic properties of matter, such as volume or pressure, from microscopic atomic particles by using statistical arguments, such as expectations, rather than resort to predictions of movement of the individual particles.

The physics of many-particle systems can be subdivided into three main disciplines [7]:

1. **Thermodynamics.** In the discipline of thermodynamics, descriptions are strictly at the macroscopic level. Valid formulas are derived based on a minimum number of postulates, without *any* detailed assumptions about the microscopic properties of the system.
2. **Statistical Mechanics.** Here, the macroscopic behavior of a multi-particle system is described based on the statistical properties of the microscopic behavior. The assumption is that the system is “in equilibrium.”

3. **Kinetic Theory.** The most popular physics approach to describing systems that are not necessarily in equilibrium is kinetic theory. Similarly to statistical mechanics, kinetic theory describes macroscopic behavior in stochastic terms. Note that kinetic theory subsumes statistical mechanics, i.e., in equilibrium, the former is the same as the latter.

The approach adopted in this paper is kinetic theory. Kinetic theory is the discipline of choice because it has a much richer and more extensive theory than thermodynamics (i.e., “relatively few statements can be made” in thermodynamics [7]), and because many multi-agent applications do not assume equilibrium.

Inspired by the huge success in applying physics to many-particle systems, we have decided to apply physics, in particular kinetic theory, principles to multi-agent swarms. Our agents are assumed to be simulated robots. Kinetic theory is used to control the agents, because fluid-like movement appears to be the most appropriate approach for a swarm of robots to achieve our task, and kinetic theory simulations are frequently used to model and study the movement of actual fluids.¹ In order to achieve a high level of predictive accuracy using kinetic theory, we use a multi-agent system that behaves like a many-particle physical system. Here, we show that using kinetic theory for both theory and simulation produces an excellent match between the two, thus providing a high degree of system behavioral assurance and safety. It is important to note that the match between theory and simulation is not achieved with a theory that predicts individual particle movements, which are generated within the simulation. Rather, the theory uses stochastic properties of the swarm as a whole to make its predictions. In other words, the theory is macroscopic and is predictive of the simulation, which is microscopic in its design and implementation.

3 The Sweeping and Obstacle Avoidance Task

The task being addressed here, which is also described in [9], consists of sweeping a large group of mobile robots through a long bounded region (a swath of land, a corridor in a building, a city sector, or an underground passageway/tunnel), to perform a search, i.e., surveillance. This requires maximum coverage. The robots have a limited sensing range for detecting other agents/objects. It is assumed that robots near the corridor boundaries can detect these boundaries, and that all robots can sense the global direction that they are to move, e.g., by using a compass. There is no other global information, and the agents behave in a distributed, non-centralized manner. As they move, the robots need to avoid large obstacles, such as large buildings. This poses a challenge because with a limited sensing range, robots on one side of a building cannot necessarily communicate with robots on the other side. The search might be for enemy mines, survivors of a collapsed building or, alternatively, the robots might be patrolling the area. It is assumed that the robots need to keep moving, because

¹ An alternative is a molecular dynamics (MD) simulation that is deterministic. For results of a physics-based approach using MD simulations, see [8].

there are not enough of them to view the entire length of the region at once. In other words, the robots begin scattered randomly at one end of the corridor and move to the opposite end (considered the “goal direction”). This is a *sweep*. A sweep terminates when all robots have reached the goal end of the corridor, or a time limit is reached. Once the robots complete one sweep, they reverse their goal direction and sweep back again. Finally, if stealth is an issue then we would like the individual robot movements to be unpredictable to adversaries. It is conjectured that the behavior of a gas is most appropriate for solving this task, i.e., each robot is modeled as a gas particle.

4 Motivation for Using a Fluid-Like Swarm

The term “fluid” refers to both liquids and gases; this paper focuses on gases in motion. Although individual atoms or molecules in a gas have unpredictable locations at any instant in time, the gas is predictable at a macroscopic level. Furthermore, when gases are placed in a container, they expand to fill the container, thereby providing outstanding spatial coverage. When not in an equilibrium state, gases can also move in bulk. The gas can be transported either by advection (due to the velocity of the ambient air in which it has been dispersed) or due to molecular diffusion, e.g., if the gas is heavier than the ambient air then it will fall slowly to the ground. Gases will also flow around obstacles, and then expand after passing the obstacles, thereby filling the space again.

These forms of coverage are precisely the ones required for excellent performance on the sweeping and obstacle avoidance task. Therefore, we use a kinetic theory particle simulation to model our agent swarm performing the task.

5 Kinetic Theory for Simulating Fluids

Our *kinetic theory (KT) simulation* is a microscopic model of individual particles, which are considered to be agents, or simulated robots. Our simulation, as well as this overview of it, borrows heavily from Garcia [10].

When modeling a gas, the number of particles is problematic, i.e., in a gas at standard temperature and pressure there are 2.687×10^{19} particles in a cubic centimeter. A typical solution is to employ a stochastic model that calculates and updates the probabilities of where the particles are and what their velocities are. This is the basis of KT. One advantage of this model is that it enables us to make stochastic predictions, such as the average behavior of the ensemble. The second advantage is that with real robots, we can implement this with probabilistic robot actions, thereby avoiding predictability of the individual agents, e.g., for stealth.

In KT, particles are treated as possessing kinetic energy but no potential energy (i.e., an ideal gas), and collisions with other particles are modeled as purely elastic collisions that maintain conservation of momentum. Using kinetic theory formulas, we can predict useful macroscopic properties of the system, such as the average speed or kinetic energy of the particles in the system. For

example, assuming k is Boltzmann's constant, where $k = 1.38 \times 10^{-23}$ J/K, m is the mass of any particle, $f(v)$ is the probability density function for speed, and T is the temperature of the system, then the average speed of any particle (in 3D) is:

$$\langle v \rangle = \int_0^\infty v f(v) dv = \frac{2\sqrt{2}}{\sqrt{\pi}} \sqrt{\frac{kT}{m}}$$

From this formula, one can see that the temperature T plays an important role in KT. In our KT simulation, T is a user-defined system parameter analogous to real temperature. In other words, increasing T in our system raises the *virtual* system heat (analogous to actual heat), for the purpose of increasing the system kinetic energy and thereby increase the particle motion, i.e., the speed of the agents in the simulation.

Our KT simulation algorithm is a variant of the particle simulations described in [10]. We substantially modified the algorithms in [10] to tailor them to simulated robots with local views. Robots, modeled as particles, behave in the aggregate like "Couette flow." Figure 1, from [10], depicts one-sided Couette flow, where a fluid moves through some environment between two walls – one wall moving with velocity v_{wall} , and the other stationary (the environment is the frame of reference). In this Couette, fluid is assumed to move in the positive y -direction (i.e., longitudinally toward the goal end of the Couette corridor), and the positive x -direction goes from the stationary wall to the moving wall (i.e., laterally across the Couette corridor). Note that the direction of virtual motion of Couette walls is determined by using a compass to sense the goal direction. In general, we have found that a Couette is useful because it introduces an external source of kinetic energy into the system and gives the agents a direction to move.

Because the fluid is Newtonian and has viscosity, there is a linear velocity distribution across the system. Fluid deformation occurs because of the shear stress, τ , and the wall velocity is transferred (via kinetic energy) because of molecular friction on the particles that strike the wall. On the other hand, the particles that strike either wall will transfer kinetic energy to that wall. This does not cause the wall to change velocity, since in a Couette flow the walls are assumed to have infinite length and depth and therefore infinite mass. We chose a Couette flow in order to introduce kinetic energy into the system and to give the particles a direction to move.




Fig. 1. Schematic for a one-sided Couette flow




Fig. 2. Schematic for a two-sided Couette flow

Our 2D simulated world models a modified (two-sided) Couette flow in which both Couette walls are moving in the same direction with the same speed (see Figure 2). We invented this variant as a means of propelling all robots in a desired general direction, i.e., the large-scale fluid motion is approximately that of the walls.

In our simulation, each agent, which is modeled abstractly as a holonomic particle, can be described by a position vector \mathbf{p} and a velocity vector \mathbf{v} . At each time step, every agent resets its position based on how far it could move in the given time step utilizing its current velocity. Particle velocities are initialized to be a random function of the (virtual) system temperature T , and these velocities remain constant unless collisions occur. Collisions are the primary mechanism for driving particle movement/acceleration in a KT simulation. (Note that with actual robots, collisions and wall motion would be virtual.) The system updates the world in discrete time steps, Δt , which occur on the order of the mean collision time for an agent.

At each time step, every agent in the system updates its position. When updating its position, a check is performed first to see if the movement would cause a (virtual) collision between the agent and a wall. If a collision would occur, then the agent selects a new velocity from a biased Maxwellian distribution, which is a function of the system temperature. If the agent is about to strike a moving wall, then some of the energy from the wall is transferred to the agent. Inter-agent (virtual) collisions are then processed. The number of collisions in any given region is a stochastic function of the number of agents in that region. In particular, the probability of a virtual collision between two agents is based on their proximity, but is independent of the angle between their velocity vectors. The new post-collision velocity vectors are based on the center of mass vector, coupled with a random component. See [10] and [11] for details. This process continues indefinitely or until a desired state is achieved.

6 The Surveillance Task Simulation

For a model of the surveillance task scenario, we have developed a 2D simulation of the task scenario, i.e., an obstacle-laden corridor with KT-driven robots flowing through it. The “two-sided” variant of the traditional Couette model is




Fig. 3. KT controllers perform a sweep. The snapshots progress in time from left to right.

used (recall this variant from Figure 2), in which both Couette walls move in the same direction with the same speed. The two-sided Couette is highly effective at driving bulk swarm movement in the goal direction. The agents begin in random locations at the top of the corridor, and sweep down the corridor in the goal direction. Typical results for a sweep are shown in Fig. 3. Because robots are not capable of distinguishing corridor walls from obstacle walls, any obstacle wall parallel to the Couette walls is considered to be “Couette” (in motion), and the (virtual) wall velocity is added to the y -component of the velocity for any agent that collides with a Couette wall. Therefore, obstacle walls, in addition to actual Couette walls, can increase the velocity of robots toward the goal.

Inter-robot collisions are processed in localized regions. Likewise, robots only (virtually) collide with walls that are in close proximity to them.

With this simulation of a swarm of agents performing the task, the question arises of how to use kinetic theory to predict swarm behavior? One option is to develop a theory that models the physics of the entire task, including obstacles. This would require considerable work, and would be especially difficult if we do not know beforehand the number, sizes, and shapes of the obstacles. An alternative option is to develop a simpler theory that makes assumptions that do not hold in the full task, and then scale up the task in simulation to see how well the theory holds when its assumptions are violated. The latter is the option adopted here, and is the topic of the following section.

7 Theoretical Predictions of Macroscopic Behavior

One of the primary advantages of physics-based swarms is the large existing body of physics-based theory for predicting system behavior and ensuring that the swarm will behave “safely.” A secondary advantage is that the theory can

be used for optimal selection of system parameter values [12]. In this section, we provide evidence that kinetic theory is predictive of the behavior of kinetic-theory-based simulations. Furthermore, we demonstrate the practical use of the theory for parameter value selection.

We focus on subtasks of the sweeping and obstacle avoidance task, described above. Recall that our objectives for the agents in this task are to sweep a corridor and avoid obstacles. The ultimate objective is to maximize coverage. Consider two types of spatial coverage: *longitudinal* (in the goal direction) and *lateral* (orthogonal to the goal direction). Longitudinal coverage can be achieved by movement of the swarm in the goal direction; lateral coverage can be achieved by a uniform spatial distribution of the robots between the side walls of the corridor. The objective of the coverage task is to maximize both longitudinal and lateral coverage in the minimum possible time (i.e., to also maximize temporal coverage).

To measure how well the robots achieve the task objective, we examine the following three metrics:

1. The degree to which the **spatial distribution** of the robots matches a uniform distribution. This is a measure of lateral coverage of the corridor and provides confirmation of equilibrium behavior.
2. The **average speed** of the robots (averaged over all robots in the corridor). This is a measure of all three types of coverage: lateral, longitudinal, and temporal. Velocity is a vector consisting of speed and direction. The type of coverage depends on the direction. To control the average swarm speed, one can directly vary the value of the system temperature, T . Therefore, our experiment explores the relationship between T and average speed.
3. The **velocity distribution** of all robots in the corridor. This is a measure of longitudinal spatial coverage, as well as temporal coverage. For example, consider the one-sided Couette in Figure 1 again, and in particular focus on the line representing the velocity distribution. The slope of that line is inversely proportional to the longitudinal spatial coverage (and the temporal coverage). In other words, for a given Couette diameter, D , if you increase the wall speed, v_{wall} , then the slope will be reduced and the longitudinal and temporal coverage will increase. Below, we run an experiment to examine the relationship between the wall speed, v_{wall} , and the velocity distribution in one-sided and two-sided Couettes. The intention is to enable the system designer to select a wall speed for optimizing the swarm velocity distribution.

The theory presented in this paper assumes simplified 2D environments with no obstacles. To develop theory for the full task simulation would require extensive theoretical physics analyses, which is beyond the scope of this paper. This will be tackled as future work.

We ran three sets of experiments, in accordance with the metrics defined above. For each experiment, one parameter was perturbed and eight different

values of the affected parameter were chosen. For each parameter value, 20 different runs through the simulator were executed, each with different random initial robot positions and velocities. The average simulation results and relative error (over the 20 runs) were computed and graphed.

For these experiments, we defined the error between the theoretical predictions and the simulation results, denoted *relative error*, to be:

$$\frac{| \textit{theoretical} - \textit{simulation} |}{\textit{theoretical}} \times 100$$

Although the theory assumes no obstacles, in the simulation we ran with six obstacle densities, ranging from 0% to 50%. Surprisingly, some of the theory holds well, despite this.

There are two complementary goals for running these experiments. The first goal is to determine how predictive the theory is. Derivations of all laws (predictive theoretical formulas) are in [9]. The second goal is to determine the relationship between parameter settings and system behavior. If a system designer understands this relationship, he/she can more easily set parameters to achieve optimal performance. Finally, and most importantly, the reason why these two goals are complementary is that if the theory is predictive of the system simulation behavior, then future system designers no longer need to run the simulation to determine the optimal parameter settings – graphs generated from theory alone will suffice. This can reduce the system design time.

8 Experiment 1: Spatial Distribution

The first experiment examines the equilibrium spatial distribution of agents within an enclosed region, i.e., a square “container. The agents begin in a tight Gaussian distribution, which then diffuses until equilibrium has been reached. During this experiment, there is no goal force or wall movement, and therefore no *externally-directed* bulk transport of the swarm.

The purpose of this experiment is to confirm the theoretically expected behavior of a KT system in equilibrium, which will thereby verify the correctness of our implementation – a big advantage of our approach. The KT gas model predicts that, upon reaching equilibrium, the particles will be spatially uniformly distributed. To confirm this prediction, we divided the square container in our KT simulator into a 2D grid of cells. Theory predicts that there should be (on average) n/c robots per cell, where n is the total number of robots and c is the total number of grid cells. We ran with six obstacle densities, ranging from 0% to 50%, to determine the sensitivity of the spatial distribution to obstacle density.

Fig. 4 shows the experimental results. Note that despite the introduction of as much as a 50% obstacle coverage, we can still predict the spatial distribution with a relative error of less than 10%, which is surprisingly low. The error is very low once the number of robots is about 300, which is surprising considering that 300 robots is far less than the 10^{19} particles typically assumed by traditional kinetic theory.

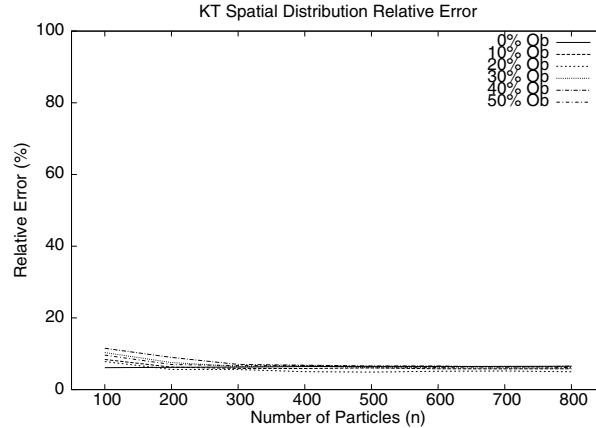


Fig. 4. Relative error for the KT spatial distribution

9 Experiment 2: Average Speed

For the second experiment, we examine the average speed of the robots in the system. Once again, there is no external force or externally-directed bulk transport of the swarm. However, recall that the agents will increase their speed if there is an increase in the system temperature, which causes an increase in kinetic energy. The objective of this experiment is to examine the relationship between the temperature, T , and the average speed of the robots. The average robot speed serves as a measure of how well the system will be able to achieve complete coverage – because higher speed translates to greater lateral and/or longitudinal coverage, depending on the velocity direction. This experiment also serves to verify that our simulation code has been implemented correctly. Note that not all applications will require maximum coverage; therefore, we want to study the general question of precisely how specific choices of speed affect coverage.

Our predictive formula for 2D is (see [9] for the mathematical derivation):

$$\langle v \rangle = \frac{1}{4} \sqrt{\frac{8\pi kT}{m}}$$

where k is Boltzmann's constant ($1.38 \times 10^{23} J/K$), m is the robot mass (assumed to be one), and T is the system temperature.

This theoretical formula is compared with the actual average speed, $\langle v \rangle$, of the robots in the simulation, after the system has converged to an equilibrium state. There were 300 robots in the simulation. Because temperature affects speed, temperature was varied from 50° Kelvin to 400° Kelvin. We ran with six obstacle densities ranging from 0% to 50%, in order to determine the sensitivity of the average speed to obstacle density.

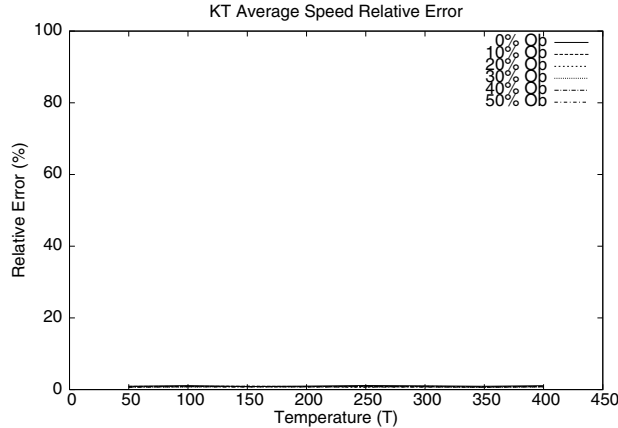


Fig. 5. Relative error for the KT average speed

The results are striking, as can be seen in Fig. 5. The difference between the theoretical predictions of the average speed and the simulated average speed results in less than a 2% error, which is outstanding considering that as much as a 50% obstacle coverage has been introduced. Finally, note that we can use our theory to not only predict swarm behavior, but also to *control* it. Specifically, by setting the temperature T , a system designer can easily achieve a desired average speed.

10 Experiment 3: Velocity Distribution

The third experiment concerns the velocity distribution of a robot swarm in a Couette. The theoretical prediction is compared with simulated behavior. Recall that in a Couette, fluid flow is in the y -direction – toward the goal. The x -direction is lateral, across the corridor. In addition to seeing how predictive the theory is, this experiment also examines the relationship between wall speed, v_{wall} , and the velocity distribution of the robots in the system.

We first focus on a subtask in which a traditional one-sided Couette flow drives the bulk swarm movement. Our predictive formula is (see [9] for the derivation):

$$v_y = \frac{x}{D} v_{wall}$$

where v_{wall} is the velocity of the Couette wall, x is the lateral distance from the stationary wall, and D is the Couette width. In other words, the theory predicts a linear velocity distribution.

We set up an experiment to measure the relative error between theory and simulation, consisting of 300 robots. The corridor was divided into eight discrete longitudinal sub-corridors. Theory predicts what the average speed will be lengthwise (in the goal direction) along each of the sub-corridors. Within each

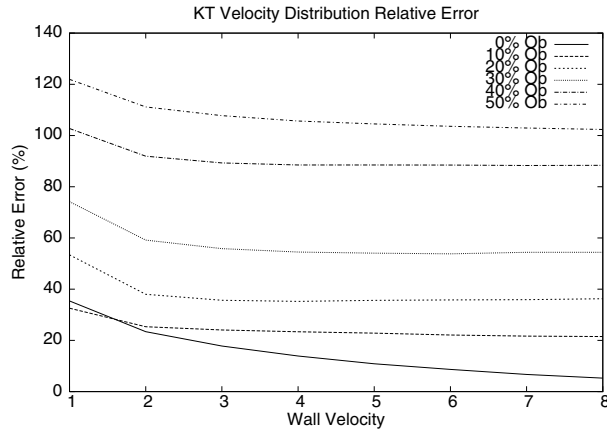


Fig. 6. Relative error for the KT velocity distribution

sub-corridor, the average y velocity of the robots is measured. The relative error between the theory and the experimental results is then calculated, for each sub-corridor. Finally, the relative error is averaged across all sub-corridors and plotted in Fig. 6 for eight different wall speeds and six different obstacle percentages. Note that although the error is reasonably low for 0% obstacles and high wall speeds, error increases dramatically as obstacles are added.

Why is there a discrepancy between the theory and experimental results? The reason is that theory predicts a certain linear velocity distribution, but assumes no obstacles. For simplicity, the theory assumes that robots never move backward (back up the corridor). In the simulator, on the other hand, robots *do* move backward, regardless of whether or not there are obstacles – because the




Fig. 7. The velocity distributions as the density of obstacles increases, for one-sided Couette flow

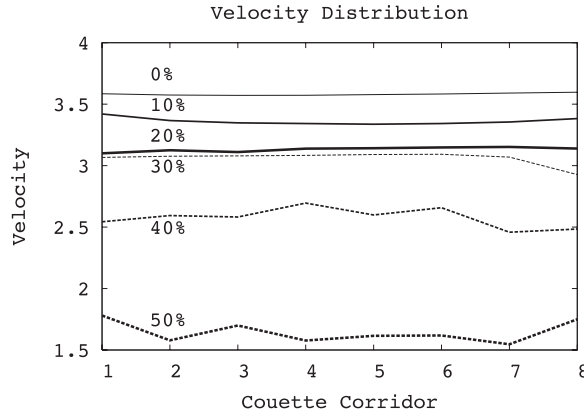


Fig. 8. The velocity distributions as the density of obstacles increases, for two-sided Couette flow

simulation has a random component. In fact, as obstacles are introduced into the simulated world, the frequency of backward moving robots increases substantially. To examine more closely the effect of obstacles, Figure 7 shows the velocity distributions themselves (where the wall velocity $v_{wall} = 4$). Even with no obstacles, the maximum robot velocity does not quite reach 4.0 (we would expect 3.75 in the sub-corridor nearest to the wall). This is caused by the backward moves. What is interesting is that the velocity distributions remain linear up to a reasonable obstacle density (30%), while the slope decreases as obstacles are added. Adding obstacles is roughly equivalent, therefore, to lowering the wall velocity v_{wall} !

To see if the correlation between obstacle density and wall velocity holds in the two-sided Couette flow, we re-ran the previous experiment, but with *both* walls having $v_{wall} = 4$. The results are shown in Figure 8. The theory predicts (see [9] for the derivation) that for the two-sided Couette, $v_y = v_{wall}$ regardless of the value of x . Note that, as theory predicts, the velocity distribution of the flow is independent of the distance from the walls – the large scale fluid motion is approximately that of the walls. Again, increasing obstacle density is very similar to decreasing wall speed.

In conclusion, without obstacles, the theory becomes highly predictive as the wall velocity increases. Furthermore, this very predictive theoretical formula can also be used to achieve a desired swarm velocity distribution, i.e., to *control* the swarm – simply set the value of v_{wall} , the virtual wall speed, to achieve the desired distribution, using the formula. On the other hand, with an increasing number of obstacles, the predictability of the theory is increasingly reduced. However, we have shown that (up to quite reasonable densities) the increase in the number of obstacles is roughly proportional to a reduction in wall velocity, v_{wall} .

11 Discussion of Theoretical Predictions

With respect to the average speed and the spatial distribution, the theory is highly predictive. Although the theory assumes no obstacles, the addition of obstacles to the simulation has a minimal effect on the results, with errors typically under 10%. Surprisingly, this level of theoretical accuracy is achieved with only hundreds of robots, which is very small from a kinetic theory perspective.

Our results for the velocity distribution are acceptable for no obstacles, generally giving errors less than 20%. As the obstacle density increases, so does the error. However, we have shown that an increase in obstacle density changes the slope of the linear velocity distribution. This is roughly equivalent to a commensurate reduction in wall speed.

In summary, we can conclude that when the actual scenario closely coincides with the theoretical assumptions (e.g., few obstacles), the theory is highly predictive. Also, we have provided an important insight into the nature of the effect that obstacle density has on the system. The most important conclusion to be drawn from these experiments is that in the future we can largely design KT systems using theory, rather than computationally intensive simulations, for the selection of optimal parameter settings. A subsidiary conclusion is that we have verified the correctness of our swarm code, which is something quite straightforward for a physics-based approach but much more difficult for alternative approaches.

12 Conclusions

KT uses a stochastic algorithm for updating particle positions; therefore KT predictions can only be approximate. Furthermore, as stated in [10], Monte Carlo simulations such as KT need very long runs and huge numbers of particles to acquire enough statistical data to produce highly accurate (theoretically predictable) results. We cannot guarantee this, since we are developing control algorithms for robotic swarms with a few to a few thousand robots under strict time limitations. Despite all of these limitations of our KT robotic swarm simulation, the theory is nevertheless highly predictive of the simulation results. The conclusion is that our approach of using KT for designing swarm-based multi-agent systems has great promise for engineering swarms that are “safe.”

13 Related and Future Work

The work that is most related consists of other theoretical analyses of swarm systems. Our comparisons are in terms of the goal and method of analysis. There are generally two goals: stability and convergence/correctness. Under stability is the work in [13, 14, 15]. Convergence/correctness work includes [13]. Other goals of theoretical analyses include time complexity [16], synthesis [17], prediction of movement cohesion [13], coalition size [14], number of instigators to switch strategies [18], and collision frequency [19].

Methods of analysis are also diverse. The most relevant is work on physics-based analyses of physics-based swarm robotics systems. We are aware of four classes of methods. The first is Lyapunov analyses, e.g., [15]. The second is force and energy analyses, e.g., [12,17]. The third develops macro-level equations describing flocking [20]. Finally, the fourth is the most related work of all – the KT research by Jantz and Doty [19]. Note that although Jantz and Doty showed that KT can be used to model multi-agent swarms with predictable behavior [19], our research extends theirs by providing a much more extensive and methodical study of the relationship between kinetic theory and simulation.

Our current research [21,22] has compared KT against behavior-based approaches and found that it performs competitively, even if the alternative algorithms have more information. The next step is to transition KT from simulation to real robot swarms. We already have substantial progress on robotic implementations [23], and the next step will be to add KT.

References

1. Perkins, T., Barto, A.: Lyapunov design for safe reinforcement learning control. In: AAAI Spring Symposium on “Safe Learning Agents” (2002)
2. Barley, M., Guesgen, H.: Towards safe learning agents. In: AAAI Spring Symposium on “Safe Learning Agents” (2002)
3. Shapiro, D.: Value-driven agents. PhD thesis, Stanford University (2001)
4. Gordon, D.: Asimovian adaptive agents. *Journal of Artificial Intelligence Research* 13 (2000)
5. Robinson, P., Hinchey, M., Clark, K.: Qu-prolog: An implementation language for agents with advanced reasoning capabilities. LNCS (LNAI), pp. 162–172 (2002)
6. Barbu, V., Sritharan, S.: Flow invariance preserving feedback controllers for the navier-stokes equation. *Journal of Mathematical Analysis and Applications* 255, 281–307 (2001)
7. Reif, F.: *Fundamentals of Statistical and Thermal Physics*. McGraw-Hill, New York (1965)
8. Spears, W., Gordon, D.: Using artificial physics to control agents. In: *IEEE International Conference on Information, Intelligence, and Systems*, pp. 281–288 (1999)
9. Kerr, W., Spears, D., Spears, W., Thayer, D.: Two formal gas models for multi-agent sweeping and obstacle avoidance. LNCS (LNAI) (2004)
10. Garcia, A.: *Numerical Methods for Physics*, 2nd edn. Prentice Hall, Englewood Cliffs (2000)
11. Kerr, W., Spears, D.: Robotic simulation of gases for a surveillance task. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2005* (2005)
12. Spears, W., Gordon-Spears, D., Hamann, J., Heil, R.: Distributed, physics-based control of swarms of vehicles. *Autonomous Robots* 17, 137–162 (2004)
13. Liu, Y., Passino, K., Polycarpou, M.: Stability analysis of m-dimensional asynchronous swarms with a fixed communication topology. *IEEE Transactions on Automatic Control* 48, 76–95 (2003)
14. Lerman, K., Galstyan, A.: A general methodology for mathematical analysis of multi-agent systems. Technical Report ISI-TR-529, USC Information Sciences (2001)

15. Olfati-Saber, R., Murray, R.: Distributed cooperative control of multiple vehicle formations using structural potential functions. In: IFAC World Congress (2002)
16. Shehory, O., Kraus, S., Yadgar, O.: Emergent cooperative goal-satisfaction in large-scale automated-agent systems. *Artificial Intelligence* 110, 1–55 (1999)
17. Reif, J., Wang, H.: Social potential fields: A distributed behavioral control for autonomous robots. In: Workshop on the Algorithmic Foundations of Robotics (1998)
18. Numaoka, C.: Phase transitions in instigated collective decision making. *Adaptive Behavior* 3(2), 185–222 (1995)
19. Jantz, S., Doty, K.: Kinetics of robotics: The development of universal metrics in robotic swarms. Technical report, Dept of Electrical Engineering, University of Florida (1997)
20. Toner, J., Tu, Y.: Flocks, herds, and schools: A quantitative theory of flocking. *Physical Review E* 58(4), 4828–4858 (1998)
21. Spears, D., Kerr, W., Spears, W.: Physics-based robot swarms for coverage problems. *International Journal on Intelligent Control and Systems* (2006)
22. Kerr, W.: Physics-based multiagent systems and their application to coverage problems. Master's thesis, University of Wyoming (2005)
23. Spears, W., Hamann, J., Maxim, P., Kunkel, T., Zarzhitsky, D., Spears, D., Karlsson, C.: Where are you? In: Şahin, E., Spears, W.M., Winfield, A.F.T. (eds.) SAB 2006 Ws 2007. LNCS, vol. 4433, pp. 129–143. Springer, Heidelberg (2006)