Microbial MAS

STUDENT NAMES REMOVED (2 UNDERGRADUATES)

In this report, we present and analyze a multi agent system (MAS) simulating the internal balance of the bloodstream. We present the design for a 'defender' agent loosely based around the function of the existing white blood cell. Here we analyze the effectiveness and efficiency of this agent model against entrenched infections and ongoing pathogenic assault. In addition, we will attempt to compare these agents against the performance of actual white blood cells in the human body.

General Terms: Design, Experimentation, Simulation, Agent
Additional Key Words and Phrases: MAS, Blood System, Pathogen, Microbe

1. INTRODUCTION

Medical technology and understanding has continued to improve steadily since man has first encountered death. From blood letting, to balance of the humors, to spontaneous generation, to genetic analysis, our understanding of why we get sick and how to combat illness has continued to evolve. However, since the invention of the computer, man has had a great ally in the battle for health. The last 50 years have seen great advances in the field of medicine where technology is involved. Surgery has become less invasive, treatments more specialized, and recovery times shorter. Yet a new age still looms over the horizon; the age of nano-machines. The promise of nano-machines is to be able to have a network of intelligent, responsive machines ubiquitous in human systems capable of analyzing, adapting to, and neutralizing any threat.

With this motivation in mind, we aim to design an effective and efficient agent that is capable of fighting, defending, and organizing a defense within the human bloodstream. To this end, we have developed an environment with supporting and opposing agents, and run extensive simulations to determine if our implementation was capable of improving on the natural design.

Throughout the paper, the motivations for and the relevance of the following hypotheses will be evaluated:

H1: If unchecked, pathogen growth will model the exponential growth displayed by true pathogen infection patterns.

H2: If not allowed to spawn, pathogens will have much more difficulty achieving emergent behavior.

This research was done as coursework for CSCE475/875 at the University of Nebraska, Lincoln.

H3: Due to exponential nature of pathogens and linear nature of white blood cells, any changes to parameters should cause a clean switch from total system failure (pathogens win) and complete system health (pathogens completely defeated).

H4: Defender agents are capable of defeating an equal number of pathogens

H5: Higher saturations of white blood cells will lead to quicker convergence and higher system utility.

H6: Higher pathogen population will result in lower overall system health.

H7: Higher communication range multipliers will allow the defenders to more ably defeat sparse pathogen threats.

H8: Emergent behavior in the system is only possible when the ratio of white blood cell density to pathogen density is over a certain value.

H9: When average pathogen virulence, E_P, and white blood cell effectiveness are related such that $E_P - E_W > T$, for some threshold T, emergent behavior will fail.

This report is composed as follows. Section 2 of this report describes the system model, the environment, and the agents. Section 3 goes into the experimental setup, while section 4 holds our experimental results. Section 5 holds our conclusions.

2. SYSTEM DESIGN

2.1 Environment

No agent is independent or truly agnostic of its environment, and therefore we must first define the world in which these agents interact. We simulate the bloodstream as a continuous, m x n grid. The system is essentially closed, as no white or red blood cells may enter or exit the system. However, pathogens are permitted to enter the system in some simulations, to represent the incursion of infectious agents.

In a true simulation of the human body, cells can be manufactured and tailored to fight specific infection over time. In addition, the continuous flow of the bloodstream constantly brings new agents into and out of a particular area. However, our model takes a closed, microcosmic view of a particular section of the body. In some ways this makes the environment simpler to understand. However, in other ways, it makes the problem facing the white blood cells much more difficult, as there are no reinforcements and no time to adapt.

The value used to track and monitor global coherence in the system was “system health.” This value was aimed to represent a quantitative state of the body at the end of

the simulation, and therefore the likeliness of serious illness or death (in the case of the more virulent pathogens). This value was a weighted sum of the total final effectiveness of all red blood cells and white blood cells left in the system. For these cells, this represents their overall capability to continue functioning and fighting off infection. The system health is further decreased by a weighted sum of the total final effectiveness of all remaining pathogens. This represents the pathogens' overall capability to continue corrupting the system. For this reason, system healthiness can become negative. In general, we consider any final healthiness at or below 0 to be a total failure of system defense.

System health can be considered the 'social good' of the system. Pathogens are attempting to replicate (and indirectly minimizing the social good of the system) using individually rational choices. Red blood cells are looking to maximize their personal utility (based on keeping their own effectiveness high). White blood cells are looking to reduce pathogen effectiveness, maintain their own, and keep red cells alive and performing their function.

2.2 Microbial Agents

In our system, we have three types of microbial agents: helpless agents, defender agents, and invader agents. Helpless agents are designed to represent the workers of the system; incapable of fighting, but necessary for bodily function. They are designed to represent red blood cells and will be referred to as such for the remainder of this report. The defender agents are the most intelligent and cooperative of the three agent types, and are the goal of the system design. We will refer to them as white blood cells. The invaders are, of course, pathogens whose only goal is to invade the system and propagate.

All agents in the system are assigned an effectiveness value when they are created. This value is an amalgamation of their life force, capabilities, and strength. This value is fixed for all red and white cells in the system and is uniformly distributed over a range for the microbes. This value is the basis of every agent's move range, sight range, communication range, fighting strength, and health. This value is an abstraction of microbial 'fitness,' or 'virulence' amongst the pathogens.

When a real microbe attacks another, the assault consists of attaching to the victim (slowing it and making it less effective), and subsequently neutralizing or consuming it over time. We simulate this behavior in our system by using an agent's effectiveness. When agents are engaged in fighting, they are bound to each other, and their current effectiveness is temporarily decreased by the effectiveness of all their opponents. So,

agents who are engaged in fighting in the system have a temporarily reduced sight, move, and communication sensing range. This is designed to simulate the effect of many microbes attaching to a single target, and reducing its ability to continue on as normal. When agents are dealt damage in battle, their effectiveness value is permanently decreased to represent depleted resources and damaged capabilities. When a microbe's effectiveness is decreased to 0, it is considered dead.

All agents in the system have the capability to move about the system at some rate based on their effectiveness. Throughout this work, an agent may be said to 'move randomly' or to 'move by vector.' For the remainder of this work, these will be considered equivalent. These terms represent the random microbial movement in the system. Essentially, all agents are inserted into the system with a random orientation. Every turn, they modify this orientation by a uniformly random but limited amount (between [-40, 40] degrees) and move along this orientation. We found that this method of locomotion appeared to resemble that of microorganisms more closely than that generated by completely random movement.

Agents in the system have a 'sight' range based off of their effectiveness value. The microbes in our system can not truly see, but this sense is an abstraction of the data gathering capabilities of agents in the system. This could be done through sensing disturbances or movements in the medium, proximity to nearby objects, or recognition of chemical patterns or cellular waste. This value abstracts an agent's capability to gather information on its surroundings, and naturally leads to more informed local decisions.

2.2.1 Red Blood Cells. As previously mentioned, red blood cells are represented by agents in the system. However, they are very simple and weak agents. They are not designed to fight or do any complex reasoning, but rather to go about their work. They cooperate with the white blood cells, but they are incapable of communicating. These agents represent the bulk of the agents (per volume) in the bloodstream, the primary target of the pathogens, and subsequently the main charge of the white blood cells.

For the most part, red blood cells simply move by vector every turn. This behavior represents the ubiquitous and seemingly random presence of red blood cells in the bloodstream. Their effectiveness is low, and therefore they have little information gathering capabilities. However, though they are weak, they are given some rational reasoning capabilities. Their first priority is to preserve themselves, and second is to go about the necessary functions of the body. Therefore, if they sense a nearby pathogen and a nearby white blood cell, they will gravitate towards the protection offered by the

white blood. If this is not the situation, or this is impossible, they will do their best to carry on their function.

2.2.2 White Blood Cells. White blood cells are the most advanced agents in the system, and are therefore the closest to what could be considered intelligent agents. We fixed the white blood cells initial effectiveness at a value markedly higher than that of the red blood cells to represent their design as a fighting microbe. The job of these agents is to continuously explore the environment, find, and engage any threats to the system health. To this end, they must weigh the cost of preserving their own life/effectiveness, defending red blood cells, assisting other white blood cells, engaging unchecked pathogens, and exploring the system.

The main advantage the defender agent has over the other types is its limited communication capacity. Essentially, this capability is an abstraction of the use of pheromones or other chemical signals. The nature of such a communication protocol is a completely pull based system (no push) in which agents can sense the presence of said signals up to a certain distance. This limited distance is again based off of the agent's current effectiveness. However, such a communication protocol is naturally discrete, allowing only a handful of different signals to be sent. The interpretation of such communication, therefore, is left up to the presence of a common understanding of signal meaning. Essentially, there were five different signals agents could transmit. Three were used in combat with a pathogen, and were emitted to represent a situation in which:

1. White blood cells are losing the fight
2. Both sides are evenly matched
3. White blood cells are winning the fight

These signals were designed to help prioritize the importance of helping the agent. The other two signals were for when an agent saw and was moving to engage a pathogen of unknown capability, and one for normal operation (patrolling).

Essentially, this is an indirect communication method because an agent simply chooses which pheromone he wants to emit, and does so without directing it at anyone. Any agents within sensing distance (if any) are capable of perceiving this signal if they wish. The signal itself is not sent (pushed) by the transmitter, but rather detected (pulled) by agents within range. The difference this establishes is that if all nearby agents are otherwise engaged, and do not wish to look for pheromones, they will not know about the signals. We felt that this is both a low power, low computation communication method, and possibly as complex of communication as microbes are capable of. For simplicity in

the system, we assume that agents can discern distance and direction of transmission sources in the medium (using signal strength and intensity). We mention as one of the future directions of the system that it could be implemented such that agents can perceive the signals without being able to discern its source and still achieve global coherence.

White blood cells also have the most complex reasoning of any agents in the system. As mentioned before, they must choose between several courses of action that they perceive as beneficial. Of course, this is further complicated, as the exact utility gain expected from any particular action is difficult to calculate. Time can often be of the utmost importance when dealing with a spreading issue. Leaving pathogens unchecked can have dire consequences, but attacking a pathogen to slow it down can end in the destruction of the white cell. In addition, preferring to attack unchecked pathogens over assisting other white cells in need of help could lead to a reduction in illness fighting effectiveness, especially if the pathogens are strong. Also, a defender must choose between traveling to a high priority target and dealing with something nearby this turn.

With these considerations in mind, we set up the white blood cells' rationale as per the following priorities:

1. Help fellow defenders within move range
2. Attack high priority threats within move range
3. Help defenders within communication range
4. Pursue high priority threats within sight range
5. Protect unprotected, low effectiveness cells if no other protectors are around
6. Move randomly (explore)

We believe that weighting choices such that this general preference ordering is observed will allow for the emergent behavior of a defended system and high system health. Within a single directive, an agent will prefer helping defenders who are doing worse, attacking targets that are eating red blood cells or who are stronger, and will generally prefer shorter travel times. When two equivalent targets are both within move range, a defender will move to the farther, as that means more exploration and possibly better choices for next turn.

2.2.3 Pathogens. Pathogens are the invaders of the system who aim to replicate themselves and preserve their own lives. They are smarter than red blood cells, and are capable of being stronger than white blood cells. However, they do not communicate, and do not cooperate with other pathogens. They are all selfish agents.

Pathogens are assigned a random, continuous effectiveness value uniformly from the range [E_{Pmin}, E_{Pmax}]. This is to simulate different strains of pathogens that might be present in the environment. In addition to the initial population of pathogens in the system, pathogens can join the system over time. The number of pathogens that join the system each turn is dictated by the sickliness factor of the system, *S*. This value $S = 8$ of an exponential distribution, where 8^{-1} is the mean of the distribution.

Pathogens attack red blood cells in the system, consuming them. When they reduce a red blood cell to 0 effectiveness successfully, it is consumed, and another pathogen is spawned. This new pathogen has the same effectiveness and position as its parent. This process of consumption and replication is meant to represent the process of pathogens consuming resources in an environment and reproducing. With a virus, this may consist of injecting genetic code, using it to produce new viruses, causing it to be mutated and destroyed. A bacterium may consume the cell for energy and reproduce through mitosis.

In addition, pathogens have a probabilistic chance to become 'hungry' each turn. This hunger is simply an abstraction of the drive causing the pathogens to attack red blood cells and reproduce. The higher a cell's virulence (effectiveness), the higher its base chance of hunger. In addition, the number of turns that have passed since it last consumed a target is weighted and used to increase the chance of 'hunger.' Once a pathogen is hungry, it will begin moving at full speed, exploring for food. It will attack the nearest red cell it finds in this way that is not already being eaten. After becoming hungry, a pathogen has *H* turns to consume food and sate itself before starving. This starvation paradigm represents the inherently limited amount of time a pathogen has before it has expended its energy and is either dead or effectively incapable of reproducing.

We find, therefore, that a pathogen's motivations are quite straightforward. It will travel around, exploring. If it becomes hungry, it will attack the first cell it sees, preferring red cells over white. If a white cell attacks it, and it is not eating, it will fight back. It seems that not all pathogens are capable of/designed to attack the immune system (white blood cells), and therefore this option is disabled in most of the tests (pathogens won't consider white blood cells food).

The pathogens in this system have their own emergent 'swarm' behavior to achieve: to consume all other agents in the system. This seems to be obviously in opposition to the emergent behavior desired by the cooperative agents (red and white). This leads us to our first hypothesis:

H1: If unchecked, pathogen growth will model the exponential growth displayed by true pathogen infection patterns.

Pathogens in this system often have a higher initial effectiveness than defenders, but we hypothesize that their capability to replenish their losses will prove to be their most influential trait.

H2: If not allowed to spawn, pathogens will have much more difficulty achieving emergent behavior.

These hypotheses will be addressed later, during the experimental results.

3. EXPERIMENTAL SETUP

The ability to control the environment, pathogens, and white blood agents allows us to simulate different 'situations' modeling different real-world examples. For example, we can run simulations in small area with high agent density or in a larger area with lower agent density (to simulate a condition like anemia). Toggling different variables in these situations allow us to see how our agents respond to different environmental conditions. Furthermore, we can explore which parameters have the most effect in different situations.

Unless otherwise noted, all experiments are run under the following guidelines. Each of these experiments is run 10 times, each time supplying a different (non-timestamp) random seed (if the reader is interested in reproducing our results, they may contact us for the seed strings). Every simulation run is considered over when either all pathogen threats have been neutralized, or all non-pathogens have been consumed. The default discrete grid size is 100 x 100. Sight, movement, and communication ranges will be based on the following multipliers (multiplied against effectiveness): M_M = 1.0 (movement), M_S = 1.5 (sight), M_C = 2.0 (white blood cell communication). Red and white cells will use the following values for fixed effectiveness: E_R = 3.0, E_W = 8.0. The range of pathogen effectiveness will be within the range defined by E_{Pmin} = 3.0, E_{Pmax} = 12.0. Pathogen starvation count, H, is set to 25 ticks.

Most of the hypotheses will be stated in section 4, near the pertinent tests done to prove/disprove them.

3.1 Pathogen Oriented Experiments

These experiments aim to analyze general pathogen behavior and traits. It is important to have a good understanding of the defending agents' opposition before attempting to analyze the system as a whole. This section is devoted to analyzing hypotheses directed

at pathogens alone. The parameters used for these tests are mentioned per data set in the experimental results.

3.2 Combating an Existing Infection

This situation is designed to test our agent's capability of neutralizing an existing pathogenic presence in an area. This type of simulation pits an initial population of each agent type against each other. For these experiments, we set default initial agent populations for red, white, and pathogen agents as follows: $P_R = 500$, $P_W = 50$, $P_P = 75$.

First, a set of experiments were run to explore the effect of altering each of the pertinent input parameters. For ease of analysis and presentation, we only change one of these parameters at a time (except when signified by "parameter$_x$ vs. parameter$_y$"). The parameters tested in this way are as follows (more information in the appropriate sections in experimental results):

1. White blood cell count (P_W)
2. Pathogen count (P_P)
3. Grid size (m, n)
4. Pathogen virulence (E_{Pmax}, E_{Pmin})
5. Communication range multiplier (M_C)
6. Pathogen count vs. white blood cell count (P_P vs. P_W)
7. Pathogen count vs. white blood cell count without pathogen spawn (P_P vs. P_W)

3.3 Defending Against Ongoing Pathogenic Assault

This situation is designed to test the defender agent's capability to adapt, detect, and neutralize pathogens inserted into the system over time. For these runs, default initial populations are fixed: $P_R = 500$, $P_W = 15$, $P_P = 0$. In addition, we set the quantity of pathogens that join the system over time, $J_P = 10$, and the system sickliness factor $S = .3$. As before, the following input parameters were tested for effect:

1. White blood cell count (P_W)
2. Pathogen join count (J_P)
3. Grid size (m, n)
4. Communication range multiplier (M_C)
5. Pathogen join vs. white blood cell count (J_P vs. P_W)
6. Pathogen join vs. white blood cell count without pathogen spawn (J_P vs. P_W)
7. System sickliness (S)

4. EXPERIMENTAL RESULTS

In this section, the results from the experiments mentioned in section 3 are given, and reasonable implications of the data are explored. Instead of listing all expectations and hypotheses centrally, they will be mentioned in each section, near the pertinent results.

4.1 Pathogen Oriented Experimental Results

Here, we're aiming to analyze the pathogen agents' ability to reach their own global coherence. In future experiments, pathogens are often opposed by the entrenched defender agents. Therefore, these runs allow us to analyze how pathogen populations act without the presence of opposition.

4.1.1 Unchecked Pathogen Growth in Finite Resource Environment. Pathogenic population in a finite resource environment follows 4 general stages: lag, log, stationary, and death[1]. In lag, population is essentially stationary, while bacteria mature and adapt to growth conditions. During log, exponential growth occurs while food is available and plentiful. During stationary, population takes on a logarithmic growth pattern as growth slows down due to resource depletion and population density. During death, population decreases rapidly due to lack of resources.

The following figure shows a representative run of the system with no white blood cells. We performed 10 tests and verified that the following is a representative of pathogen behavior in the system.

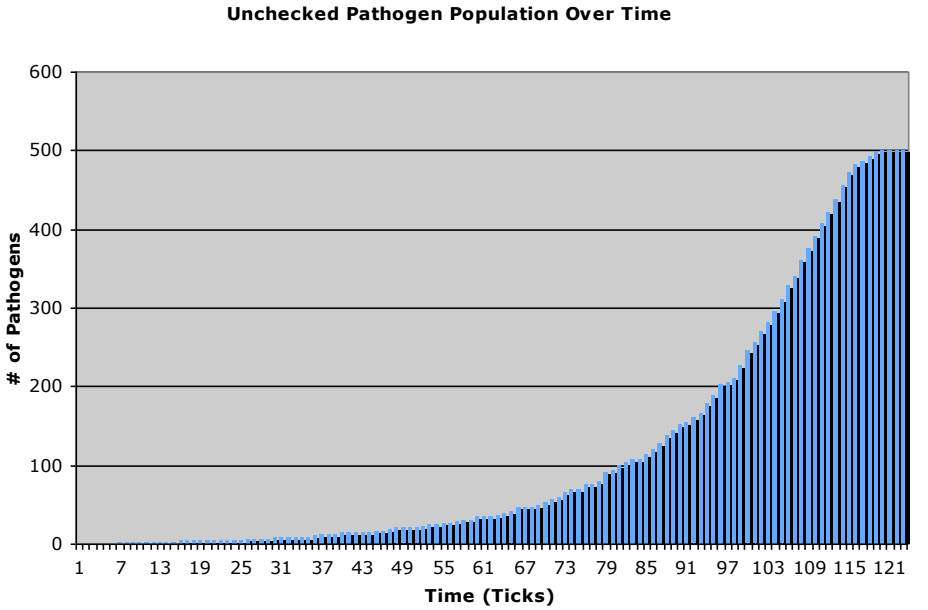

Figure 1. Representative results of unchecked pathogen growth

The pathogen population shown in this figure shows an initial delay, an exponential growth explosion, and a small logarithmic slowdown. These observations match quite closely with the expected lag, log, and stationary phases outlined previously. The lag phase here is created by the ticks required for the pathogens to explore, get hungry, and begin to consume the first red blood cells. The exponential growth experienced here is created by the ticks of consistent binary replication as food is consumed. As food becomes scarcer, this growth slows to a logarithmic curve asymptotically bounded by $P_P + P_R$. This makes sense, as binary replications are limited by food in the system (as we allow replications only at time of consumption).

However, we do not see the flat expanse expected in the stationary phase and the sharp drop of the death phase. This is because our system simulations stop when all non-pathogens were consumed. Of course, if all non-pathogens were consumed, we should expect the population to remain even or drop slightly for up to H ticks. After H ticks, some of the hungry pathogens will begin to starve, and we should see a sharp drop in population as pathogens die (less virulent agents will last longer since it takes more ticks on average to become hungry).

By these results, it would seem that H1 is reinforced. For a more conclusive set of results, the system should be altered to allow executions to continue until all pathogens are dead.

Interestingly, this shows the difficulty presented to our defender agents and real world white blood cells. The body's capability to produce cells and the white blood cells' capability to kill invaders must be linearly, or at best, polynomially bounded. In a closed system like this, with no white cell reinforcements, we hypothesize that overall capacity for white blood cells is more limited than that of pathogens. Therefore, it seems reasonable to imagine that the combined ability of the defender agents to affect system health is linearly bounded over white blood cell population. If this were true, changes to environmental and pathogen parameters would tweak the exponential curve of pathogen potential, and changes to the environmental and defender agents' parameters would alter the linear plot of defense potential. This leads to the following hypothesis:

H3: Due to exponential nature of pathogens and linear nature of white blood cells, any changes to parameters should cause a clean switch from total system failure (pathogens win) and complete system health (pathogens completely defeated).

If true, this may shine some light on the issue concerning H2, as pathogens that cannot spawn will most likely drop to a linear potential.

4.2 Combating an Existing Infection

As mentioned before, the difficulty in this situation is destroying or neutralizing some initial population of randomly located targets in the system, before their growth makes this impossible. By what we saw in H1, this could add difficulty by circumventing some of the lag phase.

Even so, it is our hypothesis that the defender agents' capacity for simple communication and teamwork makes them more capable than an equivalent number of pathogens. We formalize this with the following:

H4: Defender agents are capable of defeating an equal number of pathogens

4.2.1 White blood cell count (P_W) results. In this case, the white blood cell counts will be varied while holding all other parameters constant. We predict that the higher the saturation of white blood cells in the environment, the more capable the system will be of fighting off an infection. In addition, the more defenders present, the more defenders will be available to team up against pathogens, destroying them quicker, and leading to quicker convergence.

H5: Higher saturations of white blood cells will lead to quicker convergence and higher system utility.

The following figure shows the results of the experiment with initial white blood cell counts ranging from 10 to 100 in increments of 10.

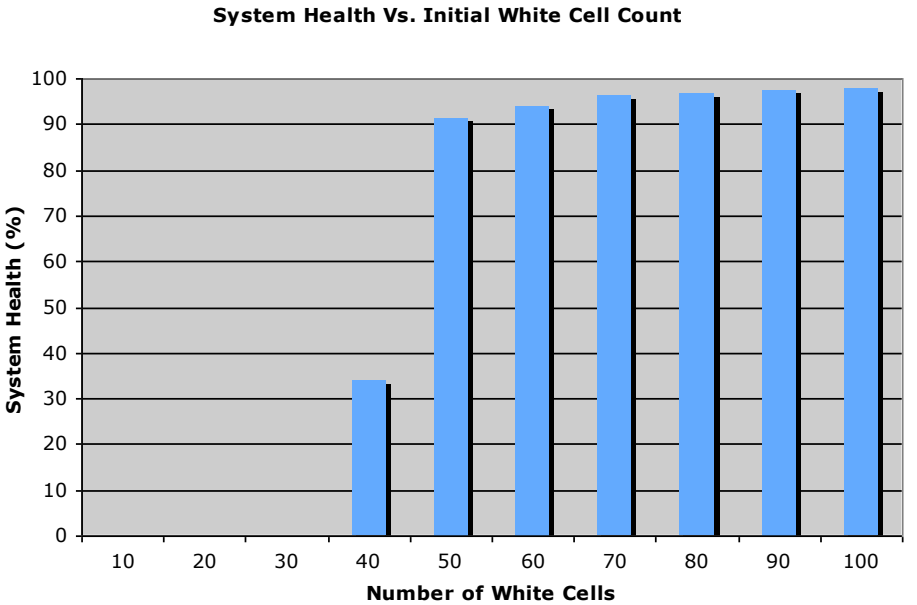

Figure 2. Average final system health vs. initial white blood cell count

The results shown in this graph seem to corroborate H5, as we have a sharp contrast on average final system health values. The graph shows a strong correspondence between initial white cells and total utility. Although $P_W = 40$ is a partial transition, the switch from 0% to near 100% system health is still quite sharp, which seems to go along with H3. It seems that there is a threshold for P_W approximately at 40. This is interesting, since $P_P = 75$. This shows that our pathogens are capable of defeating a significantly larger quantity of pathogens, even though E_{Pmax} is markedly higher than E_W. Though E_{Pmin} is lower than E_W, the more virulent pathogens reproduce more and are more significant than those of lower virulence. These results reinforce H4, but more tests are needed to be conclusive, as there are many more variables to examine.

4.2.2 Pathogen count (P_P) results. Similar to the last experiment, we vary the initial pathogen count, P_P. The last experiment showed some interesting results at $P_P = 75$, but more tests will help to see how this parameter affects the system. We held the following expectation for this experiment:

H6: Higher pathogen population will result in lower overall system health.

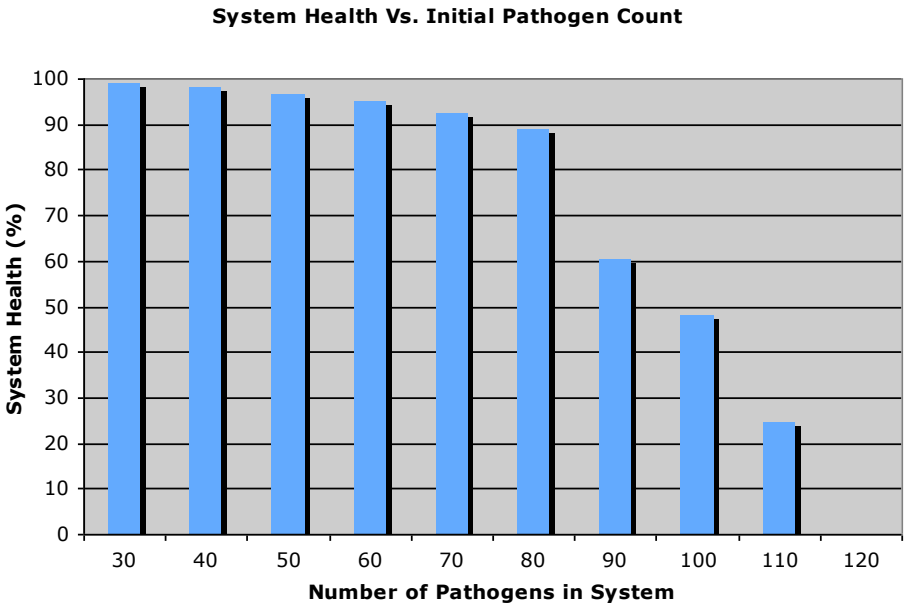

Figure 3. Average final system health vs. initial pathogen count

These results show a definite correlation between initial pathogen counts and the ability for a fixed number of white blood cells to mount a resistance. The trend shows what we expected, more pathogens results in a markedly lower performance. This seems to reinforce H6. We see a steep drop in system effectiveness over 80. If we remember that there are 50 white blood cells in this system, this is still fairly impressive (further reinforcing H4).

The interesting result here is that the switch from complete success to complete failure isn't as sharp as predicted by H3. Except for $P_P = 90$, the trend resembles an exponential decay, which is still considerably pronounced. By these observations, it could be put forward that another run between 70 and 120 with smaller increments could help show the curve more accurately.

4.2.3 Grid size (m, n) results. So far, all runs have been done in a 100 x 100 grid. In such a situation, healthy white blood cells could move in a circle of radius 8, see a radius of 12, and detect pheromones at a radius of 16. In this experiment, we change the width and height of the area over the range [25, 300] in increments of 25.

In order to properly explore H5 (and the closely related H8), we must not only increase the quantity of white blood cells in the environment, but also check the effect of a larger environment on the same populations. Since we are not varying any populations,

the ratio of white to red blood cells and white cells to pathogens remains constant, and it is the effect of density that will be in question. By H5, we would predict that higher grid sizes will result in poorer performance. This is formalized in H8: if the amount of white cells per square unit drops below some threshold *T*, system performance will degrade quickly.

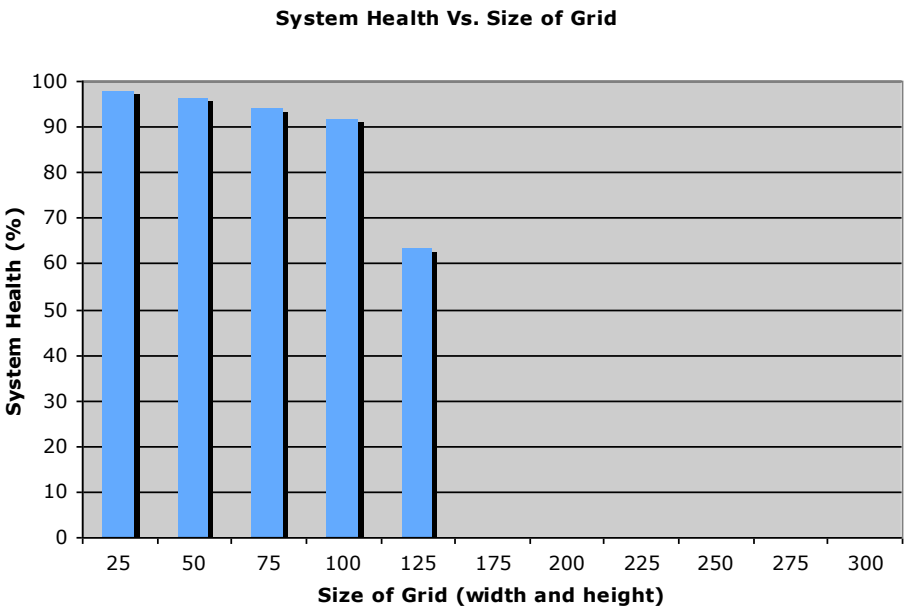

Figure 4. Average final system health vs. grid size

There is a definite correlation in these results, and it seems that as grid side length increases linearly at rate r (and grid area increases at r^2), system health drops exponentially. In fact, at $r = 125$ we find a threshold, $T = .0032$ whites/unit^2, over which our system completely fails. To explore why this is, we must realize that the area any singular defender can explore at a time is fixed, and its significance (as a ratio to total area) is decreased exponentially. Since red blood cell population is so high, pathogens do not have much trouble finding food and reproducing. However, white blood cells have a much lower chance of encountering the randomly located pathogens.

These results show that white blood cell population is very important to the success of the system, reinforcing H5. Under a certain density threshold, the system collapses into one in which pathogens are essentially unchecked. More interestingly, these results insinuate that there is some ratio of white cell density to pathogen density that acts as a threshold to system success, reinforcing the related hypothesis H8.

4.2.4 Pathogen virulence (E_{Pmax}, E_{Pmin}) results. This experiment aims to explore the effect of different pathogen effectiveness values in the system. In the experiment, both the min and max effectiveness for pathogens are varied together on the range [6, 15] at intervals of 1.

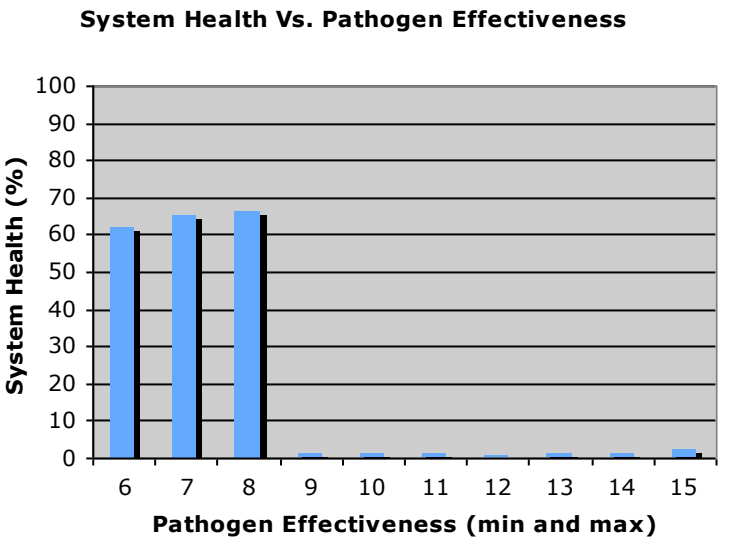

Figure 5. Average final system health vs. pathogen effectiveness

The results show the pathogens incapable of convergent behavior when E_{Pmax}, E_{Pmin} are greater than 8. The reason for this is likely the fact that all white blood cells have an effectiveness of 8. So when all pathogens' effectiveness is above this value, battle becomes more deadly (since pathogens outnumber defenders). An interesting result here is that the defenders are capable of defeating a greater quantity of equally powerful pathogens, even with pathogen spawning enabled. This supports H4, likely due to the greater intelligence and communication capabilities of the defenders.

At an effectiveness value of 9 or greater, we see a complete, consistent system failure. When pathogens are more effective than the defenders in the entrenched infection scenario, the cooperative agents are incapable of achieving global coherence. This goes along with H9, giving a threshold of 0 in this scenario.

4.2.5 Communication range multiplier (M_C) results. The results of the experiment in 4.2.3 showed that as the grid grows larger, the defender agents have a harder time finding

and neutralizing pathogens in the time it takes the pathogen population to balloon out of control. In response, we formulated the following hypothesis:

H7: Higher communication range multipliers will allow the defenders to more ably defeat sparse pathogen threats.

This hypothesis comes from the fact that as the space increased, the white blood cells were less able to form a cooperative network and deal with the threats. In this following test, we alter communication range multiplier over the range [1, 10] by increments of 1. Therefore a healthy white blood cell can detect pheromone releases from other white blood cells at a maximum radius of [8, 80] respectively. This means a transition from approximately 2% to nearly 100% communication range coverage in a 100 x 100 space.

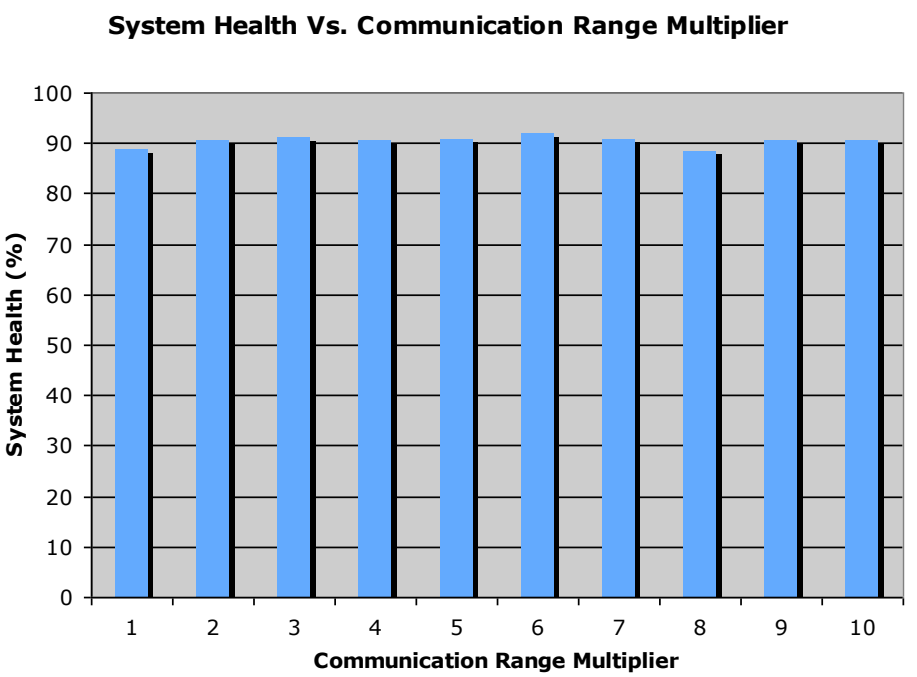

Figure 6. Average final system health vs. communication range

These results show almost no correlation between communication range and system effectiveness. This seems to counter our hypothesis, H7. Why didn't changing communication range affect the agents' performance? The first answer is that with 75 pathogens and 50 defenders, they are already reaching an asymptotically maximum performance, leaving little room for improvement.

The other possible explanation for these results is due to the nature of the environment. Since all agents are inserted randomly into the environment using a

uniform distribution, they are initially spread out. Since all pathogens are inserted as an initial population, the defenders do not need to communicate much to find them. Therefore, a different outcome may result if the model was redesigned to assign starting locations in a different way.

The first explanation could be tested by running a similar experiment with less defenders, or stronger pathogens. Therefore, such a test was formalized and conducted utilizing stronger pathogens (minimum pathogen effectiveness was increased), and the results are displayed below.

4.2.6 Pathogen count vs. white blood cell count (P_P vs. P_W) results. This experiment contrasts multiple combinations of initial pathogen and white blood cell counts to test the validity of the following:

1. H4 (defenders can defeat an equal number of pathogens)
2. H5 (higher defender count leads to higher system performance)
3. H8 (there is some threshold ratio of pathogen and defender densities to system success).
3. The implications of 4.2.3 (there exists a threshold on the ratio of defender density vs. pathogen density that strongly contributes to the coherent behavior)

In this experiment, we run 10 tests at each pair (P_P, P_W) on [10, 150] x [10, 150]. The average system health percentage resulting from each combination is displayed as a result on the z axis of the following graph.

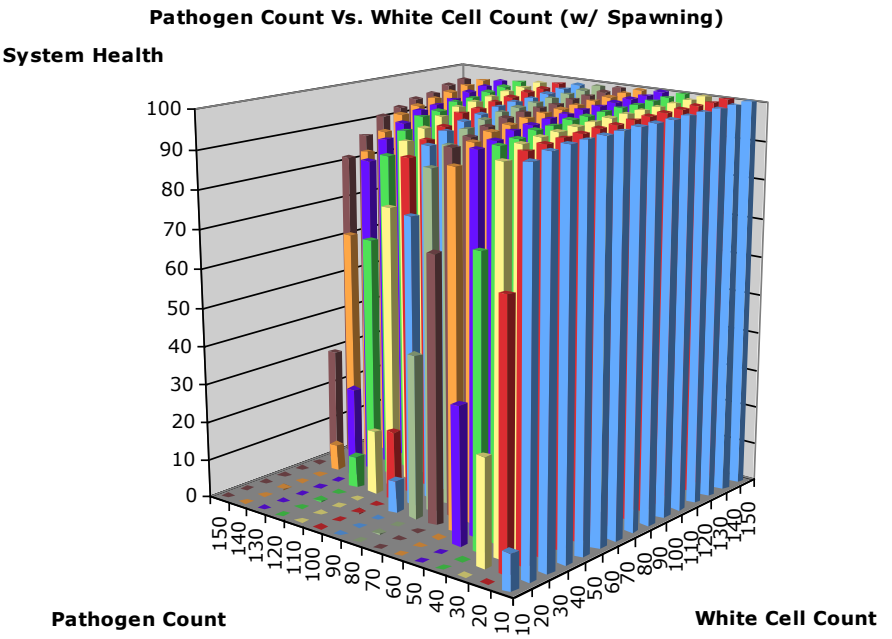

Figure 7. Average final system health over (P_P *vs.* P_W)

In this graph, we see a strong trend between P_P and P_W. Overall system health is lowest when white blood cell count is low and pathogen count is high. There is a sharp transition from success to failure as white cells are decreased, resulting in an exponential drop. This quick switch further affirms H3 (sharp transition hypothesis).

We see these trends because, keeping grid size fixed at 100 x 100, the white cell density contributes greatly to whether the system can react to the present pathogens before they grow too numerous. However, the data shows a greater weight applied to white cell density than to pathogen density.

To expand on this and test the implications of the experiments in 4.2.3, the following graph shows the minimum defender / pathogen ratio at which average system health was over 70%.

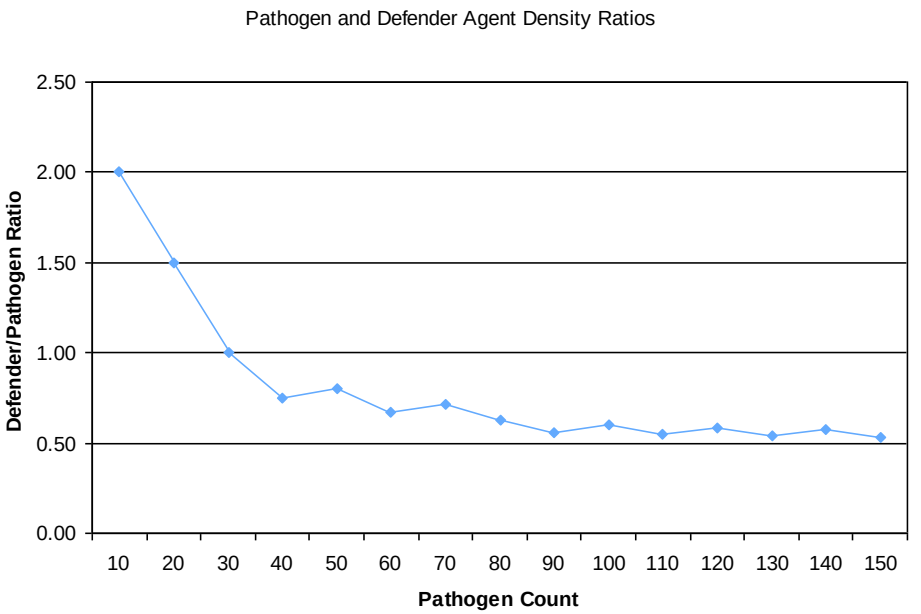

Figure 8. Defender to pathogen ratio when average system health became >= 70%

Despite the predictions of 4.2.3, there doesn't seem to be a constant density ratio threshold between pathogens and defenders; at least initially. For greater pathogen populations, the ratio seems to settle around .55 or so. This means that in 100 x 100 grid, the defenders can achieve an average system health of 70% or more with P_W only 55% of P_P. This definitely reinforces hypothesis H4. Even with a P_P of 30 or less, the white cells could still defeat the pathogens with this ratio, but not with a system health of 70% or higher.

4.2.7 Pathogen count vs. white blood cell count results without pathogen spawn (P_P vs. P_W). To continue the interesting results procured from the last experiment, another set of experiments were run pitting pathogen and white cell populations. In this set of experiments though, pathogens were not allowed to spawn. H2 predicts that the pathogens will be even less capable of achieving dominance. Subsequently, we should see a stronger dominance and quicker convergence for white blood cell global coherence.

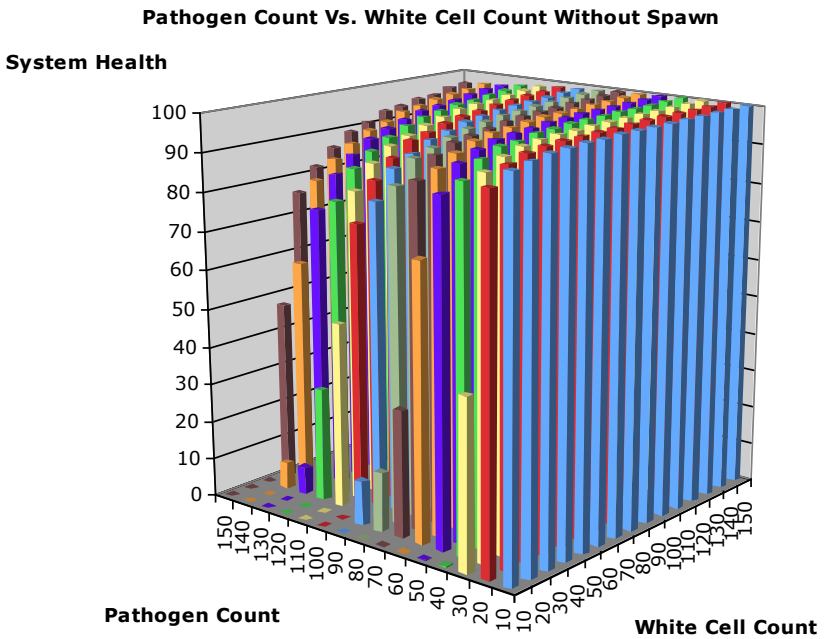

Figure 9. Average final system health over (P_P *vs.* P_W) w/o pathogen spawn

These results are still mostly positive for the cooperative agents. A similar trend is displayed in the data as well. However, the range of cases in which the pathogens caused cascading failure in the system was greatly reduced. The lack of pathogen spawn is the reason for the difference in performance in these cases vs. those in 4.2.5. This is due to the reduction in cases in which a single virulent pathogen escaped notice of the defenders and rapidly dominated the system. For reference, the following graph shows the population ratios for this experiment in which average system health was >= 70%.

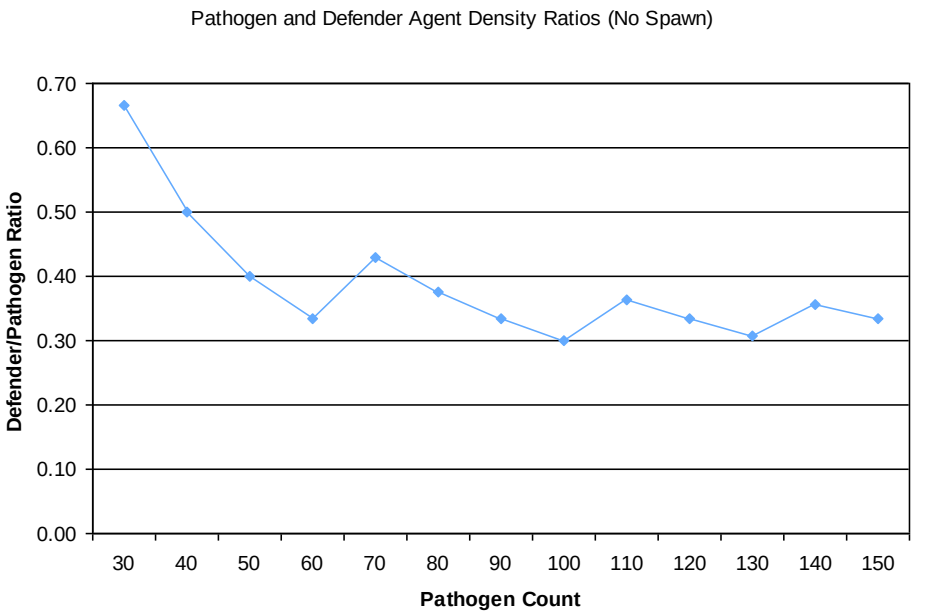

Figure 10. Defender to pathogen ratio when average system health became >= 70% (no spawn)

Data points at pathogen population of 10 and 20 were excluded because the threshold ratio could not be guessed from the data (since there was no case in which system health was < 70% for these values). However, we see the same general trend in the data. However, this time the ratio stabilizes closer to .33. This insinuates that in these experiments, good results could be achieved even with pathogens outnumbering defenders 3 to 1. This result further reinforces H4. In addition, the decrease in pathogen effectiveness in the between the results of 4.2.5 and 4.2.6 reinforces the hypothesis about the significance of pathogen spawning (H2).

4.3 Defending Against Ongoing Pathogenic Assault

As mentioned in the experimental setup, these experiments test the agents' behavior and performance when pathogens enter the system over time rather than being present initially. These experiments present a different challenge to the agents on defense. Rather than coordinating to neutralize existing pathogens, the challenge of the environment switches to quickly detecting and finding newcomers. To emphasize this, the white blood cell population was decreased for these tests.

4.3.1 White blood cell count (P_W) results. Similar to the entrenched illness test 4.2.1, this experiment alters initial white blood cell population and examines the results. White blood cell populations are tested over the range [5, 50] at increments of 5. These results should allow further inspection of H5 in a different scenario. Our hypothesis predicts higher final utility values with more white cells.

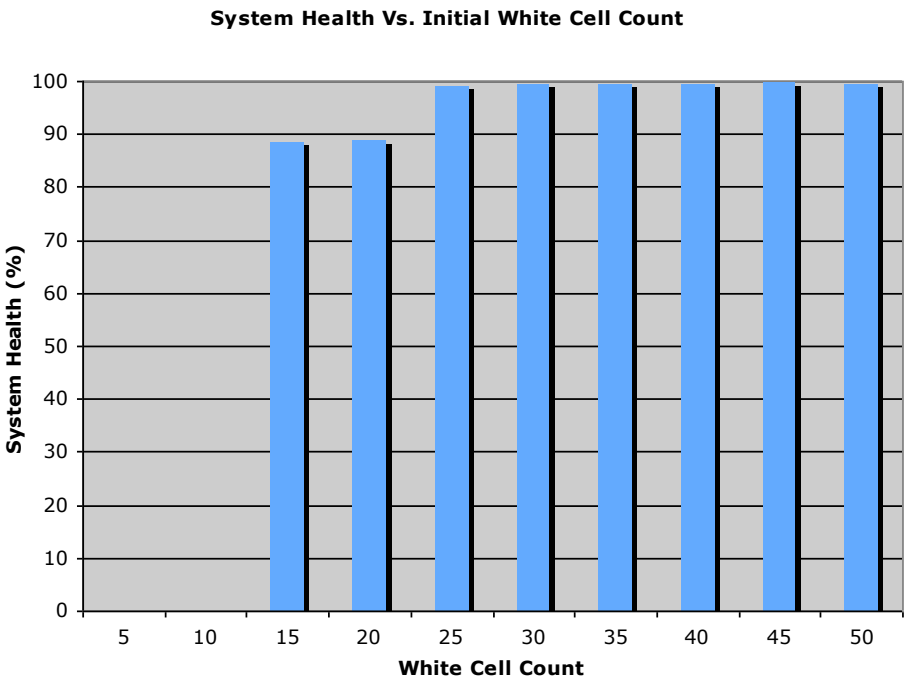

Figure 11. Average final system health vs. initial white cell count in invasion scenario

These results show a strong correlation. Below the threshold population of 15, the defenders were never capable of successfully defending the system. At 15 and above, they were much more capable. The sharp transition implies a threshold defender population density required to effectively watch the entire grid for threat emergence. This seems to strike at the heart of the invasion issue, which is that the correlation between high effectiveness and strong exploration capabilities is stronger than in previous tests.

These results reinforce H5, as we can see that white cell count contributes significantly to system health in the invasion scenario as well. In addition, this sharp transition threshold further reinforces H3.

4.3.2 Pathogen join count (J_P) results. Analogous to the previous test, this experiment analyzes the effect of varying the amount of pathogens that join the system over time. Pathogens joining the system, J_P, will be tested over the range [5, 50] at increments of 5. These results should test the validity of H6 in the invasion scenario.

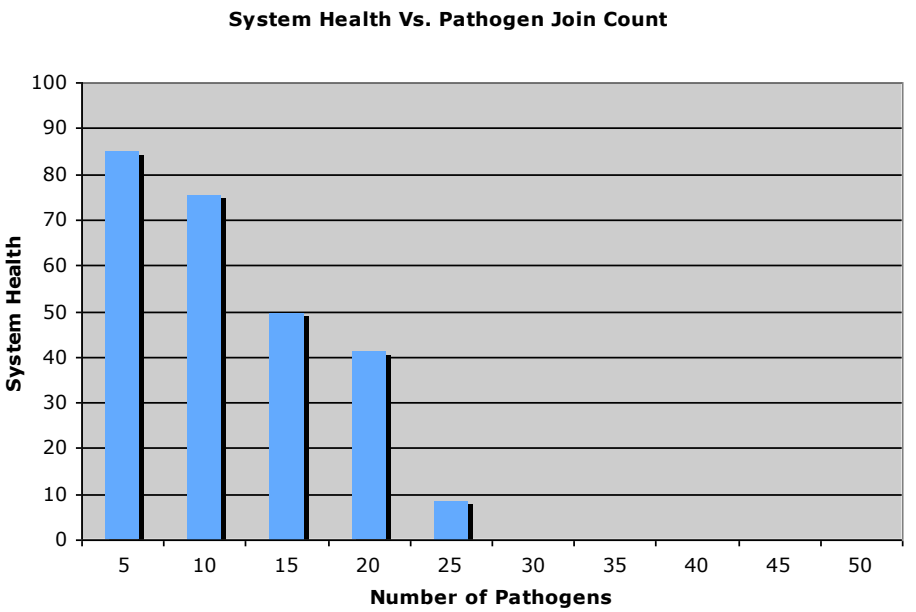

Figure 12. Average final system health vs. pathogen join count in invasion scenario

The results show a strong linear correlation between pathogen join count and the overall system health that the cooperative agents were able to protect. As more pathogens join over time, system health decreases. These results match the predictions laid out by hypothesis H6. However, no relation has been so linear thus far in the experiments, and the lack of a sharp transition casts doubt on hypothesis H3.

These results put forward the question, "Why is system performance degradation linear rather than exponential?" Unlike the existing infection scenarios, a higher pathogen count means more pathogen insertions over time, rather than a stronger force. As defenders detect and neutralize these threats, they naturally move about, bunching up and spreading out. This can result in areas being unwatched and undefended. Therefore, a higher J_P simply implies more chances for a pathogen to be inserted into an area not currently under surveillance by the defenders.

Therefore, these results further the significance of the observations made in the entrenched infection tests concerning the effect of pathogen density on system

performance. The next experiment tests the other side of this question, the effect of a varying grid size.

4.3.3 Grid size (m, n) results. Similar to the experiment of 4.2.3, the following test is conducted to examine the effect of different grid sizes on system effectiveness. Grid width and height is varied over the range [25, 300] by increments of 25. The results so far would predict higher grid size causing exponential decrease of system health.

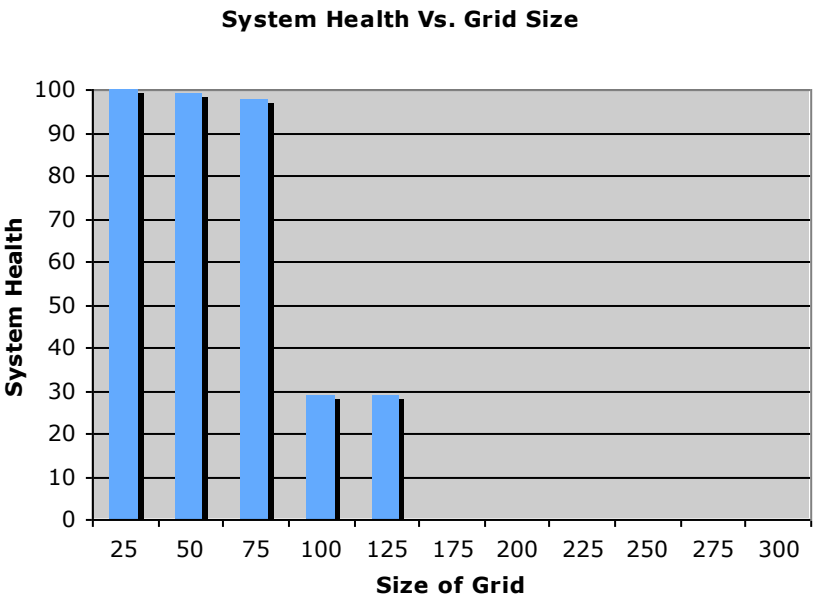

Figure 13. Average final system health vs. grid width and height in invasion scenario

These results show a definite correlation, if not quite exponential, between grid size increase and a decrease in defense capacity. It is difficult to say what type of correlation is present due to lack of measurements on the range (75, 175). Therefore, another experiment was conducted with 20 samples at each value across the range [75, 175] at an increment of 10.

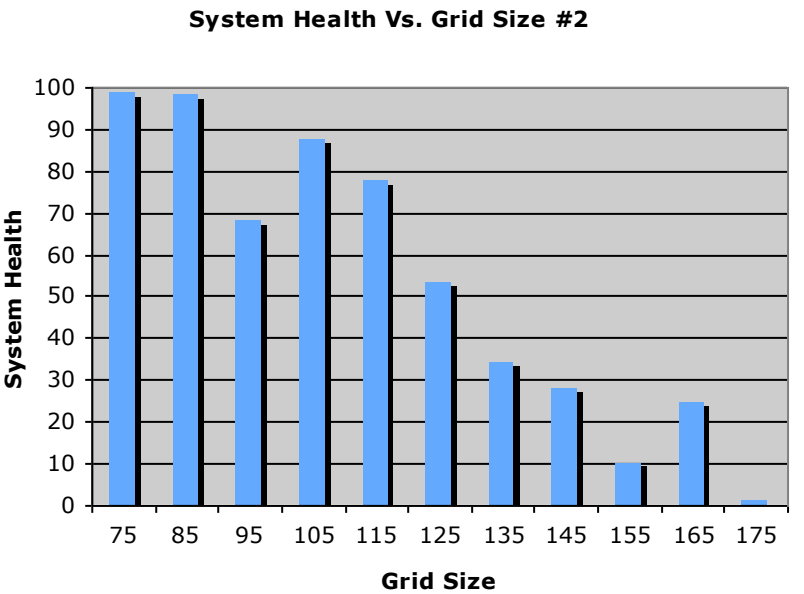

Figure 14. More specific run grid size impact analysis

This second graph shows a tighter analysis of the area of importance, with a higher run count (20 runs per value) to ensure accuracy. The trend is a lot more pronounced in the graph, yet still not very exponential. From both graphs, the results shown here lend further credence to the importance of white cell density. Therefore, we consider this more positive evidence of hypotheses H5 and H8.

4.3.4 Maximum pathogen virulence (E_{Pmax}) results. This experiment aims to test the effect of higher virulence pathogens on the invasion scenario. By raising the maximum pathogen virulence, we increase the chances of higher virulence pathogens appearing in the environment. This is designed to represent a body in a more dangerous environment (malnutrition, lack of medical attention, or exposure to more harmful contagions). This experiment varies E_{Pmax} over the range [8, 20] by increments of 2.

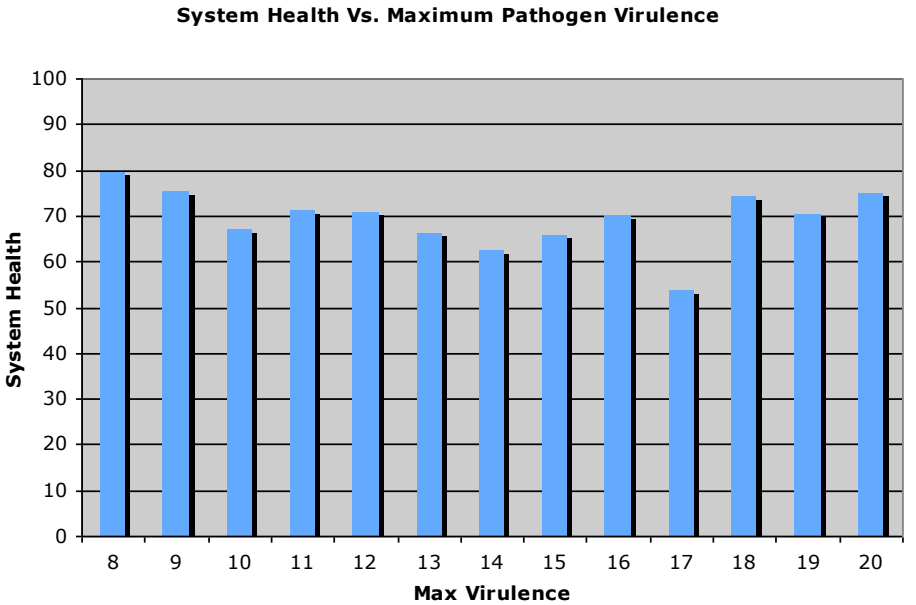

Figure 15. Average final system health vs. maximum pathogen virulence

The results of this experiment show little to no correlation between maximum pathogen effectiveness and system health. This is contrary to our predictions and intuitions. H9 would have suggested that at a certain point, average system health cascades quickly to 0. Why do the results show otherwise?

One explanation for this discrepancy is the fact that the default pathogen join count is only 10. New pathogens have their effectiveness sampled uniformly from 3 (default minimum) to the value being tested above. These parameters allow very few samples on each run, making these results more heavily influenced by pathogen join than maximum pathogen virulence. This issue could be addressed by greatly increasing the number of runs done for each value, or by increasing the default pathogen join count. As it stands, we can not draw any good conclusions on H9.

4.3.5 Communication range multiplier (M_C) results. The following experiment was designed to give communication range another try, this time in the invasion scenario. Since this scenario is inherently more sparsely populated, predictions hold that higher communication detection range will allow for quicker convergence. In addition, the invasion scenario is more about quickly adapting, while the entrenched scenario is more

about finding and neutralizing nearby threats. This could be abstracted to the difference between a need to solve global goals vs. a need to solve more local goals.

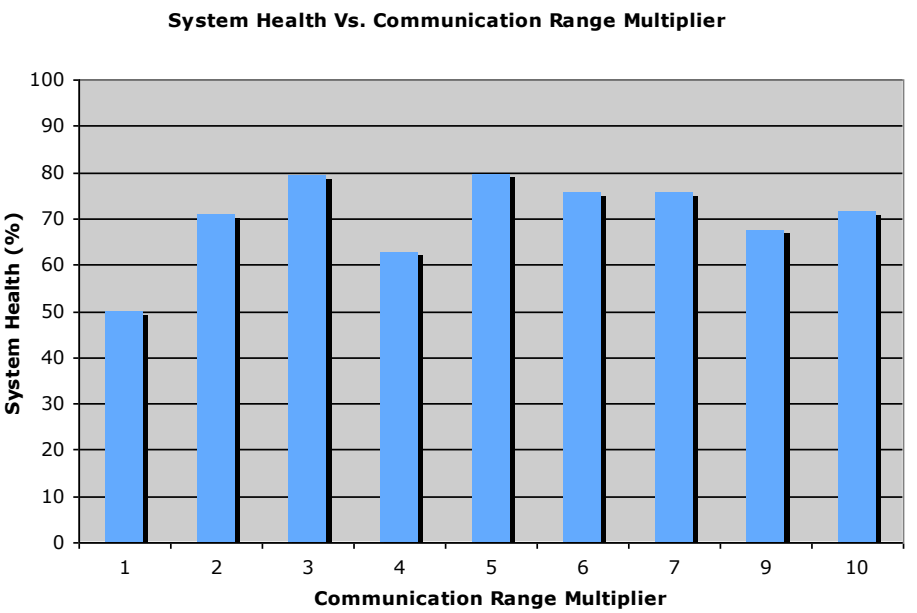

Figure 16. Average final system health vs. defender communication range multiplier (invasion scenario)

These results show very little correlation between communication range and system effectiveness. We see a sharp increase from M_C = 1 to M_C = 2, but most of the other results are inconsistent at best, and suggest a decrease in effectiveness at worst. This seems to counter our hypothesis, H7. But why are the results like this?

At extremely low multipliers (such as 0 or 1), the defenders do not gain any communication advantage over the other agents. Their decisions become based solely on their sight range, and the previously somewhat sure H4 does not hold (white blood cells may not be able to effectively defeat an equal number of pathogens). But why does a higher communication range not help the defenders? Upon visually inspecting some runs with sparse pathogens and high communication range multipliers, we noticed an interesting issue. The indirect communication protocol allows a defender to call others to his aid, but since there is no global control, no mechanism is present to organize the redistribution of agents. The effect of this is that agents would 'bunch up' to fight a pathogen. When it was defeated, they could not effectively redistribute themselves over the area.

The implications from this are serious. H7 should be true, but it is not. The more information an agent has about his surroundings, the better decision he should be able to make. These results show room for a redesign of how our defender agents utilize their communication data. For example, agents could emit a no-activity pheromone that would repulse defenders. The agents could consider all the positive and negative motivations to go in any direction, and base their move vector on the direction they deem to be 'best' from this calculation.

4.3.6 Pathogen join vs. white blood cell count (J_P vs. P_W) results. The following test is analogous to the one performed in 4.2.6, and aims to explore the effect of altering both defender and invader quantities on system health. Again, both values are varied on the range [10, 150] by increments of 10, and system health is displayed on the z axis.

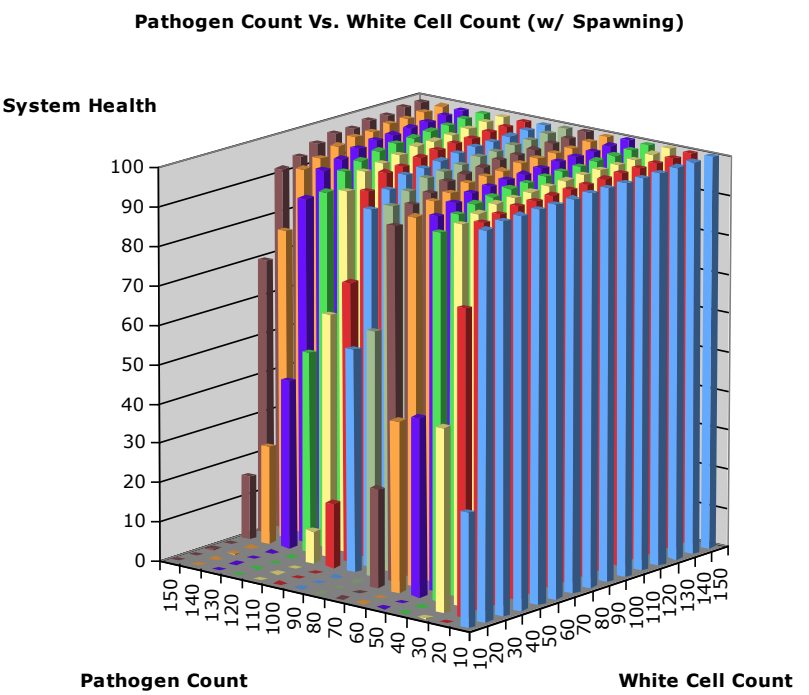

Figure 17. Average final system health as a function of defender population and pathogen join count

This graph shows a relationship between defender and invader populations similar to that observed in the entrenched scenario. However, in this scenario, several of the high P_P pathogen, low P_W pairs that led to failure in the entrenched scenario succeeded in the invasion scenario. In fact, emergent behavior is somewhat more successful across the board in this version of the experiment.

This trend is due to the fact that fewer pathogens must be dealt with at any one time in the invasion scenario. Though they must be found, these pathogens must face a greater number of white blood cells each (on average) than in the entrenched scenario. Furthermore, these results positively reinforce hypothesis H4, as the defenders were consistently able to beat higher quantities of enemies.

4.3.7 Pathogen join vs. white blood cell count without pathogen spawn (J_P vs. P_W) results. Just as in the entrenched scenario, this test analyzes the relative affect of the pathogen spawn capability. The following test is the same as the one performed for 4.3.6, but with pathogen spawning disabled.

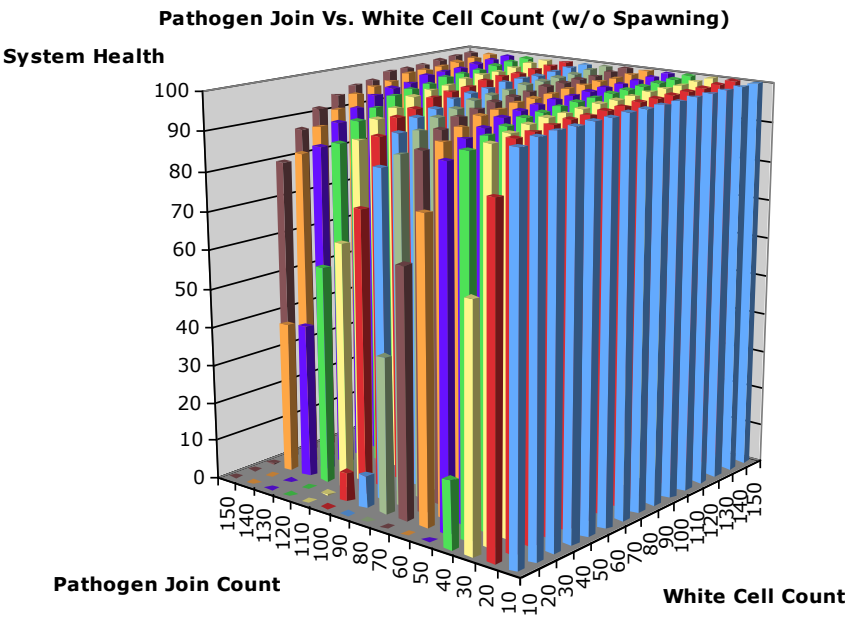

Figure 18. Average final system health as a function of defender population and pathogen join count with pathogen spawning disabled.

This graph shows the same trends and exponential drop-offs as the results in 4.3.6, but the defenders are able to achieve global coherence an even greater portion of the time. This is evidence that without spawning, the defenders are capable of fighting off potential infection with even lower populations than before. Most directly, this confirms H2, as we can see that spawning has an effect on pathogen capability in the invasion scenario as well. Though we don't analyze the population density ratio threshold value, *T*, for H8

(like in 4.2.6 and 4.2.7), the data closely resembles those results. This signifies that the results from this experiment and 4.3.6 also reinforce H3 through H6 and H8.

4.3.8 System sickliness (S) results. This experiment examines the effect of the pathogen insertion rate. As mentioned before, the system sickliness is the value 8 of an exponential distribution function, such that 8^{-1} is the mean of this distribution. *S* is varied across the range [.1, .95] at an increment of .05. From previous observations, our hypotheses predict a lower value of *S* will result in more invaders in a shorter period of time, and subsequently a lower average system health value. In this experiment, 0 can be thought of as all agents arriving at tick 0, equivalent to the entrenched illness scenario.

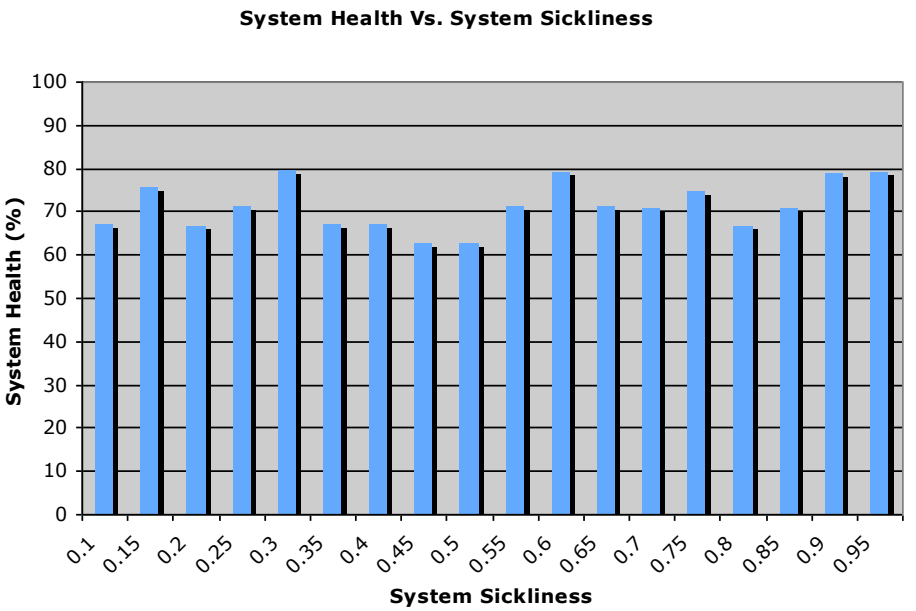

Figure 19. Average final system health vs. system sickliness

This graph shows little to no correlation between system sickliness and average system health. Though contrary to intuition, these results are likely due to the low default pathogen join value of 10. Since there are so few agents to enter the system, the rate at which they enter is underemphasized by the relatively short time over which they join and live. In addition, more of a trend may be evident if results were collected at higher values (at which point pathogens begin to join more slowly).

5. CONCLUSIONS

Throughout the experimentation process, both surprising and expected results were encountered, and many possible points for improvement were identified. Despite this, the model performed as expected in general, and provides a great staring point for future work. The following section exposes the overall outcome of the proposed hypotheses. Section 5.2 analyzes the overall success of the model and implementation, and section 5.3 presents directions this model could be taken in future applications.

5.1 Analysis of Hypotheses

The experimental data presented in this work lends great insight into the inner functionality of the system and the sensitive relationships between the different types of agents. Not only have these experiments explored the (sometimes unexpected) effects of the environmental parameters, but also presented great opportunities for furthering the field of medical technology.

Most of the initial hypotheses were affirmed, with a few exceptions. The following summary displays the general outcome of each hypothesis

5.1.1 H1 (accurate exponential pathogen growth patterns). This hypothesis was designed to judge the success of the pathogen agents' capability to achieve their own type of swarm behavior, or emergent behavior. The unchecked pathogen behavior in our model seemed to quite accurately resemble the lag, log, stationary, death sequence seen in actual microbial growth[1]. In summary, this hypothesis was strongly affirmed in section 4.1.

5.1.2 H2 (pathogen spawning is significant to pathogen swarm behavior). Through several extensive experiments run with pathogen spawning both enabled and disabled, it became evident that pathogen spawning was necessary for the most effective pathogen behavior. This hypothesis was also shown to be true.

5.1.3 H3 (the exponential nature of pathogen potential should give rise to sharp and noticeable thresholds in the system). This hypothesis was based on the expectation that the system parameters might alter the linear and exponential nature of these two agent types, but that they would remain on different orders of potential effectiveness. The expectation is that the intersection of two such potentials should be quite sharp.

In different result sets, this hypothesis met with mixed results. However, in most cases where the relationship between pathogen and defender effectiveness was being

altered directly, a noticeable, exponential transition was often present. In such cases, this hypothesis seemed consistently affirmed.

5.1.4 H4 (defenders can defeat pathogens of equal number). This hypothesis was confirmed in both of the scenarios tested here. In both fighting an entrenched infection and defending against pathogen invaders, the defender agents were able to adapt and fight off threats greatly outnumbering them.

5.1.5 H5 (more white blood cells leads to quicker convergence and higher system utility). This hypothesis was based on the simple expectation that more defenders relates to a stronger defense (and the inverse as well). The experiments affirmed this expectation and this hypothesis was shown to be true in almost all cases.

5.1.6 H6 (more pathogens leads to lower overall system health). This is a logical extension of H5, and also its logical complement. In general, this intuitive hypothesis proved to be true. As pathogens populations were increased, system health often decreased sharply.

5.1.7 H7 (higher communication ranges will facilitate more rapid convergence). This hypothesis was a logical extension of any locally minded agent's desire for more information. With a higher pheromone detection range, quicker convergence is expected (as communication approaches global). However, this hypothesis was not affirmed in the experiments conducted.

This work has presented some justifications for this behavior and subsequent changes in agent logic to change it. Essentially, white blood cell logic and communication protocols should be modified such that these agents may more fully capitalize on their communication capabilities.

5.1.8 H8 (there is a threshold in the ratio between pathogen and defender densities over which emergent behavior is possible). Experimental results generally showed that defender density was often related to emergent behavior. Though results were not completely conclusive, for a specific environment and high enough pathogen densities, such a threshold value could be found. This particular hypothesis may require more exploration before it can be said to be conclusively proven.

5.1.9 H9 (when $E_P - E_W > T$, for some threshold T, emergent behavior will fail). This hypothesis is a combination and extension of H5 and H6. This hypothesis was affirmed when experimental results were pertinent, though this paper was not always able to find said value for *T*.

5.2 Overall System Analysis

The goal of a MAS like this one is to evolve emergent behavior using the most local agent decision processes as possible. The environment we implemented was continuous, non-episodic, deterministic, inaccessible, and partially dynamic (the environment doesn't change while an agent is thinking). The inaccessibility of the system limits the knowledge an agent is capable of employing in its decision making process. This is implemented in the system as a limit to the range an agent can 'see'. Agents in this model can only gather data from their immediate proximity, and their reasoning encompasses nothing other than the qualities of their immediate surroundings. Agents cannot model much historical information about the space or other agents in it due to the ever changing nature of the environment. Every tick, agents base their next action only off the current and immediate surroundings (what they can sense). Since agent decisions in this model do not take any global knowledge into account, their reasoning can be said to be local.

In addition, agents in this system never accept more than notifications from the system or other agents; showing that they do not submit to any global control (except for the cases of agent creation and deletion). Due to the lack of any central control in this system, this model and its agents are inherently distributed. Since no agent is allowed to directly control another in this model (except for constraints they place on each other such as during a fight), it can be seen that the agent reasoning and behavior is completely autonomous.

So far, it has been shown that this model fits the requirements of good MAS, and the system hypotheses have been generally affirmed. However, this work has not yet compared the models herein to their real world analogs. [1] generally shows our pathogenic growth model to be accurate, though the transition from ticks to an actual time may require additional work.

The environment, red blood cells, and white blood cells, however, were all quite primitive in comparison to their analogs. For the sake of this paper, we have made several assumptions about the environment that are too strong for general application. If the aim of this system is possible application on nano-machines in the human

bloodstream, the environmental model must be influenced much more heavily by current understanding of microbiology.

In addition, several other considerations bar this model from implementation.

1. Even at the lowest densities, our defender agents were always granted a higher population density than that observed in the human body. In humans, a standard red blood cell to white blood cell ratio is 715:1[2] on average. In the sparsest tests, however, our system assumed at most 50:1 (and had spotty performance in these fringe cases as well).
2. This current formulation of the blood model assumes that the blood environment is closed (except for incoming pathogens). However, the true blood environment model is much more complex than the one presented here, and involves any number of outside factors, and the body's capacity to replenish depleted cells.
3. In this model, white blood cells are considered to be homogenous in capability and function. However, [2] shows that there are different types of white blood cells in the human body in different proportions, designed for different roles. Undoubtedly it is the complexity of these relationships that allows the human immune system to defeat threats so ably.

The model may be a successful MAS that meets its initial hypotheses, but considering these non-trivial issues, the system still has ground to cover before it is ready for general application. The following section analyzes some possible improvements to the model.

5.3 Future work

There are several general improvements to this model that we identified during analysis and testing. In addition to those deficiencies prohibiting the model from implementation, here are some directions that this work could be taken in.

5.3.1 Subtask assignment. In this change to the model, we introduce a utility-based improvement utilizing weighted subtasks. In this modified model, any red (and possibly white) blood cell can ask the model for a new subtask if it has none. This task would consist of a task completion reward and specifications. For red blood cells this would probably consist of going to a particular location, abstracting an oxygen delivery.

This is an improvement over the current model because it would give red blood cells a bigger effect in system performance and more realistic purpose and behavior.

5.3.2 New variable: consumption time. In this change, we implement a system parameter called consumption time. This value would represent the amount of ticks required for a microbe to consume another after it has been incapacitated. Implementing this feature could allow the user to scale the general amount of time required to consume vs. move, etc.

5.3.3 Fuzzy communication. The current implementation of this model assumes perfect agent perception. A feature that might help simulate the non-determinism of real applications is inherent sensor inaccuracy. The system could degrade agent perceptions using a normal distribution weighted by distance of the communication/perception and the medium.

5.3.4 Cascading communication protocol. Currently, the model allows for indirect communication between someone emitting a message, and nearby agents who wish to receive it. However, a possible improvement that might give white blood cells a better idea of current global conditions is to use the collection of agents as an ad hoc network.

REFERENCES

1. **bacteria**. (2007). In *Encyclopædia Britannica.* Retrieved Dec. 19, 2007, from Encyclopædia Britannica Online: http://www.britannica.com/eb/article-272364
2. **blood count**. (2007). In *Encyclopædia Britannica.* Retrieved Dec. 20, 2007, from Encyclopædia Britannica Online: http://www.britannica.com/eb/article-9015699