Organisms II: Group VII Report

Vipul Bhasin
vb2139@columbia.edu

Lawrence Wang
lcw2001@columbia.edu

Zijian Zhou
zz2107@columbia.edu

1 Introduction

Organisms II is essentially a simulation of a natural environment: it consists of a 2D virtual world in which virtual Organisms live, and our task is to build the 'brain' of these organisms to help them best cope with the environment they are in.

Our organism went through three major phases of development. In each phase, we evalutated it against other players and tried to add features and tweak rules for best performance. The resulting player performs reasonably well under most circumstances. The analysis of the problem and the design of our players are given in the following sections.

2 Problem Analysis

2.1 Environment Analysis

Understanding how the environment works is crucial to coming up with good strategies for our organisms. There are several factors affecting the conditions of the environment: p, q, m and n. p is the probability that food appears spontaneously. q is the probablility that a given unit of food in a square doubles in a turn. m and n are the dimensions of the world (board). Of course these parameters are not available to our players and we must come up with our best estimates.

p and q together decide how food is scattered on the board. If p is high, there is a a good chance that food will grow. p, along with q, decide how the food is scattered around the board. For example, if p is high but q is low, this means that food 'blows in' frequently but it does not double as frequently. In this case, there will be a lot of squares with a low amount of food. On the other hand, if p is low but q is high it means that food appears less, but once it does, there is good chance that it will double. Therefore, there will be few squares but in them there will be a high amount of food.

Parameters m and n decide how big the world is. Several strategies, as we see later, require quite a bit of space in order to succeed. Furthermore, space constraints have an effect on competition for food.

The parameters s, v, u, K, and M are the parameters whose values are known to the organisms. s is the number of energy units consumed in staying put. v is the number of energy units consumed in moving or reproducing. u is the number of energy units per food unit.K is the maximum number of food units that a square can hold. Finally M is the number of energy units that each organism has at full health. These parameters can be extremely useful in deciding how the organism should behave. For example, moving around in search of food may not be worthwhile when each food unit does not give enough energy to make up for the amount of energy spent in finding it. On the other hand, if motion is relatively cheap in terms of energy required, i.e the ratio of u and v is relatively low, it may be a better option to move more frequently in search of food. We can use these parameters in conjunction with our judgment of how good the environment is to determine if the organisms should reproduce to increase their population.

2.2 Game Analysis

The aim of the project is to come up with suitable strategies to let the organisms perform reasonably well under most circumstances. Basically this problem reduces to a number of choices that each organism must make during its life time, and they are detailed below.

1. Motion decision :

   The decision of when and where to move is an important factor in maintaining the health of an organism. For example, in abundant conditions, moving around in search of food is better than staying put to wait for the food to appear on a neighboring square. On the other hand, staying put might be a better strategy in a harsher environment.

   This decision should be based on factors such as u, v and M. The harshness of the environment must also be taken into account, since the same motion strategy may not work under all circumstances. Additionally, moving direction must not be overlooked, since nice food finding strategies can arise from them. For example, if all the organisms move in the same direction, a herd behaviour may be simulated. However, if they move in opposite directions, they may spread out and a different pattern may be formed on the board. These patterns determine how the organisms interact with each other.

2. Reproduction strategy :

   If motion is important in deciding an organism's health, then reproduction is important in deciding an organism victory over other organisms in the game. A higher population can deny other organisms access to food. But on the other hand, reproducing costs a lot of energy.

In the next sections, we will describe the strategies we have attempted and point out their positive and negative aspects. We will also explain the strategy of our final organism and its performance under different conditions.

3 Strategy & Analysis

3.1 Organism 1: The Net

3.1.1 Motivation

The first draft of our organism was written by Lawrence and tested extensively in multiorganism environments against last year's organisms. It soon became apparent that in average environments, Black Plague was the organism to beat, starting off very conservatively but building in its numbers almost monotonically to eventual dominance. We decided that in order to succeed, we would have to out-conserve Black Plague.

3.1.2 Strategy

Our organism basically follows some simple rules:

1. Don't move too much. In fact, if you're below a certain energy level or you haven't seen enough food to make moving around worth it, don't move.

2. Separate from your own kind so you don't compete against them for the same food.

3. Reproduce only if you've seen enough food that your offspring has a good chance of surviving.

A good judge of how well we're applying rules 1 and 3 is the average energy per unit; we know we're doing ok when we see 200-300 energy per unit.

Rather than design specifically for a certain set of environmental parameters, we wanted our organism to be adaptable. Therefore, it was necessary to develop a way of gauging an environment's parameters 'on-the-fly'. To do so, our organism keeps track of a number that indicates how abundant an environment is. In our code, this is avgFoodAmount. Each turn it is calculated as follows:

$$avgFoodAmount = \frac{(avgFoodAmount * (myAge - 10) + foodCount(foodpresent, enemies) * 10)}{myAge}$$

In later versions, this value is initialized in a child with the parent's current value (for the first organism it is 0); for the first version it simply starts at 0.

foodCount is a measure of the 'free food' that an organism can see on a given turn. It counts only those spaces around the organism with food and without another organism. We felt that taking other organisms into account made this a more useful measure, since food that has already been claimed is inaccessible. We also weight newer information over older information in the average, since our measure fluctuates with population and so needs to be continually updated.

1. At this point if there is food on the square that the organism is on, and our organism is healthy, ie. its amount of energy left is larger than a fixed threshold, then it should reproduce only to a square with food.

2. If there is no food on the current square, the organism will move to a neighboring square if there is food on it.

3. At this point the organism will check its energy: if it has less than the energy necessary for 8 moves, it will choose to stay put.

4. If the organism has enough energy, it will try to move away from organisms of the same 'species', which it identifies by comparing its external state.

5. Since at this point the organism knows for sure there is no food around, but if it is healthy (If the number of steps it can take is larger than the avgFoodAmount times a constant), and above certain age then it will reproduce onto a neighboring square.

6. If all else fails, the organism will try to make a move (if its energy level is above a threshold and the amount of energy needed to find an unit of food is less than or equal to the amount of energy it is worth). If all directions are blocked, it will stay put; otherwise, it will move according a preference list. In this version, organisms prefer to move south, will only move west if south is blocked, and will only move east if the first two are blocked; north is disallowed to avoid cycling.

3.1.3 Analysis

The organism we produced around these principles proved to be pleasantly powerful. Average energy level per organism soon became the standard we would use to evaluate the effectiveness of changes we made to the code.

Most changes went into the rules we used to determine when to move and reproduce. We experimented with a number of different combinations, trying to produce expressions that contained known variables in their proper positions depending on whether they should be directly or inversely proportional to a given rate; we found ourselves constantly tuning a handful of coefficients to get the best results. At this point, in the default environment, our population climbed slowly but steadily, by moving as little as possible and reproducing fairly conservatively. It was now able to consistently win against Black Plague most of the time.

The following is a screenshot of our Organisms in action by themselves.

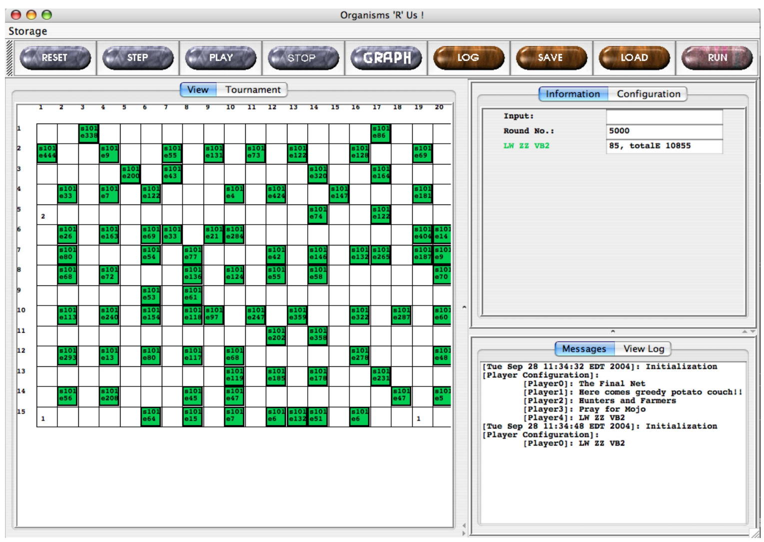

The following is a screenshot of our Organisms competing against Black Plague.

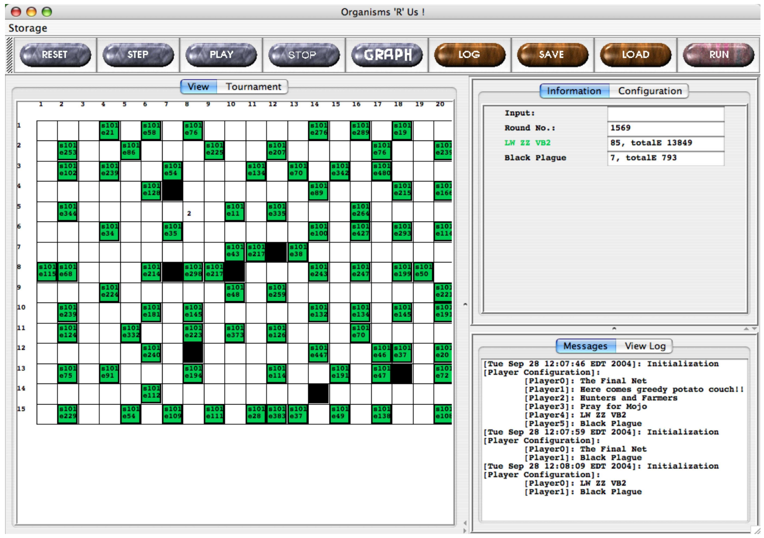

3.2 Organism 2: Mojo and The Net hybrid

3.2.1 Motivation

Our initial goal was to design an organism that excels under some conditions while surviving under most others. After competing with other organisms from previous year's class, we found that even though our organism is able of winning in an environment where food appears more often, it dies out rather quickly if the environment is low in food. We observed that our player would begin to move around, not find any food, and then simply sit still until food blew in next to it or it died of old age.

Most of the time the latter occurred; often food would be only a few steps away. We decided to use a different strategy for these environments. As a guide, we used an organism from last year's class which was often successful in harsh environments: Group 2's organism, Pray for Mojo.

3.2.2 Strategy

The hybrid organism we created tries to estimate the 'harshness' of the environment. If the environment is abundant in food, our organism will behave normally as our first prototype. But if the environment is harsh, the organism turns into Mojo and instructs its offspring to become Mojo as well.

Mojo's code is extremely simple. Mojo starts life with a preferred diagonal direction: northwest, northeast, southwest or southeast. It first checks if its health is below a certain level; if so, the organism will not move or reproduce. It will only reproduce at full health. Next, Mojo will check the current square or neighboring squares for presence of food; if an open space has food, it will move there. If there is no food, Mojo will move 25% of the time in one of the components of its preferred direction.

Our criteria for judging if the environment is harsh is as follows: For the first 100 rounds, our organism will exhibit the normal behaviour. But if the organism's age is more than 100 and it has not eaten any food, it will become Mojo and behave as described above.

3.2.3 Analysis

While this approach worked reasonably well under harsh circumstances, our organism is still unable to outlive Mojo, due to the fact that it is very hard for all organisms to reach a correct consensus, also first 100 rounds of the game proved really critical to the overall state of the game. since for the first 100 rounds we behave as if we are under abundant circumstances. by the time the organism reaches 100 it is too weak to outlive Mojo and usually dies of old age.

Also in the prosperous environment, it is possible for organisms to live until 100 rounds and not eat any food, and these organisms will incorrectly deduce that they must be living in a harsh environment and morph into Mojo, therefore influencing the behaviour of their offsprings as well. Even though this incorrect behaviour is unlikely, it reduces our organisms' overall energy and number.

The following diagram shows our player in an abundant environment competing against Black Plague. Notice the grey organisms: these organisms think that the environment is harsh and are behaving as Mojo.

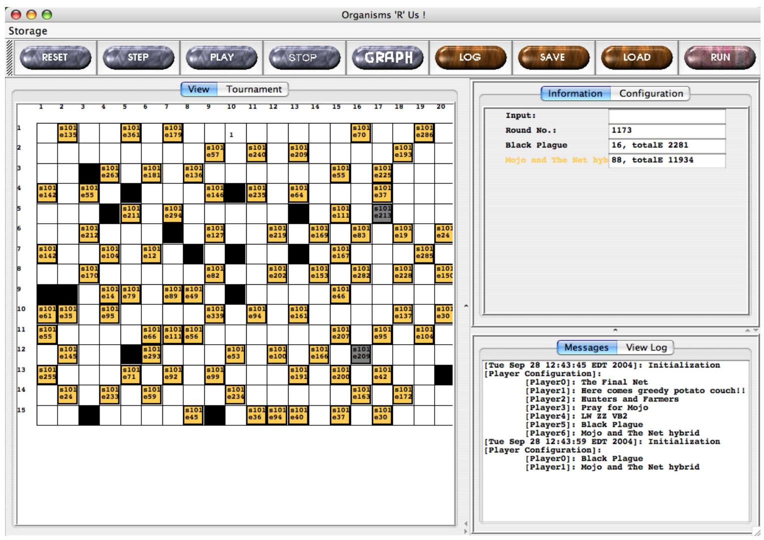

The following diagram shows our player in a harsh environment competing against Mojo. Notice because some organisms consumed food before they reached age 100, they did not morph into Mojo.

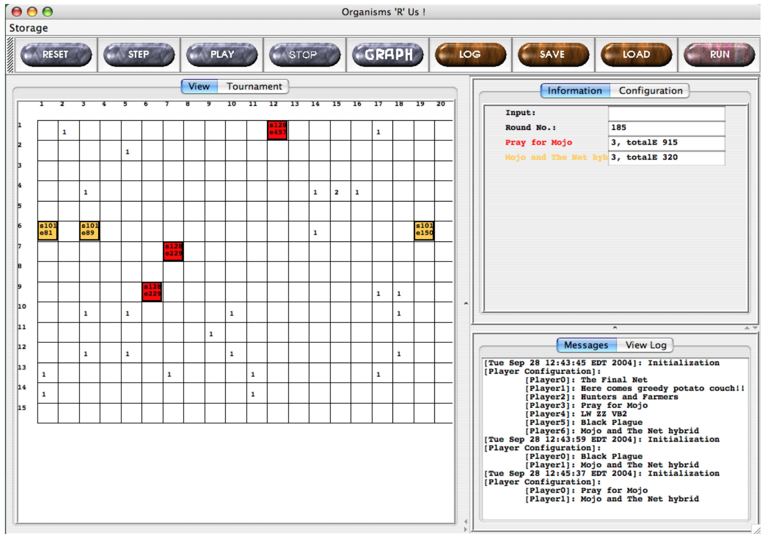

3.3 Organism 3: The Final Net

3.3.1 Motivation:

The three design rules listed earlier were originally written in an email to a member of Group 8 who wanted to know how our organism worked. After we shared our strategy with them, they updated their organism, WandererV2, one-upping us by implementing our conservative strategy for the long-term game and adding to it an aggressive early-game strategy. Also, they turned our pattern of separating from friends against us by spoofing our identifying state, with the simple trick of copying whatever states they observed in other organisms. With these new changes, they began to beat us consistently, and we went back to the drawing board.

3.3.2 Strategy

In response, we implemented our own early-game strategy. Why is our effort focused on the early part of the game? Because organisms do not know the precise values of p and q, there is generally a surge in growth at the beginning of the game, which drops abruptly when the population reaches a level sufficient to consume the incoming food. Between organisms with evenly matched long-term strategies, the one which comes out of this early period with higher energy has an advantage.

Our changes in this version are as follows: we reproduce and move at higher rates during the early turns of the game, attempting to pick up most of the food that 'blows in'. Though this higher level of activity translates into lower average energy levels, our organisms still benefit from denying other organisms access to food. Ideally, at the end of this stage, our organisms have the highest population and they are spread out across the board, giving them a good chance of picking up food as it 'blows in'.

Since Group 8 was spoofing our identifying state, whenever one of their organisms was adjacent to ours, we would try to move away from it, thinking it was a friend. This added tremendously to our movement costs, giving WandererV2 a major energy advantage. We responded at first by turning off the separate-from-friends rule entirely, but after careful consideration we made a compromise: each organism would only be allowed to 'move away from a friend' a low fixed number of times.

3.3.3 Analysis

The compromise meant that our pattern of spreading out was not as effective; however, it gave us an acceptable upper bound on the amount of energy lost to WandererV2's state spoofing. Furthermore, the effectiveness of this spoofing diminishes with the number of different organisms on the board, because it makes a direct confrontation between our organism and theirs less likely.

In large environments (30x30 board) our organisms seem to reach a dynamic equilibrium, possibly because we both reach a self-sustaining population level without interfering with each other early on. In smaller boards, one of the organisms may fall victim to bad luck with food placement early on and never recover. We also observed that playing too conservatively can be harmful; you want to keep the highest sustainable population possible, so that you are denying other organisms as much food as possible.

The following diagram illustrates that on a 30x30 board the 2 organisms maintain an equilibrium.

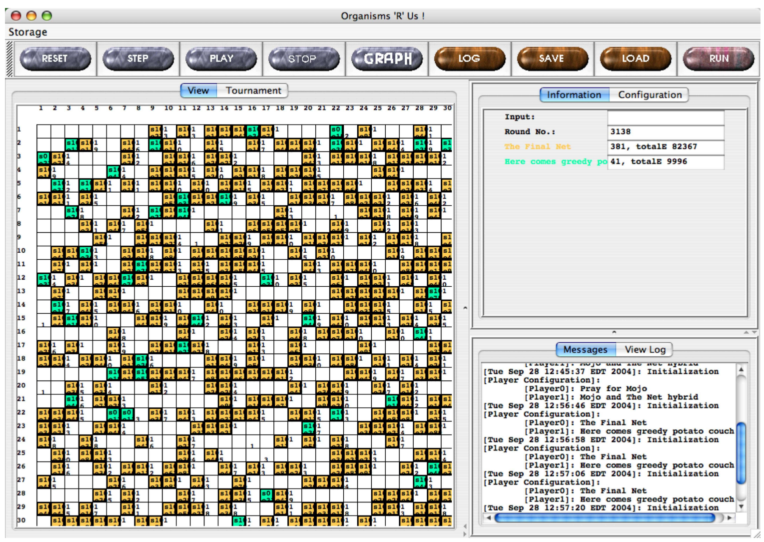

The following diagram illustrates that on a smaller board 1 of the organisms can fall victim to bad luck.

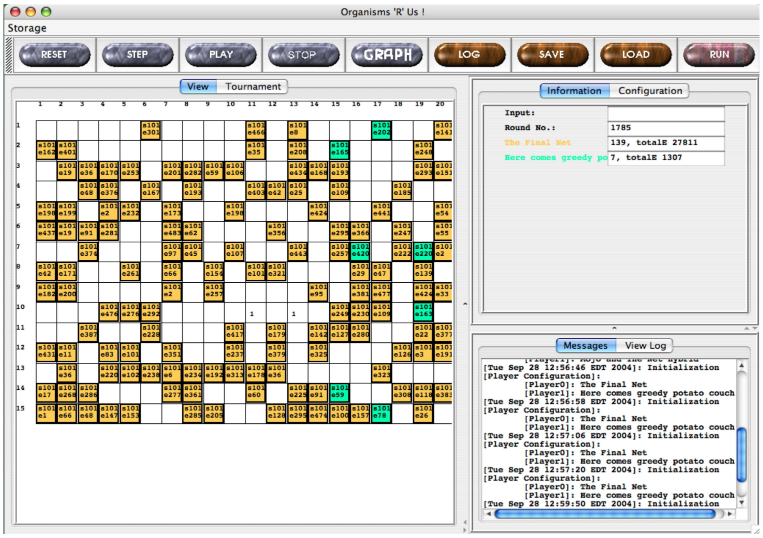

In harsh environments, our player still does not perform as well as we hoped; we were never quite able to find the level of aggressiveness that would allow our organisms to find food without moving too much.

4 Results

The following are the results from the multiorganism tournaments. Each graph shows the average number of organism after 10 games with each game lasts 5000 rounds.

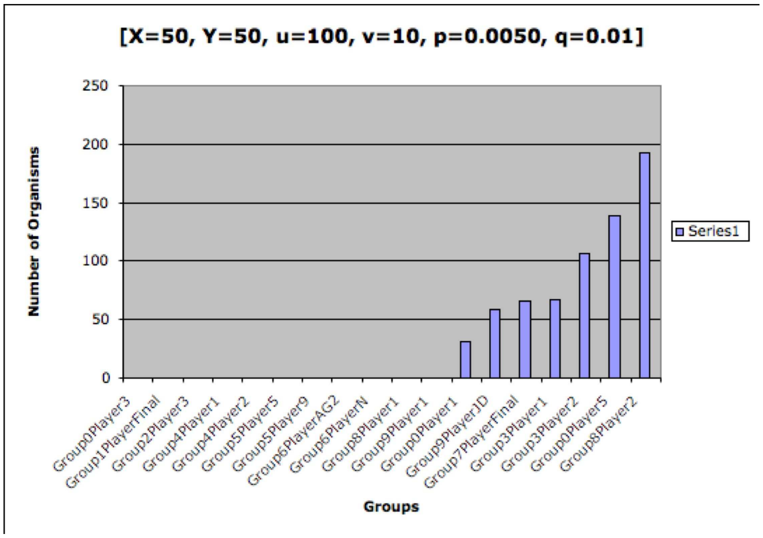

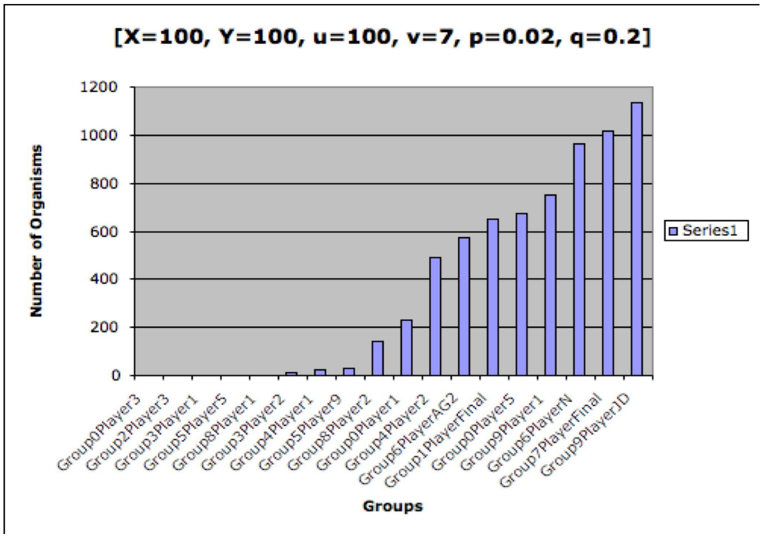
[X=100, Y=100, u=100, v=7, p=0.02, q=0.2]
Number of Organisms
Groups
Series1
Group0Player3
Group2Player3
Group3Player1
Group5Player5
Group8Player1
Group3Player2
Group4Player1
Group5Player9
Group8Player2
Group0Player1
Group4Player2
Group6PlayerAG2
Group1PlayerFinal
Group0Player5
Group9Player1
Group6PlayerN
Group7PlayerFinal
Group9Player1D

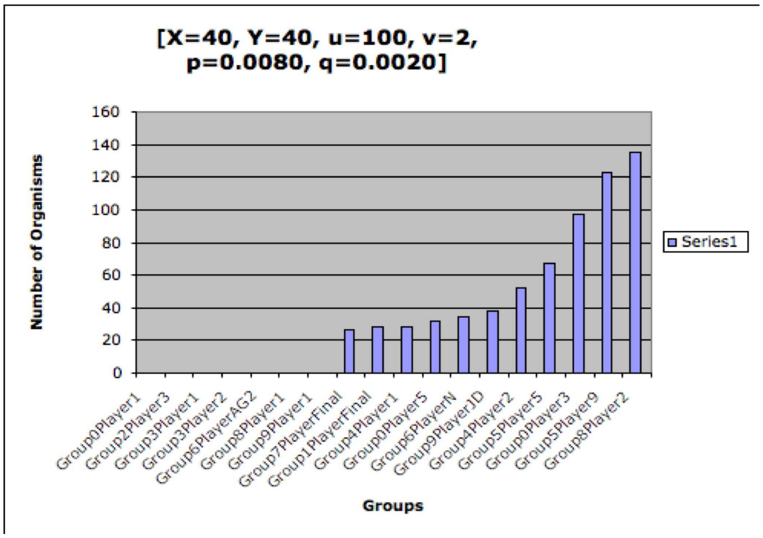
[X=40, Y=40, u=100, v=2, p=0.0080, q=0.0020]
Number of Organisms
Groups
Series1
Group0Player1
Group2Player3
Group3Player1
Group3Player2
Group6PlayerAG2
Group8Player1
Group9Player1
Group7PlayerFinal
Group1PlayerFinal
Group4Player1
Group0Player5
Group6PlayerN
Group9Player1D
Group4Player2
Group5Player5
Group0Player3
Group5Player9
Group8Player2

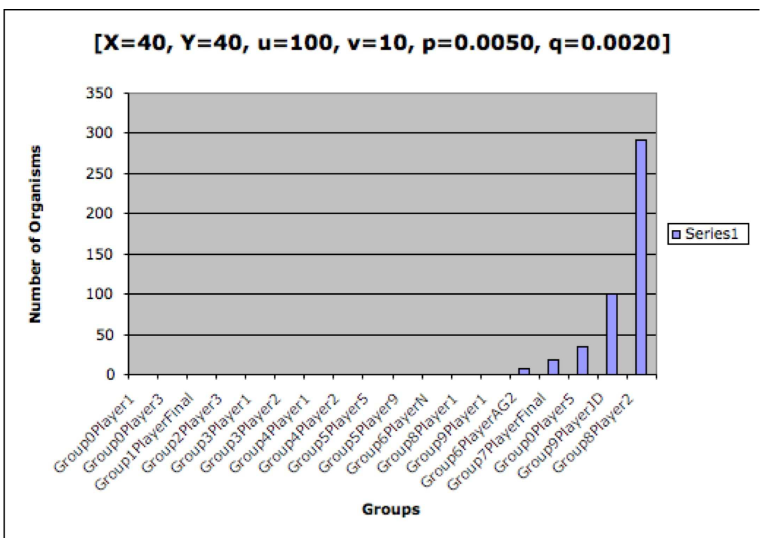
[X=40, Y=40, u=100, v=10, p=0.0050, q=0.0020]
Number of Organisms
Groups
Series1
Group0Player1
Group0Player3
Group1PlayerFinal
Group2Player3
Group3Player1
Group3Player2
Group4Player1
Group4Player2
Group5Player5
Group5Player9
Group6PlayerN
Group8Player1
Group9Player1
Group6PlayerAG2
Group7PlayerFinal
Group0Player5
Group9Player1D
Group8Player2

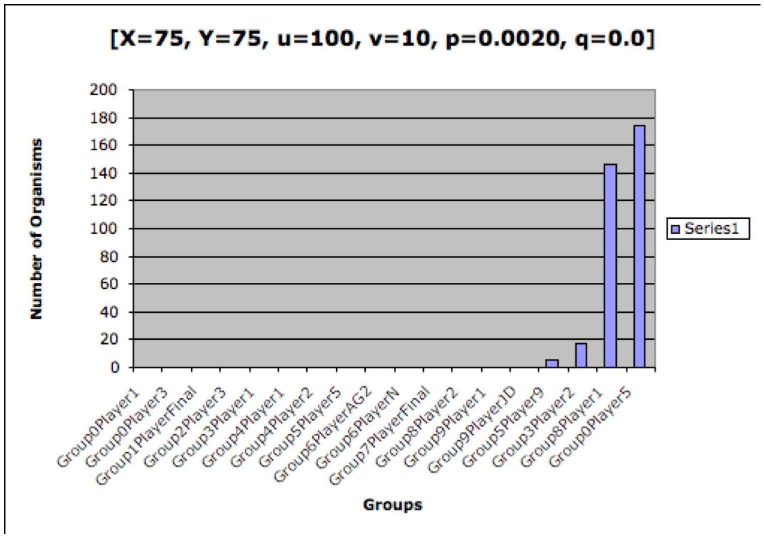
[X=75, Y=75, u=100, v=10, p=0.0020, q=0.0]
Number of Organisms
200
180
160
140
120
100
80
60
40
20
0
Group0Player1
Group0Player3
Group1PlayerFinal
Group2Player3
Group3Player1
Group4Player1
Group4Player2
Group5Player5
Group6PlayerAG2
Group6PlayerN
Group7PlayerFinal
Group8Player2
Group9Player1
Group9Player1D
Group5Player9
Group3Player2
Group8Player1
Group0Player5
Groups
Series1

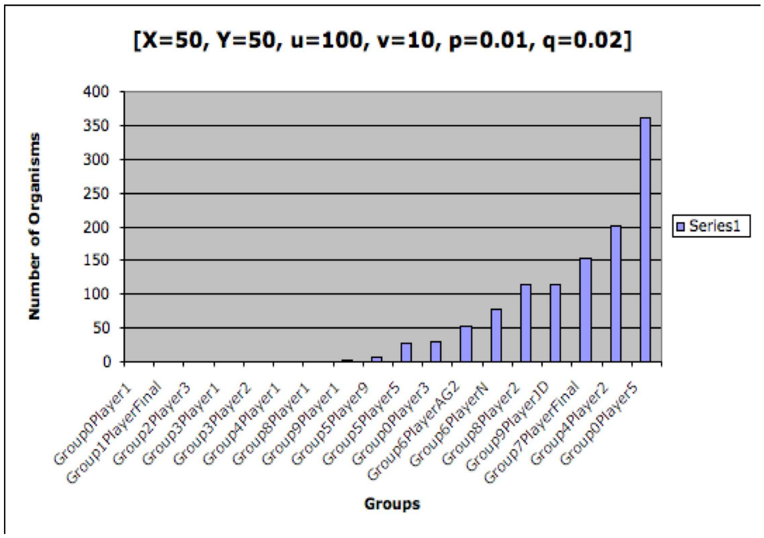
[X=50, Y=50, u=100, v=10, p=0.01, q=0.02]
Number of Organisms
400
350
300
250
200
150
100
50
0
Group0Player1
Group1PlayerFinal
Group2Player3
Group3Player1
Group3Player2
Group4Player1
Group8Player1
Group9Player1
Group5Player9
Group5Player5
Group0Player3
Group6PlayerAG2
Group6PlayerN
Group8Player2
Group9Player1D
Group7PlayerFinal
Group4Player2
Group0Player5
Groups
Series1

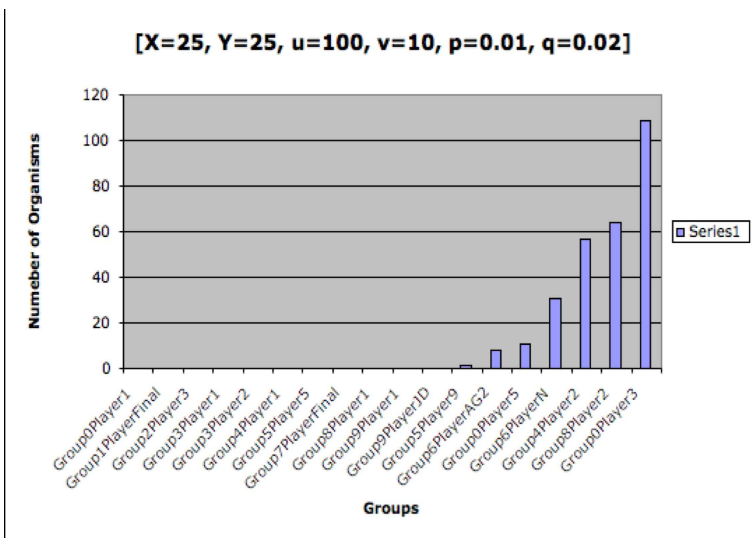
[X=25, Y=25, u=100, v=10, p=0.01, q=0.02]
Numeber of Organisms
120
100
80
60
40
20
0
Group0Player1
Group1PlayerFinal
Group2Player3
Group3Player1
Group3Player2
Group4Player1
Group5Player5
Group7PlayerFinal
Group8Player1
Group9Player1
Group9Player1D
Group5Player9
Group6PlayerAG2
Group0Player5
Group6PlayerN
Group4Player2
Group8Player2
Group0Player3
Groups
Series1

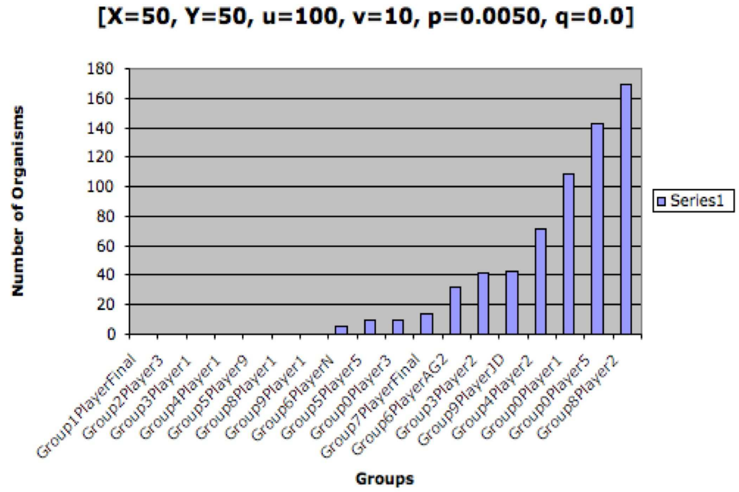

5 Conclusion

As we mentioned earlier, our player still has difficulty succeeding in harsh environments; this was evident in the tournament with parameters [X=75, Y=75, u=100, v=10, p=0.0020, q=0.0], where we died out after the 393rd turn on average. This indicates that we reproduced once or so, wandered around unsuccessfully looking for food, stopped moving once the lack of success was reflected in our measurements of avgFoodAmount, and then died of old age.

Our implementation of early-game and late-game strategies ended up being too inflexible, in contrast to a strategy like Group 8's which appears to switch between two states not solely based on the age of the game, but on observed conditions. This inflexibility is demonstrated in the tournament with the parameters [X=25, Y=25, u=100, v=10, p=0.01, q=0.02]. Though it had near-default values, because eighteen players were competing at once it was effectively a very harsh board due to crowding. As a result, our early-game foraging strategy, which we hardcoded to last for a fixed number of rounds, did more harm than good. Due to its costs, we were unable to reach a critical mass and died out near the end of the game. In contrast, the two organisms that did the best were Black Plague, which is always conservative, and Group 8's WandererV2, which we presume switched into a conservative state once observing the crowded board.

Our organism performed best when the environment is large and abundant in food, for example in the tournament with parameters [X=100, Y=100, u=100, v=7, p=0.02, q=0.2] in which our organism finishes with an average number of 1017. Or in the tournament [X=50, Y=50, u=100, v=10, p=0.01, q=0.02] in which our organism finished with a number of 153. In the cases above, the abundance of the environment lets our organism to recuperate from the energy lost in the early-game stages, when it was trying to spread its population across the board to get to a critical mass. In addition to the availability of food, space turns out to be a very important factor in deciding the success of our organism. Since the board had more space and it let the organism from each group to forage for food with less 'confrontation'.

In conclusion, we feel that our organism did quite well in the abundant environments where food is abundant and space is plenty. Since our player is quite conservative in strategy, our player will outlive most of the other players and have the board open to ourselves. However, in the harsh environments the inflexibility of our conservatism ended up doing more harm than good. Since our organism waits for food to appear and ends up dying of starvation. In retrospect, we would like to encode some genetic variation into our organism so that its offspring will exhibit different behaviours and adjust to the environment accordingly.