

CS4444: Project 4 - Organisms 2 - Group 2 Final Report

Stephen Lee, Eric Li, and Jason Winokur

December 2, 2003

1 Overview

Whereas the previous projects like Rectangles and Cluedo permitted unfettered access to global game parameters, Organisms opened a limited number of settings to each player. In the class sessions leading up to the tournament, this significant difference from previous problems resulted in organisms which specialized for different conditions. Previous groups from prior problems could either submit multiple players to work under different conditions or construct a single player flexible enough to switch between specialized sub-players. With enough time and creativity, groups could produce robust enough players which player well across the limited range of parameters.

While the Organisms tournament has a limited range of permuted global parameters, each organism does not have access to most of these values. To make matters more computationally difficult, each organism has a very simple communication mechanism which makes efficient, robust coordination amongst the same species quite hard. Moreover, each organism of each species would only be able to see locally. Ultimately, Organisms' hidden parameters, basic communication medium, limited organisms sight, and other issues would prove challenging in overcoming. Whereas previous projects required great effort as measured by amount of code, this project could become manageable with a limited amount of code. Organisms required a great deal of precision rather than bulk to create a successful being under many different definitions.

2 Methods

Although we submitted two players to the tournament, throughout the development process, we produced a number of interesting players. We eventually produced a successful series of Genetic Algorithm based players which were flexible enough to work under most games.

2.1 Genetic Algorithm

Our implementation of an Organisms player was very simple because of reliance on a genetic algorithm to evolve our species. Instead of completely relying on draconian player constants which would fail on varying board conditions and at different stages of the game, offspring of our species would have mutated genes. Theoretically, the most successful offspring would survive and live long enough to propagate its internal parameter set, albeit with some minor mutations. Hopefully, our species would survive under most global conditions as well as temporal conditions like when all the food on the board has been eaten. The following players which implement this genetic algorithm variant differ on the internal parameters used by the first organism of our species.

-	G_i	P_M	α_R
---	-------	-------	------------

Table 1: Composition of Mojo Reproduction Key

2.1.1 Mutation

The mutations are each represented by a byte and are concatenated together to form the key passed to the offspring upon reproduction. With the exception of the leftmost byte, G_i , the generation byte, the floating point value, δ_i , encompasses the percentage offset from some base hard-coded probability or threshold, depending on the internal parameters. Organisms computes its offset from the base and adds the mutation forming the new value that the offspring will use.

$$\delta_i = \frac{\beta_i}{128} \quad (1)$$

2.2 “Pray for Mojo”: A Conservative, Baseline Player

Mojo, a minor character on an episode of the Simpsons which is a reoccurring theme in Jason Winokur’s players throughout the course, is a monkey who took on the lifestyle habits of Homer Simpson, the mess of a yellow human. After some time of living like Homer, Mojo also became a bloated mess and typed out on a Stephen Hawkins speaking machine, “PRAY FOR MOJO.” (You probably had to be there.)

Anyway, Mojo is loosely symbolic of our Organisms player. Instead of Homer and the monkey’s movement and eating habits, the state of the board dictates our Mojo’s movement probability, P_M , and the threshold, α_R , at which Mojo reproduces. Thus, when the board had a lot of food, Mojo would move frequently and gorge itself. It would possibly reproduce a lot, too, if conditions warranted a population increase.

2.3 “Mojo Version 2.0”: Really Simple Player

Mojo v2.0 is a very simple conservative organism which exhibits emergent behavior. Each Mojo operates autonomously without any communication with the collective. They all follow five very simple rules. The following are the rules in order of precedence:

1. If energy is below 100, stop moving, unless adjacent to food.
2. If energy is 100%, reproduce.
3. If on food, stay there.
4. If next to food, move onto food.
5. Move in one compass quadrant direction 25% of the time.

Even though Mojo has very simple rules, it survives in a dynamic range of environments. Mojo thrives in harsh conditions because it does not split for the first time until it finds enough food to reach full energy. In low food games, this usually does not happen for many rounds. Meanwhile, the food will be growing while Mojo is grazing. Eventually, the board will be more hospitable for reproduction and Mojo reproduces. Mojoes are very greedy, always planning in their best interest, and consume all of the food they can find, thus causing the decline of the population. But, as the board thins of Mojoes, more food begins to sprout up and the cycle repeats itself.

G_i	P_F	P_M	α_R
-------	-------	-------	------------

Table 2: Composition of Quijibo Reproduction Key

We thought about stunting Mojo’s growth by coding anti- clustering, but felt it was better in the long run if we let them do their own thing. Mojo does poorly in high p and q environments because of its initial conservativeness. But when the food becomes more scarce, Mojo typically booms at the right time to starve the remaining players.

2.4 “Quijibo”: A More Aggressive, More Flexible Player

The real problem with Mojo is its initial movement probability and reproduction threshold, a ratio of its current food and its total possible food storage. When these values are set too conservatively, other species have chances to overwhelm Mojo before Mojo could reach a critical mass where it could effectively evolve. Likewise, when these values are set too aggressively, Mojo and its offspring might die due to some harsh board conditions.

Quijibo is our attempt at resolving this significant problem. In particular, Quijibo focussed on optimally computing Mojo’s movement probability. The cost to remain still remained constant throughout each game, but the cost of moving, v , and reward for moving, u varied. Thus, we based our estimate of P_M on these parameters.

$$P_M = \frac{u}{v^2} \tag{2}$$

$$\alpha_R = 0.5 \tag{3}$$

Under some conditions, while Mojo would explode and have an overwhelming population advantage over its opponents, we had noticed that players like Group 6’s Paramecia would grow conservatively and have a greater food per organism ratio than each Mojo organism. Eventually, Paramecia would overtake Mojo. Paramecia’s eating would take precedence over its reproducing. Thus, it grew slowly but healthy whereas Mojo would binge on the available food. Mojo’s change of precedence between eating and reproducing should be able to combat its eventual decline.

Though α_R is initially constant, Quijibo differs from Mojo when it actually makes use of α_R . While Mojo would always choose to reproduce even if food was available, if $\alpha_R > 0.5$, Quijibo would always choose to eat before reproducing. On the other hand, if $\alpha_R \leq 0.5$, Quijibo would have the same behavior as Mojo. This seemingly small difference would significantly curtail excessive growth and allow the Quijibo species to last much longer than the Mojo variety.

In addition to varying reproduction and eating precedence, we added a limited farming component to Quijibo. Many gluttonous players, including our own, would try to consume everything in sight. Other organisms seemed more robust to the occasional dust bowls that the board would be subjected to. Farming, under certain circumstances, seemed to help stabilize a frequently up-and-down population. Consequently, we augmented the other two mutated genes with a third, a food farming probability, P_F . Basically, if a Quijibo organism is on a spot with more than one turn’s worth of food, the Quijibo organism, depending on P_F would probabilistically stay to eat or move off. We hoped that this would lead to pockets of food accessible to members of the Quijibo population when the growth of new food sources became scarce.

2.5 Additional Player Enhancements

Each of our players have a different collection of modules which help it accomplish their greater objectives.

Frequently, organisms like DumbPlayer which implemented a random movement routine would move back and forth between two or more positions. Often none of these spots have any food. Therefore, this organisms would waste energy and ultimately accomplish nothing. To combat this behavior, our organisms would pick a random compass quadrant to move in.

In addition to the full integer organisms could pass to their offspring, organisms could read a byte which represented each others external state. Uses of this byte included friend or foe identification or states in some highly complicated coordination for some greater global configuration like farming or enemy organism entrapment. We decided to implement “good samaritan” information sharing. If an organism discovered a patch of food, it would incorporate this patch into a map which tracked its distance, size if possible, and age, the time elapsed since it was last seen. It would set its external state to the closest, largest, and youngest patch it has last seen. Thus, if another organism saw its state and understood its protocol, it could be redirected to a patch of food. Obviously, since organisms can only advertise a byte’s worth of information, the range of patches have to be limited. Yet, it still enables organisms to see beyond the one unit local area. Moreover, it is incredibly light-weight and requires no more than a single round of contact between two organisms which understand the protocol.

3 Analysis

We noticed some key factors which influenced the results of the tournament. The following enumerate the ones we felt we more significant.

3.0.1 Mobility

One of the keys to surviving that we were able to conclude empirically after trying our players with various different settings was that low mobility seems to be advantageous to long term survival. At first, we focused our attention on ways to find food or protect the food that we find or to hinder enemy organisms. What we didn’t realize until later was that all this activity requires a relatively high rate of food intake to maintain, and it’s often better for our survival to not be so active. For example, on our search for new food sources, if we happen to find an abundant source, then we’re in luck and our chances of survival become better. However, in most cases, we deplete more of our energy walking around searching than we’re able to take in by eating what we find. The net result is that we lose energy gradually, and our long term stable population decreases until there’s very little competition so that our organisms can find food often enough to sustain their active use of energy.

To remedy this problem, we limit our organisms’ frequency of movement. Since staying put always costs the minimum amount of energy, we can last a long time just by staying put. Also, instead of going out to search for food, we let food come to us by checking for the appearance of food in our immediate environment. This is much more efficient since we expend almost no effort in the process. For the actual implementation of the limits to mobility, we use a combination of an energy threshold and a random probability of movement. The energy threshold tells our organism to move only if it has enough energy, where enough is a tunable factor. Our rationale is that if our organism is about to die, moving around will only hasten the process, which we don’t want. The random movement probability uses a random number generator together with a range check to allow us to control precisely what percent of the time our organism will move when it has that option. Even when it has sufficient energy to move, we don’t want it to always move. We want our organism to find the optimal balance between staying put and moving that maximizes its chances of survival. Of course, the probability of movement is also a tunable parameter.

After extensive testing, we proved our rationale correct in that organisms with limited movement were often able to survive better than those without. So is there an optimal set of mobility parameters? The answer is that the optimal selection of mobility parameters depends on the parameters of the board, like u , v , p , and q . The relative values of these parameters will determine whether it's more efficient to actively search for food or to wait for food to come to you. For example, on harsh game board environments with high movement costs relative to the energy gained per unit of food, it is better to move less since searching will have very low food yield and very high energy costs.

3.0.2 Reproduction

Just as mobility, reproduction is also a key factor that needs to proceed at carefully controlled rates for better survival. On the one hand, reproduction multiplies the presence of our organisms on the board allowing more simultaneous searches for food potentially increasing our survival chances. On the other hand, reproduction cuts the energy, and, therefore, lifespan, of our organisms in half. This is a high risk gamble because we now can only last half as long and walk half as far before we must find food or else die off.

We have tried controlling reproduction in several different ways with varying levels of success. Our primary measure of when to reproduce is our energy level. If an organism has more than x amount of energy, e.g. three quarters of M , we allow it to reproduce. We also check for other conditions conducive to growth. For example, some of our players would stem reproduction if our organism was surrounded by too many of our own kind. This is to avoid overcrowding and competing with our own kind for food. Some of our other players would only reproduce onto space that already had food on it. The rationale here is that since reproduction is so costly in terms of energy, we want to make sure we can at least regain some of our energy and mitigate some of the dangers of reproducing by immediately having food to eat.

Through our experiments, we concluded that reproduction is only worthwhile under high food conditions. In harsh conditions, it's usually better not to reproduce until we're sure that it won't hurt us too much, i.e. we have near full energy. This way we cultivate a small but vigorous population that outlasts those players who spread themselves thin scouring the board for food to sustain them.

3.0.3 Greed

One of the obvious strategies to use is to sit on food when we find it. In the end, this was incorporated as part of players' behaviors. But initially, it might not be so obvious that this is the best thing to do. If you think about it, how do we know that by sitting on the food we find, we're not forsaking a larger supply that could sustain us longer? Also, as long as we're sitting on this food, it has no chance of growing. Is it better to move off of it when the food becomes low to let it replenish itself or is it better to consume all the food before moving on? The answer to all these questions can be found in the simple rationale that it's better to hold on to what we have for certain than to relinquish it for the possibility of a better yield later or elsewhere. Empirical observation proves just that. Especially when competing against other organisms, we found that relinquishing food is seldom a good idea, and consuming all the food that we find survives a double purpose; we bolster the health of our own organisms while eliminating a potential resource for enemy organisms. Thus, we have stuck with being just plain greedy with the food that we find.

3.0.4 Tournament Results

Screenshots, tables, and charts are available at <http://www.cs.columbia.edu/~sl2285/cs4444/p4/report/etc>.

The difference between us and the other players is that we optimized for survival. We raised the threshold for reproduction, minimized unnecessary movement, and hogged food whenever and wherever we could. In

essence, we were planning for hard conditions right from the start. On the contrary, many other players start out with aggressive expansion in an effort to crowd out the competition. But in the process, they spread themselves thinner than us, and die sooner than us if there's not sufficient food appearance and growth to sustain them, which is the case under harsh conditions. This is also the reason that our player doesn't have larger average populations and doesn't perform as well as some of the other players in high food situations. But that's alright because our goal is to survive, and we do.

By looking at the tournaments results, it is evident that our players perform moderately well under most conditions and very well under harsh conditions. In the single player cases, we almost always survive with the lowest p and q values. This is apparent in the tables of lowest p and q values at which all the players survived. We sorted the tables from lowest p and q values to highest and ranked the players in that order. From this view, we consistently ranked first. If you look at the graphs of our p and q survival values compared to those of other players, you'll notice that ours is consistently less than all other players'.

For the multiplayer games, we tallied the total number of times a player survived out of the total number of times that player was played to get the survival rate in terms of percentage of games survived. Looking at the normal Seven Player Games tournament, we can see that we have the second best survival rate. Our true capabilities, however, show through in the Seven Player Games Take Two tournament where the conditions are much harsher than in the normal Seven Player Games tournament. Here, we have the highest survival rate at 57%. The next best player has a survival rate of 10% proving that we are the more robust organisms under harsh conditions. Maybe they should put our organism on Survivor.

4 Conclusion

In conclusion, we found that despite how hard we tried, we weren't able to come up with an organism that dominated in all aspects under all environments. Our experimentation led us to the realization that while one strategy might prove successful under one condition, the opposite strategy might be best under a different situation. Thus, we shouldn't always try to cover all cases, but we should specialize instead and excel under certain conditions. In the case of our project, we made our organisms especially capable of enduring harsh environments, and that turned out to have served us quite well in the tournaments.