
A Spatial Model of Foraging Competition amongst Ant Species

Brett Calcott

Introduction

The Argentine ant has proved itself very successful at invading and displacing native ant species throughout a large range of countries (Holway 1999; Tsutsui, Suarez et al. 2000; Suarez, Holway et al. 2001). This widespread invasion depends on both an ability for long-range dispersal via human transport mechanisms and a superb ability to dominate and spread locally once established. The model I am interested in constructing is concerned only with the second of these abilities.

(Holway 1999) suggests that the Argentine ant's ability to dominate and displace native ant species results from overcoming the typical trade-off that other ant species must make between "interference" competition, and "exploitative" competition. It seems that most ant species specialise either in an ability to quickly find and retrieve food (exploitative), or an ability to defend a food source against other nests (interference). However, the Argentine ant seems to excel at both of these strategies. This ability seems, at least in part, to extend from its high population density which, in turn, is due to its unusual colony structure. The Argentine ant forms large supercolonies, where the workers of separate nests freely intermix. The lack of specific defended territories within the nests relies on a lack of intra-specific aggression. And this lack of aggression may lie in the reduced genetic variation that seems apparent in the introduced Argentine ant nests outside of their native range.

The initial aim in constructing this model was to attempt to successfully model the differences between two (or more) species of ant with differing amounts of intra-specific aggression. Additionally, I had hoped to include a range of parameters that could be adjusted to see what might be successful against such an invading species. Further investigation of ants revealed a plethora of varied behaviour in all domains including such things as scouting, predator avoidance, food storage, and mating strategies. For example, the way in which different ant species form new colonies ranges from complex synchronised mating flights amongst multiple nests to the simple vegetative budding of a small number from an already existing nest, through to the colony splitting strategy of the massively colonial army ants. And within each of these strategies there are many variations. Given these complexities and the time constraints, my attempts to model a "typical" ant species have been guided by what seems to be the most usual behaviour, and what is the simplest to model. I have concentrated my efforts and replicating at least some of the foraging behaviour, choosing to ignore the complexities involved in the nest reproduction; this approach, no doubt, will miss some critical features (i.e. I have not taken into account any probability of predation during mating). Foraging behaviour in the model can be roughly divided into four parts. Scouting, the process of discovering new food sites; recruitment, the way in which other ants become aware of this new food; harvesting, the collection and return of the food to the nest; and finally, competition, the initiation and resolution of any fights between ants from differing nests that our present at the same food site.

The Simulation

The first decision to make in the simulation is to choose the level of detail at which we are going to model the nests. One way to do this would be to ignore individual ant behaviour completely and instead model the foraging and interactions purely probabilistically. It appears that animal dispersal can be modeled as correlated random walks (Byers 2001), and this could serve as a starting point. However, this approach quickly becomes extremely complicated as, in this case, we have many other considerations to take into account to calculate the final energy contribution per ant. For example, we need to consider the following factors: interactions with other ants (including fights and possibly death), the successful gathering of the food, the time taken to harvest the food, and the travel time to and from the food.

Additionally, the successful return of a food-carrying ant to the nest is what enables recruitment of further ants to that particular food source. This type of positive feedback loop requires a very careful control to prevent it seriously distorting the results. An alternative approach would be to explicitly model the behaviour of each individual ant -- its current position, the current task it is doing, its interactions with food and other ants, and its own individual energy reserves. Given that we wish to model multiple nests, with each nest containing perhaps hundreds of ants, it is clear that this approach could quickly become computationally intractable. I have tried to steer a path between these two extremes.

Given that we are primarily interested in the foraging interactions of the ants, I model the internals of each nest as a black box for generating worker ants. Whilst inside the nest I simply maintain a figure representing the number of ants available. Once an ant leaves the nest it is represented as an individual in one of several ways. If the ant is a scout, then it exists within a particular zone radiating out from the nest, with the probability that it encounters a new food site being inversely proportionate to the size of the zone (see Figure 1). Alternatively, if the ant is currently recruited to a particular food site then it exists either at the site (waiting, harvesting, or potentially fighting), or it is maintained in a queue going to, or coming from, the food site. Modeling the ants in this way avoids the problems with a purely probabilistic approach by individually representing the ants when required, but does not become bogged down in the particulars of individual ant movements.

The simulation takes place in a two-dimensional world with nest sites and food sites occupying a single randomly allocated point in the space. Each discrete time step iterates through all nests and food sites allowing them to process what is taking place under their control. Both nests and food have a particular type that controls all of their underlying parameters. This allows us to create various scenarios involving multiple "species" and varying food types. In what follows I will expand on some of the details involved in each of these areas.

Nest Growth and Reproduction

Each nest begins with a number of workers and a surplus amount of energy. We assume a simple energy cost per worker to construct them, and that any surplus energy is transformed into workers. The rate at which this happens is constrained (representing a limit on the rate that the Queen can produce eggs), thus at the end of the time step there may still remain some surplus energy which can be used in the following time steps. Energy is gained by the successful foraging of food, and energy is lost by any scouting or foraging activities. Unsuccessful foraging and scouting may result in an energy deficit for the nest, in which case the workers are removed from the nest (i.e. they die). Given the scouting and foraging constraints supplied by the other part of the model, these simple equations result in the type of sigmoidal growth rate observed in actual ant nests (Wilson and Hölldobler 1990, p159).

The complexities involved in the dispersal and mating of the reproductive cast is effectively ignored in this model. New nests are simply created every given number of time steps a normally distributed distance from the founding colony. The founding colony does, however, incur the equivalent loss of workers and energy required to establish the new nest. One last point, new colonies share the exact same parameters as the colony that produced them. The model is clonal; it includes no variation in reproduction. All variations are between species and manually set by the parameters.

Scouting Behaviour

Each nest, depending on its type, allocates a percentage of its total workforce to scouting. The area around the nest is divided into slices of even angles, and a scout randomly chooses one of the slices when leaving the nest. During that period the ant is away from the nest, it will search for food only in this sector. This is not ad hoc as it may seem. Many species of ants forage along what are called "trunk trails" that consist of orientation pheromones and form a dendritic tree that branches out from the nest. This sector is in turn subdivided into a series of

zones whose perimeters are marked off by even subdivisions along at the radial lines of the sectors. The area encapsulated in each zone increases linearly with the distance from the nest. The size of the subdivisions is controlled by the maximum distance that this species of ant can move in one time step.

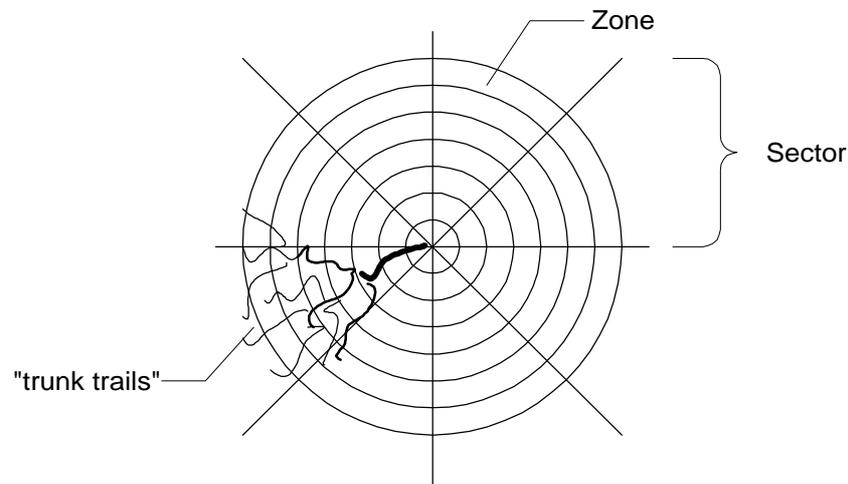


Figure 1. Sectors and zones compared with "trunk trails" around a nest site.

A scouting expedition by a particular ant consists of a stochastic series of moves from zone to zone ending when the ant reaches the nest again. Each time step the scout ant makes one or two decisions. Firstly, it decides whether it should return to the nest or continue scouting. If it decides to return to the nest, it moves one zone backwards. Alternatively, if it continues scouting, it must decide whether to stay in the current zone or advance to the next one. Each of these decisions is controlled by a probability based on a sigmoidal function. The probability that the ant turns back to the nest increases as the time it has spent scouting increases (see Figure 2), and the probability that the ant will stay in the current zone increases with the distance the ant is away from the nest (see Figure 3). Adjustment of the parameters that control these sigmoidal functions allow for scouting behaviour that spends a small amount of time in the zones close to the nest, and the majority of time in larger zones a reasonable distance from the nest, with the odd foray into some of the more extreme zones. The probability that a scout ant encounters a food site that exists within one of the zones is proportional to the size of the zone and the speed of the ant. Finding food in one of the zones close to the nests is therefore easier, as the area is smaller and the ant encounters the zone both on leaving the nest and returning to the nest.

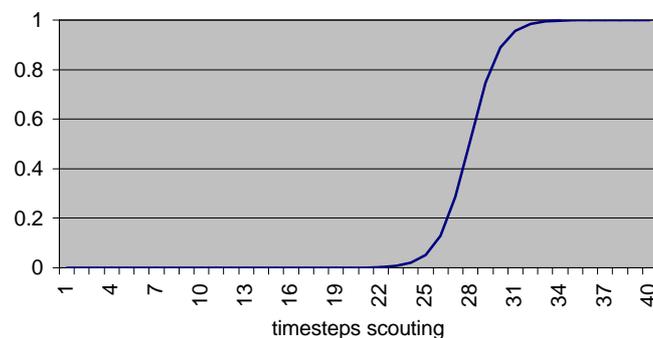


Figure 2. Example of Decision 1: $P(\text{Return Home})$

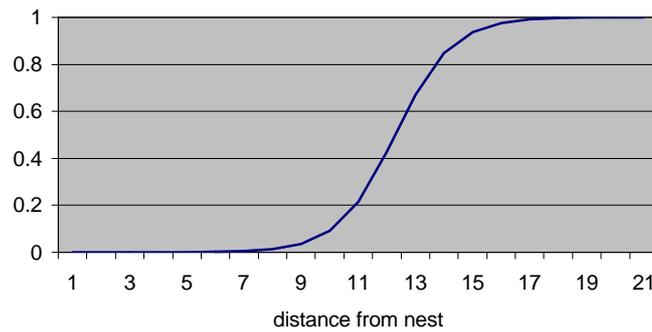


Figure 3. Example of Decision 2: $P(\text{Stay-Don't-Advance})$

This method has the advantage of accounting for the speed of the particular ant species. A faster ant can both make it further away from the nest in a smaller number of time steps, and explore the zone it is in more rapidly (via a higher probability of encountering the food site). The downside is that all parameters must be manually adjusted. A better method would be to incorporate an adaptive strategy that enlarges or shrinks the scouting zone based on encounters with ants from other nests.

Recruitment & Harvesting

Each food site contains a maximum number of parcels, and each of these parcels carries an amount of energy associated with the food type. Parcels of energy can either be harvested by ants, or dissipate with some probability for each time step (this allows undiscovered food sites to eventually disappear). Once the food site is exhausted it can magically regenerate itself again with a given probability per time step.

When a scout discovers a food site, the ant becomes "recruited" to that food site. An internal structure is created to hold all the information relating the food site to the nest site. The previous scout may already have recruited the food site; in which case the scout ant simply joins the queue of other ants attempting to harvest the food. The recruitment structure holds a number of queues of ants performing tasks related to the food site:

- ants that are leaving the nest to go to the food,
- ants waiting to harvest the food,
- ants harvesting the food,
- ants returning with food,
- ants returning to the nest empty-handed

As each ant moves from task to task it is moved from one queue to another; spending the correct number of time steps in each queue that is required to complete the task. When a scout first encounters food site, or an ant coming from the nest arrives at the food site, it is first placed in the waiting queue. Each type of food has the maximum number of ants that can harvest at any particular time (a number corresponding to the surface area available for harvesting). As soon as a space is available in the harvesting queue (which may be immediately) the ant is moved to this queue. When more than one nest is present at the site, then the ant that is moved to the harvesting queue is randomly selected with each ant, regardless of nest, having an equal chance (a nest with more ants present has a better chance of getting an ant shifted to the harvesting queue). Each food type also has a defined number of time steps that are required to harvest one energy parcel. Once the ant has finished harvesting it is moved to the "returning with food" queue, where it spends the right amount of time to return to the nest depending on the distance that the food site is from nest and the speed of the particular ant species.

On arriving at the nest the ant is absorbed back into a simple quantity of available workers and ceases to be modelled as an individual. At the same time an increase is made in a value in the recruitment structure that corresponds to the pheromone strength of the ant trail leading to the food site. This pheromone strength is decayed each time step and serves as a basis for calculating the number of the new ants from the nest that will be recruited to the food site and placed in the queue of ants going to the food. The more ants returning with food from a particular food site, the more ants that are likely to be recruited to this particular trail.

This method of laying pheromone trail down from an existing food site is one way that many actual ants use to direct their sisters to a new food site. Typically, when the new food site is encountered, the continual recruitment and dispensing of pheromone leads to an initial exponential growth in the number of ants that are recruited. However, if ants arriving at the food site are forced to wait too long, they return to nest without harvesting any food and importantly, without increasing the strength of the pheromone trail. The result is that the number of ants recruited to the food site tends to stabilise at an amount linearly proportional to the harvestable area of the particular food (Wilson and Hölldobler 1990, p271). To account for this additional behaviour in a "maximum waiting" parameter is added to each ant species. At each time step if the time spent in the waiting queue at the food site exceeds this value, then the ant is placed in the "returning empty-handed" queue. This leads back to the nest, but pheromone value is not increased.

Competition at Food Sites

Up until now, there has been no mention of competition with other nests. In this model I will only be considering competition at the food sites (I discuss the problems with this approach in the conclusion). As with scouting and harvesting, competition is controlled by several parameters that are adjustable per each ant species. At this stage the model only supports competition between those ants that are waiting at the food site. Those that are currently harvesting are not considered. The only reason for the simplification was the time constraints on finishing the model.

Within each time step, whenever there is more than one nest present at any food site, a random series of interactions between two ants of differing nests takes place. The number of these interactions is proportional to the product of the number of ants present for each of the nests. Two ants are selected from differing nests. An "aggression" parameter for each of the ant species that have been selected is then consulted to decide whether a battle will take place. Two different aggression parameters exist for each species -- one for interactions with an ant of the same species, and one for interacting with an ant of the different species. The average of the two ants' aggression parameters produces the probability that a battle takes place. Another probability decides the winner of the battle depending on the "attack strength" of the two ants. Each ant in the waiting queue has been initially issued a number of "defence points" depending on its species and the loser has one of these points subtracted. If the defence points of an ant reaches zero it is dead.

The complexity of this process provides a wealth of control over the types of interactions that can take place. The likelihood of both inter-species and intra-species aggression can be controlled. The individual strengths of each species attacking ability can be adjusted. And finally, the use of defence points allows for one nest's forces to be overwhelmed by the pure numbers of another nest, as a nest with less ants present will more than likely have each of its ants fighting multiple battles per time step.

Two additional features are added to this. Each ant species has a parameter controlling its possible behaviour after having been involved in a battle (as long as they are still alive). They can either run-away, or seek-help. Two additional queues are added to the recruitment structure to incorporate these features. An ant that runs away returns to the nest and reduces the amount of pheromone on the trail, whereas an ant that seeks help returns to the nest and increases the amount of pheromone. Neither of these delivers any the energy to the nest. Both

of these features are experimental and are not directly supported by observed ant behaviour, although there are definitely certain ants that retreat and others that escalate a battle.

There is a slight problem with the way interactions between ants are currently modelled. Ants from differing nests are randomly selected and an independent decision is made as to whether they will do battle. This assumes that the ants randomly bump into each other whilst waiting at the food site, then decide to fight. It is more likely that the aggressive species actively seek out those ants from other nests. Again, time constraints prevented the integration of this change.

Results

Pheromone strength and recruitment behaviour

The ability to recruit the optimal number of ants for any particular food site has a direct impact on nest survival. An over-allocation of ants to a food site results in ants returning empty-handed, a cost to the nest both in terms of wasting energy and the risks of predation. Ants appear to be able to elect the amount of pheromone that they dispense depending on the quality of the food, the distance it is from the nest, and how well fed the colony currently is (Wilson and Hölldobler 1990, p271). However, some of the results achieved whilst designing a model of recruitment behaviour suggest there may be another reason for the ants to elect to dispense a lower amount of pheromone when the food site is some distance from the nest, over and above the current satiation state of the nest.

Because the pheromone trail dissipates, a single ant arriving from a newly discovered food site initially recruits a few ants before the pheromone becomes an effective (remember that the pheromone is only dispensed when returning from the food site). In the case where as a large distance to travel to and from the food site, a lag exists before any successful ants return to the nest again. When they do arrive, they are bunched up and dispense a large amount of pheromone in a short time. This results in a series of waves of an exponentially increasing number of ants being recruited to the food site. The problem with this is that the wavelike behaviour takes longer to reach equilibrium, and in many cases fails to reach an equilibrium before the food source is exhausted. This results in a large number of ants being sent to a food site which no longer contains any food -- and all returning empty-handed.

Because the recruitment behaviour relies on a positive feedback mechanism, a delay in feedback because of a large distance allows the system to fluctuate out of control. The easiest solution to this problem is to reduce the amount of pheromone according to the distance that the food site is from the nest. This allows a slower and more even increase in the recruitment behaviour, which in turn allows the number of ants to smoothly reach an equilibrium. Experimentation in the model suggests that this simple change works well. It would also make sense for the amount of pheromone to be decreased based on the size (harvestable area) of the food source; however, I have not implemented this. Unfortunately, I am not widely enough read in the area of ant behaviour to know if this problem has already been recognised (there is no mention of it in Wilson and Hölldobler's 1990 bible "The Ants").

Competition between Species

Having established a set of parameters that allow a single nest to function adequately; scouting somewhat sensibly, realistically recruiting ants based on food distance, and maintaining a population relative to the local food resource, we are finally in a position to introduce multiple nests to a world. I created two separate species of ant, which were identical apart from the level of intra-species aggression that they had. Species A (or A-type) has a high degree of intra-specific aggression, and frequently attacks ants of the same species that come from a different nest when it encounters them at food sites. Species B (or B-type) has no intra-specific aggression and makes no attempt to harm ants of the same species that it

finds at a food site. Both ant species are highly aggressive to other species in all cases. In short: A attacks A, A attacks B, B attacks A, but B does not attack B. I initialised a world with an even the distribution of a single food type and place an equal number of species A and species B in the world; grouping the A species in one half and that B species in the other half. Figure 4 shows the results.

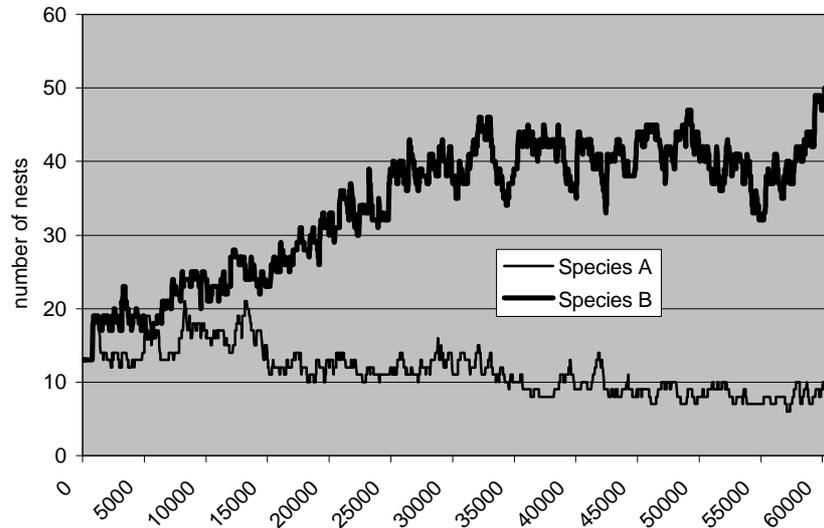


Figure 4. Species Competition

It's clear (and perhaps unsurprising) that the amount of intra-specific aggression has a significant effect on the success of a particular species. Two factors contribute to this. Firstly, for a species to spread it must be efficient at creating new nests. The most important factor in starting a successful nest is to build up a population as quickly as possible; a nest with few ants is *fragile*. Managing to do this will, in part, depend on your minimal number of scouts locating available food resources, and your ability to competitively harvest those resources; but most importantly, it will rely on your population not being wiped out in a battle at one of the food sites. A loss of five ants when your population is 500 it is not significant; however if your population is currently only 20, it can be catastrophic. Not only that, but if you end up competing at a food site with a nest that has a larger population, in all likelihood it will have sent more ants to that food site and you will thus be outnumbered, and more likely to lose. B-type ants have an advantage when their population is low: a B-type ant, when it encounters other B-type ants, only competes for resource, and that only indirectly effects their ability to increase their population. It suffers no direct and immediate loss from having its ants killed. An A-type nest, on the other hand, can suffer losses from meeting both A-type and B-type ants. Secondly, wiping out an existing nest requires reducing its population through limiting its resource and killing its workers off at food sites. In a situation where a single A-type nest is confronted with two B-type nests, the A-type nest has twice as many enemies to contend with – the B-type nests effectively 'gang up' on the A-type nests, presenting a slow moving consolidated front that steadily wipes them out. Importantly, both of these advantages rely on B-types being close to each other; if you put a single B-type nest amongst a group of A-types, it acts exactly like an A-type. So the success of the species with less intra-specific aggression also relies on it dispersing its new nests only relatively short distances.

What makes the B-type ants successful the way that they interact *with themselves*. It's worth having a closer look at each of the species in isolation. Using the same world, I created two new simulations, one beginning with 20 A-type nests; the other with 20 B-type nests. I tracked both the number of nests and the total number of ants across all nests. The results are

in Figure 5 and Figure 6. Note that each graph represents the consolidated results from 2 *separate* simulations, one with only A-type ants, the other with only B-type ants.

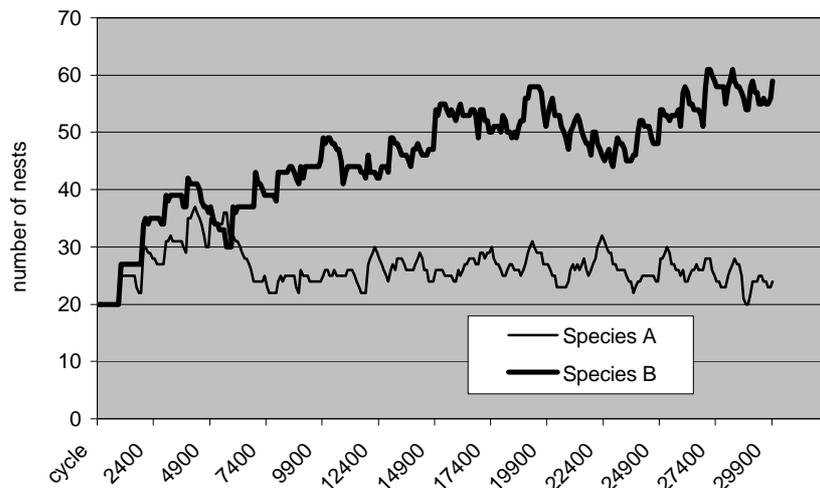


Figure 5. Nest growth in 2 simulations

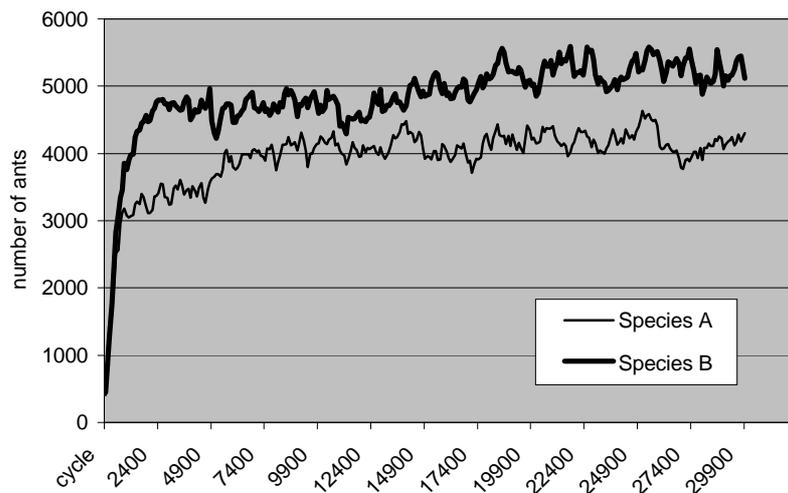


Figure 6. Global population growth in 2 simulations

Under identical food conditions, the world supports a lot more B-type nests than A-types – so B-type nests will necessarily be closer together than A-type nests. B-type nests must also, on average, contain fewer ants than A-type nests under identical conditions too (this is evident by dividing global population at any point from Figure 6 by the number of nests at the same point in Figure 5). What is most interesting though, is that having no intra-specific aggression results in a higher population density. The results parallel some relevant predictions that Wilson and Hölldobler make in a section on the co-existence of multiple species of ants.

“Interference between colonies belonging to the same species has the important effect of increasing the numbers of competing species that can coexist” (Wilson and Hölldobler 1990, p423)

This is explained by reference to the Gause-Witt theory, which predicts that when any two species interfere with one another, one of them will eventually win out *unless the population densities* of the dominant competitor somehow self-limits before total extinction of the other species. The classic case of this is when two species have only partially overlapping niches –

the population of the dominant species will be limited when it reaches the extent of its niche, and the subordinate species can coexist by existing on the non-overlapping part of its niche. Wilson and Hölldobler remark on some evidence that suggests that some species of ant control their population densities by both reducing the productivity of nearby queens and the destruction of them directly after nuptial flights. In the case of our model, it is clear that *another* way of limiting population density is simply by intra-specific aggression. This can be seen because when we *don't* have intra-specific aggression there is a clear increase in population density.

There is another way we can look at this situation. If we regard the nests as organisms, with their population being a mark of their fitness (which seems a *reasonably* justifiable claim), then B-type nests can be seen as reducing their fitness by adopting a non-aggressive strategy, and thus having to share the local resources. But although at the organism/nest level they reduce fitness, at the species level they are selected for – as their non-aggressive behaviour allows them to gang-up and swamp other species. We have already noted that the B-types ability to do this requires proximity to other B-type nests – in other words, the limited short-range dispersal of the ants provides a simple assortment mechanism for ensuring that like nests are located together. This factors seems to identify this as a clear case of species-level selection as described by the multi-level selection theory of Sober and Wilson (Sober and Wilson 1998). In short, there are two frameworks (at least) which provide an explanation for these type of results.

Conclusions

The primary goal of creating this simulation was to examine the effect of intra-specific aggression in nest competition between multiple species, whilst maintaining a fairly realistic low-level implementation of ant behaviour. To a reasonable extent, this goal has been met. Nest growth, scouting behaviour, the use of pheromone for recruitment, and somewhat realistic assumptions about interactions at food sites have all been modelled – and with a little tuning this has generated results that suggest that we are, at least, on the right track. Decreased intra-specific aggression leads to a higher population density, and when combined with limited dispersal provides the ability to invade another species that differentiates only in that it *does* practice intra-specific aggression.

The secondary goal was to modify some of the other parameters, for example different food types, different foraging strategies, or multiple numbers of species, to see how they might change the evolution of a multi-species ecology of ant nests. Unfortunately, lack of time prevented this. However, I am wary of what I think is a major deficiency in this model. Because ants only 'really' exist at the nest or at food sites, there is no possibility of interactions at any other place – even when the simulation clearly shows the recruitment trails crossing or when scouts from dissimilar nests are close by. I think this is a major problem; many interactions between species take place simply when they *encounter* each other, regardless of whether they are at a food site. Rather than progress with this model I think that what is required is a complete rewrite from the ground-up. Of course, some of the features that have worked in this model could be incorporated, but a rewrite is necessary to change the underlying world model to one that allows interactions at any location whilst scouting or foraging. Presumably some kind of cell-based approach would work, and prevent the computational tractability problems of modelling *every single* ant that I have tried to avoid.

Bibliography

- Byers, J. A. (2001). "Correlated Random Walk Equations of Animal Dispersal Resolved by Simulation." *Ecology* **82**: 1680-1690.
- Holway, D. A. (1999). "Competitive Mechanisms Underlying the Displacement of Native Ants by the Invasive Argentine Ant." *Ecology* **80**(1): 238-251.

- Sober, E. and D. S. Wilson (1998). Unto others : the evolution and psychology of unselfish behavior. Cambridge, Mass., Harvard University Press.
- Suarez, A. V., D. A. Holway, et al. (2001). "Patterns of Spread and Biological Invasions Dominated by Long-Distance Had Jumped Dispersal: Insights from Argentine Ants." Proceedings of the National Academy of Sciences **98**(3): 1095-1100.
- Tsutsui, N. D., A. V. Suarez, et al. (2000). "Reduced Genetic Variation and the Success of an Invasive Species." Proceedings of the National Academy of Sciences **97**(11): 5948-5953.
- Wilson, E. O. and B. Hölldobler (1990). The Ants, Belknap Press of Harvard University Press.