Arctic Rat Migration

John Bryden

School of Computing, University of Leeds, Leeds LS2 9JT, England
scs2jab@comp.leeds.ac.uk

Abstract. Rats migrating to Islands in the arctic are causing problems with the local flora. Scientists would like to understand their behaviour better so they can deal with this problem. The simulator described in this paper attempts to address this issue.

1 Introduction

Black rats (Rattus rattus) are migrating to three arctic islands (see Fig. 1). Biologists have been collecting rat population data for 6.5 years. A particular concern is the health of a rare species of Lichen, indigenous to Island C. Tests have shown that large concentrations of rat droppings could destroy the Lichen and conservationists are looking for methods of reducing the rat population. This paper describes the development of a tool intended to simulate and understand the growth and behaviour of the rats on the three islands and make predictions about the future.

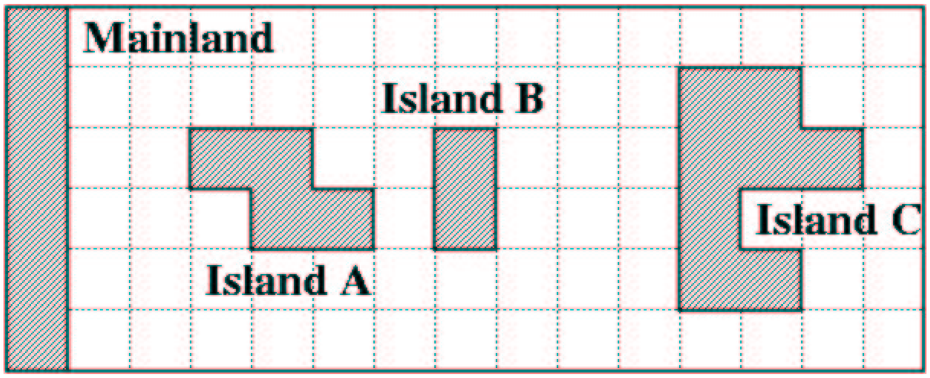

Fig. 1. The layout of the islands

Each grid-square in Fig. 1. is the equivalent of one hectare. Each grid-square can support 50 rats without major food competition. As the rats have no predators on the islands, biologists have confirmed that it is reasonable to assume that rats will only die from food starvation or from old age. The average rat would expect to live about two years.

From analysis of the real data collected by the biologists, several assumptions can be built into the simulation. At the initial stages (around month 23) when rat population is growing on Island A, the food stress reaches a certain limit then goes down (see Fig. 2). The rat population keeps climbing however. Interestingly the rat population crosses the 50 mark, the carrying capacity for one grid-square, at the same point. This evidence strongly implies that the rats are unlikely to migrate unless they are driven to do so by food stress.

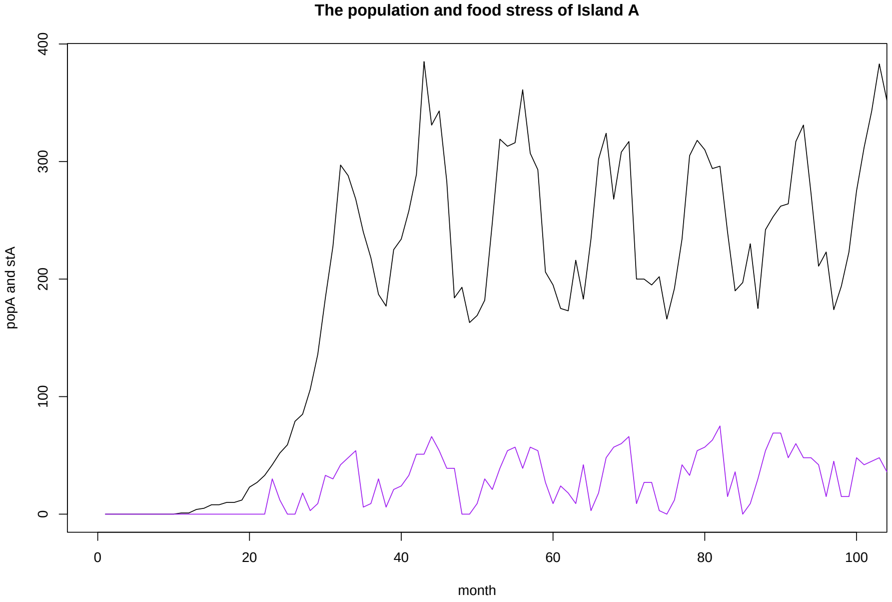

Fig. 2. Island A population in black, Island A food stress (multiplied by 300) in purple

Another assumption for the simulation was the analysis of the role of ice in the migration of the rats. Certainly the population of all three islands fluctuates in syncronisation with the ice flow (see Fig. 3). Also, note how the population on Islands A, B and C will never go above zero unless there is ice present. Island A's population clearly doesn't rise until month 11. IslandB's population rises very quickly immediately after the ice freezes over, starting on month 35. Island C's population rises little by little during months 50 to 70, but while it is at a low level it never rises unless ice is present.

By plotting the number of old rats on Island A (see Fig. 4), the very small number of rats that get to old age must imply that very few rats actually die of old age. Therefore, given the biological evidence, the majority of rats are being killed by starvation.

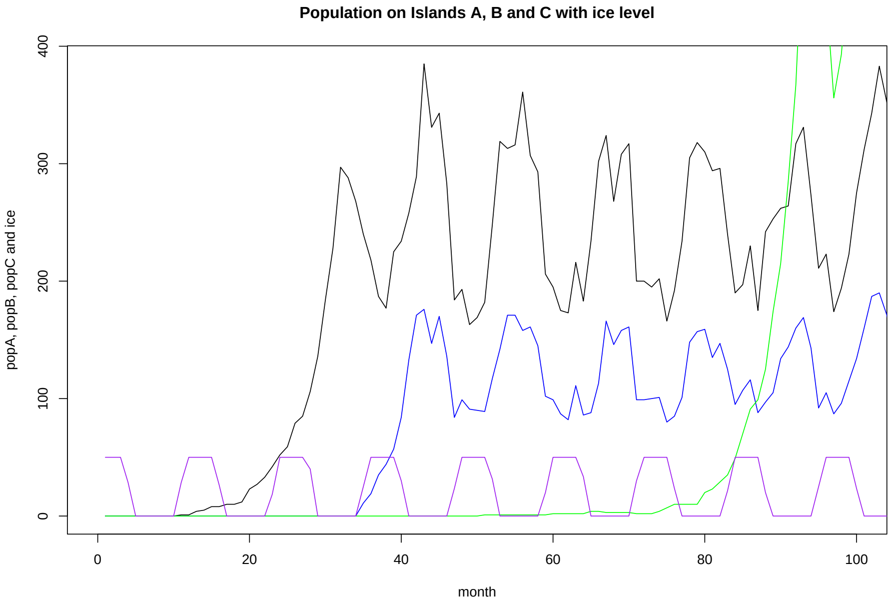

Fig. 3. Island A in black, Island B in blue, Island C in green, ice level (0-50) in purple

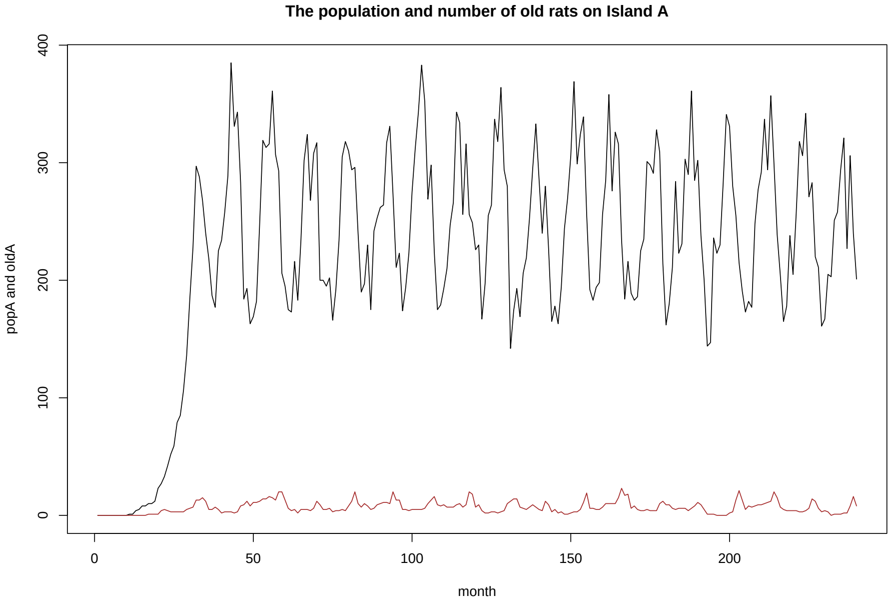

Fig. 4. Island A in black, old rat population in red

A further question put forward by the biologists regarded whether there was any evidence of the Allee effect in rat pregnancies. An Allee effect would mean that the ratio of births to females would behave differently when the population is small. Plotting the population of Island A against number of births (see Fig.5) shows that, certainly during the initial growth period, there doesn't seem to be any strange values. However there are some interesting, unexplained, spikes in the data.

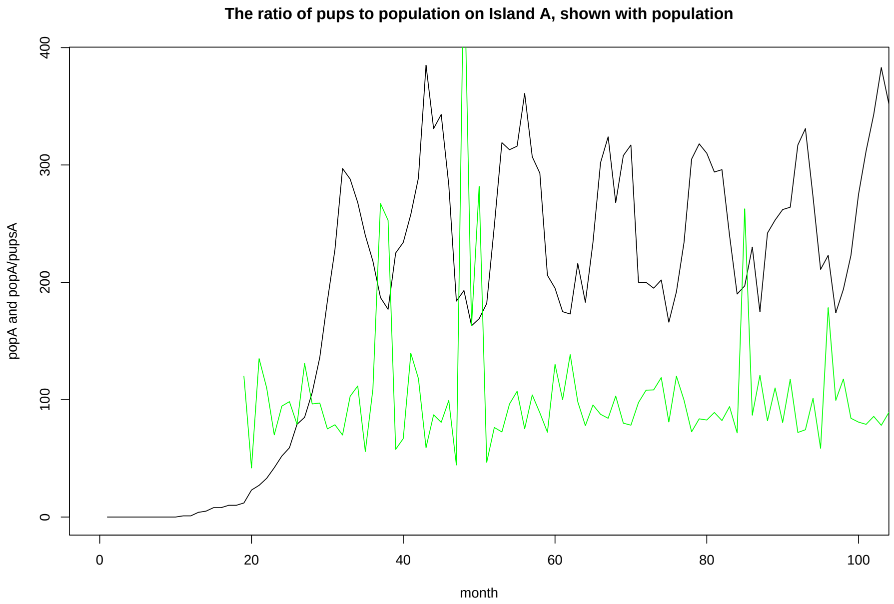

Fig. 5. Island A in black, population to pups ratio in green (muliplied by 100)

This paper aims to explore four areas. The first question regards rat migration and island carrying capacity. Some scientists have argued that the seasonal fluctuations in the data are caused by changes (of greater than 10%) in carrying capacity during summer and winter. Others argue that it is caused by increased migration due to the ice levels in winter. Both hypotheses will be explored.

Secondly, there is the question over which ages the rats migrate. It seems reasonable to assume that rats do not migrate when under the age of 1 month, however some tests may be carried to see how the simulation behaves when elder rats are denied the ability to migrate.

The third question is one of food competition between age groups. There are conflicting views concerning which age group should be fed first, the adult rats or the young rats or whether both groups should be fed equally. If the young rats were feeding first it would indicate that the older rats are protecting them.

If the opposite were true it would indicate that the older rats are using their strength to obtain food.

Fourthly, to try to save the Lichen on Island C, conservationists want to employ a rat trapper. The trapper would be expected to kill (or remove) half the rats in a particular grid-square per month. This paper will explore the use of an optimisation algorithm to generate an annual plan for the trapper to try to minimise the number of rats on Island C.

2 Method

2.1 The Simulation

The Simulation was written in Java so that a graphical demonstration could be easily developed. Each rat was modelled individually on a daily basis within the simulation.

Rats are known to reach sexual maturity after three years. Females give birth to an average of three pups (with a standard deviation of 1.0), with a gestation period of 32 months. Males and female pups are equally distributed. It was assumed that there would be a time period between birth and conception, the simulation uses an exponential distribution with an average of 30 days. Pregnant females will abort a litter if they have not been fed for two days.

From the evidence set out earlier it will be assumed that rats will only migrate when under food stress. They are known to be capable of travelling 100m per day. At the beginning of each day, in the simple simulation, 50 randomly chosen rats receive food per grid-square. Any remaining rats are said to be under food stress and will migrate. As no specific data is available covering rat migration behavioural patterns, the rats in the simulation will move randomly in any of eight directions (there will also be a one in nine probability that they won't move at all).

Data was provided outlining the time of year ice is said to be present. This is loaded into the application. The simulation was built with the hypothesis that rats migrate from island to island over ice and not by swimming. There is no food on the ice and when a rat reaches the mainland or goes too far from any island it will no longer play any part in the simulation. By observing the growth of Island A, it was estimated that at the start of every month, given ice was present, seven rats should be dropped off the coast of the mainland.

To work out the best annual plan for the rat trapper, an optimisation algorithm was chosen. The time taken to run one simulation was quite slow on a 750MHz AMD K6 computer. As genetic algorithms involve many separate individual runs of the simulator, it was decided that a 'Hill Climber' algorithm would need less processing power for the rat trapper optimisation. A list of twelve grid-square locations was chosen for each run of the simulation. A good recent solution was mutated by randomly changing three of the grid-squares to produce the next attempt. The average population on Island C over the next ten years was recorded and used as the fitness function to decide which attempts were best. The trapper was started once the population on Island C reached 10 rats - that is, at the equivalent point of 6.5 years after the start of the real data.

2.2 Verification

After the initial first few runs, the simulated population of Island A was observed to be oscillating in sync with the real data. However, it was not oscillating to sufficiently high or low levels for the simulation to be seen to be correct. After an introduction of a 10% lower carrying capacity in winter and a 10% higher carrying capacity in summer that data was observed to be a good fit, though it did not fluctuate quite as wildly as the real data (Fig. 6). The other two islands achieved similarly satisfying data.

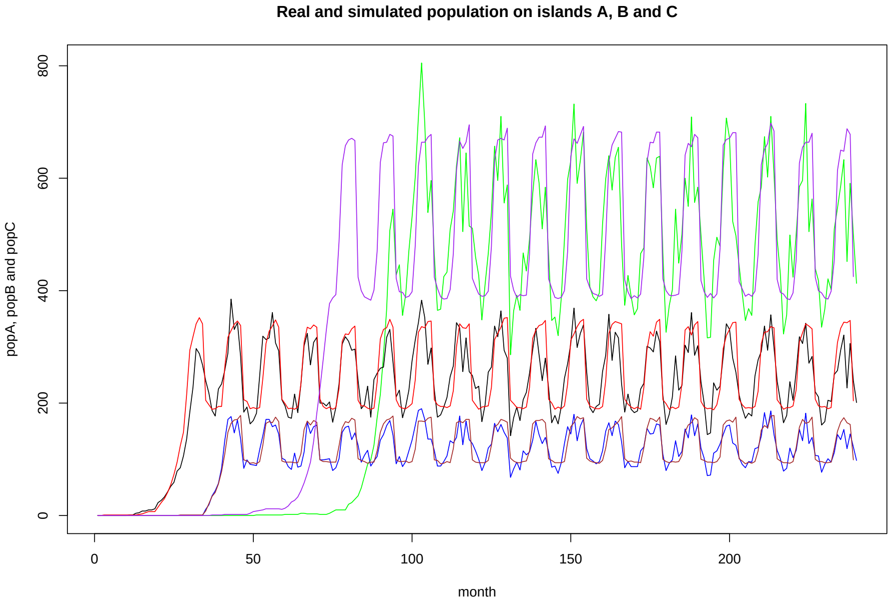

Fig. 6. Real Island A,B and C data in black, blue and green; simulated in red, brown and purple

A graphical output was developed for the application showing the population on each grid-square as a colour. The colour scheme chosen was: green for population under 50, yellow for population over 50 going to red at the highest levels. This demonstrated that the simulation was working as expected and also proved a useful tool for demonstrating the rat migration patterns. While under no food stress, the simulated population on island A rises at a very similar rate to the real data: both rise exponentially.

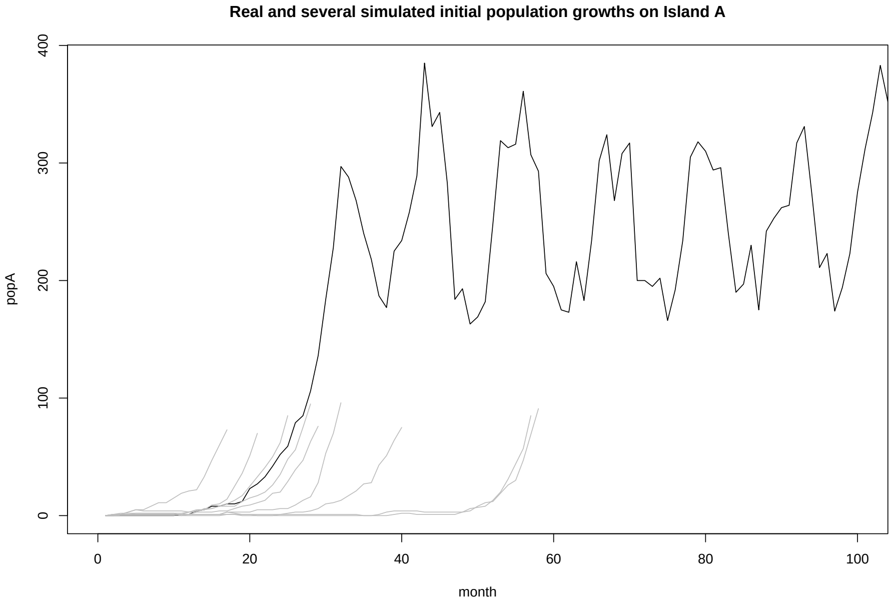

Fig. 7. Real Island A data in black, various simulated runs in grey

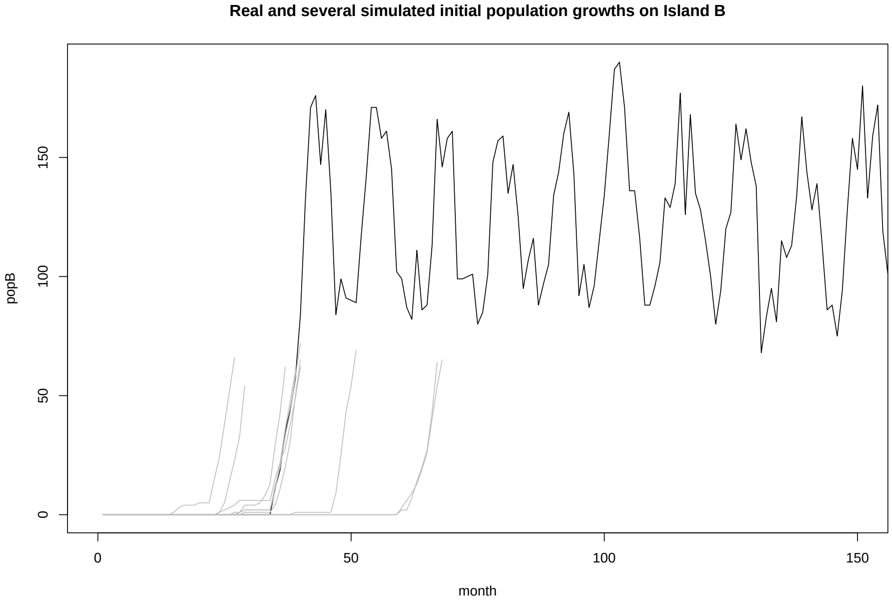

Fig. 8. Real Island B data in black, various simulated runs in grey

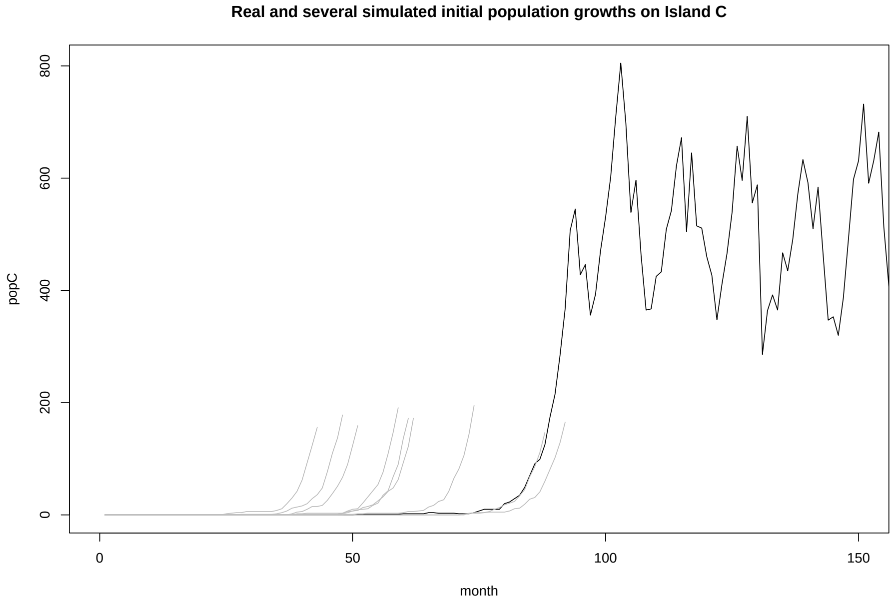

Fig. 9. Real Island C data in black, various simulated runs in grey

To check the rats migrate at a reasonable rate to Islands B and C, the simulation was run several times. The plots (Fig. 7-9) show that the real data falls within the simulated time frames for arrival of rats. The rats will only migrate from island to island when there is ice for all the simulation runs graphed.

Statistical analysis was performed on the data versus the real data, removing the initial 'warm up period' where the data can move at different speeds. The data is outlined below in Figure 10.

Island	Real/Sim	1st Quartile	Mean	3rd Quartile
Island A	Real	203.5	251.9	298.0
Island A	Sim	194.0	257.0	333.2
Island B	Real	98.5	124.5	145.0
Island B	Sim	96	128.1	167.0
Island C	Real	409.0	501.5	591.5
Island C	Sim	393.0	511.8	661.2

Fig. 10. Table showing statistical analysis for the simulated and real data

As can be observed, the 1st Quartile and the mean are close in all three islands. The 3rd Quartile is out by a factor of roughly 10%. This might suggest that the assumed change of carrying capacity of 10% may actually be an invalid assumption. The proportion of pups, young, adult and old rats was compared with the equivalent figures given by the biologists. Some dispute was found as to whether a rat at the age of three months was a young rat or an adult - as the simulated rats' age is calculated in months only this makes a significant difference. The data seemed to fit better if the rat became an adult at three months.

3 Results

To investigate whether the rats were indeed crossing over the ice, it was decided to run the simulation with permanent migration (the rats migrated at the same rate as they do during ice months all year round). The simulated level of population fluctuates on Island A in syncronisation with the real data (see Fig. 10). However it does not fluctuate as wildly. By changing the variance in carrying capacity to 30% (see Fig. 11), the data begins to look more realistic.

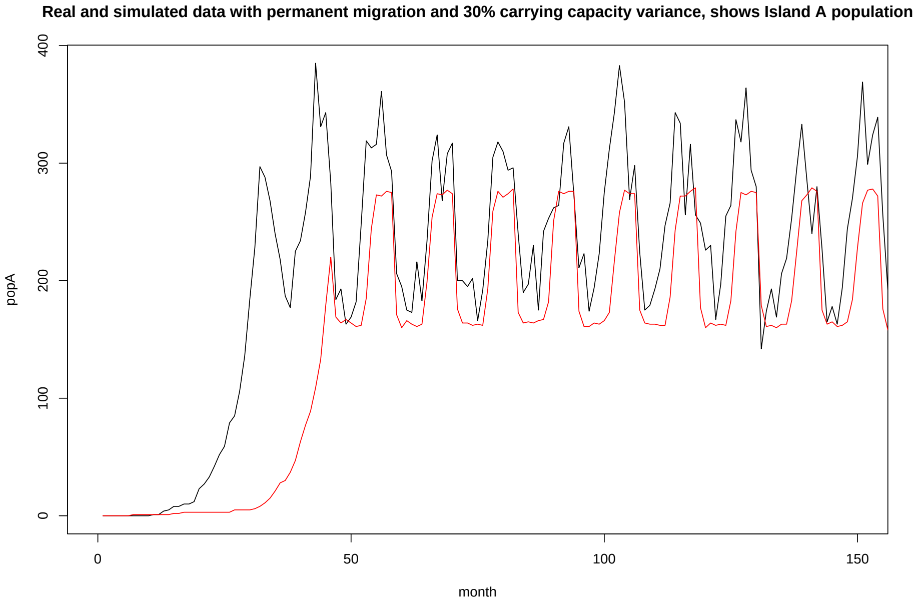

Fig. 11. Real Island A data in black, simulated data in red

In the real data, the islands only have rats arriving when there is ice present. To check the probability of this happening when the rats migrate perennially, the simulation was run 10,000 times. If a rat arrived onto a previously unpopulated island grid-square during the time where there is no ice present, the simulation was halted and a failure counted. During the 10,000 runs the simulation halted

9,976 times. If the 'no ice' hypothesis is correct, then the real data would only occur in approximately 1 in 400 times.

The population of young rats in the standard simulation run fluctuates with the ice flow; however, the population of young rats in the real data doesn't seem to fluctuate in any regular pattern (see Fig. 12). To test this, the simulation was adjusted so that rats below three months were not able to migrate (see Fig. 13). While this did not make any particular difference to the population trend as a whole, the young population did not syncronise with the ice flow. It was decided that this was a better approximation of the data so all future experiments with the simulation were done with the young rats unable to migrate.

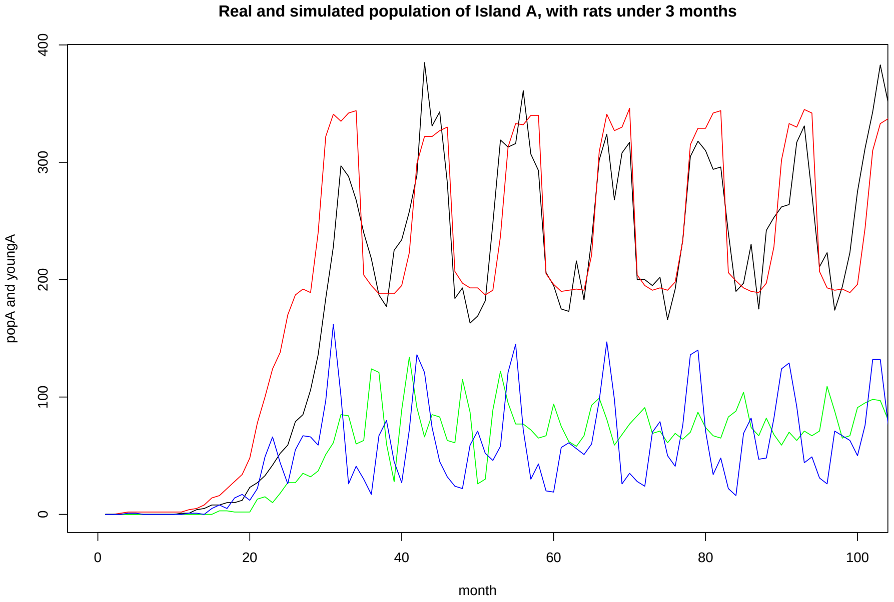

Fig. 12. Real/simulated Island A population in black/red, real/simulated young rat population in green/blue

The simulation was changed so that the rats in each grid-square were split into two groups (young and adults). The simulation was run with the entire adult group being fed first (see Fig. 14). The graph shows that the simulated population grows much more steadily and then fluctuates more moderately than normal. The older rats were starving out the young rats which were unable to migrate. The simulation was then run with the younger group being fed first (see Fig. 15). The graph shows that the simulated population grows normally but still fluctuates more moderately. The simulated population of young rats fluctuates wildly out of sync with the population of the adult rats. Clearly each

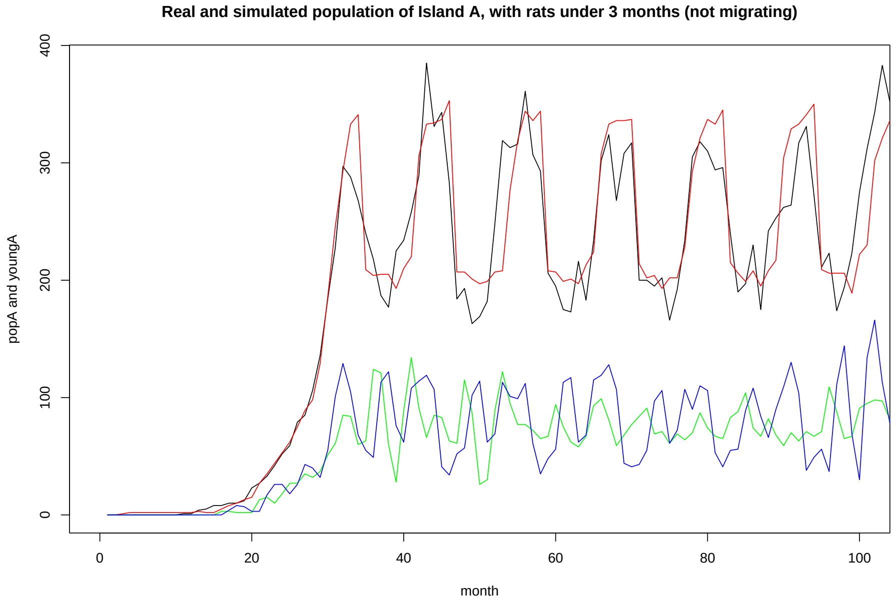

Fig. 13. Real/simulated Island A population in black/red, real/simulated young rat population in green/blue

new generation of young rats wipes out the previous generation of adult rats by eating all their food. Then, when they get to adulthood, they give birth to a new generation of young rats and the cycle repeats.

The conservationists asked that the rat trapper should start trapping when Island C is at a low population. The simulated trapper removes rats for 10 years starting when Island C gets above 10 rats. The average population on Island C was calculated over the 10 years. Each optimisation attempt was averaged over 4 runs to make sure the trapper wasn't getting too lucky. After 500 runs, the best solution had the trapper on Island C 10 months of the 12. The mean population on Island C was 397 which is approximately 70 below the mean in the real data (468 - the average over the ten years after the population on Island C rises above 10).

The system for numbering the grid squares starts in the top left corner and works down and to the right. The first (top left) sea square off the mainland has coordinate (1,1). Using this system, the optimal annual pattern for a rat trapper is:

Month 1 traps at grid-square 11, 4
Month 2 traps at grid-square 11, 2
Month 3 traps at grid-square 11, 5
Month 4 traps at grid-square 11, 3
Month 5 traps at grid-square 11, 2

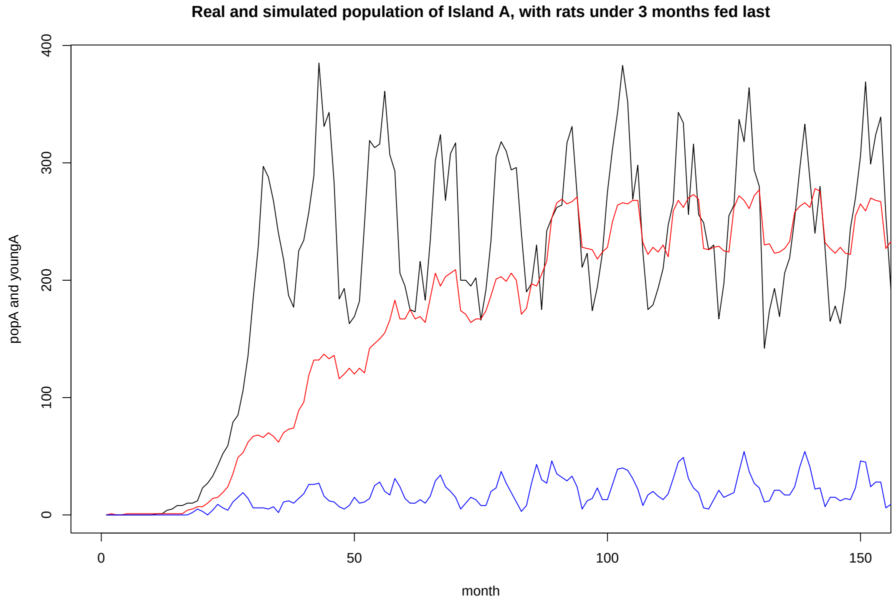

Fig. 14. Real/simulated Island A population in black/red, simulated young rat population in blue

Month 6 traps at grid-square 11, 2
Month 7 traps at grid-square 11, 5
Month 8 traps at grid-square 7, 3
Month 9 traps at grid-square 11, 2
Month 10 traps at grid-square 11, 2
Month 11 traps at grid-square 12, 5
Month 12 traps at grid-square 7, 4
It is not known if this will rescue the Lichen!

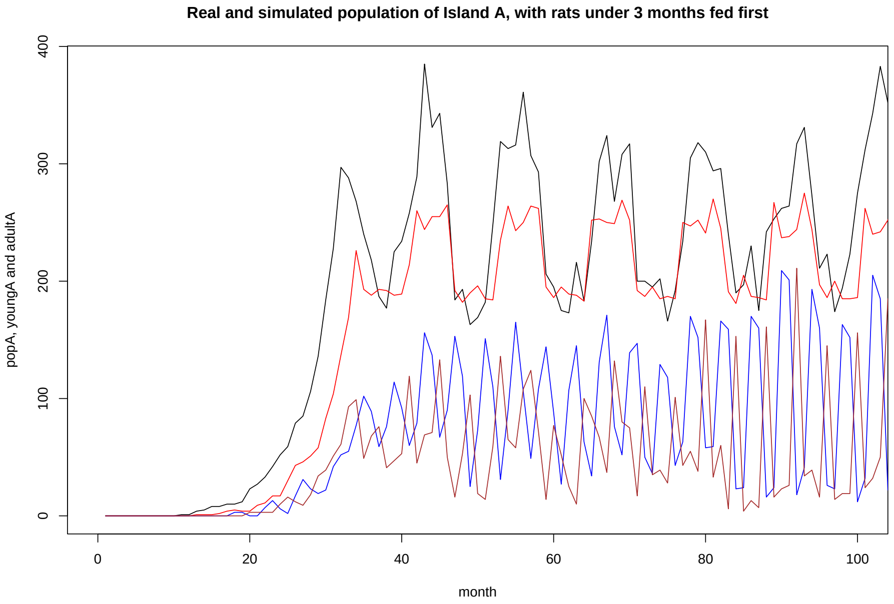

Fig. 15. Real/simulated Island A population in black/red, simulated young/adult rat population in blue/brown

4 Conclusion

The simulation developed produced data that was a good fit to the original data. It proved a valuable tool for testing the hypotheses outlined in the introduction. One of the hypotheses, questioning whether young rats migrate or not was shown to have strong evidence that they didn't migrate and therefore this functionality was incorporated into the simulation.

While it is still possible that the rats could migrate perennially, the evidence strongly suggests that they use the ice to migrate. It was shown that the real data would be very unlikely if the rats were migrating perennially.

There was no evidence to suggest that the young rats were feeding before or after the adult rats. Both cases significantly changed the behaviour of the simulation. While it is likely that parents feed their young, a whole family will miss out on food together.