

Helicopter Routing for Maintaining Remote Sites in Alaska using a Genetic Algorithm

Nicholas Armstrong-Crews, Kenrick Mock

University of Alaska Anchorage
3211 Providence Drive CAS Room 154
Anchorage, Alaska 99508
nickarmstrongcrews@hotmail.com
afkjm@uaa.alaska.edu

Background

The Department of Fish and Game keeps observational sites in remote locations for such purposes as counting fish, measuring weather, and various research projects. In Alaska, fishing and natural tourism are vital to the economy of the state, making these measurements all the more important. However, due to large geographic spread with little transportation infrastructure (none, between many villages), helicopters must be used to visit these sites. It is extremely expensive, primarily in fuel costs, to fly helicopters all over the state, so optimal or near-optimal routing of these helicopters is paramount. Currently, the “by-eye” technique is used, in which a route is chosen that simply looks like it would have the lowest total distance.

Problem Description

The well-known underlying problem at hand is the Vehicle Routing Problem (or VRP) in which a fleet of vehicles with a given capacity must make deliveries to a set of sites (Toth and Vigo 2001); this variant, however, only includes a subset of the problem specification:

- A single helicopter completes the entire circuit.
- Each site must be visited at least once.
- Some nodes are fuel depots rather than sites, allowing the helicopter to refuel; some nodes are both sites and fuel depots.
- A depot may be visited more than once.
- Helicopters have a limited fuel capacity; a helicopter must have enough fuel to reach the next node.
- Edge costs between nodes are simply geographical distance, since the helicopters fly; these also represent fuel costs.

The primary changes from the basic VRP problem are the existence of multiple depots and the use of a single vehicle

to visit all sites. A real-valued fuel constraint replaces the bin-packing constraints, essentially relaxing the bin-packing aspect of the problem. Nevertheless, the problem remains NP-complete, as it is more general than the underlying traveling salesman problem. Hence, for a large number of sites such as exist in Alaska, the problem is intractable. Note that the entire problem could be generalized to other vehicle types if edge costs were road distance rather than geographic distance; e.g., the same solution could be used for buses visiting touristic sites.

Many techniques have been developed to approximate optimal solutions for the basic and multiple-depot VRPs. These include ant algorithms (Bullnheimer, Hartl, and Strauss 1999), tabu search (Cordeau, Gendreau, and Laporte 1997), and evolutionary algorithms (Machado et al. 2002). The particular variant described above is not addressed specifically; therefore, in this project, a genetic algorithm (GA) was created with a simple genome and a novel crossover technique in an attempt to produce better solutions for this simpler variant.

Figure 1 shows an example for a small graph with $n = 10$ nodes, $k = 2$ depots, and side length $x = y = 10$. The starting point is marked zero and is enclosed asterisks; negative numbers indicate depots, while positive numbers indicate sites.




Figure 1: Simple Problem and GA Solution

DNA Encoding

The DNA for an organism in the GA consists of the path of nodes from start to finish. The DNA was generated by first randomly generating a path including all sites but no depots, then fixing that DNA by adding a depot visit whenever the fuel is exhausted. This is done by randomly choosing a time between the fuel exhaustion and the last previous depot visit and adding a visit to the nearest fuel depot at that time. This “fixing” approach weeds out impossible paths from the population, resulting in much faster solution convergence.

Selection

The organisms were sorted, using distance as the fitness function, and the most-fit half of the population was kept. These surviving individuals then reproduced to create the remaining half of the population.

Reproduction

Children were created by the following process (illustrated in Figure 2):

1. Copy DNA from Parent A to child.
2. Choose strip of DNA of random length from Parent B to donate to child.
3. Copying the strip directly would introduce duplicate genes, corresponding to multiple visits to a site (the only case we would want to revisit a site would be if it is also a fuel depot, which is still possible with this scheme via “fixing” of DNA). Instead, swap genes in the child to move genes to necessary loci to match the strip from B.
4. Fix DNA by adding depot visits to prevent fuel exhaustion. This is done by stepping through DNA, and if fuel is exhausted at any point, randomly selecting a locus between the last refueling and the exhaustion point and adding a visit to the nearest depot at that locus. The fixing process then starts again from the first gene.
5. Introduce random mutation (swapping of two gene loci) with a probability of .10.

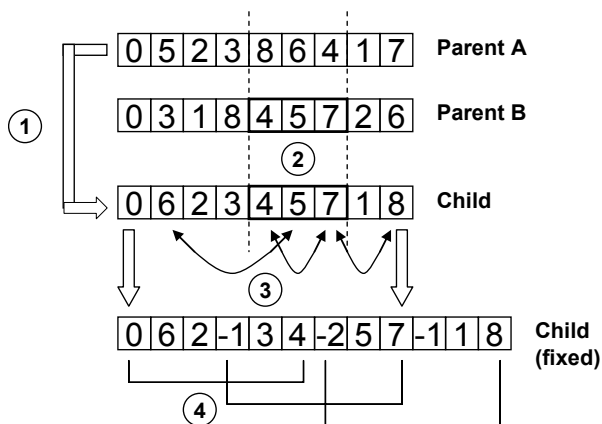


Figure 2: Genetic Crossover

Results

For the simple case of the 10-node graph presented (Figure 1), the genetic algorithm achieved a mean total distance of 32.44 over 5 runs, using a population size of 500 over 50 generations. The best solution was calculated by brute force, achieving a total distance of 31.21. A “by-eye” solution, generated by a colleague, achieved a total distance of 44.2.

Hence, in this simple case, the algorithm performed noticeably better than a human and generated a near-optimal solution. Increasing the population size and number of generations resulted in a lower mean total distance, even sometimes generating the optimal solution.

Merely increasing the problem to 20 nodes on a 20 by 20 grid (and keeping the depot-density the same), the problem becomes dizzying for the human eye. The algorithm provides a reasonable solution; however, the best solution was not generated due to prohibitive run-time and memory requirements.

Future Directions

As the number of generations was increased, higher variance was observed between runs. This suggests that perhaps a more conservative crossover technique should be used, swapping less DNA per generation and advancing the solution more slowly. Alternative selection techniques should be examined. A wider variety of site densities, depot densities, and fuel capacities should be rigorously investigated to determine if certain GA techniques work better under certain conditions. After fine-tuning, standard solutions for the more general problem should be implemented and measured against this solution to see if it is actually more efficient for this simplified situation.

References

- Bullnheimer B.; Hartl, R.; and Strauss, C. 1999. An Improved Ant System Algorithm for the Vehicle Routing Problem. *Annals of Operations Research* 89: 319-328.
- Cordeau, J.F.; Gendreau, M.; and Laporte, G. 1997. A Tabu Search Heuristic for Periodic and Multi-Depot Vehicle Routing Problems. *Networks* 30: 115-119.
- Machado, P.; Tavares, J.; Pereira, F.; and Costa, E. 2002. Vehicle Routing Problem: Doing it the Evolutionary Way. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO'2002)*, New York, USA, 9-13 July, 2002.
- Toth, P. and Vigo, D. eds. 2001. *The Vehicle Routing Problem*. Philadelphia, PA.: Siam Monographs on Mathematical Modeling.