Scatter Search And Bionomic Algorithms For The Aircraft Landing Problem

J. E. Beasley
Mathematical Sciences
Brunel University
Uxbridge UB8 3PH
United Kingdom
http://people.brunel.ac.uk/~mastjjb/jeb/jeb.html

Abstract: - The problem of deciding how to land aircraft approaching an airport involves assigning each aircraft to an appropriate runway, computing a landing sequence for each runway and scheduling the landing time for each aircraft. The objective is to achieve effective runway use. The multiple runway case of the static Aircraft Landing Problem is considered. Two heuristic techniques are presented: Scatter Search and the Bionomic Algorithm, population heuristic approaches that have not been applied to this problem before. Computational results are presented for test problems involving up to 500 aircraft and 5 runways.

Key-Words: aircraft landing, scatter search, bionomic algorithm

1 Introduction

The Aircraft Landing Problem is an important practical problem, since in many countries effective use must be made of limited runway capacity. Aircraft approaching an airport must be assigned landing times so as to respect landing time windows and too to respect the necessary separation times between successive landings. This must be done in such a way so as to ensure effective runway use. In this paper we present computational results showing that this can be done for relatively large problems.

2 Problem Formulation

In this section we (for conciseness) formulate the single runway Aircraft Landing Problem. The formulation for the multiple runway case is a simple extension of this (e.g. see [2]). Let:

- P be the number of aircraft
- E_i be the earliest landing time for aircraft i (i=1,...,P)
- L_i be the latest landing time for aircraft i (i=1,...,P)
- T_i be the target (preferred) landing time for aircraft i (i=1,...,P) .
- S_{ij} be the required separation time (≥ 0) between aircraft i landing and aircraft j landing (where aircraft i lands before aircraft j), i=1,...,P; j=1,...,P; $i \neq j$

The time window for the landing of aircraft i is hence $[E_i, L_i]$, where $E_i \leq T_i \leq L_i$.

The variables are:

x_i = the landing time for aircraft i (i=1,...,P)

δ_{ij} = 1 if aircraft i lands before aircraft j (i=1,...,P; j=1,...,P; $i \neq j$), =0 otherwise

The constraints associated with the problem are:

$$E_i \leq x_i \leq L_i \qquad i=1,...,P \qquad (1)$$

$$\delta_{ij} + \delta_{ji} = 1 \qquad i=1,...,P;\ j=1,...,P;\ j>i \qquad (2)$$

$$x_j \geq x_i + S_{ij} - M\delta_{ji} \qquad i=1,...,P;\ j=1,...,P;\ j \neq i \qquad (3)$$

where M is a large positive constant.

Equation (1) ensures that the scheduled landing time for each aircraft lies within its time window. Equation (2) ensures that either aircraft i lands before aircraft j ($\delta_{ij} = 1$) or aircraft j lands before aircraft i ($\delta_{ji} = 1$). Equation (3) ensures that the necessary separation time S_{ij} elapses between the landing of aircraft i and the landing of aircraft j when aircraft i lands before j ($\delta_{ij} = 1$, $\delta_{ji} = 0$). It is easy to show that we can replace M in equation (3) by $(L_i+S_{ij}-E_j)$, and this is convenient for computational reasons.

In this paper we use a nonlinear objective based on the difference between the scheduled landing time and the target time. We define $d_i = x_i - T_i$ as the deviation from the target time for aircraft i. If this deviation is positive, then the aircraft is landing after its target time. This is not an ideal situation and is so penalised, the corresponding contribution to the objective function for this aircraft is arbitrarily set to $-(d_i)^2$. On the other hand, if the deviation is negative, then the aircraft is landing before its target landing time which is preferred, so the corresponding contribution is set to $+(d_i)^2$. The objective is to maximise the overall aircraft contribution:

$$\text{maximise} \sum_{i=1}^{P} D_i \qquad (4)$$

where $D_i = -(d_i)^2$ if $d_i \geq 0$, $+(d_i)^2$ otherwise

3. Population heuristics

Population heuristics are based on the principles of selection and mutation, the main concepts of Darwin's theory of evolution. They mirror evolution in performing manipulations on individuals which represent possible solutions to the considered problem (e.g. see [1]). Each individual is encoded using a set of chromosomes that define the problem's variables. The fitness of an individual is evaluated with respect to the quality of the solution it represents. An initial population of individuals is generated and operators that model genetic selection, mating and other processes are defined and applied to the population individuals. The standard and most widely known population heuristics are genetic algorithms, whose general framework is:

- generate an initial population
- repeat
    - select individuals from the population to be parents
    - create new individuals as combinations of selected parents
    - optionally mutate the children
    - select the children to insert into the population
    - select the individuals to remove from the population
- until termination, whereupon report the best solution encountered

3.1 Scatter Search

The specific features of Scatter Search [5,6], as compared to a normal genetic algorithm, are:

- individuals are not limited to binary representation
- selected parents are not restricted to two
- parent mating is structured as a linear combination
- a local improvement procedure is applied to all individuals.

A general framework for Scatter Search is:

- generate the initial population called the reference set
- improve each individual in the reference set
- repeat
    - select a subset of the reference set
    - create a new individual as a linear combination of the subset
    - improve the new individual
    - update the reference set
- until termination, whereupon report the best solution encountered

3.2 Bionomic Algorithm

The Bionomic Algorithm is less well-known than Scatter Search and was first presented in [3]. As with Scatter Search, its underlying strategy involves creating new solutions as linear combinations of old solutions and procedures involved in this process are designed to use problem dependent knowledge. The specific features of a Bionomic Algorithm are:

- a maturation step to improve individuals
- structured construction of parent sets based on a graph which represents the population structure
- parent selection based on fitness and distance between individuals
- a generational approach to replace the population.

A general framework for the Bionomic Algorithm is:

- generate an initial population
- improve each individual in the initial population
- repeat
  - build a graph that represents the population structure
  - compute parent sets from this graph
  - create new individuals for each parent set as a linear combination of the members of the parent set
  - improve each new individual
  - update the population with some of the best new individuals
- until termination, whereupon report the best solution encountered

4. Algorithmic details

The main elements involved in the implementation of our Scatter Search and Bionomic Algorithm heuristics for the Aircraft Landing Problem are:

- representation and evaluation of an individual
- selection of parents
- generation of children
- local improvement of an individual

4.1 Individuals - representation and evaluation

As a key component of population heuristics are the individuals, the first step towards an implementation of such a heuristic is to define a representation adapted to the considered problem. An individual represents a possible solution to the problem. For the multiple runway Aircraft Landing Problem, an individual must provide information about the scheduled landing time and the allocated runway for each aircraft waiting to land.

Thus, the representation we define contains two variables for each aircraft: the proportion of time that elapses before the scheduled landing time and the runway allocated. This proportion y_i for aircraft i is defined by $y_i = (x_i-E_i)/(L_i-E_i)$. Below we show a representation of an individual with P aircraft where r_i is the runway allocated to aircraft i.

aircraft	1	2	…	P
proportion	y_1	y_2	…	y_P
runway	r_1	r_2	…	r_P

The fitness of an individual is evaluated with respect to its objective function value (equation (4)). Individuals who violate the constraints of the problem (equations (1)-(3)) are dealt with using unfitness [4]).

4.2 Parent selection for Scatter Search

Parent selection in our Scatter Search algorithm is a binary tournament selection scheme based on fitness. In binary tournament selection two individuals are selected at random from the population and the one with the best fitness is kept as a parent. Each execution of this selection scheme provides one individual to play the role of parent. Unlike standard genetic algorithms, the number of parents in Scatter Search can be more than two. For our implementation of Scatter Search, we set the number of parents to three. The binary tournament selection scheme is repeated three times to select three individuals. At the end of the parent selection stage for Scatter Search we have a single parent set containing three individuals.

4.3 Parent selection for the Bionomic Algorithm

Parent selection for the Bionomic Algorithm is a structured procedure where the underlying principle is to give more opportunities to individuals with a better fitness to be selected as parents and prevent individuals too close (in terms of solution structure) to each other from being selected as parents together.

Population structure is captured by an adjacency graph. Each individual of the current population is represented by a node. Nodes have an inclusion frequency value that is set to the rank (ties broken arbitrarily) of the corresponding individual when population individuals are sorted by fitness. Fitter individuals are represented by a node with a higher inclusion value and thus can appear more often in a parent set.

A distance measure is introduced to quantify, in terms of solution structure, how close two individuals are to each other. This distance measure is used to determine if an edge must exist between the two corresponding nodes in the adjacency graph.

In the Bionomic Algorithm, the adjacency graph yields parent sets through the computation of a maximal independent set. A maximal independent set is a set of maximal cardinality where none of the selected nodes are linked by an edge. The logic here is that constructing parent sets in this way means that individuals selected to play a role of parent are reasonably spanned over the problem search space, since an edge exists between two individuals i and j if and only if they have structurally similar solutions.

To illustrate the concept of a maximally independent set a possible adjacency graph for six nodes shown below. The maximal independent sets corresponding to this graph are {2,3,4,6}, {2,3,5} and {1,3,4}. In the case of parent set selection, it is sufficient to generate only one of the possible maximal independent sets. Note in particular that there is no requirement to generate the maximal independent set that contains the most nodes.

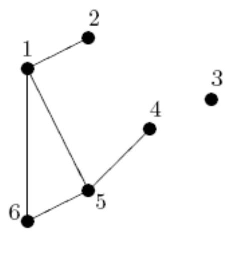

For the Bionomic Algorithm the parent selection stage generates many parent sets of various sizes. Sets are computed such that fitter individuals are more often selected than less fit individuals (by means of the inclusion frequency mechanism) and such that individuals in a given parent set are reasonably spanned over the problem search space (by means of the adjacency graph and maximal independent set mechanism).

4.4 Generation of children

The child generation procedure for both Scatter Search and the Bionomic Algorithm produces one child from each parent set, and this child may be feasible or infeasible. To create a new individual, a new set of proportion values (y_i i=1,…,P) must be computed and a new set of runway allocations (r_i i=1,...,P) must be determined. For each aircraft, the new proportion value is computed as a weighted linear combination of the corresponding parent proportion values. Random weights are used in the linear combination process in order to introduce diversity to the new individual. Runway allocation is decided in an equally weighted probabilistic manner for each aircraft.

The resulting solution may be infeasible. This does not have serious consequences to the quality of the population as the improvement procedure (see below) is applied to any new individual and, in any event, infeasible solutions can be dealt with using unfitness.

4.5 Individual local improvement

Scatter Search and the Bionomic Algorithm call for a local improvement of each individual that is based on problem dependent knowledge. For our nonlinear objective, it is possible (via a simple procedure) to compute optimal landing times for the set of aircraft allocated to each runway - provided we regard the sequence in which aircraft land as fixed. This polynomially bounded procedure is applied independently to each runway. If it computes landing times that are feasible (satisfy the time window and separation time constraints, equations (1) and (3)) then these landing times are optimal for the given landing sequence considered.

4.6 Heuristic overview

In order to provide a complete overview of our Scatter Search and Bionomic Algorithm heuristics for the Aircraft Landing Problem we give below a framework for the steps used:

- randomly generate the initial population
- improve and evaluate each individual in the population as in Sections 4.5 and 4.1
- repeat
    - generate parent sets as in Section 4.2 for Scatter Search or as in Section 4.3 for the Bionomic Algorithm
    - generate children from the parent sets as in Section 4.4
    - eliminate any children that are duplicates of population members
    - improve and evaluate each (remaining, non-duplicate) child as in Sections 4.5 and 4.1
    - add the best child to the population
- until termination, whereupon report the best solution encountered

5 Computational results

The Scatter Search and Bionomic Algorithm heuristics presented above were implemented in C++ on a 2GHz Pentium PC with 512MB of memory. Computational results are presented in this section for instances, publicly available from OR-Library (people.brunel.ac.uk/~mastjjb/jeb/info.html), involving from 10 to 500 aircraft, and up to 5 runways.

For 29 small instances ($P \leq 50$), involving up to 5 runways, Scatter Search gave an average percentage deviation from the best known solution of 0.9% in an average computation time of 13 seconds. The Bionomic Algorithm gave an average percentage deviation from the best known solution of 0.3% in an average computation time of 65 seconds.

For 23 large instances ($100 \leq P \leq 500$), involving up to 5 runways, Scatter Search gave an average percentage deviation from the best known solution of 3.4% in an average computation time of 244 seconds. The Bionomic Algorithm gave an average percentage deviation from the best known solution of 2.9% in an average computation time of 540 seconds.

We would note here that problems of the size discussed above are much larger than those that have been considered by other authors in the literature.

6 Conclusion

In this paper we have considered the Aircraft Landing Problem. Computational results indicated how the two Scatter Search and Bionomic Algorithm heuristics developed performed on instances involving up to 500 aircraft and 5 runways.

References

[1] J.E. Beasley, "Population heuristics", in "Handbook of applied optimization", P.M. Pardalos and M.G.C. Resende (eds), Oxford University Press, (2002) 138-157.

[2] J.E. Beasley , M. Krishnamoorthy, Y.M. Sharaiha and D. Abramson, “Scheduling aircraft landings - the static case", Transportation Science, 34 (2000) 180-197.

[3] N. Christofides, "The Bionomic Algorithm", Paper presented at the AIRO'94 conference, Savona, Italy (1994).

[4] P.C. Chu and J.E. Beasley, “Constraint handling in genetic algorithms: the set partitioning problem", Journal of Heuristics, 4 (1998) 323-357.

[5] F. Glover, "Genetic algorithms and scatter search: unsuspected potentials", Statistics and Computing, 4 (1994) 131-140.

[6] F. Glover, "Scatter search and star-paths: beyond the genetic metaphor", OR Spektrum, 17 (1995) 125-137.