

Scheduling On Demand Passenger Air Service

TRISTAN VI

June 2007

George L. Nemhauser

School of Industrial and Systems Engineering

Georgia Institute of Technology

Abstract

A non-scheduled airline provides regional "on demand" air transportation on small jet planes having a capacity of 3 passengers. A request for travel specifies an origin, destination, earliest departure time, and latest arrival time. Based on requests already accepted for that day, the accept/reject problem is to determine whether the new request can be accommodated. At the beginning of the day an optimal schedule that minimizes flying time is created for all of the accepted requests for that day.

DayJet will begin providing this service in 2007 using the Eclipse 500, which costs about one million dollars, is fuel efficient and has a range of over a thousand miles. They expect to have several hundred jets covering overlapping regions of the U.S. in a few years. This service is especially useful for areas that are not well served by large airports. By using small airports, they eliminate the hassles associated with long drives to the airport, packed parking lots, security lines, etc. For many travelers, this service will yield huge time savings in comparison to the alternatives of a scheduled airline or driving.

In this talk, I will discuss the optimization models and algorithms that we have developed for scheduling DayJet's forthcoming service. These include a multi-commodity flow model with side constraints for solving small instances, which is imbedded into a local search algorithm with an asynchronous parallel implementation for solving real-life instances.