From: AAAI-99 Proceedings. Copyright © 1999, AAAI (www.aaai.org). All rights reserved.

Worldwide Aeronautical Route Planner

Charles B. McVey David P. Clements Barton C. Massey Andrew J. Parkes

Computational Intelligence Research Laboratory
University of Oregon
Eugene, OR 97403-1269
{mcvey, clements, bart, parkes}@cirl.uoregon.edu

Introduction

We consider the common problem of calculating routes from a starting point to a destination through a given space. This process typically involves discretizing the navigational space into a graph of intermediate waypoints linked together through transitions. A search for a solution that fits desired criteria can then be performed. Typical criteria include: minimum distance, minimum transit time, minimum fuel use, and maximum agent safety.

Our work in this area focuses on the rapid determination of minimal fuel routes for aircraft to fly from any source point to any destination point on the earth. The Worldwide Aeronautical Route Planner (WARP) is our prototype demonstration of technology developed to perform optimal real-time (under one minute per route) flight planning. WARP uses high-fidelity aircraft flight models, standard rules of flight operation, actual time-dependent worldwide weather data, and a variant of A* (Multi-Pass A*) which retains A*'s completeness and optimality properties. WARP discovers the flight plan that guarantees minimal use of fuel. Typical routes are generated by WARP in under thirty seconds, with fuel savings on the order of one to eight percent compared to great circle (minimum distance) routes.

Modeling

WARP's model of airspace represents aircraft position (in time, altitude, and latitude/longitude coordinates) as well as weather. In direct flight, with no zones of avoidance or predefined airways which must be followed, the USAF Air Mobility Command's (AMC) standard operating procedure (SOP) allows aircraft to change heading only when over integral latitude/longitude points in which either the latitude is evenly divisible by five or the longitude is evenly divisible by ten. The route generated by WARP must conform to this restriction, leading to a waypoint discretization of five degree by ten degree zones upon the surface of the earth. On the borders of these sections, every integral degree point is a potential waypoint for WARP to consider during its search. Depending upon whether a waypoint is located on a five degree latitude line, ten degree longitude line, or an intersection of the two, the number of two dimensional neighbors for that waypoint will be forty, fifty, or sixty respectively.

Transitions in altitude are allowed among eighteen discrete levels, and are subject to another SOP rule: an aircraft may only change its altitude if it has flown at least 400 miles since its last altitude change. This has the benefit of restricting the search space. The aircraft performance models also restrict altitude changes by not allowing the aircraft to ascend or descend at a rate which is not physically possible, or to ascend to an altitude which the aircraft at current weight is not capable of reaching.

Weather representation within WARP is discretized into cells. Each cell is 2.5 degrees in latitude by 2.5 degrees in longitude, and contains the temperature, wind direction, and wind speed for a specific altitude at a specific point in time. Forecast data are available covering fifteen different altitude ranges at ten different times over a three day period. Since the weather cells are of different size than, and offset from, the latitude/longitude and altitude discretized grid, weather effects are calculated using a weighted interpolation scheme during the search.

Multi-Pass A*

A* is a well-known and well-studied best-first search algorithm (Hart, Nilsson, and Raphael 1968). It typically searches outward from the starting node until it reaches the goal node. At each step the current fringe node with the minimum combined actual cost from the start to the fringe node plus expected remaining cost (the heuristic cost) to get from the fringe node to the goal node is expanded. If the heuristic always underestimates the cost from any node to the goal, then it is said to be admissible, and A* will find the optimal solution. However, if the heuristic is too optimistic then A* will expand too many nodes and may run out of time before a solution is found. If the heuristic is pessimistic then a solution will likely be

---

found quickly, but it will almost certainly be suboptimal. In this work, we must generate optimal routes as quickly as possible, so our heuristic function must be optimistic, but as close to the real cost to the goal as possible.

Most A* work has assumed that the cost of evaluating an edge in the search graph is minimal, and that the primary goal is to expand as few nodes as possible. However, in realistically modeled flight planning problems, the number of nodes is relatively small (on the order of one quarter million), but the cost of evaluating edges is extremely expensive. We therefore seek to keep the number of expensive edge evaluations done to a minimum. This is related, but not identical, to minimizing the number of nodes expanded.

Fortunately, a “cheap” heuristic is available to guide the search; it assumes that the remaining distance from a node to the goal will be flown under the best weather conditions that can possibly be reached on the current flight. The heuristic is admissible, but is typically optimistic by approximately twenty percent. While the heuristic successfully prunes much of the search space, using it with A* and the exact flight performance models still requires excessive computation time; it needs to be less optimistic.

To overcome this problem, we developed and implemented a two-pass version of A* (Multi-Pass A*) which uses approximate, cached “cheap heuristic” flight performance values during an initial *backward* pass (it searches from the destination to the point of departure), and exact values during a final *forward* pass. We determine the initial time and fuel weight for the backward search by first “flying” a great circle route from the origin to destination and using the resulting time and weight at arrival.

During the backwards search, we *decrement* time and *increment* fuel in the aircraft. The cost of flying backwards from the destination to each waypoint is recorded as the new heuristic estimate for the cost to reach the destination from that particular waypoint. Nodes are expanded in minimum cost order, with the total great-circle-route cost serving as a bound to prune those nodes with a larger cost plus heuristic estimate (Dijkstra 1959). The results of this initial backward pass then guide a second, much tighter, forward pass. In this forward search, we utilize the expensive exact flight-model equations. However, far fewer nodes are explored since the heuristic values were well tightened by the backward seeding.

This approach does not decrease the number of edges evaluated, but it greatly reduces the average cost of edge evaluation. The majority (typically over eighty percent) of nodes expanded during the entire search process are expanded during the backward pass in which cheaper approximate values are used. As a result, the forward search runs in a fraction of the time taken by the backward search, despite its use of expensive performance models.

The existing system for route planning in use by the AMC employs search that is bounded by a narrow window around the great circle route. This approach can obviously never guarantee to find optimal solutions. In comparisons with WARP in which the optimal route was near the great circle, both systems produce similar results. However, in many cases WARP finds optimal routes that deviate enough from the great circle that the existing system’s search method never discovered them. The improvements in these cases translate into enormous savings when one considers the amount of fuel used by the AMC annually.

The Demonstration

WARP’s user interface displays a plot of the earth’s surface, allowing the user point-and-click selection of source and destination coordinates, as well as straightforward specification of departure time, starting fuel (or ending fuel), and payload weight. The user may also pick a variety of different search strategies, including multi-pass A* or single pass A*, and use of approximate or exact aircraft modeling values. During route generation, the display first shows the nodes expanded by the backward bounded Dijkstra pass, followed by the nodes expanded by the highly optimized forward A* search.

After route generation, the solution is displayed along with the great-circle-route from source to destination, allowing easy visual comparison. Statistics are also presented, including the fuel usage and distance traveled along both the optimal and the great-circle routes, the CPU time required to generate the solution, and the number of nodes and edges expanded during each phase of the search.

The route consists of a list of waypoints and headings to follow during the flight, along with information such as time, altitude, and fuel load at each waypoint. A list of waypoints is shown in a list panel on the display, and an altitude profile for the flight is shown in the bottom panel of the interface. The user may click on various parts of the display to solicit detailed information about the solution.

Acknowledgments

We wish to thank the members of CIRL for their suggestions and support. We also thank Majors Adrian Hayes and Mark Weiser (USAF) and John Will (NCI) for sharing expertise in this domain. Efforts sponsored by DARPA and AFRL, Air Force Materiel Command, USAF, under agreement number F30602-95-1-0294. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, AFRL, or the U.S. Government.

References

Dijkstra, E. 1959. A note on two problems in connection with graphs. *Numerische Mathematik* 1:269-271.

Hart, P., Nilsson, N., and Raphael, B. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on SCC* 4:100-107.