

Probabilistic Roadmap Based Path Planning for an Autonomous Unmanned Aerial Vehicle

Per Olof Pettersson and Patrick Doherty

Linköping University

Department of Computer and Information Science

581 83 Linköping, Sweden

{peope, patdo}@ida.liu.se

Abstract

The emerging area of intelligent unmanned aerial vehicle (UAV) research has shown rapid development in recent years and offers a great number of research challenges for artificial intelligence. For both military and civil applications, there is a desire to develop more sophisticated UAV platforms where the emphasis is placed on development of intelligent capabilities. Imagine a mission scenario where a UAV is supplied with a 3D model of a region containing buildings and road structures and is instructed to fly to an arbitrary number of building structures and collect video streams of each of the building's respective facades. In this article, we describe a fully operational UAV platform which can achieve such missions autonomously. We focus on the path planner integrated with the platform which can generate collision free paths autonomously during such missions. It is based on the use of probabilistic roadmaps. The path planner has been tested together with the UAV platform in an urban environment used for UAV experimentation.

In Proceedings of the Workshop on Connecting Planning and Theory with Practice. ICAPS'2004.

Introduction

The emerging area of intelligent unmanned aerial vehicle (UAV) research has shown rapid development in recent years and offers a great number of research challenges for artificial intelligence. Much previous research has focused on low-level control capability with the goal of developing controllers which support the autonomous flight of UAVs from one way-point to another at high altitudes. The most common type of mission scenario involves placing sensor payloads in position for data collection tasks where the data is eventually processed off-line or in real-time by ground personnel. Use of UAVs and mission tasks such as these have become increasingly more important in recent conflict situations and are predicted to play increasingly more important roles in any future conflicts.

Intelligent UAVs will play an equally important role in civil applications. For both military and civil applications, there is a desire to develop more sophisticated UAV platforms where the emphasis is placed on development of intelligent capabilities. Focus in research has moved from

Copyright © 2004, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

low-level control towards a combination of low-level and decision-level control integrated in sophisticated software architectures. These should also integrate well with larger net-centric based C⁴I² systems. Such platforms are a prerequisite for supporting the capabilities required for the increasingly more complex mission tasks on the horizon and an ideal testbed for the development and integration of AI technologies.

The WITAS¹ Unmanned Aerial Vehicle Project (Doherty *et al.* 2000) is a long-term basic research project whose main objectives are the development of an integrated hardware/software VTOL (Vertical Take-Off and Landing) platform for fully autonomous missions and its deployment in applications such as traffic monitoring and surveillance, emergency services assistance, photogrammetry and surveying.

Basic and applied research in the project covers a wide range of topics which include the development of a distributed architecture for autonomous unmanned aerial vehicles. In addition to the software architecture, many AI technologies have been developed such as path planners, task planners, chronicle recognition and situational awareness techniques. The architecture supports modular and distributed integration of these and any additional functionalities added in the future.

An experimental version of the WITAS UAV hardware/software platform has been developed and successfully used in a VTOL system capable of achieving a number of complex autonomous missions flown in a challenging *urban* environment populated with building and road structures. In one mission, our UAV autonomously tracked a moving vehicle for up to 20 minutes. In another, several building structures in the test area were arbitrarily chosen as survey targets and our UAV autonomously generated collision free path plans to fly to each and take photographs of each of the building's facades. This and similar missions have been successfully executed.

Figure 1 shows an aerial photo of our primary flight test area located in Revinge, Sweden. An emergency services training school is located in this area and consists of a col-

¹WITAS (pronounced *vee-tas*) is an acronym for the Wallenberg Information Technology and Autonomous Systems Laboratory at Linköping University, Sweden.




Figure 1: Aerial photo over Revinge, Sweden

lection of buildings, roads and even makeshift car and train accidents. This provides an ideal flight test area for experimenting with traffic surveillance, photogrammetric and surveying scenarios, in addition to scenarios involving emergency services. We have also constructed an accurate 3D model for this area which has proven invaluable in simulation tests and as a visualization tool. Parts of the model are integrated in the on-board geographic information system (GIS) and are used by many of the services in the architecture including the path planner which will be considered in detail in this paper.

In the remainder of the paper, we will concentrate on a solution to path planning for our UAV based on the use of probabilistic roadmaps. This path planning module is implemented and used in the on-board system. Before providing details, the hardware and software platforms for the WITAS UAV will be considered briefly in order to provide a context for understanding how path planning is integrated with the system.

The VTOL and Hardware Platform

The WITAS Project UAV platform we use is a slightly modified Yamaha RMAX (figure 2). It has a total length of 3.6 m (incl. main rotor), a maximum take-off weight of 95 kg, and is powered by a 21 hp two-stroke engine. Yamaha equipped the radio controlled RMAX with an attitude sensor (YAS) and an attitude control system (YACS). Figure 3 shows a high-level schematic of the hardware platform that we have built and integrated with the RMAX platform. The hardware platform consists of three PC104 embedded computers (figure 3).²

²The primary flight control (PFC) system consists of a PIII (700Mhz) processor, a wireless modem (serial line RS232C) and the following sensors: an integrated INS/DGPS (serial), a barometric altitude sensor (analog), a sonar and infrared altimeter (analog), and a compass (serial). It is connected to the YAS and YACS (serial), the image processing computer (serial) and the deliberative computer (Ethernet). The image processing (IP) system consists of a second PC104 embedded computer (PIII 700MHz), a color CCD camera (S-VIDEO, serial interface for control) mounted on



Figure 2: The WITAS RMAX Helicopter

The Software Platform

CORBA³ has been chosen as a basis for the design and implementation of a loosely coupled distributed software architecture for the WITAS aerial robotic system (Doherty *et al.*). It is believed that this is a good choice which enables us to manage the complexity of a deliberative/reactive (D/R) software architecture with as much functionality as we require for our applications. It also ensures clean and flexible interfacing to the deliberative and control components in addition to the hardware platform via the use of IDL (Interface Definition Language).

In short, CORBA (Common Object Request Broker Architecture) is middleware that establishes client/server relationships between objects or components. A component can be a complex piece of software such as a path planner, or something less complex such as a task procedure which is used to interface to helicopter or camera control. Objects or components can make requests to, and receive replies from, other objects or components located locally in the same process, in different processes, or on different processors on the same or separate machines. In our case, we have three onboard PC104s in addition to ground station computers.

Many of the functionalities which are part of the architecture can be viewed as CORBA objects or collections of objects, where the communication infrastructure is provided by CORBA facilities and other services such as real-time and standard event channels. This architectural choice provides us with an ideal development environment and versatile run-time system with built-in scalability, modularity, software relocatability on various hardware configurations, performance (real-time event channels and schedulers), and

a pan/tilt unit (serial), a video transmitter (composite video) and a recorder (miniDV). The deliberative/reactive (D/R) system runs on a third PC104 embedded computer (PIII 700MHz) which is connected to the PFC system with Ethernet using CORBA event channels and standard CORBA method calls.

³TAO/ACE (Object Computing, Inc. 2000) is currently being used. The Ace Orb is an open source implementation of CORBA 2.6.




Figure 3: On-Board Hardware Schematic

support for plug-and-play software modules.

Figure 4 depicts an (incomplete) high-level schematic of some of the software components used in the architecture. Each of these may be viewed as a CORBA server/client providing or requesting services from each other and receiving data and events through both real-time and standard event channels. Each of these functionalities has been implemented and all are being used and developed in our applications. A great deal of effort has gone into the devel-




Figure 4: Some deliberative, reactive and control services

opment of a control system for the WITAS UAV which incorporates a number of different control modes and includes a high-level interface to the control system. This enables other parts of the architecture to call the appropriate control modes dynamically during the execution of a mission. The ability to switch modes contingently is a fundamental functionality in the architecture and can be programmed into the task procedures associated with the reactive component in the architecture. We have developed and tested the following autonomous flight control modes:

- take-off (TO-Mode) and landing via visual navigation (L-Mode)
- hovering (H-Mode)

- dynamic path following (DPF-Mode)
- reactive flight modes for interception and tracking (RTF-Mode).

These modes and their combinations have been successfully demonstrated in a number of missions at the Revinge flight test area.

A Path Planning Mission Scenario

One application area we have focused on is surveying and data collection. The WITAS UAV contains an on-board GIS which includes a terrain model of the Revinge area accurate to within decimeters in the x , y , z directions and equally accurate models of building and road structures. A plan template has been designed using a task procedure which can be parameterized with an arbitrary number of building structure identifiers from Revinge. When this task procedure is called, it iterates through each of the building identifiers and generates a set of waypoints (one for each building facade) where the UAV should be positioned to take photos of each respective facade. The waypoints are computed using additional functionality built into the GIS. In addition to the waypoint positions for the UAV, camera positions are also computed.

The UAV must fly from building structure to building structure and while completing the surveying task for a building, it must also fly to each of the waypoints associated with the facades, hover and yaw appropriately and take pictures of the facade before moving onto the next. To fly from one waypoint to another it calls a path planning service which generates a collision free path between waypoints. The resulting 3D paths between any two waypoints may consist of a segmented trajectory. Each of these segments is passed to functional units in the primary flight controller which call an appropriate flight mode to follow the trajectories. Progress is monitored by the calling task procedure.

This is a complex scenario which must use deliberative (path planner), reactive (task procedures) and control (trajectory following, hovering modes) functionalities concurrently in an integrated manner. What makes the problem even more difficult is that in general, computing a collision free path between two waypoints is intractable.

In the next two sections, we show how this problem has been solved efficiently using probabilistic roadmaps as a basis for the path planner. We then consider actual survey missions similar to the above that have been executed successfully by our UAV.

The Path Planner

The path planner used for the helicopter is an adaptation of probabilistic roadmap (PRM) algorithms (Kavraki *et al.* 1996), to our application domain. The problem of finding an optimal path between two robot configurations in a configuration space such as Revinge is intractable. For static configuration spaces where one assumes no dynamic objects except cars on the ground and a helicopter, PRM algorithms *hedge* the intractability problem by outputting non-optimal paths and working in two phases, one off-line and the other

during runtime. The main processing stages for our adaptation of the PRM algorithm are shown in figure 5. In the




Figure 5: Generating a Flight Trajectory

offline phase, a roadmap graph is generated for the area of interest (e.g.. Revinge). This process takes a 3D polygonal model of the area and the helicopter kinematics as input. Helicopter configurations⁴ are randomly generated and checked for collisions with the model. An attempt is then made to connect collision-free configurations using a local path planner which generally takes into account the kinematic and dynamic constraints of the helicopter. Each of the local paths generated also have to be checked for collisions. The collision checker, used to check whether a given curve or line intersects any obstacle in the environment, is based on the OBBTree-algorithm (Gottschalk, Lin, & Manocha 1996a), that uses a tree of oriented bounding boxes as its central component.

There are a number of choices that can be made for the local path planner at this stage dependent on how much or little work is chosen to be done during run-time. In the first case, which is more in keeping with the original PRM algorithm, the roadmap is generated in the off-line phase with spline-curves and non-holonomic constraints are taken into account. Another alternative, the one described in the diagram, is to initially ignore the non-holonomic constraints of the helicopter in the off-line phase and add them as refinements to the plan in smoothing and curve replacement phases during run-time (Sekhavat *et al.* 1998). Additional constraints may also influence which parts of the roadmap are legally usable.

We have experimented with variations of both approaches. Using these techniques, roadmaps can be generated with high-levels of coverage even in tight spaces. Using straight lines, an initial roadmap for Revinge was generated containing 5000 nodes with 97% coverage. Using spline-curves, an alternative roadmap was generated using 15000

⁴A helicopter configuration currently consists of three coordinates for position and one angle for direction. The orientation of the helicopter is omitted and is handled independently by the control system.

nodes with 70% coverage.

During the mission or run-time phase, the path planning service is called with an initial and goal helicopter configuration. An attempt is made to connect the two configurations to the previously generated roadmap using the local path planner and, if successful, an A* search is used on the graph to generate a multi-segmented trajectory. The resulting straight lines are smoothed and if possible replaced with spline curves using the local path planner. One may also add additional runtime constraints.

Some Modifications to the Standard Probabilistic Roadmap Algorithm

One of the most important modifications we have made to the standard PRM algorithm is to delay some of the processing of constraints normally done off-line in the roadmap generation stage and instead do the processing during run-time when path plans are actually being generated. As long as it is not too expensive (meaning slow) to process certain types of constraints during runtime, this approach is of great benefit since one can dynamically apply constraints rather than building them into the roadmap in the off-line phase. We consider two modifications to the PRM algorithm which have been incorporated into our path planner.

Multi-level non-holonomic roadmap planning The standard probabilistic roadmap algorithm is formulated for fully controllable systems only. This assumption is more or less true for a helicopter flying at low speed in a hovering mode. However, when the speed is increased the helicopter is no longer able to negotiate turns of too small a radius, which imposes similar demands on the planner as the non-holonomic constraints on car-like robots.

There are a number of proposals for dealing with non-holonomic constraints in the probabilistic roadmap paradigm. One approach taken in (Sekhavat, P. Svestka, & Overmars 1996), is to first solve a relaxed problem using only the holonomic constraints, and then refine the solution by adding non-holonomic constraints one at a time. For robots respecting certain topological properties (e.g. all locally controllable robots as well as car-like robots), it is always possible to upgrade the solution so as to take the non-holonomic constraints into account as long as there exists a non-zero margin between the solution to the relaxed problem and the obstacles.

Inspired by this approach, the path planner for our UAV initially generates a piecewise linear plan, which is later refined to the preferable piecewise cubic curve required for smooth high speed flight. If the planner is unable to do this curve replacement, the original linear path segment is left as is, which will require the helicopter to go into hover mode at the sharp corners that arise if two or more linear segments remain in succession. In practice, this situation rarely arises.

Delayed Constraint Handling The delayed handling of constraints can also be used for dealing with constraints that are not known during roadmap construction time. This approach is used with our path planner. New constraints can be introduced at the time of a query for a path plan. These con-

straints should be efficiently evaluable since rapid response time is so much more important during the runtime query phase. Currently we have implemented runtime constraints pertaining to maximum and minimum altitude, forbidden regions and limits on ascent-/descent-rate, which are useful when setting up UAV missions in the field. In fact, this extension to PRMs is necessary in practical applications. Mission constraints rapidly change and one can not afford to recompile the roadmap to incorporate such constraints. This extension to roadmaps grew out of a practical necessity as we experimented with PRMs in the field and found the approach lacking in this respect.

Local Path Planner and Helicopter Controller

The control mode that carries out the plan produced by the path planner takes a cubic polynomial space curve parameterized from 0 to 1 as its input. The curve is derived from the two end positions and the direction of flight through these points. This leaves one degree of freedom at each end in the magnitude of the derivative, which is currently set to the distance between the two points and which generally produces nice curves. In the future, the magnitude of this vector can be used to adjust the curvature of the curve segment to suit the requested flight speed.

Even if the trajectory following controller is able to fly most well-behaved curves, there are a number of limits in the physical platform and the current controller implementation that makes different curve-forms more or less effective to fly. The responsibility of staying inside these limits is shared between the controller, the task procedure that calls the path planner and the path planner, while most of the constraints are handled by the controller.

The limitations include maximum acceleration which puts a limit on how fast the helicopter can fly along a curve with a certain radius as well as ascent rate. The descent rate is even more limited due to aerodynamic effects. Limitations that stem from the system architecture include a timeout limit which is a safety feature in the controller that requires the next curve segment to be ready some time in advance in order to be able to stop in time if no next segment arrives.

Currently the path planner is only aware of the timeout limitation on speed while the other speed limits are imposed dynamically by the controller. Using A* search, the path planner currently optimizes only on shortest distance but we are planning to incorporate the limitations mentioned in a flight time estimate so that the path planner can also optimize on flight time. The trajectory-following control-mode used is described in more detail in Conte (Conte, Duranti, & Merz 2004) and the low-level control architecture is described in Merz (Merz 2004).

Collision Checker

Both during roadmap-construction and online path planning queries, possible paths have to be tested for collision.

The collision checking algorithm used for the path planner in the WITAS-project is based on the OBB-tree algorithm presented in (Gottschalk, Lin, & Manocha 1996b). The OBBTree-algorithm constructs a tree of bounding boxes around the obstacles in the environment by including all

polygons in the root-box and then recursively dividing the polygons into smaller and smaller boxes. The orientations of the bounding boxes are determined by doing a principal component analysis on the vertices.

However, in addition to checking for overlap between two stationary objects, the collision checker must also be able to check that a full path between two helicopter configurations does not collide with any obstacles.

Since the helicopter is quite small compared to the environment, a simplifying assumption has been made where the helicopter is regarded as point-object. The actual size of the helicopter is instead added to a safety margin that is placed around the obstacles in the 3D Revinge world model. If the helicopter is regarded as a point-object, it suffices to check if the cubic polynomial describing a flight path intersects an OBB. This can be done by analytically solving the intersection points for the cubic curve and the bounding planes of the OBB.

The 3D Revinge model we are using is large. It covers an area of $800 \times 800 \text{ m}^2$ and consists of roughly 140,000 polygons of which 120,000 polygons represent ground terrain. The remaining 20,000 polygons represent other types of obstacles, mainly buildings and trees.

Since the majority of polygons in the model represent the ground terrain and the majority of missions are not performed very close to the ground, the depth of the OBB-tree for terrain can be limited to 10. Using this cut-off-value, the generation of the OBB-tree takes approximately 200 seconds, while collision-checking against a typical curve-segment of 40 m length takes 22 milliseconds. In the case of take-off and landing which obviously takes place near terrain, visual navigation and GPS techniques are used for maneuvering close to the ground.

The Mission Planning procedure

The path planning module is used in the following manner during missions. At any one time during flight, a number of task procedures are running which call control and other functionalities as needed. During missions over Revinge when the UAV needs to fly from one point to another, the task procedure in control will call the path planning module with its current position and the position to which it wants to fly. In addition, the task procedure may also provide additional runtime constraints on the path such as minimum or maximum altitude or forbidden flight areas.

The path planning module then attempts to resolve the request by first connecting the start and end position to the roadmap which is stored onboard and then applying the A* algorithm with an appropriate search policy such as shortest distance. The run-time constraints are resolved by eliminating nodes and edges that violate them during the search.

The resulting piecewise linear curve is then subjected to a series of smoothing steps in order to reduce the jagginess that is often the result of randomized path planning algorithms. All these smoothing steps are done by applying the smoothing operation on a certain segment or waypoint and then checking if the resulting path is still collision-free.

The most important of these are the replacement of the linear segments with cubic space curves that makes it pos-

sible to join the segments with a continuous first derivative which is required to fly through the waypoints at any greater speed. In principle C^2 -continuity (continuous 2nd-order derivative) could also be achieved but this would make all segments of the curve (if we restrict ourselves to cubic functions) interdependent which would make the replacement more complicated and demanding to apply.

Other smoothing operations that are applied are elimination of waypoints which results in fewer and longer path segments, and stepwise alignment of the waypoints to make the path straighter.

Missions

The path planner described in this paper is fully integrated and part of our on-board software architecture. It is possible to use the path planner in a fully autonomous manner or to use it in combination with a ground operator for interactive missions. Although, we focus on flight in Revinge, it is straightforward to add and use 3D models from other regions and take advantage of this technology. In this section, we briefly describe some of the missions that have been flown using the path planner functionality.

The first mission involves starting at home base and flying to two buildings in Revinge and photographing each of its facades. For any given building structure in Revinge, the on-board GIS has functionality to generate optimal waypoints to fly to for each facade while taking contextual information such as sun position into account. In figure 6 and 7, the path generated by the path planning module (white) is shown together with the actual plotted log-data from the real flight (black).




Figure 6: Autonomous survey mission to building1 and then building2 photographed from each side.

The second mission is of an interactive character where the operator determines where the UAV should fly and where its camera should point by sending commands from the ground. In between commands the helicopter will simply hover autonomously.

Figures 8 and 9 depict both the planned path (white) and actual logged data (black) from a mission where the UAV starts at home base and then flies to a building with a garage door. The UAV is instructed to point the camera at the garage




Figure 7: The same mission plotted against a skeleton map.

door. The ground operator then decides to continue flight to a junk yard to gather additional data. The path planner is automatically called and generates a collision free path to fly to this point. Additional commands are sent to point the camera autonomously. Before returning to home base, the operator commands the UAV to fly to a point on the other side of a nearby building. In all these cases the helicopter dynamically finds a collision free flight trajectory that is shown to the operator while the helicopter is hovering autonomously. If the operator acknowledges the trajectory as acceptable, the flight is executed by the helicopter.



Figure 8: An interactive mission with flight in clockwise direction from building1 with garage to junkyard to behind building2 and then back to home base.

Conclusions

In this paper, we described the use of the probabilistic roadmap path planning paradigm for an unmanned aerial vehicle application. During the development and experimental use of our PRM path planning prototype, a number of limitations to the standard PRM approach were identified.




Figure 9: The same mission plotted against a skeleton map.

One of the most important issues that arose relates to how one deals with non-holonomic constraints which arise in the interaction between the path planner and the low-level flight controller. In this particular case, the most problematic aspect has been to respect the non-holonomic constraints required to achieve smooth transitions between different trajectory segments, while retaining an efficient planning procedure. We have observed that it is useful to postpone the non-holonomic constraints to the run-time stage in order to reduce the dimensionality of graph generation during the off-line stage.

Another important issue that arose is in regard to the efficiency tradeoff between what is done off-line and what is done on-line. It appears that flexibility as regards adjustment to contingent changes in the UAV environment during plan execution time is traded off against efficient runtime planning based on a PRM generated offline. Since many of our mission scenarios would involve dealing with runtime contingencies such as new no-fly zones or additional visibility constraints, we have moved towards a lazy PRM philosophy.

Ideally, one would like a flexible and uniform means of adding mission constraints both during the off-line and on-line stage. The extent to which this is done may very well be mission dependent and based on the preferences of ground operators. The interaction between constraints that are known in advance and run-time constraints and how both can be incorporated in a planner is an issue that we are currently pursuing.

Acknowledgments

This work is funded in part by grants from the Wallenberg Foundation and COMPAS NFFP nr-539.

References

Conte, G.; Duranti, S.; and Merz, T. 2004. Dynamic 3D path following for an autonomous helicopter. In *Proc. of the 5th IFAC Symposium on Intelligent Autonomous Vehicles*.

Doherty, P.; Haslum, P.; Heintz, F.; Merz, T.; Nyblom, P.; Persson, T.; and Wingman, B. A distributed architecture for autonomous unmanned aerial vehicle experimentation. Submitted 2004.

Doherty, P.; Granlund, G.; Kuchcinski, K.; Nordberg, K.; Sandewall, E.; Skarman, E.; and Wiklund, J. 2000. The WITAS unmanned aerial vehicle project. In *Proceedings of the 14th European Conference on Artificial Intelligence*, 747–755.

Gottschalk, S.; Lin, M.; and Manocha, D. 1996a. Obbtree: A hierarchical structure for rapid interference detection. In *Proc. of the 23rd Int'l. Conf. on Computer graphics and interactive techniques*, 171–180.

Gottschalk, S.; Lin, M.; and Manocha, D. 1996b. Obbtree: A hierarchical structure for rapid interference detection. Technical Report TR96-013, Dept. of Computer Science, Univ. of N. Carolina, Chapel Hill.

Kavraki, L. E.; Svestka, P.; Latombe, J.-C.; and Overmars, M. 1996. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Trans. on Robotics and Automation* 12(4):566–580.

Merz, T. 2004. Building a system for autonomous aerial robotics research. In *Proc. of the 5th IFAC Symposium on Intelligent Autonomous Vehicles*.

Object Computing, Inc. 2000. *TAO Developer's Guide, Version 1.1a*. See also <http://www.cs.wustl.edu/~schmidt/TAO.html>.

Sekhavat, S.; Švestka, P.; Laumond, J.-P.; and Overmars, M. H. 1998. Multi-level path planning for nonholonomic robots using semi-holonomic subsystems. *Int. Journal of Robotics Research*.

Sekhavat, S.; P. Švestka, J.-P. L.; and Overmars, M. H. 1996. Multi-level path planning for nonholonomic robots using semi-holonomic subsystems. *The int'l journal of robotics research* 17:840–857.