Optimizing Rotorcraft Approach Trajectories with Acoustic and Land Use Models

Robert Morris
NASA Ames Research Center
robert.a.morris@nasa.gov

Matthew Johnson
Florida Institute for Human and Machine Cognition
mjohnson@ihmc.us

Kristen Brent Venable
Tulane University/Florida Institute for Human and Machine Cognition
kvenabl@tulane.edu

Mateo Rodriguez
Tulane University
mrodrig@tulane.edu

Abstract

Recent increase in interest in using rotorcraft (helicopters and tilt-rotor craft) for public transportation has spurred research in making rotorcraft less noisy, particularly as they land. The ground noise associated with landing trajectories followed by rotorcraft depends in part on the changes in altitude and velocity of the rotorcraft during flight. Acoustic models of ground noise taking altitude and velocity effects into account can be used in an optimization process to determine a set of potentially quieter pilot operations. However, optimizing solely for acoustic properties produces patterns that abstract away from the environment in which the trajectory is flown. A quiet procedure flown over a residential area can create considerable annoyance. To overcome this limitation of acoustic-based optimization we propose a hybrid cost model for optimization that combines acoustic criteria with a land use model that views noise-sensitive areas around landing facilities as weighted obstacles. The result is a 3D route planning problem with obstacles. We introduce a system, called NORA (Noise Optimization for Rotorcraft Approach) that allows for the computation of trajectories that simultaneously solve for acoustically quiet patterns that also avoid land sensitive areas.

Introduction

Recently, there has been greater interest within the aeronautics community in using rotorcraft (including helicopters and tilt-rotorcraft) for commercial air transportation. A significant obstacle to implementing vertical lift transportation that in large part has deterred commercial rotorcraft use is the significant amount of ground noise produced, especially during takeoff and landing. While noise is very difficult to accurately model, it does have some non-linear dependence on the changes in altitude and velocity of the rotorcraft during flight that can be mathematically modeled. On the most basic level, ground noise increases as a result of a decrease in altitude or an increase in velocity.

The relationships between ground noise, altitude, and velocity can be used to develop a model that predicts the amount of ground noise produced by different rotorcraft. While useful as a tool for determining generic recipes for quiet operations for pilots, it is limited in not considering the landing environment in which the operations are performed. Performing an acoustically quiet trajectory might still create considerable noise because unbeknownst to the pilot the route flies over a land-sensitive area (e.g. retirement community or hospital).

Our goal in this research is to create a planning tool that can be used operationally as part of an air traffic control (ATC) tool, and, eventually, as an application that can be used in the cockpit, to provide real-time advice to a pilot. The contribution of this paper is to introduce a model of noise that combines an acoustic model with a model of land use. The hybrid model is integrated into a system called NORA (Noise Optimization for Rotor Approach).

The remainder of this paper is organized as follows. We introduce the architecture of the NORA system, following by a section describing the hybrid noise models used to evaluate trajectories, and one on the optimization approaches used. We finish with a discussion of the analysis and visualization tools in NORA, and a discussion of future work.

NORA Architecture

The overall NORA system architecture combining optimization and trajectory evaluation is displayed in Figure 1. The optimizer algorithm searches through the space of flyable approach trajectories and generates candidate trajectories for evaluation by a cost function derived from a combined acoustics model and land use model (both described in more detail below). The result is a noise profile (‘heat map’), which is input into a cost function that aggregates the noise data into a single value that summarizes the overall annoyance. This value then guides the optimizer for further search.

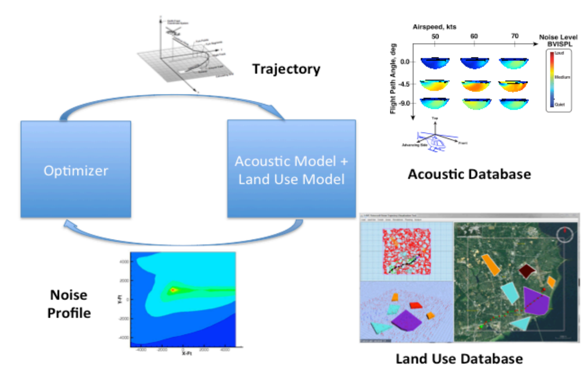

Figure 1. The NORA Architecture. An optimizer generates a valid trajectory that is evaluated using a hybrid acoustic/land use model. A noise profile is created, which is fed into a cost function that creates a single score rating the relative noise of the trajectory.

Noise Models

Rotorcraft Noise Model

Helicopter noise sources include the main rotor, the tail rotor, the engine(s), and the drive systems. The most noticeable acoustical property of helicopters is referred to as BVI (Blade Vortex Interaction) noise. This impulsive noise occurs during high-speed forward flight as a result of blade thickness and compressible flow on the advancing blade. A common noise measure is the Sound Exposure Level (SEL). SEL provides a comprehensive way to describe noise events for use in modeling and comparing noise environments. The average SEL value over the plane is called the SEL average.

One challenge in performing a systematic study of approach trajectories for optimization is the cost of verifying results. The most accurate means of verification is through field tests, but these are too costly and time-consuming to perform on a casual basis. Fortunately, there are a number of robust noise models that allow for the evaluation of trajectories through simulation. One such modeling tool is the Rotorcraft Noise Model (RNM) (RNM 2007), a simulation program that predicts how the sound of a rotorcraft propagates through the atmosphere and accumulates on the ground. The core of the RNM method is a database of vehicle source noises defined as sound semi-spheres. Semi-spheres are obtained through measured test data or through models. The semi-spheres allow for a representation of the 3D noise directivity patterns associated with the operating rotorcraft. A sphere is associated with one noise source and one flight condition (a value for flight path angle and airspeed). Each sphere represents constant airspeed conditions for a given flight path angle. The sound source properties are extracted from the sphere database using a linear interpolation of both required speed and flight path angle.

The input to RNM consists of a set of computational parameters, including identity of rotorcraft, and the dimensions and resolution of a grid that will display output noise and a specification of the flight trajectory, including position, velocity and orientation. RNM outputs a time history and the effective SEL (or similar metrics) at any location along the path. The result is usually displayed as a contour plot (Figure 2) over a ground plane. Each color corresponds to a dB level (redder and lighter colors noisier).

One important parameter that controls performance of RNM is *grid resolution*. The grid resolution specifies the distance between grid points, and thus the size of the area of the ground surface (in the SEL computation) being evaluated. Higher resolution means more accurate prediction of ground noise, but at the cost of slower run times for simulation. Experience shows a dramatic degradation of performance with increased resolution (RNM 2007).

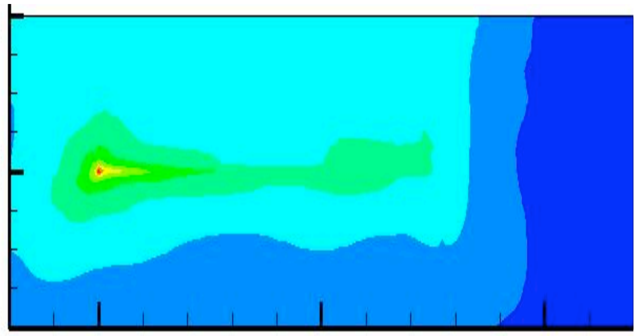

Figure 2. Noise profile (heat map) generated by RNM simulator. Lighter and redder colors represent heavier ground noise (SEL) values.

Land Use Model

The land use model is implemented within the NASA World Wind virtual globe, using the World Wind Java Software Development Kit (SDK), with plug-ins to identify latitude-longitude for a region of interest (for our application, we're clearly interested in urban areas around airports where approach planning would occur) as well as urban features (locations of potentially land-sensitive objects). Urban features can be categorized into land-sensitive areas (which include hospitals, schools, and residential areas) and noise-tolerant areas (like vacant lots or agricultural areas). Land sensitive areas can be viewed as either *obstacles* or *weighted areas:* a feasible path is precluded to fly through an obstacle, but only incurs a penalty, called a *sensitivity weight* if it flies through weighted area. Users of the model can assign weights to weighted area to estimate relative annoyance levels. An optimal plan might include flying through a weighted area based on other factors, such as the overall annoyance level that would incur by any alternative path that avoided it.

The final component of the hybrid land use/acoustics model is the incorporation of the RNM acoustic contour map into the land use profile. Given an approach path, the land use model determines penalties for violating land use constraints. The segments in which the path intersects a sensitive land use are assigned penalties corresponding to the sensitivity weights. RNM is then run to generate its acoustic profile, and the segments infringing on a land use area are assigned a 'boosted' value corresponding to increased annoyance.

Optimizer Cost Function

In (Morris et al. 2012) a method for aggregating the information in a contour map into a scalar value to be used for noise minimization algorithms was proposed. The method is based on a Binning function, *Bin*. Given trajectory t, and a SEL value for each grid point (x, y), denoted SEL(t, x, y) , a sequence of decreasing ranges, $<r^1, \ldots, r^n>$ partitioning the SEL values of the grid points is defined. Moreover, given $S_i(t) = \{(x, y) | SEL(t, x, y) \in r^i\}$, the set of grid values within the range r^i, the vector $b(t) = <b_1(t), b_2(t), \ldots, b_n(t)\}$, where $b_i(t) = |S_i(t)|$, represents the number of grid points with a corresponding SEL value within each range. The bin-score of solution t is then defined as $Bin(t) = \sum_{i=1\ldots n} w_i b_i(t)$ where w_i is the weight associated to the i-th bin, $w_i > w_{i+1}$ and $\sum w_i = 1$. Thus a solution that assigns lower levels of noise to larger regions of the grid has a lower *BIN* score. Weights can be tuned in various ways to penalize the presence of noisy regions in the grid.

3D Trajectory Noise Optimization

The problem to be solved is a multi-dimensional trajectory optimization problem. Trajectory planning has been defined as the simultaneous planning of path and velocity for robotic systems (LaValle, 2006). Solutions are described by means of a configuration space C, divided into free space and obstacle space. A planning instance is defined by the pair (c_s, c_f) of an initial and a final configuration. A feasible solution for an instance is a sequence of configurations in C, $(c_1, \ldots c_n)$ such that each configuration is in free space, and satisfies differential constraints on velocity.

A solution requires a specification of a configuration space, a set of dynamic constraints, an approach to discretization of the search space, and an approach to search.

Configuration Space and Constraint Model

The state variables for a rotorcraft consist of position (x, y, z (altitude) in 3D space), velocity (in feet per second), heading and angle of bank (both in radians). Thus the configuration space is in 6D, and the optimization problem is to solve a query: given two configurations corresponding to start and goal, compute a quiet, feasible path. For our domain, the goal is fixed to a landing configuration, but the start configuration can vary. In addition, there are three control variables, flight path angle, turn rate and velocity, which determine transitions between states (Figure 3, top).

State	Control
X	
Y	
Z	Flight path angle
Heading	
Angle of bank	Turn Rate
Velocity	Velocity

State and Control Variables

Neighborhood Function

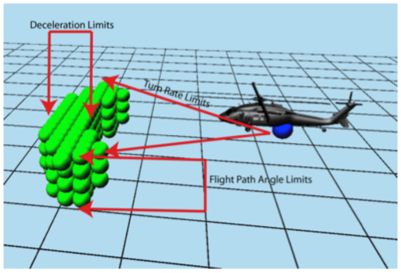

Figure 3. State Model and Neighborhood Function for Trajectory Planning. There are 6 variables in the state model and 3 in the control model. Dynamic constraints define dependencies between changes in control and changes in state.

Path feasibility is enforced by a set of dynamic constraints. Constraints define allowable initial conditions (e.g. velocity as approach commences), set boundaries on state (e.g. limits on velocity and bank angle) and implement differential constraints (e.g. determining the effects of change in turn rate on the current state). Control actions over a continuous space are discretized through the imposition of a grid. To avoid confusion with the notion of grid associated with the noise data, we refer to the grid imposed on the configuration space as the *search grid* with an associated search resolution. A search grid can be defined as a range (upper and lower bounds) and a resolution, or number of grid points. Figure 3 (bottom) shows how a search grid is applied to generate a finite set of discrete successor states from a given state. Intuitively, the search grid defines how many control choices a pilot (or autopilot) has at an arbitrary point: the finer the resolution, the more choices.

A* Optimization Using PRM

Sampling-based path planning techniques emerged out of the need to address the challenge of high-dimensional search typical in realistic path planning problems. The main idea behind sampling-based approaches is to generate and organize a sequence of samples from free space into a graph, where the edges are labeled with distance, usually a metric on C.

Probabilistic Road Maps (PRMs) (Kavraki et al. 1996) is a simple sampling-based algorithm for path planning that consists of a road map construction (learning) phase and a query phase in which the roadmap is searched for planning purposes. During the learning phase, configurations are randomly sampled and edges added between configurations in the roadmap that are close to one another, and collision-free (in our application, edges are allowed to pass through weighted obstacles). To solve a planning problem, the initial and final configurations are added to the roadmap. The effectiveness of the planner depends on how well the roadmap captures the connectivity of the configuration space. In general, the larger the roadmap and the higher the connectivity, the better the chances of finding a good, feasible path.

Typically, planning on roadmaps is done to find obstacle-free feasible paths. Because we're interested in optimization, we chose A* (Choset et al. 2005) for the planning phase. A* is a best-first search algorithm for discrete planning that is based on incrementally expanding a partial solution s through the use of a function f(s) which estimates the minimal cost path from the start through s. f(s) is the sum of the cost g(s) to reach the state plus a heuristic h(s) that estimates the lowest cost path to achieve the goal state that passes through s. A* is complete and optimal provided the h never overestimates the cost of achieving the goal through s. In previous experiments it was shown that a reasonable relaxation of the problem for the purpose of defining a heuristic value for a given node is the cost of a trajectory in which altitude is maintained until the goal, but airspeed is reduced as fast as possible to the minimum velocity. This *fly high and slow* heuristic has been empirically confirmed to be admissible. Notice however, that A* requires that f be factorable into the sum of g and h; in our case, g is the noise generated so far by the partial path, and h is an prediction of the noise from the current state to the goal. Unfortunately, the contour map that represents the input to the BIN cost function is the result of averaging noise levels (in decibels) over space and time over the path flown. Therefore, it is not possible to decompose the aggregate noise into a sum of noises of the component paths. Nonetheless, previous tests have shown that approximating the cost function by assuming decomposability works fairly well for our purposes.

Figure 4 shows a run of A* to solve the problem of finding an optimal (minimal noise) approach and landing path at Pensacola airport, using the land use model for this city, represented by multi-colored polygons (the colors signifying different categories of land use). The roadmap is shown in the upper left quadrant, and the path is displayed on the World Wind map on the right.

A* Optimization on a fixed grid

To explore and compare different approaches to discretizing a continuous configuration space, we allow A* to be applied over a fixed grid, defined by a specifying a distance metric between configurations. Control actions over a continuous space are discretized through the imposition of a grid on the configuration space. To avoid confusion with the notion of grid size associated with the simulator, we refer to the grid imposed on the configuration space as the *search grid* with an associated search resolution. A search grid can be defined as a range (upper and lower bounds) and a resolution, or number of grid points.

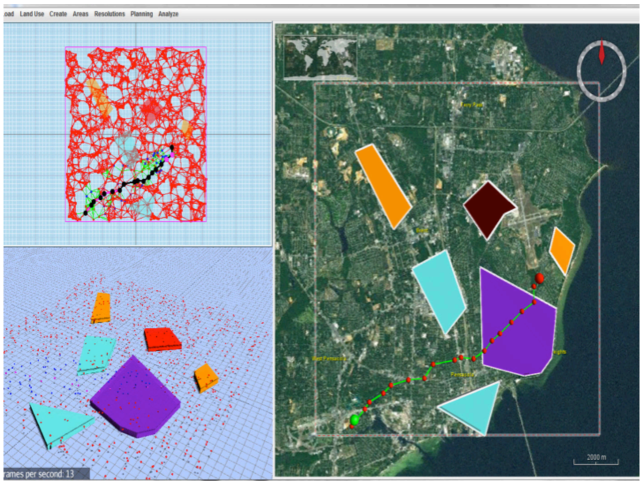

Figure 4. Path planning problem defined using land use model of Pensacola Florida. The optimal 3D approach path to Pensacola airport found by A* is displayed on the right figure. The roadmap appears on the upper left figure, and the weighted land use obstacles are shown schematically in the bottom left figure.

Adjacent configurations in the search grid are defined by identifying the range of control decisions that are available in a particular state (e.g. the upper and lower bounds of changes to velocity) then defining the number of discrete values that will be generated between the boundary values. For example, a velocity resolution of 10 means that the optimizer will consider, for any state, 10 values of (change in) velocity within the range defined by the upper and lower bounds, while computing a path. The bottom image in Figure 3 shows how adjacent configurations on the grid are generated through defining the boundaries and resolution of the three control variables. The collection of green balls represents feasible extensions of the path from the current state (blue ball). Intuitively, the higher the search resolution, the larger the set of green balls that is produced during path expansion. There is an obvious potential tradeoff in specifying search resolution between quality of solution and search time. Figure 5 shows a visualization of A* performed on a fixed grid.

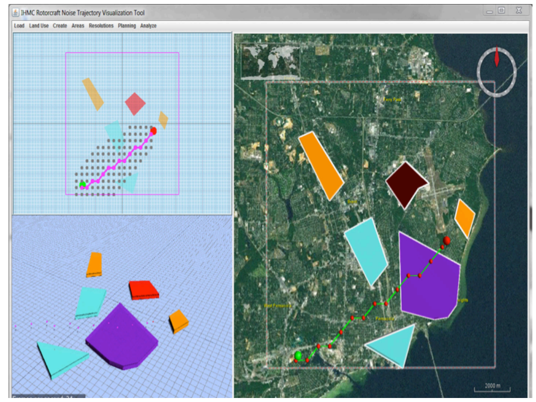

Figure 5. A* search on a fixed grid illustrated in upper left image. The configurations expanded are shown in grey, along with optimal path.

A final variation in search options emerges out of the ability in the RNM simulator to designate "points of interests" and limit the computation of noise estimates to those points. This is useful in cases where the interest in noise is limited to specific known regions (for example, in military applications in remote mostly uninhabited regions, the interest might be limited to known hostile areas). Points-of-interest planning also provides computational savings over running the simulator over an entire grid. Figure 6 illustrates planning around points-of-interest. The balls on the upper half image represent points on the grid where noise predictions are made by RNM. The balls are labeled with their SEL values.

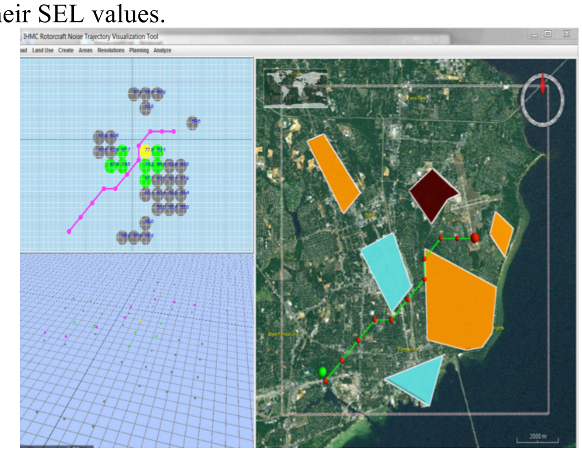

Figure 6. Planning using points of interest. The land sensitive regions, represented by the polygons on the right figure, are mapped into points of interest in RNM. Noise predictions for a given trajectory are limited to the regions shown as balls.

Analysis and Visualization Tool

Figures 4-6 show screenshots of the NORA planning interface. To specify a planning problem, the user interface allows the following to be specified:

- Loading a land use database (e.g. of Pensacola)

- Filter the land-sensitive categories to those of interest (e.g. limited to only residential areas and hospitals)
- Specify a bounding box in 3D space within which the planning will be conducted
- Designate whether the land sensitive areas are weighted or strictly obstacles
- Specify grid and search resolution (if search is to be done on grid)
- Specify cost function (including BIN, but also other cost functions not discussed in this paper, such as minimum distance)
- Specify search function, currently limited to A* using PRM or fixed grid, but other search functions (such as Rapidly Expanding Random Trees) are currently under development
- Discretization method (currently, PRM or fixed grid, but other discretization methods are possible)
- Start and end configurations of the planning problem.

The output of a run of the NORA planning tool is shown in Figures 4-6. In addition, there are tools for analyzing and comparing different methods. The analysis includes runtime and quality-of-solution statistics, for example comparing A*+RNM to A*+fixed grid using these metrics. The user can also compare different path profiles to extract useful pilot guidance. For example, for a particular rotorcraft, the path profiles might indicate that for certain routes it is quieter to descend quickly early in the approach at a steady velocity, and then stay level and reduce speed before landing. We are currently conducting experiments of these sorts to be summarized in future reports.

Summary

This paper has described a framework for 3D path planning for designing quiet approach trajectories. The framework, called NORA, utilizes an acoustics model for rotorcraft noise based on a simulator called RNM. It also incorporates land use data from NASA World Wind in order to identify land sensitive areas, which then can be treated as regions to be avoided during planning. NORA planning is based on a 6D state space and 3D control space, and utilizes a constraint model for defining feasible trajectories. NORA supports path-planning using A* searching over a graph generated probabilistically using PRM, or on a fixed grid. An array of cost functions can be used; for the acoustic profile generated from RNM, an aggregation of noise data as a weighted sum of noise values over different ranges of annoyance is the default. Users of the NORA interface can specify a planning problem by setting a number of parameters related to regions of interest, land use categories, and start and end configurations. NORA visualizes path information and other statistics of interest to trajectory designers. Although not discussed in this paper, the planner's performance can be tuned in such a way that a solution can be generated quickly, thus opening up the possibility that NORA can be used by air traffic controllers to upload quiet plans to pilots in order to provide real-time guidance. This option will be explored in future research.

References

RNM, 2007. Rotorcraft Noise Model Technical Reference and User Manual (Version 7). NASA Langley Research Center.

Morris, R.; Venable, K. B.; Pegoraro, M.; and Lindsay, J. 2012. Local search for designing noise-minimal rotorcraft approach trajectories. Proceedings of Innovative Applications of Artificial Intelligence (IAAI).

Kavraki, L.; P., S.; J.-C., L.; and M.H., O. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. IEEE Transactions on Robotics and Automation Volume: 12 .

LaValle, S. M. 2006. Planning Algorithms. Cambridge University Press.