Demonstration of the Emergency Landing Planner

Elif Kurklu* **Nicolas Meuleau*** **Christian Neukom***
Christian Plaunt **David E. Smith** **Tristan Smith†**

Intelligent Systems Division
NASA Ames Research Center
Moffet Field, California 94035-1000
{elif.kurklu, nicolas.f.meuleau, christian.neukom, christian.j.plaunt, david.smith, tristan.b.smith}@nasa.gov

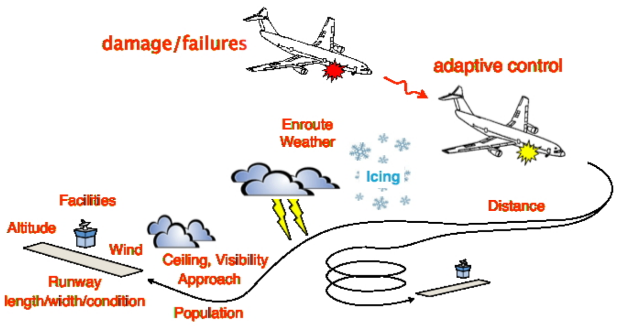

Figure 1: Basic Scenario

Abstract

We describe an Emergency Landing Planner (ELP) designed to assist pilots in choosing the best emergency landing site when damage occurs to an aircraft. In 2010, we integrated the ELP into the cockpit of a 6 DOF full-motion simulator for transport category aircraft, and performed experiments to evaluate the software using crews of professional airline pilots. We briefly describe how the ELP works, and show how it was integrated into the avionics of the simulator.

1. The Scenario

Figure 1 illustrates the type of scenario that the Emergency Landing Planner (ELP) addresses. When damage or failures occur in an aircraft an adaptive controller takes over to help stabilize and control the aircraft. The pilots then invoke the ELP using the Flight Management Computer. The ELP provides the pilots with a ranked set of possible emergency landing sites.

Fundamentally, the ELP solves a 3D path planning problem with dynamics. It does this by constructing a probabilistic roadmap of points and edges that includes the current position of the aircraft and an approach point to every possible runway within a viable range. (This may cover hundreds of airports for an aircraft at high altitude.) A sophisticated model of risk is used to assess the probability of success for each edge in the roadmap. This model of risk takes into account:

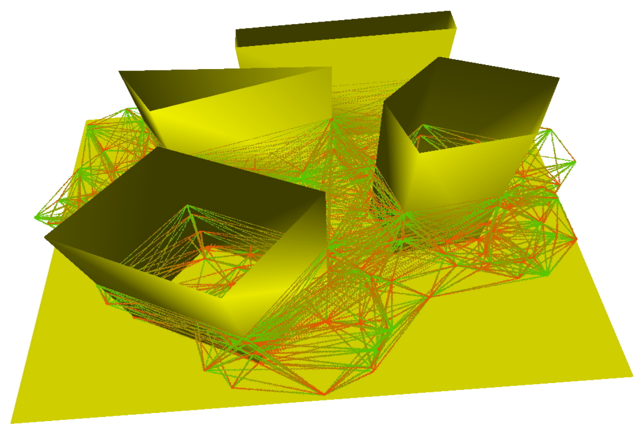

Figure 2: An example roadmap for an ELP scenario. The vertical polygons are areas of thunderstorm or other weather activity. Terrain obstacles (lower) are not shown.

- Control capabilities of the (damaged) aircraft
- Weather conditions in the area (e.g. thunderstorms, turbulence, icing)
- Ceiling, visibility and winds at each possible landing site
- Instrument approaches available at the site (if any)
- Characteristics of the landing site (runway length, width, condition)
- Emergency facilities at the site (fire, medical)
- Danger to population along the approach path

The flight envelope plays a key role in the assessment of risk for the different options. For example, if a damaged aircraft must maintain a higher airspeed than normal, additional runway length is needed, and finding a runway with a strong headwind is important in order to lower ground speed at touchdown. Similarly, if the aircraft has limited ability to bank to the right, a right crosswind or gusty conditions will be problematic, as will paths that require sharp turns to the right.

*Stinger Ghaffarian Technologies
†Mission Critical Technologies

Figure 3: The Advanced Concepts Flight Simulator (ACFS).

A^* is used to search the roadmap to find the best landing options. The g value for a given path is the product of the probabilities of success for the legs in the path. The heuristic value h is estimated using the probability of success for a direct path to the approach point for each runway (assuming no weather) times the probability of success for the approach and landing on that runway. This heuristic value is admissible, and fairly informative.

The performance of the ELP is largely a function of the number of points and edges in the roadmap. Currently, we generate 2000 points and connect them to their 200 nearest neighbors, which results in a roadmap with 400,000 edges. The A^* search typically expands about 20 percent of those edges for the scenarios we considered. With this sized roadmap, the ELP produces an ordered list of options for the pilot in under 10 seconds. This list can therefore be refreshed and updated as often as desired, to account for the aircraft movement, weather updates, or additional failures.

Our experience has been that paths generated from probabilistic roadmaps of this density can be far from optimal, and just don't look very good when displayed. This problem can be addressed by dramatically increasing the density of points and edges, but this approach also significantly increases search time. The more practical solution is to use local search to shorten and smooth paths. We do this local search by constructing a second roadmap consisting only of points along the path just found, creating a dense network of edges among those points, and re-running A^* on this reduced graph. The resulting paths are shorter, smoother, and seem more natural when displayed.

More detail about the ELP, the path planner, the risk model, and the local search can be found in (Meuleau et al. 2009; 2011a; 2011b)

Figure 4: The cockpit of the ACFS.

2. Integration

Figures 3 and 4 show the Advanced Concepts Flight Simulator (ACFS) at NASA Ames Research Center. The simulator is representative of modern glass cockpit twin engine commercial transport aircraft such as the Boeing 757, 767, and Airbus A320. The ELP was integrated into the avionics of this simulator in order to conduct experiments with teams of professional pilots with different damage scenarios and weather conditions. Unfortunately, the ACFS is not very portable, so for demonstration purposes, we use a reasonably high fidelity simulator that runs on a laptop. It includes the Primary Flight Display, Navigation Display, and Flight Management System common to modern glass cockpit aircraft. Just as in the ACFS, this simulator incorporates an adaptive controller, and has several damage models available.

A typical scenario involves starting the aircraft in cruise flight following a flight plan like that shown on the Navigation Display in Figure 5. A failure is then introduced as illustrated in the surface position display shown in Figure 6. In the case illustrated in Figure 6, the left wing is damaged and the left aileron has failed (red).

When the damage occurs, the adaptive controller takes over and stabilizes the aircraft. In the example illustrated, the adaptive controller is adding right up aileron (blue) and right spoilers (blue) to keep the aircraft from rolling left. To help the pilots understand the control limitations of the damaged aircraft, color bands are shown on the primary flight display as illustrated in Figure 7. These color bands indicate safe ranges for airspeed, bank angle and vertical speed. In this case, the aircraft must maintain a much higher speed than normal to keep sufficient airflow over the remaining aileron. The ability to bank right is also very limited.

Pilots access the ELP from the Departure/Arrival page of the Flight Management Computer as shown in Figure 8. After a brief splash screen, a set of "Emergency Pages" is displayed, showing the options ordered from lowest to highest risk. Figure 9 shows the first of four emergency pages for this scenario. Each entry shows an airport, runway, run-

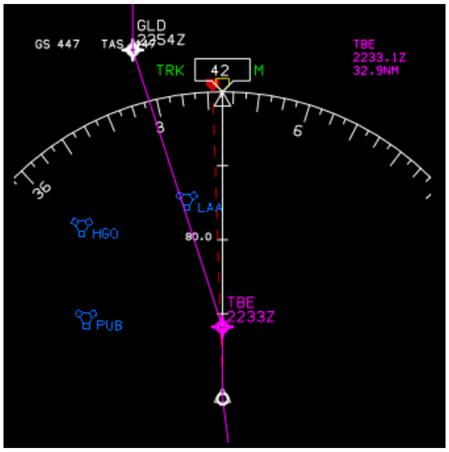

Figure 5: The Navigation Display showing the current route.

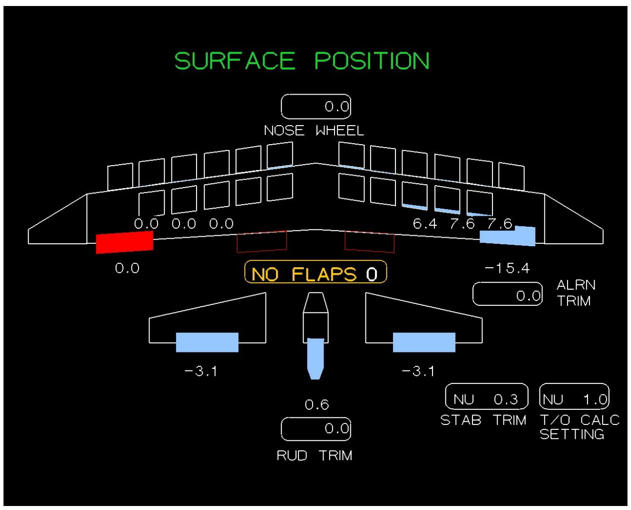

Figure 6: Surface position display showing status and deflection of control surfaces.

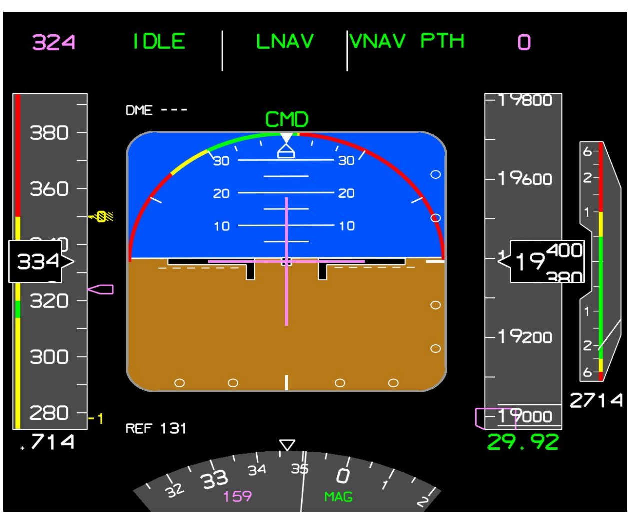

Figure 7: The Primary Flight Display (PFD) showing bank angle, pitch, airspeed, vertical speed, altitude and heading.

Figure 8: The display for the Flight Management Computer showing the Departures/Arrivals page for Denver (KDEN). The emergency prompt appears next to button 6R at the lower right.

Figure 9: The first of four emergency pages for a scenario.

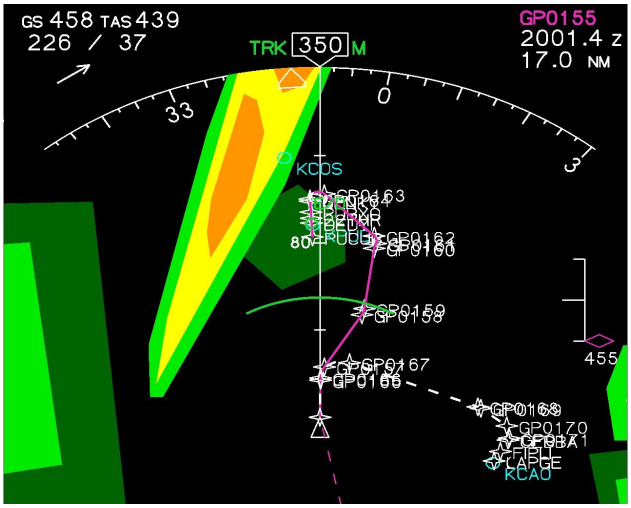

Figure 10: The Navigation Display showing both the current route (magenta) and the new route being considered (dashed white). Green, yellow, and orange areas indicate rain and thunderstorm activity.

Figure 11: An Airport Information page showing runways and current weather for KCAO.

way length, distance, and direction (magnetic bearing). The smaller symbols below each entry indicate the principle risks associated with that option; for example, RL indicates runway length is an issue, and CE indicates that the cloud ceiling is close to the minimums for the best approach to that runway. To select an entry, the button to the left of the entry is pressed. In this case, the first entry has been selected by pressing button 1L, which causes the route for that option to show up as a dashed white line on the Navigation Display, as shown in Figure 10. Pressing the EXEC key would cause the route to become the current route (solid magenta). The pilots can page through the options using the NEXT PAGE and PREV PAGE buttons as desired. To see more information about a particular option, the pilots can press the button to the right of the option, which brings up an airport information page showing runway information and the current weather at the airport (Figure 11).

Acknowledgments

This work was supported by the NASA Aviation Safety Program and the American Reinvestment and Recovery Act. We thank John Kaneshige, Stefan Campbell, Shivanjli Sharma, Captain Mietek Steglinski, Matt Gregory and the staff of the NASA Ames Crew Vehicle Systems Research Facility for their work in helping to prepare for and conduct the simulator experiment.

References

Meuleau, N.; Plaunt, C.; Smith, D.; and Smith, T. 2009. An emergency landing planner for damaged aircraft. In *Proceedings of the Twenty First Innovative Applications of Artificial Intelligence Conference*. AAAI Press.

Meuleau, N.; Neukom, C.; Plaunt, C.; Smith, D.; and Smith, T. 2011a. The Emergency Landing Planner experiment. In *ICAPS-11 Scheduling and Planning Applications Workshop (SPARK)*.

Meuleau, N.; Neukom, C.; Plaunt, C.; Smith, D.; and Smith, T. 2011b. The Emergency Landing Planner experiment. Technical report, NASA Ames Research Center.