The Emergency Landing Planner Experiment

Nicolas Meuleau* and **Christian Neukom*** and **Christian Plaunt** and **David E. Smith** and **Tristan Smith***
Intelligent Systems Division
NASA Ames Research Center
Moffet Field, California 94035-1000
{nicolas.f.meuleau, christian.neukom, christian.j.plaunt, david.smith, tristan.b.smith}@nasa.gov

Abstract

In previous work, we described an Emergency Landing Planner (ELP) designed to assist pilots in choosing the best emergency landing site when damage or failures occur in an aircraft. In this paper, we briefly describe the system, but focus on the integration of this system into the cockpit of a 6 DOF full-motion simulator and a study designed to evaluate the ELP. We discuss the results of this study, the lessons learned, and some of the issues involved in advancing this work further.

In a previous paper (Meuleau et al. 2009b), we described 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 our planning software into the cockpit of a 6 DOF full-motion simulator for 757/767 category transport aircraft, and performed experiments to evaluate the software using crews of professional airline pilots. In this paper we briefly review the Emergency Landing Planner (ELP), but focus on three topics:

- Integration of the software into the aircraft avionics
- Design and results of an experiment to evaluate the system
- Challenges to further advancing and fielding the technology

1. The Emergency Landing Planner

Figure 1 illustrates the type of scenario that the ELP addresses; when damage or failures occur in an aircraft an adaptive controller takes over to help stabilize and control the aircraft. The ELP then provides the pilot with a ranked set of possible emergency landing sites. Fundamentally, the ELP is solving 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:

- Control capabilities of the (damaged) aircraft

*Mission Critical Technologies

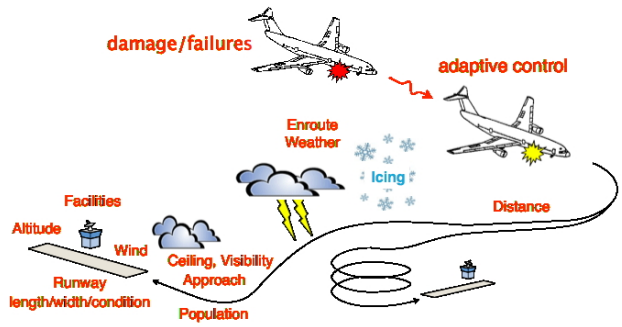

Figure 1: Basic Scenario

- 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

A^* search is used to search the roadmap to find the best options. The heuristic used to guide A^* is a combination of the risk associated with flying the remaining (Euclidean) distance to each runway, and the risk associated with approach and landing at that runway. For each path that is constructed by A^*, local search is used to optimize and smooth the path.

Currently, the ELP only considers officially recognized airports and runways (large and small). However, there is no fundamental reason that additional sites could not be considered, including fields, highways, and waterways. Such sites should probably not be considered unless the airport options are exhausted or are too risky. The ELP makes two additional assumptions:

1. Real time weather information is available to the aircraft
2. The flight envelope for the (damaged) aircraft is known[1].

[1]For our purposes, the flight envelope of an aircraft is the four dimensional space of airspeeds, bank angles, vertical speeds, and altitudes in which the aircraft can operate.

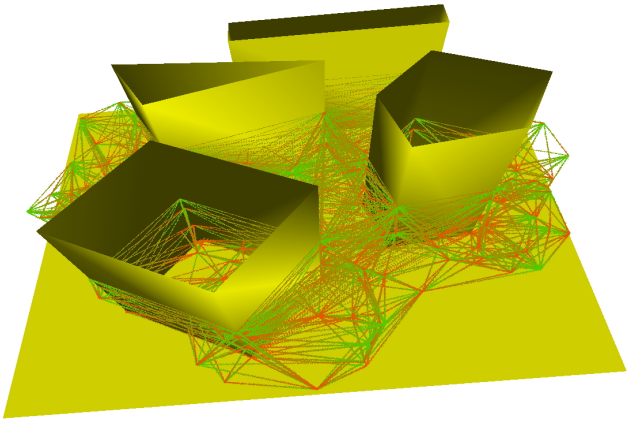

Figure 2: An example roadmap for an ELP scenario. The vertical polygons are areas of thunderstorm or other weather activity. Terrain obstacles (lower) are not shown.

The first of these assumptions is quite reasonable, given the availability of satellite weather services and internet connectivity for large aircraft. The second assumption is more optimistic. We will discuss this more in Section 5. The flight envelope does play a key role in the assessment of risk for various 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 to lowering 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.

The performance of the ELP is largely a function of the number of points and edges in the roadmap. Currently, we generate 1000 points and connect them to their 100 nearest neighbors, which results in a roadmap with 100,000 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 6 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 dramatically 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 details about the risk model the path planning, and the local search can be found in (Meuleau et al. 2009b; 2009a; 2011)

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

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. Unlike most large commercial flight simulators, the code of this simulator has been "exposed" to allow for experimentation with adaptive control software, damage models, and experimental pilot aids and displays.

In normal operations, pilots view, enter, and modify destination, route and approach information using a pair of keypads and displays (CDUs) located just above, and on either side of the throttles (Figure 5). Information entered on a CDU is communicated to the aircraft's Flight Management System (FMS), which interfaces with the autopilot and with the various displays in the cockpit. When route information

is entered on a CDU, the route shows up as a dashed white line on the pilot and co-pilots Navigation Displays (Figure 6). Once executed, the previous route disappears, and the route becomes solid magenta.

Figure 5: A CDU showing the Departures/Arrivals page for Denver (KDEN) airport. The emergency prompt appears next to button 6R at the lower right.

To integrate the ELP into the aircraft cockpit, we needed to make it accessible through the CDUs and make it communicate route information to the FMS, so that the emergency routes would appear on the Navigation Displays. Furthermore, we wanted the pilots to be able to edit or change an emergency route just as they can with any other route. As a result, the ELP had to be fully integrated with the CDUs and the FMS. In addition, we wanted to make the style of the interface reasonably intuitive and consistent with existing CDU pages.

The ELP is accessed using button 6R from the Departure/Arrivals page (Figure 5). After a brief splash screen, a set of "Emergency Pages" is displayed, showing the options ordered from lowest to highest risk. Figure 7 shows the first of four emergency pages for a scenario. Each entry shows an airport, runway, runway 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 Displays, as shown in Figure 6. 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

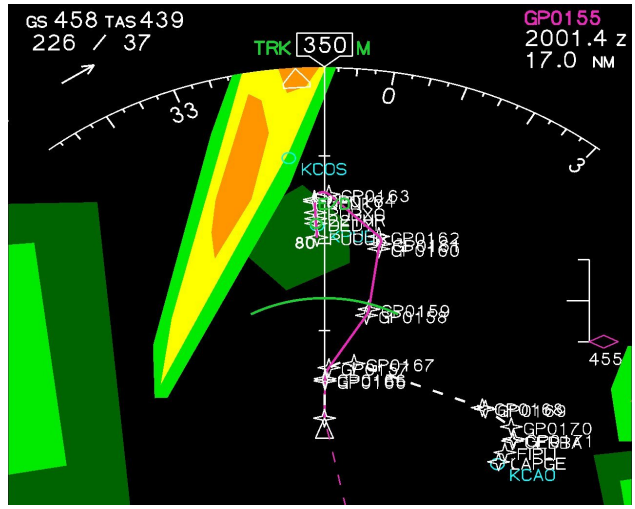

Figure 6: The Navigation Display showing both the current route (magenta) and the new route being considered (dashed white). Green, yellow, and orange areas indicate rain and thunderstorm activity.

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 8).

The screen size and lack of color on the CDUs limited the amount of information we could convey for each option. With greater screen real estate, we could display winds, ceiling, and visibility information for each option. With color, we could show the severity of the principal risks. It seems likely that the displays and interfaces of future aircraft will not be quite so limited.

In a previous study comparing different adaptive controllers, Campbell et al 2010b; 2010a found that because of the assistance of the adaptive controller, pilots were unaware of when they were approaching the boundaries of the flight envelope. For example, pilots would slow the aircraft too much on final approach, not recognizing that in doing so they were nearing saturation of one or more control surfaces. On reaching saturation, the nose of the aircraft would suddenly drop, or the aircraft would roll inverted causing them to lose control and crash. As a result, we felt that it was essential to give the pilots some additional guidance on the limitations of the flight envelope. To do this, we added color bands to the primary flight display to indicate safe airspeeds, bank angles, and vertical speeds as shown in Figure 9. If the airspeed, bank angle and vertical speed remain in the green regions the aircraft can be readily controlled. However, as airspeed decays down into the yellow (not yet visible in the figure), the green regions for bank and climb rate shrink, ultimately to nothing. The size of the regions is dictated by the 4-dimensional model of the flight envelope, which varies depending on the damage or failure. The green regions for bank can be asymmetric, as is the case when there is damage to a wing or aileron.

Figure 7: The first of four emergency pages for a scenario. Each page can show up to five options.

Figure 8: An Airport Information page showing runways and current weather for KCAO.

3. The Experiment

To evaluate the ELP, we developed a set of scenarios involving different locations, different flight plans, different weather conditions and different damage models. There were three different locations, two flight plans for each location, two different weather severities, and three different damage models, for a total of 36 possibilities. The damage models were previously developed at NASA Langley through a combination of vortex lattice code and wind tunnel testing. In addition to the scenarios, we needed a baseline with which to compare the ELP. We therefore developed a simple aid for the pilots that just listed the nearby airports grouped by runway length. We also developed an intermediate aid that evaluated runways using our risk model, but did not consider en route weather, an did not generate a path for the pilots. The matrix of testing possibilities is summarized in Table 10. The number of scenarios we could consider was necessarily limited by: the number of damage models available in the simulator; the number of realistic weather models we could develop; and the time required to test all the possible emergency flight plans and approaches that might be produced for each location.

To carry out the experiments, we employed 5 teams of professional pilots for two days each. All of these pilots were either current or recently retired airline pilots with experience in glass cockpit aircraft of the appropriate type. Each pilot team was briefed on the functioning of the ELP and baseline aids, and conducted several short training flights to ensure they were comfortable with the systems and handling of the simulator. The team was then subjected to 16 of the possible scenario/aid combinations. Each run began in cruise flight. Damage was introduced after 1-3 minutes, resulting in a master caution alarm in the cockpit, and indications of the failures on a display of the control surfaces shown in Figure 11. The pilots would then utilize the aid provided, chose an emergency landing site, and fly the aircraft until touchdown or loss of control. In some cases we also terminated the run after a decision was reached, because of time limitations. A typical run lasted about 35-40 minutes. At the end of the run, the pilots were asked to fill out a brief questionnaire about their decisions and about their assessment of the aid provided by the software. At the end of the two day period, the pilots were asked to fill out a longer questionnaire giving their overall impressions, criticisms, and suggestions for the emergency aids and interfaces.

During the runs, we observed the pilot performance from a control room with video screens of all major cockpit instruments (Figure 12). We collected multiple data streams including: video and audio from the cockpit; aircraft state at 10 Hz (location, altitude, airspeed, pitch, bank, control settings, etc.); keystrokes and display from the CDUs; and video of the Primary Flight Display (PFD) and Navigation Display (ND).

From the outset, we were aware that there were some serious limitations with the study:

1. The number of possible runs was limited because of both time and cost.
2. The number of different scenarios was limited because of the amount of data required and the difficulty of constructing the scenarios. As a result, the pilots could become familiar with scenarios and damage models as the study progressed.
3. The pilots could become fatigued, particularly later in the day.

The first of these limitations makes it difficult to draw sta-

Pilot Aid	Damage	Weather	Location
Nearest Airports	Vert. Stabilizer	Mild Overcast	Arizona LAS → STL, ABQ → SEA
Ranked Airports	Horiz. Stabilizer		Idaho GEG → DEN, GTF → SFO
ELP	Left Wing	Thunderstorms Low Ceilings	New Mexico COS → SAT, ABQ → MSP

Figure 10: Experiment test matrix.

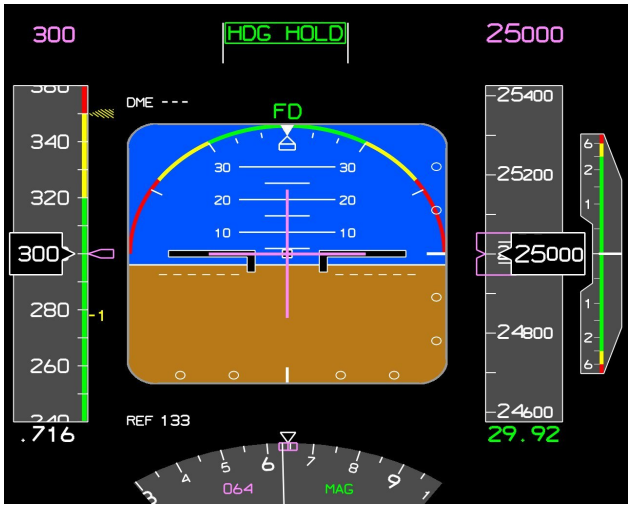

Figure 9: The Primary Flight Display (PFD) showing bank angle, pitch, airspeed, vertical speed, altitude and heading.

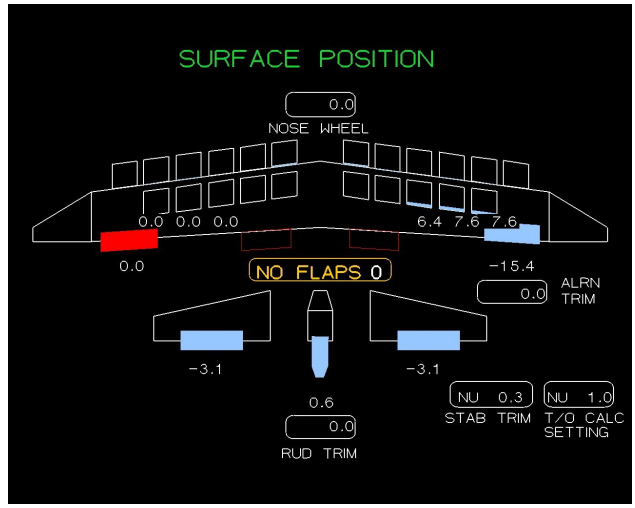

Figure 11: Surface position display showing status and deflection of control surfaces. In this case, the left wing is damaged and the left aileron has failed (red). As a result, the adaptive controller is using right up aileron (blue) and right spoilers (blue) to keep the aircraft from rolling left. When a control surface is saturated (at its limits) it turns yellow.

tistically significant conclusions. In any study dealing with human subjects, there is a great deal of variability and randomness, so large sample sizes are needed. The cost of the simulator and pilots makes this impractical.

The second limitation, the limited number of scenarios, meant that the pilots became increasingly "contaminated" as the study progressed. We tried to minimize this by mixing up the different damage models, weather conditions, locations, and flight plans. However, the pilots clearly became more familiar with the terrain and airports in each region, and their skill with the different damage models improved over time. To attempt to average out these effects, we ordered the scenarios differently for the different crews.

The third limitation, pilot fatigue, seemed to show up primarily during the afternoon of the second day of testing. We noticed it because there were some cases where the pilots lost control of the aircraft and crashed during easier scenarios.

The combination of these limitations means that many of our results are anecdotal, are based on small sample sizes, or are the results of subjective feedback from the pilots.

4. Results

Figure 13 is a trajectory plot showing the options considered by the pilots for one particular run. The red dot indicates the position of the aircraft at the time damage occurs. The black line is the aircraft's actual trajectory. Yellow lines indicate other options considered by the pilots, and the green line indicates the route provided by the ELP at the time they finally made a decision. As can be seen from the plot, the pilots made a tighter turn to the left (back towards the airport) than the ELP recommended. They also chose to intercept and get established on the final approach course further from the airport. In this run, damage was to the left wing and aileron, making it more difficult to turn right. In addition the weather was challenging, with larger airports in the area having low ceilings, poor visibility, or difficult crosswinds. In this case, KCAO runway 02 was the highest ranked option provided by the ELP, and it proved to be one of few choices for which pilots had any success in getting the aircraft on the ground. The blue path shows the route that would have been recommended if the pilots had made their decision instantly after the damage occurred. By the time the decision was made, it was no longer practical to make a right turn towards the chosen runway, given the control characteristics of the aircraft.

Figure 14 shows a run for a different scenario in the same general area. In this run, damage was to the horizontal stabilizer and elevator so turning was not difficult, but a higher

Figure 12: The ACFS control room.

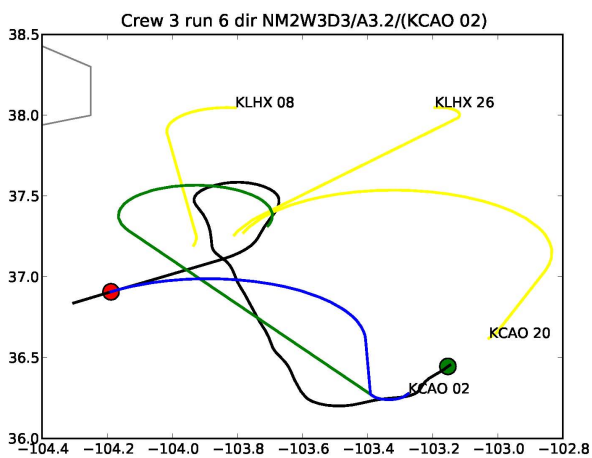

Figure 13: A trajectory plot for a left wing damage scenario.

airspeed had to be maintained to preserve enough airflow over the remaining elevator. In this case, pilots were tempted by long runways at lower ranked Colorado Springs (KCOS) and Cannon Air Force Base (KCVS), but winds and weather did not favor the available runways. They ended up choosing a more highly ranked option with a shorter runway (KCVN 04), because of the strong headwind straight down the runway. In this case, KCAO 02 would have also been a good choice for the same reason.

In analyzing the data, we considered the time required to make a decision, and the pilots success rate as a function of the damage model, weather conditions, and location for each different emergency aid. Our initial hypothesis was that that the ELP would prove helpful to the pilots in cases where either the damage or weather was severe, but that the pilots would do just fine with the baseline emergency aid when the weather and damage were both benign. This hypothesis is only partially correct; weather severity was a factor,

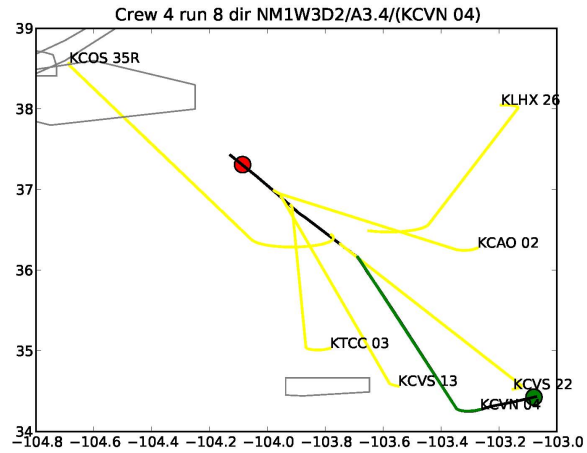

Figure 14: A second trajectory plot for a scenario with elevator damage.

damage severity was not. For the scenarios involving mild weather conditions, the ELP does not seem to offer any objective improvement in pilot performance over the two simpler emergency aids. However, when the weather was poor, the ELP generally led to quicker and better decisions. The reason for this is that when the weather was mild, one of the nearby large airports with a long runway was usually the best choice. Pilots could find this choice easily enough using only the simpler emergency aid, and could choose the most appropriate runway by looking at the airport information page for that airport. In contrast, when the weather was poor the pilots would be forced to look at many different options before finding one with acceptable weather conditions. In a few particularly difficult cases, pilots took more than 20 minutes to reach a decision using the baseline emergency aid. With the ELP, decisions for these same scenarios were made in 4-5 minutes.

From the objective data, the paths constructed by the ELP did not seem to offer a significant advantage to the pilots either in terms of decision making time, or in terms of decision quality. When intervening weather was not an issue, a direct route to a point about 10 miles out on the final approach course was appropriate and was relatively easy for them to construct. Even when the weather was more severe, the pilots were able to construct their own paths, although it took them longer to do so. For this reason, pilots subjectively reported that it was a lot easier to have the assistance of the ELP in all cases, both because of the ranking of options, and because the route was constructed automatically and guaranteed terrain clearance. Thus, even for mild weather conditions, pilots preferred the ELP, and felt that it reduced workload.

In contrast to our hypothesis, the severity of the damage does not seem to be correlated with whether or not the ELP provides an advantage. For mild weather, but severe damage, there is no objective evidence that the ELP provides an advantage over the baseline emergency aid. Likewise, for

severe weather, the advantage of the ELP over the baseline aid does not become any more pronounced if the damage is severe. We speculate that this is largely due to the adaptive controller – it does such a good job of stabilizing the aircraft that the pilots have time to investigate options and construct routes manually. Without the adaptive controller the workload is much higher; it is much more of a struggle to maintain control of the aircraft, and the pilots would likely not have time to consider multiple options, or construct routes manually.

Overall, we were both surprised and thrilled with the enthusiastic response we received from the pilots, as illustrated by this quote:

> *As a Captain for the past twenty some years I've trained for emergencies frequently and the most difficult part is selecting and getting the aircraft on the ground when a immediate landing is called for. Your software program alleviates the uncertainty about finding a suitable landing site and also reduces workload so the Crew can concentrate on "flying" the aircraft.*

Although the technology was designed for next generation aircraft, several pilots indicated that they wanted to see this capability in their existing aircraft and suggested that it would be particularly valuable in time critical situations like cargo fires, loss of engine power, rapid fuel or hydraulic fluid loss, or medical emergencies. Ironically, during our study a brand new UPS 747 crashed in Dubai as the result of a cargo fire. The pilots chose to return to the takeoff airport, although closer options were available. This proved to be a fatal mistake, as the smoke became so thick that the pilot could no longer see his instruments.

We have two final observations relevant to pilot training. The first is that there were significant differences in pilot performance. Two of the three damage scenarios required approach and landing at considerably higher speeds than the pilots were accustomed to. Those pilots with experience in high speed military aircraft did much better at this than those without that experience. Although rare, several actual damage incidents have had this characteristic, so regular simulator training in high speed landings would potentially be valuable.

The second observation is that many pilots preferred long runways with poor weather and wind conditions to shorter runways with better weather and winds. For example, several teams were seduced by the 13,000 ft runway at Colorado Springs, even though it was ranked low because the visibility was 1 mile in blowing snow, with a strong 70 degree crosswind. This proved fatal in almost every case. In contrast the top ranked option only had a 7000 ft runway, but had good visibility and a strong headwind straight down the runway. The crews that chose this option were generally successful. This lends some credence to our risk model, and suggests that pilots should be trained to favor better wind and weather conditions over longer runways in emergency situations.

5. Challenges and Regrets

As we expected, the experiment made us aware of many ways in which the ELP could be improved. Some of these are concerned with the robustness of the communication interface between the ELP and the cockpit CDUs and FMS, some are improvements to the user interface and information layout on the CDUs. The most important involves improvements to the risk model and to the path planner. For the risk model, we recognized that we need to increase the risk for crosswinds and gusts in cases where the aircraft has limited yaw control (rudder damage). We also recognized that ground speed at touchdown should be weighted more heavily due to the likelihood of tires blowing at high speeds. Finally, we did not consider the terrain roughness along and in the vicinity of the approach path – this was clearly a factor that the pilots considered when choosing options, particularly when controllability was limited.

For the path planner, we found that the pilots tended to prefer a gentler turn to intercept the final approach course and a longer final approach course, as illustrated by Figure 13. This was particularly true when controllability was poor. While these are all relatively simple improvements, they illustrate the need for further testing and refinement of the models. They also indicate that the path planning needs to incorporate more "knowledge" about flying; when close to the ground, more precision is required, causing the pilots to prefer gentler turns, and gentle course intercepts.

If we had it to do over again, we would split the experiment into two phases: in the first phase, we would remove the decision making aspect and have pilots fly approaches to many different airports with various damage models and weather conditions. We could then use this information to improve the risk model and path planning. In a second phase we would then evaluate the role of the ELP in helping the pilots to make quicker and better decisions.

Unfortunately, getting software like the ELP into the cockpit of commercial transport aircraft is a difficult process. The certification process for commercial avionics is both time consuming and costly. It's also not clear how to verify that the ELP will give the best recommendations in all cases, which opens up the manufacturers to additional liability concerns. Although we are beginning to talk to avionics manufacturers, there are other possible ways of fielding some of this technology that may prove much easier. An increasing number of general aviation pilots are now using handheld devices in the cockpit for maps, charts, and GPS navigation. Such devices range from specially designed units like the Garmin GPSMAP 695/696 to Aviation apps like Foreflight HD for the Apple iPad. Much of the ELP's capability could be incorporated into such a unit, with the advantage that certification is not required. The disadvantage is that tight integration with the aircraft avionics and autopilot are not possible with this solution. A second possibility is to work with an airline or freight carrier to make the technology available through the ACARS system. ACARS is a datalink system that allows communication of data between dispatching centers and aircraft cockpits. The information is accessed through special pages on the cockpit CDU. Using this approach, the ELP could be based on the

ground, and recommendations would be sent to the CDU through ACARS. This also has the advantage that certification is unnecessary, and that the dispatching center would be aware of and could assist with the emergency. We are just beginning to explore these possible avenues for fielding this technology, but hope to forge a partnership with one or more of these players.

The biggest assumption behind this work is that the flight envelope for the damaged aircraft is completely known. For certain categories of failures such as engine failures and control surface failures, the flight envelope can be computed and tested in advance and stored in a library. However, the effects of arbitrary damage are more difficult to predict and it therefore seems unreasonable to suppose that complete models for these conditions are available in a library. In stabilizing the aircraft, the adaptive controller explores portions of the flight envelope, and learns how deflections of the control surfaces affect the aircraft. As a result, the adaptive controller can provide a partial model of the flight envelope as well as some knowledge about areas of the envelope that are likely to exceed control limits. However, there may still be areas of the flight envelope that are only partially known. The most conservative approach is to only consider solutions that remain within the known portion of the flight envelope. However, further exploration of the envelope might result in the ability to produce better solutions (for example, slowing the aircraft could allow a shorter runway to be used). Of course, exploration of the flight envelope involves risk, which must be balanced by any potential gains. In general, this problem becomes a POMDP since we have beliefs about the flight envelope, and we can refine those beliefs through actions that explore the flight envelope. But those actions could also throw us into an undesirable and unrecoverable state. Fortunately, flight envelopes do not appear to be this ill-behaved as a rule. As one moves into a particular state in the control envelope, one gains knowledge of the surrounding states, by virtue of how close the control surfaces are to saturation. It is therefore possible to explore the boundaries of the known control envelope and learn what additional states can be explored without undue risk. We have developed a prototype planner that can generate conditional plans that explore portions of the flight envelope and select different landing options based on the outcome of those exploration actions. Surprisingly, this planner is proving to be much more efficient than we expected and we now believe it may be possible to do this in practice. A more detailed description of this work can be found in (Meuleau and Smith 2011; Meuleau et al. 2010). This approach does raise a difficult user interface issue: how does one depict conditional plans of this sort for pilots? Perhaps the best approach is to only display the most probable path and landing site with some indication that there are decision points along the way.

A second assumption we have made is that the flight envelope remains constant once damage has occurred. It is always possible that additional failures may occur, causing the flight envelope to change again. Unless these subsequent failures are predictable, the best that can be done is to run the ELP again when the failure occurs. A more difficult situation is when the flight envelope changes continuously over time. As an example consider the situation where the left wing is damaged and fuel is leaking out at a rapid rate. Initially, there is loss of lift on the left wing and the aircraft has a tendency to roll to the left. As fuel continues to leak out, the left wing becomes lighter counteracting the loss of lift. As more fuel leaks out, the right wing becomes heavy and the aircraft develops a tendency to roll to the right. Whether or not one prefers a left or a right crosswind on landing therefore depends on how long it will take to get to the runway. To our knowledge, the problem of path planning with continuously changing dynamics has not been addressed in the literature. We think that our approach of doing A^* search over a probabilistic roadmap should still be effective, but the heuristic must take into account the estimate of the time that will be required to reach the runway, so that the landing risk can be evaluated using a reasonably correct estimate of what the flight envelope will be at the time.

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

Campbell, S.; Kaneshige, J.; Nguyen, N.; and Krishnakumar, K. 2010a. An adaptive control simulation study using pilot handling qualities evaluations. In *AIAA Guidance, Navigation, and Control Conference*.

Campbell, S.; Kaneshige, J.; Nguyen, N.; and Krishnakumar, K. 2010b. Implementation and evaluation of multiple adaptive control technologies for a generic transport aircraft simulation. In *AIAA Infotech@Aerospace 2010 Conference*.

Meuleau, N., and Smith, D. 2011. Optimal motion planning with uncertain dynamics. In *Submitted to Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI-11)*.

Meuleau, N.; Plaunt, C.; Smith, D.; and Smith, T. 2009a. A comparison of risk sensitive path planning methods for aircraft emergency landing. In *ICAPS-09: Proceedings of the Workshop on Bridging The Gap Between Task And Motion Planning*, 71–80.

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

Meuleau, N.; Plaunt, C.; Smith, D.; and Smith, T. 2010. A pomdp for optimal motion planning with uncertain dynamics. In *ICAPS-10: POMDP Practitioners Workshop*.

Meuleau, N.; Neukom, C.; Plaunt, C.; Smith, D.; and Smith, T. 2011. The emergency landing planner experiment. Technical report, NASA Ames Research Center.