# PlayMaker:
# An Application of Case-Based Reasoning to Air Traffic Control Plays*

Kenneth R. Allendoerfer and Rosina Weber

Drexel University, College of Information Science and Technology
Philadelphia, PA 19104, USA
{kenneth.allendoerfer, rosina.weber}@drexel.edu

**Abstract**. When events such as severe weather or congestion interfere with the normal flow of air traffic, air traffic controllers may implement plays that reroute one or more traffic flows. Currently, plays are assessed and selected based on controllers' experience using the National Playbook, a collection of plays that have worked in the past. This paper introduces PlayMaker, a CBR prototype that replicates the Playbook and models how controllers select plays. This paper describes the PlayMaker design, a model validation, and discusses developments necessary for a full-scale CBR tool for this application.

## 1 Introduction

The Air Traffic Control System Command Center (ATCSCC) is responsible for establishing nationwide responses to situations that affect air traffic control (ATC) operations [1]. Controllers at the ATCSCC develop and communicate plans to local ATC facilities for implementation. Over time, the ATCSCC has collected these plans, called *plays*, to standardize them and make them easier to communicate. In this paper, we describe the development and performance of PlayMaker, a CBR prototype designed to represent and recommend ATC plays.

### 1.1 ATC National Playbook

Every day, thousands of aircraft follow routes between their departure and arrival airports. Routes are *roads in the sky* that minimize flight time between airports and balance overall congestion. For example, not every aircraft flying from the northeastern United States to Florida can fly precisely the same route because of congestion and capacity. Instead, there are many routes between cities but some are more direct and desirable than others. Controllers at the ATCSCC continually monitor the routes and intervene when situations disrupt the normal flow.

*In Funk, Peter; González Calero, Pedro A. (Eds.), Advanced in Case-Based Reasoning (Lecture Notes in Artificial Intelligence , Vol. 3155): Springer-Verlag.

Bad weather and congestion are the most common situations requiring controller intervention. Large aircraft can fly safely through nearly any weather, but small aircraft cannot, and all pilots prefer to avoid bad weather if possible. When bad weather interferes with normal routes, controllers reroute traffic to avoid it. They select new routes based on where the bad weather is located, the route's original path, and conditions in areas where the aircraft may be sent. For example, if a storm interferes with traffic headed from the northeastern United States to Florida, controllers may reroute traffic to the west over the Appalachian Mountains. However, if the Appalachian routes are already congested, controllers may reroute some aircraft over the Atlantic Ocean. However, aircraft that are not equipped for flight over water would need to be rerouted even farther west or delayed on the ground until the congestion eases or the storm weakens. Other situations that may affect flows include military operations, equipment outages, and national emergencies.

Over time, controllers have identified numerous recurring situations and have collected their solutions into plays. A play contains the reroutes that are necessary to handle a situation. Figure 1 shows a play called Snowbird 5, which reroutes southbound traffic away from the Carolina coast and over the Appalachian Mountains. The reroutes in a play have been selected and negotiated with stakeholders to minimize their impact on airlines, airports, and ATC facilities.

When a situation arises that requires reroutes, controllers select and initiate the play that they believe is best for the situation and the overall ATC system. Controllers then contact system stakeholders such as the airlines, local ATC facilities, and the military to inform them which play has been selected. There is some negotiation during these discussions but ultimately, the decision regarding which play to implement comes from the ATCSCC. Because the plays are published ahead of time, the communication burden is much lower than if controllers had to communicate every reroute to every stakeholder individually. The ATCSCC has collected their plays into the National Playbook, which is available online and updated regularly [2].

### 1.2 Motivating Problem

First, the National Playbook is currently maintained manually. As of early 2004, there are over 130 plays in the Playbook and many are only subtly different. Controllers and other stakeholders must rely on their knowledge of how plays performed in the past to identify the best one for the current situation. This requires extensive training and experience. A system that could recommend plays could potentially assist less-experienced personnel and could be used as a training tool.

Second, many current controllers were hired shortly after the 1981 strike when President Ronald Reagan fired more than 11,000 controllers. Because of this, many experienced controllers will be eligible to retire starting about 2006. With the possibility of such a large number of retirements imminent, there is a need to begin capturing the collective knowledge of these experts in a form that could be used by a new generation of controllers.

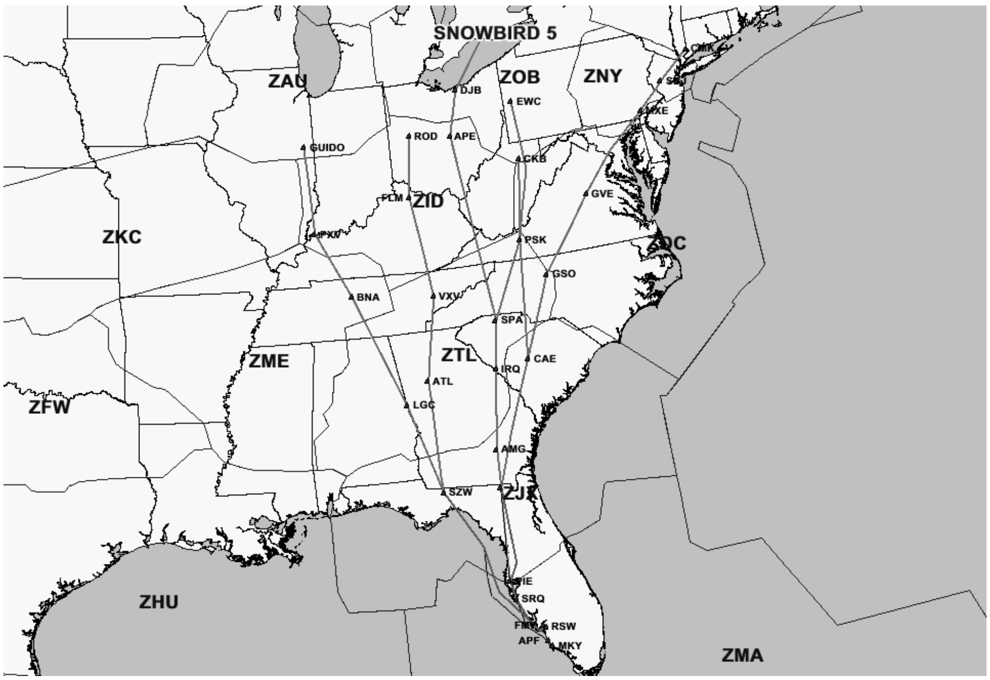


**Fig. 1**. A map showing the play known as Snowbird 5. Dark lines show routes of southbound traffic heading for western Florida. Note that traffic is avoiding the Carolina coast where an event, typically weather, is occurring. From the National Playbook [2]

Third, when a play is selected, it is disseminated to stakeholders such as airlines. When a play involves substantial reroutes, which cause delays, annoy customers, and cost money, stakeholders increasingly demand explanations and justifications. In addition, the ATCSCC is interested in increasing the transparency and accountability of its processes. When controllers implement a play, it would be beneficial if a system could retrieve concrete examples of occasions when the play worked well to serve as justification.

Finally, novel situations do occasionally arise for which no play exists. A system that could process similarities across many previous events might be able to detect similarities that humans might not immediately identify. A system that could suggest candidate plays could help organize, focus, and expedite the play design process. A system that could represent the creative process by which experts design new plays would be relevant to cognitive science in general.

The prototype described in this paper, PlayMaker, addresses the first three of these goals; we have intentionally left the process of developing new plays to later versions. We selected to implement PlayMaker using CBR because it seemed to fit the procedures and practices already used by controllers. Patterns and regularities between situations became apparent over time and controllers recorded them for use later as plays. This is very analogous to the CBR cycle [3] and should provide good face validity and acceptance from the users. CBR systems can provide explanations by analogy which would be required for any operational use of PlayMaker. A more advanced version of PlayMaker using CBR could be designed to derive new plays and thereby meet all the goals.

### 1.3 Past Research

ATC is widely studied in AI because of its complexity, the expertise of its workers, and its safety and economic implications. Bonzano and Cunningham [4] and Bonzano, Cunningham, and Meckiff [5] have developed several CBR systems to support ATC decision making. They focused on identifying and avoiding conflicts, which is a critical aspect of tactical ATC. A conflict occurs when two or more trajectories bring aircraft closer than allowed. The most important controller tactical task is to identify potential conflicts and resolve the situation before safety is jeopardized. Bonzano, Cunningham, and Meckiff [5] created a CBR system that was able to provide good resolutions to several types of conflicts.

However, their system targeted reasoning at a tactical level, not strategic. Strategic controllers do not identify or resolve conflicts. They do not work at the individual flight level; they work with whole flows of traffic. In tactical ATC, a controller sees a conflict and immediately takes action to change one or both trajectories. It is not a negotiation; a controller gives commands directly to pilots that are specific to the conflict at hand. In strategic ATC, however, all stakeholders know the possible actions long ahead of time because plays already have been negotiated, validated, and disseminated. The task, then, is to identify which known action best applies to the current situation and inform others of the decision. The amount of tailoring is intentionally minimized.

Recent research in the field of highway traffic operations centers, which respond to events like accidents and congestion, is related to the strategic ATC task. Logi and Ritchie [6] have developed an innovative decision support system for managing highway congestion across multiple jurisdictions. In their system, intelligent agents develop solutions to problems and then negotiate to develop cooperative overall responses.

Logi and Richie's [6] approach might apply to strategic ATC as well because there are many similarities between the domains. But unlike the highway traffic situations, the ATCSCC does not truly negotiate plays in real time. Instead, controllers select from a collection of pre-determined plays. Some real-time modifications may occur but only when absolutely necessary. The plays are designed to minimize impact to stakeholders and to be as straightforward as possible to implement. Dissemination is a major obstacle for strategic ATC. There may be thousands of flights already in progress. There are hundreds of ATC facilities that may need to learn about the change. Without having a pre-determined Playbook, it may be impossible to reach all stakeholders with accurate and timely information. The Logi and Ritchie [6] approach may be useful when modeling the play design process but less so when modeling how plays are used and implemented.

We know of no research in the AI literature that examines how controllers assess and select strategic plays. Because these strategic decisions may affect hundreds of flights and thousands of passengers, the economic importance of these decisions is clear. Because of important differences between related literature and our area of interest, we believe a different approach is warranted. In particular, we hope to develop an intelligent system that closely matches current strategic ATC practice and that has good face validity for the controllers.

# 2 PlayMaker Design

PlayMaker is a feature-vector CBR system that we prototyped using the Esteem CBR shell [7]. We conceptualize it as a recommendation system that examines existing plays and suggests ones that may be useful for the current situation. It does not generate new reroutes nor does it attempt, at this point, to adapt existing plays. In this section, we describe the features and the similarity metric we used to develop the prototype. In section 4, we propose a more complete implementation of PlayMaker and discuss the changes that would be necessary to create a system for deployment.

## 2.1 Case Description

We developed a set of features that describe situations that cause controllers to institute plays by interviewing a supervisory ATC specialist from Jacksonville Air Route Traffic Control Center. Each of these is described in the following paragraphs.

*Location* is the most important feature of a strategic ATC situation. Location determines which routes are affected and which must be changed. Location determines where rerouted traffic can be sent. It determines which facilities must implement the play and it is one factor that determines how bad a situation is. For the purpose of the prototype, we have defined Location as two features, *East-West* and *North-South*, according to a grid of the United States. One unit on the grid represents approximately 50 nautical miles, with 0, 0 located over northern Kansas. We selected this system to avoid the complexities of latitude-longitude geometry. An operational implementation of PlayMaker would require proper geometry and a more sophisticated method for specifying locations (see section 4).

*Severity* describes the level of operational impact of a situation. For the purpose of selecting plays, controllers do not need to distinguish between causes of events. A thunderstorm that airplanes cannot fly through is the same severity as a military operation that airplanes are not allowed to fly through. The feature Severity, then, describes a situation in terms of the portion of aircraft that would normally fly through an area that cannot because of the event. A low score means that most airplanes that normally come through an area still can. For example, if certain navigation beacons are out of service, older aircraft might not be able to fly through the area whereas newer aircraft would still be able to. In this case, the event is considered low Severity. On the other hand, a military operation might require that all civilian aircraft avoid the area—a complete roadblock. This would receive a high score because all aircraft must be rerouted. As coded in our prototype, each case receives a Severity score between 1 (almost no aircraft must be rerouted) to 5 (all aircraft must be rerouted).

*Direction of Majority* describes the predominant flow through the affected area. Plays are usually developed for traffic coming from one general direction. This is coded in the prototype according to compass directions: N, NE, E, SE, S, SW, W, NW.

*Time of Day* describes when the situation is occurring. Controllers may implement different plays (or no play at all) depending on when a situation occurs because traffic levels change throughout the day. Our domain expert explained that situations occurring during the heavier daytime hours are sometimes handled differently than if the same event occurred in the evening or overnight. A play designed for a daytime situation would not be appropriate for an overnight situation and vice versa. In the prototype, we have broken the day into three categories: 0000 to 0659, 0700 to 1859, and 1900 to 2359.

*Effect on Major Facilities* describes how many major facilities are affected by the situation. Because of the interdependencies of the ATC system, a situation may have effects all over the country but events that affect major facilities deserve special consideration. If two plays could address the same situation equally well, controllers would like to select the play with the least effect on major facilities. For the purpose of our prototype, we defined major facilities as New York (known to controllers as ZNY), Washington (ZDC), and Indianapolis (ZID) as major due to the number of large cities that they serve and their high traffic volumes. We focused on these areas because our case base focuses on these areas. A full implementation of PlayMaker would require including major facilities from across the country. As coded in the prototype, the feature can vary from 0 (no major facilities affected) to 3 (ZNY, ZDC, and ZID affected).

### 2.2 Similarity Assessment

The prototype uses a weighted similarity metric based on priorities elicited from our domain expert and aimed at the intended use of the system. The most important index features, according to our domain expert, are the Direction of Majority and the two Location features. If a retrieved case is not a close match on these three features, it will very likely not be an appropriate play. The Severity, Time of Day, and Effect on Major Facilities features are more useful when selecting among several good candidates. These judgments of operational importance by our expert were used to determine the weights for the similarity metric. The matching features and weights are shown in Table 1.

The features Location, Severity, and Effect on Major Facilities use the Absolute Fuzzy Range numeric matching function provided by Esteem [7], which returns a number between 0 and 1 according to the formula MAX (0, ABS(value2-value1)/Range). This allows the retrieval of cases whose values are close to the original. This is especially important for Location because of the imprecision of the method we used to specify location. Time of Day uses the Exact text matching function of Esteem, which returns a 0 or 1 if the text matches precisely. Direction of Majority uses the Partial, Case Indifferent text matching function of Esteem, which returns a value between 0 and 1, depending on the number of differences in the letters. This allows directions "NE" and "N" to be treated more similarly than "S" and "N."

**Table 1**. Features, matching functions, and weights of the playmaker similarity metric

| Feature | Type of Feature Matching* | Weight |
|---|---|---|
| Location East-West | Absolute Fuzzy Range: 3 | 0.190 |
| Location North-South | Absolute Fuzzy Range: 3 | 0.190 |
| Severity | Absolute Fuzzy Range: 2 | 0.095 |
| Time of Day | Exact | 0.048 |
| Direction of Majority | Partial (case indifferent) | 0.381 |
| Effect on Major Facilities | Absolute Fuzzy Range: 1 | 0.095 |
| *The types of feature matching are functions defined in Esteem® [7] | | |

### 2.3 Cases

For our prototype, we selected 20 plays from 135 available in the National Playbook [2]. We focused on plays affecting the southeastern United States because our ATC domain expert was most knowledgeable about this area and its traffic flows. A full implementation of PlayMaker would require input from multiple controllers who work in various parts of the country.

In CBR, a case is a representation of an experience including a situation and a solution. In ATC, a play is representation of a solution only. To build the cases for PlayMaker, we interviewed a domain expert about each play to determine the situations to which it applied. Our domain expert examined each play and described the situations where it is normally used. This process resulted in 32 cases based on 20 plays because several plays could be applied to multiple situations. For example, the *Florida to NE 1* play can only be applied to situations occurring during the day whereas the *Florida to NE 2* can be applied to situations occurring in the day, evening, or overnight. We created cases for each situation to which a play can be applied. We then presented all the situations to our expert and asked him to characterize each based on the features.

**Example.** Northbound traffic from Florida follows a path that takes them from the east coast of Florida, off the coast of Georgia and South Carolina, then over North Carolina, Virginia, Maryland and Delaware. Today, however, a strong storm will reach North Carolina and Virginia at about 10:00 am, closing the area to traffic. Controllers at the ATCSCC look at the weather forecast and choose a play to address the situation. A likely choice is a play called *Florida to NE 3* that reroutes this traffic west over the

Appalachians and Ohio River Valley. This play is quite severe and affects hundreds of flights. A representation of this situation as a case in PlayMaker is shown in Table 2.

**Table 2**. An example problem situation described using the PlayMaker features

| Feature | Value | Comments |
|---|---|---|
| Location East-West | 16 | These Location values yield the southeast of Virginia and the northeast of North Carolina. |
| Location North-South | -2 | |
| Severity | 5 | No traffic may come through the area |
| Time of Day | 0700 to 1859 | Storm forecast to arrive at 10:00 am |
| Direction of Majority | N | Northbound traffic primarily |
| Effect on Major Facilities | 2 | New York and Washington are both affected |

# 3 Prototype Validation

We performed a validation of the PlayMaker prototype by comparing its responses to novel situations to those of an ATC domain expert. Our goal is to assess whether the Playmaker model can make recommendations like a human expert.

## 3.1 Methodology

The dataset we used included all the cases we coded in the Playmaker and a set of 6 new situations (first column Table 3) that we designed to test the prototype. Because our case base included only cases in the south and southeastern United States, we developed our test situations from this region as well. We constructed six test situations using the following procedure. First, we identified several general locations that were represented in our case base: the Carolinas, the Mid-Atlantic, the Appalachians, and Dallas. We identified one or two cases in each location and changed one or more features. We developed a verbal description of the situation and presented these to our expert. We also developed PlayMaker target cases for each situation and entered these in the PlayMaker prototype.

**Table 3**. Test situastions and recommend plays from the domain expert and PlayMaker

| Test Situation | Domain Expert | PlayMaker: Top 3 retrievals (with ties) | Sim. Score |
|---|---|---|---|
| 1. Hurricane out in the Atlantic Ocean. Effects (heavy rain, high winds) will start being felt in the Outer Banks of North Carolina around 1000 today. | There is no play in the playbook for this situation. It would be negotiated. | Inland Atlantic Route – Northbound | 65 |
| | | MGM 1 – Morning | 62 |
| | | Florida to NE3 - Morning | 58 |
| | | Florida to NE3 - Evening | 58 |
| 1.1 Move location of Test Situation 1 to near Wilmington, NC. All other features identical. | Florida to NE 1 | Florida to NE 1 - Morning | 65 |
| | | Inland Atlantic Route – Northbound | 65 |
| | | Florida to NE 2 - Morning | 62 |
| 2. Strong thunderstorms north and northeast of Dallas starting around 1500 today. | DFW East 1 | DFW East 1 – Morning | 80 |
| | | DFW East 1 – Evening | 75 |
| | | Florida to NE 2 - Morning | 42 |
| 3. Moderate storms over southern North Carolina and all of South Carolina, affecting north and southbound routes, starting around 2000 tonight and continuing until 0300 tomorrow. | Florida to NE 2 | Florida to NE 1 - Morning | 63 |
| | | Inland Atlantic Route – Northbound | 60 |
| | | Florida to NE 2 - Evening | 58 |
| 4. Two navaids out of service along J75 from 0000 tonight until 0600 tomorrow. | No J75 1 or No J75 3, depending on the level of congestion in ZDC. | A761 | 79 |
| | | Inland Atlantic Route – Southbound | 79 |
| | | No J75 3 | 73 |
| | | No J75 1 | 73 |
| | | No WHITE/No WAVEY | 73 |
| 5. Very heavy congestion expected all day today along northbound routes in ZDC. | Florida to NE 3 | Inland Atlantic Route - Northbound | 90 |
| | | Snowbird 7 | 65 |
| | | Florida to NE 3 - Morning | 58 |
| 6. Moderate weather affecting all Appalachian Mountain areas all day today, coupled with heavy congestion in ZDC during the day. | Florida to NE 3 | Inland Atlantic Route | 65 |
| | | MGM 1 – Morning | 58 |
| | | MGM 1 - Evening | 53 |

### 3.2 Results

The full validation results are presented in Table 3. For four of the six test situations, the play recommended by the domain expert was also one of the top three plays retrieved by PlayMaker. However, the similarity scores on PlayMaker were low (58-80% match). This indicates that while our case base and metrics are on the right track, more work is needed to increase the similarity scores for the top retrieved cases. In particular, because of their importance to the match, small discrepancies or inconsistencies in the features Location and Direction of Majority have major effects on the score. For example, on Test Situation 2, changing Direction of Majority from W to SW (which would be another reasonable interpretation of the direction) changes the similarity score from 80% to 65%. More precision in specifying these features is necessary.

**Mismatches.** The two cases where PlayMaker did not match the expert's recommendations at all reveal two interesting complexities of the system that must be incorporated in later versions.

First, Test Situation 1 was unexpectedly difficult for PlayMaker to handle. We unintentionally created a situation for which controllers in the field do not use a Playbook play. In the test situation, controllers at the ATCSCC would develop an ad hoc play to precisely address the situation rather than institute a stored play. PlayMaker, on the other hand, retrieved several cases that seem to fit. The similarity scores were low but were in the range of our other test situations.

To continue within the validation, we modified Test Situation 1 and moved it several units southwest to Wilmington, NC. This modification is listed as Test Situation 1.1 in Table 3. In this case, the domain expert recommended the same play as PlayMaker. The most important factor, then, appears to be the location of Test Situation 1. When this situation affects the Outer Banks, controllers do not use a Playbook play but do when the same situation occurs farther south. For this situation, our Location features and similarity metric (which uses a the Fuzzy Range function provided by Esteem) are not precise enough to make this distinction easily. Apparently, in some locations, small differences are important but in other areas, they are not.

The best way to make PlayMaker handle Test Situation 1 is to improve the precision of Location and increase the overall retrieval threshold. Because so many of our similarity scores were generally low, we set a low threshold (50%) in order to retrieve at least a few cases. However, in situations like Test Situation 1, it appears that small differences in Location can be operationally significant. Better representations of Location would provide higher match scores and allow us to increase our threshold. Because no stored cases would be similar to Test Situation 1, PlayMaker's retrieved cases would fall below the threshold and PlayMaker would essentially respond "No Play."

The second mismatch, Test Situation 6, was more complex than our set of features and similarity metric could handle. We could specify either the location of the storm (Appalachian Mountains) or the location of the congestion (ZDC) but not both. As such, PlayMaker's best match avoided the storms in the Appalachians but rerouted the traffic directly into ZDC, which would compound an already bad situation there. The expert recognized this complexity and selected a play that reroutes some of the traffic through the weather in the Appalachians and avoids the most congested areas of ZDC. This is allowable because the severity of the weather was moderate. Future versions of PlayMaker should provide some way to handle multiple simultaneous situations within a single case.

## 4 Future Work

The PlayMaker prototype presents a first step toward developing a CBR system to replicate the process of using plays in the ATC National Playbook. While it does not demonstrate the complete concept, it helps demonstrate that the idea is feasible and makes clear the areas that need urgent attention. It was not worthwhile to start directly developing the system from scratch without before having this first assessment. The next steps will consist of the development of a prototype developed from scratch to this application including all elements required to demonstrate the feasibility of transitioning the tool for deployment.

### 4.1 Improvements to Features

First, the representation of Location is very crude. We created a two-dimensional grid and used it to represent an entire event location. In the actual operation, however, an event could be a three-dimensional weather system hundreds of miles long and changing over time. Characterizing such a complex event using our Location features lacks a great deal of information and makes the feature susceptible to error when coding cases or presenting them to experts.

An operational version of PlayMaker would allow users to specify locations as polygons, using latitude-longitude geometry, and would contain altitude and movement parameters. This is similar to how controllers define what they call Flow Constrained Areas in their automation system. In Figure 2, we illustrate how such a feature might look.

The next version of PlayMaker should represent this feature with graphs. Consequently, we will need to employ graph matching methods [8] to assess the similarity between these locations. This would allow much more accurate matching of locations than the current method which is very subjective and error prone.

Second, we must similarly add precision to the Direction of the Majority metric. Currently, an entire set of complex routes must be distilled to a single direction. This is subjective and introduces mistakes. For example, the play known as *No WHITE No WAVEY* is generally southbound but its routes are actually sometimes southbound, sometimes southeast, and sometimes southwest bound.

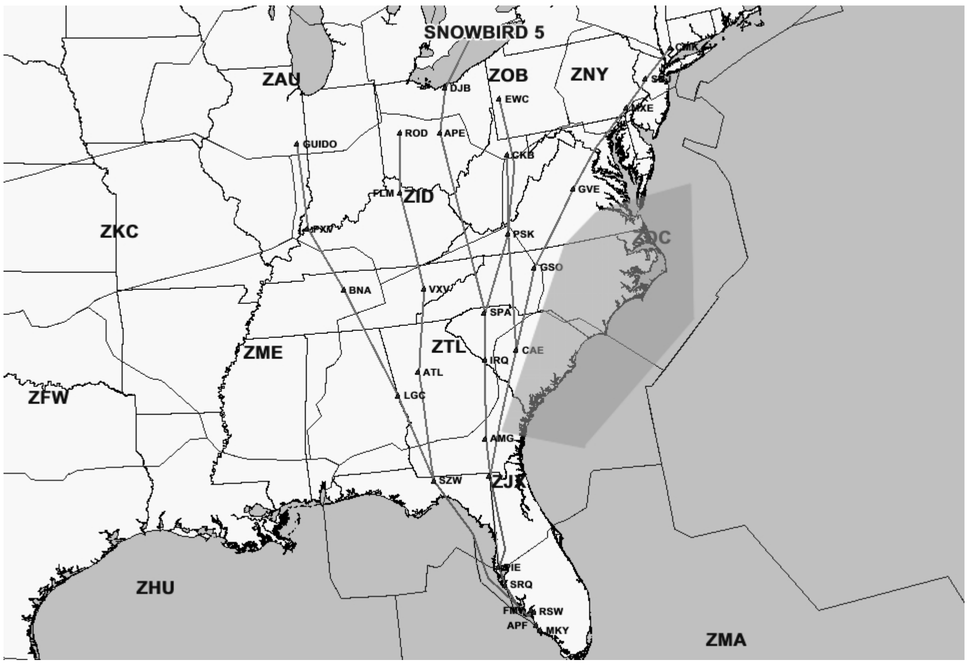


**Fig 2**. A map showing the play known as Snowbird 5 with a more complex and precise representation of Location, shown as a gray polygon

When the situation is entered into PlayMaker, the user would need to judge how to characterize the flows. Because of the importance of the Direction feature, a small discrepancy on this feature can have a major effect on the match. A better method is to characterize the flows according to categories rather than directions. For example, instead of categorizing traffic coming to Florida as southbound, it could be characterized as *East Coast Florida Arrivals*. This more closely matches how controllers talk about this feature and it would reduce the chance of discrepancy or error.

Third, the system must be able to represent complex situations. In particular, it must be able to represent situations with multiple event locations. For example, our Test Situation 6 had two locations (Appalachian Mountains and ZDC) and PlayMaker could not provide a useful recommendation. One way to implement this improvement is to specify a list of events (e.g., primary, secondary) events based on their severity. For example, in Test Situation 6, the storm would be the primary event because few aircraft could use the area. The congestion over ZDC would be a secondary event.

Finally, current controller automation tools can model traffic flows and can compute which traffic would be affected by a situation in a given three-dimensional location. PlayMaker could use these to compute Severity rather than having to ask experts to estimate it. Examining our cases, there was not a wide range of responses and it seemed that our expert was not confident making the severity judgments. The Severity feature would still represent the same aspect of the situation. That is, how much of the traffic that normally traverses the area will be able to do so when the event occurs? Severity, in this case, could change from a categorical feature to a proportion. Additionally, we may investigate the benefit of incorporating fuzzy labels into this representation. Given the presence of uncertainty in the acquisition of the feature, we may rely on the use of fuzzy rules as employed by Bandini and Manzoni [9]. Main and Dillon [10] have already combined fuzzy and crisp values using a retrieval method based on neural networks. We shall examine its extension also to combine graph-based features.

### 4.2 Improvements to the Case Base

The PlayMaker case base currently contains only a small portion of the existing cases and is focused on the southeastern United States. Expanding the case base to include all areas of the country may reveal additional insights about the features and similarity assessment. This will involve knowledge elicitation sessions with controllers working in numerous different areas.

In addition, our case base was developed by using the existing plays and asking an expert to describe the situations to which the play applies. An alternative method might be to use ATCSCC records of real situations. This would significantly increase the size of the case base because a play may be used for many situations. By basing the case base on real situations rather than typical ones, we might remove some of the subjectivity inherent in determining the characteristics of a typical situation.

### 4.3 Adaptation

PlayMaker currently contains no methods for adaptation but a system that could derive new plays from existing ones likely would be useful. The difficulty is that the current method by which plays are designed is extremely complex and not well understood. By involving experts in the CBR adaptation process itself, we could leverage their rich knowledge knowledge of the airspace, existing traffic flows, and political and economic implications of different decisions. For example, experts could be asked to manually adapt recommended plays in PlayMaker which would help us understand their process.

Deriving new plays requires understanding the effects of reroutes on the overall system. As it stands, PlayMaker uses only plays that have already been designed and agreed upon by controllers and stakeholders. Any recommended play is internally sound because it would not exist in PlayMaker otherwise. In this way, PlayMaker reflects the knowledge of the experts without explicitly containing it. For derived plays, however, no such assurance would exist because the derived plays would not exist in the Playbook and would not have undergone the negotiation and validation process. To make use of an already successful manual process, PlayMaker could be built to derive plays offline without waiting for situations to occur and present the derived plays to the controllers and stakeholders. The experts could then negotiate and improve the play as needed and enter the final play into PlayMaker. In this way, derived plays would continue to have some assurance of validity without modeling all aspects of the play development process.

### 4.4 Maintenance

Watson [11] discusses methods for maintaining CBR systems when they are fielded. If PlayMaker were fielded, it could quickly develop a large base of actual cases because a play is implemented almost every day during the summer. As Watson discusses, redundant cases will need to be identified and eliminated. In addition, as traffic flows in the ATC system change, plays and their associated cases will become obsolete and will require modification or removal. For example, when a new airline moves into an airport (or goes out of business), the volume and characteristics of routes from that airport change. Affected plays may need to be reviewed and revised to reflect the new traffic flows. Once adaptation is included in PlayMaker, adaptation rules will also need to be monitored and refined over time as well.

## 5 Conclusions

Though this is an initial study, it suggests that the PlayMaker concept is promising. To create a useful system that can recommend plays, next enhancements should focus on improving the case base and the index features. Then, enhancements should seek to add adaptation functions to move the tool toward supporting the design of new plays.

Finally, PlayMaker highlights some of the complexities associated with converting information that is already in a case-like form into a CBR system. The experts who designed the plays in the Playbook did not include all their knowledge and experience in the plays. They still hold deep knowledge about the history of individual plays, the past success of plays, and the implications of plays that is not captured in the Playbook itself. Interviewing the experts reveals some of these complexities but others remain undocumented, as the mismatches we encountered demonstrate. Perhaps by encouraging the experts to document their cases in different ways, future versions of PlayMaker can be improved.

## Acknowledgements

We wish to thank Kotb (Sam) Elbialy and Pieter Linden, Master’s students in the College of Information Science and Technology at Drexel University, for their assistance conceptualizing the features of PlayMaker. We also wish to thank David Conley of the Federal Aviation Administration for serving as our ATC domain expert. Rosina Weber is supported in part by the National Institute for Systems Test and Productivity at USF under the USA Space and Naval Warfare Systems Command grant no. N00039-02-C-3244, for 2130 032 L0, 2002.

## References

[1] National Research Council: Flight to the Future: Human Factors in Air Traffic Control. National Academy Press, Washington, DC (1997)

[2] Federal Aviation Administration: Air Traffic Control System Command Center National Severe Weather Playbook. Retrieved February 29, 2004 from the World Wide Web: http://www.fly.faa.gov/PLAYBOOK/pbindex.html. (2004)

[3] Aamodt, A., Plaza, E.: Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. Artificial Intelligence Communications 7 (1994) 39-59

[4] Bonzano, A., Cunningham, P.: Hierarchical CBR for Multiple Aircraft Conflict Resolution in Air Traffic Control. Proceedings of the European Conference on Artificial Intelligence, Vol. 13. John Wiley and Sons, Hoboken, NJ (1998) 58-62

[5] Bonzano, A., Cunningham, P., Meckiff, C.: ISAC: A CBR System for Decision Support in Air Traffic Control. In: Smith, I., and Faltings, B. (eds.): Advances in Case-Based Reasoning. Lecture Notes in Artificial Intelligence, Vol. 1186. Springer-Verlag, Berlin Heidelberg New York (1996) 44-57

[6] Logi, F., Ritchie, S. G.: A Multi-Agent Architecture for Cooperative Inter-Jurisdictional Traffic Congestion Management. Transportation Research Part C 10 (2002) 507-527

[7] Esteem Software Incorporated: Esteem Version 1.4. Stottler Henke Associates, San Mateo, CA. (1994)

[8] Bunke, H., Messmer, B.: Recent Advances in Graph Matching. International Journal of Pattern Recognition and Artificial Intelligence 11 (1997) 169-203

[9] Bandini, S., S. Manzoni, S.: Application of Fuzzy Indexing and Retrieval In Case Based Reasoning For Design. In: Proceedings of the 16th ACM SAC Symposium on Applied Computing. ACM Press, New York, NY (2001) 462-466

[10] Main, J., Dillon, T: A Neuro-Fuzzy Methodology for Case Retrieval and an Object-Oriented Case Schema for Structuring Case Bases and their Application to Fashion Footwear Design. In: Pal, S. K., Dillon, T. S., Yeung, D. S., (eds.): Soft Computing in Case Based Reasoning. Springer-Verlag, London (2001)

[11] Watson, I. A Case Study of Maintenance of a Commercially Fielded Case-Based Reasoning System. Computational Intelligence 17 (2001) 386-398